

# Permutation-Invariant Feature Restructuring for Correlation-Aware Image Set-based Recognition

Xiaofeng Liu<sup>1,2</sup>, Zhenhua Guo<sup>1,4</sup>, Site Li<sup>1</sup>, Lingsheng Kong<sup>3</sup>,  
 Ping Jia<sup>3</sup>, Jane You<sup>5</sup>, B.V.K. Vijaya Kumar<sup>1\*</sup>

<sup>1</sup>Carnegie Mellon University; <sup>2</sup>Harvard University; <sup>3</sup>CIOMP, Chinese Academy of Sciences;

<sup>4</sup>Graduate School at Shenzhen, Tsinghua University; <sup>5</sup>The Hong Kong Polytechnic University

\*Corresponding author: kumar@ece.cmu.edu

## Abstract

We consider the problem of comparing the similarity of image sets with variable-quantity, quality and un-ordered heterogeneous images. We use feature restructuring to exploit the correlations of both inner&inter-set images. Specifically, the residual self-attention can effectively re-structure the features using the other features within a set to emphasize the discriminative images and eliminate the redundancy. Then, a sparse/collaborative learning-based dependency-guided representation scheme reconstructs the probe features conditional to the gallery features in order to adaptively align the two sets. This enables our framework to be compatible with both verification and open-set identification. We show that the parametric self-attention network and non-parametric dictionary learning can be trained end-to-end by a unified alternative optimization scheme, and that the full framework is permutation-invariant. In the numerical experiments we conducted, our method achieves top performance on competitive image set/video-based face recognition and person re-identification benchmarks.

## 1. Introduction

Many research studies focus on using a single image [33, 35, 32] or video [27, 30] as shown in Fig. 1. However, in many practical applications, a set of images of a subject (consisting of still images, or frames from a video, or a mixture of both) can usually be collected from different checkpoints, segments in videos, mugshot history of a criminal and lifetime enrollment images for identity documents *etc* [28]. These can contain extreme rotations, complex expressions and illumination variations [25, 40, 61]. This setting is more similar to the real-world biometric scenarios [15].

Compared to the case of a single image [24, 37, 31, 38, 29, 26], richer and complementary information can be expected in a set, because the samples are captured from mul-

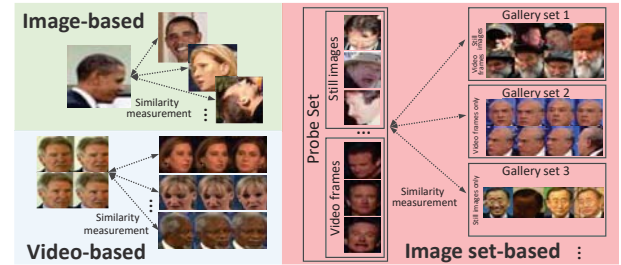


Figure 1. Illustration of the single image-based, video sequence-based and image set-based open-set face identification.

tiples views [36]. However, it also poses several challenges including a) *variable number* of samples within a set, b) larger *inner-set variability* than its video-based recognition counterpart, and c) *order-less data*.

The conventional approaches to set-based recognition aggregate all images  $X = \{x_n\}_{n=1}^N$  in a set with  $N$  images to a single feature vector. The aggregation is achieved through max/average pooling [21, 43, 5]. More formally,  $f(X) = \rho(\max/\text{average pooling}(\{\varphi(x_1), \dots, \varphi(x_N)\}))$ , where  $\rho, \varphi$  are any appropriate mapping functions. Considering that the importance of each image maybe different, their weights are calculated independently by  $\varphi$  as the soft attention score in pooling operation [61, 40]. By doing this, the high quality (*i.e.*, clear and frontal) faces are favored in these models. These methods are not sensitive to the order of images within a set (*i.e.*, permutation-invariant), but their weight selection procedures for each image in the set do not pay attention to the other images. Without inner-set correlation can result in redundant weights as in Fig. 2.

The inter-set correlation is also important for similarity measurements [62]. If we are given a query set consisting of clear profile face images, it should be more reasonable to emphasize the profile face images in the gallery set rather than the high-quality frontal face images.

How to model the inner/inter-set correlations in image

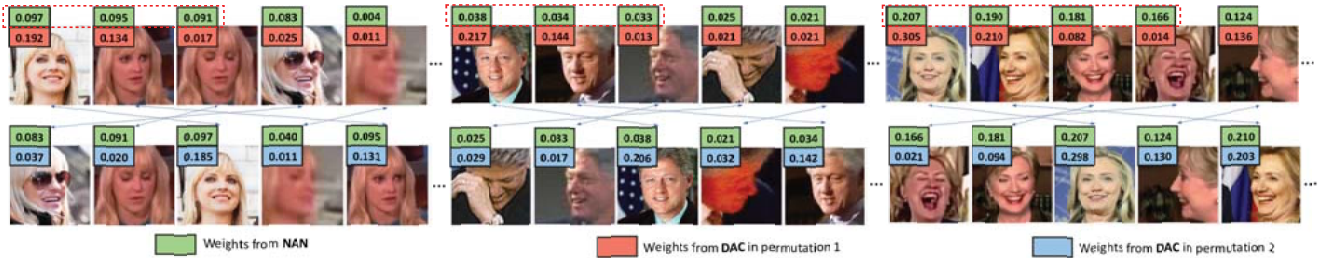


Figure 2. Examples in IJB-A dataset showing the weights of images calculated by independent image quality assessment (NAN [61]) that does not take into account inner-set correlations and RL-based inner-set relationship exploring (DAC [25]). The NAN suggests a similar score for a discriminative image and its blurry version in the same pose (dashed red box). However, there is little additional information in the latter. When we change the order of images from the 1st row to the 2nd row, the output of DAC is different. If the inferior samples are processed ahead by DAC, their weights can be larger, and the performance would be degraded.

sets with arbitrary number of order-less images has long been a challenging [25].

For exploring the inner-set correlations, the typical permutation-sensitive methods, *e.g.*, recurrent neural networks, are totally not applicable for order-less data [25]. The reinforcement learning (RL) was introduced to estimate the weights of each image sequentially [25]. Fig. 2 shows that RL is still sensitive to the order.

For exploring the inter-set correlations, the parametric pair-wise input network (takes two sets as input and output the similarity score using a sigmoid unit directly) usually computationally inefficient for open-set identification [56, 61]. [25] proposes to represent and compare frontal and profile faces separately. When we start using separate representations for expressions, lighting, occlusion, makeup *etc.*, the enormous number of final representations and hyper-parameters needed to balance them will be uncontrollable.

To address the challenge of taking inner-set and inter-set correlations into account in a computationally efficient way, we propose the *permutation-invariant feature restructuring* (PIFR) framework to reduce the redundancy and maintain the benefit of diversity information. It fully considers the inner/inter-set correlation by restructuring a sample in the probe set with the samples in probe/gallery set to explore their complementarity and substitutability.

Specifically, a parametric redundancy-eliminating self-attention (RSA) module allows all samples in a set to contribute to produce a refined feature vector based on feature pair-wise affinity. Its residual attention mechanism not only learns to emphasize discriminative samples while repelling inferior ones, but it also explicitly reduces the redundancy. Moreover, spatial clues are incorporated using the fully convolutional feature extractor and Gaussian similarity metric.

To take all possible inter-set variations into account and scale to open-set identification, we develop a non-parametric dependency-guided feature alignment (DFA) module for inter-set interaction. It assumes that each probe image feature vector can be sparsely/collaboratively restructured on the basis of gallery feature set [34]. In testing,

Method	permutation invariant	inner-set correlation	inter-set correlation	verification	open-set identification
AvePool	yes			yes	yes
NAN[61]	yes			yes	yes
DAC[25]		yes	pose-only	yes	yes
CNet[56]	yes	pose-only	pose-only	yes	
PIFR	yes	yes	yes	yes	yes

Table 1. Comparisons between proposed PIFR and other state-of-the-art approaches for set-based recognition.

the gallery features can be pre-computed and stored, and its reconstruction coefficients can be computed in parallel.

The parametric deep learning-based RSA and non-parametric dictionary learning-based DFA can be alternatively optimized in our PIFR. Table 1 compares the capabilities of the proposed PIFR to other set-based approaches. The main contributions of this paper are

- We introduce PIFR that takes into account the correlations among variable number of order-less samples by restructuring the feature with the other inner&inter-set features in a permutation-invariant manner. It emphasizes the discriminative images, reduces the redundancy, adaptively aligns the probe-gallery set, and is efficient for both verification and open-set identification tasks.
- We show how self-attention and dictionary learning can be applied to set-based recognition. The spatial cues and explicit redundancy-reduction are further incorporated in RSA. A novel reconstruction-based sparse/collaborative similarity and loss are proposed in DFA.
- The proposed PIFR with alternate optimization can be a general framework to integrate the parametric deep learning and non-parametric dictionary learning.

We apply the proposed method to several image set-based verification/identification tasks and empirically demonstrate its effectiveness and generality.

## 2. Related Works

**Representation of sets.** The most straightforward permutation-invariant operations are the max/average pooling over a set of embeddings  $X=\{x_n\}_{n=1}^N$  [21, 43, 5]. Zaheer *et al.* [63] has proven that the function that can be represented the

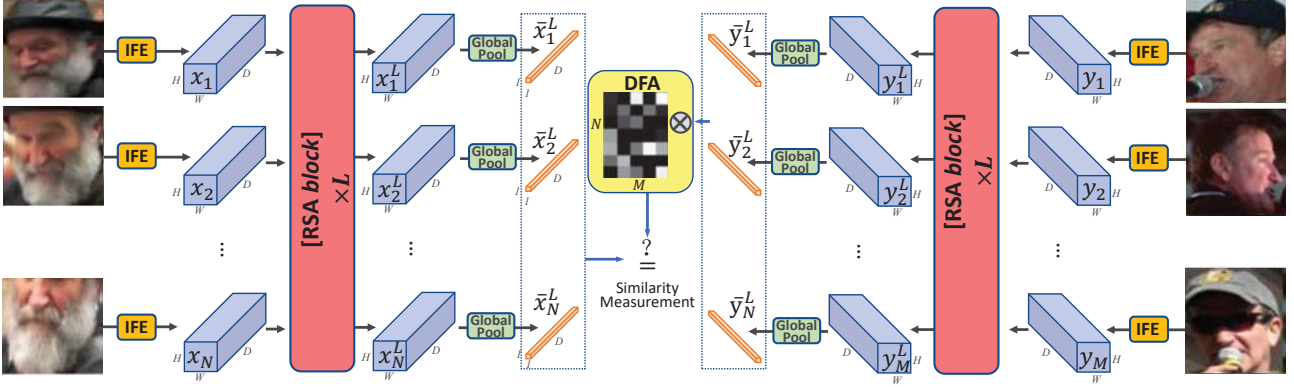


Figure 3. Our *permutation-invariant feature restructuring (PIFR)* framework for image set-based recognition, which consists of the fully convolutional image feature extractor (IFE), redundancy-eliminating self-attention blocks for inner-set restructuring (RSA $\times L$ ), global pooling layer and sparse/collaborative representation-based dependency-guided feature alignment (DFA) module for inter-set restructuring.

following form is permutation-invariant:

$$f(X) = \rho\left(\sum_{x \in X} (\varphi(x))\right) \quad (1)$$

However,  $\varphi$  acts independently on each element, which does not allow an exploration of the inter-sample relationships. This limitation is also in the methods in [13, 10].

**Image set based recognition.** A video can be considered as an image set with ordered images and has been actively studied, while a general image set is orderless and usually contains more challenging variations. Thus the methods relying on temporal dynamics will not be covered here. Traditional methods usually model the set/video as appearance subspaces or manifolds [51, 19, 49].

Along a different axis, [61, 40] are based on computing a score for each image with neural image assessment modules. Then, a set of features are aggregated to a fixed size feature vector via weighted pooling. Without inner-set interactions, this may result in redundancy and not use the information diversity present in a set.

[25] proposes to exploit the inner-set relationship using the RL. However, it is permutation-variant since the weight of the first image is computed based on  $[x_1; \frac{x_2 + \dots + x_N}{N-1}]$ , while if we move that image to the second position, its weight will be based on the  $[x_2; \frac{a_1 x_1 + x_3 + \dots + x_N}{N-2+a_1}]$ , where  $[\cdot; \cdot]$  denotes the concatenation and  $a_1$  is the updated weight of the first image. Besides, the RL itself is usually unstable [17]. The feature dimension is compressed from 1024 in its GoogleNet backbone [45] to 128 which unavoidably weakens its representation ability. Different from these extensions, we address the inner-set redundancy by constructing a feature vector that takes into account all of its neighbors and results in a set of redundancy-reduced features.

Using a single set-level feature vector is inadequate for representing inter-set interactions. [25, 56] target the inter-set pose-aware alignment relying on pose/landmark detec-

tion, and the parametric pair-wise input network in [56] is computationally inefficient for open-set identification. Our PIFR is end-to-end trained given only the set-level identity annotations, and is compatible with multiple inter-set variations and is scalable for open-set identification.

**Self-attention and non-local filtering.** As attention models grow in popularity, [48] develops a self-attention mechanism for machine translation. It calculates the response at one position as a weighted sum of all positions in the sentences. A similar idea is also inherited in the non-local algorithm [2], which is a classical image denoising technique. The interaction networks were also developed for modeling pair-wise interactions [1, 18, 53, 60]. Moreover, [52] proposes to bridge self-attention to the more general non-local filtering operations. [65] proposes to learn temporal dependencies between video frames at multiple time scales. Inspired by above works, we further adapt this idea to the set-based problem, the difference with the aforementioned studies includes, but is not limited to incorporating spatial clues with Gaussian similarity matrix, explicit redundancy-eliminating residual term and single-layer pairwise affinity.

**Sparse/collaborative representation based classification (SRC/CRC).** SRC [54] and CRC [64], and their many extensions [34, 7, 55] have been widely studied. They achieve promising single image and video-based [6] face recognition performance under occlusions and illumination variations. Our DFA is inspired by them, but using a novel similarity measurement instead of reconstruction coefficients.

### 3. Methodology

We consider the image-set based recognition as matching the probe set  $P$  containing  $N$  images (*i.e.*,  $P = \{p_n\}_{n=1}^N$ ) with the gallery set  $G$  which consist of  $M$  images (*i.e.*,  $G = \{g_m\}_{m=1}^M$ ),  $n, m$  are the image/video frames indexes. As illustrated in Fig. 3, our permutation-invariant feature restructuring (PIFR) framework consists of three ma-



	Layers		Output size
	conv	pool	
Image Feature Extractor	conv	7 × 7, 64, stride 2	112 × 112 × 64 × N
	pool	max pool, 3 × 3, stride 2	56 × 56 × 64 × N
	1-3 res blocks	conv 1 × 1, 64 conv 3 × 3, 64 conv 1 × 1, 256	56 × 56 × 256 × N
	4-7 res blocks	conv 1 × 1, 128 conv 3 × 3, 128 conv 1 × 1, 512	28 × 28 × 512 × N
	8-13 res blocks	conv 1 × 1, 256 conv 3 × 3, 256 conv 1 × 1, 1024	14 × 14 × 1024 × N (H × W × D × N)
RSA	RSABlock × L		14 × 14 × 1024 × N
	Global pool, 14 × 14		1 × 1 × 1024 × N

Table 2. The structure details of our IFE and RSA modules. We adopt the pre-trained ResNet-50 backbone, and utilize the output of  $3^{rd}$  resblock as our  $\{x_n\}_{n=1}^N$  on IJB-A&Celebriety-1000.

major components: fully convolutional image feature extractor (IFE), a group of redundancy-eliminating self attention (RSA) blocks and a dependency-guided feature alignment (DFA) module. To represent the spatial correlations, we choose the convolution layers of a deep convolutional neural network (CNN) networks with high-end recognition performances (e.g., GoogleNet[45], ResNet50[16]) as our IFE, instead of either hand-crafted extractors or CNN with fully connected layers used in [25, 61]. It embeds the images into latent space independently and generates the corresponding feature sets  $X = \{x_n\}_{n=1}^N$  and  $Y = \{y_m\}_{m=1}^M$  from  $P$  and  $G$  respectively.  $x_n, y_m \in \mathbb{R}^{H \times W \times D}$  where  $H, W$  and  $D$  are the height, width and channel dimension, respectively of our representation. Its parameters are fixed after pre-training by a large scale image-based recognition datasets.

### 3.1. Redundancy-eliminating Self Attention

As shown in Fig. 4, we cascade  $L \times$  RSA blocks, which works as a residual self-attention module receiving all extracted features  $\{x_n\}_{n=1}^N$  and restructuring them based on inner-set correlations. Since the features  $\{x_n\}_{n=1}^N$  are deterministically computed from the probe set  $\{p_n\}_{n=1}^N$ , they also inherit and display large variations and redundancy. Ideally, the RSA should display four attributes: Ability to handle arbitrary number of samples, inner-set correlation-awareness, permutation invariance to the features, and sensitivity to relative locations in the image plane.

We propose to harvest the spatially-aware inner-set correlations by exploiting the affinity of point-wise feature vectors. We use  $i = 1, \dots, H \times W$  to index the position in HW plane and the  $j$  is the index for all  $D$ -dimensional feature vectors other than the  $i^{th}$  vector ( $j = 1, \dots, H \times W \times N - 1$ ). Specifically, our RSA block can be formulated as

$$x_{n,i}^l = x_{n,i}^{l-1} + \frac{\Omega^l}{C_{n,i}} \sum_{\forall n,j} \omega(x_{n,i}^0, x_{n,j}^0) (x_{n,j}^{l-1} - x_{n,i}^{l-1}) \Delta_{i,j} \quad (2)$$

$$C_{n,i} = \sum_{\forall n,j} \omega(x_{n,i}^0, x_{n,j}^0) \Delta_{i,j}$$

where  $\Omega^l \in \mathbb{R}^{1 \times 1 \times D}$  is the weight vector to be learned,  $l = 0, 1, \dots, L$  with  $L$  being the number of stacked sub-

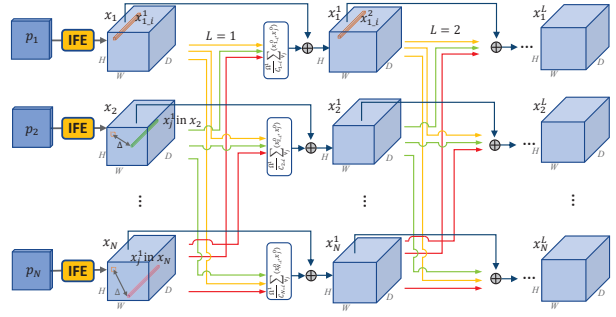


Figure 4. Redundancy-eliminating self attention blocks.

self attention blocks and  $x_n^0 = x_n$ . The pairwise affinity  $\omega(\cdot, \cdot)$  is a scalar. The response is normalized by  $C_{n,i}$ .

To explore the spatial clues, we propose to use a Gaussian similarity metric  $\Delta_{i,j} = \exp(\frac{\|hw_{n,i} - hw_{n,j}\|_2^2}{\sigma})$ , where  $hw_{n,i}, hw_{n,j} \in \mathbb{R}^2$  represent the position of  $i^{th}$  and  $j^{th}$  vector in the HW-plane of  $x_n$ , respectively.

The residual term is the difference between the neighboring feature (i.e.,  $x_{n,j}^{l-1}$ ) and the computed feature  $x_{n,i}^{l-1}$ . If  $x_{n,j}^{l-1}$  incorporates complementary information and has better imaging/content quality compared to  $x_{n,i}^{l-1}$ , then RSA will erase some information of the inferior  $x_{n,i}^{l-1}$  and replaces it by the more discriminative feature representation  $x_{n,j}^{l-1}$ . Compared to the method of using only  $x_{n,j}^{l-1}$  [52, 65], our setting shares more common features with diffusion maps [9, 46], graph Laplacian [8] and non-local image processing [3, 14]. All of them are non-local analogues [12] of local diffusions, which are expected to be more stable than its original non-local counterpart [2, 52] due to the nature of its inherit Hilbert-Schmidt operator [12].

The operation of  $\omega$  in Eq. (2) is not sensitive to many function choices [52, 65]. We simply choose the embedded Gaussian given by

$$\omega(x_{n,i}^0, x_{n,j}^0) = e^{\psi(x_{n,i}^0)^T \phi(x_{n,j}^0)} \quad (3)$$

where  $\psi(x_{n,i}^0) = \Psi x_{n,i}^0$  and  $\phi(x_{n,j}^0) = \Phi x_{n,j}^0$  are two embeddings, and  $\Psi, \Phi$  are matrices to be learned.

Different from [52], our affinity  $\omega$  is pre-computed given the input feature set  $\{x_n\}_{n=1}^N$  and stays the same through the propagation within RSA, which reduces the computational cost while  $\omega(x_{n,i}^0, x_{n,j}^0)$  can still represent the affinity between  $x_{n,i}^{l-1}$  and  $x_{n,j}^{l-1}$  to some extent. This design reduces the number of parameters and speeds up the learning.

**Definition 1.** Let  $\pi$  be an arbitrary permutation function for a sequence. We say a function  $f(X)$  is *permutation-invariant* iff for  $\forall \pi; f(X) = f(\pi(X))$ .

**Property 1.** The RSA is *permutation-invariant*, due to the fact that all features in  $\{x_n\}_{n=1}^N$  are considered in Eq. (2) in an order-less manner (sum operation).

The formulation in Eq. (2) supports the inputs of variable number of features, and maintains the corresponding

number and dimension of features in the output tensor  $X^L$  after several blocks are calculated consecutively. We apply the global pooling (GP) [23] to summarize the information in HW-plane, and get the final representations:

$$\begin{cases} \bar{X}^L = \{\bar{x}_1^L, \dots, \bar{x}_N^L\}, \bar{x}_n^L = \text{GP}(x_n^L) \in \mathbb{R}^{1 \times 1 \times D} \\ \bar{Y}^L = \{\bar{y}_1^L, \dots, \bar{y}_M^L\}, \bar{y}_m^L = \text{GP}(y_m^L) \in \mathbb{R}^{1 \times 1 \times D} \end{cases} \quad (4)$$

RSA learns to aggregate inner-set correlations for each individual position of a sample's feature adaptively. It can be flexibly attached to any pre-trained IFE (by initializing  $\Omega^l$  as zero) and possible to be fine-tuned jointly with IFE.

### 3.2. Dependency-guided Feature Alignment

Beyond specifically aligning the poses, illuminations, expressions *etc.*, [25, 62], we propose to measure the dependency-guided similarity between a pair of feature sets with different feature numbers leveraging the non-parametric dictionary learning methods [34].

$\bar{x}_n^L$  is expected to be represented using a linear combination of  $\{\bar{y}_1^L, \dots, \bar{y}_M^L\}$ , which means that we can search similar feature vectors in  $\bar{Y}^L$  to reconstruct  $\bar{x}_n^L$ . It transforms the matching task into a sparse/collaborative representation problem:

$$\min_{a_n} \|a_n\|_p \quad \text{s.t.} \quad \bar{x}_n^L = \bar{Y}^L a_n \quad (5)$$

where  $a_n \in \mathbb{R}^{M \times 1}$  is the sparse/collaborative coding vector of  $\bar{x}_n^L$  w.r.t.  $\bar{Y}^L$ . The  $\ell_p$ -norm is used to constrain  $a_n$ , which leads to fewer features in  $\bar{Y}^L$  to reconstruct the  $\bar{x}_n^L$ .  $p = 1$  and  $2$  are adopted in sparse and collaborative representation-based classification, respectively. However, SRC/CRC is proposed to reconstruct a probe image based on all of the gallery images and uses the coefficient  $a_n$  to rank similar samples. Our method is different from that, since the objectives are the samples in the sets and we do not use  $\|a_n\|_p$  directly for similarity measurement. We can relax the constraint term and rewrite Eq. (5) as

$$\min_{a_n} \|\bar{x}_n^L - \bar{Y}^L a_n\|_2^2 + \frac{\lambda}{M} \|a_n\|_p \quad (6)$$

where  $\|\bar{x}_n^L - \bar{Y}^L a_n\|_2^2$  is the  $\ell_2$  distance of  $\bar{x}_n^L$  and  $\bar{Y}^L$ , and  $\lambda$  controls the sparsity of coefficients  $a_n$ . For a probe set with  $N$  samples, we need a sparse/collaborative reconstruction coefficient matrix  $A = \{a_1, \dots, a_N\} \in \mathbb{R}^{M \times N}$ , and the set-to-set verification is reformulated to calculate the set-level similarity, which can be measured by  $\frac{1}{N} \|\bar{X}^L - \bar{Y}^L A\|_2^2$ . It is essentially different from previous metric learning-based algorithms for two sets.

As such, DFA is essentially accounting for all possible variations to further boost the performance at the extremes of these variations, and there is no need for addi-

---

#### Algorithm 1: Bi-level training of PIFR

---

**Input:** Training sets and labels;  $\lambda$ ; Pre-trained IFE

**Output:** Updated RSA parameters  $\theta$

---

```

1 Level 1: Pre-train RSA;
2 for all IFE extracted training sets X and Y do
3   Reconstruct  $\{x_n\}_{n=1}^N, \{y_m\}_{m=1}^M$  to  $X^L, Y^L$  using RSA
4    $\bar{X}^L = \text{Global pool}(X^L)$ ;  $\bar{X}^L = \text{Average pool}(\bar{X}^L)$ 
5   Compute cross entropy loss and update RSA by SGD
6 end
7 Level 2: Alternative optimization with RSA and DFA;
8 for all X and Y pairs do
9   1) Compute index function  $\alpha$  with training set labels
10  2) Compute  $\{a_n\}_{n=1}^N$  in DFA following Eq. (8)
11  3) Compute the reconstruction error loss in Eq. (7)
12  4) Update the parameters  $\theta$  by SGD following Eq. (9)
13 end

```

---

tional pose/landmark annotations. Since there are no additional constraints for the relations of the elements in A, it is also permutation-invariant.

In the testing stage of an open-set identification task with  $K$  gallery sets, we need  $K$  times verification, where  $K$  can be a few hundreds (*e.g.*, IJB-A) or even millions in real world applications. The pair-wise input network [56] for inter-set interaction need to process  $K$  query-gallery image set pairs. Considering the limited memory of GPUs, we usually cannot put in  $2K$  image sets and several network models. Its processing time equals to  $K$  times of verification, which is impractical. However, our gallery  $\bar{Y}^L$ s can be pre-computed and stored. We only need to extract  $\bar{X}^L$  from a query set, and the  $K$  times similarity measurements is applied to  $\bar{X}^L$  and  $\bar{Y}^L$  vector set pairs instead of original image. Unlike parametric network-based scheme, our DFA does not take a lot of memory. Therefore, we can process numerous  $\bar{X}^L$  and  $\bar{Y}^L$  pairs in parallel. Actually,  $a_1, \dots, a_N$  are also solved using a parallel algorithm [20, 64]. The separated and non-parametric DFA enables our framework to be practicable for open-set identification.

### 3.3. Bi-level Training and Alternative Optimization

We note that the DFA is nonparametric and does not have a training phase. For the closed-set identification, there are no inter-set interaction, and the RSA only model is applied. The nonparametric frame-wise average pooling is applied to  $\bar{X}^L$ , which aggregates features into a fixed-size representation  $\tilde{X}^L \in \mathbb{R}^{1 \times 1 \times D}$ , followed by a softmax output layer to calculate the cross-entropy loss. For the verification or open-set identification, the RSA only model is trained by contrastive loss as in [61] using  $\tilde{X}^L$ , while the PIFR is trained following a bi-level scheme using paired  $\bar{X}^L$  &  $\bar{Y}^L$ .

First, we pre-train the RSA module following closed-set




Figure 5. Examples of DFA. The coefficients almost align the pose variation, but essentially take all possible inter-set variations into account. Following [25], we rotate images to frontal/right face in pre-processing stage. The value of red points  $a_n$  represents the weight of each gallery feature and is a row of the DFA matrix in Fig. 3. The CRC has fewer 0 values than SRC.

identification setting. In the second stage, the DFA module is added, and the dictionary learning similarity is used as the fine-tuning signal. The loss function is defined as

$$\mathcal{L}(\theta, A) = \frac{\alpha}{N} \|\bar{X}^L - \bar{Y}^L A\|_2^2 + \frac{\lambda}{N \times M} \|A\|_p \quad (7)$$

where  $\theta(\Omega^l, \Psi, \Phi)$  denotes the parameters in RSA and implements as in [52, 65],  $\alpha$  is an indicator function which takes value 1 when the probe and gallery set have the same identity and value -1 for different identity.  $\lambda$  is usually set to 1 for simply. This objective explicitly encourages the feature sets from the same subject to be close while the feature sets of different subjects are far away from each other.

To minimize the loss in Eq. (7), we use the following alternative optimization algorithm to incorporate a parametric free module into deep models:

**a) Fix (initialize)  $\theta$  and minimize  $\mathcal{L}$  w.r.t.  $A$ .** The task in this step is to solve the optimal sparse/collaborative coding matrix  $A$ . In practice, this is achieved by solving  $a_1, \dots, a_N$  with parallel computing mode. Eq. (6) can be further rewritten as follows:

$$\min_{a_n} \frac{1}{2} a_n^T \bar{Y}^{LT} \bar{Y}^L a_n - \bar{x}_n^T \bar{Y}^L a_n + \frac{\lambda}{M} \|a_n\|_p \quad (8)$$

The feature-sign searching algorithm [20, 64] is utilized to find the optimal  $a_n$ .

**b) Fix  $A$  and minimize  $\mathcal{L}$  w.r.t.  $\theta$ .** One of the remarkable properties of  $\mathcal{L}$  is differentiable. We update the parameters of RSA using the standard Stochastic Gradient Descent (SGD) by the following gradients of  $\mathcal{L}$  w.r.t.  $\bar{X}$  and  $\bar{Y}$

$$\begin{cases} \frac{\partial \mathcal{L}(\theta)}{\partial \bar{X}} = \frac{2\alpha}{N} (\bar{X}^L - \bar{Y}^L A) \\ \frac{\partial \mathcal{L}(\theta)}{\partial \bar{Y}} = -\frac{2\alpha}{N} (\bar{X}^L - \bar{Y}^L A) A^T \end{cases} \quad (9)$$

The overall bi-level training is outlined in Algorithm 1. The similarity value learned through PIFR are self-adaptive and sensitive to spatial location and inner/inter set variations. The benefits of utilizing both of the RSA and DFA restructuring modules are demonstrated in experiments.

**Analysis.** Since IFE extracts  $\{x_n\}_{n=1}^N$  from  $\{p_n\}_{n=1}^N$  independently, RSA used to reconstruct the representation set is permutation-invariant and the computing of the  $a_n$  in DFA is a permutation-invariant operation. Combined with the fact that the similarity measurement is permutation-invariant as well, we have the following:

**Proposition 1.** The PIFR is permutation-invariant.

Note that we can deconstruct PIFR into two operation stages to align with Eq. (1). In RSA, the residual term is our  $\varphi$  and sum operation, which does not act independently on each item but considers the whole set to obtain the embeddings. In the second stage, the RSA can be regarded as  $\varphi$ . The sum operator is inherent in  $\|\cdot\|_2^2$  and  $\ell_p$ -norm. Then  $\rho$  is simply the *sqrt* function. Additionally, instead of a fixed function such as pooling, our inner-set relationship function is parameterized and can adapt to the task at hand.

## 4. Numerical Experiments

In this section, we first explore the meaning of the coefficients in DFA (see Fig. 5). Based on the above understanding, we evaluate PIFR on several benchmarks via exhaustive ablation studies. RSA is implemented on our Titan Xp GPU. In testing stage, our gallery  $\bar{Y}^L$ s are pre-computed and stored. The average verification time of collaborative DFA is only 93ms for IJB-A, when we use a single Xeon E5 v4 CPU. The detailed architectures for the individual modules are provided in the Supplementary file.

### Unconstrained Set-based face recognition

The IARPA Janus Benchmarks (IJB)-A is a set-based face recognition dataset, where all images and videos are collected in unconstrained environments with large variants in imaging condition and viewpoints. It contains 5,397 still images and 20,412 video frames from 500 identities, and the number of samples in a set ranges from 1 to 190 with an average of 11.4 images and 4.2 videos per subject.

To take advantage of the long-tested image-based network structure and millions of available training images, the backbone of IFE is adapted from the GoogleNet [45] /ResNet50 [16] and trained on 3M Indoor face images [61] or VGGFace2 [4] (3.31M images) respectively. We choose the same training details as [25, 57] for fair comparisons.

In the baseline methods, CNN+Mean $\ell_2$  measures the  $\ell_2$  similarity of all  $\bar{X}^L$  and  $\bar{Y}^L$  pairs. The CNN+AvePool is the average pooling along each feature dimension for aggregation, followed by a linear classifier [61, 25].

For the 1:1 *verification* task, the performances are shown using the receiver operating characteristics (ROC) curves in Fig. 6. We also report the true accept rate (TAR) *vs.* false positive rates (FAR) in Table 3.

The performance of 1:N *identification* is characterized by the cumulative match characteristics (CMC) curves shown in Fig. 6. It is an information retrieval metric, which

Method	IFE Backbone	Dataset	1:1 Verification TAR			1:N Identification TPIR		
			FAR=0.001	FAR=0.01	FAR=0.1	FPIR=0.01	FPIR=0.1	Rank-1
Pose-model[41]	VGGNet	WebFace	0.652±0.037	0.826±0.018	-	-	-	0.840±0.012
Masi <i>et al.</i> [42]	VGGNet	WebFace	0.725	0.886	-	-	-	0.906
Adaptation[11]	VGGNet	VGGNet	0.836±0.027	0.939±0.013	0.979±0.004	0.774±0.049	0.882±0.016	0.928±0.010
DR-GAN[47]	CASIA-Net	WebFace	0.539±0.043	0.774±0.027	-	-	-	0.855±0.015
QAN[40]	-	-	0.893±0.039	0.942±0.015	0.980±0.006	-	-	0.955±0.006
NAN[61]	GoogleNet	Indoor(3M)	0.881±0.011	0.941±0.008	0.978±0.003	0.817±0.041	0.917±0.009	0.958±0.005
DAC[25]	GoogleNet	Indoor(3M)	0.893±0.010	0.954±0.010	0.981±0.008	0.855±0.042	0.934±0.009	0.973±0.011
<b>RSA</b>	GoogleNet	Indoor(3M)	0.926±0.011	0.964±0.010	0.986±0.008	0.882±0.042	0.954±0.009	0.984±0.011
<b>DFA(<math>\ell_1</math>)</b>	GoogleNet	Indoor(3M)	0.912±0.017	0.957±0.011	0.983±0.005	0.873±0.045	0.940±0.015	0.976±0.009
<b>DFA(<math>\ell_2</math>)</b>	GoogleNet	Indoor(3M)	0.914±0.013	0.958±0.006	0.984±0.003	0.875±0.037	0.940±0.010	0.978±0.008
<b>PIFR(<math>\ell_1</math>)</b>	GoogleNet	Indoor(3M)	0.933±0.015	0.966±0.007	0.987±0.005	0.894±0.039	<b>0.958±0.013</b>	0.985±0.010
<b>PIFR(<math>\ell_2</math>)</b>	GoogleNet	Indoor(3M)	0.937±0.016	<b>0.968±0.005</b>	<b>0.988±0.004</b>	0.896±0.035	<b>0.958±0.007</b>	<b>0.988±0.006</b>
<b>PIFR(<math>\ell_2</math>)<math>\times 2</math></b>	GoogleNet	Indoor(3M)	<b>0.938±0.013</b>	<b>0.968±0.006</b>	0.988±0.006	<b>0.897±0.031</b>	<b>0.958±0.009</b>	<b>0.988±0.007</b>
VGG Face2[4]	ResNet50	VGGFace2	0.895±0.019	0.950±0.005	0.980±0.003	-	-	-
MNet[57]	ResNet50	VGGFace2	0.920±0.013	0.962±0.005	0.989±0.002	-	-	-
<b>RSA</b>	ResNet50	VGGFace2	0.943±0.008	0.976±0.006	0.991±0.002	0.886±0.041	0.962±0.009	0.987±0.005
<b>PIFR(<math>\ell_2</math>)</b>	ResNet50	VGGFace2	<b>0.955±0.010</b>	<b>0.983±0.004</b>	<b>0.993±0.003</b>	<b>0.908±0.028</b>	<b>0.969±0.009</b>	<b>0.990±0.005</b>

Table 3. Evaluation on the IJB-A dataset. For verification, the true accept rates (TAR) vs. false positive rates (FAR) are presented. For identification, the true positive identification rate (TPIR) vs. false positive identification rate (FPIR) and the Rank-1 accuracy are reported.

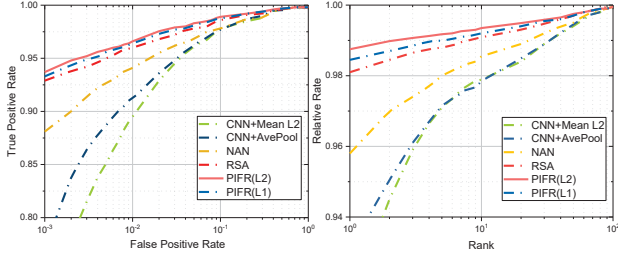


Figure 6. Average ROC (Left) and CMC (Right) curves of the proposed method (with GoogleNet backbone) and its baselines on the IJB-A dataset over 10 splits.

plots identification rates corresponding to different ranks. A rank- $k$  identification rate is defined as the percentage of probe searches whose gallery match is returned within the top- $k$  matches. The true positive identification rate (TPIR) *vs.* false positive identification rate (FPIR), as well as the rank-1 accuracy, are shown in Table 3.

The GoogleNet-based RSA, DFA( $\ell_2$ ), PIFR( $\ell_2$ ) outperforms DAC [25] by 2.6%/3.3%, 2.0%/2.1%, **4.2%/4.5%** in terms of TPIR@FPIR=10<sup>-2</sup>/TAR@FAR=10<sup>-3</sup> respectively. We note that the training and testing time of PIFR( $\ell_2$ ) are 3 $\times$  and 1.5 $\times$  smaller than DAC [25]. These results not only show that our model improve both the verification and identification performance significantly, but also indicates the RSA and DFA can be well combined in our alternative optimization framework and are complementary to each other. The improvement over the previous state-of-the-art *w.r.t.* rank-1 accuracy is also more notable than of recent works [61, 25]. Using the ResNet50 backbone, our PIFR( $\ell_2$ ) achieves state-of-the-art 99.0% and outperforms [57] by 3.5% for TAR@FAR=10<sup>-3</sup>.

The performance of PIFR( $\ell_2$ ) slightly exceeds PIFR( $\ell_1$ ), indicating that the collaborative representation is better for set feature reconstructions. A possible reason is that SRC usually assigns the coefficient strictly to 0 (see Fig. 5), which is similar to hard rather than soft attention [58].

The PIFR( $\ell_2$ ) $\times 2$  represents the symmetric dictionary learning which reconstructs  $\bar{x}_n^L$  using  $\{\bar{y}_1^L, \dots, \bar{y}_M^L\}$ ,




Figure 7. On IJB-A dataset, selecting the output layer of IFE (left) and the number of RSA blocks (right) matters. The low-level features are less discriminative, while the higher-level features lose too much spatial information. The performance usually saturated with more than 5 RSA blocks.

and reconstructs  $\bar{y}_m^L$  using  $\{\bar{x}_1^L, \dots, \bar{x}_N^L\}$  simultaneously. However, the doubled complexity does not lead to improved performance. Although the GoogleNet and ResNet50 are designed in different structures and trained on different datasets, PIFR can be applied on top of either of them.

Then PIFR’s sensitivity to the level of feature extraction and the number of RSA blocks is analyzed in Fig. 7.

#### Unconstrained video-based face recognition

The Celebrity-1000 dataset is a video-based unconstrained face identification dataset, but used as a typical set-based benchmark. It contains 2.4M frames from 159,726 videos of 1,000 subjects. We follow the protocol in [61, 25] and report the performances in Table 4&5 and Fig. 8.

In the closed-set testing, we choose the restructured representations from the 5<sup>th</sup> RSA blocks, and use the softmax output of our RSA only model. The RSA outperforms all of the previous methods consistently following the setting without fine-tuning the IFE module on Celebrity-1000 dataset. The video sequences are obviously redundant, and the inner-set correlation modeling does contribute to the improvements over [61]. The low quality of the several former frames also have negative impact of [25].

For the open-set protocol, we further measure the collaborative representation similarity to rank the gallery subjects. We see that our GoogleNet based PIFR outperforms the average pooling baseline and DAC by more than 9% and 3%



Method	Temp info	Number of Subjects( <i>closed</i> )			
		100	200	500	1000
NAN(GoogleNet)[61]	no	0.9044	0.8333	0.8227	0.7717
DAC(GoogleNet)[25]	no	0.9137	0.8783	0.8523	0.8353
RSA(GoogleNet)	no	<b>0.9385</b>	<b>0.9047</b>	<b>0.8859</b>	<b>0.8572</b>
RSA(ResNet50)	no	<b>0.9458</b>	<b>0.9140</b>	<b>0.9943</b>	<b>0.8671</b>

Table 4. Closed-set identification performance (rank-1 accuracies), on the Celebrity-1000 dataset. This setting does not have inter-set comparison, and only RSA is necessary.

Method	Temp info	Number of Subjects( <i>open</i> )			
		100	200	500	800
CNN+Mean $\ell_2$ (GoogleNet)	no	0.8488	0.7988	0.7676	0.7067
AvePool(GoogleNet)	no	0.8411	0.7909	0.7840	0.7512
NAN(GoogleNet)[61]	no	0.8876	0.8521	0.8274	0.7987
DAC(GoogleNet)[25]	no	0.9004	0.8715	0.8428	0.8264
RSA(GoogleNet)non-local	no	0.8987	0.8680	0.8376	0.8251
RSA(GoogleNet)w/o $\Delta$	no	0.9104	0.8812	0.8545	0.8420
RSA(GoogleNet)- $x_{n-j}^{l-1}$	no	0.9097	0.8791	0.8603	0.8388
RSA(GoogleNet)- $(\Psi\Phi)^L$	no	0.9297	0.8990	0.8764	0.8556
RSA(GoogleNet)	no	0.9297	0.8991	0.8763	0.8558
DFA( $\ell_2$ )(GoogleNet)	no	0.9105	0.8807	0.8562	0.8383
PIFR( $\ell_2$ )(GoogleNet)	no	<b>0.9355</b>	<b>0.9189</b>	<b>0.8876</b>	<b>0.8604</b>
PIFR( $\ell_2$ )(ResNet50)	no	<b>0.9497</b>	<b>0.9308</b>	<b>0.8922</b>	<b>0.8713</b>

Table 5. Open-set identification performance (rank-1 accuracies) on the Celebrity-1000 dataset

at all measurements. More appealingly, the use of ResNet50 can further enhance our framework and achieve the state-of-the-art results with more than 4% improvement over [25].

We did not manage to get competitive results using original non-local [52] in our RSA. We note that non-local is also first introduced to set-based recognition in this paper, and can be regarded as a baseline. For ablation study, in RSAw/o $\Delta$ , we excise  $\Delta$  in Eq. (2).  $x_{n-j}^{l-1}$  refers to using  $x_{n-j}^{l-1}$  as residue term instead of the difference in Eq. (2). Their inferior performance demonstrates the effectiveness of our choices.  $(\Psi\Phi)^L$ , denotes using embedded Gaussian pair-wise affinity for every RSA blocks, which has similar performance but takes more than  $2\times$  training time.

### Person re-identification

To verify the effectiveness of PIFR on full body images, we carry out additional experiments on the iLIDS-VID [50] benchmark which was created at an airport arrival hall under a multi-camera CCTV network, and contains 300 people in total. Its image sequences (rangings from 23 to 192 images) were accompanied by clothing similarities among people, lighting and viewpoint variations, cluttered background and occlusions. We follow the evaluation protocol from [22] for 10-fold cross evaluation. We sample the frames with a stride of 5, and the remaining frames, in the end, are simply discarded. Notice that since the data pre-processing, training setting and network structure vary in different state-of-the-art methods, we only list recent best-performing methods in the tables just for reference. For Re-ID, the widely used Cumulative Match Curve (CMC) is adopted.

We use the same IFE as in [22], and set  $L = 5$ . The PIFR with ResNet50 backbone is comparable to the TCP which is based on more powerful ResNet101 and extra training




Figure 8. CMC curves of different methods on Celebrity 1000. (left) Close-set identification on 1000 subjects, (right) Open-set identification on 800 subjects.

Method	Temp info	iLIDS-VID			
		CMC-1	CMC-5	CMC-10	CMC-20
TAM+SRM[66]	yes	0.552	0.865	-	0.970
ASTPN[59]	yes	0.620	0.860	0.940	0.980
QAN[40]	no	0.680	0.868	0.954	0.974
RQEN[44]	no	0.771	0.932	0.977	<u>0.994</u>
STAN(ResNet50)[22]	no	<u>0.802</u>	-	-	-
TCP(ResNet101)[39]	no	0.717	<b>0.951</b>	<b>0.983</b>	0.993
PIFR( $\ell_2$ )(ResNet50)	no	<b>0.825</b>	0.950	0.980	<b>0.997</b>

Table 6. Experimental results of the proposed and other comparisons on iLIDS-VID re-identification dataset. The best are in bold while the second best are underlined.

data. Benefiting from the more efficient inner-set complementary/diversity exploring and additional inter-set correlation module, PIFR achieves 82.5% rank-1 accuracy dispense with the optical flow information (see Table 6). It outperforms the STAN and TCP by 2.3% and 10.8% respectively, and  $2\times$  faster than STAN *w.r.t.* training.

## 5. Conclusions

We present a novel permutation-invariant feature restructuring framework (PIFR) for set-based representation and similarity measurement. Both the RSA and DFA module inherit the restructuring idea (in a parametric/non-parametric manner) to take into account the inner/inter-set correlations respectively. It can emphasize the discriminative images, reduces the redundancy, adaptively aligns the probe-gallery set and is scalable to open-set identification. The proposed alternative optimization can be a unified way to train the deep learning and dictionary learning jointly. The PIFR can be a general framework for the task with variable number of order-less samples. We plan to apply it to vehicle re-id, action recognition, relation reasoning etc., in the future.

## 6. Acknowledgement

The funding support from Fanhan Information Tech, Hong Kong Government General Research Fund GRF (No.152202/14E), Natural Science Foundation of China (NSFC) (Nos. 61772296, 61627819, 61727818, U1713214) and Shenzhen fundamental research fund (Grant No. JCYJ20170412170438636) are greatly appreciated.



## References

- [1] Peter Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, et al. Interaction networks for learning about objects, relations and physics. In *Advances in neural information processing systems*, pages 4502–4510, 2016. 3
- [2] Antoni Buades, Bartomeu Coll, and J-M Morel. A non-local algorithm for image denoising. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 2, pages 60–65. IEEE, 2005. 3, 4
- [3] Antoni Buades, Bartomeu Coll, and Jean-Michel Morel. A review of image denoising algorithms, with a new one. *Multiscale Modeling & Simulation*, 4(2):490–530, 2005. 4
- [4] Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *Automatic Face & Gesture Recognition (FG 2018), 2018 13th IEEE International Conference on*, pages 67–74. IEEE, 2018. 6, 7
- [5] Jun-Cheng Chen, Rajeev Ranjan, Amit Kumar, Ching-Hui Chen, Vishal M Patel, and Rama Chellappa. An end-to-end system for unconstrained face verification with deep convolutional neural networks. In *IEEE CVPRW*, pages 118–126, 2015. 1, 2
- [6] Yi-Chen Chen, Vishal M Patel, P Jonathon Phillips, and Rama Chellappa. Dictionary-based face recognition from video. In *European conference on computer vision*, pages 766–779. Springer, 2012. 3
- [7] Yuejie Chi and Fatih Porikli. Classification and boosting with multiple collaborative representations. *IEEE transactions on pattern analysis and machine intelligence*, 36(8):1519–1531, 2014. 3
- [8] Fan RK Chung and Fan Chung Graham. *Spectral graph theory*. Number 92. American Mathematical Soc., 1997. 4
- [9] Ronald R Coifman and Stéphane Lafon. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30, 2006. 4
- [10] Andrew Cotter, Maya Gupta, Heinrich Jiang, James Muller, Taman Narayan, Serena Wang, and Tao Zhu. Interpretable set functions. *arXiv preprint arXiv:1806.00050*, 2018. 3
- [11] Nate Crosswhite, Jeffrey Byrne, Chris Stauffer, Omkar Parkhi, Qiong Cao, and Andrew Zisserman. Template adaptation for face verification and identification. In *FG*, pages 1–8. IEEE, 2017. 7
- [12] Qiang Du, Max Gunzburger, Richard B Lehoucq, and Kun Zhou. Analysis and approximation of nonlocal diffusion problems with volume constraints. *SIAM review*, 54(4):667–696, 2012. 4
- [13] Harrison Edwards and Amos Storkey. Towards a neural statistician. *arXiv preprint arXiv:1606.02185*, 2016. 3
- [14] Guy Gilboa and Stanley Osher. Nonlocal linear image regularization and supervised segmentation. *Multiscale Modeling & Simulation*, 6(2):595–630, 2007. 4
- [15] Patrick Grother and Mei Ngan. Face recognition vendor test (frvt). *Performance of face identification algorithms*, 2014. 1
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4, 6
- [17] Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. Deep reinforcement learning that matters. *arXiv preprint arXiv:1709.06560*, 2017. 3
- [18] Yedid Hoshen. Vain: Attentional multi-agent predictive modeling. In *Advances in Neural Information Processing Systems*, pages 2701–2711, 2017. 3
- [19] Zhiwu Huang, Jiqing Wu, and Luc Van Gool. Building deep networks on grassmann manifolds. *arXiv preprint arXiv:1611.05742*, 2016. 3
- [20] Honglak Lee, Alexis Battle, Rajat Raina, and Andrew Y Ng. Efficient sparse coding algorithms. In *Advances in neural information processing systems*, pages 801–808, 2007. 5, 6
- [21] Haoxiang Li, Gang Hua, Xiaohui Shen, Zhe Lin, and Jonathan Brandt. Eigen-pep for video face recognition. In *Asian Conference on Computer Vision*, pages 17–33. Springer, 2014. 1, 2
- [22] Shuang Li, Slawomir Bak, Peter Carr, and Xiaogang Wang. Diversity regularized spatiotemporal attention for video-based person re-identification, 2018. 8
- [23] Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013. 5
- [24] Xiaofeng Liu. Research on the technology of deep learning based face image recognition. In *Thesis*, 2019. 1
- [25] Xiaofeng Liu, Kumar B.V.K, Chao Yang, Qingming Tang, and Jane You. Dependency-aware attention control for unconstrained face recognition with image sets. In *European Conference on Computer Vision*, 2018. 1, 2, 3, 4, 5, 6, 7, 8
- [26] Xiaofeng Liu, Yi Ge, Lingsheng Kong, Ping Jia, You Jane, and Jun Lu. Unimodal regularized neuron stick-breaking for ordinal regression. In *ArXiv*, 2019. 1
- [27] Xiaofeng Liu, Yubin Ge, Chao Yang, and Ping Jia. Adaptive metric learning with deep neural networks for video-based facial expression recognition. *Journal of Electronic Imaging*, 27(1):013022, 2018. 1
- [28] Xiaofeng Liu, Zhenhua Guo, Jane Jia, and B.V.K. Kumar. Dependency-aware attention control for imageset-based face recognition. In *ArXiv*, 2019. 1
- [29] Xiaofeng Liu, Xu Han, Yukai Qiao, Yi Ge, Site Li, and Jun Lu. Unimodal-uniform constrained wasserstein training for medical diagnosis. In *ArXiv*, 2019. 1
- [30] Xiaofeng Liu, Lingsheng Kong, Zhihui Diao, and Ping Jia. Line-scan system for continuous hand authentication. *Optical Engineering*, 56(3):033106, 2017. 1
- [31] Xiaofeng Liu, BVK Vijaya Kumar, Yubin Ge, Chao Yang, Jane You, and Ping Jia. Normalized face image generation with perceptron generative adversarial networks. In *2018 IEEE 4th International Conference on Identity, Security, and Behavior Analysis (ISBA)*, pages 1–8, 2018. 1
- [32] Xiaofeng Liu, BVK Vijaya Kumar, Ping Jia, and Jane You. Hard negative generation for identity-disentangled facial expression recognition. *Pattern Recognition*, 88:1–12, 2019. 1

- [33] Xiaofeng Liu, Site Li, Lingsheng Kong, Wanqing Xie, Ping Jia, Jane You, and BVK Kumar. Feature-level frankenstein: Eliminating variations for discriminative recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 637–646, 2019. 1
- [34] Xiaofeng Liu, Zhaofeng Li, Lingsheng Kong, Zhihui Diao, Junliang Yan, Yang Zou, Chao Yang, Ping Jia, and Jane You. A joint optimization framework of low-dimensional projection and collaborative representation for discriminative classification. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 1493–1498. 2, 3, 5
- [35] Xiaofeng Liu, BVK Vijaya Kumar, Jane You, and Ping Jia. Adaptive deep metric learning for identity-aware facial expression recognition. In *CVPR Workshops*, pages 20–29, 2017. 1
- [36] Xiaofeng Liu, Yang Zou, Tong Che, Jane You, and Kumar B.V.K. Conservative wasserstein training for pose estimation. In *ICCV*, 2019. 1
- [37] Xiaofeng Liu, Yang Zou, Lingsheng Kong, Zhihui Diao, Junliang Yan, Jun Wang, Site Li, Ping Jia, and Jane You. Data augmentation via latent space interpolation for image classification. In *24th International Conference on Pattern Recognition (ICPR)*, pages 728–733, 2018. 1
- [38] Xiaofeng Liu, Yang Zou, Yuhang Song, Chao Yang, Jane You, and BV K Vijaya Kumar. Ordinal regression with neuron stick-breaking for medical diagnosis. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018. 1
- [39] Yu Liu, Guanglu Song, Jing Shao, Xiao Jin, and Xiaogang Wang. Transductive centroid projection for semi-supervised large-scale recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 70–86, 2018. 8
- [40] Yu Liu, Junjie Yan, and Wanli Ouyang. Quality aware network for set to set recognition. In *Proc. IEEE Int. Conf. Comput. Vision Pattern Recognit.*, pages 5790–5799, 2017. 1, 3, 7, 8
- [41] Iacopo Masi, Stephen Rawls, Gérard Medioni, and Prem Natarajan. Pose-aware face recognition in the wild. In *IEEE CVPR*, pages 4838–4846, 2016. 7
- [42] Iacopo Masi, Anh Tun Trn, Tal Hassner, Jatuporn Toy Leksut, and Gérard Medioni. Do we really need to collect millions of faces for effective face recognition? In *ECCV*, pages 579–596. Springer, 2016. 7
- [43] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, et al. Deep face recognition. In *BMVC*, volume 1, page 6, 2015. 1, 2
- [44] Guanglu Song, Biao Leng, Yu Liu, Congrui Hetang, and Shaofan Cai. Region-based quality estimation network for large-scale person re-identification, 2017. 8
- [45] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016. 3, 4, 6
- [46] Yunzhe Tao, Qi Sun, Qiang Du, and Wei Liu. Nonlocal neural networks, nonlocal diffusion and nonlocal modeling. *arXiv preprint arXiv:1806.00681*, 2018. 4
- [47] Luan Tran, Xi Yin, and Xiaoming Liu. Disentangled representation learning gan for pose-invariant face recognition. In *CVPR*, volume 3, page 7, 2017. 7
- [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017. 3
- [49] Ruiping Wang, Huimin Guo, Larry S Davis, and Qionghai Dai. Covariance discriminative learning: A natural and efficient approach to image set classification. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2496–2503. IEEE, 2012. 3
- [50] Taiqing Wang, Shaogang Gong, Xiatian Zhu, and Shengjin Wang. Person re-identification by video ranking. *Lecture Notes in Computer Science*, page 688703, 2014. 8
- [51] Wen Wang, Ruiping Wang, Zhiwu Huang, Shiguang Shan, and Xilin Chen. Discriminant analysis on riemannian manifold of gaussian distributions for face recognition with image sets. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2048–2057, 2015. 3
- [52] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3, 4, 6, 8
- [53] Nicholas Watters, Daniel Zoran, Theophane Weber, Peter Battaglia, Razvan Pascanu, and Andrea Tacchetti. Visual interaction networks: Learning a physics simulator from video. In *Advances in Neural Information Processing Systems*, pages 4539–4547, 2017. 3
- [54] John Wright, Allen Y Yang, Arvind Ganesh, S Shankar Sstry, and Yi Ma. Robust face recognition via sparse representation. *IEEE transactions on pattern analysis and machine intelligence*, 31(2):210–227, 2009. 3
- [55] Yang Wu, Vansteenberge Jarich, Masayuki Mukunoki, and Michihiko Minoh. Collaborative representation for classification, sparse or non-sparse? *arXiv preprint arXiv:1403.1353*, 2014. 3
- [56] Weidi Xie, Li Shen, and Andrew Zisserman. Comparator networks. In *European Conference on Computer Vision*, 2018. 2, 3, 5
- [57] Weidi Xie and Andrew Zisserman. Multicolumn networks for face recognition. 2018. 6, 7
- [58] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057, 2015. 7
- [59] Shuangjie Xu, Yu Cheng, Kang Gu, Yang Yang, Shiyu Chang, and Pan Zhou. Jointly attentive spatial-temporal pooling networks for video-based person re-identification. *2017 IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. 8
- [60] Flood Sung Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. 2018. 3
- [61] Jiaolong Yang, Peiran Ren, Dongqing Zhang, Dong Chen, Fang Wen, Hongdong Li, and Gang Hua. Neural aggregation

- network for video face recognition. In *IEEE CVPR*, pages 4362–4371, 2017. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [62] Xi Yin and Xiaoming Liu. Multi-task convolutional neural network for pose-invariant face recognition. *IEEE Transactions on Image Processing*, 2017. [1](#), [5](#)
- [63] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan R Salakhutdinov, and Alexander J Smola. Deep sets. In *Advances in Neural Information Processing Systems*, pages 3391–3401, 2017. [2](#)
- [64] Lei Zhang, Meng Yang, and Xiangchu Feng. Sparse representation or collaborative representation: Which helps face recognition? In *Computer vision (ICCV), 2011 IEEE international conference on*, pages 471–478. IEEE, 2011. [3](#), [5](#), [6](#)
- [65] Bolei Zhou, Alex Andonian, and Antonio Torralba. Temporal relational reasoning in videos. In *ECCV*, 2018. [3](#), [4](#), [6](#)
- [66] Zhen Zhou, Yan Huang, Wei Wang, Liang Wang, and Tieniu Tan. See the forest for the trees: Joint spatial and temporal recurrent neural networks for video-based person re-identification. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jul 2017. [8](#)