

Multi-Source Domain Adaptation for Visual Sentiment Classification

Chuang Lin,^{1*} Sicheng Zhao,^{2*} Lei Meng,^{1†} Tat-Seng Chua¹

¹NExT++, National University of Singapore

²University of California, Berkeley, USA

clin@u.nus.edu, schzhao@gmail.com, {lmeng, dcscts}@nus.edu.sg

Abstract

Existing domain adaptation methods on visual sentiment classification typically are investigated under the single-source scenario, where the knowledge learned from a source domain of sufficient labeled data is transferred to the target domain of loosely labeled or unlabeled data. However, in practice, data from a single source domain usually have a limited volume and can hardly cover the characteristics of the target domain. In this paper, we propose a novel multi-source domain adaptation (MDA) method, termed Multi-source Sentiment Generative Adversarial Network (MSGAN), for visual sentiment classification. To handle data from multiple source domains, it learns to find a unified sentiment latent space where data from both the source and target domains share a similar distribution. This is achieved via cycle consistent adversarial learning in an end-to-end manner. Extensive experiments conducted on four benchmark datasets demonstrate that MSGAN significantly outperforms the state-of-the-art MDA approaches for visual sentiment classification.

Introduction

It is increasingly popular for people to share experiences, express opinions, and record activities by posting images on social networks like Instagram and Twitter. It offers a great opportunity for visual sentiment analysis that infers their emotional behaviors and provide personalized services (Zhao et al. 2018b), such as blog recommendation (Borth et al. 2013) and tourism (Alaei, Becken, and Stantic 2019).

Recent advances in deep learning have significantly improved the state-of-the-art performance in visual sentiment classification (Yang et al. 2018a; 2018c; Katsurai and Satoh 2016) or image emotion distribution learning (Yang et al. 2017). However, considering the difficulties in the acquisition of sentiment labels due to the high subjectivity in the human perception process, training a model on a labeled source domain that can well generalize to another new domain is necessary. Because of the presence of *domain shift* or *dataset bias* (Torralba and Efros 2011), even a slight de-

*Equal Contribution.

†Corresponding Author.

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

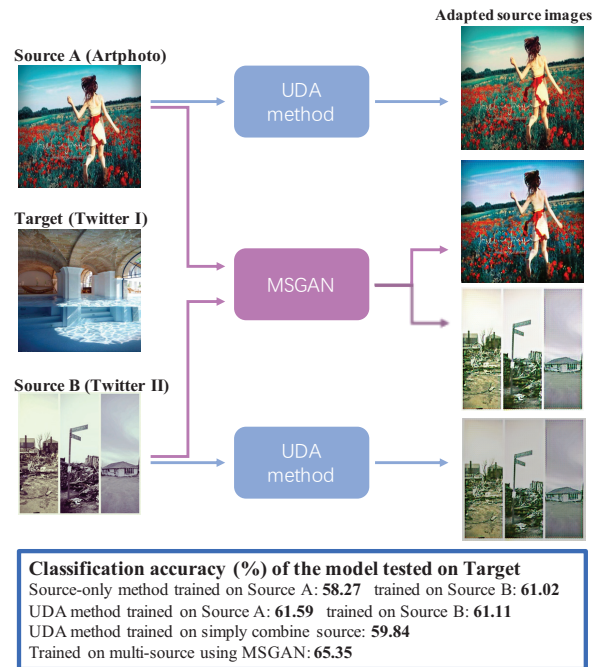


Figure 1: An example of *domain shift* in sentiment classification using training data from multiple sources.

parture from a network’s training domain can lead to incorrect predictions and significantly reduce its performance. Therefore, establishing knowledge transfer from a labeled source domain to an unlabeled target domain for visual sentiment analysis has attracted significant attention (Zhao et al. 2018c; 2019b). Recently, most existing deep unsupervised domain adaptation (UDA) methods assume that there is only one single source domain and the labeled source data are implicitly sampled from the same underlying distribution. In practice, there are obvious biases in different domains, even for the images of the same sentiment. For example, as shown in Figure 1, the styles of artistic photos, abstract paintings, and natural images are quite different. Simply combining different sources into one source and directly

employing a state-of-the-art single-source UDA model (Zhu et al. 2017) lead to a drop in classification accuracy from 61.59% (if trained on the single-best source) to 59.84% (if trained on simply combined source). This is because images from different source domains may interfere with each other during the learning process (Riemer et al. 2019). Comparing to source-only method (Peng et al. 2015), UDA model can alleviate domain shift between single source domain and the target domain to some extent. However, it still exists between multiple source domains and target domain, leading to multi-source domain adaptation (MDA) for visual sentiment classification. Despite of the rapid progress in UDA, no study has been investigated on MDA for visual sentiment classification. This task is much more challenging due to the following reasons. First, the multiple sources are different not only from the target but also from each other. Existing MDA methods only align each source and target pair. Although different sources are matched towards the target, there may exist significant mis-alignment across different sources. Second, these MDA methods focus on matching the visual features but ignore the semantic labels, which hardly characterize the consistency of image sentiment. Third, current methods typically require multiple classifiers or transfer structures for different source domains. This leads to a high model complexity and low learning stability when learning from small-scale sentiment data.

In this paper, we study the multi-source unsupervised domain adaptation problem for visual sentiment classification. Specifically, we propose a novel adversarial framework, termed Multi-source Sentiment Generative Adversarial Network (MSGAN), which is composed of three pipelines, *i.e.* image reconstruction, image translation, and cycle-reconstruction. The image reconstruction and cycle-reconstruction pipelines learn a unified sentiment space where data from the source and target domains share similar distributions. Subsequently, the image translation pipeline restricted by emotional semantic consistency learns to adapt the source domain images to appear as if they were drawn from the target domain, while preserving the annotation information. Notably, thanks to the unified sentiment latent space, MSGAN requires a single classification network to handle data from different source domains.

In summary, the contributions of this paper are threefold:

1. We propose to adapt visual sentiment from multiple source domains to a target domain in an end-to-end manner. To the best of our knowledge, this is the first multi-source domain adaptation work on visual sentiment classification.

2. We develop a novel adversarial framework, termed MSGAN, for visual sentiment classification. By the joint learning of image reconstruction, image translation, and cycle-reconstruction pipelines, images from multiple sources and one target can be mapped to a unified sentiment latent space. Meanwhile, the semantic consistency loss in image translation pipeline preserves the semantic information of images.

3. We conduct extensive experiments on the Art-Photo (Machajdik 2010), FI (You et al. 2016), Twitter I (You et al. 2015), and Twitter II (You et al. 2016) datasets, and the results demonstrate the superiority of the proposed MSGAN model compared with the state-of-the-art MDA approaches.

Related Work

Visual Sentiment Classification: Recently, with the great success of convolutional neural network (CNN) on many computer vision tasks, CNN has also been employed in sentiment classification. You et al. (2015) proposed a progressive CNN architecture to make use of noisily labeled data for binary sentiment classification. Yang et al. (2018b) employed deep metric learning to optimize both retrieval and classification tasks by jointly optimizing cross-entropy loss and a novel sentiment constraint. Different from improving global image representations, several methods (You, Jin, and Luo 2017; Yang et al. 2018a) consider the local information for sentiment classification. All the above methods employ a supervised manner to learn the mapping between image content and sentiments. In this paper, we study how to adapt the models from multiple labeled source domain to an unlabeled target domain for visual sentiment classification.

Single-source Domain Adaptation: Since data from the source and target domains have intrinsically different distributions, the key problem in single-source UDA is how to reduce the domain shift. Discrepancy-based methods explicitly measure the discrepancy between the source and target domains on corresponding activation layers of the two network streams (Sun, Feng, and Saenko 2017; Zhuo et al. 2017); Adversarial generative models combine the domain discriminative model with a generative component generally based on GANs (Goodfellow et al. 2014). The Coupled Generative Adversarial Networks (CoGAN) (Liu and Tuzel 2016) can learn a joint distribution of multi-domain images with a tuple of GANs; Reconstruction based methods incorporate a reconstruction loss to minimize the difference between the input and the reconstructed input (Ghifary et al. 2015; 2016). Zhao et al. (2019b) proposed CycleEmotionGAN for image emotion classification by adapting source domain images to have similar distributions to the target ones by enforcing emotional semantic consistency. However, none of them can handle data from multiple source domains, which is the target of this paper.

Multi-source Domain Adaptation: Compared with single source UDA, multi-source domain adaptation (MDA) assumes that training data from multiple sources are available (Zhao et al. 2019a). Early efforts on this task used shallow models (Sun and Shi 2013). MDA also develops with theoretical supports. Blitzer et al. (2008) provided the first learning bound for MDA. Mansour, Mohri, and Rostamizadeh (2009) claimed that an ideal target hypothesis can be represented by a distribution of a weighted combination of source hypotheses. In the more applied works, Deep Cocktail Network (DCTN) (Xu et al. 2018) proposed a k-way domain discriminator and category classifier for digit classification and real-world object recognition. Zhao et al. (2018a) proposed new generalization bounds and algorithms under both classification and regression settings for MDA. Peng et al. (2018) directly matched all the distributions based on moments and provided a concrete proof of why matching the moments of multiple distributions works for MDA. Different from these methods, we learn a unified sentiment latent space which jointly aligns data from all source and target domains.

Problem Setup

Suppose we have M source domains S_1, S_2, \dots, S_M and one target domain T . In the unsupervised multi-source domain adaptation (MDA) scenario, S_1, S_2, \dots, S_M are labeled and T is fully unlabeled. For the i th source domain S_i , the observed images and corresponding sentiment labels drawn from the source distribution $p_i(\mathbf{x}, \mathbf{y})$ are $X_i = \{\mathbf{x}_i^j\}_{j=1}^{N_i}$ and $Y_i = \{\mathbf{y}_i^j\}_{j=1}^{N_i}$, where N_i is the number of images in S_i . The target images drawn from the target distribution $p_T(\mathbf{x}, \mathbf{y})$ are $X_T = \{\mathbf{x}_T^j\}_{j=1}^{N_T}$ without label observation, where N_T is the number of target images.

The main idea of MSGAN is to learn a mapping that can align the images from both the multiple source and target domains to have similar distributions in a unified sentiment space. As shown in Figure 2, images from both the source and target domains are mapped to have similar distributions, and their information is preserved by the reconstruction loss.

To achieve this, we introduce the encoders $E_{S_1}, E_{S_2}, \dots, E_{S_M}, E_T$ and generators G_{ST}, G_{TS} . To obtain the sentiment latent space, we enforce a weight-sharing constraint on the encoders and generators. Specifically, we share the weights of the last block layers of $E_{S_1}, E_{S_2}, \dots, E_{S_M}$ and E_T to extract the high-level representations of the input images from the multiple sources and target. Similarly, we share the weights of the first block layers of G_{ST} and G_{TS} to decode their high-level representations for reconstructing the input images. The unified sentiment space is learned from the cycle-consistency constraint (Liu et al. 2017): $\mathbf{x}_i = F_{T \rightarrow S}^*(F_{S \rightarrow T}^*(\mathbf{x}_i))$ and $\mathbf{x}_T = F_{S \rightarrow T}^*(F_{T \rightarrow S}^*(\mathbf{x}_T))$, where $F_{S \rightarrow T}^*(\mathbf{x}_i) = G_{ST}(E_{S_i}(\mathbf{x}_i))$, $F_{T \rightarrow S}^*(\mathbf{x}_T) = G_{TS}(E_T(\mathbf{x}_T))$. With the latent space established, we do not need multiple transfer structures.

Multi-source Sentiment Generative Adversarial Network

Our framework, as illustrated in Figure 3, is based on variational autoencoders (VAEs) and generative adversarial networks (GANs). It consists of three pipelines: image reconstruction, image translation, and cycle-reconstruction. Image reconstruction pipeline includes multiple domain image encoders $E_{S_1}, \dots, E_{S_M}, E_T$ encoding images to a unified sentiment space and two image generators G_{ST}, G_{TS} reconstructing input images. Image translation pipeline includes two generative adversarial networks: $GAN_{ST} = \{D_T, G_{ST}\}$ and $GAN_{TS} = \{D_S, G_{TS}\}$ learning the mapping between multiple source domains and the target domain. The cycle reconstruction pipeline is used to learn a unified sentiment space and ensure that features of images from different domains preserved the information of their original images.

Image Reconstruction Pipeline

The image reconstruction pipeline is achieved by multiple encoder-generator pairs $\{E_{S_i}, G_{ST}\}$, each of which maps an input image \mathbf{x}_i to a code in a latent space Z via E_{S_i} and then decodes a random-perturbed version of the code to reconstruct the input image via G_{ST} . We assume the

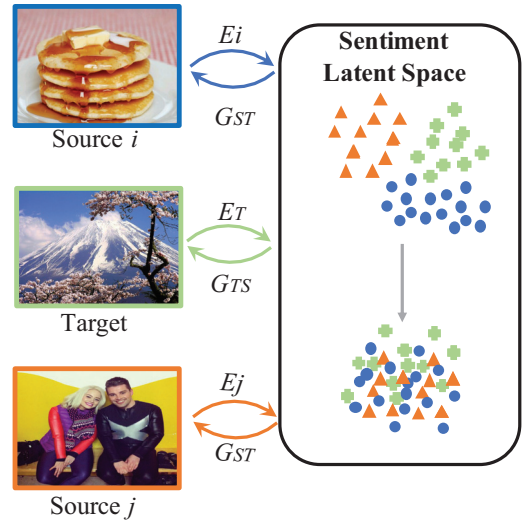


Figure 2: The unified sentiment latent space. MSGAN can map source i , source j and target domain images to a unified sentiment latent space. E_i, E_j and E_T are encoding functions, mapping images to latent codes. G_{ST} and G_{TS} are two generation functions, mapping latent codes to images.

components in the latent space Z are conditionally independent Gaussian distributions with unit variance. The encoder outputs the latent code $\mathbf{z}_i = E_{S_i}(\mathbf{x}_i)$. The reconstructed image is $\mathbf{x}'_i = G_{TS}(\mathbf{z}_i, \eta)$, where η is a noise has the same distribution of \mathbf{z}_i . Similarly, $\{E_T, G_{TS}\}$ constitutes a VAE for the target domain, where the reconstructed image is $\mathbf{x}'_T = G_{ST}(E_T(\mathbf{x}_T), \eta)$.

MVAE training aims to minimize a variational upper bound as follows:

$$\mathcal{L}_{MVAE_i}(E_{S_i}, G_{TS}) = \lambda_0 KL(q_i(\mathbf{z}_i | \mathbf{x}_i) \| p_\eta(z)) - \lambda_1 \mathbb{E}_{\mathbf{z}_i \sim q_i(\mathbf{z}_i | \mathbf{x}_i)} [\log p_{G_{TS}}(\mathbf{x}_i | \mathbf{z}_i)], \quad (1)$$

$$\mathcal{L}_{MVAE_T}(E_T, G_{ST}) = \lambda_0 KL(q_T(\mathbf{z}_T | \mathbf{x}_T) \| p_\eta(z)) - \lambda_1 \mathbb{E}_{\mathbf{z}_T \sim q_T(\mathbf{z}_T | \mathbf{x}_T)} [\log p_{G_{ST}}(\mathbf{x}_T | \mathbf{z}_T)]. \quad (2)$$

where the hyper-parameters λ_0 and λ_1 control the weights of the objective terms and the KL divergence terms penalize the deviation of the distribution of the latent code from the prior distribution. $p_{G_{TS}}$ and $p_{G_{ST}}$ are modeled using Laplacian distributions, respectively.

Image Translation Pipeline

As aforementioned, the unified sentiment space allows to use only one generator G_{ST} to adapt multi-source images indistinguishable from the target domain. G_{ST} can generate two types of images: (1) images from the reconstruction pipeline $\mathbf{x}'_T = G_{ST}(\mathbf{z}_T \sim q_T(\mathbf{z}_T | \mathbf{x}_T))$ and (2) images from the translation pipeline $\mathbf{x}_i^{S_i \rightarrow T} = G_{ST}(\mathbf{z}_i \sim q_i(\mathbf{z}_i | \mathbf{x}_i))$. A similar processing is applied to G_{TS} . Meanwhile, two discriminators D_T and D_S are used to distinguish between \mathbf{X}_T and $G_{ST}(E_i(\mathbf{X}_{S_i}))$, and \mathbf{X}_{S_i} and $G_{TS}(E_T(\mathbf{X}_T))$, respectively.

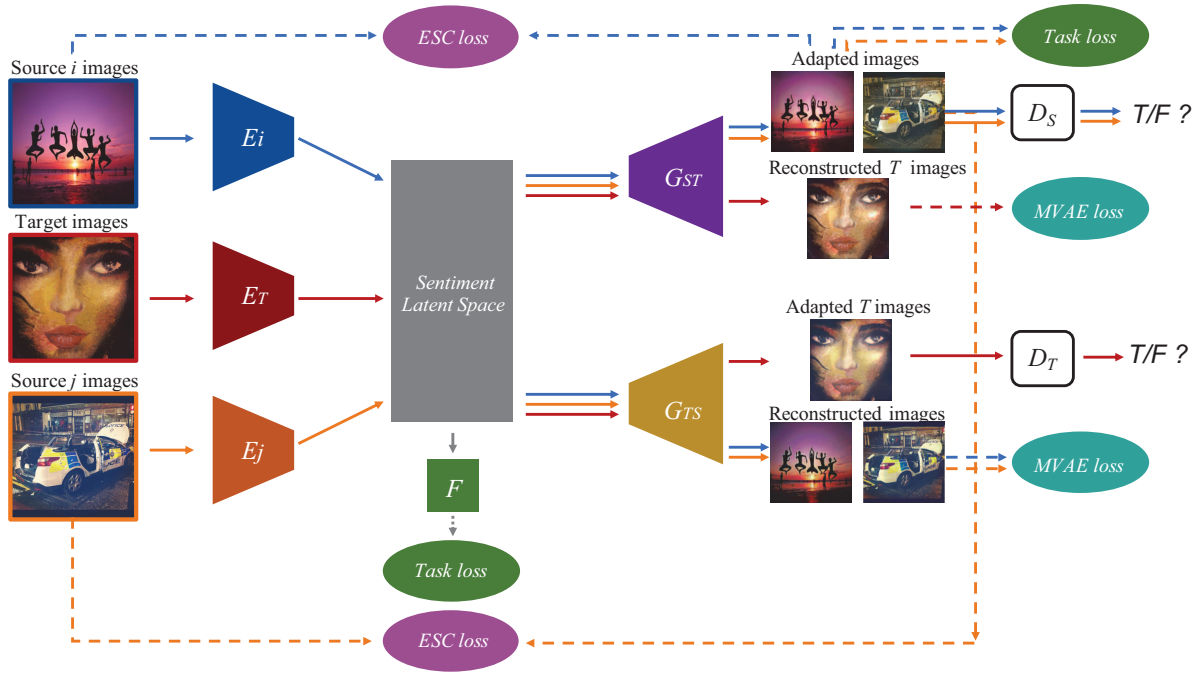


Figure 3: The framework of the proposed Multi-source Sentiment Generative Adversarial Network (MSGAN) for visual sentiment classification. The solid arrowed lines indicate the operations in the training stage. The dot dash arrowed lines correspond to different losses. For clarity the cycle-reconstruction pipeline is omitted. In the testing phase, testing data from the target domain are forwarded through the target encoder and classifier to make the final prediction.

The GAN objective functions are defined by:

$$\mathcal{L}_{GAN_i}(E_{S_i}, G_{ST}, D_T) = \mathbb{E}_{\mathbf{x}_T \sim P_{X_T}} \log[1 - D_T(\mathbf{x}_T)] + \mathbb{E}_{\mathbf{x}_i \sim P_{X_i}} \log D_T(G_{ST}(E_{S_i}(\mathbf{x}_i))), \quad (3)$$

$$\mathcal{L}_{GAN_T}(E_T, G_{TS}, D_S) = \mathbb{E}_{\mathbf{x}_i \sim P_{X_i}} \log[1 - D_S(\mathbf{x}_i)] + \mathbb{E}_{\mathbf{x}_T \sim P_{X_T}} \log D_S(G_{TS}(E_T(\mathbf{x}_T))). \quad (4)$$

To preserve the semantics of the adapted images, generated by G_{ST} , from source to the target domain, an emotional semantic consistency loss is used, defined by:

$$\mathcal{L}_{ESC}(E_{S_i}, E_T, G_{ST}) = \mathbb{E}_{\mathbf{x}_{S_i} \sim P_{S_i}} KL(F(E_i(\mathbf{x}_{S_i})) \| F(E_T(\mathbf{x}_i^{S_i \rightarrow T}))), \quad (5)$$

where $\mathbf{x}_i^{S_i \rightarrow T} = G_{ST}(E_i(\mathbf{x}_i))$ and $KL(\cdot \| \cdot)$ is the KL divergence between two distributions.

Cycle Reconstruction Pipeline

The cycle reconstruction pipeline is used to learn a unified sentiment latent space and ensure that features of images from different domains preserve the information of their original images. According to (Liu et al. 2017), a VAE-like objective function is used to model the cycle-consistency constraint, defined by:

$$\mathcal{L}_{cyc_i}(E_{S_i}, G_{ST}, E_T, G_{TS}) = \lambda_2 KL(q_T(\mathbf{z}_T | \mathbf{x}_i^{S_i \rightarrow T}) \| p_\eta(z)) - \lambda_3 \mathbb{E}_{\mathbf{z}_T \sim q_T(\mathbf{z}_T | \mathbf{x}_i^{S_i \rightarrow T})} [\log p_{G_{TS}}(\mathbf{x}_i | \mathbf{z}_T)], \quad (6)$$

$$\mathcal{L}_{cyc_T}(E_T, G_{TS}, E_{S_i}, G_{ST}) = \lambda_2 KL(q_i(\mathbf{z}_i | \mathbf{x}_T^{T \rightarrow S_i}) \| p_\eta(z)) - \lambda_3 \mathbb{E}_{\mathbf{z}_i \sim q_i(\mathbf{z}_i | \mathbf{x}_T^{T \rightarrow S_i})} [\log p_{G_{ST}}(\mathbf{x}_T | \mathbf{z}_i)]. \quad (7)$$

The hyper-parameters λ_2 and λ_3 control the weights of the two different objective terms.

Therefore, the augmented MVAE-GAN loss is:

$$\begin{aligned} \mathcal{L}_{MVAE-GAN}(E_{S_1}, \dots, E_{S_M}, E_T, G_{TS}, G_{ST}, D_T, D_S) \\ = \sum_{i=1}^M [\mathcal{L}_{MVAE_i}(E_{S_i}, G_{TS}) + \mathcal{L}_{MVAE_T}(E_T, G_{ST}) + \\ \mathcal{L}_{GAN_i}(E_{S_i}, G_{ST}, D_T) + \mathcal{L}_{cyc_i}(E_{S_i}, G_{ST}, E_T, G_{TS}) + \\ \mathcal{L}_{GAN_T}(E_T, G_{TS}, D_S) + \mathcal{L}_{cyc_T}(E_T, G_{TS}, E_{S_i}, G_{ST})]. \end{aligned} \quad (8)$$

Sentiment Classification with Adapted Images

After the joint learning image reconstruction, image translation, and cycle-reconstruction pipelines, both the source images and adapted images of different domains $S_i (i = 1, 2, \dots, M)$ can be mapped to a same latent representation in a unified sentiment latent space. Meanwhile, the semantic consistency loss in the image translation pipeline ensures the semantic information, *i.e.* the corresponding sentiment labels, is preserved before and after image translation.

Generally, multiple classification models are needed to correspond to different domains for the final classification task. To the contrary, thanks to the unified latent space, the

proposed MSGAN is augmented with a single classifier optimized by minimizing the following cross-entropy loss:

$$\begin{aligned} \mathcal{L}_{task}(E_{S_i}, E_T, F) = & \\ & \lambda_4 \mathbb{E}_{(\mathbf{x}_i, y_i) \sim P_{\mathbf{x}_i}} \sum_{l=1}^L \mathbb{1}_{[l=y_i]} \log(\sigma(F^{(l)}(E_{S_i}(\mathbf{x}_i)))) \\ & + \lambda_5 \mathbb{E}_{(\mathbf{x}_i, y_i) \sim P_{\mathbf{x}_i}} \sum_{l=1}^L \mathbb{1}_{[l=y_i]} \log(\sigma(F^{(l)}(E_T(\mathbf{x}_i^{S_i \rightarrow T})))) \end{aligned} \quad (9)$$

where σ is the SoftMax function, $\mathbb{1}$ is an indicator function and $\mathbf{x}_i^{S_i \rightarrow T} = G_{ST}(E_i(\mathbf{x}_i))$. The hyper-parameters λ_4 and λ_5 control the weights of the two different objective terms.

MSGAN Learning

We jointly solve the learning problems of the $MVAE_i$, $MVAE_T$, GAN_S , GAN_T ($i = 1, 2, \dots, M$), F for the image reconstruction, the image translation, and the cycle-reconstruction pipelines and classification task.

Inheriting from GAN, training the proposed MSGAN framework results in solving a mini-max problem where the optimization aims to find a saddle point. It can be seen as a two player zero-sum game. The first player is a team consisting of the encoders and generators. The second player is a team consisting of the adversarial discriminators. In addition to defeating the second player, the first player has to minimize the MVAE losses and the cycle-consistency losses. We apply an alternating gradient update scheme similar to the one described in (Goodfellow et al. 2014).

The procedure is summarized in Algorithm 1, where $\alpha_{S_1}, \dots, \alpha_{S_M}, \alpha_T, \theta_{ST}, \theta_{TS}, \phi_S, \phi_T$, and φ are the parameters of $E_{S_1}, \dots, E_{S_M}, G_{ST}, G_{TS}, D_S, D_T$, and F .

Experiments

This section presents the experimental analysis of MSGAN. First, the detailed experimental setups are introduced, including the datasets, baselines, and evaluation metrics. Second, the performance of MSGAN and the state-of-the-art algorithms in MDA is reported. Finally, an in-depth analysis on MSGAN, including a parameter sensitivity analysis and an ablation study, is presented.

Datasets

We evaluate our framework on four public datasets including the Flickr and Instagram (**FI**) (You et al. 2016), Artistic (**ArtPhoto**) dataset (Machajdik 2010), **Twitter I** (You et al. 2015) and **Twitter II** (You et al. 2016) datasets. FI is collected by querying with eight sentiment categories as keywords from social websites. 22,700 images are included in the FI dataset. ArtPhoto dataset consists of 806 artistic photographs from a photo sharing site searched by emotion categories. We combine excitement, amusement, awe, contentment as positive images and disgust, anger, fear, sadness as negative images (Mikels et al. 2005). The Twitter I and Twitter II datasets are collected from the social websites and labeled with sentiment polarity (i.e. positive, negative) labels, which consist of 1,269 and 603 images, respectively.

Algorithm 1: Adversarial training procedure of the proposed MSGAN model

Input: Sets of M sources images with sentiment labels $(\mathbf{x}_1, y_1) \in S_1, \dots, (\mathbf{x}_M, y_M) \in S_M$, and target images $\mathbf{x}_T \in \mathbf{X}_T$, the maximum number of steps T

Output: Predicted sentiment label of target image \mathbf{x}_T

```

1 for  $t \leftarrow 1$  to  $T$  do
2   Sample a mini-batch of target images  $\mathbf{x}_T$ ;
3   for  $i \leftarrow 1$  to  $M$  do
4     Sample a mini-batch of source images  $\mathbf{x}_i$ ;
4     /* Updating  $\phi_S, \phi_T$  when fixing
4      $\alpha_{S_i}, \alpha_T, \theta_{ST}, \theta_{TS}$ , and  $\varphi$  */
5     Compute  $E_{S_i}(\mathbf{x}_i; \alpha_{S_i}), G_{ST}(\mathbf{z}_i; \theta_{ST})$  with
5     current  $\alpha_{S_i}, \theta_{ST}$ ;
6     Compute  $E_T(\mathbf{x}_T; \alpha_T), G_{TS}(\mathbf{z}_T; \theta_{TS})$  with
6     current  $\alpha_T, \theta_{TS}$ ;
7     Update  $\phi_S$  and  $\phi_T$  by taking a Adam step on
7     mini-batch loss  $\mathcal{L}_{GAN_i}, \mathcal{L}_{GAN_T}$  in Eq. (3) and
7     Eq. (4), respectively;
7     /* Updating  $\alpha_{S_i}, \alpha_T, \theta_{ST}$  and  $\theta_{TS}$ 
7     when fixing  $\phi_T, \phi_S$  and  $\varphi$  */
8     Update  $\alpha_{S_i}, \alpha_T, \theta_{ST}$  and  $\theta_{TS}$  by taking a
8     Adam step on mini-batch loss  $\mathcal{L}_{MVAE-GAN},$ 
8      $\mathcal{L}_{ESC}$  in Eq. (8) and Eq. (5), respectively;
8     /* Updating  $\alpha_{S_i}, \alpha_T, \varphi$  when
8     fixing  $\phi_T, \phi_S, \theta_{ST}$ , and  $\theta_{TS}$  */
9     Update  $\alpha_{S_i}, \alpha_T, \varphi$  by taking a Adam step on
9     mini-batch loss  $\mathcal{L}_f$  in Eq. (9);
10  end
11 end
12 return  $E_T(\mathbf{x}_T; \alpha_T), F(\mathbf{x}_T; \varphi)$ .
```

Baselines

To the best of our knowledge, MSGAN is the first work on multi-source domain adaptation for classifying visual sentiment. We compare MSGAN with three types of baseline algorithms, termed **Source-only**, **Single-source DA**, and **Multi-source DA**. The **Source-only** methods are trained on source images and directly test their classification performance on the target images. The **Single-source DA** methods include CycleGAN (Zhu et al. 2017) and CycleEmotionGAN (Zhao et al. 2019b). For CycleGAN, we extend the original transfer network, i.e. first adapt the source images to the adapted ones cycle-consistently, and then train the classifier on the adapted source images with the emotion labels from corresponding source images. Since those methods perform in single-source setting, we employ two MDA standards: (1) single-best, i.e. performing adaptation on each single source, and we choose single best performance from three source results; (2) source-combined, i.e. all source domains are combined into a traditional single source. For resnet-simple-extend, our encoder and classifier can be seen as a simply extension of Resnet-18 (He et al. 2016). We train a classifier with the same network on the source combined. Additionally, we introduce two methods,

Table 1: Classification accuracy (%) comparison of the proposed MSGAN method to the state-of-the-art approaches. We separately use FI, Artphoto, Twitter I and Twitter II as the target domain and the rest three datasets as source domains. The best method trained on the source domains is emphasized in bold.

Standards	Method	FI	Artphoto	Twitter I	Twitter II
Source-only	single-best	64.64	61.11	61.41	71.9
	source-combined	65.2	59.87	60.99	70.24
	resnet-simple-extend	65.71	58.02	58.66	69.34
Single-best DA	CycleGAN (Zhu et al. 2017)	63.87	61.11	61.59	70.24
	CycleEmotionGAN (Zhao et al. 2019b)	67.48	61.72	62.38	71.9
Source-combined DA	CycleGAN (Zhu et al. 2017)	66.05	60.49	59.84	71.07
	CycleEmotionGAN (Zhao et al. 2019b)	67.78	61.96	61.77	72.72
Multi-source DA	DCTN (Xu et al. 2018)	65.31	62.34	62.59	66.94
	MDAN (Zhao et al. 2018a)	63.92	60.36	61.59	71.9
	MSGAN(Ours)	70.63	63.58	65.35	72.72
Oracle (Train on sufficiently-labeled target data)		75.24	64.81	68.5	72.72

DCTN (Xu et al. 2018), MDAN (Zhao et al. 2018a) as the **Multi-source DA** baselines. For comparison, we also report the results of an oracle setting, where the classifier is both trained and tested both on the target domain.

Comparison with State-of-the-art

The performance comparisons between the proposed MSGAN model and the state-of-the-art approaches measured by classification accuracy are shown in Table 1. From the results, we have the following observations:

(1) The source-only methods can't handle the *domain shift* or *dataset bias*, where a rough combination is used to transfer all source data. Different domains represent the diverse joint probability distributions of observed images and emotion labels, so a simple combination of all source data for training is harmful to the model performance.

(2) Both adaptation methods, CycleGAN (Zhu et al. 2017) and CycleEmotionGAN (Zhao et al. 2019b), are superior to the source-only methods, while CycleEmotionGAN performs better. This result demonstrates the effectiveness of CycleEmotionGAN for unsupervised domain adaptation in classifying image emotions. Obviously, in source-combined settings, adaption methods transfer the negative samples across the multiple source domains, which indicates multiple sources domain adapter should not be modeled in the same way with single source domain adapter.

(3) The proposed MSGAN achieves superior performance over the state-of-the-art approaches. The improvements benefit from four aspects: image reconstruction, image translation, cycle-reconstruction pipelines and the unified latent space. Firstly, compared with source-combined DA methods, our MSGAN further improves the classification performance, which demonstrates the proposed sentiment latent space can bridge the gap of multiple sources more effectively. Especially, the images in Artphoto dataset, such as abstract oil painting, are far different from other source datasets. However, MSGAN can better distinguish them. Thus, MSGAN further improves the classification performance than other adaption methods. Secondly, compared to single-source DA, MSGAN utilizes more useful information from multiple sources. Thirdly, while other multi-

Table 2: Ablation study on different components in MSGAN. Baseline denotes using GAN with multiple encoders network, MVAE+GAN denotes using both MVAE loss and GAN loss, +CC denotes using the cycle-consistency loss, +ESC denotes using the emotional semantic consistency loss.

Method	FI	Artphoto
baseline	65.71	58.02
MVAE+GAN	67.2	60.99
MVAE+GAN+CC	69.03	62.16
MVAE+GAN+ESC	68.76	60.7
MVAE+GAN+ESC+CC	70.63	63.58

source DA methods only consider the alignment between each source domain and target domain, MSGAN attempts to align all source and target domains jointly. In addition, existing DA methods, such as CycleGAN (Zhu et al. 2017) and MDAN (Zhao et al. 2018a), focus on matching the visual features but ignore the semantic labels. Therefore, they may not well characterize the consistency of visual sentiment. As a result, these methods may not preserve the mappings between visual content and the corresponding sentiment. However, all the DA methods with sentiment consistency loss, such as single-source method (Zhao et al. 2019b) and our multi-source method, significantly outperform the source-only approach, which demonstrates the effectiveness of preserving the sentiments of the adapted images for visual sentiment classification.

(4) The oracle method, *i.e.* testing on the target domain using the model trained on the same domain, achieves the best performance. However, this model is trained using the ground truth sentiment labels from the target domain, which are actually unavailable in unsupervised domain adaptation.

Visualization We visualize the results of image-space adaptation from Artphoto, FI, Twitter II to Twitter I in Figure 4. We can see that with our final proposed MSGAN method (d), the styles of the images are similar to FI while the emotion semantic information is well preserved.




Figure 4: Visualization of image translation. From left to right are: (a) original source images, (b) MVAE+GAN, (c) MVAE+GAN+CC, (d) MVAE+GAN+CC+ESC, (e) target TwitterI image. The source images of (1)(2)(3) are from Artphoto, FI, and TwitterII datasets respectively. With adding the components of MSGAN, the adapted images look more visually similar to the target images than the source images.




Figure 5: Sensitivity of λ in Eq. (1)(2)(6)(7) of the proposed MSGAN model.

Ablation Study

We incrementally investigate the effectiveness of different components in MSGAN. The results are shown in Table 2. We can observe that: (1) MVAE+GAN can obtain better performance by making different adapted domains more closely aggregated; (2) adding the cycle-consistency loss could further improve the accuracy, again demonstrating the effectiveness of the unified sentiment latent space; (3) ESC loss also contributes to the visual sentiment adaptation task; (4) the modules are orthogonal to each other, since adding each one of them does not introduce performance degradation.

Parameter Sensitivity. we analyze the impact of the hyper-parameter values in Eq. (1)(2)(6)(7) to the sentiment

classification accuracy. For different weight values on the negative log likelihood terms (*i.e.*, λ_1, λ_3), we computed the achieved classification accuracy over different weight values on the KL terms (*i.e.*, λ_0, λ_2) for both $\{\text{Artphoto, Twitter I, Twitter II}\} \rightarrow \text{FI}$ and $\{\text{FI, Twitter I, Twitter II}\} \rightarrow \text{Artphoto}$. The results are reported in Figure 5. We can observe that, in general, a larger weight value on the negative log likelihood terms yields a better result. We also find that setting the weights of the KL terms to 0.1 can result in consistently good performance. We hence set $\lambda_0 = \lambda_2 = 10$, $\lambda_1 = \lambda_3 = 0.1$.

Conclusion

In this paper, we tackle the problem of multi-source domain adaptation (MDA) in visual sentiment classification. A novel framework, termed Multi-source Sentiment Generative Adversarial Network (MSGAN), is proposed to learn a unified sentiment latent space such that data from both the source and target domains share a similar distribution. Such mappings are learned via three pipelines, including image reconstruction, image translation, and cycle reconstruction. Extensive experiments conducted on four benchmark datasets demonstrate that MSGAN significantly outperforms the state-of-the-art MDA approaches for visual sentiment classification. For further studies, we plan to extend the MSGAN model to other image emotion recognition tasks, such as emotion distribution learning (Zhao et al. 2017). We will also investigate methods that can improve the

intrinsic problem of GAN variants in training stability.

Acknowledgement

NExT++ research is supported by the National Research Foundation, Prime Minister's Office, Singapore under its IRC@SG Funding Initiative.

References

- Alaei, A. R.; Becken, S.; and Stantic, B. 2019. Sentiment analysis in tourism: capitalizing on big data. *JTR* 58(2):175–191.
- Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Wortman, J. 2008. Learning bounds for domain adaptation. In *NIPS*, 129–136.
- Borth, D.; Ji, R.; Chen, T.; Breuel, T.; and Chang, S.-F. 2013. Large-scale visual sentiment ontology and detectors using adjective noun pairs. In *ACM MM*, 223–232.
- Ghifary, M.; Bastiaan Kleijn, W.; Zhang, M.; and Balduzzi, D. 2015. Domain generalization for object recognition with multi-task autoencoders. In *ICCV*, 2551–2559.
- Ghifary, M.; Kleijn, W. B.; Zhang, M.; Balduzzi, D.; and Li, W. 2016. Deep reconstruction-classification networks for unsupervised domain adaptation. In *ECCV*, 597–613.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *NIPS*, 2672–2680.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Katsurai, M., and Satoh, S. 2016. Image sentiment analysis using latent correlations among visual, textual, and sentiment views. In *ICASSP*, 2837–2841.
- Liu, M.-Y., and Tuzel, O. 2016. Coupled generative adversarial networks. In *NIPS*, 469–477.
- Liu, M.-Y.; Breuel, T.; Kautz; and Jan. 2017. Unsupervised image-to-image translation networks. In *NIPS*, 700–708.
- Machajdik, Jana, H. A. 2010. Affective image classification using features inspired by psychology and art theory. In *ACM MM*, 83–92.
- Mansour, Y.; Mohri, M.; and Rostamizadeh, A. 2009. Domain adaptation with multiple sources. In *NIPS*, 1041–1048.
- Mikels, J. A.; Fredrickson, B. L.; Larkin, G. R.; Lindberg, C. M.; Maglio, S. J.; and Reuter-Lorenz, P. A. 2005. Emotional category data on images from the international affective picture system. *BRM* 37(4):626–630.
- Peng, K.-C.; Sadovnik, A.; Gallagher, A.; and Chen, T. 2015. A mixed bag of emotions: Model, predict, and transfer emotion distributions. In *CVPR*, 860–868.
- Peng, X.; Bai, Q.; Xia, X.; Huang, Z.; Saenko, K.; and Wang, B. 2018. Moment matching for multi-source domain adaptation. *arXiv:1812.01754*.
- Riemer, M.; Cases, I.; Ajemian, R.; Liu, M.; Rish, I.; Tu, Y.; and Tesauro, G. 2019. Learning to learn without forgetting by maximizing transfer and minimizing interference. *ICLR*.
- Sun, S.-L., and Shi, H.-L. 2013. Bayesian multi-source domain adaptation. In *ICMLC*, volume 1, 24–28.
- Sun, B.; Feng, J.; and Saenko, K. 2017. Correlation alignment for unsupervised domain adaptation. In *Domain Adaptation in Computer Vision Applications*. 153–171.
- Torralba, A., and Efros, A. A. 2011. Unbiased look at dataset bias. In *CVPR*, 1521–1528.
- Xu, R.; Chen, Z.; Zuo, W.; Yan, J.; and Lin, L. 2018. Deep cocktail network: Multi-source unsupervised domain adaptation with category shift. In *CVPR*, 3964–3973.
- Yang, J.; She, D.; Sun; and Ming. 2017. Joint image emotion classification and distribution learning via deep convolutional neural network. In *IJCAI*, 3266–3272.
- Yang, J.; She, D.; Lai, Y.-K.; Rosin, P. L.; and Yang, M.-H. 2018a. Weakly supervised coupled networks for visual sentiment analysis. In *CVPR*, 7584–7592.
- Yang, J.; She, D.; Lai, Y.; and Yang, M.-H. 2018b. Retrieving and classifying affective images via deep metric learning. In *AAAI*.
- Yang, J.; She, D.; Sun, M.; Cheng, M.-M.; Rosin, P. L.; and Wang, L. 2018c. Visual sentiment prediction based on automatic discovery of affective regions. *IEEE TMM* 20(9):2513–2525.
- You, Q.; Luo, J.; Jin, H.; and Yang, J. 2015. Robust image sentiment analysis using progressively trained and domain transferred deep networks. In *AAAI*, 381–388.
- You, Q.; Luo, J.; Jin, H.; and Yang, J. 2016. Building a large scale dataset for image emotion recognition: The fine print and the benchmark. In *AAAI*, 308–314.
- You, Q.; Jin, H.; and Luo, J. 2017. Visual sentiment analysis by attending on local image regions. In *AAAI*, 231–237.
- Zhao, S.; Yao, H.; Gao, Y.; Ji, R.; and Ding, G. 2017. Continuous probability distribution prediction of image emotions via multi-task shared sparse regression. *IEEE TMM* 19(3):632–645.
- Zhao, H.; Zhang, S.; Wu, G.; Moura, J. M.; Costeira, J. P.; and Gordon, G. J. 2018a. Adversarial multiple source domain adaptation. In *NIPS*, 8559–8570.
- Zhao, S.; Yao, H.; Gao, Y.; Ding, G.; and Chua, T.-S. 2018b. Predicting personalized image emotion perceptions in social networks. *IEEE TAFMC* 9(4):526–540.
- Zhao, S.; Zhao, X.; Ding, G.; and Keutzer, K. 2018c. Emotiongan: unsupervised domain adaptation for learning discrete probability distributions of image emotions. In *ACM MM*, 1319–1327.
- Zhao, S.; Li, B.; Yue, X.; Gu, Y.; Xu, P.; Hu, R.; Chai, H.; and Keutzer, K. 2019a. Multi-source domain adaptation for semantic segmentation. In *NIPS*, 7285–7298.
- Zhao, S.; Lin, C.; Xu, P.; Zhao, S.; Guo, Y.; Krishna, R.; Ding, G.; and Keutzer, K. 2019b. Cycleemotiongan: Emotional semantic consistency preserved cyclegan for adapting image emotions. In *AAAI*, 2620–2627.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2242–2251.
- Zhuo, J.; Wang, S.; Zhang, W.; and Huang, Q. 2017. Deep unsupervised convolutional domain adaptation. In *ACM MM*, 261–269.