

Intensity-Aware Loss for Dynamic Facial Expression Recognition in the Wild

Hanting Li, Hongjing Niu, Zhaoqing Zhu, Feng Zhao*

University of Science and Technology of China, Hefei, China
 {ab828658, sasori, zhaopingzhu}@mail.ustc.edu.cn, fzhao956@ustc.edu.cn

Abstract

Compared with the image-based static facial expression recognition (SFER) task, the dynamic facial expression recognition (DFER) task based on video sequences is closer to the natural expression recognition scene. However, DFER is often more challenging. One of the main reasons is that video sequences often contain frames with different expression intensities, especially for the facial expressions in the real-world scenarios, while the images in SFER frequently present uniform and high expression intensities. Nevertheless, if the expressions with different intensities are treated equally, the features learned by the networks will have large intra-class and small inter-class differences, which are harmful to DFER. To tackle this problem, we propose the global convolution-attention block (GCA) to rescale the channels of the feature maps. In addition, we introduce the intensity-aware loss (IAL) in the training process to help the network distinguish the samples with relatively low expression intensities. Experiments on two in-the-wild dynamic facial expression datasets (i.e., DFEW and FERV39k) indicate that our method outperforms the state-of-the-art DFER approaches. The source code will be available at <https://github.com/muse1998/IAL-for-Facial-Expression-Recognition>.

Introduction

Facial expressions are one of the most common human behaviors and play an important role in interpersonal communications (Darwin and Prodger 1998). Facial expression recognition is the fundamental task of the intervention of mental diseases (Bisogni et al. 2022), human-computer interaction (Liu et al. 2017b), and driver safety monitoring (Wilhelm 2019). In recent years, static facial expression recognition (SFER) has become a popular topic among affective computing tasks. Many SFER methods have gradually improved the accuracy of expression recognition based on a single image (Li and Deng 2020; Li et al. 2021; She et al. 2021; Wang et al. 2020a; Xue, Wang, and Guo 2021; Zhao, Liu, and Zhou 2021; Liu et al. 2017a). However, the facial expression is dynamic, and video-based data can provide richer facial information. Therefore, dynamic facial expression recognition (DFER) has increasingly attracted




Figure 1: In-the-wild expression video sequences with different intensities. In both rows, the facial expression images in the middle have lower intensities, and the facial images on both sides have higher expression intensities.

more and more attentions. Unlike SFER, the DFER aims to classify a video into several discrete expressions (e.g., happiness, sadness, neutral, anger, surprise, disgust, and fear).

According to different data scenarios, the DFER datasets can be mainly divided into two categories: lab-controlled and in-the-wild. The lab-controlled datasets are collected in the laboratory settings, e.g., MMI (Pantic et al. 2005) and CK+ (Lucey et al. 2010), so most sequences follow a relatively fixed pattern and have similar intensities (i.e., shifting from a neutral facial expression to a peak expression). Researchers have proposed various techniques to effectively improve the performance of DFER methods in the laboratory scenarios (Yu et al. 2018; Jeong, Kim, and Dong 2020). Compared with lab-controlled DFER datasets, the in-the-wild ones are closer to the natural facial events and can provide more diverse data by collecting video sequences from the internet, such as Aff-Wild (Zafeiriou et al. 2017), DFEW (Jiang et al. 2020), and FERV39k (Wang et al. 2022). As shown in Figure 1, the video sequences in the real world with different expression intensities could result in the problem that the inter-class distance becomes smaller than the intra-class distance. However, the supervised information of discrete-label datasets does not contain intensity-related priors, e.g., DFEW and FERV39k, which leads to the fact that low-intensity expression sequences are often more likely to be misclassified. To address this issue, some datasets annotate video sequences with continuous labels along the

*Corresponding author.

arousal-valence axes, such as OMG (Barros et al. 2018) and AMIGOS (Miranda-Correa et al. 2018), and the arousal dimension can be used to measure the intensities of expressions. Nevertheless, annotating these datasets can be very expensive, and it is difficult to unify the annotation criteria of different annotators. Therefore, it is necessary to consider the impact of expression intensity on classifying low-cost datasets annotated by discrete labels.

For DFER in the wild, the early works are mainly designed based on the hand-crafted features, like LBP-TOP (Dhall et al. 2013), STLMBP (Huang et al. 2014), and HOG-TOP (Chen et al. 2014). In recent years, with the development of parallel computing hardware and collection of large-scale DFER datasets (Wang et al. 2022; Jiang et al. 2020), deep learning-based methods have gradually replaced the algorithms based on hand-crafted features and achieved state-of-the-art performance on the in-the-wild DFER datasets. For instance, vision transformer (ViT) (Dosovitskiy et al. 2020) has obtained promising results on many computer vision tasks, which inspires many researchers to build DFER models based on ViT (Zhao and Liu 2021; Ma, Sun, and Li 2022). Since transformer has strong robustness against severe occlusion and disturbance (Naseer et al. 2021), these transformer-based approaches mostly deal with various interferences in practical scenarios (e.g., variant head poses, poor illumination, and occlusions) by utilizing both spatial transformer and temporal transformer.

However, previous works on discrete-label DFER ignore the differences in the expression intensities between the samples of the same class (i.e., giving equal-status supervisory to all video sequences). To solve this problem, we design a novel global convolution-attention block (GCA) to rescale the channels of the feature maps, which helps the network heighten critical channels in low-intensity samples (i.e., pulling in the features of different-intensity sequences to reduce the intra-class differences) and suppresses less useful channels like other attention mechanisms do at the same time. Besides, we propose a novel intensity-aware loss (IAL) to help the network learn more information from hard-classified video sequences (i.e., low-intensity sequences).

In summary, this paper has the following contributions:

- By aggregating information in each channel with a global convolution, we devise a plug-and-play global convolution-attention block that can rescale the channels of the feature maps. Specially, the GCA block can not only suppress the less important channel as done in other attention mechanisms (e.g., SE and CBAM), but also enhance the target-related channel for low-intensity expressions. To the best of our knowledge, this is the first work focusing on the problem of expression intensity in the discrete-label DFER task.
- We propose a simple but effective loss function called intensity-aware loss, which can force the network to pay extra attention to the most confusing class of low-intensity expression sequences. As a result, the network can learn more precise classification boundaries.
- Extensive ablation studies and visualization results demonstrate the effectiveness of our method. It out-

performs the baseline model significantly and achieves state-of-the-art results on two popular in-the-wild DFER benchmarks.

Related Work

DFER in the Wild

Before the rise of deep learning, hand-crafted features occupied the mainstream in DFER, including LBP-TOP, STLMBP, and HOG-TOP (Dhall et al. 2013; Chen et al. 2014; Huang et al. 2014). Besides, Liu et al. utilized different Riemannian kernels to measure the similarity/distance between sequences (Liu et al. 2014). In recent years, deep learning-based approaches have gradually surpassed these techniques using hand-crafted features. So we mainly discuss methods based on deep neural networks.

Unlike SFER that classifies expressions based on a single image, the DFER needs to explore the dependencies between frames of a video sequence. Therefore, an intuitive way is to add a network to model the temporal relation between the spatial features of frames. Under this paradigm, many methods (Lee et al. 2019b; Lu et al. 2018; Ouyang et al. 2017) use CNN-based models like ResNet (He et al. 2016) and VGG (Simonyan and Zisserman 2014) to extract spatial features from each frame and then apply RNN-based models such as LSTM (Hochreiter and Schmidhuber 1997) and GRU (Chung et al. 2014) to model the temporal dependencies between frames. Specially, Wang et al. introduced a multiple attention fusion network (MAFN) to fuse multimodal data by simulating human expression recognition mechanisms (Wang, Wu, and Hoashi 2019). Zhang et al. proposed an end-to-end STRNN to jointly integrate spatial and temporal dependencies (Zhang et al. 2018). In addition to these “two-stage” frameworks, 3D-CNN-based methods adopt 3D convolution to extract spatial and temporal features jointly (Fan et al. 2016; Lee et al. 2019a; Vielzeuf, Pateux, and Jurie 2017; Hara, Kataoka, and Satoh 2018). Particularly, Jiang devised a novel EC-STFL loss that boosts the performance of multiple 3DCNN and 2DCNN-based models (Jiang et al. 2020).

In recent years, with the introduction of transformer-based models from natural language processing (Vaswani et al. 2017) to computer vision tasks (Dosovitskiy et al. 2020), some researchers began to use transformer to extract spatial and temporal features from the video sequences. Specifically, Zhao et al. devised a dynamic facial expression recognition transformer (Former-DFER) consisting of CS-Former and T-Former for learning spatial and temporal features, respectively (Zhao and Liu 2021). Ma et al. proposed a spatio-temporal transformer (STT), which can get the spatial and temporal information jointly by a transformer-based encoder (Ma, Sun, and Li 2022). Additionally, Li et al. introduced a NR-DFERNet for suppressing the impact of noisy frames in video sequences (Li et al. 2022).

Although the above solutions have designed many effective models to tackle the discrete-label DFER problem, they all ignore the effect of the expression intensities on the DFER task. Therefore, we propose IAL to force the network to pay extra attention to samples with low intensities.

Attention Mechanisms

Attention mechanisms have been widely used in various computer vision tasks in recent years. By simulating the attention mechanism in the human vision system (Tootell et al. 1998), the networks can focus on task-relevant information and suppress task-irrelevant information. Among them, the SE block (Hu, Shen, and Sun 2018) squeezes global spatial information into a channel descriptor by doing global average pooling operation and models channel-wise relationships using two fully connected (FC) layers. The ECA-Net (Wang et al. 2020b) replaces the FC layer in the SE block with a 1D convolution filter to reduce the model complexity. In addition to modeling channel-wise dependencies, CBAM (Woo et al. 2018) learns a spatial-attention mechanism. Similarly, SGE (Li, Hu, and Yang 2019) combines spatial and channel attentions. Especially, DANet (Fu et al. 2019) utilizes a two-branch structure to learn spatial attention and channel attention separately.

Most of the above attention mechanisms use global spatial pooling (e.g., global average pooling or global max pooling) to aggregate the channel information before learning the channel dependencies and then apply the sigmoid function to constrain the channel-attention weights.

However, the global pooling can lose spatial location information, which is vital for FER (e.g., the wrinkle-related features extracted from eyebrows or mouth can make a big difference to results). Besides, since the value range of the sigmoid function belongs to $(0, 1)$, the generated channel-attention weights can only be used to suppress useless channels, which does not meet our original intention of enhancing the channel features of low-intensity samples. Due to the above two reasons, we propose a global convolution-attention block to model the channel dependencies and enhance the critical channels.

Method

Overview

As shown in Figure 2, the facial expression sequence with the length of T is dynamically sampled from the raw video as an input. Then the input clips $X_{in} \in \mathbb{R}^{T \times 3 \times H_{in} \times W_{in}}$ are fed into several building blocks (He et al. 2016) to extract the frame-level features. Then the proposed GCA block is used to rescale the channels of the feature maps (i.e., suppress the less important channels and enhance the target-related channels). Then the dynamic-static fusion module (DSF) proposed in (Li et al. 2022) is applied for learning spatial features by fusing both static features of each frame and dynamic features between adjacent frames. The frame-level features that incorporate both dynamic and static features are flattened before being fed into a temporal transformer to learn the long-distance dependencies between frames. Subsequently, the mean value of the token sequence is fed into a FC layer to obtain the recognition result. Both the proposed IAL and cross-entropy loss are used to optimize the network.

Input Clips

Our method takes a clip X_{in} consisting of T RGB images as input. The input clips are dynamically sampled from the raw

video as did in previous works. Specifically, for the training clips, the raw sequences are divided into U segments equally and then randomly pick V frames from each segment. As for the test clips, we first split all frames into U segments and then select V frames in the mid of each segment. Therefore, the length of the sampled clip is $T = U \times V$ for both the training set and testing set.

Global Convolution-attention Block

SE block (Hu, Shen, and Sun 2018) was proposed to learn the channel-wise dependencies of the features. As shown in Figure 3(a), the attention weights $S_{se} \in (0, 1)$ generated by SE block are used to suppress the less important channels. For an input $X \in \mathbb{R}^{T \times C \times H \times W}$, where T is the number of frames, C is the number of the channels, while H and W are the height and width of the feature, respectively. The output \tilde{X} of SE block can be formulated as,

$$\tilde{X} = S_{se} \otimes X, \quad (1)$$

with

$$S_{se} = \sigma(W_2 \delta(W_1 Z)), \quad (2)$$

$$Z_c = \frac{1}{H \times W} \sum_{i=1}^H \sum_{j=1}^W X_c(i, j). \quad (3)$$

Where $Z \in \mathbb{R}^C$ stands for the channel descriptor generated by a global average pooling operation on input X , and Z_c and X_c are the c -th element of Z and X , respectively. δ refers to the ReLU function, and σ represents the sigmoid function. $W_1 \in \mathbb{R}^{\frac{C}{r} \times C}$ and $W_2 \in \mathbb{R}^{C \times \frac{C}{r}}$ are the weights of the first and the second FC layers, respectively. r denotes a dimensionality-reduction ratio, and \otimes is a broadcast element-wise multiplication.

However, since the value of S_{se} belongs to $(0, 1)$, the SE block can only suppress the less important channel. A suitable attention mechanism for DFER is also supposed to enhance the key channels, which is crucial for strengthening the target-related features of low-intensity samples. Besides, location information is also critical for the DFER task (e.g., wrinkle-related features appearing on the corners of the mouth or between the eyebrows can indicate a completely different expression), but the channel descriptor Z generated by global average pooling loses the location information of the input features.

Based on these viewpoints, we propose the GCA block as shown in Figure 2(b). For a same input $X \in \mathbb{R}^{T \times C \times H \times W}$, the output of GCA block can be calculated by,

$$\tilde{X} = S_{gca} \otimes X, \quad (4)$$

with

$$S_{gca} = 2\sqrt{s \otimes GAP_t(s)}, \quad (5)$$

$$s = \sigma(W_2 \delta(W_1 Z)), \quad (6)$$

and

$$Z_c = \sum_{i=1}^H \sum_{j=1}^W X_c(i, j) \times \hat{W}_c(i, j). \quad (7)$$

Where S_{gca} is the attention weights generated by GCA block, and $\hat{W}_c \in \mathbb{R}^{H \times W}$ stands for the global convolution

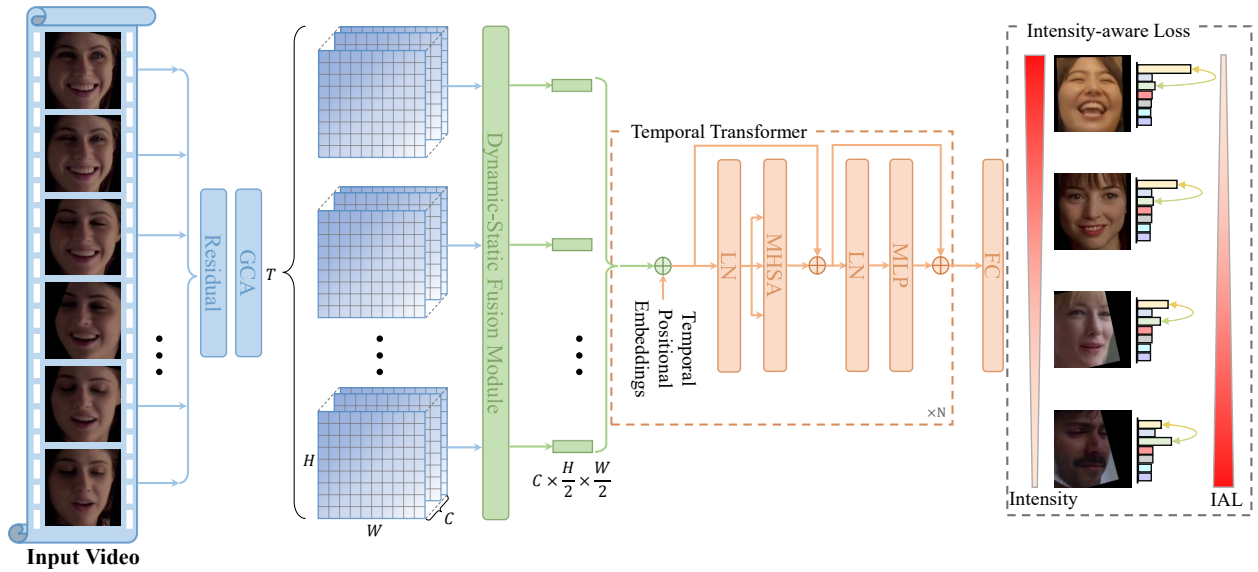


Figure 2: Overview of our method for DFER. The dynamic-static fusion module was proposed in (Li et al. 2022). Since previous methods usually use ResNet18, which contains less parameters than ResNet50, as the backbone of DFER network, we use the building block of ResNet18 to avoid over-fitting. “Residual” stands for the building block of ResNet18 (He et al. 2016). GCA and IAL represent the proposed global convolution-attention block and intensity-aware loss, respectively. LN, MHSA, and MLP separately denote the layer normalization, multi-heads self-attention, and multi-layer perceptron, while HA, SA, NE, AN, SU, DI, and FE individually indicates happiness, sadness, neutral, anger, surprise, disgust, and fear.

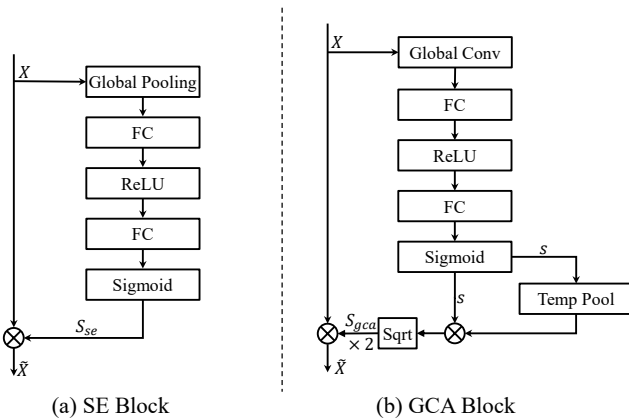


Figure 3: Diagram of the SE and proposed GCA blocks. (a) SE block and (b) our GCA block.

kernel of the c -th channel. GAP_t denotes a global average pooling on temporal dimensions.

Since S_{gca} belongs to $(0, 2)$, the GCA block can not only inhibit less important channels but also enhances critical channels. At the same time, the global convolution can retain the location information of the features, which is essential for DFER. Besides, a temporal global average pooling offers a larger inter-frame field of view.

Intensity-aware Loss

Expression intensity is an important attribute of human facial expressions. When we define $intensity \in (0, 1)$, it is ob-

vious that all non-neutral expressions tend to approach neutral expressions when the intensity converges to zero, which can be defined as,

$$\lim_{intensity \rightarrow 0} NNE = NE. \quad (8)$$

Where NNE and NE are the non-neutral and neutral expressions, respectively. It seems that taking facial expressions as a regression task is more appropriate. However, since DFER datasets with continuous labels are expensive to annotate, the scale of such datasets is often limited. Besides, the intensity criteria for annotating are difficult to unify. To tackle this, we propose intensity-aware loss to reduce the effect brought by low-intensity samples to the DFER task.

Based on the assumption that a low-intensity sample is likely to be confused with low-intensity samples from other classes (as shown in Figure 1), the network is supposed to pay extra attention to the most confusing category of each sample. Therefore, the proposed intensity-aware loss can be formulated as,

$$\mathcal{L}_{IA} = -\log(P_{IA}), \quad (9)$$

with

$$P_{IA} = \frac{e^{x_t}}{e^{x_t} + e^{x_{max}}}. \quad (10)$$

Where x_t denotes the logits of the target class and x_{max} stands for the largest logits excluding the target class. In summary, the proposed IAL gives extra attention to the class that is most likely to cause confusion as shown in Figure 2. Finally, with the commonly used cross-entropy loss, the overall loss function is defined as follows,

$$\mathcal{L} = \mathcal{L}_{CE} + \lambda \mathcal{L}_{IA}, \quad (11)$$

where λ is a hyper-parameter controlling loss coefficients. It is worth noting that when \mathcal{L}_{CE} converges to zero, \mathcal{L}_{IA} also converges to zero.

Experiments

We carry out extensive experiments on two popular in-the-wild DFER datasets (i.e., DFEW (Jiang et al. 2020) and FERV39k (Wang et al. 2022)). In this section, we first introduce the datasets and implementation details. Then we explore the effectiveness of GCA block and intensity-aware loss. Subsequently, we compare the proposed method with several SOTAs and give some visualizations to demonstrate the effectiveness of our method.

Datasets

DFEW (Jiang et al. 2020) consists of 16,372 video clips from more than 1,500 movies. Each video clip is individually annotated by ten well-trained annotators under professional guidance and assigned to one of seven basic expressions (i.e., happiness, sadness, neutral, anger, surprise, disgust, and fear). Consistent with the previous works, we only conduct experiments on 12,059 video clips, which can be clearly assigned to a specific single-labeled emotion category. All the samples have been split into five same-size parts (fd1~fd5) without overlap. The division of 5 folds is originally provided when obtaining the DFEW dataset. We choose 5-fold cross-validation as the evaluation protocol.

FERV39k (Wang et al. 2022) is current the largest in-the-wild DFER dataset and contains 38,935 video clips collected from 4 scenarios, which can be subdivided into 22 fine-grained scenes (e.g., daily life, talk show, business, and crime). Each video clip is independently annotated by 30 professional annotators to ensure high-quality labels and assigned to one of the seven primary expressions as with DFEW. In FERV39k Benchmark, video clips of all scenes are randomly shuffled and split into training (80%/31,088 clips) and testing (20%/7,847 clips) sets without overlapping. Therefore for a fair comparison, we directly use the training set and the testing set divided by FERV39k.

Implementation Details

Training Setting: In our experiments, all facial images are resized to the size of 112×112 . Random cropping, horizontal flipping, rotation and color jittering are employed to avoid over-fitting. We use SGD (Robbins and Monro 1951) to optimize our model with a batch size of 40. For the DFEW dataset, the learning rate is initialized to 0.001 and decreased at an exponential rate in 80 epochs for intensity-aware and cross-entropy loss function. For the FERV39k dataset, the learning rate is also initialized to 0.001, decreased at an exponential rate in 100 epochs using the same loss function. For both datasets, models are trained from scratch. As for sampling, the length of the dynamically sampled sequence is 16 ($U=8, V=2$ for FERV39K and $U=16, V=1$ for DFEW). The number of the self-attention heads and the temporal transformer encoders are set at 4 and 2. By default, the dimensionality-reduction ratio r is set at 16, and the loss coefficients λ is set at 0.1. All the experiments are conducted on a single NVIDIA RTX 3090 card with PyTorch toolbox.

Setting	Method		DFEW (%)		FERV39k (%)	
	GCA	IAL	UAR	WAR	UAR	WAR
a			53.56	67.42	32.79	45.24
b	✓		54.12	68.03	34.67	46.90
c		✓	54.58	68.57	34.96	47.66
d	✓	✓	55.71	69.24	35.82	48.54

Table 1: Evaluation of the proposed GCA block and intensity-aware loss. The best results are in bold.

Method	DFEW (%)		FERV39k (%)	
	UAR	WAR	UAR	WAR
SE (w/o aux)	53.10	66.98	32.85	45.56
CBAM (w/o aux)	53.31	67.14	33.41	45.76
GCA (w/o aux)	53.97	67.74	34.11	46.45
SE	53.47	67.32	33.16	45.90
CBAM	53.78	67.64	33.89	46.15
GCA	54.12	68.03	34.67	46.90

Table 2: Comparison of the proposed GCA block, SE block (Hu, Shen, and Sun 2018), and channel-attention module of CBAM block (Woo et al. 2018). The auxiliary classifier is used by default. “w/o aux” means no auxiliary classifier is used to optimize the network. The best results are in bold.

Validation Metrics: Consistent with previous methods, we take the unweighted average recall (UAR, i.e., the accuracy per class divided by the number of classes without considering the number of instances per class) and weighted average recall (WAR, i.e., accuracy) as the metrics.

Ablation Studies

We conduct ablation studies on DFEW and FERV39k to demonstrate the effectiveness of each component in our method (i.e., GCA and IAL). We take the model consisting of building blocks of ResNet18 (He et al. 2016), dynamic-static fusion module of NR-DFERNet (Li et al. 2022) and vanilla transformer (Dosovitskiy et al. 2020; Vaswani et al. 2017) as a baseline in our experiments (i.e., setting (a) in Table 1). By default we add auxiliary classifiers connected to the intermediate attention blocks (e.g., SE, CBAM, and GCA) as did in (Szegedy et al. 2015) to help the attention mechanism generate attention weights according to the target then pull in the features of different-intensity samples.

Evaluation of GCA Block and Intensity-Aware Loss: We first study the effectiveness of GCA and IAL. As shown in Table 1, replacing the building block used in ResNet18 with the proposed GCA block (setting (b)) and introducing IAL to force the network to pay extra attention to the easily-confused class (setting (c)) can both boost the overall performance on two benchmarks. Specifically, when we deploy both GCA and IAL, our method setting (d) exceeds the baseline setting (a) by 2.15%/1.82% of UAR/WAR on DFEW and 3.03%/3.30% of UAR/WAR on FERV39k, which fully

Setting		DFEW (%)		FERV39k (%)	
λ	r	UAR	WAR	UAR	WAR
0.1	8	54.58	68.64	35.08	47.60
0.15	8	54.43	68.57	35.27	47.79
0.3	8	54.16	68.19	35.02	47.47
0.1	16	55.71	69.24	35.56	48.42
0.15	16	55.24	69.02	35.82	48.54
0.3	16	54.82	68.76	35.21	48.11

Table 3: Evaluation of different dimensionality-reduction ratios r and loss coefficients λ . The best results are in bold.

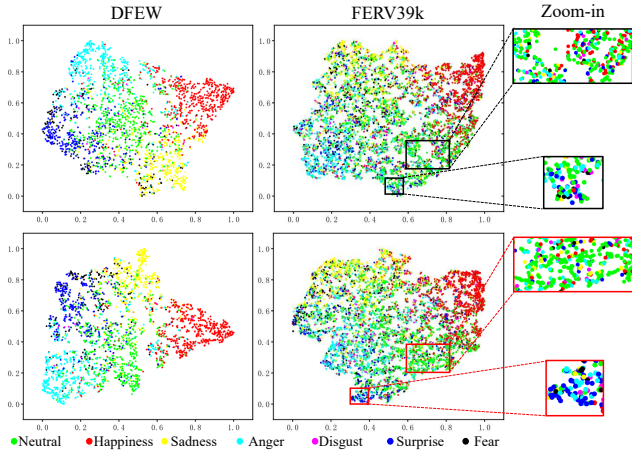


Figure 4: Visualization of the feature distribution (t-SNE (Van der Maaten and Hinton 2008)) learned by the baseline (top) and our method (bottom) on two datasets.

indicates effectiveness of the proposed modules.

Comparison with Other Channel-Attention Modules:

We compare our GCA block with several classic channel-attention modules in Table 2. In addition, we also explore the effectiveness of the auxiliary classifiers (Szegedy et al. 2015). It can be seen that the performance of all attention mechanisms is improved after deploying auxiliary classifiers to generate attention weights based on the target label. Besides, among different channel-attention blocks, our GCA block exceeds both SE block and CBAM block on two benchmarks. Specifically, GCA outperforms SE by 0.65%/0.71% of UAR/WAR on DFEW and 1.51%/1.00% of UAR/WAR on FERV39k. The main advantage of GCA is that the location information of features is preserved by replacing global pooling with global convolution.

Evaluation of Hyper-Parameters: We conduct ablation studies on several key hyper-parameters (i.e., the loss coefficients λ of IAL and the dimensionality-reduction ratio r of GCA block) of our method to explore their impact on performance. As shown in Table 3, smaller loss coefficients tend to have a better performance, whose reason is that a large loss coefficient will make the network pay too much attention to the most confusing category and ignore the distinction between other categories. As for the dimensionality-reduction

ratio, larger r , which corresponds to smaller parameter sizes, tends to have a better performance on two datasets. In summary, all the results in Table 3 are significantly improved compared to the baseline in Table 1, which also shows the robustness of our method.

Comparison with State-of-the-Arts

In this section, we compare our best results with current state-of-the-art methods on the DFEW and FERV39k benchmarks to demonstrate the superiority of our methods.

Results on DFEW: Consistent with the previous works, the experiments are conducted under 5-fold cross-validation. The experimental results are shown in Table 4. As can be seen, our methods achieve the best results both in WAR and UAR. Specifically, the proposed method outperforms NR-DFERNet (Li et al. 2022) by 1.50%/1.05% of UAR/WAR. We also show the performance on each expression in Table 4. It is obvious that “fear” and “disgust” get poorer results than other classes, which are mainly on account of the insufficient training samples (the proportion of “disgust” and “fear” is 1.22% and 8.14%, respectively).

Results on FERV39k: As the current largest in-the-wild DFER dataset, FERV39k benchmark contains diverse data sources (22 fine-grained scenes), which poses challenges for DFER tasks. As shown in Table 5, our method significantly outperforms the compared methods in WAR and achieves comparable results with Former-DFER (Zhao and Liu 2021) in UAR. Specifically, the proposed method exceeds Former-DFER by 1.69% of WAR. It should be noticed that FERV39k also has a imbalanced data distribution. The proportions of “disgust” and “fear” sequences are 5.89% and 5.4%, which is the reason why our method achieve a relatively low performance in UAR.

Visualization

We utilize t-SNE (Van der Maaten and Hinton 2008) to analyze the feature distribution learned by the baseline and our methods in Figure 4.

For the DFEW dataset, the features of our method are more closely distributed, and the boundaries between the different classes are also more pronounced. It worth noting that the features of neutral sequences learned by the baseline are distributed at the center of the remaining categories, which is consistent with our analysis in Eq. (8) (i.e., non-neutral expressions tend to approach neutral expressions when the intensity converges to zero). Since there is currently no unified quantitative definition of expression intensity, we use Eq. (10) to measure the intensity. From Figure 4, we can see that non-neutral samples surround the neutral ones without using IAL (samples at class boundaries correspond to low-intensity samples). After using IAL, the distribution of neutral samples is no longer at the center, thereby reducing the overlap with other non-neutral expressions, indicating that low-intensity samples at the boundaries are better learned. As for the FERV39k dataset, it can be seen that our method learns more discriminative features than the baseline, especially for the “neutral” and “surprise” samples. Since the overall performance on FERV39k is still

Method	Accuracy of Each Emotion (%)							Metrics (%)		FLOPs (G)
	HA	SA	NE	AN	SU	DI	FE	UAR	WAR	
C3D [†] (Tran et al. 2015)	75.17	39.49	55.11	62.49	45.00	1.38	20.51	42.74	53.54	38.57
P3D [†] (Qiu, Yao, and Mei 2017)	74.85	43.40	54.18	60.42	50.99	0.69	23.28	43.97	54.47	N/A
R(2+1)D18 [†] (Tran et al. 2018)	79.67	39.07	57.66	50.39	48.26	3.45	21.06	42.79	53.22	42.36
3D Resnet18 [†]	73.13	48.26	50.51	64.75	50.10	0.00	26.39	44.73	54.98	8.32
I3D-RGB [†] (Carreira and Zisserman 2017)	78.61	44.19	56.69	55.87	45.88	2.07	20.51	43.40	54.27	6.99
VGG11+LSTM [†]	76.89	37.65	58.04	60.70	43.70	0.00	19.73	42.39	53.70	31.65
ResNet18+LSTM [†]	78.00	40.65	53.77	56.83	45.00	4.14	21.62	42.86	53.08	7.78
3D R.18+Center Loss [†] (Wen et al. 2016)	78.49	44.30	54.89	58.40	52.35	0.69	25.28	44.91	55.48	8.32
EC-STFL [†] (Jiang et al. 2020)	79.18	49.05	57.85	60.98	46.15	2.76	21.51	45.35	56.51	8.32
3D Resnet18 [‡]	76.32	50.21	64.18	62.85	47.52	0.00	24.56	46.52	58.27	8.32
ResNet18+LSTM [‡]	83.56	61.56	68.27	65.29	51.26	0.00	29.34	51.32	63.85	7.78
Resnet18+GRU [‡]	82.87	63.83	65.06	68.51	52.00	0.86	30.14	51.68	64.02	7.78
Former-DFER [‡] (Zhao and Liu 2021)	84.05	62.57	67.52	70.03	56.43	3.45	31.78	53.69	65.70	9.11
STT [‡] (Ma, Sun, and Li 2022)	87.36	67.90	64.97	71.24	53.10	3.49	34.04	54.58	66.45	N/A
NR-DFERNet [‡] (Li et al. 2022)	88.47	64.84	70.03	75.09	61.60	0.00	19.43	54.21	68.19	6.33
GCA+IAL [‡] (Ours)	87.95	67.21	70.10	76.06	62.22	0.00	26.44	55.71	69.24	9.63

Table 4: Comparison with state-of-the-art methods on DFEW. [†] denotes time interpolation (Zhou, Zhao, and Pietikäinen 2011; Zhou et al. 2013). [‡] indicates dynamic sampling used in (Zhao and Liu 2021; Ma, Sun, and Li 2022; Li et al. 2022). HA, SA, NE, AN, SU, DI, FE denote happiness, sadness, neutral, anger, surprise, and fear, respectively. The bold denotes the best.

Method	Metrics (%)		FLOPs (G)
	UAR	WAR	
C3D (Tran et al. 2015)	22.68	31.69	38.57
P3D (Qiu, Yao, and Mei 2017)	23.20	33.39	N/A
I3D-RGB	30.17	38.78	6.99
R(2+1)D (Tran et al. 2018)	31.55	41.28	42.36
3D ResNet18	26.67	37.57	8.32
VGG16+LSTM	30.43	41.47	N/A
ResNet18+LSTM	30.92	42.95	7.78
Two ResNet18+LSTM	31.28	43.2	15.56
Former-DFER (Zhao and Liu 2021)	37.20	46.85	9.11
NR-DFERNet (Li et al. 2022)	33.99	45.97	6.33
GCA+IAL (Ours)	35.82	48.54	9.63

Table 5: Comparison with state-of-the-art methods on FERV39k. All experiments use dynamic sampling strategy.

weak, the t-SNE result still looks a little bit messy, we zoom in on critical areas to give a more intuitive display.

Discussions

Whether from the facial images in Figure 1 or the visualization results in Figure 4, we can find the challenges that low-intensity samples pose to the discrete-label DFER, which is different from the static expression recognition task based on a single image with high intensity. Therefore, we first raise this concern on discrete-label DFER and propose the IAL to deal with these low-intensity samples, and the proposed GCA block is used to enhance their features.

However, there is still much room for improvement in accurately discovering these low-intensity samples and en-

hancing their features without additional supervision. Here, we provide several technical routes that can be tried for future work on dealing with low-intensity expression samples:

- Find the low-intensity samples based on the output of the teacher network through knowledge distillation methods.
- Build a facial expression space with the neutral expression as the origin, and measure the intensity by the offset between the input and the origin.

We hope this work can inspire researchers to propose more effective solutions to this issue.

Conclusion

In this paper, we develop a plug-and-play module called global convolution-attention block and a simple but effective intensity-aware loss for in-the-wild DFER. Our global convolution-attention block is designed to rescale the channels of feature maps, so that the target-related features of low-intensity sequences can be enhanced and the less important features are suppressed. The proposed intensity-aware loss helps the networks pay extra attention to the most likely confusing class of each sample (low-intensity samples often have a likely confusing class). The experiments and visualization results demonstrate the effectiveness and superiority of our method.

Specifically, we focus on dealing with the low-intensity samples in in-the-wild dynamic facial expression recognition task. To the best of our knowledge, this is the first work focusing on the expression intensity problem in discrete-label DFER. We hope that more researchers can note this issue and provide more interesting solutions in the future.

Acknowledgements

This work was supported by the Anhui Provincial Natural Science Foundation under Grant 2108085UD12, and the JKW Research Funds under Grant 20-163-14-LZ-001-004-01. We acknowledge the support of GPU cluster built by MCC Lab of Information Science and Technology Institution, USTC.

References

- Barros, P.; Churamani, N.; Lakomkin, E.; Siqueira, H.; Sutherland, A.; and Wermter, S. 2018. The OMG-emotion behavior dataset. In *IJCNN*, 1–7. IEEE.
- Bisogni, C.; Castiglione, A.; Hossain, S.; Narducci, F.; and Umer, S. 2022. Impact of deep learning approaches on facial expression recognition in healthcare industries. *IEEE Transactions on Industrial Informatics*, 18(8): 5619–5627.
- Carreira, J.; and Zisserman, A. 2017. Quo vadis, action recognition? A new model and the Kinetics dataset. In *CVPR*, 6299–6308.
- Chen, J.; Chen, Z.; Chi, Z.; and Fu, H. 2014. Emotion recognition in the wild with feature fusion and multiple kernel learning. In *ICMI*, 508–513.
- Chung, J.; Gulcehre, C.; Cho, K.; and Bengio, Y. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Darwin, C.; and Prodger, P. 1998. *The Expression of the Emotions in Man and Animals*. Oxford University Press, USA.
- Dhall, A.; Goecke, R.; Joshi, J.; Wagner, M.; and Gedeon, T. 2013. Emotion recognition in the wild challenge 2013. In *ICMI*, 509–516.
- Dosovitskiy, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Fan, Y.; Lu, X.; Li, D.; and Liu, Y. 2016. Video-based emotion recognition using CNN-RNN and C3D hybrid networks. In *ICMI*, 445–450.
- Fu, J.; Liu, J.; Tian, H.; Li, Y.; Bao, Y.; Fang, Z.; and Lu, H. 2019. Dual attention network for scene segmentation. In *CVPR*, 3146–3154.
- Hara, K.; Kataoka, H.; and Satoh, Y. 2018. Can spatiotemporal 3D CNNs retrace the history of 2D CNNs and ImageNet? In *CVPR*, 6546–6555.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural Computation*, 9(8): 1735–1780.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *CVPR*, 7132–7141.
- Huang, X.; He, Q.; Hong, X.; Zhao, G.; and Pietikainen, M. 2014. Improved spatiotemporal local monogenic binary pattern for emotion recognition in the wild. In *ICMI*, 514–520.
- Jeong, D.; Kim, B.-G.; and Dong, S.-Y. 2020. Deep joint spatiotemporal network (DJSTN) for efficient facial expression recognition. *Sensors*, 20(7): 1936.
- Jiang, X.; Zong, Y.; Zheng, W.; Tang, C.; Xia, W.; Lu, C.; and Liu, J. 2020. DFEW: A large-scale database for recognizing dynamic facial expressions in the wild. In *ACM MM*, 2881–2889.
- Lee, J.; Kim, S.; Kim, S.; Park, J.; and Sohn, K. 2019a. Context-aware emotion recognition networks. In *ICCV*, 10143–10152.
- Lee, M. K.; Choi, D. Y.; Kim, D. H.; and Song, B. C. 2019b. Visual scene-aware hybrid neural network architecture for video-based facial expression recognition. In *FG*, 1–8. IEEE.
- Li, H.; Sui, M.; Zhao, F.; Zha, Z.; and Wu, F. 2021. MVT: mask vision transformer for facial expression recognition in the wild. *arXiv preprint arXiv:2106.04520*.
- Li, H.; Sui, M.; Zhu, Z.; et al. 2022. NR-DFERNet: Noise-robust network for dynamic facial expression recognition. *arXiv preprint arXiv:2206.04975*.
- Li, S.; and Deng, W. 2020. Deep facial expression recognition: A survey. *IEEE Transactions on Affective Computing*, 13(3): 1195–1215.
- Li, X.; Hu, X.; and Yang, J. 2019. Spatial group-wise enhance: Improving semantic feature learning in convolutional networks. *arXiv preprint arXiv:1905.09646*.
- Liu, M.; Shan, S.; Wang, R.; and Chen, X. 2014. Learning expressionlets on spatio-temporal manifold for dynamic facial expression recognition. In *CVPR*, 1749–1756.
- Liu, X.; Vijaya Kumar, B.; You, J.; and Jia, P. 2017a. Adaptive deep metric learning for identity-aware facial expression recognition. In *CVPR Workshops*, 20–29.
- Liu, Z.; Wu, M.; Cao, W.; Chen, L.; Xu, J.; Zhang, R.; Zhou, M.; and Mao, J. 2017b. A facial expression emotion recognition based human-robot interaction system. *IEEE/CAA Journal of Automatica Sinica*, 4(4): 668–676.
- Lu, C.; Zheng, W.; Li, C.; Tang, C.; Liu, S.; Yan, S.; and Zong, Y. 2018. Multiple spatio-temporal feature learning for video-based emotion recognition in the wild. In *ICMI*, 646–652.
- Lucey, P.; Cohn, J. F.; Kanade, T.; Saragih, J.; Ambadar, Z.; and Matthews, I. 2010. The extended Cohn-Kanade dataset (CK+): A complete dataset for action unit and emotion-specified expression. In *CVPR Workshops*, 94–101. IEEE.
- Ma, F.; Sun, B.; and Li, S. 2022. Spatio-temporal transformer for dynamic facial expression recognition in the wild. *arXiv preprint arXiv:2205.04749*.
- Miranda-Correa, J. A.; Abadi, M. K.; Sebe, N.; and Patras, I. 2018. AMIGOS: A dataset for affect, personality and mood research on individuals and groups. *IEEE Transactions on Affective Computing*, 12(2): 479–493.
- Naseer, M. M.; Ranasinghe, K.; Khan, S. H.; Hayat, M.; Shahbaz Khan, F.; and Yang, M.-H. 2021. Intriguing properties of vision transformers. In *NeurIPS*, 34: 23296–23308.
- Ouyang, X.; Kawai, S.; Goh, E. G. H.; Shen, S.; Ding, W.; Ming, H.; and Huang, D.-Y. 2017. Audio-visual emotion recognition using deep transfer learning and multiple temporal models. In *ICMI*, 577–582.

- Pantic, M.; Valstar, M.; Rademaker, R.; and Maat, L. 2005. Web-based database for facial expression analysis. In *ICME*. IEEE.
- Qiu, Z.; Yao, T.; and Mei, T. 2017. Learning spatio-temporal representation with pseudo-3D residual networks. In *ICCV*, 5533–5541.
- Robbins, H.; and Monro, S. 1951. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3): 400–407.
- 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 *CVPR*, 6248–6257.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. In *CVPR*, 1–9.
- Tootell, R. B.; Hadjikhani, N.; Hall, E. K.; Marrett, S.; Vanduffel, W.; Vaughan, J. T.; and Dale, A. M. 1998. The retinotopy of visual spatial attention. *Neuron*, 21(6): 1409–1422.
- Tran, D.; Bourdev, L.; Fergus, R.; Torresani, L.; and Paluri, M. 2015. Learning spatiotemporal features with 3D convolutional networks. In *ICCV*, 4489–4497.
- Tran, D.; Wang, H.; Torresani, L.; Ray, J.; LeCun, Y.; and Paluri, M. 2018. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, 6450–6459.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(86): 2579–2605.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *NeurIPS*, 30: 6000–6010.
- Vielzeuf, V.; Pateux, S.; and Jurie, F. 2017. Temporal multi-modal fusion for video emotion classification in the wild. In *ICMI*, 569–576.
- Wang, K.; Peng, X.; Yang, J.; Lu, S.; and Qiao, Y. 2020a. Suppressing uncertainties for large-scale facial expression recognition. In *CVPR*, 6897–6906.
- Wang, Q.; Wu, B.; Zhu, P.; Li, P.; Zuo, W.; and Hu, Q. 2020b. ECA-Net: Efficient channel attention for deep convolutional neural networks. In *CVPR*, 11531–11539.
- Wang, Y.; Sun, Y.; Huang, Y.; Liu, Z.; Gao, S.; Zhang, W.; Ge, W.; and Zhang, W. 2022. FERV39k: A large-scale multi-scene dataset for facial expression recognition in videos. In *CVPR*, 20922–20931.
- Wang, Y.; Wu, J.; and Hoashi, K. 2019. Multi-attention fusion network for video-based emotion recognition. In *ICMI*, 595–601.
- Wen, Y.; Zhang, K.; Li, Z.; and Qiao, Y. 2016. A discriminative feature learning approach for deep face recognition. In *ECCV*, 499–515. Springer.
- Wilhelm, T. 2019. Towards facial expression analysis in a driver assistance system. In *FG*, 1–4. IEEE.
- Woo, S.; Park, J.; Lee, J.-Y.; and Kweon, I. S. 2018. CBAM: Convolutional block attention module. In *ECCV*, 3–19.
- Xue, F.; Wang, Q.; and Guo, G. 2021. Transfer: Learning relation-aware facial expression representations with transformers. In *ICCV*, 3601–3610.
- Yu, Z.; Liu, G.; Liu, Q.; and Deng, J. 2018. Spatio-temporal convolutional features with nested LSTM for facial expression recognition. *Neurocomputing*, 317: 50–57.
- Zafeiriou, S.; Kollias, D.; Nicolaou, M. A.; Papaioannou, A.; Zhao, G.; and Kotsia, I. 2017. Aff-Wild: Valence and arousal ‘in-the-wild’ challenge. In *CVPR Workshops*, 34–41.
- Zhang, T.; Zheng, W.; Cui, Z.; Zong, Y.; and Li, Y. 2018. Spatial-temporal recurrent neural network for emotion recognition. *IEEE Transactions on Cybernetics*, 49(3): 839–847.
- Zhao, Z.; and Liu, Q. 2021. Former-DFER: Dynamic facial expression recognition transformer. In *ACM MM*, 1553–1561.
- Zhao, Z.; Liu, Q.; and Zhou, F. 2021. Robust lightweight facial expression recognition network with label distribution training. In *AAAI*, volume 35, 3510–3519.
- Zhou, Z.; Hong, X.; Zhao, G.; and Pietikäinen, M. 2013. A compact representation of visual speech data using latent variables. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(1): 1–1.
- Zhou, Z.; Zhao, G.; and Pietikäinen, M. 2011. Towards a practical lipreading system. In *CVPR*, 137–144. IEEE.