

POSTER: A Pyramid Cross-Fusion Transformer Network for Facial Expression Recognition

Ce Zheng, Matias Mendieta, Chen Chen

Center for Research in Computer Vision, University of Central Florida

{ce.zheng,matias.mendieta}@ucf.edu, chen.chen@crcv.ucf.edu

Abstract

Facial expression recognition (FER) is an important task in computer vision, having practical applications in areas such as human-computer interaction, education, healthcare, and online monitoring. In this challenging FER task, there are three key issues especially prevalent: inter-class similarity, intra-class discrepancy, and scale sensitivity. While existing works typically address some of these issues, none have fully addressed all three challenges in a unified framework. In this paper, we propose a two-stream Pyramid crOss-fuSion TransformerER network (POSTER), that aims to holistically solve all three issues. Specifically, we design a transformer-based cross-fusion method that enables effective collaboration of facial landmark features and image features to maximize proper attention to salient facial regions. Furthermore, POSTER employs a pyramid structure to promote scale invariance. Extensive experimental results demonstrate that our POSTER achieves new state-of-the-art results on RAF-DB (92.05%), FERPlus (91.62%), as well as AffectNet 7 class (67.31%) and 8 class (63.34%). Code is available at <https://github.com/zczcwh/POSTER>.

1. Introduction

Facial expression recognition (FER) is a critical task in computer vision, as it plays a crucial role in understanding human emotions and intentions. FER has various practical applications in fields such as human-computer interaction, education, healthcare, and online monitoring. Thus, it has received increasing interest in recent years, as highlighted in the comprehensive survey by Li et al. [21].

Traditional approaches [46, 47] utilize handcrafted features such as Histogram of oriented Gradients (HOG) [9], Local Binary Patterns (LBP) [32], and SIFT [26] for FER [4, 3, 31]. However, these handcrafted features are often not sufficiently robust and accurate. More recent approaches have focused on leveraging the power of deep

learning, as evidenced by the popularity of methods such as RAN [40], SCN [39], and KTN [20]. These approaches have benefited from large-scale datasets that provide sufficient training data from challenging real-world scenarios and have shown significant performance improvement over traditional methods.

Despite the great progress made so far, there are still several challenges remain in FER:

- Inter-class similarity: Similar images with subtle changes between them can be classified into different expression categories. As illustrated in Fig. 1 (a), A small change in a specific region of an image, such as the mouth, can determine the expression category, even when the overall appearance remains largely unchanged. Due to the subtlety of these differences, current methods may not be sufficiently robust to differentiate between such images.
- Intra-class discrepancy: Within the same expression category, images can have significant differences, such as the skin tone, gender, and age of a person varies across samples, as well as image background appearance. As shown in Fig. 1 (b), two images both represent the expression of happiness but have very different visual appearances.
- Scale sensitivity: When naively applied, deep-learning based networks are often sensitive to image quality and resolution. The image sizes within FER datasets and with in-the-wild testing images vary considerably. Therefore, ensuring consistent performance across scales is critical in FER [37].

In light of these issues, some early works, some earlier works [17, 15] (prior to 2018) attempted to improve the attention to detail and promote invariance to intra-class discrepancies by incorporating facial landmarks into their deep-learning based FER methods. These facial landmarks refer to a set of keypoints on the human face image that provide a sparse representation of key facial regions to complement direct image features. However, these methods typically only use simple concatenation to combine the image and landmark features before the last fully connected layer or prior to a set of basic blocks. While incorporating facial landmarks with image features has potential to allevi-

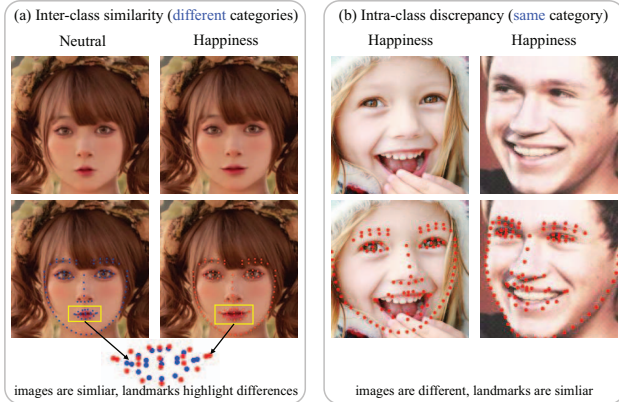


Figure 1. Inter-class similarity and intra-class discrepancy. (The facial landmarks are detected by [16].)

ate intra-class discrepancies and inter-class similarities, just simple concatenation is insufficient for exploring the correlations between landmarks and image features. Moreover, recent works [30, 42] address only intra-class discrepancy and inter-class similarity issues using image features, while another [37] addresses only scale sensitivity, *with none fully addressing all three challenges in FER*. Notably, landmark features, widely used in face-related tasks, have been largely ignored by recent FER approaches.

Therefore, in this paper, we propose a Pyramid crOss-fuSion Transformer network (POSTER) to tackle all three challenges of inter-class similarity, intra-class discrepancy, and scale sensitivity in FER. POSTER is a two-stream architecture that consists of an image stream and a landmark stream. Landmarks pinpoint salient regions and can serve as a guide for focusing feature attention, helping alleviate the inter-class similarity issue. As shown in Fig. 1 (a), the salient region (mouth area) indicating the happiness expression is more easily captured through the discrepancy of landmarks rather than general image features. Also, the sparse landmark features can help to reduce the intra-class discrepancy because they are less sensitive in regard to various skin tones, genders, ages, and background appearance (image features may be affected significantly). On the other hand, image features contain global information of the entire image that landmarks do not cover (such as cheeks, forehead, or a tears drop). Motivated by these intuitions, we propose to explore the correlations between landmark and image features with POSTER. Specifically, we design a transformer-based cross-fusion block that effectively allows the two streams to guide each other, and enables global correlation across features through attention. Experiments validate that the proposed cross-fusion transformer mechanism successfully alleviates inter-class similarity and intra-class discrepancy. Furthermore, to solve the scale-sensitivity for FER, we incorporate a pyramid architecture [24] with the cross-fusion transformer network to capture various resolutions of the extracted feature maps with different informa-

tion granularities. **With our cross-fusion transformer design and integration of a feature pyramid structure, we fully address all three issues in a unified framework to bridge this research gap and achieve new SOTA results on several popular benchmarks.**

Overall, our contributions are summarized as follows:

- We propose a Pyramid crOss-fuSion Transformer network (POSTER) to alleviate inter-class similarity, intra-class discrepancy, and scale sensitivity issues in the FER.
- The cross-fusion transformer structure ensures that image features can be guided by landmark features with prior attention to salient facial regions, while landmark features can utilize global information provided by image features beyond landmarks.
- We extensively validate the efficiency and effectiveness of our proposed POSTER. We show that POSTER outperforms previous state-of-the-art methods on three commonly used datasets. (92.05% on RAF-DB, 67.31% on AffectNet, and 91.62% on FERPlus.)

2. Related Work

Deep learning in FER: With the rapid progress of deep learning in computer vision, such techniques have been found increasingly useful for the challenging FER task [44, 45]. Wang et al. [40] proposed a Region Attention Network (RAN) to capture facial regions for occlusion and pose variant FER. Farzaneh and Qi [12] introduced a Deep Attentive Center Loss (DACL) method to estimate the attention weight for the features for enhancing the discrimination. A sparse center loss was designed to achieve intra-class compactness and inter-class separation with the weighted features. Wang et al. [39] proposed a Self-Cure Network (SCN) to suppress uncertainty, which prevents the network from overfitting incorrectly labeled samples. Shi et al. [34] designed an Amending Representation Module (ARM) to reduce the weight of eroded features and decompose facial features to simplify representation learning.

Facial landmarks in FER: Facial landmark detection aims to estimate the location of predefined keypoints on the human face. The detected facial landmarks are used in many face analysis tasks such as face recognition [35, 25], face tracking [19], and emotion recognition [17, 15, 29]. Significant progress has been made by employing deep learning techniques in the facial landmark detection task, and many accurate facial landmark detectors such as [38, 41, 5, 16] have been proposed. Taking advantage of these off-the-shelf detectors, researchers have paid more attention to utilizing facial landmarks as informative features for the FER task. Jung et al. [17] proposed two networks, where one receives the image as input and another receives facial landmark as input. The output of the two networks is integrated by weighted summation. Hassani et al. [15] proposed a 3D Inception-ResNet where facial landmarks are

multiplied with image features at certain layers. Khan [18] used the facial landmarks to crop small regions first, then generate features as the input for the neural networks. However, existing methods that utilize facial landmarks ignore the correlations of landmark features and image features. Among SOTA methods [33, 34, 42] in the FER task, none of them use facial landmarks.

Vision transformer: The breakthrough of transformer networks in Natural Language Processing (NLP) has sparked great interest in the computer vision domain. In NLP, the transformer is designed to model long sequence inputs. When adapting to the computer vision task, ViT [11] split the image into patches, then the self-attention mechanism in the transformer can capture long-range dependencies across these patches. In FER, Aouayeb et al. [1] directly applied the ViT structure by adding a SE block before the MLP-head for the FER task. Xue et al. [42] proposed a transformer-based method named TransFER. After extracting feature maps with a backbone CNN, the local CNN blocks were designed to locate diverse local patches. Then, a transformer encoder explored the global relationship between these local patches with a multi-head self-attention dropping module.

Compared to previous works, our POSTER employs a two-stream pyramid cross-fusion transformer network to explore the correlation of image features and landmark features to tackle inter-class similarity, intra-class discrepancy, and scale sensitivity issues simultaneously in FER.

3. Methodology

3.1. Baseline

First, we will describe our baseline architecture as shown in Fig. 2 (a). We define an input image $X \in \mathbb{R}^{H \times W \times 3}$, where H and W are the height and width of the image, respectively. In the network, we begin with an image backbone (e.g. IR50 [10]) to get image features $X_{img} \in \mathbb{R}^{P \times D}$. An off-the-shelf facial landmark detector (e.g. MobileFaceNe [7]) is used to obtain the landmark features $X_{lm} \in \mathbb{R}^{P \times D}$, where P is the number of patches (same as the number of landmark keypoints) and D is the feature dimension, respectively. During training, the image backbone is fine-tuned, while the off-the-shelf facial landmark detector is frozen to maintain proper landmark outputs.

After obtaining the image features X_{img} and landmark features X_{lm} , an intuitive solution is to aggregate X_{img} with X_{lm} as the fused features $X_{fuse} \in \mathbb{R}^{2P \times D}$ (concatenate in patch dimension P). The transformer architecture utilizes its self-attention mechanism to capture the correlations across patches. Thus, we directly apply transformer encoders (illustrated in Fig. 3 (a)) to operate on the fused features X_{fuse} .

The self-attention mechanism is achieved by the Multi-

head Self-Attention Layer (MSA) in the transformer architecture. The input $X_{fuse} \in \mathbb{R}^{2P \times D}$ is first mapped to three matrices: the query matrix Q , key matrix K and value matrix V by three linear transformations:

$$Q = X_{fuse}W_Q, \quad K = X_{fuse}W_K, \quad V = X_{fuse}W_V, \quad (1)$$

where W_Q, W_K and $W_V \in \mathbb{R}^{D \times D}$.

The vanilla transformer attention block is illustrated in Fig. 4 (a) can be described as the following mapping function:

$$\text{Attention}(Q, K, V) = \text{Softmax}(QK^\top / \sqrt{d})V, \quad (2)$$

where $\frac{1}{\sqrt{d}}$ is the scaling factor for appropriate normalization to prevent extremely small gradients.

Next, the vanilla transformer encoder architecture consisting of MSA and MLP with a layer normalization operator is shown in Fig. 3 (a). The encoder output $X_{fuse_out} \in \mathbb{R}^{2P \times D}$ keeps the same size as the encoder input $X_{fuse} \in \mathbb{R}^{2P \times D}$, and is represented as follows:

$$X'_{fuse} = \text{MSA}(Q, K, V) + X_{fuse}, \quad (3)$$

$$X_{fuse_out} = \text{MLP}(\text{Norm}(X'_{fuse})) + X'_{fuse}, \quad (4)$$

where $\text{MSA}(\cdot)$ represents the Multi-head Self-Attention block, $\text{Norm}(\cdot)$ is the normalization operator, and $\text{MLP}(\cdot)$ denotes the multilayer perceptron. Finally, an MLP head returns the predicted emotion label $Y \in \mathbb{R}^N$ where N is the number of classes.

Although the transformer encoders can inherently model the image features and landmark features jointly to some extent with simple concatenation of X_{img} with X_{lm} , the strong correlations between image features and landmark features are not fully exploited by the vanilla transformer since features are concatenated together. In the next section, we describe POSTER which improves the correlation and representation capability of the model.

3.2. POSTER

The facial landmarks locate a set of keypoints in the human face region, which pinpoints salient regions related to facial expression. At the same time, global information beyond landmarks is also important for recognizing human expression (e.g. cheeks, wrinkled forehead). Motivated by this, we design a two-stream network that consists of both an image and a landmark stream. We swap the key matrices of the two streams, creating a cross-fusion operation to facilitate feature collaboration. Specifically, by performing this operation, we enable the image features to be guided by some prior knowledge of salient regions from the landmarks. Likewise, the representations of the landmark stream are provided with global context from the image features while moving through the block operations. In this

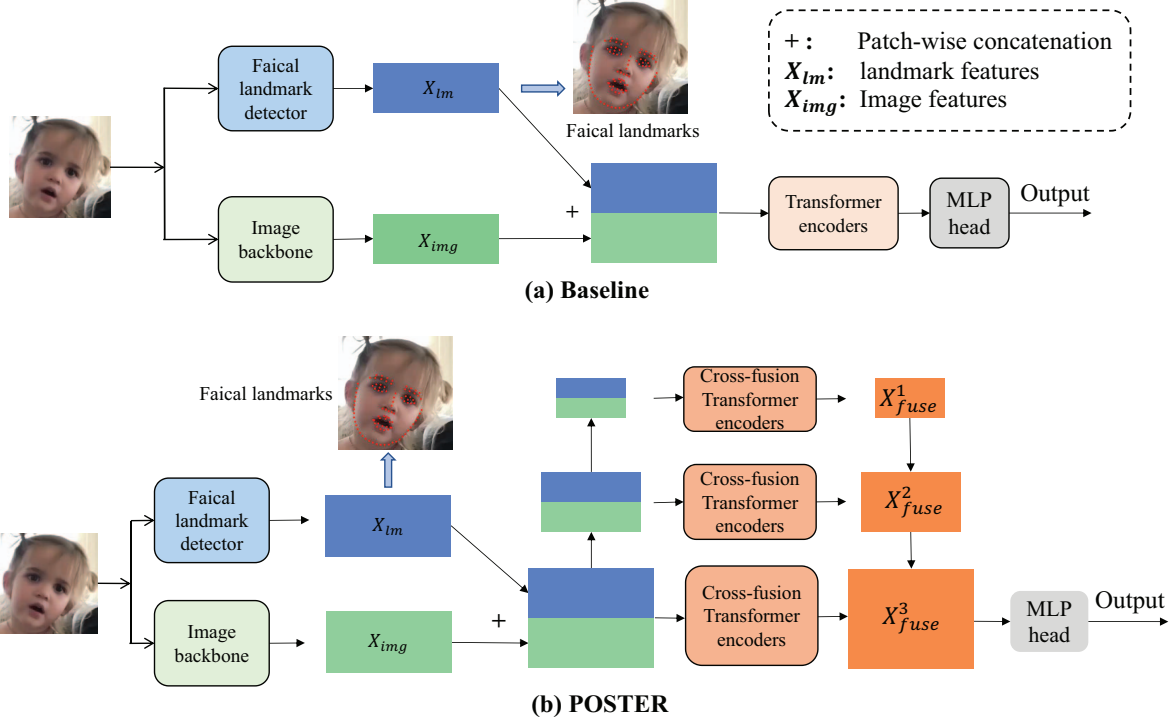


Figure 2. The architectures of baseline (a) and our proposed POSTER (b) for Facial Expression Recognition (FER). A facial landmark detector (MobileFaceNet [7]) is applied to obtain landmark features X_{lm} . The image backbone (IR50 [10]) is used to extract image features X_{img} . "+" denotes patch-wise concatenation operation.

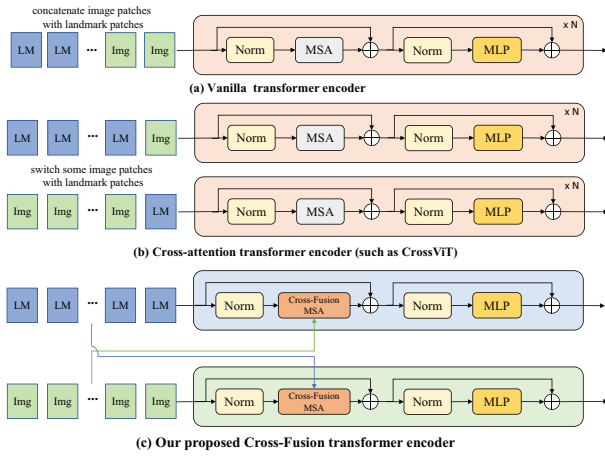


Figure 3. (a) The vanilla transformer encoder (from ViT [11]). (b) The cross-attention transformer encoder (such as CrossViT) (c) Our proposed cross-fusion transformer encoder.

way, we foster improved contextual understanding to alleviate intra-class discrepancy and inter-class similarity. Therefore, given the extracted image features X_{img} and landmark features X_{lm} , we design the cross-fusion MSA blocks as shown in Fig. 4 (b) to achieve our goal.

Cross-fusion transformer encoder: For the MSA in the image stream, the input $X_{img} \in \mathbb{R}^{P \times D}$ is mapped to three image matrices: query matrix Q_{img} , key matrix K_{img} and value matrix V_{img} by three linear transformations. For the

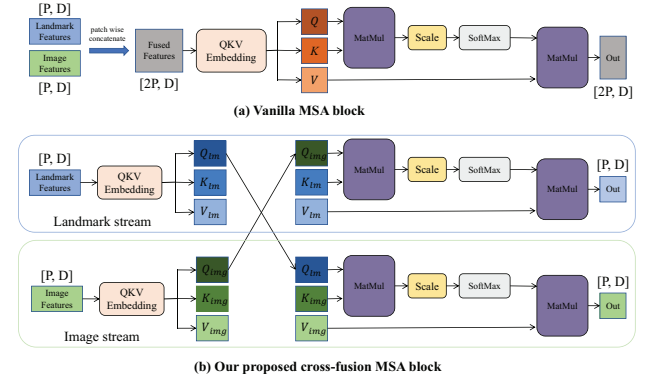


Figure 4. (a) The vanilla MSA block in transformer encoder. (b) Our proposed cross-fusion MSA block in cross-fusion transformer encoder. P denotes the number of patches and D is the embedding dimension.

landmark stream, the input $X_{lm} \in \mathbb{R}^{P \times D}$ is mapped to three landmark matrices: query matrix Q_{lm} , key matrix K_{lm} and value matrix V_{lm} by three linear transformations:

$$Q_{img} = X_{img}W_{Q1}, K_{img} = X_{img}W_{K1}, V_{img} = X_{img}W_{V1}, \quad (5)$$

$$Q_{lm} = X_{lm}W_{Q2}, K_{lm} = X_{lm}W_{K2}, V_{lm} = X_{lm}W_{V2}, \quad (6)$$

where $W_{Q1}, W_{Q2}, W_{K1}, W_{K2}, W_{V1}$ and $W_{V2} \in \mathbb{R}^{D \times D}$. The cross-fusion transformer block illustrated in Fig. 3

(b) can be described as the following mapping functions:

$$\text{Attention}_{(img)} = \text{Softmax}(Q_{lm}K_{img}^\top/\sqrt{d})V_{img}, \quad (7)$$

$$\text{Attention}_{(lm)} = \text{Softmax}(Q_{img}K_{lm}^\top/\sqrt{d})V_{lm}, \quad (8)$$

where $\frac{1}{\sqrt{d}}$ is the scaling factor for appropriate normalization to prevent extremely small gradients. The queries Q_{img} and Q_{lm} are swapped between the image and landmark streams. By doing this, the image features are empowered with salient regions provided by landmark features. On the other hand, the landmark features are provided with global information from the image features.

The cross-fusion transformer encoder is shown in Fig. 3 (b). The outputs of the cross-fusion transformer encoder X_{img_out} and X_{lm_out} given the image features X_{img} and landmark features X_{lm} are represented as follows:

$$X'_{img} = \text{CFMSA}_{img}(Q_{lm}, K_{img}, V_{img}) + X_{img}, \quad (9)$$

$$X_{img_out} = \text{MLP}(\text{Norm}(X'_{img})) + X'_{img}, \quad (10)$$

$$X'_{lm} = \text{CFMSA}_{lm}(Q_{img}, K_{lm}, V_{lm}) + X_{lm}, \quad (11)$$

$$X_{lm_out} = \text{MLP}(\text{Norm}(X'_{lm})) + X'_{lm}, \quad (12)$$

where $\text{CFMSA}(\cdot)$ represents the Cross-Fusion MSA block in the image stream or the landmark stream, $\text{Norm}(\cdot)$ is the normalization operator, and $\text{MLP}(\cdot)$ denotes the multilayer perceptron.

Feature pyramid structure: In FER, the image quality and resolution often vary considerably. Therefore, to support applicability across image scales, we adapt the well-known and effective technique of feature pyramid structure [24] for producing a multi-scale feature representation.

Specifically, we construct large/medium/small levels of extracted features as shown in Fig. 2 (b). The large/medium/small features are fed to separate cross-fusion transformer encoders to capture specific feature scales. The output of the three cross-fusion transformer encoders is aggregated to form the emotion feature. Finally, an MLP head returns the predicted emotion label $Y \in \mathbb{R}^N$, where N is the number of classes.

Distinction with other cross-attention design: There are some methods proposed cross-attention transformer such as CrossViT [8]. The goal is to switch the patches from two different modalities as illustrated in Fig 3 (b). Different from this CrossViT-style transformer encoder, the motivation of our proposed cross-fusion transformer encoder is the feature collaboration. We enable the image features to be guided by some prior knowledge of salient regions from the landmarks. Meanwhile, the representations of the landmark stream are provided with global context from the image features while moving through the block operations.

Take-away: Overall, we revitalize the idea of a two-stream architecture with our unique transformer-based cross-fusion

method to alleviate inter-class similarity and intra-class discrepancy. Furthermore, we employ pyramid structures in POSTER to tackle scale-sensitivity, and therefore intelligently formulate a unified framework that addresses the three key challenges of FER. We verify the effectiveness of POSTER and its components in the following section.

4. Experiments

4.1. Datasets

RAF-DB: Real-world Affective Faces Database (RAF-DB) [22] is a large-scale facial expression dataset with 29,672 real-world facial images. All images are collected from the Internet with great variability in the subject’s age, gender, ethnicity, lighting conditions, occlusions, etc. For the FER task, there are 15,339 facial expression images utilized (12,271 images are used for training and 3,068 images are used for testing) with seven basic expressions (happiness, surprise, sadness, anger, disgust, fear, and neutral).

FERPlus: FERPlus [2] is extended from FER2013 [13] as used in the ICML 2013 Challenges. It is a large-scale dataset collected by APIs in the Google search engine. For the FER task, it contains 28,709 training images, 3,589 validation images, and 3,589 testing images. FERPlus relabeled the FER2013 into eight emotion categories (seven basic expressions plus contempt). Following [42] and [20], we report the overall accuracy on the test set.

AffectNet: AffectNet [27] is one of the largest publicly available datasets for FER task. It is a large-scale in-the-wild dataset that contains more than 1M facial images collected from the Internet by querying three major search engines using 1250 emotion-related keywords in six different languages. Images are labeled into eight emotion categories (seven basic expressions plus contempt).

4.2. Implementation Details

We implemented POSTER with Pytorch [28] on two NVIDIA RTX 3090 GPUs. in an end-to-end manner. We utilized IR50 [10] as the image backbone, which is pre-trained on Ms-Celeb-1M dataset [14]. The weights of the image backbone are updated during training. For the facial landmark detector, we select MobileFaceNet [7] with all of the weights frozen to ensure it outputs landmark features. In the feature pyramid structure, we produce large/medium/small extracted features with embedding dim $D_H = 512$, $D_M = 256$, and $D_L = 128$, respectively. For the cross-fusion transformer encoders, each level of encoders consists of $depth = 8$ transformer encoders. The mlp ratio and drop path rate in transformer encoders are 2 and 0.01, respectively. We set the batch size to 100 with a learning rate of 4×10^{-5} . Unlike many methods [12, 20] that rely on complicated loss, we use the standard label smoothing cross-entropy loss.

Table 1. Comparison on RAF-DB, AffectNet, and FERPlus datasets. “mean Acc” denotes mean class accuracy and “cls” is the abbreviation of classes.

Method	Year	RAF-DB		AffectNet		FERPlus
		Acc	mean Acc	Acc (7 cls)	Acc (8 cls)	Acc
SCN[39]	CVPR 2020	87.03	-	-	60.23	89.39
PSR[37]	CVPR 2020	88.98	80.78	63.77	60.68	-
RAN[40]	TIP 2020	86.90	-	-	-	89.16
DACL[12]	WACV 2021	87.78	80.44	65.20	-	-
KTN[20]	TIP 2021	88.07	-	63.97	-	90.49
DMUE[33]	CVPR 2021	89.42	-	63.11	-	-
FDRL[30]	CVPR 2021	89.47	-	-	-	-
ARM[34]	arXiv 2021	90.42	82.77	65.20	61.33	-
TransFER[42]	ICCV 2021	90.91	-	66.23	-	90.83
Face2Exp[43]	CVPR 2022	88.54	-	64.23	-	-
EAC[45]	ECCV 2022	89.99	-	65.32	-	89.64
FER-former[23]	arXiv 2023	91.30	-	-	-	90.96
Baseline	-	91.00	84.64	65.06	60.94	90.91
POSTER	-	92.05	86.03	67.31	63.34	91.62

4.3. Comparison with State-of-the-art Results

Evaluation on RAF-DB: Table 1 compares POSTER with previous methods on the RAF-DB dataset. POSTER outperforms the SOTA methods both in terms of accuracy (the accuracy of all samples) and mean accuracy (the average of the accuracy of each category). Our POSTER yields the highest accuracy of 92.05 %, which is 1.14 % better than the second-best method (TransFER [42]). Our POSTER also achieves the highest score of 86.03 % in mean accuracy, which is 3.26 % better than the second-best method (ARM [34], as TransFER did not report the mean accuracy).

Evaluation on AffectNet: Table 1 reports the results of POSTER with previous methods on the AffectNet dataset. AffectNet is the largest publicly available dataset with challenging facial expressions. The training set is extremely unbalanced (134,415 happiness images but only 3,803 disgust images; the gap is over 35×). In terms of seven expression categories, our POSTER achieves the best accuracy of 67.31 %, which is 1.08 % higher than the second-best method (TransFER [42]). When considering the Contempt category (total eight classes), Our POSTER also yields the highest accuracy of 63.34 %, which is 2.01 % better than the second-best method (ARM [34]).

Evaluation on FERPlus: Table 1 evaluates POSTER with previous methods on FERPlus dataset. All images in FERPlus are grayscale images with small resolutions (48×48). Our POSTER still achieves the best accuracy of 91.62 %, which is 0.79 % higher than the second-best method (TransFER [42]). Overall, the superior performance on all three datasets has demonstrated the effectiveness of POSTER.

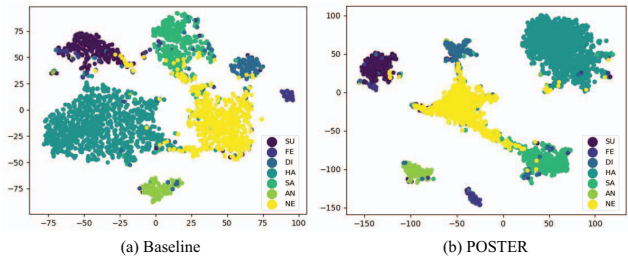


Figure 5. Visualisation of high dimensional features using t-SNE[36] for the baseline and POSTER on the RAF-DB dataset.

4.4. Results analysis

In order to better understand the effectiveness of POSTER and its design, we perform further analysis on individual class performance. Specifically, we first analyze the class-specific evaluations of our baseline and POSTER on RAF-DB, FERPlus, and AffectNet (7 classes and 8 classes) in Tables 2 and 3. For **RAF-DB**, the accuracy of Neutral, Happy, Sad, and Surprise categories of POSTER is higher than 90% as shown in Table 2. However, the accuracy of Fear category is relatively low. This is due to insufficient training samples on this category (281 of Fear images, while others such as Neutral, Happy, Sad, and Surprise are more than 1,000). Similar to the RAF-DB dataset, the accuracy of Disgust and Contempt categories in the **FERPlus** are relatively low as shown in Table 2. This is also due to insufficient training samples on these two categories (only 116 of Disgust and 120 of Contempt, more than 20× less than other categories). For **AffectNet**, all images are collected from the internet, most of the training and testing samples are from in-the-wild settings. The extremely unbalanced training set and challenging in-the-wild samples make the performance quite lower than in RAF-

Table 2. Class-wise accuracy on RAF-DB, AffectNet, and FERPlus datasets. The blue color indicates the intra-class discrepancy has been reduced.

Dataset	Method	cls Acc								mean Acc
		Neutral	Happy	Sad	Surprise	Fear	Disgust	Anger	Contempt	
RAF-DB	Baseline	90.44	96.71	90.38	87.23	62.16	76.25	88.89	-	84.58
	POSTER	92.35	96.96	91.21	90.27	67.57	75.00	88.89	-	86.04
AffectNet (7 cls)	Baseline	64.00	87.80	62.60	63.60	65.40	56.00	56.40	-	65.11
	POSTER	67.20	89.00	67.00	64.00	64.80	56.00	62.60	-	67.23
AffectNet (8 cls)	Baseline	57.20	74.60	61.20	63.00	63.40	61.40	52.00	54.71	60.94
	POSTER	59.40	80.20	66.60	63.60	63.60	59.80	58.80	54.71	63.34
FERPlus	Baseline	92.52	95.52	91.88	85.60	91.08	50.00	62.79	13.33	72.84
	POSTER	93.26	96.08	92.39	85.86	91.45	43.75	68.60	33.33	75.59

Table 3. Prediction details given the target class on RAF-DB dataset. The blue color indicates the intra-class discrepancy has been reduced. The red color denotes the inter-class similarity has been alleviated.

	Ground truth	Prediction Percentage							
		Neutral	Happy	Sad	Surprise	Fear	Disgust	Anger	
Baseline	Neutral	90.44	2.50	4.56	1.62	0.00	0.74	0.15	
POSTER	Neutral	92.35	1.91	4.26	1.32	0.00	0.15	0.00	
Baseline	Fear	2.70	5.41	10.81	14.86	62.16	2.70	1.35	
POSTER	Fear	4.05	2.70	9.46	12.16	67.57	2.70	1.35	
Baseline	Surprise	5.47	0.91	0.91	87.23	1.52	1.82	2.13	
POSTER	Surprise	2.74	0.91	0.30	90.27	2.43	1.82	1.52	

DB and FERPlus datasets. The accuracy of Happy category achieves 80% in Table 2 since there are 134,415 happiness samples (almost 50 % of the total training images).

When comparing the class-wise accuracy in Table 2, POSTER significantly increases the accuracy of most of the categories on three datasets, which indicates that POSTER reduces intra-class discrepancy for FER (marked in blue color if the class accuracy has been improved). In Table 3, we show the percentage of prediction given the target categories. For example, given a total of 680 Neutral testing images, 628 images are correctly classified as Neutral by POSTER (92.35 %) and 13 images are classified incorrectly as Happy (1.91 %). The percentage of the wrong prediction into certain categories also decreases as marked in red color, which shows that POSTER alleviates intra-class discrepancy for FER. Moreover, we visualize the high dimensional features before final output of our baseline and POSTER by t-SNE [36]. The improved separation across classes and density with class clusters also verifies that POSTER alleviates the intra-class discrepancy and inter-class similarity.

4.5. Ablation Study

We conduct the ablation study on RAF-DB and AffectNet datasets to verify the contribution of the proposed structures and the impact of hyperparameters on performance. More ablation studies are provided in the supplementary.

Effectiveness of the architecture design: We investigate the different architectures of utilizing image features and landmark features and report the results in Table 4. In Fig. 6 (a) and (b), we only use the landmark features and image features, respectively. Utilizing image features achieves

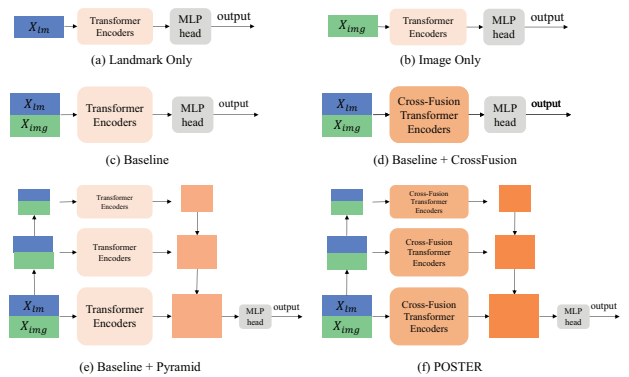


Figure 6. (a) Only use landmark features X_{lm} for classification. (b) Only use image features X_{img} for classification. (c) Our baseline in Sec. 3.1. (d) Our baseline with replacing vanilla transformer encoders with proposed cross-fusion transformer encoders. (e) Our baseline with adding a pyramid structure. (f) Our proposed POSTER in Sec. 3.2.

Table 4. Ablation study on the proposed components.

Components	RAF-DB		AffectNet	
	Acc	Acc(mean)	7 cls	8 cls
(a) Landmark only	80.08	72.21	49.88	45.34
(b) Image only	90.51	82.73	64.95	56.60
(c) Baseline	91.00	84.64	65.06	60.94
(d) Baseline+pyramid	91.27	85.66	66.50	62.36
(e) Baseline+cross_fusion	91.63	85.01	65.35	61.87
(f) POSTER	92.05	86.03	67.31	63.34

better results than landmark features. When concatenating image features with landmark features as a baseline in Fig. 6 (c) (discussed in Sec. 3.1), the performance

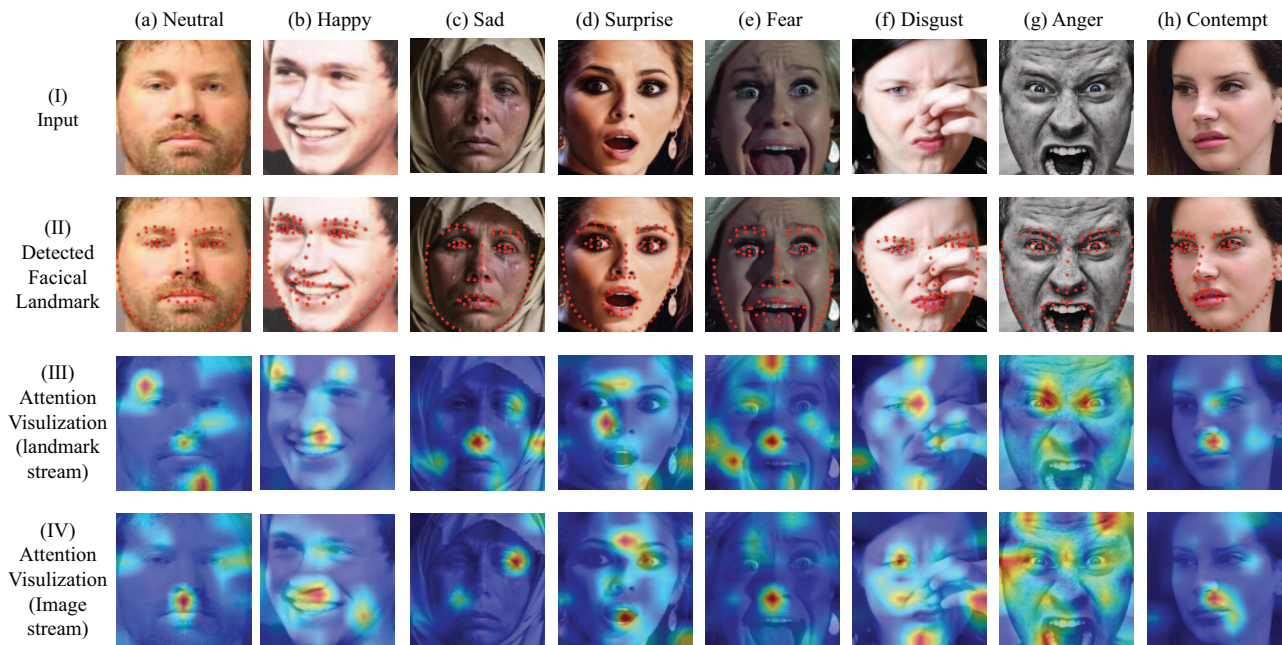


Figure 7. Attention visualization on facial images of different categories (images are from the AffectNet dataset).

is better than image features only and landmark features only. When replacing the vanilla transformer attention blocks with our proposed cross-fusion transformer attention blocks as shown in Fig. 6 (d), the performance can be boosted, which verifies the effectiveness of our proposed cross-fusion design. Next, we report the performance with our pyramid structure illustrated in Fig. 6 (e). When comparing (c) with (d), and (e) with (f), the performance has improved because the pyramid structure can alleviate scale sensitivity issues given various input sizes and quality of training samples. Finally, our POSTER in Fig. 6 (f) achieves the best results.

Attention Visualization: In order to visualize the parts of the facial image that contributes to the category clarification, we apply the visualization method of [6] to visualize the attention maps in the transformer. The relevance score is computed for each attention head in each layer of the transformer encoders, and these scores are then integrated by incorporating both relevancy and gradient information. More details are discussed in [6].

The attention visualizations and detected facial landmarks for a set of facial images are shown in Fig. 7. The images of 8 categories are selected from the in-the-wild AffectNet dataset. Since our proposed cross-fusion transformer attention block has two streams (landmark and image), we visualize the attention map of each stream separately. The attention learned by these two streams are different but complementary.

For the landmark stream attention map, the highlight areas are commonly from the landmark areas. However, it

also obtains some global attention beyond landmarks from the image features, such as cheek and forehead areas in (iii, a), (iii, e), and (iii, h). For the image stream attention maps, the highlight areas indicate the discriminate patterns that led to a certain classification. For example, the tears drop in (iv, c) is discovered, which indicates a sad expression. The forehead wrinkles in (iv, g) are highlighted which relates to an angry expression. Because landmark features are cross-fused as guidance, the attention can focus on important regions related to landmarks. Moreover, some areas such as between the eyebrows and above the brow are also captured by image-guided features to improve the indication of the expression. By integrating attention from both landmark and image streams, POSTER can effectively recognize facial expressions with excellent performance.

5. Conclusions

In this paper, we have proposed a Pyramid crOss-fuSion Transformer network (POSTER) for the FER task. As a two-stream network, landmark features are detected by an off-the-shelf facial landmark detector, and image features are extracted by a backbone CNN. The correlations of image features and landmark features are fully exploited in POSTER by our proposed cross-fusion transformer architecture. POSTER tackles all three challenges of inter-class similarity, intra-class discrepancy, and scale sensitivity in FER. Extensive experiments on three commonly used FER datasets have demonstrated that POSTER outperforms state-of-the-art methods.

References

- [1] Mouath Aouayeb, Wassim Hamidouche, Catherine Soladie, Kidiyo Kpalma, and Renaud Seguier. Learning vision transformer with squeeze and excitation for facial expression recognition. *arXiv preprint arXiv:2107.03107*, 2021. [3](#)
- [2] Emad Barsoum, Cha Zhang, Cristian Canton Ferrer, and Zhengyou Zhang. Training deep networks for facial expression recognition with crowd-sourced label distribution. In *ACM International Conference on Multimodal Interaction (ICMI)*, 2016. [5](#)
- [3] Stefano Berretti, Boulbaba Ben Amor, Mohamed Daoudi, and Alberto Del Bimbo. 3d facial expression recognition using sift descriptors of automatically detected keypoints. *The Visual Computer*, 27(11):1021–1036, 2011. [1](#)
- [4] Pierluigi Carcagnì, Marco Del Coco, Marco Leo, and Cosimo Distantè. Facial expression recognition and histograms of oriented gradients: a comprehensive study. *SpringerPlus*, 4(1):1–25, 2015. [1](#)
- [5] Prashanth Chandran, Derek Bradley, Markus Gross, and Thabo Beeler. Attention-driven cropping for very high resolution facial landmark detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5861–5870, 2020. [2](#)
- [6] Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 782–791, 2021. [8](#)
- [7] Cunjian Chen. PyTorch Face Landmark: A fast and accurate facial landmark detector, 2021. [3](#), [4](#), [5](#)
- [8] Chun-Fu Richard Chen, Quanfu Fan, and Rameswar Panda. Crossvit: Cross-attention multi-scale vision transformer for image classification. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 357–366, 2021. [5](#)
- [9] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*. Ieee, 2005. [1](#)
- [10] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019. [3](#), [4](#), [5](#)
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. [3](#), [4](#)
- [12] Amir Hossein Farzaneh and Xiaojun Qi. Facial expression recognition in the wild via deep attentive center loss. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2402–2411, 2021. [2](#), [5](#), [6](#)
- [13] Ian J Goodfellow, Dumitru Erhan, Pierre Luc Carrier, Aaron Courville, Mehdi Mirza, Ben Hamner, Will Cukierski, Yichuan Tang, David Thaler, Dong-Hyun Lee, et al. Challenges in representation learning: A report on three machine learning contests. In *International conference on neural information processing*, pages 117–124. Springer, 2013. [5](#)
- [14] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *European conference on computer vision*, pages 87–102. Springer, 2016. [5](#)
- [15] Behzad Hasani and Mohammad H Mahoor. Facial expression recognition using enhanced deep 3d convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2017. [1](#), [2](#)
- [16] Haibo Jin, Shengcai Liao, and Ling Shao. Pixel-in-pixel net: Towards efficient facial landmark detection in the wild. *International Journal of Computer Vision*, 2021. [2](#)
- [17] Heechul Jung, Sihaeng Lee, Junho Yim, Sunjeong Park, and Junmo Kim. Joint fine-tuning in deep neural networks for facial expression recognition. In *Proceedings of the IEEE international conference on computer vision*, 2015. [1](#), [2](#)
- [18] Fuzail Khan. Facial expression recognition using facial landmark detection and feature extraction via neural networks. *arXiv preprint arXiv:1812.04510*, 2018. [3](#)
- [19] Muhammad Haris Khan, John McDonagh, and Georgios Tzimiropoulos. Synergy between face alignment and tracking via discriminative global consensus optimization. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 3811–3819. IEEE, 2017. [2](#)
- [20] Hangyu Li, Nannan Wang, Xinpeng Ding, Xi Yang, and Xinbo Gao. Adaptively learning facial expression representation via cf labels and distillation. *IEEE Transactions on Image Processing*, 30:2016–2028, 2021. [1](#), [5](#), [6](#)
- [21] Shan Li and Weihong Deng. Deep facial expression recognition: A survey. *IEEE transactions on affective computing*, 13(3):1195–1215, 2020. [1](#)
- [22] Shan Li, Weihong Deng, and JunPing Du. Reliable crowd-sourcing and deep locality-preserving learning for expression recognition in the wild. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2584–2593. IEEE, 2017. [5](#)
- [23] Yande Li, Mingjie Wang, Minglun Gong, Yonggang Lu, and Li Liu. Fer-former: Multi-modal transformer for facial expression recognition. *arXiv preprint arXiv:2303.12997*, 2023. [6](#)
- [24] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. [2](#), [5](#)
- [25] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Spheraface: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 212–220, 2017. [2](#)
- [26] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004. [1](#)
- [27] Ali Mollahosseini, Behzad Hasani, and Mohammad H Mahoor. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing*, 10(1):18–31, 2017. [5](#)

- [28] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. [5](#)
- [29] Yinghong Qiu and Yi Wan. Facial expression recognition based on landmarks. In *2019 IEEE 4th Advanced Information Technology, Electronic and Automation Control Conference (IAEAC)*. IEEE, 2019. [2](#)
- [30] Delian Ruan, Yan Yan, Shenqi Lai, Zhenhua Chai, Chunhua Shen, and Hanzi Wang. 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, 2021. [2](#), [6](#)
- [31] Caifeng Shan, Shaogang Gong, and Peter W McOwan. Robust facial expression recognition using local binary patterns. In *IEEE International Conference on Image Processing 2005*, volume 2, pages II–370. IEEE, 2005. [1](#)
- [32] Caifeng Shan, Shaogang Gong, and Peter W McOwan. Facial expression recognition based on local binary patterns: A comprehensive study. *Image and vision Computing*, 2009. [1](#)
- [33] Jiahui She, Yibo Hu, Hailin Shi, Jun Wang, Qiu Shen, and Tao Mei. 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, 2021. [3](#), [6](#)
- [34] Jiawei Shi, Songhao Zhu, and Zhiwei Liang. Learning to amend facial expression representation via de-albino and affinity. *arXiv preprint arXiv:2103.10189*, 2021. [2](#), [3](#), [6](#)
- [35] Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, 2014. [2](#)
- [36] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. [6](#), [7](#)
- [37] Thanh-Hung Vo, Guee-Sang Lee, Hyung-Jeong Yang, and Soo-Hyung Kim. Pyramid with super resolution for in-the-wild facial expression recognition. *IEEE Access*, 8:131988–132001, 2020. [1](#), [2](#), [6](#)
- [38] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *TPAMI*, 2019. [2](#)
- [39] Kai Wang, Xiaojiang Peng, Jianfei Yang, Shijian Lu, and Yu Qiao. Suppressing uncertainties for large-scale facial expression recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6897–6906, 2020. [1](#), [2](#), [6](#)
- [40] Kai Wang, Xiaojiang Peng, Jianfei Yang, Debin Meng, and Yu Qiao. Region attention networks for pose and occlusion robust facial expression recognition. *IEEE Transactions on Image Processing*, 29:4057–4069, 2020. [1](#), [2](#), [6](#)
- [41] Qingzhong Wang, Pengfei Zhang, Haoyi Xiong, and Jian Zhao. Face. evolve: A high-performance face recognition library. *arXiv preprint arXiv:2107.08621*, 2021. [2](#)
- [42] Fanglei Xue, Qiangchang Wang, and Guodong Guo. Transfer: Learning relation-aware facial expression representations with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3601–3610, October 2021. [2](#), [3](#), [5](#), [6](#)
- [43] Dan Zeng, Zhiyuan Lin, Xiao Yan, Yuting Liu, Fei Wang, and Bo Tang. Face2exp: Combating data biases for facial expression recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20291–20300, 2022. [6](#)
- [44] Yuhang Zhang, Chengrui Wang, and Weihong Deng. Relative uncertainty learning for facial expression recognition. *Advances in Neural Information Processing Systems*, 34:17616–17627, 2021. [2](#)
- [45] Yuhang Zhang, Chengrui Wang, Xu Ling, and Weihong Deng. Learn from all: Erasing attention consistency for noisy label facial expression recognition. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVI*, pages 418–434. Springer, 2022. [2](#), [6](#)
- [46] Guoying Zhao and Matti Pietikainen. Dynamic texture recognition using local binary patterns with an application to facial expressions. *IEEE transactions on pattern analysis and machine intelligence*, 2007. [1](#)
- [47] Lin Zhong, Qingshan Liu, Peng Yang, Bo Liu, Junzhou Huang, and Dimitris N Metaxas. Learning active facial patches for expression analysis. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2012. [1](#)