

QCS: Feature Refining from Quadruplet Cross Similarity for Facial Expression Recognition

Chengpeng Wang¹, Li Chen³, Lili Wang³, Zhaofan Li⁴, Xuebin Lv^{2*}

¹Wisefsoft Inc., Chengdu, China

²School of Computer Science Sichuan University, Chengdu, China

³Chinese PLA General Hospital, Beijing, China

⁴Medical School of Chinese PLA, Beijing, China

Abstract

Facial expression recognition faces challenges where labeled significant features in datasets are mixed with unlabeled redundant ones. In this paper, we introduce Cross Similarity Attention (CSA) to mine richer intrinsic information from image pairs, overcoming a limitation when the Scaled Dot-Product Attention of ViT is directly applied to calculate the similarity between two different images. Based on CSA, we simultaneously minimize intra-class differences and maximize inter-class differences at the fine-grained feature level through interactions among multiple branches. Contrastive residual distillation is utilized to transfer the information learned in the cross module back to the base network. We ingeniously design a four-branch centrally symmetric network, named Quadruplet Cross Similarity (QCS), which alleviates gradient conflicts arising from the cross module and achieves balanced and stable training. It can adaptively extract discriminative features while isolating redundant ones. The cross-attention modules exist during training, and only one base branch is retained during inference, resulting in no increase in inference time. Extensive experiments show that our proposed method achieves state-of-the-art performance on several FER datasets.

Code — <https://github.com/birdwcp/QCS>

Extended version — <https://arxiv.org/abs/2411.01988>

Introduction

Facial expression is crucial for machines to comprehend human emotions. Facial expression recognition (FER) holds vast research potential and application worth in human-computer interaction, emotion analysis, and intelligent healthcare. Over decades, deep learning advancements in computer vision have propelled significant progress in facial expression recognition.

However, FER is still a complex task facing many difficulties such as Inter-class similarity and intra-class variances. Faces with varying expressions retain the common features that constitute a human face, yet same expressions can vary widely in pose, identity, illumination, etc. Traditional object classification methods struggle with complex face datasets

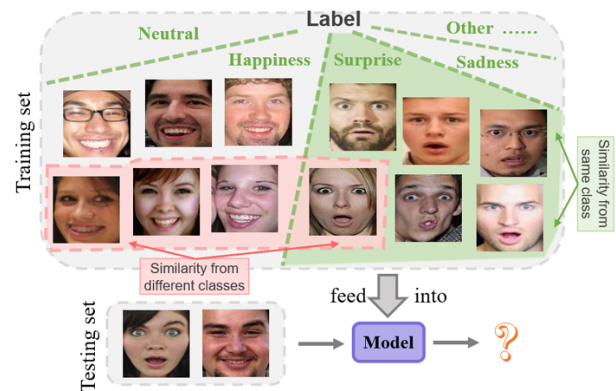


Figure 1: We often assume labeled features exhibit significant and dominant distributions within the dataset; however, unlabeled redundant features may impact this. Feature similarity can be leveraged to distinguish target features from redundant ones, based on whether they originate from the same label group or different label groups.

due to assumptions of labeled feature dominance. Dividing datasets by labels for training can result in the learning of unlabeled redundant features if they're non-randomly distributed, impacting classification accuracy. Fig. 1 illustrates one such impact. In reality, interference features are diverse, numerous, and unpredictable, exerting varying degrees of influence.

How to extract discriminative features and how to eliminate the interference of redundant features are the core issues of FER. Using self-attention, some deep learning methods (Wang et al. 2020; Xue, Wang, and Guo 2021; Xue et al. 2022; Ma, Sun, and Li 2021) focus on key features, enhancing model discrimination. However, model performance hinges on the high-quality of training data. Some methods (Xie, Hu, and Wu 2019; Zhang et al. 2021; Cai et al. 2021; Ruan et al. 2021, 2022) study separating redundant from expression features using GANs or transformers. They extract expression information via decomposition and reconstruction, but often need redundant feature labels, such as age, gender, race, or rely on identity recognition modules. Additionally, interference decoupling FER (Li, Deng, and Du 2017) relies on one or several interference feature

*Corresponding Author. Email: lvxb@scu.edu.cn
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

labels and can only address limited and specific interference factors. However, interference features are often unclear in type and quantity, and these methods that rely on labeled data often fail to achieve good performance when the types, quantities, and quality of labeled data are limited. Furthermore, when the expression feature labels are ambiguous (She et al. 2021; Zhang, Wang, and Deng 2021), it is unrealistic to clearly label all interference features.

Inspired by self-supervised contrastive learning methods (Yang et al. 2022; Gao and Patras 2024) that extract feature representations from unlabeled data, we attempt to conduct similarity comparisons between image pairs of different classes in order to isolate these redundant and interfering features, even though they are unlabeled. Additionally, we harness the similarity of same-class image pairs to extract more discriminative common features. Our method differs from those (Gao et al. 2020; Zhuang, Wang, and Qiao 2020) that utilize channel similarity and only perform comparisons between two images. It is also distinct from approaches (Yang et al. 2022; Vemulapalli and Agarwala 2019) that solely contrast high-level features among triplets.

Firstly, we propose Cross Similarity Attention (CSA), specifically tailored for mining fine-grained feature similarity between different images. We have identified a limitation when the Scaled Dot-Product Attention (SDPA) from ViT (Dosovitskiy et al. 2021) is directly applied to calculate image similarity: the similarity weights derived from the correlation between Q (query) and K (key) directly act on the feature vectors of the paired image, rather than the output feature vectors of the branch, resulting in attention that cannot obtain direct feedback (See Fig. 3b for details). Our CSA focuses solely on the importance of fine-grained similarity for the K (key), and directs attention directly between the input and output feature vectors of one branch. To our knowledge, we are the first to implement global spatial attention in the spatial dimension, analogous to the method of SENet (Hu, Shen, and Sun 2018) in the channel dimension. Previous works, such as BAM (Park et al. 2018) and CBAM (Woo et al. 2018), were limited to capturing only local spatial attention. The self-attention in ViT does not output a complete attention map after applying the softmax function; instead, it outputs results for each individual feature point.

By leveraging CSA for the interaction module, we simultaneously minimize intra-class feature differences and maximize inter-class differences, mimicking the effect of Triplet loss (Schroff, Kalenichenko, and Philbin 2015) at the fine-grained feature level. In this way, our method enhances discriminative features and separates redundant ones. To effectively integrate the refined information from cross modules, we introduce residual connections between the cross classifier and the base classifier. To avoid gradient conflicts arising from multiple interaction modules across the anchor, positive, and negative branches, we introduce an additional image of the same category as the negative to construct a four-branch centrally symmetric closed-loop network. This framework reduces the computation of two interaction modules and enables a stable joint training. During inference, the multi-branch interaction modules are removed, leaving only a single base classifier, without increasing the inference

time.

In summary, our main contributions are summarized as follows:

- We introduce Cross Similarity Attention (CSA), a global spatial attention mechanism specifically designed to mine fine-grained feature similarities across different images. It addresses the issue of indirect attention between input and output vector of one branch when computing similarity within the cross attention of ViT.
- We propose Quadruplet Cross Similarity (QCS), a novelty four-branch, centrally symmetric closed-loop joint training framework based on CSA. It refines features by adaptively mining discriminative features within the same class while simultaneously separating interfering features across different classes.
- Compared to methods that introduce additional landmark information, our method achieves state-of-the-art performance on several FER datasets by mining richer intrinsic information.

Related Work

Facial Expression Recognition

With the development of deep learning, FER has become increasingly popular in recent years, as more and more researchers focus on human-computer interaction. RAN (Wang et al. 2020) proposes a Region Attention Network that adaptively captures the importance of different region features. SCN (Ma, Sun, and Li 2020) introduces the Self-Cure Network, which suppresses uncertainty in facial expression datasets by excluding samples with ambiguous annotations. DMUE (She et al. 2021) leverages image pair similarity comparisons to mine real label confidence and potential distribution information. RUL (Zhang, Wang, and Deng 2021) computes relative uncertainty through image comparisons, addressing uncertainty arising from blurry images and inaccurate labels.

Transfer (Xue, Wang, and Guo 2021) is the first attempt to apply Vision Transformer (ViT) to adaptively characterize the relations between different facial parts. APViT (Xue et al. 2022) proposes two Attentive Pooling modules with Vision Transformer to focus on the most discriminative features and discard the less relevant ones. POSTER (Zheng, Mendieta, and Chen 2023), POSTER V2 (Mao et al. 2023), LA-Net (Wu and Cui 2023), and S2D (Chen et al. 2024) utilize ViT cross-fusion to integrate information from a facial landmark detector MobileFaceNet (Chen 2021), while FMAE (Ning, Salah, and Ertugrul 2024) is pretrained on a dataset of 9 million facial images (termed Face9M) for tasks related to AU detection and FER.

Attention Mechanisms

Self-attention primarily captures dependencies within a sequence to focus on important information. SENet (Hu, Shen, and Sun 2018) introduced a plug-and-play channel attention module. BAM (Park et al. 2018) and CBAM (Woo et al.

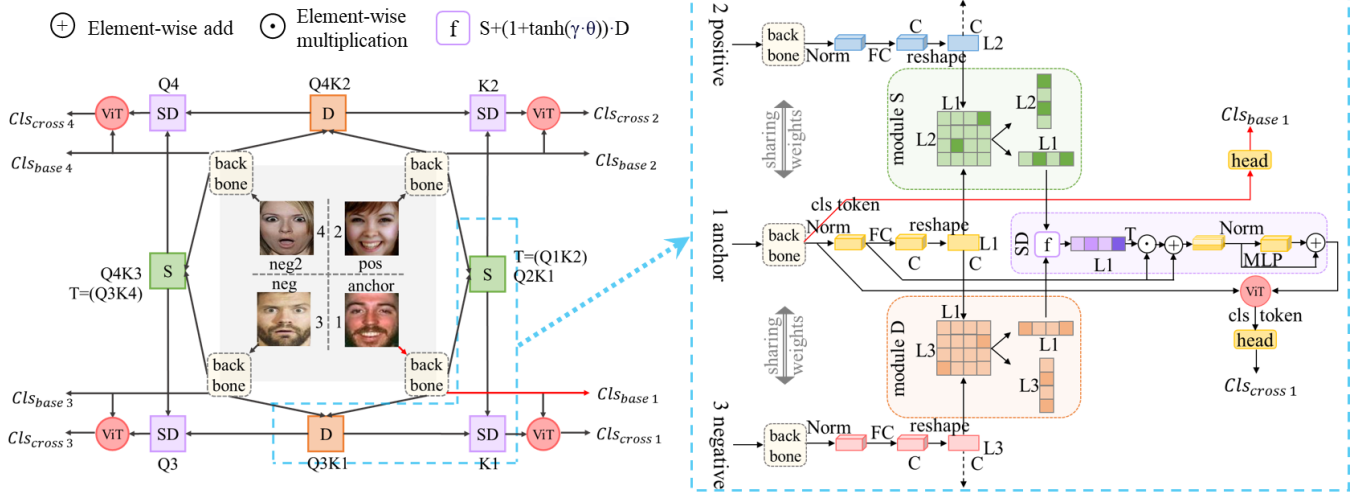


Figure 2: The framework of Quadruplet Cross Similarity Network. Joint training is performed on 4 classifiers Cl_{base} based on a weight-shared backbone and 4 classifiers Cl_{cross} enhanced by cross module, with only the red branch retained during inference. Anchor and pos are in the same category, so are neg and neg2. The matrix SD performs attention refinement on matrix S and matrix D by rows or columns.

2018) introduced spatial attention modules based on convolutions. However, these convolution-based spatial attention mechanisms can only perform local information interaction. The Non-Local Block(Wang et al. 2018) can capture long-range dependencies between any two positions in images. Subsequently, the Transformer(Vaswani et al. 2017) based on Scaled Dot-Product Attention became the state-of-the-art for natural language processing (NLP) tasks. ViT(Dosovitskiy et al. 2021) was the first to implement a pure transformer structure in the field of computer vision.

Using distinct features for computing interaction relationships constitutes cross-attention. In Transformer, K and V sequences stem from the source encoder, while Q sequence originates from the target decoder. In Few-shot classification tasks, cross-attention in CEC and CAN(Lai et al. 2023; Hou et al. 2019) can extract more discriminative features even when the sample size is limited. Fine-grained image classification methods CIN(Gao et al. 2020) and API(Zhuang, Wang, and Qiao 2020) propose channel cross-attention, which is used to extract channel similarities between image pairs. CrossViT(Chen, Fan, and Panda 2021) unifies self-attention features across different scales by fusing information from small-patch and large-patch tokens. In multimodal feature fusion(Wei et al. 2020; Nagrani et al. 2021; Zhang et al. 2023), cross-attention is used to exchange information between different modalities through early, middle, or late fusion.

Contrastive Learning

Contrastive learning is a type of Self-Supervised Learning method(Chen et al. 2020; Chen, Xie, and He 2021). Supervised Contrastive Learning(Khosla et al. 2020) introduces labeled and different data of the same class for contrastive learning, improving the discriminative ability of the model. Triplet loss(Schroff, Kalenichenko, and Philbin

2015) is a commonly used contrastive learning loss function in face recognition tasks. FEC(Vemulapalli and Agarwala 2019) learns a compact and continuous expression embedding space by comparing the similarity of three images using Triplet loss.

Methods

Cross Similarity Attention

Intuitively, comparing the similarities and differences between images can make it easier to judge the category of an image. This similarity comes not only from high-level semantic features but also from relatively low-level fine-grained features. Based on the non-local block(Wang et al. 2018), we propose a spatial attention called Cross Similarity Attention (CSA) to focus on the similarity between different images. This is significantly different from the Scaled Dot-Product Attention in ViT(Dosovitskiy et al. 2021), and it does not require concat on spatial dimension of different images like the non-local block.

Firstly, In standard ViT qkv self-attention (as shown in Fig. 3a), the query Q , key K , and value V are all feature vectors derived from the same image. The attention weights acting on the queried vectors are equivalent to acting on the output vectors.

However, if the query Q comes from the features of one image p , while the key K and values V are derived from the features of a different image a , as shown in Fig. 3b, some changes occur. The output feature points in image a are different from the query feature points in image p which makes the output feature points unrelated to the calculation of the weights. The weight w applies attention to the values V and outputs to a feature point in the vector Out . The attention distribution within the vector Out is difficult to assess through the weights.

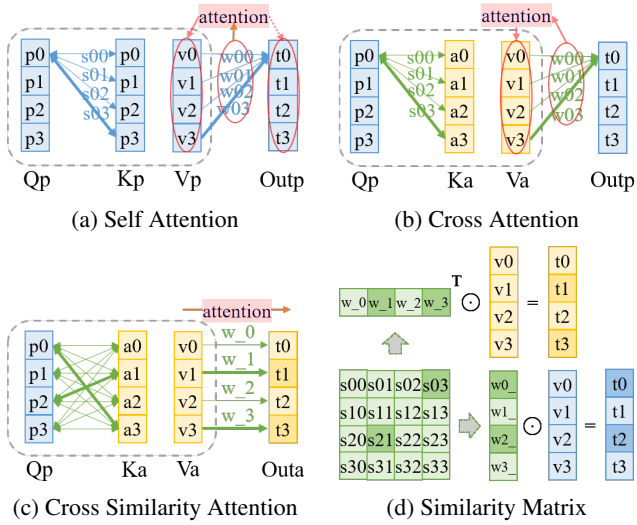


Figure 3: Different types of attention on feature vectors $h \times w \times c$, where $h=w=2$, and c denotes channels. Blue and yellow represent features from different images, and each element corresponds to a $1 \times 1 \times c$ vector. Double-headed lines denote the interaction values between corresponding two elements, and thicker lines indicate higher weights. The 4×4 green box in (d) represents the interaction matrix of (c).

Therefore we consider only quantifying the importance of the positions in the queried image (key K), regardless of where the query Q comes from. Using these weights, we directly assign weight to all the features of values V and output them to the vector Out of the same branch, as shown in Fig. 3c 3d.

Secondly, the self-attention in the standard ViT tends to capture the interactions rather than the similarity between vectors. The D_h in $W_s(Q, K) = \text{softmax}(QK^T / \sqrt{D_h})$ is a unified scaling factor. If we emphasize similarity, when calculating self-attention(SA), the vector at a certain spatial position in the same image should have a higher response with itself, as shown in Fig. 4a 4c. The yellow on the diagonal of the cosine similarity or Euclidean distance similarity matrix is more prominent. However, the dot product (Fig. 4b) calculation used in ViT weakens the response intensity on the diagonal of the interaction matrix. Meanwhile, the similarity matrix (Fig. 4d) between different images based on cross-attention(CA) lacks a highlighted diagonal, where $S_{cs}(Q, K)_{ij} \neq S_{cs}(Q, K)_{ji}$.

Meanwhile, although the self-attention mechanism of the original ViT module can capture similarities among concatenated feature vectors, it also introduces interference from the similarities among the elements within the feature representation of a single image. The concat operation enlarges the interaction matrix. In contrast, our CSA (Figure 3c) computes similarities between different feature maps more efficiently.

$$D_{cs-pa}(Q_p, K_a) = \|Q_p - K_a\| \in \mathbb{R}^{l \times l} \quad (1)$$

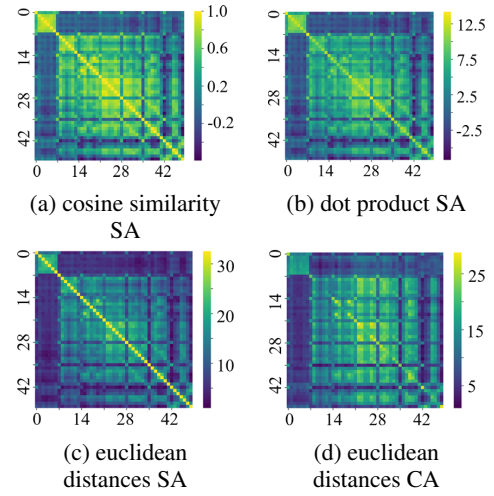


Figure 4: Interaction matrix of $7 \times 7 \times c$ ($h=7$, $w=7$, $\text{channel}=c$) image features in the spatial dimension ($h \times w=49$).

$$S_{cs-pa}(Q_p, K_a) = D_{cs-pa}(Q_p, K_a)_{\max} - D_{cs-pa}(Q_p, K_a) \in \mathbb{R}^{l \times l} \quad (2)$$

We calculate the similarity interaction matrix using Euclidean distance as shown in Eq. 1 2, $D_{cs}(Q, K)_{ij} = \sqrt{\sum_{c=1}^C (Q_{ic} - K_{jc})^2}$, $Q, K, V \in \mathbb{R}^{l \times c}$ where $l = h \times w$ and separately assess the importance of the queried positions in both images, thereby emphasizing the weights of common features with high similarity across different spatial locations within each image. Inspired by GCNet(Cao et al. 2019), which observed that attention maps calculated for different query positions are nearly identical, we aggregate the interaction matrix by rows or columns.

$$S_{cs-a}(Q_p, K_a) = \sum_{i=1}^l \left(\text{norm} \left((S_{cs-pa}(Q_p, K_a))_{i,:} \right) \right)_{ij} \in \mathbb{R}^{l \times l} \quad (3)$$

$$A_{cs-a}(Q_p, K_a) = \text{softmax}(S_{cs-a}(Q_p, K_a))^T \cdot V_a \in \mathbb{R}^{l \times 1} \quad (4)$$

Eq. 3 4 calculate the attention for image a . Here, $\text{norm} \left((S_{cs}(Q, K))_{i,:} \right)$ normalizes the similarity matrix by row using the L_2 norm to eliminate absolute differences in similarity calculated from query features at different positions, then sums the values for the same queried spatial position across columns. Finally, a mutually exclusive softmax function is employed to directly compute the weight distribution along the spatial dimension, where $\sum_{j=1}^l \left(\text{softmax}(S_{cs-a}(Q, K))^T \right)_j = 1$. We use softmax in the global spatial dimension because it operates on the similarity map of two different images, rather than on the feature map of a single image. Thus, there is no concern that applying Eq. 4 will disrupt the semantic structure of the image. Eq. 5 6 symmetrically calculate the attention for image p .

$$S_{cs-p}(Q_p, K_a) = \sum_{j=1}^l \left(\text{norm} \left((S_{cs-pa}(Q_p, K_a))_{:,j} \right) \right)_{ij} \in \mathbb{R}^{l \times 1} \quad (5)$$

$$A_{cs-p}(Q_p, K_a) = \text{softmax}(S_{cs-p}(Q_p, K_a)) \cdot V_p \in \mathbb{R}^{l \times 1} \quad (6)$$




Figure 5: The framework of Dual Cross Similarity Network.




Figure 6: Cross similarity heatmap. Yellow represents higher similarity, and blue represents greater distance, i.e., lower similarity. The cross similarity vectors $S_{cs,p}$ (1×49) and $S_{cs,a}$ (49×1) are reshaped to 7×7 heatmaps.

Quadruplet Cross Similarity Network

By referencing the method of interactive comparison between two images (Gao et al. 2020; Zhuang, Wang, and Qiao 2020), we utilize Cross Similarity Attention to enhance fine-grained features that exhibit high similarity between two images belonging to the same category. The resulting classifier, Cls_{cross} , outperforms the backbone-based Cls_{base} in discrimination. This approach is termed the Dual Cross Similarity (DCS) Network (Fig. 5).

However, merely relying on Cross Similarity Attention for feature weighting in same-category image pairs has limitations. Highly similar features may not be discriminative, e.g., flat edge areas in most face images. Unlabeled features like gender, skin tone, etc., may exhibit similarity between images of the same or different categories. Leveraging distribution differences with labeled features via cross similarity attention can mine redundant interfering features, separating them from similar same-category features, refining cleaner features. The cross similarity heatmaps between the anchor image and images of the same or different categories are shown in Fig. 6.

To achieve the joint cross similarity matrix from both same and different categories through addition or multiplication operations, we have adopted Euclidean distance instead of cosine similarity. We directly utilize the distance matrix D of different categories for coarse alignment with the similarity matrix S of the same category in terms of their numerical values, as shown in Eq. 7. By applying L2 normalization to the interaction matrix row-wise or column-wise, as in Eq. 8, we further achieve fine alignment.

$$D_{cs,na}(Q_n, K_a) = D_{cs,na}(Q_n, K_a) - D_{cs,na}(Q_n, K_a)_{min} \in \mathbb{R}^{l \times l} \quad (7)$$

$$D_{cs,a}(Q_n, K_a) = \sum_{i=1}^l \left(norm \left(D_{cs,na}(Q_n, K_a)_{i,:} \right) \right)_{ij} \in \mathbb{R}^{1 \times l} \quad (8)$$

We found that adding matrices (Eq. 9) led to better performance than multiplying them. γ is a hyperparameter and θ is a learnable parameter. The nonlinear activation function $\tanh(\gamma \cdot \theta)$ is used to limit the value to $(-1, 1)$, and more effective gradient propagation is obtained near the initial value $\theta = 0$, enabling adaptive enhancement or suppression of matrix D on S . Finally, the vector $SD_{cs,a}$ is used to replace $S_{cs,a}$ in Eq. 4.

$$SD_{cs,a}(Q_p, K_a, Q_n) = S_{cs,a}(Q_p, K_a) + (1 + \tanh(\gamma \cdot \theta)) \cdot D_{cs,a}(Q_n, K_a) \in \mathbb{R}^{1 \times l} \quad (9)$$

The Triplet loss (Schroff, Kalenichenko, and Philbin 2015) relies on comparing distances between embedding representations. Our method utilizes CSA in conjunction with the supervised classifier Cls_{cross} to achieve two objectives at the fine-grained feature level: it brings the distances between images of the same class closer (enhancing fine-grained features with high similarity) and pushes the distances between images of different classes farther away (weakening fine-grained features with high similarity, i.e., small distances).

The backbone of the anchor and pos branches are directly connected to the Cls_{base} for supervised learning. If the classifier is not added to the neg branch, gradient conflicts will arise among the anchor, pos, and neg branches. Considering that the cross similarity modules S and D learn different content, the neg branch cannot be frozen during training. Through various experimental attempts, it has been difficult to train effective models using networks constructed with triplet.

To avoid gradient conflicts, the pos, anchor, and neg branches need to have equal status throughout the network, with each branch connected to the supervised classifier Cls_{base} . Ultimately, we ingeniously added another negative image, forming the four branches into a centrally symmetric closed-loop network called Quadruplet Cross Similarity (QCS) Network, thereby avoiding gradient conflicts and enhancing training stability, as shown in Fig. 2.

Meanwhile, the QCS offers computational efficiency. It allows us to calculate only two S matrices for the same category and two D matrices for different categories, with each interaction matrix utilized by two adjacent images simultaneously. The features obtained from the backbone without the class token go through an FC layer to produce vectors QK and V . When calculating the interaction matrices with the features of two adjacent images, the weights of QK are shared and act as either query Q or key K . Consequently, the interaction matrices exhibit symmetry, where $(Q2K1)^T = Q1K2$ and $(Q3K4)^T = Q4K3$. When SD fuses the matrices S and D by rows or columns, it reduces the need for two interactive calculations.

Contrastive Residual Distillation

The training loss curve in Fig. 7 shows a significant performance gap between Cls_{base} and the enhanced Cls_{cross} by the cross module. During inference, only one Cls_{base} (red

branch in Fig. 2) is retained to predict results, posing a challenge in effectively transferring cross module information into the inference network.

A straightforward approach is to introduce residual connections between Cls_{base} and Cls_{cross} for direct gradient feedback. We explored diverse connection methods, such as element-wise addition followed by Global Average Pooling (GAP) and Bilinear pooling (BP) to fuse and pool the feature vectors. Additionally, we considered treating Cls_{cross} as the teacher model and Cls_{base} as the student model, introducing a Kullback-Leibler (KL) divergence loss or L2 loss to construct a self-distillation learning (Zhang et al. 2019; Liu, Tang, and Wu 2020). We further integrated a ViT module into Cls_{cross} to align with Cls_{base} , minimizing gradient conflicts on feature maps. We term the transfer of information from interactive contrastive learning in the cross module to the base network as Contrastive Residual Distillation.

$$\mathcal{L}_{total} = \sum_{i=1}^C \mathcal{L}_{ce_{base,i}} + \lambda_1 \cdot \mathcal{L}_{ce_{cross,i}} + [(\lambda_2 \cdot \mathcal{L}_{kl_i}) \text{ or } (\lambda_3 \cdot \mathcal{L}_{2_i})] \quad (10)$$

The joint training loss, utilizing cross entropy for all classifiers, is defined in Eq. 10. C is set to 2 for DCS and 4 for QCS, λ_1 is set to 1, The [] indicates an optional element.

Experiments

Datasets

RAF-DB (Li, Deng, and Du 2017) is a real-world affective face database which contains 29,672 facial images labeled with basic or compound facial expressions. In the experiment, 15,331 images with 7 basic expressions (i.e. surprise, fear, disgust, happiness, sadness, anger, neutral) are chosen, of which 12271 are used for training and 3068 for testing.

FERPlus (Fer 2016) is extended from FER2013 (consists of 28,709 training images and 3,589 test images) and relabeled with 8 expression labels (neutral, happiness, surprise, sadness, anger, disgust, fear, contempt) by 10 taggers. In the experiment, we follow the code (Fer 2016) to handle majority voting, which keeps a single label for each image and ignores the samples with unknown or NF labels, then measure prediction accuracy on the test set.

AffectNet (Mollahosseini et al. 2017) is a large-scale FER dataset in the wild, which contains about 440,000 manually annotated facial images. We utilized AffectNet-7/8 with 7/8 expressions, comprising 283,901/287,651 training and 3,500/4,000 validation images. Balance-sampling (Imb 2019) is used since the training set exhibits a long-tailed distribution, while the validation set is balanced.

Implementation Details

We adopted POSTER++ removing the landmark branch as the baseline and constructed our model upon it. All face images are resized to 224×224 pixels for training and testing. We use IR50 (Deng et al. 2019), which is pretrained on the Ms-Celeb-1M dataset (Guo et al. 2016), as the backbone. We employ an Adam optimizer with Sharpness Awareness Minimization (Foret et al. 2020) to train the model (200 epochs

	IR50 ViT	SDPA CSA	\oplus GAP	BP \oplus ViT	KL	RAF-DB	FERPlus
1	✓					92.05	91.06
2	✓		✓			91.98	90.83
3	✓	✓	✓			91.95	90.93
4	✓		✓			92.14	91.02
5	✓			✓		92.24	91.25
6	✓				✓	92.57	91.22
7	✓			✓	✓	92.21	90.93

Table 1: Accuracy (%) comparison of various components on DCS. SDPA denotes Scaled Dot-Product Attention in ViT, GAP denotes global average pooling, BP denotes bilinear pooling, KL denotes KL divergence loss. \oplus denotes element-wise addition.

for RAF-DB, 150 epochs for FERPlus and 80 epochs for AffectNet). For the DCS, the batch size is set to 48, the learning rate is initialized as $9e-6$ and an exponential decay learning rate scheduling with a gamma of 0.98 is employed. For the QCS, the batch size and the initial learning rate are reduced to around half. It is trained end-to-end on a single Nvidia RTX3090 GPU via Pytorch. Each experiment is repeated five times, then the maximum accuracy is reported.

Ablation Analysis

Initially, we trained a classification network utilizing the IR50+ViT backbone as a baseline. Row 1 of Table 1 presents the accuracy of the baseline. Then, the DCS Network with shared weight branches was constructed on the backbone to evaluate the performance impact of each component in our method through ablation studies on RAF-DB and FERPlus.

Evaluation of cross attentions. Rows 3 and 4 show that our CSA outperforms SDPA in the comparison. For a fair comparison, all structures remain consistent except for varying core interaction approaches, with Multi-Head Attention and class token removed for maximum structural simplicity. Features ($7 \times 7 \times c$) extracted from 3 abstraction levels of the backbone interact individually and are concatenated on the spatial dimension. CSA avoids computing similarities between features of different levels, and SDPA adopts the same approach. The FC and MLP (Multi-Layer Perceptron) of the interaction module share weights across the interaction branches. The visualization results are shown in the Row 2 and Row 3 of Fig. 9.

Evaluation of residual connection. We observe the performance gap between Cls_{base} and Cls_{cross} through their training loss curves (Fig. 7), yet we are unable to obtain the accuracy of Cls_{cross} during the inference stage. The model accuracy in row 2 of Table 1 is lower than the baseline (row 1), indicating the gradient differences after multiple layers may have distorted the baseline model’s original feature representation. The improvement from rows 2 to 4 of Table 1 demonstrates the information transfer through residual connections. Rows 4-6 suggest superior performance from minimizing gradient conflicts in connections, that BP retains

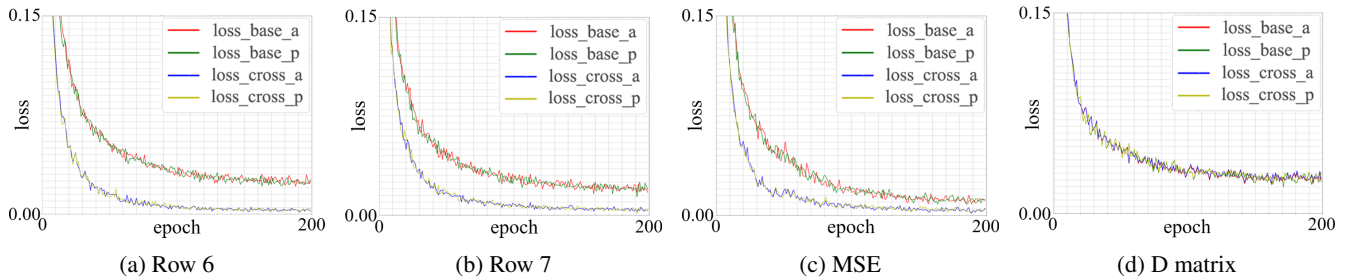


Figure 7: The training loss curves of $Cl_{s_{base}}$ and $Cl_{s_{cross}}$. (a)(b) correspond to the models in rows 6-7 of Table 1, respectively. (c)(d) correspond to the QCS, displaying only the anchor and pos branches.

	S	D	SD
RAF-DB	92.18	92.11	92.47
FERPlus	91.12	90.96	91.21

Table 2: Comparison(%) of different cross matrices.

fine-grained features while ViT module align $Cl_{s_{cross}}$ with $Cl_{s_{base}}$. The experiments below adopt the ViT connections.

Evaluation of cross similarity module overfitting. In Eq. 10, one of λ_2 and λ_3 is set to 1, while the other is set to 0. As shown in Table 1(row 7) and Fig. 7b, adding the KL divergence loss narrows the losses gap between $Cl_{s_{base}}$ and $Cl_{s_{cross}}$, but the accuracy of the model decreases. This phenomenon becomes even more evident in the more complex QCS network using MSE loss, as shown in Fig. 7c. It is possible that the cross module exhibits overfitting, influencing the base model’s generalization capacity. We set the dropout ratio ranging from 0.1 to 0.4 in the ViT modules to mitigate overfitting during training on AffectNet.

Evaluation of intra-class similarity and inter-class dissimilarity. To explore how intra-class similarity and inter-class differences influence effective feature extraction, we conducted experimental analysis leveraging the QCS Network. In Table 2, S represents retaining only the intra-class similarity matrix and D represents retaining only the inter-class distance matrix. SD utilizes Eq. 9 to fuse the interaction matrices S and D, $\gamma=1$.

Table 2 shows S contributes more to improving the model performance compared to D. As can be seen from Fig. 7d, $Cl_{s_{cross}}$, which is enhanced by D, does not show loss reduction compared to $Cl_{s_{base}}$. Fig. 6 shows that the inter-class similarity matrix has more dark blue areas (i.e., more bright yellow highlights in the inter-class distance matrix), indicating more dispersed attention. We infer that although weakening the features with small distances (high similarity) across categories helps separate redundant features, the simultaneously enhanced features with large distances do not necessarily contribute to classification. The SD matrix, resulting from the adaptive fusion of S and D, harnesses the strengths of both S and D to extract more clearly features.

Methods	RAF-DB	FERPlus	AffectNet-7	AffectNet-8
SCN(2020)	87.03	88.01	-	60.23
RUL(2021)	88.98	-	-	-
DMUE(2021)	89.42	89.51	-	63.11
TransFER(2021)	90.91	90.83	66.23	-
Face2Exp(2022)	88.54	-	64.23	-
EAC(2022)	89.99	89.64	65.32	-
APViT(2022)	91.98	90.86	66.91	-
TAN(2023)	90.87	91.00	66.45	-
POSTER(2023)	92.05	<u>91.62</u>	67.31	63.34
POSTER++(2023)	92.21	-	67.49	63.77
S2D‡(2024)	<u>92.57</u>	91.17	67.62	63.06
DCS(Ours)	<u>92.57</u>	91.41	<u>67.66</u>	64.40
QCS(Ours)	92.50	91.41	67.94	<u>64.30</u>
QCS(Ours)‡	93.02	91.85	-	-

Table 3: Performance comparison (%) on RAF-DB, FERPlus, and AffectNet. ‡means pre-trained on the AffectNet-8.

Comparison with the State of the Art

Taking into account the potential overfitting issue in the cross module, we incorporated data augmentation such as random cropping, rotation and color jitter. We also optimized the sampling strategy for positive and negative samples, resulting in further improvements in model performance. (See the Appendix in Extended version for details.)

Table 3 compares our best results with the current state-of-the-art methods on RAF-DB, FERPlus, and AffectNet-7/8. APViT and the other methods listed below are Transformer-based methods, while POSTER, POSTER++, and S2D incorporate additional landmark information. Our proposed method achieves these results solely by leveraging self-supervised contrastive learning to mine richer intrinsic information, without incorporating any extra landmark information.

While QCS effectively separates redundant features, its performance does not demonstrate a significant advantage compared to DCS. This may be due to the fact that more complex interactive modules may pose challenges during training, where over-sampling of underrepresented classes can exacerbate overfitting, ultimately resulting in a performance trade-off. Nevertheless, our method outperforms recent state-of-the-art methods on those datasets.

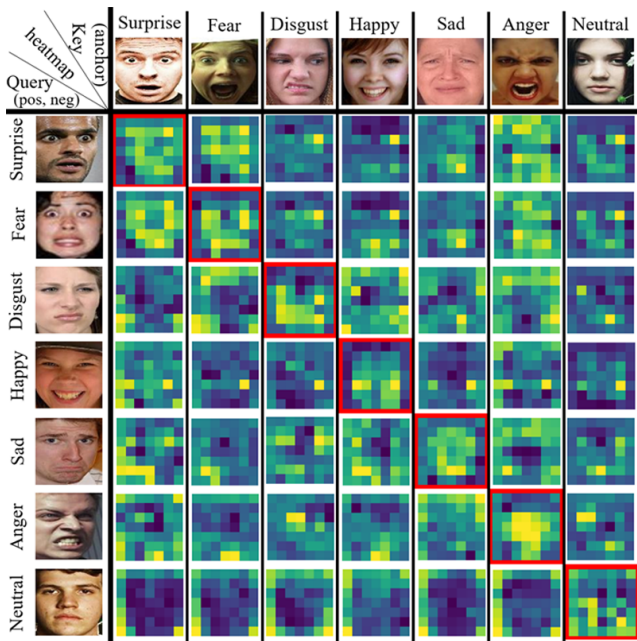


Figure 8: The confusion matrix of Cross Similarity heatmaps on RAF-DB. The anchor image located above a certain column is used as the key to calculate cross similarity on that anchor. The yellow highlighted area represents high similarity. During adaptive fusion through Eq. 9, high-similarity features on the confusion matrix diagonal are enhanced, while those outside the diagonal are suppressed.

Visualization

Cross Similarity Visualization To visually demonstrate the effectiveness of QCS in extracting distinctive features and separating redundant features, we calculated the cross similarity between every pair of the 7 expressions. We computed the similarity matrices between the anchor and positive as well as between the anchor and negative. Fig. 8 only displays the similarity heatmaps of the anchor that is used as Key, and the heatmaps for positive and negative are not shown.

The heatmaps in diagonal positions show that the distinctive features extracted from images of the same category mainly concentrate on the areas of eyes and mouth corners. The heatmaps in non-diagonal positions show the redundant features separated from comparing the images in a certain column (Key) with various expressions. As can be seen from the last row of heatmaps, Neutral, which does not have the facial features of expressions, mainly separates out feature regions from various expressions in the corners of the images.

Attention Visualization We conducted a visual comparison of the cross attentions in the Ablation Analysis. Due to the removal of the class token for cross module, we present SDPA by averaging the attention weights from all other tokens. From Row 2 of Fig. 9, it can be observed that the attention areas of SDPA are not prominent, and a significant number of redundant features in the corner regions of the

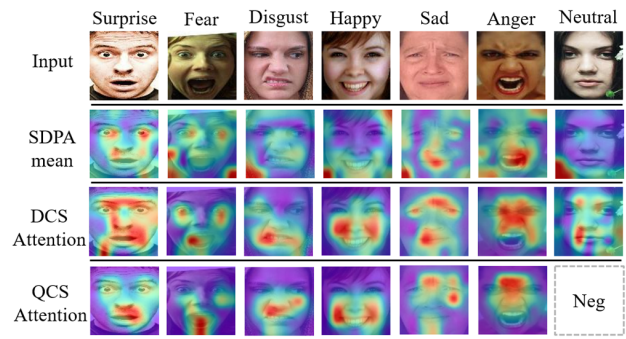


Figure 9: Cross Attention visualization on RAF-DB. The QCS uniformly employs the Neutral image as the input of negative branch.

images receive greater emphasis during information transmission. In contrast, CSA directly acts between the input and output feature vector of its own branch, enabling more direct feedback. As clearly shown in Row 3 of Fig. 9, our method achieves effective attention capture, with the effects mainly concentrated on key areas such as the facial features, cheeks, and eyebrows.

Rows 3 and 4 of Fig. 9 demonstrate the attentions of the cross module S in the DCS Network and the cross module SD in the QCS Network during training, respectively. It can be observed that the highlighted areas in QCS Attention are more concentrated, the local areas of the eyes are weakened, more attention is being paid to the movement of facial muscles. For instance, col. 4 (Happy) emphasizes the cheek area, while col. 2 (Fear), col. 3 (Disgust), and col. 5 (Sad) focus on the cheekbone area. Additionally, both col. 6 (Anger) and col. 5 (Sad) highlight the point between the eyebrows.

Conclusion

This paper proposes a feature refining network named Quadruplet Cross Similarity (QCS) to mine richer intrinsic information for Facial Expression Recognition (FER). The QCS comprises a closed-loop symmetric network for stable joint training. It adaptively extracts discriminative features within the same class and synchronously separates interfering features between different classes. The cross module is based on Cross Similarity Attention (CSA), a novel global spatial attention designed for different image pairs, which directly assigns similarity attention between the input and output feature vectors of a branch. Extensive experiments on several public datasets indicate that our QCS effectively refines the features into a purer form and achieves state-of-the-art performance. Furthermore, effectively transferring the performance of the cross module back to the basic inference network while avoiding overfitting, and enhancing the stability of joint training, remains a crucial task.

Acknowledgments

This work was supported by funding from the National Natural Science Foundation of China (U21B2035, 82372218).

References

2016. Fer+ emotion label. <https://github.com/Microsoft/FERPlus>.
2019. Imbalanced Dataset Sampler. <https://github.com/ufoyim/imbalanced-dataset-sampler>.
- Cai, J.; Meng, Z.; Khan, A. S.; Reilly, J. O.; Li, Z.; Han, S.; and Tong, Y. 2021. Identity-Free Facial Expression Recognition Using Conditional Generative Adversarial Network. In *IEEE International Conference on Image Processing (ICIP)*.
- Cao, Y.; Xu, J.; Lin, S.; Wei, F.; and Hu, H. 2019. GC-Net: Non-local Networks Meet Squeeze-Excitation Networks and Beyond. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 0–0.
- Chen, C. 2021. PyTorch Face Landmark: A Fast and Accurate Facial Landmark Detector. <https://github.com/cunjian/pytorch-face-landmark>.
- Chen, C.-F. R.; Fan, Q.; and Panda, R. 2021. CrossViT Cross-Attention Multi-Scale Vision Transformer for Image Classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 357–366.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *Proceedings of the 37th International Conference on Machine Learning (PMLR)*, 1597–1607.
- Chen, X.; Xie, S.; and He, K. 2021. An Empirical Study of Training Self-Supervised Vision Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 9640–9649.
- Chen, Y.; Li, J.; Shan, S.; Wang, M.; and Hong, R. 2024. From Static to Dynamic: Adapting Landmark-Aware Image Models for Facial Expression Recognition in Videos. *IEEE Transactions on Affective Computing*, 1–15.
- Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4690–4699.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; and Gelly, S. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *9th International Conference on Learning Representations (ICLR)*, 3–7.
- Foret, P.; Kleiner, A.; Mobahi, H.; and Neyshabur, B. 2020. Sharpness-Aware Minimization for Efficiently Improving Generalization. arXiv:2010.01412.
- Gao, Y.; Han, X.; Wang, X.; Huang, W.; and Scott, M. 2020. Channel Interaction Networks for Fine-Grained Image Categorization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 10818–10825.
- Gao, Z.; and Patras, I. 2024. Self-Supervised Facial Representation Learning with Facial Region Awareness. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2081–2092.
- Guo, Y.; Zhang, L.; Hu, Y.; He, X.; and Gao, J. 2016. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 87–102.
- Hou, R.; Chang, H.; MA, B.; Shan, S.; and Chen, X. 2019. Cross Attention Network for Few-shot Classification. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-Excitation Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7132–7141.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised Contrastive Learning. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Lai, J.; Yang, S.; Zhou, J.; Wu, W.; Chen, X.; Liu, J.; Gao, B.-B.; and Wang, C. 2023. Clustered-patch Element Connection for Few-shot Learning. arXiv:2304.10093.
- Li, S.; Deng, W.; and Du, J. 2017. Reliable crowdsourcing and deep locality-preserving learning for expression recognition in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2852–2861.
- Liu, J.; Tang, J.; and Wu, G. 2020. Residual Feature Distillation Network for Lightweight Image Super-Resolution. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 41–55.
- Ma, F.; Sun, B.; and Li, S. 2020. Suppressing Uncertainties for Large-Scale Facial Expression Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6897–6906.
- Ma, F.; Sun, B.; and Li, S. 2021. Facial expression recognition with visual transformers and attentional selective fusion. *IEEE Transactions on Affective Computing*, 14: 1236–1248.
- Ma, F.; Sun, B.; and Li, S. 2023. Transformer-augmented network with online label correction for facial expression recognition. *IEEE Transactions on Affective Computing*, 15: 593–605.
- Mao, J.; Xu, R.; Yin, X.; Chang, Y.; Nie, B.; and Huang, A. 2023. POSTER V2 A simpler and stronger facial expression recognition network. arXiv:2301.12149.
- Mollahosseini, A.; Hasani, B.; Mahoor, M. H.; and Mahoor, M. H. 2017. Affectnet: A database for facial expression and valence and arousal computing in the wild. *IEEE Transactions on Affective Computing*, 10: 18–31.
- Nagrani, A.; Yang, S.; Arnab, A.; Jansen, A.; Schmid, C.; and Sun, C. 2021. Attention Bottlenecks for Multimodal Fusion. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Ning, M.; Salah, A. A.; and Ertugrul, I. O. 2024. Representation Learning and Identity Adversarial Training for Facial Behavior Understanding. arXiv:2407.11243.
- Park, J.; Woo, S.; Lee, J.-Y.; and Kweon, I. S. 2018. BAM Bottleneck Attention Module. arXiv:1807.06514.

- 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: 455–477.
- 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 (CVPR)*, 7660–7669.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. FaceNet: A Unified Embedding for Face Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 815–823.
- 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 (CVPR)*, 6244–6253.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 6000–6010.
- Vemulapalli, R.; and Agarwala, A. 2019. A Compact Embedding for Facial Expression Similarity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5683–5692.
- Wang, C.; Chen, L.; Wang, L.; Li, Z.; and Lv, X. 2024. QCS: Feature Refining from Quadruplet Cross Similarity for Facial Expression Recognition. <https://arxiv.org/abs/2411.01988>.
- Wang, K.; Peng, X.; Yang, J.; Meng, D.; and Qiao, Y. 2020. Region attention networks for pose and occlusion robust facial expression recognition. In *IEEE Transactions on Image Processing*, 4057–4069.
- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local Neural Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7794–7803.
- Wei, X.; Zhang, T.; Li, Y.; Zhang, Y.; and Wu, F. 2020. Multi-Modality Cross Attention Network for Image and Sentence Matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10941–10950.
- Woo, S.; Park, J.; Lee, J.-Y.; and Kweon, I. S. 2018. CBAM: Convolutional Block Attention Module. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 3–19.
- Wu, Z.; and Cui, J. 2023. LA-Net: Landmark-Aware Learning for Reliable Facial Expression Recognition under Label Noise. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 20698–20707.
- Xie, S.; Hu, H.; and Wu, Y. 2019. Deep multi-path convolutional neural network joint with salient region attention for facial expression recognition. *Pattern Recognition*, 92: 177–191.
- Xue, F.; Wang, Q.; and Guo, G. 2021. Transfer: Learning relation-aware facial expression representations with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 3601–3610.
- Xue, F.; Wang, Q.; Tan, Z.; Ma, Z.; and Guo, G. 2022. Vision transformer with attentive pooling for robust facial expression recognition. *IEEE Transactions on Affective Computing*, 14: 3244–3256.
- Yang, C.; An, Z.; Cai, L.; and Xu, Y. 2022. Mutual Contrastive Learning for Visual Representation Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Zeng, D.; Lin, Z.; Yan, X.; Liu, Y.; Wang, F.; and Tang, B. 2022. Face2Exp: Combating Data Biases for Facial Expression Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 20291–20300.
- Zhang, L.; Song, J.; Gao, A.; Chen, J.; Bao, C.; and Ma, K. 2019. Be Your Own Teacher: Improve the Performance of Convolutional Neural Networks via Self Distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 3713–3722.
- Zhang, W.; Ji, X.; Chen, K.; Ding, Y.; and Fan, C. 2021. Learning a Facial Expression Embedding Disentangled from Identity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6759–6768.
- Zhang, Y.; Gong, K.; Zhang, K.; Li, H.; Qiao, Y.; Ouyang, W.; and Yue, X. 2023. Meta-Transformer A Unified Framework for Multimodal Learning. arXiv:2307.10802.
- Zhang, Y.; Wang, C.; and Deng, W. 2021. Relative Uncertainty Learning for Facial Expression Recognition. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Zhang, Y.; Wang, C.; Ling, X.; and Deng, W. 2022. Learn from all: Erasing attention consistency for noisy label facial expression recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 418–434.
- Zheng, C.; Mendieta, M.; and Chen, C. 2023. Poster: A pyramid cross-fusion transformer network for facial expression recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 3146–3155.
- Zhuang, P.; Wang, Y.; and Qiao, Y. 2020. Learning Attentive Pairwise Interaction for Fine-Grained Classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 13130–13137.