

Hypergraph-Guided Disentangled Spectrum Transformer Networks for Near-Infrared Facial Expression Recognition

Bingjun Luo¹, Haowen Wang¹, Jinpeng Wang¹, Junjie Zhu¹, Xibin Zhao^{1*}, Yue Gao^{1*}

¹BNRist, KLISS, School of Software, Tsinghua University
luobingjun@gmail.com, {wang-hw21,wjp21,zhujj18}@mails.tsinghua.edu.cn, {zxb, gaoyue}@tsinghua.edu.cn

Abstract

With the strong robusticity on illumination variations, near-infrared (NIR) can be an effective and essential complement to visible (VIS) facial expression recognition in low lighting or complete darkness conditions. However, facial expression recognition (FER) from NIR images presents a more challenging problem than traditional FER due to the limitations imposed by the data scale and the difficulty of extracting discriminative features from incomplete visible lighting contents. In this paper, we give the first attempt at deep NIR facial expression recognition and propose a novel method called near-infrared facial expression transformer (NFER-Former). Specifically, to make full use of the abundant label information in the field of VIS, we introduce a Self-Attention Orthogonal Decomposition mechanism that disentangles the expression information and spectrum information from the input image, so that the expression features can be extracted without the interference of spectrum variation. We also propose a Hypergraph-Guided Feature Embedding method that models some key facial behaviors and learns the structure of the complex correlations between them, thereby alleviating the interference of inter-class similarity. Additionally, we construct a large NIR-VIS Facial Expression dataset that includes 360 subjects to better validate the efficiency of NFER-Former. Extensive experiments and ablation studies show that NFER-Former significantly improves the performance of NIR FER and achieves state-of-the-art results on the only two available NIR FER datasets, Oulu-CASIA and Large-HFE.

Introduction

FER is an important biometric technology in the field of affective computing, which has been widely used in many areas, including health care and social networks. The goal of FER is to encode emotional information from facial geometric and texture details. However, under low lighting or complete darkness conditions, VIS light sensors are not suitable for acquiring high-quality images, which inevitably hinders the accuracy of subsequent FER. Fortunately, the NIR spectrum has been proven to be more robust to visible light illumination variations, showing great potential to be an effective complement to VIS FER. Hence, how to perform NIR FER can be a critical problem for real-world systems.

*Corresponding author.

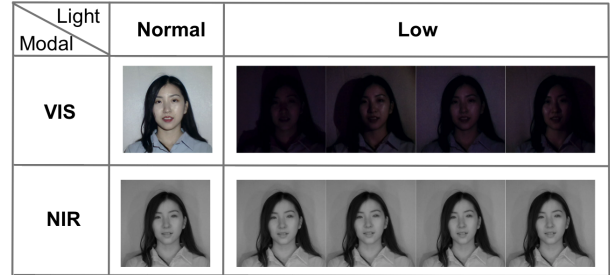


Figure 1: Comparison of VIS and NIR expression images under different lighting conditions. VIS expression images exhibit richer appearance information when captured in normal lighting conditions. However, in the case of low-light conditions, NIR expression images demonstrate superior stability and preservation of visual qualities.

As we know, having as many labeled training data as possible is extremely important for designing a deep FER system. Due to the limitation of training samples, NIR FER presents more challenging issues that cannot be solved by simply extending the successful VIS FER method to the NIR spectrum. To acquire sufficient VIS expression images, researchers have made tremendous efforts. From the earliest CK dataset, which collected 486 videos, to Multi-PIE, which collected 755,370 images, both the scale and quality are rapidly improving. In contrast, up until now, there has been only one available NIR expression dataset, Oulu-CASIA, which contains 480 valid NIR image sequences. This is far from sufficient for obtaining a well-performing NIR FER model.

Since collecting large-scale facial expression images is expensive and time-consuming, starting from scratch may not be feasible. A more ideal solution would be to leverage the rich annotation information from the VIS domain to help alleviate the problem of insufficient data in the NIR domain. To achieve this goal, there are two main challenges as follows: (1) High intra-class variation raised by different geometric and textural details between VIS and NIR images. Since VIS images (380–750 nm) and NIR images (780–1100 nm) are captured through different sensing modalities, even the same person with the same expression in these two modalities will show different patterns. VIS

images lack the characteristic spectral information found in NIR images, particularly when it comes to extreme lighting conditions. Additionally, NIR images frequently fail to capture some visual appearances or contents such as the pixels encompassing the cheek and hair. (2) High inter-class similarity. As indicated by Ekman’s Facial Action Coding System (FACS) [43], different expressions often share some key facial behaviors or action units (AU), showing the same visually discernible facial movements. Moreover, individuals in different expression classes also exhibit highly similar facial appearance characteristics.

To address these challenges, we propose a feasible method termed near-infrared facial expression transformer (NFER-Former), which can better extract modality-invariant emotional features from NIR-VIS heterogeneous samples and model high-order correlations among different expression classes for NIR FER. **Firstly**, the self-attention orthogonal decomposition (SAOD) module is proposed to decompose the expression and spectrum information from the input image. To avoid the limitation of external orthogonality constraint in previous methods, the SAOD module proposes an endogenously orthogonal decomposition strategy based on the linear algebra theorem, and then adapts the Transformer encoder layer to the task of orthogonal subspace decomposition. With the support of the Transformer network, the modality-invariant feature is extracted from heterogeneous image samples. **Secondly**, the hypergraph-guided feature embedding (HGFE) module is proposed to model the complex inter-class correlations among expression classes. HGFE module explores the AU-expression co-occurrence relationships and represents them as an AU-based knowledge hypergraph. With the embedding of hypergraph neural networks (HGNN), knowledge-guided expression features are learned to alleviate the interference of inter-class similarity. **Finally**, to further validate the performance of NFER-Former, we collect a new dataset called Large Heterogeneous Facial Expression (Large-HFE). Large-HFE dataset consists of 360 subjects ranging in age from 5 to 70, representing various cultural backgrounds.

In summary, our main contributions are three-fold:

(1) We develop a novel method, i.e. NFER-Former, to improve NIR FER performance by leveraging rich annotation information from VIS modality. To the best of our knowledge, this is the first effort to utilize NIR-VIS heterogeneous samples for the improvement of NIR FER performance.

(2) We collect a new dataset named Large-HFE for better evaluation of NIR FER. The number of subjects in Large-HFE is 4.5 times larger than the previous NIR facial expression dataset. A new benchmark on Large-HFE and Oulu-CASIA is also established to quantitatively evaluate the improvement of NIR FER performance.

(3) We conduct extensive experiments on Oulu-CASIA and Large-HFE. The results demonstrate that the proposed method significantly outperforms state-of-the-art methods.

Related Work

Facial Expression Recognition. With the advancement of deep learning, the accuracy of FER has increased rapidly. Researchers have shifted their focus from designing manual

features to improving the FER system’s robustness in real-world scenarios. Various challenges, such as occluded facial image (Xing et al. 2022), noisy labels (Wang et al. 2022), cross-dataset generalization (Li and Deng 2020), dynamic FER (Zhao and Liu 2021), have been addressed through numerous efforts. However, FER in extreme lighting conditions remains an open problem. NIR imaging has shown great potential in outperforming traditional VIS FER in uncontrolled or non-visible light conditions due to its robustness to visible light variations. In this paper, we make the first attempt to utilize NIR-VIS heterogeneous samples to alleviate the overfitting problem on small-scale data, so that better NIR FER performance can be obtained.

Orthogonal Subspace Decomposition. To utilize heterogeneous samples, many efforts have been made to reduce the domain discrepancy between NIR and VIS images. Among them, one mainstream solution is to extract modality-invariant features by mapping the input into a common subspace, which was widely used in heterogeneous face recognition (He et al. 2017, 2018; Hu and Hu 2019, 2022). Recently, deep networks are usually proposed to map both NIR and VIS images to a compact Euclidean space and further divide the feature space into two orthogonal subspaces of modality-invariant and modality-specific information respectively by adding additional orthogonality constraints. However, it is hard for such soft orthogonality constraints to enforce orthonormality in models (Li, Li, and Todorovic 2020). The extra regularization also bring challenges to model optimization (Huang, Zhou, and He 2022). In this paper, we propose a novel endogenously orthogonal decomposition strategy to disentangle the modality-invariant and modality-specific information in the feature space, which is more effective and friendly for training.

Problem Formulation

Given a NIR-VIS heterogeneous facial expression recognition dataset $\mathcal{D} = \{(\mathbf{x}^{(i)}, l^{(i)}, y^{(i)})\}_{i=1}^n$, where $\mathbf{x}^{(i)}$ is the image data, $l^{(i)} \in \{N, V\}$ is the modality label (NIR or VIS), $y^{(i)}$ is the corresponding expression label, and n is the number of training samples. From the given dataset \mathcal{D} , The goal of NIR facial expression recognition is to learn a feature extractor $f(\cdot, \Theta_f)$ that encodes modality-invariant features from both NIR and VIS facial images, and a classifier $g(\cdot, \Theta_g)$ that predicts the expression label from the encoded feature, where Θ_f and Θ_g are the parameters of the feature extractor and classifier respectively. During the test process, the recognition performance of the learned feature extractor and classifier is expected to be improved with the joint training of both NIR and VIS facial expression samples.

Our Model

To resolve the intra-class variation and inter-class similarity challenge in the NIR-VIS FER task, we propose a transformer-based network named NFER-Former, including Self-Attention Orthogonal Decomposition (SAOD) and Hypergraph-Guided Feature Embedding (HGFE). In this section, we present the overview of our model and the details of the proposed SAOD and HGFE modules.

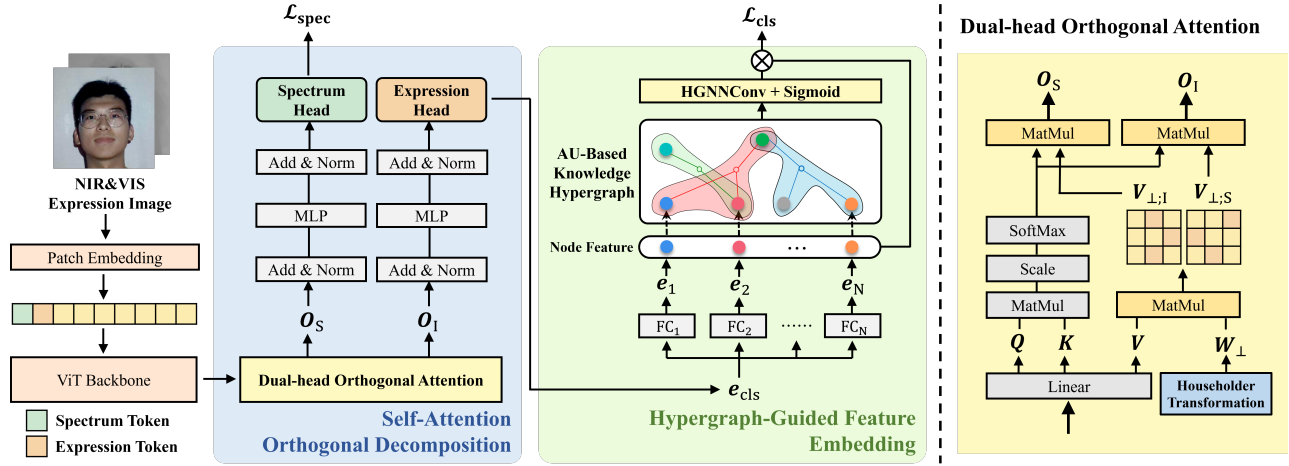


Figure 2: The main framework of the proposed NFER-Former network for near-infrared facial expression recognition. The framework consists of two modules, including a self-attention orthogonal decomposition module (the blue part) and a hypergraph-guided feature embedding module (the green part). O_S and O_I denote the modality-specific output and modality-invariant output of the dual-head orthogonal attention block defined in Eq. (3).

Model Overview

As shown in Fig. 2, the proposed NFER-Former consists of two components: Self-Attention Orthogonal Decomposition (SAOD) and Hypergraph-Guided Feature Embedding (HGFE). **Firstly**, a self-attention orthogonal decomposition mechanism is proposed to remove the spectrum information from the heterogeneous facial expression samples by decomposing the information into orthogonal subspaces. In the SAOD module, the dual-head orthogonal attention is constructed to guarantee the orthogonality of the decomposed features via linear algebra theorem. **Secondly**, a hypergraph-guided feature embedding method is proposed to alleviate the interference of inter-class similarity by modeling the complex inter-class correlations using hypergraph. In the HGFE module, the AU-based knowledge hypergraph is constructed to explore the AU-expression co-occurrence relationships and then fed into hypergraph neural networks (HGNN) to learn the knowledge-guided expression features. The whole framework is trained in an end-to-end manner.

Self-Attention Orthogonal Decomposition

In cross-modality FER task, the appearance of facial expressions may be affected by various factors, such as spectrum, identity, pose, etc., which brings large intra-class variations. Similar to the challenge in heterogeneous face recognition (Hu and Hu 2019), the spectrum information usually takes up the main part among these factors. Therefore, the key challenge is to disentangle the expression information and spectrum information from the input image, so that the expression features can be extracted without the interference of spectrum variation.

According to Related Work, a mainstream solution is to decompose the feature map of a backbone network into two orthogonal subspaces, one of which is the modality-specific subspace and the other is the modality-invariant subspace. However, the previous orthogonal decomposition methods

rely on external orthogonality constraints, which are inefficient and unfriendly to the optimization of deep neural networks. Furthermore, the previous methods are either based on or heavily relied on the property of CNN backbone. Therefore, it is still an open problem to incorporate simple yet effective orthogonal decomposition modules with a Vision Transformer network for heterogeneous FER.

To solve the above problem, we propose a self-attention orthogonal decomposition (SAOD) module based on the encoder block of Vision Transformer (Vaswani et al. 2017; Dosovitskiy et al. 2020). To conduct effective orthogonal decomposition in the self-attention mechanism, we construct an endogenously orthogonal matrix via linear algebra theorem to parameterize the projection matrix, and further develop a dual-head structure to decompose the input feature into two orthogonal subspaces without imposing any extra constraints. Besides, in the input embedding process we append an additional *[spectrum]* token, which is utilized to indicate the spectrum information of the input image and perform the spectrum classification task. The details are presented as follows.

In each transformer encoder block, the input feature map $Z \in \mathbb{R}^{N \times d}$ is linearly transformed to queries, keys, and values $Q, K, V \in \mathbb{R}^{N \times d}$. Given the matrices of query, key, and value, the attention output matrix O is then computed as $O = \text{softmax}(\frac{QK^T}{\sqrt{d}})V$. Inspired by the multi-head mechanism (Vaswani et al. 2017), a dual-head mechanism is employed in the encoder block. The value matrices V are split into two parallel heads as $V = [V_S, V_I]$ where $V_S, V_I \in \mathbb{R}^{N \times \frac{d}{2}}$ are the modality-specific value matrix and modality-invariant value matrix respectively. In this way, the input feature Z can be decomposed into two representation subspaces using V_S and V_I respectively. However, these subspaces are not orthogonal to each other, since V_S and V_I can be optimized as arbitrary matrices.

To guarantee the orthogonality of the decomposed subspaces, we further introduce an orthogonal projection $\mathbf{W}_\perp \in \mathbb{R}^{d \times d}$ on the value matrices $\mathbf{V}_S, \mathbf{V}_I$ using Householder transformation as

$$\mathbf{W}_\perp = \prod_{i=1}^N (\mathbf{I} - 2 \frac{\mathbf{v}_i \mathbf{v}_i^\top}{\|\mathbf{v}_i\|^2}) \quad (1)$$

where $\mathbf{v}_i \in \mathbb{R}^N (i = 1, \dots, d)$ are learnable parameters. Denote $\mathbf{W}_\perp = (\mathbf{w}_1, \dots, \mathbf{w}_d)^\top$. By the projection of \mathbf{W}_\perp , \mathbf{V}_S and \mathbf{V}_O are transformed as

$$\mathbf{V}_{\perp;S} = \mathbf{V}_S(\mathbf{w}_1, \dots, \mathbf{w}_{d/2})^\top, \mathbf{V}_{\perp;I} = \mathbf{V}_I(\mathbf{w}_{d/2+1}, \dots, \mathbf{w}_d)^\top \quad (2)$$

Given the transformed value matrices, we can derive the modality-specific output \mathbf{O}_S and modality-invariant output \mathbf{O}_I as

$$\mathbf{O}_i = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d}}\right) \mathbf{V}_{\perp;i}, i \in \{S, I\} \quad (3)$$

We will prove the orthogonality of these two decomposed output subspaces in the following.

Firstly, we formulate the proof goal as $\mathbf{O}_S \mathbf{O}_I^\top = \mathbf{0}$. Denote $\mathbf{O}_S = (\mathbf{o}_{S;1}, \dots, \mathbf{o}_{S;d})^\top$, $\mathbf{O}_I = (\mathbf{o}_{I;1}, \dots, \mathbf{o}_{I;d})^\top$. For any $i, j = 1, \dots, d$, row vectors $\mathbf{o}_{S;i}, \mathbf{o}_{I;j}$ are any representations in the modality-specific and modality-invariant output subspaces respectively. Therefore, the orthogonality of the two subspaces can be proved by showing that $\mathbf{o}_{S;i} \mathbf{o}_{I;j} = 0$, which is equivalent to $\mathbf{O}_S \mathbf{O}_I^\top = \mathbf{0}$.

Secondly, we start from the property of \mathbf{W}_\perp and derive the proof goal $\mathbf{O}_S \mathbf{O}_I^\top = \mathbf{0}$. It has been proved that \mathbf{W}_\perp is endogenously orthogonal and can be optimized as arbitrary orthogonal matrices (Huang, Zhou, and He 2022). Therefore, the row vectors of \mathbf{W}_\perp are orthogonal to each other, i.e., $\mathbf{w}_i^\top \mathbf{w}_j = 0 (i \neq j)$. By multiplying $\mathbf{V}_{\perp;S}$ and the transpose of $\mathbf{V}_{\perp;I}$, we have

$$\begin{aligned} \mathbf{V}_{\perp;S}^\top \mathbf{V}_{\perp;I} &= \mathbf{V}_S(\mathbf{w}_1, \dots, \mathbf{w}_{d/2})^\top (\mathbf{w}_{d/2+1}, \dots, \mathbf{w}_d) \mathbf{V}_I^\top \\ &= \mathbf{0} \end{aligned} \quad (4)$$

Hence, $\mathbf{O}_S \mathbf{O}_I^\top = \mathbf{0}$.

From the above, it is proved that the input feature map \mathbf{Z} is decomposed into two orthogonal subspaces \mathbf{O}_S and \mathbf{O}_I through the proposed SAOD module. Since the orthogonality is guaranteed endogenously without any extra constraints, the proposed module is more flexible and friendly to the optimization of deep neural networks compared with the existing regularization-based decomposition method (Huang, Zhou, and He 2022). Through the great representation power of the self-attention mechanism, orthogonal feature maps are extracted from the input feature map and then fed into the following layers for further processing.

For better disentanglement of expression and spectrum information, we further introduce the spectrum labels to supervise the model besides the original expression labels. Specifically, in the patch embedding process (Dosovitskiy et al. 2020), we append a spectrum token [*spectrum*] to the existing [*class*] token. Given the modality-specific output \mathbf{O}_S , a spectrum head implemented by linear and sigmoid layers

is proposed to transform the representation of [*spectrum*] token into the spectrum classification score \hat{l} . The spectrum classification loss is defined as the binary cross-entropy loss between the predicted spectrum score \hat{l} and the ground-truth spectrum label l as $\mathcal{L}_{\text{spectrum}}$.

Hypergraph-Guided Feature Embedding

After extracting the modality-invariant feature \mathbf{O}_I by the SAOD module, the framework still faces a great challenge of large inter-class similarity, which lies in the fact that different expressions often share some key facial behaviors (Xue, Wang, and Guo 2021). For example, the expressions *happy*, *fearful* and *surprised* usually share the same behavior of parted lips in the mouth, while *sad* and *fearful* share the same behavior of lowered brows in the eyes (Du, Tao, and Martinez 2014). The sharing of key behaviors forms complex correlations among different expressions and makes it difficult for the framework to learn discriminative features. Therefore, it is necessary to model these key facial behaviors and learn the structure of the complex correlations to alleviate the interference of inter-class similarity.

Fortunately, facial action units (AU) can be utilized to represent the key facial behaviors and further construct the complex correlation structure among different expressions. According to the facial action coding system (Ekman and Friesen 1978), AU is defined to encode the subtle facial muscle representations and movements, which are widely utilized in the FER field and believed to have strong correlations with facial expression analysis (Lien et al. 1998; Tian, Kanade, and Cohn 2001).

Therefore, we propose a hypergraph-guided feature embedding (HGFE) module to address the above challenge of inter-class similarity. The details are described as follows.

Firstly, we introduce the preliminary knowledge of hypergraph. Hypergraph is a generalization of graph where one hyperedge can connect two or more vertices. A hypergraph is defined as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of vertices and \mathcal{E} is the set of hyperedges. The hypergraph \mathcal{G} can be also represented as an incidence matrix $\mathbf{H} \in \{0, 1\}^{|\mathcal{V}| \times |\mathcal{E}|}$, with each entry $h(e, v)$ as the indicator of whether hyperedge v contains vertex e . If the vertex v is connected by the hyperedge e , then $h(v, e) = 1$, otherwise $h(v, e) = 0$.

Secondly, we formally construct the AU-based knowledge hypergraph. Hypergraph explicitly models the high-order correlations among vertices and edges rather than pairwise relationships, which is suitable for the representation of complex AU-expression relationships. Therefore, we construct an AU-based knowledge hypergraph to represent the co-occurrence relationships between expressions and prototypical AUs according to (Du, Tao, and Martinez 2014). Specifically, the AU-based knowledge hypergraph \mathcal{G}_k is constructed as $\mathcal{G}_k = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{v_1, v_2, \dots, v_N\}$ denotes the N prototypical AUs involved in the expression dataset, and $\mathcal{E} = \{e_1, e_2, \dots, e_M\}$ denotes the M expression classes ($M = 6$ for six basic expressions). The entry $h(v, e)$ of the incidence matrix \mathbf{H} is defined as

$$h(v, e) = \begin{cases} 1, & \text{if } v \text{ is consistently active in } e \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

Thirdly, we decompose the modality-invariant feature and feed it into HGNN (Feng et al. 2019) convolutional layers with the guidance of the constructed knowledge hypergraph. The $[class]$ embedding $e_{cls} \in \mathbb{R}^d$ is obtained from the modality-invariant feature O_I of the SAOD module and then mapped into N branches with fully connected layers as

$$e_i = FC_i(e_{cls}), i = 1, 2, \dots, N \quad (6)$$

where $e_i \in \mathbb{R}^{d^{(0)}}$ can be regarded as the vertex feature of v_i in the knowledge hypergraph \mathcal{G}_k . The vertex features are concatenated to form the vertex feature matrix $\mathbf{E}^{(0)} = (e_1, \dots, e_N)^T \in \mathbb{N}^{N \times d^{(0)}}$ at the input layer. Given the vertex matrix $\mathbf{E}^{(l)}$ at layer l and the hypergraph \mathcal{G}_k , the result $\mathbf{E}^{(l+1)}$ at HGNN layer l is calculated as

$$\mathbf{E}^{(l+1)} = \text{HGNNConv}(\mathbf{E}^{(l)}, \mathcal{G}_k; \Theta^{(l)}) \quad (7)$$

where $\Theta^{(l)}$ is the learnable parameter matrix at layer l .

Finally, we obtain the output of the HGNNConv layer sequences and utilize it to aggregate the multiple branches for final feature embedding. Through the HGNNConv layer sequences of length L , the dimension of the vertex feature is gradually reduced to 1. The output $\mathbf{E}^{(L)} \in \mathbb{R}^{N \times 1}$ of the last HGNNConv layer is fed into a sigmoid activation layer to produce the vertex-wise attention weight. Based on this attention weight, the decomposed branches e_1, \dots, e_N are further aggregated using weighted sum to produce the final feature embedding $\mathbf{E}_{\text{agg}} = \sum_{i=1}^N \text{Sigmoid}(\mathbf{E}_i^{(L)}) \cdot e_i$.

Based on the feature embedding \mathbf{E}_{agg} , an expression classification head implemented by a single fully connected layer is attached to obtain the final expression prediction result $\hat{y} = FC_{\text{cls}}(\mathbf{E}_{\text{agg}})$. The expression classification loss \mathcal{L}_{cls} is defined as the cross-entropy loss between the prediction result \hat{y} and the ground truth y .

Loss Function

In our method, the spectrum classification loss $\mathcal{L}_{\text{spectrum}}$ from the SAOD module and the expression classification loss \mathcal{L}_{cls} from the HGFE module are jointly optimized in an end-to-end manner. The overall loss function is defined as

$$\mathcal{L} = \mathcal{L}_{\text{cls}} + \lambda \mathcal{L}_{\text{spectrum}} \quad (8)$$

where λ is the balance parameter between the two parts.

Experiment Setup

In this section, we present the details of datasets, task settings, baselines, and implementation.

Datasets

We conduct experiments on two NIR&VIS facial expression datasets: Oulu-CASIA (Zhao et al. 2011) and the proposed Large-HFE. There are rather few publicly available datasets in the field of NIR-VIS heterogeneous FER (Yu et al. 2022). To the best of our knowledge, **Oulu-CASIA** is the only publicly available dataset before this paper, which consists of NIR&VIS video sequences captured in the lab environment with 6 basic facial expressions and 80 subjects. Among

the subjects, there are 21 females and 59 males. Following (Ruan et al. 2021), we select the last 3 frames in each sequence of strong illumination to construct the image dataset.

Large-HFE is a new NIR&VIS facial expression dataset proposed in this paper. Compared with Oulu-CASIA, the proposed dataset has a larger size in subjects, and more diversity in gender and age distribution. In Large-HFE, the images are also captured in the laboratory-controlled environment, with 6 basic expressions plus neutral and 360 subjects at the age of 5-70. Among the subjects, there are 184 females and 176 males. For each expression of each subject, 1-2 images captured in the normal illumination are selected to construct the dataset.

To enhance the training set with extra VIS samples, we select 2 VIS facial expression datasets: KDEF (Lundqvist, Flykt, and Öhman 1998) and CFEE (Du, Tao, and Martinez 2014). **KDEF** accommodates 70 subjects with 6 basic expressions plus neutral, in which images from the front view are utilized in the experiments. **CFEE** consists of 230 subjects with basic and compound expressions. For comparability with the NIR&VIS datasets, the selected datasets are both from the laboratory-controlled environment and only samples of 6 basic expressions are selected.

Baselines and Task Settings

To verify the effectiveness of the proposed method, we select baselines from two related tasks: single-modality FER methods and heterogeneous face recognition methods. For the single-modality FER methods, we select four baselines including IL-CNN (Cai et al. 2018), SCN (Wang et al. 2020), MA-Net (Zhao, Liu, and Wang 2021), and FDRL (Ruan et al. 2021). For the heterogeneous face recognition (HFR) methods, we select two subspace learning methods including DSVNs (Hu and Hu 2019) and DFD (Hu and Hu 2022), and two image synthesis methods including HiFaceGAN (Yang et al. 2020) and T2V-DDPM (Nair and Patel 2023).

In the experiments, the training set of each NIR&VIS facial expression dataset (original set) and the full set of the above 2 VIS facial expression datasets (extra set) are mixed to form the enhanced training set (enhanced set). For the train set of **NIR only**, only the NIR samples in the original set are utilized to train the single-modality FER models. For the train set of **NIR + VIS**, both NIR and VIS samples from the original set is utilized to train the single-modality FER models and subspace learning HFR models. For the train set of **NIR + VIS + Extra VIS**, both NIR and VIS samples from the enhanced set are utilized to train all the models. The test set of NIR&VIS facial expression dataset is used as the test set for all settings. The classification accuracy and macro F1 score are evaluated on the NIR part of the test set.

Implementation Details

The proposed method is implemented with *PyTorch* 1.12. The backbone is Vision Transformer (Dosovitskiy et al. 2020). Before training, the facial areas of the image samples are detected and cropped by the face detector in *dlib* library. The cropped images are then resized and randomly cropped to 224×224 and the batch size is set to 64. The

Train Set	Method	Oulu-CASIA		Large-HFE	
		Acc	F1	Acc	F1
NIR only	ResNet (He et al. 2016) †	76.39±0.55	76.32±0.71	74.23±0.67	74.07±1.12
	IL-CNN (Cai et al. 2018) †	72.71±0.21	72.36±0.52	72.08±0.44	71.94±0.69
	FDRL (Ruan et al. 2021) †	70.29±0.38	69.71±0.48	71.95±1.55	70.75±1.40
	MA-Net (Zhao, Liu, and Wang 2021) †	71.46±1.73	71.40±1.16	73.96±0.38	73.70±0.60
	SCN (Wang et al. 2020) †	78.86±0.69	78.78±0.69	79.34±0.62	79.23±0.62
NIR + VIS	ResNet (He et al. 2016) †	78.75±0.76	78.91±0.83	77.28±0.91	77.30±0.90
	IL-CNN (Cai et al. 2018) †	76.25±0.49	76.23±0.59	76.18±0.34	76.35±0.70
	FDRL (Ruan et al. 2021) †	76.67±0.31	76.63±0.49	77.15±1.71	77.49±1.59
	MA-Net (Zhao, Liu, and Wang 2021) †	68.68±1.47	68.74±1.40	70.40±0.53	70.34±0.57
	SCN (Wang et al. 2020) †	80.04±0.95	79.97±0.51	80.04±0.68	79.79±0.33
	DSVNs (Hu and Hu 2019) ‡	72.57±1.14	73.08±1.12	71.01±0.72	70.07±0.79
	DFD (Hu and Hu 2022) ‡	64.58±1.74	63.94±1.78	77.78±1.32	77.68±1.36
NFER-Former (proposed)	81.95±0.60	81.86±0.65	80.49±0.90	80.11±1.00	
NIR + VIS + VIS Extra	ResNet (He et al. 2016) †	77.64±0.78	77.83±0.90	76.43±0.91	76.41±0.91
	IL-CNN (Cai et al. 2018) †	79.17±0.68	79.15±0.75	76.59±0.63	76.74±0.72
	FDRL (Ruan et al. 2021) †	78.82±0.61	78.69±0.61	79.12±1.45	78.92±1.55
	MA-Net (Zhao, Liu, and Wang 2021) †	73.54±1.40	73.70±1.42	72.67±0.36	72.48±0.36
	SCN (Wang et al. 2020) †	80.53±0.71	80.61±0.72	79.48±0.55	79.23±0.54
	DSVNs (Hu and Hu 2019) ‡	76.39±1.07	61.56±1.06	72.81±0.52	72.32±0.51
	DFD (Hu and Hu 2022) ‡	66.67±1.43	66.21±1.48	79.83±1.04	79.82±1.13
	HiFaceGAN (Yang et al. 2020) ‡	75.83±1.21	75.76±1.19	77.31±1.20	77.36±1.27
	T2V-DDPM (Nair and Patel 2023) ‡	73.15±1.30	73.02±1.30	74.80±1.44	74.70±1.04
	NFER-Former (proposed)	84.03±0.65	83.82±0.71	82.17±1.08	81.99±1.18

Table 1: The performance on two NIR&VIS facial expression datasets in %. *NIR only* means only using the NIR part of the dataset as the train set. *NIR + VIS* means using both NIR and VIS parts of the dataset as the train set. *NIR + VIS + Extra VIS* means using the dataset and extra samples from VIS facial expression datasets as the train set. † denotes single-modality FER methods. ‡ denotes heterogeneous face recognition (HFR) methods. The best results are highlighted in bold.

training process is stopped after 40 epochs. The AdamW optimizer is adopted to optimize the backbone with the learning rate set to 10^{-4} and a one-cycle scheduler. The weight decay is set to 5×10^{-4} . The hyper-parameter λ is set to 0.1 for both Oulu-CASIA and Large-HFE. For each experiment, subject-independent five-fold cross-validation protocol is adopted and the average results are reported.

Results and Analysis

In this section, we present the experimental results and analysis. First, we compare the proposed method with the baselines on the two NIR&VIS facial expression datasets. Then, we conduct ablation study to verify the effectiveness of each component in the proposed method.

Comparison with State-of-the-art Methods

Table 1 shows the performance of the proposed method and several state-of-the-art baselines on Oulu-CASIA and Large-HFE under different settings of train set. The compared baselines are categorized into two groups, including facial expression recognition baselines and heterogeneous face recognition baselines. From the results, we can obtain the following observations.

(1) It is significant that the proposed framework achieves the best performance on both datasets under all settings. Compared with the best baseline with NIR + VIS + Extra VIS setting, the proposed method improves the accuracy by

3.50% and 2.34%, and the macro F1 score by 3.21% and 2.17% on Oulu-CASIA and Large-HFE respectively. The results demonstrate the effectiveness of the proposed method for NIR FER with the augmentation of visible facial expression datasets. This is mainly due to the proposed SAOD module and HGFE module, which can perform effective disentanglement of spectrum information using vision transformer and extract the discriminative feature embeddings with the guidance of knowledge-hypergraph. Therefore, the proposed method can make full use of the visible samples to enhance the performance of NIR FER.

(2) Among the compared baselines, FER methods tend to show higher performance than HFR methods. FER baselines including SCN, IL-CNN and FDRL achieve relatively high performance on both datasets, since they are specially designed to address the challenges in FER task, such as inter-class similarity and noisy labels. The performance of HFR methods such as DSVNs and T2V-DDPM is relatively less satisfactory in most situations since the consideration for subtle facial representations and high-order inter-class correlations is lacking for facial expression modeling.

Ablation Study

In order to verify the effectiveness of each component in the proposed method, ablation study is conducted on Oulu-CASIA and Large-HFE with different combinations of the key modules. Specifically, 4 combinations of key modules

Method	Oulu-CASIA		Large-HFE	
	Acc	F1	Acc	F1
Baseline	81.21	81.22	79.11	78.89
SAOD	82.71	82.75	80.51	80.20
HGFE	82.15	82.03	80.78	80.59
SAOD + HGFE	84.03	83.82	82.17	81.99

Table 2: The ablation study results on Oulu-CASIA and Large-HFE in %. The best results are highlighted in bold.

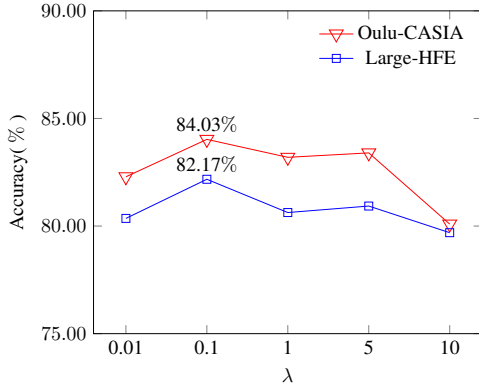


Figure 3: Ablation study results for different values of hyper-parameter λ on Oulu-CASIA and Large-HFE.

are compared in the ablation study as the followings:

- **Baseline:** The basic ViT backbone.
- **SAOD:** The basic model with SAOD module.
- **HGFE:** The basic model with HGFE module.
- **SAOD + HGFE:** The proposed NFER-Former method.

As shown in Table 2, the ViT backbone performs the worst, which is limited by the modality gap and subtle inter-class variation in NIR FER. Incorporating the SAOD module into ViT backbone significantly improves the performance, which indicates that the SAOD module effectively disentangles the spectrum information from the modality-invariant feature extraction. Moreover, the performance is further improved by incorporating the HGFE module, which demonstrates the effectiveness of the high-order inter-class correlation from knowledge hypergraph. The proposed framework that combines SAOD and HGFE modules makes full use of the above advantages and achieves the best performance.

In order to verify the recognition performance of the proposed method with different values of hyper-parameter λ , we conduct ablation study of the hyper-parameter $\lambda = 0.01, 0.1, 1, 5, 10$ on Oulu-CASIA and Large-HFE, as shown in Fig. 3. It is observed that our method achieves the best performance when $\lambda = 0.1$ for both datasets. Therefore, we set the values of λ to 0.1 in the experiments.

Visualization

In order to demonstrate the performance on the challenges of intra-class variation and inter-class similarity, we visual-

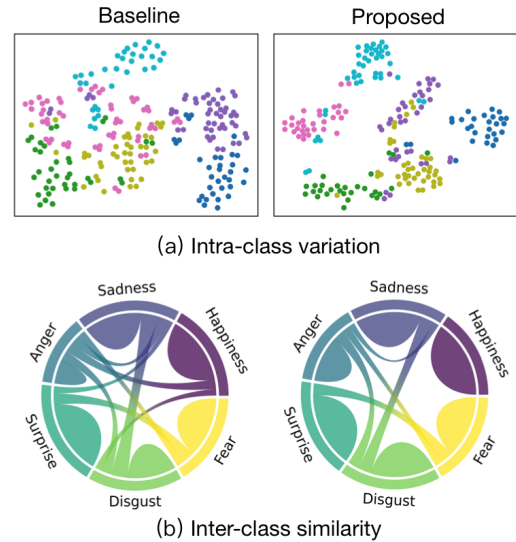


Figure 4: Visualization of intra-class variation and inter-class similarity of the expression features learned by baseline and proposed method. Different colors denote different classes of expressions.

ize the expression features extracted by the baseline and the proposed NFER-Former on Large-HFE dataset in Fig. 4.

As shown in Fig. 4.a, the intra-class variations of the extracted features are illustrated in the scatter plot, where the features are reduced to 2D space using t-SNE (Van der Maaten and Hinton 2008). From the figure, we can see that the features learned by the baseline are not compact for each expression class. In contrast, the features learned by the proposed method show more compact intra-class variation.

As shown in Fig. 4.b, the inter-class similarities of the extracted features are illustrated in the chord diagram, where each expression class is regarded as a node and the feature similarity is regarded as the edge weight between two classes. The similarity between two classes is calculated from the average distance from the features of one class to the features of another class. For clarity, the edges with similarity score less than a constant threshold are removed. It is observed that features extracted by baseline method have higher inter-class similarity, such as *happiness* and *anger*, *sadness* and *surprise*. In contrast, the proposed method effectively alleviates the inter-class similarity problem.

Conclusion

In this paper, we propose a novel NFER-Former method to utilize NIR-VIS heterogeneous facial expression samples for effective NIR FER. We introduce Self-Attention Orthogonal Decomposition and Hypergraph-Guided Feature Embedding modules to learn the modality-invariant discriminative features from NIR-VIS heterogeneous samples. We construct a new dataset named Large-HFE with 4.5 times more subjects than the previous dataset. The experimental results show that our NFER-Former method achieves leading performance on existing NIR FER datasets.

Acknowledgments

This research is sponsored in part by the NSFC Program (No. U20A6003, 62076146, 62021002, U19A2062, U1911401, 6212780016) and Industrial Technology Infrastructure Public Service Platform Project "Public Service Platform for Urban Rail Transit Equipment Signal System Testing and Safety Evaluation" (No. 2022-233-225), Ministry of Industry and Information Technology of China.

References

- Cai, J.; Meng, Z.; Khan, A. S.; Li, Z.; O'Reilly, J.; and Tong, Y. 2018. Island loss for learning discriminative features in facial expression recognition. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, 302–309. IEEE.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Du, S.; Tao, Y.; and Martinez, A. M. 2014. Compound facial expressions of emotion. *Proceedings of the national academy of sciences*, 111(15): E1454–E1462.
- Ekman, P.; and Friesen, W. V. 1978. Facial action coding system. *Environmental Psychology & Nonverbal Behavior*.
- Feng, Y.; You, H.; Zhang, Z.; Ji, R.; and Gao, Y. 2019. Hypergraph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 3558–3565.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- He, R.; Wu, X.; Sun, Z.; and Tan, T. 2017. Learning invariant deep representation for nir-vis face recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.
- He, R.; Wu, X.; Sun, Z.; and Tan, T. 2018. Wasserstein CNN: Learning invariant features for NIR-VIS face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 41(7): 1761–1773.
- Hu, W.; and Hu, H. 2019. Disentangled spectrum variations networks for NIR–VIS face recognition. *IEEE Transactions on Multimedia*, 22(5): 1234–1248.
- Hu, W.; and Hu, H. 2022. Domain-Private Factor Detachment Network for NIR-VIS Face Recognition. *IEEE Transactions on Information Forensics and Security*, 17: 1435–1449.
- Huang, H.; Zhou, X.; and He, R. 2022. Orthogonal transformer: An efficient vision transformer backbone with token orthogonalization. *Advances in Neural Information Processing Systems*, 35: 14596–14607.
- Li, J.; Li, F.; and Todorovic, S. 2020. Efficient Riemannian Optimization on the Stiefel Manifold via the Cayley Transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(12): 3500–3513.
- Li, S.; and Deng, W. 2020. A deeper look at facial expression dataset bias. *IEEE Transactions on Affective Computing*, 13(2): 881–893.
- Lien, J. J.; Kanade, T.; Cohn, J. F.; and Li, C.-C. 1998. Automated facial expression recognition based on FACS action units. In *Proceedings third IEEE international conference on automatic face and gesture recognition*, 390–395. IEEE.
- Lundqvist, D.; Flykt, A.; and Öhman, A. 1998. Karolinska directed emotional faces. *PsychTESTS Dataset*, 91: 630.
- Nair, N. G.; and Patel, V. M. 2023. T2V-DDPM: Thermal to Visible Face Translation using Denoising Diffusion Probabilistic Models. In *2023 IEEE 17th International Conference on Automatic Face and Gesture Recognition (FG)*, 1–7. IEEE.
- 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*, 7660–7669.
- Tian, Y.-I.; Kanade, T.; and Cohn, J. F. 2001. Recognizing action units for facial expression analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 23(2): 97–115.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, K.; Peng, X.; Yang, J.; Lu, S.; and Qiao, Y. 2020. Suppressing uncertainties for large-scale facial expression recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6897–6906.
- Wang, L.; Jia, G.; Jiang, N.; Wu, H.; and Yang, J. 2022. EASE: Robust Facial Expression Recognition via Emotion Ambiguity-SEnsitive Cooperative Networks. In *Proceedings of the 30th ACM International Conference on Multimedia*, 218–227.
- Xing, Z.; Tan, W.; He, R.; Lin, Y.; and Yan, B. 2022. Co-Completion for Occluded Facial Expression Recognition. In *Proceedings of the 30th ACM International Conference on Multimedia*, 130–140.
- 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*, 3601–3610.
- Yang, L.; Wang, S.; Ma, S.; Gao, W.; Liu, C.; Wang, P.; and Ren, P. 2020. Hifacegan: Face renovation via collaborative suppression and replenishment. In *Proceedings of the 28th ACM international conference on multimedia*, 1551–1560.
- Yu, S.; Han, H.; Shan, S.; and Chen, X. 2022. CMOS-GAN: Semi-Supervised Generative Adversarial Model for Cross-Modality Face Image Synthesis. *IEEE Transactions on Image Processing*, 32: 144–158.
- Zhao, G.; Huang, X.; Taini, M.; Li, S. Z.; and Pietikäinen, M. 2011. Facial expression recognition from near-infrared videos. *Image and vision computing*, 29(9): 607–619.

Zhao, Z.; and Liu, Q. 2021. Former-dfer: Dynamic facial expression recognition transformer. In *Proceedings of the 29th ACM International Conference on Multimedia*, 1553–1561.

Zhao, Z.; Liu, Q.; and Wang, S. 2021. Learning deep global multi-scale and local attention features for facial expression recognition in the wild. *IEEE Transactions on Image Processing*, 30: 6544–6556.