Disentangled Source-Free Personalization for Facial Expression Recognition with Neutral Target Data

Masoumeh Sharafi[1], **Emma Ollivier**[1], **Muhammad Osama Zeeshan**[1], **Soufiane Belharbi**[1],
Marco Pedersoli[1], **Alessandro Lameiras Koerich**[2], **Simon Bacon**[3,4], and **Eric Granger**[1]

[1] LIVIA, ILLS, Dept. of Systems Engineering, ETS Montreal, Canada

[2] LIVIA, Dept. of Software and IT Engineering, ETS Montreal, Canada

[3] Dept. of Health, Kinesiology & Applied Physiology, Concordia University, Montreal, Canada

[4] Montreal Behavioural Medicine Centre, CIUSSS Nord-de-l'Ile-de-Montréal, Canada

`{masoumeh.sharafi.1,emma.ollivier.1,muhammad-osama.zeeshan.1}@ens.etsmtl.ca`
`{soufiane.belharbi,Marco.Pedersoli,Alessandro.LameirasKoerich,eric.granger}@etsmtl.ca`
`Simon.Bacon@concordia.ca`

ABSTRACT

Facial Expression Recognition (FER) from videos is a crucial task in various application areas, such as human-computer interaction and health monitoring (e.g., pain, depression, fatigue, and stress). Beyond the challenges of recognizing subtle emotional or health states, the effectiveness of deep FER models is often hindered by the considerable variability of expressions among subjects. Source-free domain adaptation (SFDA) methods are employed to adapt a pre-trained source model using only unlabeled target domain data, thereby avoiding data privacy and storage issues. Typically, SFDA methods adapt to a target domain dataset corresponding to an entire population and assume it includes data from all recognition classes. However, collecting such comprehensive target data can be difficult or even impossible for FER in healthcare applications. In many real-world scenarios, it may be feasible to collect a short neutral control video (displaying only neutral expressions) for target subjects before deployment. These videos can be used to adapt a model to better handle the variability of expressions among subjects. This paper introduces the Disentangled Source-Free Domain Adaptation (DSFDA) method to address the SFDA challenge posed by missing target expression data. DSFDA leverages data from a neutral target control video for end-to-end generation and adaptation of target data with missing non-neutral data. Our method learns to disentangle features related to expressions and identity while generating the missing non-neutral target data, thereby enhancing model accuracy. Additionally, our self-supervision strategy improves model adaptation by reconstructing target images that maintain the same identity and source expression. Experimental results[1] on the challenging BioVid and UNBC-McMaster pain datasets indicate that our DSFDA approach can outperform state-of-the-art adaptation methods. Our code is available at: github.com/MasoumehSharafi/DSFDA-for-Pain-Assessment.

---

[1] Our code is provided in suppl. materials and will be made public.

1 Introduction

Facial expression recognition (FER) has attracted much attention in computer vision due to the breadth of its applications, such as pain assessment, stress estimation, and driver fatigue monitoring. Although state-of-the-art (SOTA) deep FER models [33, 40, 41, 59] have made significant progress, developing applications and models capable of characterizing expression remains a challenge. Deep learning models [31], [32], [50] have shown promising performance on various FER tasks. However, their performance can degrade significantly when there is a shift in data distributions between the source (lab) and target (operational) domains. Various unsupervised adaptation (UDA) methods [25, 29, 56, 61] have been proposed to address the distribution shifts between the source and the target domain. For instance, the work by Zeeshan et al. [56] introduces a subject-based multi-source domain adaptation (MSDA) method specifically tailored for FER tasks, where each subject is considered as an individual domain for fine-grained adaptation to the variability of expressions among subjects.

State-of-the-art UDA methods [5, 19, 25, 26, 61] rely on the availability of the labeled source data during adaptation, which is impractical in real-world scenarios due to concerns for data privacy, data storage, transmission cost, and computation burden. Therefore, to reduce the reliance on source data during the adaptation process, several source-free (unsupervised) domain adaptation (SFDA) methods have been proposed [11, 24, 28, 54, 57].

SFDA methods typically focus on the scenario where unlabeled target data from all classes is available for the adaptation process (see Fig. 1(a)). However, in real-world video-based health applications, person-specific data representing non-neutral expressions is typically costly or unavailable. A short

Figure 1: Illustration of SFDA settings using pain estimation as an example. (a) A standard subject-based SFDA setting where the source pre-trained model F_S is adapted through a pseudo-labeling approach using unlabeled target video data from all classes (neutral and pain levels). (b) Our proposed subject-based disentangled SFDA (DSFDA) setting with neutral control target data. The generator G is employed to synthesize the missing classes in the target domain based on the neutral control video. (c) Comparison of accuracy on BioVid data [46] of a source-only pre-trained Inception-V3 CNN (F_s before adaptation), after adaptation with only neutral control data, and with our DSFDA method on neutral and generated pain data. (Locks denote that labeled source data is not available during adaptation.)

neutral control video may, however, be collected for target individuals and used to personalize a model to the variability of an individual's diverse expressions. In practice, collecting neutral target data for adaptation is generally easier than gathering non-neutral emotional data. Annotating neutral expressions is more straightforward and less subjective than interpreting emotional expressions, which can vary considerably among expert annotators. Creating scenarios to elicit specific emotions is complex, time-consuming, and raises ethical and privacy concerns. In contrast, capturing neutral expressions is less intrusive and may encourage more willing participation, as it avoids the need to act out emotional states.

This paper introduces a new setting for SFDA, where data from all classes is available for source domain subjects to pre-train a deep FER model, while only neutral control data is available for target domain subjects to adapt this model. In the standard SFDA setting, data from all classes is available for the source and target domains. In contrast, our proposed setting assumes that only neutral data is available for adaptation. At inference time, however, the FER model will process data from all classes. Fig. 1(c) shows the limitation of using a source-only pre-trained deep FER model without adaptation (in blue) or when fine-tuning with only neutral control target data (in orange). The accuracy of models is compared across neutral and pain classes (all pain levels combined). A model adapted only to neutral data becomes biased toward the neutral class and struggles to generalize to non-neutral expression. This is shown by its significantly lower accuracy on pain expressions compared to the source-only model.

In this paper, a new personalized SFDA method is proposed for FER with our new setting (see Fig. 1(b)). The Disentangled SFDA (DSFDA) method is proposed to address the challenges of personalized SFDA using neutral control target data by generating the missing non-neutral expressive image data for target subjects. Using a dual-generator architecture pre-trained on source subjects, a neutral image is transformed into one that exhibits the missing non-neutral expression. By combining features from two input images, the method synthesizes expressive images that belong to the missing non-neural target expressions, enabling accurate FER predictions for target subjects even without explicit data for those expressions. Our approach is further enhanced using self-supervised learning, which ensures the reconstruction of input target images with the same identity and source expression. This process improves the model's adaptability and prediction accuracy in the absence of source data. As shown in Fig. 1(c), DSFDA overcomes the challenges of domain shifts and incomplete target data, ensuring a high level of performance across diverse target domains.

Our contributions can be summarized as follows. **(1)** A novel SFDA setting is introduced, where all classes in the source data are utilized to pre-train the source model. For target

adaptation, however, only neutral control target data is available to adapt the model for a target subject. **(2)** We propose the DSFDA method, a one-stage approach tailored to our novel setting, that combines data generation and target domain adaptation to simultaneously generate non-neutral target image data and adapt a FER model. **(3)** An extensive set of experimental results indicating that the FER model adapted using our DSFDA method outperforms state-of-the-art SFDA methods on the BioVid and UNBC-McMaster datasets. The method scales well to the number of subjects and is robust across various subjects, target domains, and datasets.

2 Related Work

2.1 Source-Free Domain Adaptation

UDA methodologies [5, 19, 26, 61] require simultaneous access to both labeled source domain and unlabeled target domain data during training. Some facial analysis datasets are characterized by concerns about privacy and sensitivity. SFDA methods have been proposed to address the challenges of privacy and memory storage. SOTA methods SFDA can be classified into two distinct groups [11]: data-based methods [10, 17, 18, 23, 43, 51] and the model-based methods [12, 28, 52].

Data-based methods: [18] used BN statistics to generate source-like images, which transfer the style by aligning batchwise feature statistics of generated source-like image features to that stored in the BN layer. Instead of directly generating source-like data, other studies focus on approximating and generating the underlying data distribution. For example, [10] utilized source classifier weights and target pseudo-labels computed through spherical k-means clustering to estimate the source feature distribution. Next, the proxy source data can be extracted from the approximated source distribution, and a conventional DA approach is used to achieve cross-domain alignment.

Model-based methods: Conti et al. [8] introduced a cluster purity score to obtain reliable pseudo-labels using k-means clustering. The target feature extractor is then trained using two augmentations of input images instead of raw input data. Then, the reduced target dataset is trained using the obtained reliable pseudo-labels with the self-supervised feature extractor. Guo et al. [13] incorporate Tsallis entropy into the information maximization constraint to enhance the confidence of prediction labels while reducing confidence bias. Moreover, the method utilizes virtual adversarial training by employing perturbed data alongside original target domain data to improve the model's domain generalization ability.

Unlike the promising results of SFDA methods in benchmark classification tasks, current SFDA approaches are typically designed under the assumption of an equal number of classes in both the source and target domains. However, in FER tasks, non-neutral expression data (missing classes) may be unavailable due to the high costs of data collection and annotation, with only a short neutral video available for adaptation. This leads to a label shift between the target adaptation and testing phases, in addition to the domain shift between the source and target domains, ultimately degrading classification performance. To mitigate the shift between the source and target domains, generative models are utilized to synthesize the missing classes within the target domain, thereby enriching the diversity of the target data. These models aim to produce high-quality images or videos that accurately represent the emotional expressions absent in the original target dataset. In the literature, a variety of techniques have been explored for this purpose, ranging from traditional generative adversarial networks (GANs) and variational autoencoders (VAEs) to more advanced frameworks such as conditional GANs, diffusion models, and disentanglement methods.

2.2 Facial Expression Generation and Disentanglement

Generative modeling, a branch of machine learning (ML), concentrates on understanding the underlying distribution of training data to enable the generation of new samples that mimic the statistical features of that data. Over time, significant progress has been made in this field. Most common techniques like GANs [2, 3, 4, 7, 58], VAEs [22, 37, 44], and diffusion models [9, 21, 34, 38] have significantly improved the quality and variety of the samples generated. For facial expression generation, [42] developed FExGAN-Meta for meta-humans, while ImaGINator [48] employed CGANs for creating facial expression video sequences. Niinuma et al. [35] used a GAN-based approach to synthesize facial expressions for AU intensity estimation, improving performance on imbalanced and domain-shifted datasets through personalized networks. Recent works have emphasized domain generalization and adaptation for cross-corpus facial expression recognition. Yin et al. [55] introduced the FATE model, which uses contrastive learning and facial animation warping for superior cross-domain emotion recognition. Furthermore, DeepFN [15] proposed deep face normalization for robust action unit recognition across individuals. Kim et al. [53] leveraged single-image temporal cues for action unit detection, enhancing the understanding of subtle expressions.

Disentanglement in ML refers to the process of learning representations where different factors of variation in the data are isolated into distinct components [16]. It is helpful for tasks like data augmentation, image synthesis, and improving ML model performance. In the literature, numerous methods of disentanglement are used for different tasks. In their study, [47] suggest classifying these techniques according to several criteria. Disentanglement holds significant relevance

for facial expression recognition, where expression features often exhibit a high correlation with the identities of subjects. Popularized by [6], GANs are frequently employed in FER disentanglement studies to effectively separate expression features from identity features, thereby improving the performance and interpretability of models [1, 50]. Additionally, disentangling other factors is possible as well. The IPD-FER method utilizes three encoders to disentangle identity, pose, and expression factors [20]. Furthermore, Halawa et al. [14] propose a GAN-based disentanglement approach that achieves the separation of expression features without relying on identity labels.

Beyond GAN, other techniques are also explored for disentanglement, including the VAEs methods, which is another popular method for disentanglement [60]. Moreover, Li et al. [27] propose to use the Hilbert-Schmidt independence criterion to disentangle expression from identity features, demonstrating the application of mathematical criteria in disentanglement tasks. In the same way, mutual information minimization is a strategy for achieving this objective [39]. These studies highlight diverse methodologies aimed at achieving robust and interpretable disentangled representations in facial expression recognition.

Given only neutral control videos in the target domain, we generate video frames that represent non-neutral expressions. This two-stage approach converts the SFDA problem with only neutral target videos into a standard SFDA scenario where both domains share the same classes. First, a generator (like VAEs, GANs, or DMs) is trained on the source domain to create images with non-neutral expressions, while a source model is pre-trained on the source data. Then, during inference, the generator produces images or videos for the target data, which initially lacks these expressions. These synthetic images are combined with the target's neutral images to adapt the source model to the target domain, improving expression recognition. However, current generative methods often struggle to capture subtle pain intensity levels due to the complexity of nuanced expressions.

SOTA methods for SFDA assume the same classes for target adaptation and testing across all settings, such as closed-set, partial-set, and open-set. However, in FER tasks, non-neutral expression data may not be available due to the challenges and costs of data collection and annotation, leaving only neutral control target video for adaptation. This results in a label shift between the target adaptation and testing phases, in addition to the domain shift between the source and target domains, ultimately degrading classification performance. While generative models are used to produce non-neutral expressions for target data, they often struggle with subtle expressions, such as varying levels of pain intensity. This limitation arises because these deep FER models find it challenging to handle the nuances required for accurate expression generation. As

described in the next section, our DSFDA method addresses this challenge by employing a disentanglement process that effectively separates expression information from other facial identity attributes, resulting in a higher overall performance.

3   Proposed Method

Fig. 2 shows the training and testing architecture of our Disentangled SFDA (DSFDA) method for adapting a FER model using only neutral control video per target subject. To produce the non-neutral expressions for the target data and align with the source domain, a disentanglement method is used to simultaneously generate and adapt the target data by reducing the perceptual and reconstruction loss between the input image and the generated image.

3.1   Notation

In the proposed method, $\mathbf{I}_{id}$ denotes the identity image with its corresponding one-hot label y_{id}, and $\mathbf{I}_{exp}$ denotes the expression image with its corresponding one-hot label y_{exp}. The generated image produced by the model is represented as $\mathbf{I}_g$, while y_g is the one-hot label of the generated image $\mathbf{I}_g$. The number of identity classes is denoted as c_{id}, and the number of expression classes is represented as c_{exp}. The distribution of identity images is indicated by p_{id}, the distribution of expression images by p_{exp}, and the distribution of generated images by p_g. The reconstructed identity image is denoted as $\hat{\mathbf{I}}_{id}$, and the reconstructed expression image is represented as $\hat{\mathbf{I}}_{exp}$. The identity representation obtained from the identity encoder is indicated as $\mathbf{d}_{id}$, while $\mathbf{d}_{g\_id}$ refers to the identity representation of the generated image. Additionally, n represents the number of training samples, and $\hat{y}_i^j$ is the j-th output value from the softmax classifier for the i-th sample.

3.2   Model Architecture

DSFDA is inspired by the architecture in [50], where expressions are disentangled from other facial information to improve expression classification performance. DSFDA relies on two encoders to disentangle identity and expression information from facial images for improved adaptation. It consists of a generator, with two encoders (E_{id}, E_{exp}) and a decoder (D), and two discriminators (D_{id}, D_{exp}) for identity and expression classification. The generator creates synthesized images the combining the identity of one image with the expression of another, and these images are verified by the discriminators. The model is trained using reconstruction, perceptual, and adversarial losses, along with an expression classifier (C_{exp}) trained on real data.

The training and adaptation architecture is a GAN model used originally to disentangle identity and expression from identity.
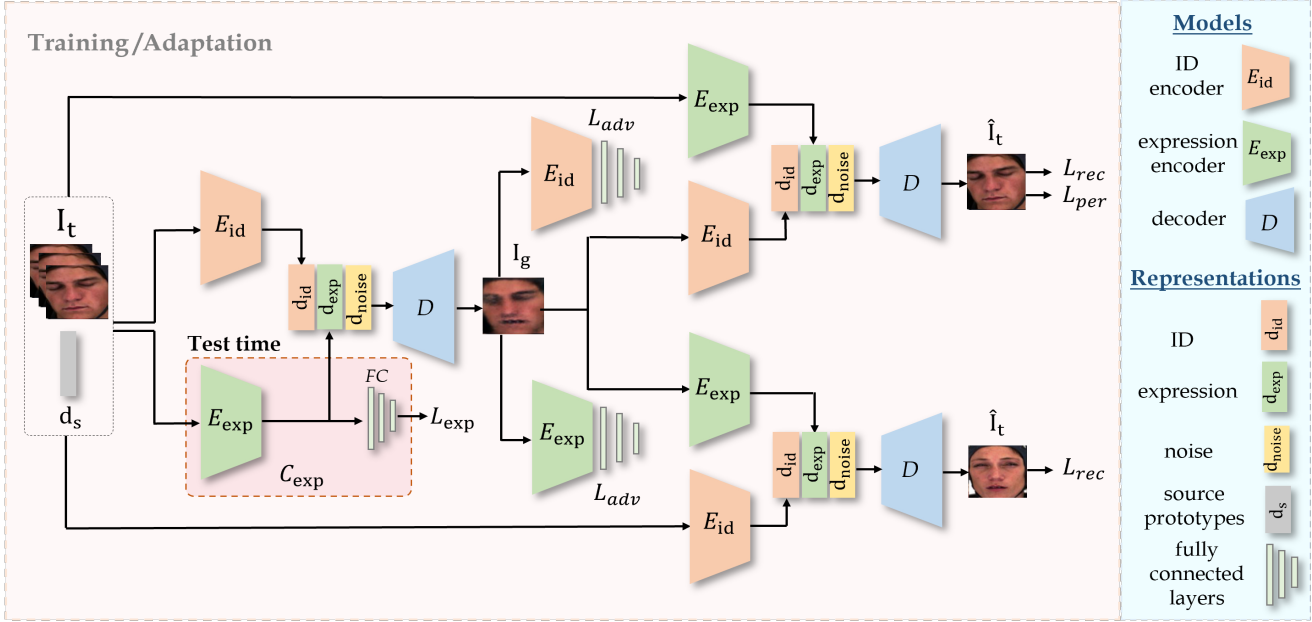
Figure 2: Training and testing architecture of our DSFDA method. In the training and adaptation part (*pink*), the identity and expression features of input images are used to generate images with the identity of the target image and the expression of the source prototype embeddings through adversarial training. The model learns to disentangle the expression and identity information through generation. At test time (*red*), the expression encoder with fully connected layers is used to classify expressions.

Our architecture is thus made of different components: one generator, G, and two discriminators, D_{id} and D_{exp}, respectively, for identity and expression. The generator is composed of one decoder (D) and two encoders; the first one specializes in extracting identity features, E_{id}, and the second one specializes in extracting expression features E_{exp}. A pair of images feed them $\mathbf{I}_{id}$ and $\mathbf{I}_{exp}$ from which a fake image $\mathbf{I}_g$ is generated with the identity of $\mathbf{I}_{id}$ and the pain level of $\mathbf{I}_{exp}$. Both discriminators D_{id} and D_{exp} are fed with this fake image to ensure that the correct identity and the correct expression have been generated. The generator is then once again used to generate, from $\mathbf{I}_g$ and the two real images $\mathbf{I}_{id}$ and $\mathbf{I}_{exp}$, two other fake images $\hat{\mathbf{I}}_{id}$ and $\hat{\mathbf{I}}_{exp}$, which are respectively the copies of $\mathbf{I}_{id}$ and $\mathbf{I}_{exp}$.

3.3 Objective Function

(1) Classifier: Fully connected layers of the classifier C_{exp} are trained using real images $\mathbf{I}_{exp}$ with cross-entropy loss:

$$\mathcal{L}_{exp} = -\mathbb{E}_{(\mathbf{I}_{exp}, y_{exp}) \sim p_{exp}} \sum_{e=1}^{c_{exp}} \mathbb{1}[y_{exp} = e] \log(C_{exp}(\mathbf{I}_{exp})) \quad (1)$$

(2) Discriminator: To operate the adversarial method, two discriminators are trained with real images $\mathbf{I}_{id}$ and $\mathbf{I}_{exp}$ to recognize either identity or pain expression. We employ an

expression discriminator D_{exp} with c_{exp} classes. Similarly, an identity discriminator D_{id}, configured with $c_{id} + 1$ classes, guides the generator to produce realistic images with the specified identities. The loss functions for the identity discriminator $\mathcal{L}_{dis\_id}$ and the expression discriminator $\mathcal{L}_{dis\_exp}$ are defined by:

$$\mathcal{L}_{dis\_id} = -\mathbb{E}_{(\mathbf{I}_{id}, y_{id}) \sim p_{id}} \sum_{i=1}^{c_{id}+1} \mathbb{1}[y_{id} = i] \log(D_{id}(\mathbf{I}_{id}))$$

$$- \mathbb{E}_{(\mathbf{I}_g, y_g) \sim p_g} \sum_{g=1}^{c_g} \mathbb{1}[y_g = g] \log(D_g(\mathbf{I}_{id})) \quad (2)$$

$$\mathcal{L}_{dis\_exp} = -\mathbb{E}_{(\mathbf{I}_{exp}, y_{exp}) \sim p_{exp}} \sum_{e=1}^{c_{exp}} \mathbb{1}[y_{exp} = e] \log(D_{exp}(\mathbf{I}_{exp}))$$

$$(3)$$

where y_g is the label of the generated image, p_{id}, p_{exp}, and p_g represent the distributions of identity images, expression images, and generated images, respectively and c_{id}, c_{exp}, and c_g are the number of identity, expression, and generated image classes.

Given that the generated images should replicate the overall appearance of the identity input image, an additional class for real and fake images is included in the identity discriminator. Consequently, only the identity discriminator processes the generated images. The expression classifier's role is to

classify the expressions of the images. Using real expression data for the expression discriminator during training is therefore advantageous. Using generated images with incorrect emotions can destabilize the training process.

(3) Generator: The generator consists of a decoder and two encoders: one for expression and the other for identity. During training, all components of the generator are updated simultaneously. A reconstruction loss is employed to ensure the generation of accurate images, where the reconstructed identity image $\hat{\mathbf{I}}_{\text{id}}$ closely resembles the input identity image $\mathbf{I}_{\text{id}}$, and the reconstructed expression image $\hat{\mathbf{I}}_{\text{exp}}$ closely matches the input expression image $\mathbf{I}_{\text{exp}}$:

$$\mathcal{L}_{\text{rec}} = \mathbb{E}_{\mathbf{I}_{\text{id}} \sim p_{\text{id}}} \|\hat{\mathbf{I}}_{\text{id}} - \mathbf{I}_{\text{id}}\|_1 + \mathbb{E}_{\mathbf{I}_{\text{exp}} \sim p_{\text{exp}}} \|\hat{\mathbf{I}}_{\text{exp}} - \mathbf{I}_{\text{exp}}\|_1 \quad (4)$$

In addition, perceptual loss, defined by MSE loss, is used to ensure that identity features produced by the discriminator are the same for real and generated images, $\mathbf{I}_{\text{id}}$ and $\mathbf{I}_{\text{g}}$:

$$\mathcal{L}_{\text{per}} = \mathbb{E}_{\mathbf{I}_{\text{id}} \sim p_{\text{id}}} \|\mathbf{d}_{\text{id}}(\mathbf{I}_{\text{g}}) - \mathbf{d}_{\text{id}}(\mathbf{I}_{\text{id}})\|_2 \quad (5)$$

where $\mathbf{d}_{\text{id}}(.)$ denotes features extracted by the discriminator specialized in identity recognition.

Discriminators D_{id} and D_{exp} are trained to recognize identity and expression. To make the generator fool both discriminators, the following adversarial loss is applied:

$$\mathcal{L}_{\text{adv}} = -\lambda_3 \mathbb{E}_{(\mathbf{I}_{\text{id}}, y_{\text{id}})} \sum_{i=1}^{c_{\text{id}}} \mathbb{1}[y_{\text{id}} = i] \log(D_{\text{id}}(\mathbf{I}_{\text{d}}))$$
$$-\lambda_4 \mathbb{E}_{(\mathbf{I}_{\text{exp}}, y_{\text{exp}})} \sum_{e=1}^{c_{\text{exp}}} \mathbb{1}[y_{\text{exp}} = e] \log(D_{\text{exp}}(\mathbf{I}_{\text{g}})) \quad (6)$$

where $D_{\text{id}}(.)$ and $D_{\text{exp}}(.)$ are, respectively, the predicted class distribution by identity and expression discriminators, and y_{exp} and y_{id} denote labels of expression and identity for real images.

The total loss function for the generator is:

$$\mathcal{L}_{\text{gen}} = \lambda_1 \mathcal{L}_{\text{rec}} + \lambda_2 \mathcal{L}_{\text{per}} + \mathcal{L}_{\text{adv}} \quad (7)$$

In Eqs. (6) and (7), λ_1 to λ_4 are hyperparameters used to adjust the impact of individual losses.

4   Results and Discussion

4.1   Experimental Methodology

We evaluate the performance of our model using the BioVid Heat Pain dataset [46] and the UNBC-McMaster Shoulder Pain dataset [30].

Datasets. We evaluate the performance of our model using the BioVid Heat Pain dataset [46] and the UNBC-McMaster Shoulder Pain dataset [30].

The BioVid dataset is divided into five distinct parts, each varying in terms of subjects, labeling, modalities, and tasks. In our experimentation, we use Part A of the dataset, which includes data from 87 subjects, encompassing four pain levels (PA1, PA2, PA3, and PA4) in addition to a neutral (BL1). The modalities in Part A consist of frontal video recordings and biological signals. In this study, we focus on the highest pain level and the neutral state, using only the frames extracted from frontal video recordings. Each subject has 20 videos per class, each lasting 5.5 seconds. According to [49], the PA4 dataset does not show any facial activity in the first two seconds of the videos and has demonstrated that initial pain intensities do not elicit any facial activities. They suggest focusing solely on 'no pain' and the highest pain intensities. Therefore, we only consider frames after the first two seconds to exclude the initial part of the sequence that did not indicate any response to pain. In summary, we used 87 subjects with two classes (no pain and pain), 20 videos per class, and 75 frames per video.

The UNBC-McMaster dataset comprises 200 pain videos from 25 subjects. Each video frame is annotated with a pain intensity score based on the PSPI scale in a range of 0 to 15. To address the significant imbalance in pain intensity levels, we adopt the quantization strategy proposed in [36], categorizing pain intensities into five discrete levels: 0 (neutral), 1 (1), 2 (2), 3 (3), 4 (4–5), and 5 (6–15).

Experimental Protocol. In our experiments, we divide both datasets into two domains: source and target. The *first experiment* is conducted using the BioVid dataset, where the source domain comprises 77 subjects and the target domain includes the remaining 10 subjects, with each subject treated as an individual target domain. To evaluate the performance of our approach, we follow a three-phase methodology: source training, target adaptation, and target testing. During the source training phase, the model is trained using the 77 source subjects, each with both neutral and pain level 4 images. In the adaptation phase, the model is adapted to the target data, consisting of neutral videos of the 10 remaining subjects. Finally, in the target testing phase, the model is tested on each target subject using 12 videos representing both neutral and pain expressions.

In the *second experiment*, we use the UNBC-McMaster dataset. In this case, 20 subjects are designated as the source domain, while the remaining five subjects form the target domain. Similar to the BioVid experiment, the model is first trained on the source subjects, adapted to the target data, and tested on each target subject individually.

Experimental Settings. To validate our approach, several settings are explored: (1) Source only, where a pre-trained model on the source domain is directly tested on the target domain without any adaptation. It serves as a baseline for evaluating

Table 1: Accuracy (%) of the proposed method and SOTA methods on the BioVid dataset for 10 target subjects with all 77 sources. Bold text shows the highest accuracy. Note: SFDA (77-shot) indicates that 77 prototypes were used from the source data.

Setting	Methods	Sub-1	Sub-2	Sub-3	Sub-4	Sub-5	Sub-6	Sub-7	Sub-8	Sub-9	Sub-10	Avg.
Source only	Source model	66.11	55.55	86.36	85.11	87.11	59.00	75.66	70.44	52.66	48.22	68.62
SFDA	SHOT [28]	55.78	47.76	41.05	52.44	54.67	48.89	54.46	49.22	54.86	44.44	50.35
	NRC [54]	59.33	39.38	84.00	72.89	79.67	42.67	63.92	53.95	54.89	42.47	60.31
	DSFDA (ours)	**77.00**	**75.11**	**90.89**	**85.67**	**88.11**	**67.48**	**89.22**	**87.56**	**69.22**	**72.11**	**80.24**
SFDA (77-shot)	SHOT [28]	**85.67**	75.11	**90.89**	85.36	87.11	59.11	76.86	84.22	60.33	**86.22**	79.09
	NRC [54]	82.78	71.40	90.91	85.22	87.56	59.31	83.73	85.89	65.89	82.11	79.48
	DSFDA (ours)	**89.56**	**75.11**	**91.00**	**85.56**	**87.78**	**68.33**	**86.78**	**88.33**	**72.22**	74.11	**81.88**
Oracle	Fine-tune	97.11	91.43	96.76	97.89	96.11	92.30	90.01	95.09	98.22	97.00	95.19

how well the model generalizes to unseen data; (2) SFDA, where only neutral images from the target subjects are used for adaptation, while testing is performed on both neutral and pain images. This setting is introduced to investigate how well the model performs when adapted using only neutral images and to compare its results to a scenario where both neutral and pain images are available for adaptation; (3) SFDA (k-shot), where we utilized target neutral data and k prototypes (images) from the source domain. The image prototypes from the source data were selected using k-means clustering, with k being the number of subjects in the source domain. Therefore, one prototype was selected for each subject. The goal here is to investigate whether pain prototypes, drawn from the source domain, can effectively guide the model in recognizing pain in the target domain; and (4) Oracle, an upper-bound scenario, achieved by fine-tuning the source model with labeled neutral and pain images from the target data.

Table 2: Accuracy (%) of the proposed method and SOTA methods on the UNBC-McMaster dataset for five target subjects with all 19 sources. Bold text shows the highest accuracy. Note: SFDA (19-shot) indicates that 19 prototypes were used from the source data.

Setting	Methods	Sub-1	Sub-2	Sub-3	Sub-4	Sub-5	Avg.
Source only	Source model	68.56	78.32	60.37	65.54	33.72	61.30
SFDA	SHOT [28]	69.93	70.84	42.24	63.56	34.41	56.19
	NRC [54]	63.33	69.13	52.24	63.28	33.24	56.24
	DSFDA (ours)	**78.23**	**84.39**	**89.63**	**78.99**	**66.56**	**79.56**
SFDA (19-shot)	SHOT [28]	77.51	79.42	42.99	65.65	36.47	60.41
	NRC [54]	77.26	78.56	64.61	77.89	43.53	68.37
	DSFDA (ours)	**78.23**	**89.45**	**89.63**	**78.85**	**65.23**	**80.27**
Oracle	Fine-tune	95.63	98.44	98.11	97.54	98.25	97.59

Implementation Details. In our experiments, we used the Inception-V3 CNN with Instance Normalization as the backbone for the encoders. We trained the entire model on the source data using a learning rate of 10^{-5} for 100 epochs. To adapt the model to the target subjects, we trained with a learning rate of 10^{-4} for 25 epochs, employing a batch size 32. For the expression transfer, we utilized two inputs corresponding to the expression and identity labels. In our experiments, we maintained consistency by using the same dataset for both the identity and expression encoders. For image preprocessing, we resized the input images to 128×128 pixels and applied data augmentation techniques, including random cropping, horizontal flipping, and color adjustments during the training phase. Throughout our experiments, we set the hyperparameters empirically as $\lambda_1 = 5$, $\lambda_2 = 1$, $\lambda_3 = 0.2$, and $\lambda_4 = 0.8$. For k-shot SFDA, we set $k = 77$ for the BioVid dataset and $k = 19$ for the McMaster dataset, which were also determined empirically.

4.2   Comparison with SOTA Methods

We compared the performance of our DSFDA method with SOTA approaches on the BioVid dataset under various experimental settings, as detailed in Table 1. The *Source-only model* serves as a lower-bound baseline, where the pre-trained source model is directly tested on target domain data without any adaptation. In the *SFDA*, the source model is adapted to the target data using only neutral video frames from the target domain. In *SFDA (k-shot)*, real neutral images with representative pain prototypes from the source domain are used to address the missing classes in the adaptation phase. The goal is to investigate whether pain prototypes drawn from the source domain can effectively guide the model in recognizing pain in the target domain.

Table 3: Accuracy (%) of the proposed method and state-of-the-art methods on the BioVid dataset for 10 target subjects with all 19 sources from UNBC-McMaster. Bold text shows the highest accuracy.

Setting	Method (Source→Target)	Sub-1	Sub-2	Sub-3	Sub-4	Sub-5	Sub-6	Sub-7	Sub-8	Sub-9	Sub-10	Avg.
Source only	Source model	50.00	50.00	50.11	50.33	57.89	52.78	58.00	79.44	65.11	48.22	56.48
SFDA	SHOT [28]	63.33	46.48	52.22	38.11	56.67	51.33	56.13	42.22	46.75	55.33	50.86
	NRC [54]	50.33	51.17	50.11	45.56	56.56	51.00	57.93	50.78	51.33	58.22	52.30
	DSFDA (ours)	**71.78**	**62.56**	**78.11**	**74.22**	**77.33**	**62.33**	**90.56**	**92.33**	**65.00**	**65.89**	**74.01**

As illustrated in Table 1, the performance of the *source only* model, which does not utilize any adaptation, is significantly lower compared to the results achieved after adaptation using either synthetic pain images or pain prototypes. This highlights the importance of domain adaptation in addressing the gap between the source and target domains. As expected, the *SFDA* adaptation setting, which utilizes only short neutral video frames for adaptation, significantly degrades performance, yielding even worse results than the source-only model in existing SOTA methods. When critical expression classes (pain) are missing from the target data, two SOTA methods, SHOT [28], and NRC [54], experience substantial drops in accuracy, with reductions to 18.27% and 8.31%, respectively. These findings highlight the limitations of relying solely on neutral data for adaptation and the challenges models face when key emotional classes are unavailable during adaptation. We used SFDA methods for comparison, which closely align with our proposed method. Our proposed DSFDA method achieves 11.62% higher accuracy than the source-only model due to the generation and adaptation of missing classes in one step through a disentanglement framework. These methods help separate identity-specific and expression-related features, allowing the model to focus on subtle variations in expressions. This leads to better generalization and improved performance, even when pain-specific data is missing. In *SFDA (77-shot)*, we introduced source domain prototypes alongside neutral target data for adaptation, selecting representative samples from the source subjects to address the absence of pain data in the target domain. Our proposed DSFDA method outperformed both SHOT and NRC when using these prototypes, demonstrating the effectiveness of leveraging them to bridge the domain gap. Currently, we use a single image prototype for each subject from the source data. The final setting, referred to as the *Oracle*, involves adapting the source model using both real neutral and pain images from each target subject. This setting serves as the upper bound for adaptation performance, providing the model full access to all relevant class information in the target domain. By utilizing both neutral and pain data during adaptation, the Oracle demonstrates the model's maximum potential when no data is missing, serving as a benchmark for comparison against more challenging settings where class information is incomplete or synthetic.

The performance of the proposed DSFDA method is compared with SOTA approaches on the UNBC-McMaster shoulder pain dataset under four experimental settings in Table 5. DSFDA consistently outperforms both the source-only baseline and other SOTA methods, achieving an average accuracy of up to 78.75% in the prototype setting. Notably, DSFDA attains 77.70% accuracy in the neutral-only setting, surpassing SHOT and NRC by over 20%.

Next, we evaluated the performance of the proposed DSFDA method and SOTA methods on the BioVid dataset for 10 target subjects using all 19 sources from UNBC-McMaster. As shown in Table 3, the *source-only model* achieved an average accuracy of 56.48%. Among the SFDA methods, SHOT [28, 42] and NRC [28, 48] achieved average accuracy of 50.86% and 52.30%, respectively. Notably, the proposed DSFDA method obtained the highest average accuracy of 74.01% and outperformed the other methods across most subjects. For more detailed information on the generation and adaptation processes within a two-stage framework, including visual results and the impact of various clustering methods on selecting source image prototypes, please refer to the supplementary material.

4.3   Ablation Studies

Impact of each loss function. This section analyzes the impact of each loss function on the adaptation of the pre-trained source model to the target subjects. We aim to determine how each loss term contributes to the overall target accuracy by evaluating the model's performance on target data. When only the perceptual and adversarial losses ($\mathcal{L}_{per}$ and $\mathcal{L}_{adv}$) are applied, the model achieves 76.28% accuracy on target data, indicating that these losses promote adaptation but lack the reconstruction constraints needed for better performance. When the reconstruction and adversarial losses ($\mathcal{L}_{rec}$ and $\mathcal{L}_{adv}$) are used without perceptual loss, there is a slight improvement in target accuracy (0.28%). This suggests that reconstruction helps the model better adapt to the target domain. Applying both reconstruction and perceptual losses without adversarial loss significantly decreases performance on target data (66.00%), highlighting the importance of adversarial training for effective target domain adaptation. When

all three losses are combined, the model achieves the best performance, with 81.88% accuracy on target data. The combination of reconstruction, perceptual, and adversarial losses facilitates effective adaptation to the target domain.

Impact of different generative models. To evaluate the impact of using different generative models (FExGAN vs. ImaGINator) for pain-image synthesis, we conduct an ablation study across three methods, SHOT, NRC, and DSFDA, on 10 target BioVid subjects. In Fig. 3, each subject is represented on the x-axis, and the corresponding target accuracies are plotted as grouped bars. Within each group, bars reflect the chosen method paired with either FExGAN or ImaGINator. Among generative models, ImaGINator variants (e.g., SHOT-ImaGINator, NRC-ImaGINator) often surpass their FExGAN-based counterparts (SHOT-FExGAN, NRC-FExGAN), suggesting that ImaGINator generates pain expressions with greater diversity or fidelity. Furthermore, DSFDA demonstrates a consistent performance boost across most subjects, underlining the benefits of its domain-discriminative approach. Specifically, DSFDA-ImaGINator often leads to the highest accuracy, indicating strong synergy between our domain adaptation framework and a robust generative model.

Impact of generation methods on pseudo-labels. We evaluated the reliability of pseudo-labels produced by SHOT and NRC methods when using FExGAN and ImaGINator to generate the missing classes in the target data. ImaGINator consistently produced a higher ratio of reliable pseudo-labels compared to FExGAN. Specifically, the NRC method achieved 75% reliability with FExGAN and improved to 78% with ImaGINator, while the SHOT method resulted in 70% with FExGAN and reached 80% with ImaGINator. Notably, the combination of SHOT with FExGAN yielded the lowest
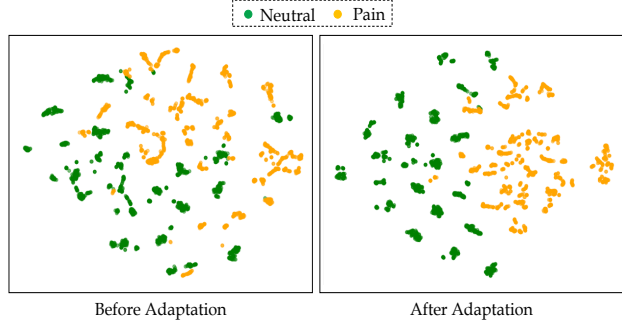
Figure 4: t-SNE visualizations of embedding from source to target subjects for sub-1. Different colors represent different classes (green for neutral and orange for pain). (left) the source-only setting, without adaptation; (right) our DSFDA method after adaptation.

pseudo-label ratio (70%), suggesting that SHOT is less effective when paired with FExGAN compared to ImaGINator. Furthermore, the NRC-FExGAN configuration outperformed our DSFDA-FExGAN method, likely because the generated pain images—often resembling neutral expressions—received reliable pseudo-labels. Since most real or generated pain images do not visibly express pain, assigning pain labels indiscriminately can introduce misleading representations, ultimately diminishing the model's performance during testing.

4.4 Visualization of Feature Distributions

Fig. 4 shows the t-SNE visualizations [45] of the feature representations generated by the proposed DSFDA method for sub-a from the BioVid dataset. Fig. 4 (left) displays the feature distribution before adaptation, while Fig. 4 (right) shows the distribution after adaptation. In the source-only scenario, the features for the two classes (neutral and pain) are poorly aligned, indicating significant domain discrepancy. After applying our DSFDA method, the features are better aligned across the two classes. This improvement is due to the disentanglement strategy employed by our method, which separates domain-specific variations from class-relevant information, leading to more effective adaptation and improved alignment of feature representations.

5   Conclusion

This paper introduces a novel subject-based SFDA framework that specializes in FER in health monitoring applications in the practical case when only neutral expression control target data is available for adaptation. Unlike conventional SFDA settings, our approach effectively adapts a pre-trained source model to target subjects by leveraging a disentangled generative strategy. We propose a one-stage Disentangled SFDA


Figure 3: Per-subject target accuracies using target neutral and synthetic pain images for 10 subjects in the BioVid dataset. Each group of six bars represents the three methods (SHOT, NRC, and DSFDA) combined with two generative models (FExGAN and ImaGINator). The legend at the top identifies each method-model pairing.

(DSFDA) framework that integrates data generation with domain adaptation, improving model generalization while mitigating the practical constraints of FER data collection in real-world clinical environments. Experimental results on the BioVid and UNBC-McMaster pain assessment datasets show that DSFDA can significantly outperform SOTA SFDA methods and highlight its effectiveness in handling cross-domain variability and subject-specific adaptations.

Acknowledgement

This supplementary material encompasses the following sections:

- **A Two-Stage Methodology:** A detailed description of the generation and adaptation stages.

- **Experimental Results of the Two-Stage Framework:** Comprehensive evaluation and analysis of the framework's performance.

- **Impact of prototype selection methods:** Source pain prototype selection using different clustering methods.

A   Two-Stage Methodology

The architecture of the proposed disentangled SFDA (DSFDA) method is presented in the main paper. Algorithm 1 summarizes the DSFDA training process.

In this two-stage framework, the primary goal is to address the challenge of missing pain intensity classes in the target domain by generating video frames that represent varying levels of pain. This approach effectively converts the problem of SFDA with missing classes in the target data into a standard SFDA scenario, where the source and target domains share the same set of classes, as illustrated in Fig. 5.

To achieve this, we leverage several advanced generative models, including VAEs, GANs, and diffusion models, to synthesize images or videos that depict a wide range of pain intensities. The research begins by thoroughly exploring different conditional GANs (cGANs) to create realistic images or videos that correspond to various pain levels. These generative models are trained on source domain data, where pain intensity labels are available, to ensure that the generated frames accurately reflect the intended pain levels.

Algorithm 1 DSFDA Training Process.

1: **Input:**
2:   $\mathbf{I}_{\text{id}}$: target neutral identity
3:   $\mathbf{I}_{\text{exp}}$: target neutral expression
4:   C_{exp}: pre-trained expression classifier
5:   D_{id}: identity discriminators
6:   D_{exp}: expression discriminator
7:   E_{id}: identity encoder
8:   E_{exp}: expression encoder
9:   D: decoder
10:   Hyperparameters: $\lambda_1, \lambda_2, \lambda_3, \lambda_4$
11: **Output:** Adapted generator G, Adapted expression classifier C_{exp}
12: **Initialization:** Generator G, Discriminators $D_{\text{id}}, D_{\text{exp}}$, Classifier C_{exp}
13: **for** $t = 1 : T$ **do**
14:     **Train Generator:**
15:     **for** each pair of images $\mathbf{I}_{\text{id}}, \mathbf{I}_{\text{exp}}$ **do**
16:         Extract features $\mathbf{z}_{\text{id}} = E_{\text{id}}(\mathbf{I}_{\text{id}})$, $\mathbf{z}_{\text{exp}} = E_{\text{exp}}(\mathbf{I}_{\text{exp}})$
17:         Generate fake image $\mathbf{I}_{\text{g}} = G(\mathbf{z}_{\text{id}}, \mathbf{z}_{\text{exp}})$
18:         Generate reconstructions $\hat{\mathbf{I}}_{\text{id}}, \hat{\mathbf{I}}_{\text{exp}}$
19:     **end for**
20:     **Train Discriminators:** Update $\mathcal{L}_{\text{dis\_id}}, \mathcal{L}_{\text{dis\_exp}}$
21:     **Train Classifier:** Apply cross-entropy loss $\mathcal{L}_{\text{exp}}$
22:     **Update Generator:** Minimize generator loss $\mathcal{L}_{\text{gen}}$
23: **end for**
24: **Return:** Adapted expression classifier C_{exp}

In the first stage, Fig. 5(left), the generator and the source model is trained on the 77 source subjects with two classes (neutral and pain). In the second stage, Fig. 5(right), the pretrained generator is utilized to generate missing classes (pain) for the target subject. Moreover, the pre-trained source model is used for adaptation on the real neutral and generated pain images/videos for the ten target subjects.

B   Experimental results

B.1   Qualitative Results

The visual results of the generated images using ImaGINator and FExGAN-Meta are illustrated in Fig.s 6 and 7, respectively. In addition to the visual results, we compare them using numerical results using three metrics, including structural similarity index measure (SSIM), peak signal-to-noise ratio (PSNR), and fréchet inception distance (FID). SSIM is used as a measure to evaluate the quality of generated images using generative models compared to the ground truth or target images. PSNR is used to assess how closely the generated images match the ground truth images. FID measures the similarity between the distributions of feature representations
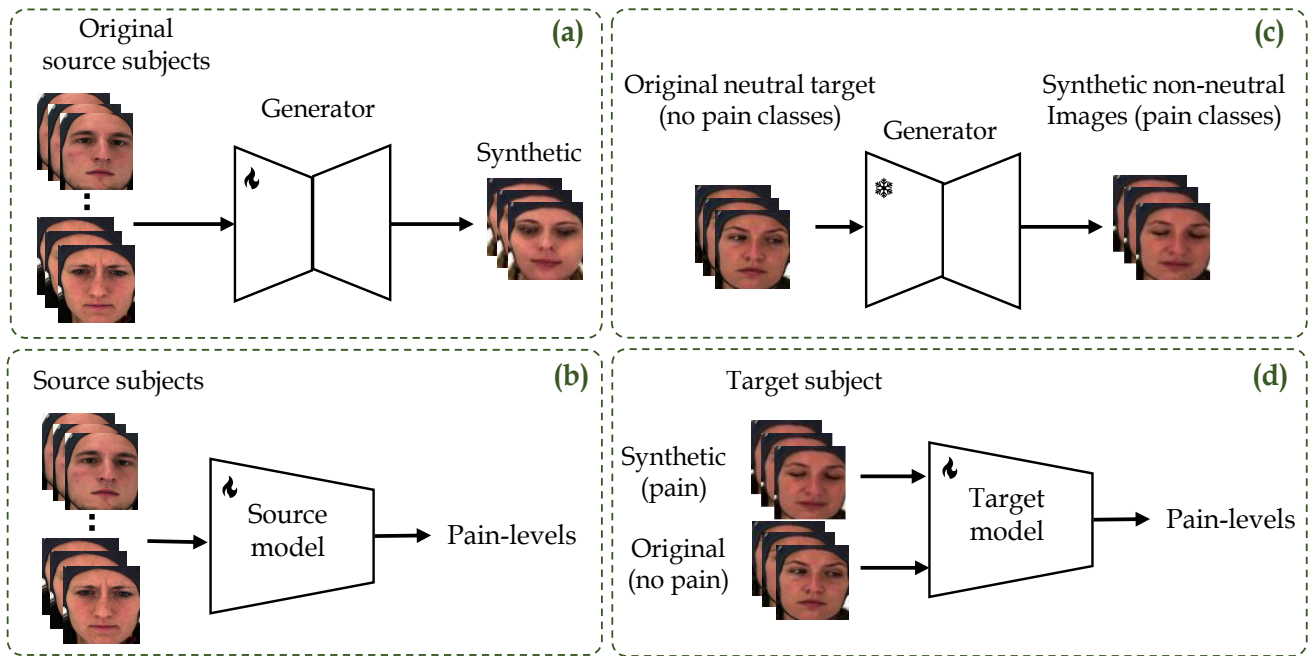
Figure 5: The two-stage method. The first stage (left): (a) the generator is trained using source subjects to generate images with different pain levels, (b) the source model is pre-trained on source subjects. Second stage (right): (c) Pain images/videos for the target subject are generated using a generative model in inference time, (d) the generated pain images/videos are combined with original neutral images for the target subject for source model adaptation.

of real and generated images. In Fig. 8, the generated pain images using our Dis-SFDA method are shown.
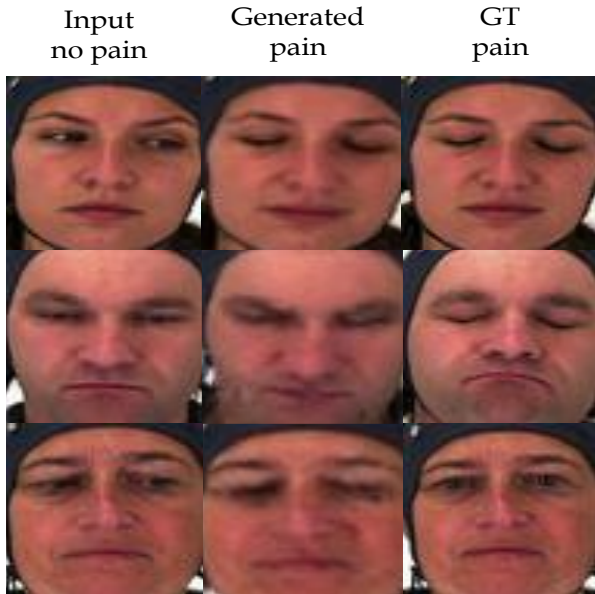

Figure 6: The visual results of generated pain images using the ImaGINator model on the BioVid dataset.

Figure 7: The visual results of generated pain images using the FExGAN-Meta model on the BioVid dataset.

Input
no pain

Generated
pain

GT
pain

Figure 8: The visual results of generated pain images using our Dis-SFDA model on the BioVid dataset.

B.2   Quantitative Results

The quantitative performance of the generative models used to produce synthetic pain images is presented in Table 4. ImaGI-Nator achieves superior results compared to FExGAN-Meta, with higher SSIM (0.82 vs. 0.69) and PSNR (23.00 vs. 22.14), as well as a lower FID score (99.50 vs. 110.58), indicating that ImaGINator generates more realistic and high-quality images, resulting in improved adaptation performance.

We compared the results for the same settings on the UNBC-McMaster dataset using the two-stage framework in Table 5. In the Source Only setting, the pre-trained source model achieves an average accuracy of 61.30%. Using synthetic pain images in the Org. Neutral & Syn. Pain setting improves the results of adaptation significantly compared to the Source Only setting. In this category, SHOT (68.27%) and NRC (68.37%) show comparable results, while the proposed method, SFDA-DE (ours), outperforms both with an average accuracy of 70.41%, demonstrating the effectiveness of our approach. Finally, in the Oracle setting, where the model is fine-tuned on labeled target data, the framework achieves the highest overall accuracy of 97.59%, setting the performance ceiling under ideal conditions.

C   Impact of prototype selection methods.

We evaluated the effectiveness of various image prototype selection strategies from the source domain by comparing random, match, DBSCAN, and k-means clustering methods

and their impact on adaptation performance, as shown in Fig. 9. For DBSCAN, we varied the hyperparameters epsilon (the maximum distance between two points for them to be considered neighbors) and min_samples (the minimum number of points required to form a dense region). These changes result in different clustering outcomes. In DB-All and DB-Sub, the algorithm was applied to all 77 subjects collectively and each subject individually. In DB-D, the most dense cluster for each subject was selected, while in DB-Cls, we selected the cluster closest to the target subject's cluster. For the K-means we used two configurations: K=77 (KM-All), where the closest point to the centroid was chosen for all subjects, and K=1 (KM-Sub), where the closest point to the centroid was selected for each subject individually. Based on the results, subject-specific clustering significantly improved performance in both cases. Applying DBSCAN to individual subjects outperformed the all-subject approach by 3.72%, while for K-means, the improvement was 4.83%. Based on the average results, utilizing matched videos from the source domain significantly enhances performance more than using a random video. This improvement can be attributed to the shift in expression between different subjects; each individual expresses pain in distinct ways and with varying intensity. For instance, facial cues may be more pronounced in women and younger subjects compared to men and older subjects. Consequently, the model can better disentangle expression from other information when using videos that align with the age and gender of the target subjects.
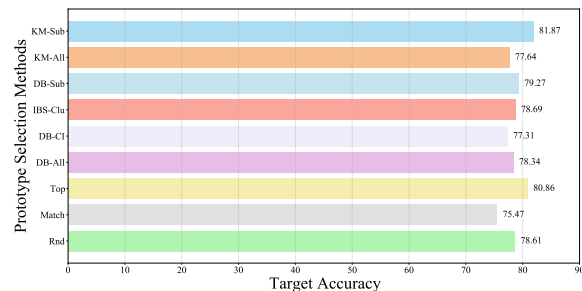

Figure 9: Performance comparison of various prototype selection methods applied to the source data. Methods include random selection (Rnd), direct matching (Match), correctly classified images (Top), DBSCAN clustering across all subjects (DB-All), K-means clustering across all subjects (KM-All), DBSCAN clustering per subject (DB-Sub), K-means clustering per subject (KM-Sub), DBSCAN with the densest cluster (DB-D), and DBSCAN with the closest cluster to the target subject (DB-Cls).

Table 4: Quantitative results of the generative models on BioVid dataset.

Method	SSIM ↑	PSNR ↑	FID ↓
ImaGINator	**0.82**	**23.00**	**99.50**
FExGAN-Meta	0.69	22.14	110.58

Table 5: Accuracy (%) of the proposed method and state-of-the-art methods on the UNBC-McMaster dataset for 5 target subjects with all 20 sources. Bold text shows the highest accuracy.

Setting	Methods	Sub-1	Sub-2	Sub-3	Sub-4	Sub-5	Avg.
Source only	Source model	68.56	78.32	60.37	65.54	33.72	**61.30**
Org. neutral & syn. pain	SHOT [28]	78.24	84.39	60.37	77.19	41.18	**68.27**
	NRC [54]	78.24	84.39	61.00	77.05	41.18	**68.37**
	SFDA-DE (ours)	78.73	84.39	67.88	79.00	42.05	**70.41**
Oracle	Fine-tune	95.63	98.44	98.11	97.54	98.25	**97.59**

References

[1] Ali, K. and Hughes, C. E. (2021). Facial expression recognition by using a disentangled identity-invariant expression representation. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 9460–9467. IEEE.

[2] Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR.

[3] Azari, B. and Lim, A. (2024). Emostyle: One-shot facial expression editing using continuous emotion parameters. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6385–6394.

[4] Brock, A., Donahue, J., and Simonyan, K. (2018). Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*.

[5] Chen, T., Pu, T., Wu, H., Xie, Y., Liu, L., and Lin, L. (2021). Cross-domain facial expression recognition: A unified evaluation benchmark and adversarial graph learning. *IEEE transactions on pattern analysis and machine intelligence*, 44(12):9887–9903.

[6] Chen, X., Duan, Y., Houthooft, R., Schulman, J., Sutskever, I., and Abbeel, P. (2016). Infogan: interpretable representation learning by information maximizing generative adversarial nets. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, page 2180–2188. Curran Associates Inc.

[7] Choi, Y., Uh, Y., Yoo, J., and Ha, J.-W. (2020). Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8188–8197.

[8] Conti, A., Rota, P., Wang, Y., and Ricci, E. (2022). Cluster-level pseudo-labelling for source-free cross-domain facial expression recognition. *arXiv preprint arXiv:2210.05246*.

[9] Dhariwal, P. and Nichol, A. (2021). Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794.

[10] Ding, N., Xu, Y., Tang, Y., Xu, C., Wang, Y., and Tao, D. (2022). Source-free domain adaptation via distribution estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7212–7222.

[11] Fang, Y., Yap, P.-T., Lin, W., Zhu, H., and Liu, M. (2022). Source-free unsupervised domain adaptation: A survey. *arXiv preprint arXiv:2301.00265*.

[12] Guichemerre, A., Belharbi, S., Mayet, T., Murtaza, S., Shamsolmoali, P., McCaffrey, L., and Granger, E. (2024). Source-free domain adaptation of weakly-supervised object localization models for histology. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 33–43.

[13] Guo, Z., Liu, Y., Liu, X., Pan, Z., Liu, S., and Fan, Y. (2023). Ltval: Label transfer virtual adversarial learning framework for source-free facial expression recognition. *Multimedia Tools and Applications*, pages 1–22.

[14] Halawa, M., Wöllhaf, M., Vellasques, E., Sanz, U. S., and Hellwich, O. (2020). Learning disentangled expression representations from facial images. *arXiv preprint arXiv:2008.07001*.

[15] Hernandez, J., McDuff, D., Rudovic, O. O., Fung, A., and Czerwinski, M. (2022). Deepfn: towards generalizable facial action unit recognition with deep face normalization. In *2022 10th International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 1–8. IEEE.

[16] Higgins, I., Amos, D., Pfau, D., Racaniere, S., Matthey, L., Rezende, D., and Lerchner, A. (2018). Towards a definition of disentangled representations. *arXiv preprint arXiv:1812.02230*.

[17] Hong, J., Zhang, Y.-D., and Chen, W. (2022). Source-free unsupervised domain adaptation for cross-modality abdominal multi-organ segmentation. *Knowledge-Based Systems*, 250:109155.

[18] Hou, Y. and Zheng, L. (2020). Source free domain adaptation with image translation. *arXiv preprint arXiv:2008.07514*.

[19] Ji, Y., Hu, Y., Yang, Y., Shen, F., and Shen, H. T. (2019). Cross-domain facial expression recognition via an intra-category common feature and inter-category distinction feature fusion network. *Neurocomputing*, 333:231–239.

[20] Jiang, J. and Deng, W. (2022). Disentangling identity and pose for facial expression recognition. *IEEE Transactions on Affective Computing*, 13(4):1868–1878.

[21] Kim, M., Liu, F., Jain, A., and Liu, X. (2023). Dcface: Synthetic face generation with dual condition diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12715–12725.

[22] Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.

[23] Kurmi, V. K., Subramanian, V. K., and Namboodiri, V. P. (2021). Domain impression: A source data free domain adaptation method. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 615–625.

[24] 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*.

[25] Li, S. and Deng, W. (2018). Deep emotion transfer network for cross-database facial expression recognition. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 3092–3099. IEEE.

[26] Li, S. and Deng, W. (2020). A deeper look at facial expression dataset bias. *IEEE Transactions on Affective Computing*, 13(2):881–893.

[27] Li, Y., Lu, Y., Gong, M., Liu, L., and Zhao, L. (2022). Dual-channel feature disentanglement for identity-invariant facial expression recognition. *Information Sciences*, 608:410–423.

[28] 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 *International conference on machine learning*, pages 6028–6039. PMLR.

[29] Liu, X., Yoo, C., Xing, F., Oh, H., El Fakhri, G., Kang, J.-W., Woo, J., et al. (2022). Deep unsupervised domain adaptation: A review of recent advances and perspectives. *APSIPA Transactions on Signal and Information Processing*, 11(1).

[30] Lucey, P., Cohn, J. F., Prkachin, K. M., Solomon, P. E., and Matthews, I. (2011). Painful data: The unbc-mcmaster shoulder pain expression archive database. In *2011 IEEE International Conference on Automatic Face & Gesture Recognition (FG)*, pages 57–64. IEEE.

[31] Ma, F., Sun, B., and Li, S. (2021). Facial expression recognition with visual transformers and attentional selective fusion. *IEEE Transactions on Affective Computing*, 14(2):1236–1248.

[32] Minaee, S., Minaei, M., and Abdolrashidi, A. (2021). Deep-emotion: Facial expression recognition using attentional convolutional network. *Sensors*, 21(9):3046.

[33] Mohan, K., Seal, A., Krejcar, O., and Yazidi, A. (2021). Fer-net: facial expression recognition using deep neural net. *Neural Computing and Applications*, 33(15):9125–9136.

[34] Nichol, A. Q. and Dhariwal, P. (2021). Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pages 8162–8171. PMLR.

[35] Niinuma, K., Onal Ertugrul, I., Cohn, J. F., and Jeni, L. A. (2022). Facial expression manipulation for personalized facial action estimation. *Frontiers in Signal Processing*, 2:861641.

[36] Rajasekhar, G. P., Granger, E., and Cardinal, P. (2021). Deep domain adaptation with ordinal regression for pain assessment using weakly-labeled videos. *Image and Vision Computing*, 110:104167.

[37] Razavi, A., Van den Oord, A., and Vinyals, O. (2019). Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32.

[38] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695.

[39] Ruan, D., Mo, R., Yan, Y., Chen, S., Xue, J.-H., and Wang, H. (2022). Adaptive deep disturbance-disentangled learning for facial expression recognition. *International Journal of Computer Vision*, 130(2):455–477.

[40] Ruan, D., Yan, Y., Lai, S., Chai, Z., Shen, C., and Wang, H. (2021). Feature decomposition and reconstruction learning for effective facial expression recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7660–7669.

[41] She, J., Hu, Y., Shi, H., Wang, J., Shen, Q., and Mei, T. (2021). Dive into ambiguity: Latent distribution mining and pairwise uncertainty estimation for facial expression recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6248–6257.

[42] Siddiqui, J. R. (2022). Fexgan-meta: Facial expression generation with meta humans. *arXiv preprint arXiv:2203.05975*.

[43] Tian, Q., Ma, C., Zhang, F.-Y., Peng, S., and Xue, H. (2021). Source-free unsupervised domain adaptation with sample transport learning. *Journal of Computer Science and Technology*, 36(3):606–616.

[44] Vahdat, A. and Kautz, J. (2020). Nvae: A deep hierarchical variational autoencoder. *Advances in neural information processing systems*, 33:19667–19679.

[45] Van der Maaten, L. and Hinton, G. (2008). Visualizing data using t-sne. *Journal of machine learning research*, 9(11).

[46] Walter, S., Gruss, S., Ehleiter, H., Tan, J., Traue, H. C., Werner, P., Al-Hamadi, A., Crawcour, S., Andrade, A. O., and da Silva, G. M. (2013). The biovid heat pain database data for the advancement and systematic validation of an automated pain recognition system. In *2013 IEEE international conference on cybernetics (CYBCO)*, pages 128–131. IEEE.

[47] Wang, X., Chen, H., Tang, S., Wu, Z., and Zhu, W. (2022). Disentangled representation learning. *arXiv preprint arXiv:2211.11695*.

[48] Wang, Y., Bilinski, P., Bremond, F., and Dantcheva, A. (2020). Imaginator: Conditional spatio-temporal gan for video generation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1160–1169.

[49] Werner, P., Al-Hamadi, A., and Walter, S. (2017). Analysis of facial expressiveness during experimentally induced heat pain. In *2017 Seventh international conference on affective computing and intelligent interaction workshops and demos (ACIIW)*, pages 176–180. IEEE.

[50] Xie, S., Hu, H., and Chen, Y. (2020). Facial expression recognition with two-branch disentangled generative adversarial network. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(6):2359–2371.

[51] Yang, C., Guo, X., Chen, Z., and Yuan, Y. (2022). Source free domain adaptation for medical image segmentation with fourier style mining. *Medical Image Analysis*, 79:102457.

[52] Yang, G., Tang, H., Zhong, Z., Ding, M., Shao, L., Sebe, N., and Ricci, E. (2021a). Transformer-based source-free domain adaptation. *arXiv preprint arXiv:2105.14138*.

[53] Yang, H. and Yin, L. (2019). Learning temporal information from a single image for au detection. In *2019 14th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2019)*, pages 1–8. IEEE.

[54] Yang, S., Van de Weijer, J., Herranz, L., Jui, S., et al. (2021b). Exploiting the intrinsic neighborhood structure for source-free domain adaptation. *Advances in neural information processing systems*, 34:29393–29405.

[55] Yin, Y., Lu, L., Xiao, Y., Xu, Z., Cai, K., Jiang, H., Gratch, J., and Soleymani, M. (2021). Contrastive learning for domain transfer in cross-corpus emotion recognition. In *2021 9th International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 1–8. IEEE.

[56] Zeeshan, M. O., Aslam, M. H., Belharbi, S., Koerich, A. L., Pedersoli, M., Bacon, S., and Granger, E. (2024). Subject-based domain adaptation for facial expression recognition. In *2024 IEEE 18th International Conference on Automatic Face and Gesture Recognition (FG)*, pages 1–10. IEEE.

[57] Zhang, N., Lu, J., Li, K., Fang, Z., and Zhang, G. (2023). Source-free unsupervised domain adaptation: Current research and future directions. *Neurocomputing*, page 126921.

[58] Zhang, Z., Song, Y., and Qi, H. (2017). Age progression/regression by conditional adversarial autoencoder. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5810–5818.

[59] Zhou, L., Fan, X., Tjahjadi, T., and Das Choudhury, S. (2022). Discriminative attention-augmented feature learning for facial expression recognition in the wild. *Neural Computing and Applications*, 34(2):925–936.

[60] Zhu, Q., Gao, L., Song, H., and Mao, Q. (2021). Learning to disentangle emotion factors for facial expression recognition in the wild. *International Journal of Intelligent Systems*, 36(6):2511–2527.

[61] Zhu, R., Sang, G., and Zhao, Q. (2016). Discriminative feature adaptation for cross-domain facial expression recognition. In *2016 International Conference on Biometrics (ICB)*, pages 1–7. IEEE.