

Bridge Then Begin Anew: Generating Target-Relevant Intermediate Model for Source-Free Visual Emotion Adaptation

Jiankun Zhu¹, Sicheng Zhao², Jing Jiang¹, Wenbo Tang³, Zhaopan Xu¹, Tingting Han⁴, Pengfei Xu³, Hongxun Yao^{1*}

¹School of Computer Science and Technology, Harbin Institute of Technology, Harbin, China

²BNRist, Tsinghua University, Beijing, China

³NavInfo, Beijing, China

⁴Hangzhou Dianzi University, Hangzhou, China

{23s003059, 23s003045, 20b903054}@stu.hit.edu.cn, h.yao@hit.edu.cn, schzhao@gmail.com, ttinghan@hdu.edu.cn, {xupengfei11850, tangwenbo11883}@navinfo.com

Abstract

Visual emotion recognition (VER), which aims at understanding humans' emotional reactions toward different visual stimuli, has attracted increasing attention. Given the subjective and ambiguous characteristics of emotion, annotating a reliable large-scale dataset is hard. For reducing reliance on data labeling, domain adaptation offers an alternative solution by adapting models trained on labeled source data to unlabeled target data. Conventional domain adaptation methods require access to source data. However, due to privacy concerns, source emotional data may be inaccessible. To address this issue, we propose an unexplored task: source-free domain adaptation (SFDA) for VER, which does not have access to source data during the adaptation process. To achieve this, we propose a novel framework termed Bridge then Begin Anew (BBA), which consists of two steps: domain-bridged model generation (DMG) and target-related model adaptation (TMA). First, the DMG bridges cross-domain gaps by generating an intermediate model, avoiding direct alignment between two VER datasets with significant differences. Then, the TMA begins training the target model anew to fit the target structure, avoiding the influence of source-specific knowledge. Extensive experiments are conducted on six SFDA settings for VER. The results demonstrate the effectiveness of BBA, which achieves remarkable performance gains compared with state-of-the-art SFDA methods and outperforms representative unsupervised domain adaptation approaches.

Introduction

With the rapid development of social networks, people have become used to posting images to express their feelings (Zhao et al. 2017). Visual emotion recognition (VER) aims to identify human emotions elicited by images (Zhao et al. 2022b), attracting increasing attention and playing an essential role in various applications, such as depression detection (Bokolo and Liu 2023) and opinion mining (Razali et al. 2021). Advanced technologies based on supervised deep neural networks have been proposed to improve the VER performance (Deng et al. 2024; Cen et al. 2024). Existing methods mainly follow a supervised pipeline requiring

sufficient emotion annotations. However, labeling reliable large-scale datasets for VER tasks is challenging in practical applications because of the intrinsic properties of emotions, such as subjectivity, complexity, and ambiguity (Zhao et al. 2023). To reduce the annotation burden, much attention has been devoted to unsupervised domain adaptation (UDA) for VER, which enables the models trained on labeled source data to generalize well on the unlabeled target data (Zhao et al. 2019b). However, existing UDA methods require access to the source data during adaptation (Li, Guo, and Ge 2023; Zhao et al. 2024; Jiang et al. 2024). Regarding privacy and security concerns, source-free domain adaptation (SFDA) has attracted much interest in dealing with the situation where labeled source data is unavailable (Li et al. 2024).

On the one hand, although prior SFDA methods are effective for standard classification tasks, they face specific challenges and performance decreases when directly applied to VER. This is because of the large affective gap (Zhao et al. 2014) between VER datasets, which arises from annotator variations and the scope of data collection. The distribution gap significantly affects the feature alignment, which in turn reduces the confidence of pseudo-labels. As shown in Figure 1 (a), we use SHOT (Liang, Hu, and Feng 2020), a typical SFDA method, to generate pseudo-labels for VER datasets EmoSet \rightarrow FI and standard classification datasets Real World \rightarrow Clipart, respectively. The results show that the classification scores for emotional pseudo-labels are relatively low, which increases the risk of misalignment between features and labels, and can lead to a confirmation bias due to the accumulation of errors (Ding et al. 2023). To address this issue, bridging the inter-domain variation to generate reliable pseudo-labels is necessary.

On the other hand, traditional SFDA methods are fine-tuned on the source model. However, the unclear class features of the VER dataset make the source model lack robustness. For a better illustration, as shown in Figure 1 (b), we train two ResNet-101 models (He et al. 2016) with the same number of categories on EmoSet and Clipart from the Office-Home (Venkateswara et al. 2017), respectively. The results indicate that the VER dataset has lower interclass distinction and higher intraclass variability. Consequently, the source model will overfit on noise, which is unrelated

*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

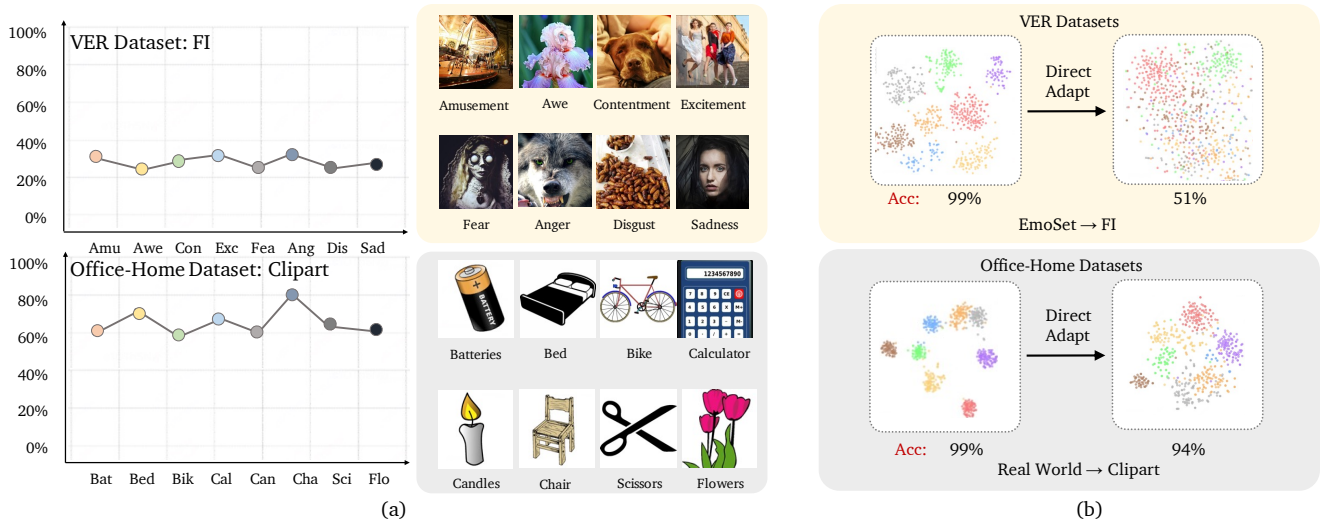


Figure 1: Illustration of challenges for SFDA-VER tasks. (a) shows the pseudo-label confidence generated by SHOT on the VER dataset and the standard classification dataset OfficeHome. Differently colored dots indicate prediction probabilities for images in different categories. (b) shows the distributions of ResNet-101 features on the VER and OfficeHome datasets.

to the discriminative class features. As a result, when directly adapted to the target dataset, *i.e.*, EmoSet \rightarrow FI and Real World \rightarrow Clipart, without further improvement, standard classification model tends to outperform VER model by a large margin. This suggests that the VER dataset lacks explicit class characterization, making the source domain model prone to overfitting. To eliminate the effects of source domain noise, it is necessary for the target model to learn from the target data itself, instead of fine-tuning from the source model.

We summarize the challenges above: the large affective gap leads to incorrect pseudo-labels, and lack of clarity in the distribution of the VER dataset can lead to overfitting the source domain model. To address these two challenges, we propose a novel SFDA framework for VER, termed Bridge then Begin Anew (BBA). BBA contains two steps: Domain-bridged model generation (DMG) and target-related model adaptation (TMA). To improve the confidence of pseudo-labels, DMG generates a bridge model to align emotional data across various domains, thus avoiding the challenges associated with directly fine-tuning emotional data. Moreover, we introduce a clustering-based pseudo-label post-processing and masking strategy to constrain the pseudo-label distribution and explore richer semantic contexts. Then, to eliminate the effects of overfitting the source model, TMA starts training anew, which allows the target model to learn feature relations independently, thus extending the exploration of the target-specific features. Furthermore, in order to better learn about the emotional features, we introduce polarity constraints to enhance the target model’s discriminative ability for emotion categories.

The main contributions of our work are outlined below:

- We propose a new task, *i.e.*, source-free domain adaptation for visual emotion recognition (SFDA-VER), which focuses on the cross-domain transfer of VER without ac-

cessing source data during adaptation.

- We propose a two-step framework, comprising DMG and TMA, to improve the reliability of pseudo-labels and eliminate the effects of source model overfitting.
- We conduct experiments on six SFDA-VER settings, and the results show that our approach outperforms existing state-of-the-art methods by an average of +3.03%.

Related Works

Visual Emotion Recognition

With the advent of deep learning, convolutional neural networks (CNNs) have shown great success in many machine learning tasks, including VER (Zhao et al. 2023). In the beginning, each region in an image is treated equally, and global features are extracted (You et al. 2015; Zhu et al. 2017; Yang et al. 2018a; Rao, Li, and Xu 2020). Later, based on the theory that emotional content is always present in certain vital regions, more studies have focused on extracting local features and exploring the associations between these features (Yang et al. 2018b; Zhao et al. 2019a; Rao et al. 2019; Yao et al. 2019; Zhang et al. 2020). Recently, with the rise of language-supervised pre-training models, some studies have introduced CLIP to the VER task (Deng et al. 2022, 2024; Cen et al. 2024).

Unsupervised Domain Adaptation for VER

Due to the significant domain bias in VER datasets, Zhao et al. (2019b) introduced CycleEmotionGAN, an unsupervised approach for cross-domain emotion adaptation using emotional consistency loss. Zhao et al. (2022a) enhanced this with multi-scale similarity and improved emotional consistency. Considering the privacy constraints that limit direct access to source data for emotion analysis, our study presents a new source-free domain adaptation for VER

In order to improve interclass distinction and reduce intra-class variability in VER features, we consider mitigating the pseudo-label mismatch problem caused by domain differences. Inspired by SHOT (Liang, Hu, and Feng 2020), we propose a fused distance clustering method as a label optimization strategy, which imposes additional constraints on the model outputs. Specifically, we apply weighted k-means clustering to compute the centroid of each class within the target domain:

$$c_k^{(0)} = \frac{\sum_{x_t \in X_t} \sigma_k(\phi_s(x_t)) g_s(x_t)}{\sum_{x_t \in X_t} \sigma_k(\phi_s(x_t))}, \quad (1)$$

where $\sigma_k(a) = \frac{\exp(a_k)}{\sum_i \exp(a_i)}$ denotes the k -th element in the softmax output of a K -dimensional vector a , $\phi_s = f_s \circ g_s$ indicates source model, g_s is the source feature extractor and $X_t = \{x_t^i\}_{i=1}^{N_T}$. After calculating the centroid of each class, we assign each sample to the nearest class as follows:

$$\hat{y}_t' = \arg \min_k D_f(g_s(x_t), c_k^{(0)}), \quad (2)$$

where $D_f(P, Q)$ is the distance measure used to calculate the distance between each sample and centroid. Considering different distance metrics focus on various parts, we adopt a fused distance metric, *i.e.*, we calculate Euclidean (D_{eu}), Cosine (D_{cos}), and Manhattan (D_{man}) distances between two features:

$$D_f(P, Q) = D_{eu}(P, Q) + D_{cos}(P, Q) + D_{man}(P, Q). \quad (3)$$

Then, we update the category centroids as follows, where 1 is the indicator function:

$$c_k^{(1)} = \frac{\sum_{x_t \in X_t} 1(\hat{y}_t' = k) g_s(x_t)}{\sum_{x_t \in X_t} 1(\hat{y}_t' = k)}. \quad (4)$$

To reduce intra-class variability, we consider the difference between VER images and traditional classification images: unlike standard classification tasks, where different blocks in an image correspond to different semantics, different blocks in VER images are usually associated with a unified emotion label. As such, we propose a masking-enhanced framework to leverage emotional consistency across different contexts. As shown in Figure 3, models are encouraged to focus on local features to more effectively extract emotionally consistent features for better classification. Our method uses the random mask technique. Specifically, given a 2D image I of size $H \times W$, and the number of blocks to be masked n , which is randomly selected from a given range or distribution, the mask M is defined as:

$$M = \{m_{ij}\}_{i=1, j=1}^{H, W}, \text{ with } \sum_{i=1}^H \sum_{j=1}^W m_{ij} = nb^2, \quad (5)$$

where $m_{ij} \in \{0, 1\}$ is the mask value at position (i, j) , and if $(i, j) \in$ randomly selected blocks, $m_{ij} = 1$. Each block has a size of $b \times b$. The masked image x_t^M is obtained by the element-wise multiplication of M and x_t , denoted as:

$$x_t^M = M \odot x_t. \quad (6)$$

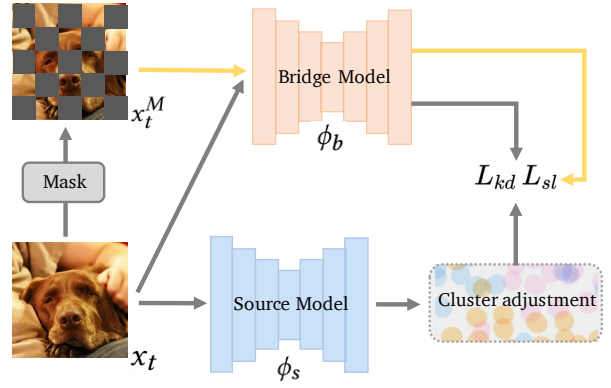


Figure 3: Illustration of our masking strategy for feature enhancement. The bridge model ϕ_b , guided by the source model ϕ_s , predicts both original and masked target domain images to compute self-labeling loss \mathcal{L}_{sl} in Eq. (8) and distillation loss \mathcal{L}_{kd} in Eq. (9), respectively.

After masking, the classifier is forced to use sparse residual cues, to learn local and global affective information. Moreover, to provide a consistent and stable guidance for bridge model, our source model adopts the exponential moving average (EMA) strategy for parameter updates (Tarvainen and Valpola 2017). The bridge model $\phi_b = f_b \circ g_b$ inherits parameters from the pre-trained source model ϕ_s . To better generalize to the target domain, the feature extractor g_b is updated while the classifier f_b is fixed. Both g_b and g_s are updated at this step, g_b is updated by backpropagation as described in Eq. (10), and g_s is updated by EMA:

$$\theta_{t+1}^s \leftarrow \alpha_e \theta_t^s + (1 - \alpha_e) \theta_t^b, \quad (7)$$

where t is the training step, α_e is the smoothing factor, θ_t^s , θ_t^b are the parameters of g_s , g_b . Then, we train the masked image together with the original target domain image to participate in the computation of the self-labeling loss and the knowledge distillation loss:

$$\begin{aligned} \mathcal{L}_{sl} = & -\mathbb{E}_{x_t \in X_t} \sum_{k=1}^K 1(\hat{y}_t' = k) \log \sigma(\phi_b(x_t)) \\ & -\mathbb{E}_{x_t \in X_t} \sum_{k=1}^K 1(\hat{y}_t' = k) \log \sigma(\phi_b(x_t^M)), \end{aligned} \quad (8)$$

$$\mathcal{L}_{kd} = -\mathbb{E}_{x_t \in X_t} \hat{y}_t^s \log \phi_b(x_t) - \mathbb{E}_{x_t \in X_t} \hat{y}_t^s \log \phi_b(x_t^M), \quad (9)$$

where $\hat{y}_t^s = D_f(g_s(x_t), c_k)$ denotes the soft pseudo output of the teacher feature extractor g_s after clustering, and c_k denotes the final centroid of each class. The total loss function for the DMG is expressed as:

$$\mathcal{L}_{dmg} = \mathcal{L}_{kd} + \lambda \mathcal{L}_{sl}, \quad (10)$$

where λ is the weight for \mathcal{L}_{sl} , controlling the importance of \mathcal{L}_{sl} during DMG step.

Target-related Model Adaptation

As stated earlier, noise in the source data that is unrelated to the discriminative features can cause the source model

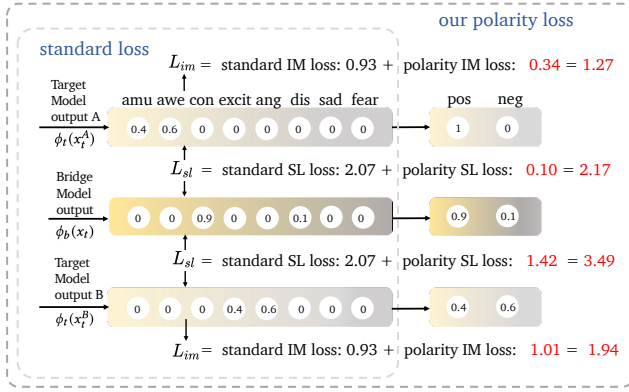


Figure 4: Numerical examples to validate the effectiveness of the proposed emotion polarity loss. The addition of emotion polarity loss to IM loss \mathcal{L}_{im} and SL loss \mathcal{L}_{sl} allows the model to differentiate samples more finely.

to overfit, potentially constraining its representation of target features. To overcome it, we argue that the target model should be given sufficient latitude to explore the target feature space from scratch, enabling it to identify novel features. Following recent advances in self-distillation (Kim et al. 2021; Laine and Aila 2016; Liang et al. 2022), we train the target model ϕ_t from scratch using target data alone, guided by the bridge model ϕ_b , which produces more reliable pseudo-labels for improved supervision:

$$\mathcal{L}_{align} = \mathbb{E}_{x_t \in X_t} D(\phi_b(x_t), \phi_t(x_t)), \quad (11)$$

where $D(\cdot)$ is a function that measures the distance between two emotion samples. In this paper, $D(\cdot)$ represents the KL divergence. ϕ_b updates pseudo-labels by mixing ϕ_t outputs, where α_t is a momentum hyper-parameter:

$$\phi_b(x_t) \leftarrow \alpha_t \phi_b(x_t) + (1 - \alpha_t) \phi_t(x_t), \forall x_t \in X_t. \quad (12)$$

In addition, we note that emotions naturally contain polarized information; categories with matching emotion polarities (positive or negative) show closer associations. To explore global and hierarchical polarity features when learning the target structure, we combine emotion polarity with the loss of self-labeling (SL) and the loss of information maximization (IM). Specifically, the distribution of emotion polarity is detailed as follows:

$$p_{pos} = \sum_{k \in C_{pos}} \sigma(\phi_t(x_t))_k, p_{neg} = \sum_{k \in C_{neg}} \sigma(\phi_t(x_t))_k, \quad (13)$$

where C_{pos} and C_{neg} represents the positive and negative emotion categories. p_{pos} and p_{neg} symbolize the aggregate predicted likelihood across all positive and negative emotion categories for each sample, and they are collectively referred to p_e . Then, we propose a polarized IM loss that incorporates the emotion polarity distribution:

$$\mathcal{L}_{im} = -\mathbb{E}_{x_t \in X_t} (\mathcal{L}_e + \mathcal{L}_d), \quad (14)$$

$$\mathcal{L}_e = \mathbb{E}_{k \in K} [p_k \log(p_k)] + \mathbb{E}_{e \in \{pos, neg\}} [p_e \log(p_e)], \quad (15)$$

$$\mathcal{L}_d = -\mathbb{E}_{k \in K} [p_k \log \mathbb{E}_k(p_k)] - \mathbb{E}_{e \in \{pos, neg\}} [p_e \log \mathbb{E}_e(p_e)], \quad (16)$$

where \mathcal{L}_e and \mathcal{L}_d denote the entropy and diversity loss respectively. p_k represents the predicted probability of the k^{th} class, represented by $\sigma(\phi_t(x_t))_k$. Moreover, we also incorporate the emotion polarity distribution into the SL loss, resulting in the following polarized SL loss:

$$\mathcal{L}_{sl} = -\mathbb{E}_{x_t \in X_t} \sum_{k=1}^K \hat{y}_t \log p_k - \mathbb{E}_{x_t \in X_t} \sum_{e \in \{pos, neg\}} \hat{y}_e \log p_e, \quad (17)$$

where $\hat{y}_e = \arg \max p_e$ denotes the polarity pseudo-label. To illustrate the necessity and effectiveness of the proposed emotion polarity loss, we cite a simple example in Figure 4. We can see that the standard IM loss \mathcal{L}_{im} gives the same loss value for different outputs $\phi_t(x_t^A)$ and $\phi_t(x_t^B)$, which are both 0.93. This is because the standard IM loss only focuses on the entropy of the output vector and overlooks the category relations. However, the polarity loss, which has values of 0.34 and 1.01, imposes an extra constraint on the model output. Furthermore, for the standard SL loss \mathcal{L}_{sl} , it is unchanged between the outputs $\phi_t(x_t^A)$, $\phi_b(x_t)$ and $\phi_t(x_t^B)$, $\phi_b(x_t)$, remaining at 2.07. After including the polarity loss, which are 0.10 and 1.42, the total loss differs. This highlights how the polarity loss considers an additional correlation across categories.

Overall, the total loss for the TMA phase can be expressed as:

$$\mathcal{L}_{tma} = \mathcal{L}_{align} + \mathcal{L}_{pol}, \quad (18)$$

where $\mathcal{L}_{pol} = \gamma \mathcal{L}_{sl} + \delta \mathcal{L}_{im}$, γ and δ are the weights for \mathcal{L}_{sl} and \mathcal{L}_{im} .

BBA Learning

For training, in the DMG step, we use clustered post-processed pseudo-labels to guide the training of both the target and masked target data. This approach aims to improve interclass separability and intraclass compactness, thereby generating more reliable pseudo-labels. The loss function is \mathcal{L}_{dmg} in Eq. (10). In the TMA step, to avoid the effects of source domain model overfitting, the target model is randomly initialized and guided by DMG-corrected pseudo-labels to learn the target structure. The loss function is \mathcal{L}_{tma} in Eq. (18). For inference, we only use the target model, and the model size is maintained the same as the source model.

Experiments

Experimental Settings

Extensive experiments are conducted on four VER datasets: FI (You et al. 2016), EmoSet (Yang et al. 2023), and Art-Photo (Machajdik and Hanbury 2010) are categorized into Mikels' eight emotion categories (Zhao et al. 2016); In contrast, Emotion6 is classified according to the Ekman model's six emotion categories (Peng et al. 2015). We employ six SFDA-VER settings for experiments: FI \leftrightarrow EmoSet, FI \leftrightarrow ArtPhoto, and FI \leftrightarrow Emotion6, with the latter considered as binary classification tasks due to their distinct categorical frameworks.

Method	SF	EmoSet \rightarrow FI							FI \rightarrow EmoSet						
		Acc	macro avg			weighted avg			Acc	macro avg			weighted avg		
			P	R	F1	P	R	F1		P	R	F1	P	R	F1
Source only	-	51.57	48.84	51.14	48.71	54.02	51.57	51.66	50.87	54.40	50.30	49.77	54.07	50.87	50.04
Oracle	-	67.32	62.71	59.44	60.55	66.50	67.32	66.55	71.96	72.60	72.40	72.41	72.12	71.96	71.95
SHOT	✓	50.38	47.38	50.78	47.74	53.25	50.38	50.68	53.31	54.01	53.40	52.28	54.42	53.31	52.54
SHOT++	✓	50.42	48.00	52.17	48.25	55.30	50.42	51.31	54.15	54.17	54.68	53.66	54.70	54.15	53.71
G-SFDA	✓	50.47	47.95	52.09	48.22	55.22	50.47	51.36	54.40	55.36	54.29	53.66	55.32	54.40	53.81
DaC	✓	50.24	47.41	50.42	47.55	53.57	50.24	50.62	51.55	52.13	51.69	51.49	52.21	51.55	51.44
AaD	✓	49.90	48.01	51.96	47.74	55.71	49.90	50.88	52.68	52.54	53.46	52.37	53.33	52.68	52.40
DINE	✓	55.25	51.87	54.67	52.19	58.11	55.25	55.70	56.36	59.14	56.01	55.42	59.04	56.36	55.66
TPDS	✓	51.45	48.52	51.51	48.65	54.21	51.45	51.64	55.88	58.52	55.64	55.14	58.36	55.88	55.27
BBA (ours)	✓	57.36	52.69	55.67	53.27	59.92	57.36	57.76	58.76	60.32	58.83	58.52	59.62	58.76	58.19
MCC	×	56.58	52.65	54.49	52.74	59.19	56.58	57.16	56.26	58.76	55.81	55.90	58.38	56.26	55.97
ELS	×	56.42	53.82	53.94	52.80	59.16	56.42	56.97	56.41	59.10	56.50	56.49	58.61	56.41	56.23
MIC	×	56.26	52.62	54.61	52.53	59.85	56.26	57.41	54.31	54.26	54.77	53.74	54.73	54.31	53.78

Table 1: Performance comparison between BBA and state-of-the-art approaches on FI \leftrightarrow EmoSet.

Method	SF	Emotion6 \rightarrow FI							FI \rightarrow Emotion6						
		Acc	macro avg			weighted avg			Acc	macro avg			weighted avg		
			P	R	F1	P	R	F1		P	R	F1	P	R	F1
Source only	-	68.18	69.55	73.55	67.16	78.70	68.18	69.57	70.66	70.96	72.78	70.14	75.11	70.66	71.26
Oracle	-	89.20	88.06	85.21	86.45	89.03	89.20	88.99	82.99	81.90	80.54	67.16	82.80	82.99	82.81
SHOT	✓	67.22	68.97	72.80	66.26	78.24	67.22	68.63	67.90	70.15	71.52	67.72	74.93	67.90	68.41
SHOT++	✓	68.71	70.44	74.56	67.79	79.74	68.71	70.06	67.25	69.79	71.05	67.09	74.65	67.25	67.73
G-SFDA	✓	65.39	68.64	72.10	64.71	78.37	65.39	66.75	66.89	69.62	70.81	66.75	74.53	66.89	67.35
DaC	✓	64.43	67.66	70.94	63.73	77.33	64.43	65.82	65.93	69.80	70.58	65.88	75.01	65.93	66.24
AaD	✓	70.04	71.04	75.39	68.98	80.06	70.04	71.36	67.84	69.92	71.33	67.64	74.65	67.84	68.37
DINE	✓	71.71	71.31	75.76	70.30	79.79	71.71	72.99	<u>71.32</u>	<u>71.28</u>	<u>73.15</u>	<u>70.72</u>	74.97	<u>71.32</u>	<u>71.91</u>
TPDS	✓	72.42	72.07	76.68	71.06	80.55	72.42	73.67	63.23	69.04	68.92	63.23	74.65	63.23	63.18
BBA (ours)	✓	78.12	73.84	75.64	74.55	80.59	78.12	78.52	73.23	72.49	74.38	72.43	76.15	73.23	73.77
MCC	×	74.04	70.54	73.70	<u>71.19</u>	77.36	<u>74.04</u>	<u>74.96</u>	68.68	68.35	69.95	67.92	72.26	68.68	69.32
ELS	×	73.12	70.57	74.36	70.85	78.00	73.13	74.24	71.32	70.14	71.71	70.25	73.62	<u>71.32</u>	71.85
MIC	×	71.94	71.49	75.97	70.52	79.94	71.94	73.21	65.21	68.56	69.46	65.13	73.58	<u>65.21</u>	65.59

Table 2: Performance comparison between BBA and state-of-the-art approaches on FI \leftrightarrow Emotion6.

We compare BBA with the following baselines: **1) Source only**, which refers to a basic approach where the model is trained on the source domain and directly applied to the target domain. **2) SFDA methods**, which include SHOT (Liang, Hu, and Feng 2020), SHOT++ (Liang et al. 2021), G-SFDA (Yang et al. 2021), AaD (Yang et al. 2022), DaC (Zhang et al. 2022b), DINE (Liang et al. 2022), and TPDS (Tang et al. 2024). **3) UDA methods**, which can utilize source domain data during the adaptation process, including MCC (Jin et al. 2020), ELS (Zhang et al. 2022a), and MIC (Hoyer et al. 2023). **4) Oracle**, which is the ideal scenario where the model is trained and tested within the target domain, representing the upper bound of performance.

In BBA, the EMA parameter α_e of the source network is set to 0.999. The loss weights λ , γ , and δ are set to 0.9, 1, and 0.3 for all six settings, respectively. More information about evaluation metrics and implementation details can be found in the supplementary material.

Comparisons with State-of-the-art Methods

The results on FI \leftrightarrow EmoSet and FI \leftrightarrow Emotion6 are comprehensively detailed in Table 1 and Table 2, respectively. More results for the other two settings, consisting of the

Method	Accuracy(%)					
	E \rightarrow F	F \rightarrow E	A \rightarrow F	F \rightarrow A	E6 \rightarrow F	F \rightarrow E6
baseline	51.59	52.63	27.40	33.95	66.81	70.42
+cluster	51.66	53.22	28.08	34.57	73.86	72.28
+mask	54.50	57.00	28.73	36.42	76.65	73.05
+ \mathcal{L}_{align}	56.79	57.38	28.91	37.04	78.12	73.23
+ \mathcal{L}_{pol}	57.36	58.76	29.63	37.65	-	-

Table 3: Accuracy of different variants of our BBA on six SFDA settings. The methods are all added to the previous one. E, F, E6 means EmoSet, FI, and Emotion6 dataset, respectively.

small ArtPhoto dataset, can be found in the supplementary material. Compared to the prior state-of-the-art SFDA methods, BBA achieves average improvements ranging from +2.26% to +7.17% on FI \leftrightarrow EmoSet and from +3.81% to +11.30% on FI \leftrightarrow Emotion6. The increase is due to our ‘bridge then begin anew’ training framework, which improves the pseudo-labeling accuracy while learning from the target data distribution. Despite the theoretical advantage of the UDA method in achieving better results by having ac-

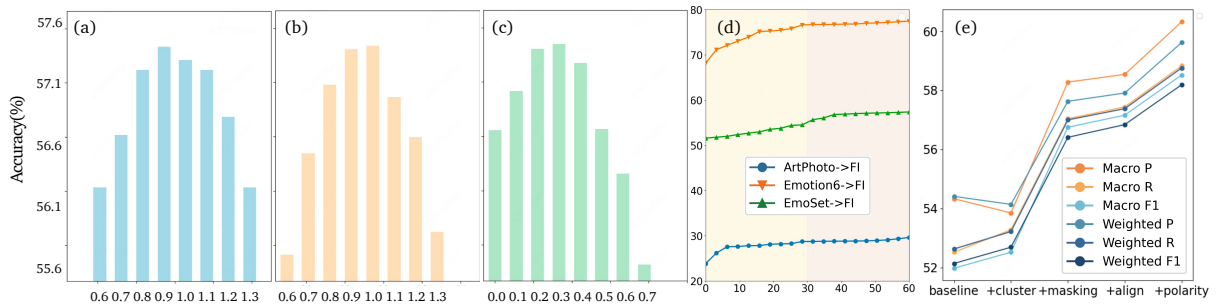


Figure 5: Parameter analysis on EmoSet \rightarrow FI (best viewed in color). (a): Analysis on λ . (b): Analysis on γ . (c): Analysis on δ . (d): the accuracy during the training process. (e): Ablation study on FI \rightarrow EmoSet.




Figure 6: t-SNE visualizations on FI \rightarrow EmoSet. Different colors represent different classes.

cess to the source domain data during the adaptation process, BBA outperforms the UDA method in accuracy, with average improvements of +1.57% to +2.78% on FI \leftrightarrow EmoSet and +3.00% to +7.10% on FI \leftrightarrow Emotion6, due to its targeted solutions to the challenges inherent in emotion data.

Ablation Study

In our ablation study, shown in Table 3, we analyze the effectiveness of each component. We reveal that all modules are effective, with clustering and masking strategies bridging the domain gaps, enhancing feature discrimination, and generating more reliable pseudo-labels. Furthermore, \mathcal{L}_{align} and \mathcal{L}_{pol} enable the model to train directly on target domain data and explore the hierarchical features inherent in emotional images. Panels (a), (b), and (c) of Figure 5 provide a sensitivity analysis for the loss weights λ , γ , and δ . They show that the performance of BBA is relatively stable across different hyper-parameters. Panel (d) details the performance gains of the model during training, with different colors representing different stages of training. Panel (e) visualizes the ablation results with different evaluation metrics on FI \rightarrow EmoSet, further illustrating the effectiveness of the components proposed in BBA.

Visualization

Figure 6 visually analyzes the results before and after adaptation on FI \rightarrow EmoSet. Panel (a) and Panel (f) illustrate the feature distribution on EmoSet for models trained on FI and EmoSet, respectively, showing the initial discrepancy between the source and target domains. Panels (b) and (c) show the feature distributions obtained by applying the SFDA methods SHOT and its extension SHOT++. Both methods

retain substantial source domain knowledge due to their fundamental reliance on the source domain model.

Panel (d) visualizes the results after our DMG step. Compared with panels (b) and (c), DMG increases intra-class compactness and inter-class separability, making the feature distribution clearer and enabling the generation of high-confidence pseudo-labels. Panel (e) visualizes the results after our TMA step, highlighting its ability to effectively learn emotion feature representations from the target domain. This is achieved by learning the target domain model from scratch, closely mirroring the Oracle results shown in Panel (f). It demonstrates the superior ability of our method to capture the nuanced emotional features of the target domain, thereby setting a new state-of-the-art in SFDA-VER.

Conclusion

This paper introduces a new task termed source-free domain adaptation for visual emotion recognition (SFDA-VER). To address this task, we propose a novel method called BBA, which consists of two steps: domain-bridged model generation (DMG) and target-related model adaptation (TMA). First, the DMG step bridges the source and target domains. The clustering post-processing strategy enhances inter-class separability, while the masking strategy increases intra-class compactness. These improvements bolster the separability of categories and generate more reliable pseudo-labels. Subsequently, the TMA step focuses on the target structure to begin training anew. The emotion polarity loss enhances the model’s capabilities in emotion recognition. In summary, BBA addresses the unique challenges of the SFDA task within the VER dataset.

Acknowledgments

This research was supported by the National Science and Technology Major Project (2021ZD0110901) and the National Science Foundation of China (No. 62476069).

References

- Ahmed, W.; Morerio, P.; and Murino, V. 2022. Cleaning noisy labels by negative ensemble learning for source-free unsupervised domain adaptation. In *WACV*, 1616–1625.
- Bokolo, B. G.; and Liu, Q. 2023. Deep learning-based depression detection from social media: Comparative evaluation of ml and transformer techniques. *Electronics*, 12(21): 4396.
- Cen, J.; Qing, C.; Ou, H.; Xu, X.; and Tan, J. 2024. MASANet: Multi-Aspect Semantic Auxiliary Network for Visual Sentiment Analysis. *IEEE Transactions on Affective Computing*.
- Deng, S.; Wu, L.; Shi, G.; Xing, L.; Hu, W.; Zhang, H.; and Xiang, Y. 2022. Simple but powerful, a language-supervised method for image emotion classification. *IEEE Transactions on Affective Computing*, 14(4): 3317–3331.
- Deng, S.; Wu, L.; Shi, G.; Xing, L.; Jian, M.; Xiang, Y.; and Dong, R. 2024. Learning to compose diversified prompts for image emotion classification. *Computational Visual Media*, 1–15.
- Ding, Q.; Yin, J.; Zhang, D.; and Gao, J. 2023. Combating Confirmation Bias: A Unified Pseudo-Labeling Framework for Entity Alignment. *arXiv:2307.02075*.
- Du, Y.; Yang, H.; Chen, M.; Luo, H.; Jiang, J.; Xin, Y.; and Wang, C. 2023. Generation, augmentation, and alignment: A pseudo-source domain based method for source-free domain adaptation. *Machine Learning*, 1–21.
- Hao, Y.; Guo, Y.; and Yang, C. 2021. Source-free Unsupervised Domain Adaptation with Surrogate Data Generation. In *BMVC*, 198.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Hoyer, L.; Dai, D.; Wang, H.; and Van Gool, L. 2023. MIC: Masked image consistency for context-enhanced domain adaptation. In *CVPR*, 11721–11732.
- Huang, Y.; Yang, X.; Zhang, J.; and Xu, C. 2022. Relative alignment network for source-free multimodal video domain adaptation. In *ACM MM*, 1652–1660.
- Jeon, H.; Lee, S.; and Kang, U. 2021. Unsupervised multi-source domain adaptation with no observable source data. *PloS one*, 16(7): e0253415.
- Jiang, J.; Zhao, S.; Zhu, J.; Tang, W.; Xu, Z.; Yang, J.; Xu, P.; and Yao, H. 2024. Multi-source Domain Adaptation for Panoramic Semantic Segmentation. *arXiv:2408.16469*.
- Jin, Y.; Wang, X.; Long, M.; and Wang, J. 2020. Minimum class confusion for versatile domain adaptation. In *ECCV*, 464–480. Springer.
- Kim, K.; Ji, B.; Yoon, D.; and Hwang, S. 2021. Self-knowledge distillation with progressive refinement of targets. In *ICCV*, 6567–6576.
- Kothandaraman, D.; Chandra, R.; and Manocha, D. 2021. SS-SFDA: Self-supervised source-free domain adaptation for road segmentation in hazardous environments. In *ICCV*, 3049–3059.
- Kurmi, V. K.; Subramanian, V. K.; and Namboodiri, V. P. 2021. Domain impression: A source data free domain adaptation method. In *WACV*, 615–625.
- Laine, S.; and Aila, T. 2016. Temporal Ensembling for Semi-Supervised Learning. In *ICLR*.
- Li, J.; Yu, Z.; Du, Z.; Zhu, L.; and Shen, H. T. 2024. A comprehensive survey on source-free domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Li, R.; Jiao, Q.; Cao, W.; Wong, H.-S.; and Wu, S. 2020. Model adaptation: Unsupervised domain adaptation without source data. In *CVPR*, 9641–9650.
- Li, Y.; Guo, L.; and Ge, Y. 2023. Pseudo labels for unsupervised domain adaptation: A review. *Electronics*, 12(15): 3325.
- Liang, J.; Hu, D.; and Feng, J. 2020. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *ICML*, 6028–6039.
- Liang, J.; Hu, D.; Feng, J.; and He, R. 2022. Dine: Domain adaptation from single and multiple black-box predictors. In *CVPR*, 8003–8013.
- Liang, J.; Hu, D.; Wang, Y.; He, R.; and Feng, J. 2021. Source data-absent unsupervised domain adaptation through hypothesis transfer and labeling transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11): 8602–8617.
- Machajdik, J.; and Hanbury, A. 2010. Affective image classification using features inspired by psychology and art theory. In *ACM MM*, 83–92.
- Mao, H.; Du, L.; Zheng, Y.; Fu, Q.; Li, Z.; Chen, X.; Han, S.; and Zhang, D. 2024. Source Free Graph Unsupervised Domain Adaptation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, 520–528.
- Peng, K.-C.; Chen, T.; Sadovnik, A.; and Gallagher, A. C. 2015. A mixed bag of emotions: Model, predict, and transfer emotion distributions. In *CVPR*, 860–868.
- Rao, T.; Li, X.; and Xu, M. 2020. Learning multi-level deep representations for image emotion classification. *Neural Processing Letters*, 51: 2043–2061.
- Rao, T.; Li, X.; Zhang, H.; and Xu, M. 2019. Multi-level region-based convolutional neural network for image emotion classification. *Neurocomputing*, 333: 429–439.
- Razali, N. A. M.; Malizan, N. A.; Hasbullah, N. A.; Wook, M.; Zainuddin, N. M.; Ishak, K. K.; Ramli, S.; and Sukardi, S. 2021. Opinion mining for national security: techniques, domain applications, challenges and research opportunities. *Journal of Big Data*, 8: 1–46.
- Tang, S.; Chang, A.; Zhang, F.; Zhu, X.; Ye, M.; and Zhang, C. 2024. Source-free domain adaptation via target prediction distribution searching. *International Journal of Computer Vision*, 132(3): 654–672.

- Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, volume 30.
- Venkateswara, H.; Eusebio, J.; Chakraborty, S.; and Panchanathan, S. 2017. Deep hashing network for unsupervised domain adaptation. In *CVPR*, 5018–5027.
- Wang, R.; Wu, Z.; Weng, Z.; Chen, J.; Qi, G.-J.; and Jiang, Y.-G. 2022. Cross-domain contrastive learning for unsupervised domain adaptation. *IEEE Transactions on Multimedia*, 25: 1665–1673.
- Xie, B.; Yuan, L.; Li, S.; Liu, C. H.; Cheng, X.; and Wang, G. 2022. Active learning for domain adaptation: An energy-based approach. In *AAAI*, 8, 8708–8716.
- Yang, J.; Huang, Q.; Ding, T.; Lischinski, D.; Cohen-Or, D.; and Huang, H. 2023. EmoSet: A large-scale visual emotion dataset with rich attributes. In *ICCV*, 20383–20394.
- Yang, J.; She, D.; Lai, Y.-K.; and Yang, M.-H. 2018a. Retrieving and classifying affective images via deep metric learning. In *AAAI*, 1.
- Yang, J.; She, D.; Sun, M.; Cheng, M.-M.; Rosin, P. L.; and Wang, L. 2018b. Visual sentiment prediction based on automatic discovery of affective regions. *IEEE Transactions on Multimedia*, 20(9): 2513–2525.
- Yang, S.; Jui, S.; van de Weijer, J.; et al. 2022. Attracting and dispersing: A simple approach for source-free domain adaptation. In *NeurIPS*, 5802–5815.
- Yang, S.; van de Weijer, J.; Herranz, L.; Jui, S.; et al. 2021. Exploiting the intrinsic neighborhood structure for source-free domain adaptation. In *NeurIPS*, volume 34, 29393–29405.
- Yao, X.; She, D.; Zhao, S.; Liang, J.; Lai, Y.-K.; and Yang, J. 2019. Attention-aware polarity sensitive embedding for affective image retrieval. In *ICCV*, 1140–1150.
- You, Q.; Luo, J.; Jin, H.; and Yang, J. 2015. Robust image sentiment analysis using progressively trained and domain transferred deep networks. In *AAAI*, 1.
- 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*.
- Zhang, J.; Chen, M.; Sun, H.; Li, D.; and Wang, Z. 2020. Object semantics sentiment correlation analysis enhanced image sentiment classification. *Knowledge-Based Systems*, 191: 105245.
- Zhang, Y.; Liang, J.; Zhang, Z.; Wang, L.; Jin, R.; Tan, T.; et al. 2022a. Free Lunch for Domain Adversarial Training: Environment Label Smoothing. In *ICLR*.
- Zhang, Z.; Chen, W.; Cheng, H.; Li, Z.; Li, S.; Lin, L.; and Li, G. 2022b. Divide and contrast: Source-free domain adaptation via adaptive contrastive learning. In *NeurIPS*, 5137–5149.
- Zhao, S.; Chen, X.; Yue, X.; Lin, C.; Xu, P.; Krishna, R.; Yang, J.; Ding, G.; Sangiovanni-Vincentelli, A. L.; and Keutzer, K. 2022a. Emotional semantics-preserved and feature-aligned cyclegan for visual emotion adaptation. *IEEE Transactions on Cybernetics*, 52(10): 10000–10013.
- Zhao, S.; Gao, Y.; Ding, G.; and Chua, T.-S. 2017. Real-time multimedia social event detection in microblog. *IEEE Transactions on Cybernetics*, 48(11): 3218–3231.
- Zhao, S.; Gao, Y.; Jiang, X.; Yao, H.; Chua, T.-S.; and Sun, X. 2014. Exploring principles-of-art features for image emotion recognition. In *ACM MM*, 47–56.
- Zhao, S.; Hong, X.; Yang, J.; Zhao, Y.; and Ding, G. 2023. Toward Label-Efficient Emotion and Sentiment Analysis. *Proceedings of the IEEE*, 111(10): 1159–1197.
- Zhao, S.; Jia, Z.; Chen, H.; Li, L.; Ding, G.; and Keutzer, K. 2019a. PDANet: Polarity-consistent deep attention network for fine-grained visual emotion regression. In *ACM MM*, 192–201.
- Zhao, S.; Jiang, J.; Tang, W.; Zhu, J.; Chen, H.; Xu, P.; Schuller, B. W.; Tao, J.; Yao, H.; and Ding, G. 2024. Multi-source multi-modal domain adaptation. *Information Fusion*, 102862.
- 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.
- Zhao, S.; Yao, H.; Gao, Y.; Ji, R.; Xie, W.; Jiang, X.; and Chua, T.-S. 2016. Predicting personalized emotion perceptions of social images. In *ACM MM*, 1385–1394.
- Zhao, S.; Yao, X.; Yang, J.; Jia, G.; Ding, G.; Chua, T.-S.; Schuller, B. W.; and Keutzer, K. 2022b. Affective image content analysis: Two decades review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10): 6729–6751.
- Zhu, X.; Li, L.; Zhang, W.; Rao, T.; Xu, M.; Huang, Q.; and Xu, D. 2017. Dependency Exploitation: A Unified CNN-RNN Approach for Visual Emotion Recognition. In *IJCAI*, 3595–3601.