

GroupFace: Learning Latent Groups and Constructing Group-based Representations for Face Recognition

Yonghyun Kim¹ Wonpyo Park² Myung-Cheol Roh¹ Jongju Shin¹

¹Kakao Enterprise, ²Kakao Corp.

Abstract

In the field of face recognition, a model learns to distinguish millions of face images with fewer dimensional embedding features, and such vast information may not be properly encoded in the conventional model with a single branch. We propose a novel face-recognition-specialized architecture called GroupFace that utilizes multiple group-aware representations, simultaneously, to improve the quality of the embedding feature. The proposed method provides self-distributed labels that balance the number of samples belonging to each group without additional human annotations, and learns the group-aware representations that can narrow down the search space of the target identity. We prove the effectiveness of the proposed method by showing extensive ablation studies and visualizations. All the components of the proposed method can be trained in an end-to-end manner with a marginal increase of computational complexity. Finally, the proposed method achieves the state-of-the-art results with significant improvements in 1:1 face verification and 1:N face identification tasks on the following public datasets: LFW, YTF, CALFW, CPLFW, CFP, AgeDB-30, MegaFace, IJB-B and IJB-C.

1. Introduction

Face recognition is a primary technique in computer vision to model and understand the real world. Many methods and enormous datasets [3, 10, 16, 25, 32, 41] have been introduced, and recently, methods that use deep learning [7, 9, 14, 19, 34, 35, 44] have greatly improved the face recognition accuracy, but it still falls short of expectations.

To reduce the shortfall, most of the recent research in face recognition focused on improving the loss function. The streams from CenterLoss [35], CosFace [34], ArcFace [7] and RegularFace [44] all tried to minimize the intra-class variation and maximize the inter-class variation. These methods are effective and have gradually improved the accuracy by elaborating the objective of learning.

Despite the development of loss functions, general-

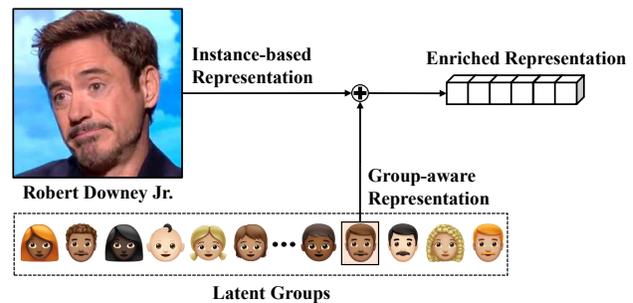


Figure 1. Conceptual scheme of the proposed method. The proposed method enhances the quality of the embedding feature by supporting the instance-based representation of **Robert Downey Jr.** using the group-aware representation considering **Brown-beard Man Group**.

purpose networks, not a network devised for a face recognition, can have difficulty in enabling effective training of the network to recognize a huge number of person identities. Unlike common problems such as classification, in the evaluation stage, a face-recognition model encounters new identities, which are not included in the training set. Thus, the model has to embed nearly 100k identities [10] in the training set and also consider a huge number of unknown identities. However, most of the existing methods just attach several fully-connected layers after widely-used backbone networks such as VGG [25] and ResNet [12] without any designs for the characteristics of face recognition.

Grouping is a key idea to efficiently-and-flexibly embed a significant number of people and briefly describe an unknown person. Each person has own characteristics in his or her face. At the same time, they have common ones shared in a group of people. In the real world, group-based description (man with deep, black eyes and red beard) that involves common characteristics in the group, can be useful to narrow down the set of candidates, even though it cannot identify the exact person. Unfortunately, explicit grouping requires manual categorizing on huge data and may be limited by the finite range of descriptions by human knowledge. However, by adopting the concept of grouping, the recogni-

tion network can reduce the search space and flexibly embed a significant number of identities into an embedding feature.

We propose a novel face-recognition architecture called GroupFace that learns multiple latent groups and constructs group-aware representations to effectively adopt the concept of grouping (Figure 1). We define Latent Groups, which are internally determined as latent variables by comprehensively considering facial factors (*e.g.*, hair, pose, beard) and non-facial factors (*e.g.*, noise, background, illumination). To learn the latent groups, we introduce a self-distributed grouping method that determines group labels by considering the overall distribution of latent groups. The proposed GroupFace structurally ensembles multiple group-aware representations into the original instance-based representation for face recognition.

We summarize the contributions as follows:

- GroupFace is a novel face-recognition-specialized architecture that integrates the group-aware representations into the embedding feature and provides well-distributed group-labels to improve the quality of feature representation. GroupFace also suggests a new similarity metric to consider the group information additionally.
- We prove the effectiveness of GroupFace in extensive experiments and ablation studies on the behaviors of GroupFace.
- GroupFace can be applied many existing face-recognition methods to obtain a significant improvement with a marginal increase in the resources. Especially, a hard-ensemble version of GroupFace can achieve high recognition-accuracy by adaptively using only a few additional convolutions.

2. Related Works

Face Recognition has been studied for decades. Many researchers proposed machine learning techniques with feature engineering [1, 5, 6, 28, 17, 23, 31, 38, 42]. Recently, deep learning methods have overcome the limitations of traditional face-recognition approaches with public face-recognition datasets [3, 10, 16, 25, 32, 41]. DeepFace [30] used 3D face frontalization to achieve a breakthrough in face recognition methods that use deep learning. FaceNet [26] proposed triplet loss to maximize the distance between an anchor and its negative sample, and to minimize the distance between the same anchor and its positive sample. CenterLoss [35] proposed center loss to minimize the distance between samples and their class centers. Marginal-Loss [8] adopted the concept of margin to minimize intra-class variations and to keep inter-class distances with margin. RangeLoss [43] used long-tailed data during the train-

ing stage. RingLoss [47] constrained a feature's magnitude to be a certain number. NormFace [33] proposed to normalize features and fully connected layer weights; verification accuracy increased after normalization. SphereFace [19] proposed angular softmax (A-Softmax) loss with multiplicative angular margin. Based on A-Softmax, CosFace [34] proposed an additive cosine margin and ArcFace [7] applies an additive angular margin. The authors of RegularFace [44] and UniformFace [9] argued that approaches that use angular margin [7, 19, 34] concentrated on intra-class compactness only, then suggested new losses to increase the inter-class variation. These previous methods, in general, focused on how to improve loss functions to improve face recognition accuracy with conventional feature representation. A slight change such as adding a few layers or increasing the number of channels, commonly did not bring a noticeable improvement. However, GroupFace improves the quality of feature representation and achieves a significant improvement by adding a few more layers in parallel.

Grouping or clustering methods such as k-means internally categorize samples by considering relative metrics such as a cosine similarity or Euclidean distance without explicit class labels. In general, these clustering methods attempt to construct well-distinguished categories by preventing the assignment of most images to one or a few clusters. Recently, several methods that used deep learning [4, 24, 40] have been introduced. These methods are effective, however, they use full batches as in previous methods, not mini-batches as in deep learning. Thus, these methods are not readily incorporate deeply and end-to-end in an application framework. To efficiently learn the latent groups, our method introduces a self-distributed grouping method that considers an expectation-normalized probability in a deep manner.

3. Proposed Method

Our GroupFace learns the latent groups by using a self-distributed grouping method, constructs multiple group-aware representations and ensembles them into the standard instance-based representation to enrich the feature representation for face recognition.

3.1. GroupFace

We discuss that how the scheme of latent groups are effectively integrated into the embedding feature in GroupFace.

Instance-based Representation. We will call a feature vector in conventional face recognition scheme [7, 34, 35, 44] as an *Instance-based Representation* in this paper (Figure 2). The instance-based representation is commonly trained as an embedding feature by using softmax-based loss (*e.g.*, CosFace [34] and ArcFace [7]) and is used to

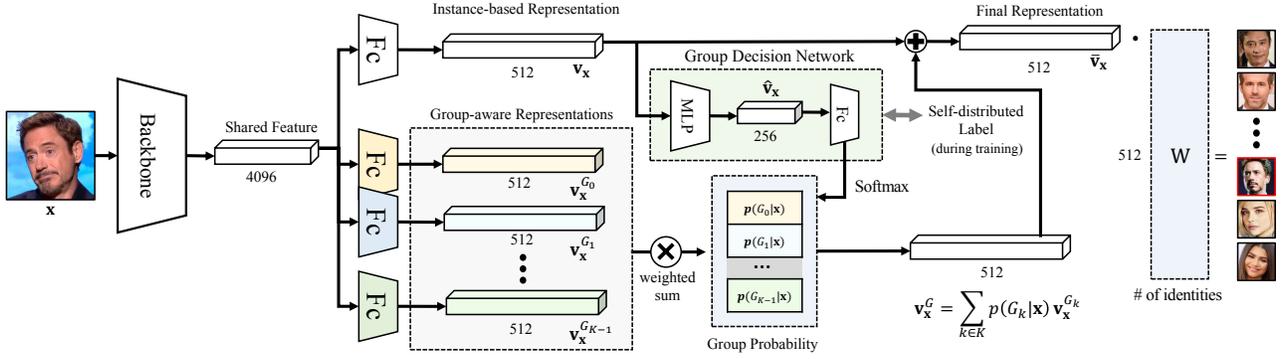


Figure 2. GroupFace generates a shared feature of 4096 dimension and deploys a fully-connected layer for an instance-based representation \mathbf{v}_x and K fully-connected layers for group-aware representations \mathbf{v}_x^G for a given sample \mathbf{x} . A group-decision-network, which is supervised by the self-distributed labeling, outputs a set of group probabilities $\{p(G_0|\mathbf{x}), p(G_1|\mathbf{x}), \dots, p(G_{K-1}|\mathbf{x})\}$ from the instance-based representation. The final representation of 512 dimension is an aggregation of the instance-based representation and the weighted sum \mathbf{v}_x^G of the group-aware representations with the group probabilities. W is a weight of the function g .

predict an identity as:

$$p(\mathbf{y}_i|\mathbf{x}) = \text{softmax}_k(g(\mathbf{v}_x)), \quad (1)$$

where \mathbf{y}_i is an identity label, \mathbf{v}_x is the instance-based representation of a given sample \mathbf{x} , and g is a function which projects an embedding feature of 512 dimension into M dimensional space. M is the number of person identities.

Group-aware Representation. GroupFace uses a novel *Group-aware Representation* as well as the instance-based representation to enrich the embedding features. Each group-aware representation vector is extracted by deploying fully-connected layers for each corresponding group (Figure 2). The embedding feature ($\bar{\mathbf{v}}_x$, *Final Representation* in Figure 2) of GroupFace is obtained by aggregating the instance-based representation \mathbf{v}_x and the weighted-summed group-aware representation \mathbf{v}_x^G . GroupFace predicts an identity by using the enriched final representation $\bar{\mathbf{v}}_{x_i}$ as:

$$\begin{aligned} p(\mathbf{y}_i|\mathbf{x}) &= \text{softmax}_k(g(\bar{\mathbf{v}}_x)), \\ &= \text{softmax}_k(g(\mathbf{v}_x + \mathbf{v}_x^G)), \end{aligned} \quad (2)$$

where \mathbf{v}_x^G is an ensemble of multiple group-aware representations with group probabilities.

Structure. GroupFace calculates and uses instance-based representation and group-aware representations, concurrently. The instance-based representation is obtained by the same procedures that are used in conventional face recognition methods [7, 34, 35, 44], and the K group-aware representations are obtained similarly by deploying a fully-connected layer. Then, the group probabilities are calculated from the instance-based representation vector by deploying a Group Decision Network (GDN) that consists of three fully-connected layers and a softmax layer. Using the

group probabilities, the multiple group-aware representations are sub-ensembled in a soft manner (S-GroupFace) or a hard manner (H-GroupFace).

1. S-GroupFace aggregates multiple group-aware representations with corresponding probabilities of groups as weights, and is defined as:

$$\mathbf{v}_x^G = \sum_{k \in K} p(G_k|\mathbf{x}) \mathbf{v}_x^{G_k}. \quad (3)$$

2. H-GroupFace selects one of the group-aware representations for which the corresponding group probability has the highest value, and is defined as:

$$\mathbf{v}_x^G = \arg \max_{p(G_k|\mathbf{x})} \mathbf{v}_x^{G_k}. \quad (4)$$

S-GroupFace provides a significant improvement of recognition accuracy with a marginal requirement for additional resources, and H-GroupFace is more suitable for practical applications than S-GroupFace, at the cost of a few additional convolutions. The final representation $\bar{\mathbf{v}}_x$ is enriched by aggregating both the instance-based representation and the sub-ensembled group-aware representation.

Group-aware Similarity. We introduce a group-aware similarity that is a new similarity considering both the standard embedding feature and the intermediate feature of GDN in the inference stage. The group-aware similarity is penalized by a distance between intermediate features of two given instances because the intermediate feature is not trained on the cosine space and just describes the group identity of a given sample, not the explicit identity of a given sample. The group-aware similarity S^* between the i^{th} image I_i and the j^{th} image I_j is defined as:

$$S^*(\mathbf{x}_i, \mathbf{x}_j) = S(\bar{\mathbf{v}}_{x_i}, \bar{\mathbf{v}}_{x_j}) - \beta D(\hat{\mathbf{v}}_{x_i}, \hat{\mathbf{v}}_{x_j})^\gamma, \quad (5)$$

where S is a cosine similarity metric, D is a distance metric, $\hat{\mathbf{v}}_{\mathbf{x}}$ denotes the intermediate feature of GDN and, β and γ are a constant parameter. The parameters are determined empirically to be $\beta = 0.1$ and $\gamma = 1/3$.

3.2. Self-distributed Grouping

In this work, we define a group as a set of samples that share any common visual-or-non-visual features that are used for face recognition. Such a group is determined by a deployed GDN. Our GDN is gradually trained in a self-grouping manner that provides a group label by considering the distribution of latent groups without any explicit ground-truth information.

Naïve Labeling. A naïve way to determine a group label is to take an index that has the maximum activation of softmax outputs. We build a GDN f to determine a belonging group G^* for a given sample \mathbf{x} by deploying MLP and attaching a softmax function:

$$p(G_k|\mathbf{x}) = \text{softmax}_k(f(\mathbf{x})), \quad (6)$$

$$G^*(\mathbf{x}) = \arg \max_k p(G_k|\mathbf{x}), \quad (7)$$

where G_k is the k^{th} group. The lack of the consideration for the group distribution can cause the naïve solution to assign most of samples to one or few groups.

Self-distributed Labeling. We introduce an efficient labeling method that utilizes a modified probability regulated by a prior probability to generate uniformly-distributed group labels in a deep manner. We define an expectation-normalized probability \tilde{p} to balance the number of samples among K groups:

$$\tilde{p}(G_k|\mathbf{x}) = \frac{1}{K} \{p(G_k|\mathbf{x}) - E_{\mathbf{x} \sim \text{data}} [p(G_k|\mathbf{x})]\} + \frac{1}{K}, \quad (8)$$

where the first $1/K$ bounds the normalized probability between 0 and 1. Then, the expectation of the expectation-normalized probability is computed as:

$$\begin{aligned} & E_{\mathbf{x} \sim \text{data}} [\tilde{p}(G_k|\mathbf{x})], \\ &= \frac{1}{K} \{E_{\mathbf{x} \sim \text{data}} [p(G_k|\mathbf{x})] - E_{\mathbf{x} \sim \text{data}} [p(G_k|\mathbf{x})]\} + \frac{1}{K}, \\ &= \frac{1}{K}. \end{aligned} \quad (9)$$

The optimal self-distributed label is obtained as:

$$G^*(\mathbf{x}) = \arg \max_k \tilde{p}(G_k|\mathbf{x}). \quad (10)$$

The trained GDN estimates a set of group probabilities that represent the degree to which the sample belongs to the latent groups. As the number of samples approaches infinity, the proposed method stably outputs the uniform-distributed labels (Figure 3).

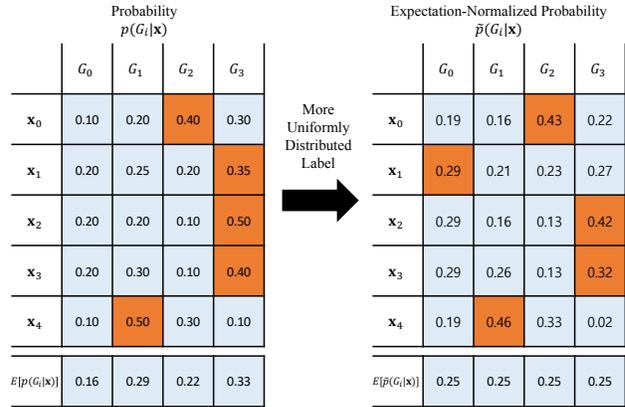


Figure 3. Conceptual scheme of self-distributed labeling. When group probabilities of 5 samples on 4 groups is obtained by GDN, the expectation-normalized probability is computed by subtracting the corresponding expectation of the group probabilities from the group probabilities. The group identifier of the maximum probability is selected as the label of a given sample and the consideration of the expectation makes the distribution of labels more uniform.

3.3. Learning

The network of GroupFace is trained by both the standard classification loss, which is a softmax-based loss to distinguish identities, and the self-grouping loss, which is a softmax loss to train latent groups, simultaneously.

Loss Function. A softmax-based loss L_1 (ArcFace [7] is mainly used in this work) is used to train a feature representation for identities and is defined as:

$$L_1 = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{s(\cos(\theta_{\mathbf{y}_i} + m))}}{e^{s(\cos(\theta_{\mathbf{y}_i} + m))} + \sum_{j=1, j \neq \mathbf{y}_i}^n e^{s(\cos(\theta_j))}}, \quad (11)$$

where N is a number of samples in a mini-batch, θ is the angle between a feature and the corresponding weight, s is a scale factor, m is a marginal factor. To construct the optimal group-space, a self-grouping loss, which reduces the difference between the prediction and the self-generated label, is defined as:

$$L_2 = -\frac{1}{N} \sum_{i=1}^N \text{CrossEntropy}(\text{softmax}(f(\mathbf{x}_i)), G^*(\mathbf{x}_i)). \quad (12)$$

Training. The whole network is trained by using the aggregation of two losses:

$$L = L_1 + \lambda L_2, \quad (13)$$

where the parameter λ balances the weights of different losses and is empirically set to 0.1. Thus, GDN can learn

the group, which is an attribute beneficial to face recognition.

4. Experiments

We describe implementation details and extensively perform experiments and ablation studies to show the effectiveness of GroupFace.

4.1. Implementation Details

Datasets. For the train, we use MSCeleb-1M [10] which has contain about 10M images for 100K identities. Due to the noisy labels of MSCeleb-1M original dataset, we use the refined version [7] which contains 3.8M images for 85k identities. For the test, we conduct our experiments with nine commonly used datasets as follows:

- LFW [13] which contains 13,233 images from 5,749 identities and provides 6000 pairs from them. CALFW [46] and CPLFW [45] are the reorganized datasets from LFW to include higher pose and age variations.
- YTF [37] which consists of 3,425 videos of 1,595 identities.
- MegaFace [16] which composed of more than 1 million images from 690K identities for challenge 1(MF1).
- CFP-FP [27] which contains 500 subjects, each with 10 frontal and 4 profile images.
- AgeDB-30 [22] which contains 12,240 images of 440 identities.
- IJB-B [36] which contains 67,000 face images, 7,000 face videos and 10,000 non-face images.
- IJB-C [21] which contains 138,000 face images, 11,000 face videos and 10,000 non-face images.

Metrics. We compare the verification-accuracy for identity-pairs on LFW [13], YTF [37], CALFW [46], CPLFW [45], CFP-FP [27], AgeDB-30 [22] and MegaFace [16] verification task. MegaFace [16] identification task is evaluated by rank-1 identification accuracy with 1 million distractors. We compare a True Accept Rate at a certain False Accept Rate (TAR@FAR) from 1e-4 to 1e-6 on IJB-B [36] and IJB-C [21].

Experimental Setting. We construct a normalized face image [7, 19, 34] (112×112) by warping a face-region using five facial points from two eyes, nose and two corners of mouth. We employ the ResNet-100 [12] as the backbone network similar to the recent works [7, 14]. We vectorize the activation and reduced # activations to 4096 (shared feature in Figure 2) by a block of BN-FC. Our GroupFace

	TAR	
	FAR=1e-6	FAR=1e-5
Baseline [7]	0.3828	0.8933
- 4 Groups	0.4395	0.8962
- 16 Groups	0.4435	0.8993
- 32 Groups	0.4678	0.9115

(a) Number of Groups

	TAR@FAR=1e-6
Baseline [7]	0.3828
- without Loss	0.4468
- Naive Labeling	0.4535
- Self-distributed Labeling	0.4678

(b) Learning for GDN

	TAR@FAR=1e-6	GFLOPS
Baseline [7]	0.3828	24.2G
- H-GroupFace	0.4439	24.4G
- S-GroupFace	0.4678	24.5G

(c) Hard vs. Soft Ensemble

	TAR	
	FAR=1e-6	FAR=1e-5
Baseline [7]	0.3828	0.8933
- Concat	0.4745	0.8999
- Aggregation	0.4678	0.9115

(d) Aggregation vs. Concatenataion

	TAR@FAR=1e-6
Baseline [7]	0.3828
- GroupFace	0.4678
- GroupFace [†]	0.5212

(e) Group-aware Similarity

	TAR		
	FAR=1e-6	FAR=1e-5	FAR=1e-4
ResNet-100 [7]	0.3828	0.8933	0.9425
ResNet-34 [7]	0.3962	0.8669	0.9308
+ GroupFace	0.4361	0.8820	0.9316
+ GroupFace [†]	0.4823	0.8842	0.9354

(f) Lightweight Model

Table 1. Ablation studies for the proposed GroupFace on IJB-B dataset. The baseline is a recognition-model trained by ArcFace [7] and [†] denotes an evaluation procedure using the group-aware similarity (Eq. 5).

is attached after *res5c* in ResNet-100, where its activation dimension is $512 \times 7 \times 7$. The MLP in GDN consists of two blocks of BN-FC and a FC for group classification. We follow [7, 34] to set the hyper-parameters of the loss function.

Learning. We train the model with 8 synchronized GPUs and a mini-batch involving 128 images per GPU. To stable the group-probability, the network of GroupFace is trained from the pre-trained network that is trained by only the softmax-based loss [7, 34]. We used a learning rate of 0.005 for the first 50k, 0.0005 for the 20k, and 0.00005 for 10k with a weight decay of 0.0005 and a momentum of 0.9 with stochastic gradient descent (SGD). We compute the expectation of group probabilities by computing the group probabilities of 128×8 samples on all GPUs and averaging the expectations over the recent B -batches to accurately estimate the expectation of the group probabilities on the population; B between 32 and 128 empirically shows a similar performance.

4.2. Ablation Studies

To show the effectiveness of the proposed method, we perform the ablation studies on the it’s behaviors. For all experiments, we also use the same network structure with the hyper-parameters mentioned earlier. To clearly show the effect of each ablation study, TAR@FAR of the models are compared on IJB-B dataset [36]; all models in the ablation studies shows around 99.85% on LFW.

Number of Groups. We compare the recognition performance according to the number of groups (Table 1a). As the number of groups grows, the performance increases steadily. In particular, a few initial groups can benefit greatly, and by deploying more groups, significant improvement of performance can be obtained.

Learning for GDN. We compare the learning method for GDN (Table 1b): (1) without loss (adopt the group-aware network structure only), (2) naive labeling, and (3) self-distributed labeling. Just by applying our novel network structure, the recognition performance is greatly improved. In particular, the performance is further increased by adjusting the proposed self-distributed labeling method.

Hard vs. Soft. S-GroupFace shows a high improvement in the performance because it uses all group-aware representations comprehensively with a reasonable additional resource (Table 1c). Since H-GroupFace uses only one strongest group-aware representation even if many groups are deployed, the burden of increasing the number of groups is fixed to a slight amount of additional resource. Thus, H-GroupFace can be applied immediately for high performance gains in practical applications.

Aggregation vs. Concatenation. We compare how to combine the instance-based representation and the group-aware representations into an one embedding feature (Table 1d): (1) aggregation and (2) concatenation. Concatenation-based GroupFace shows a better TAR@FAR= $1e-6$ by 0.67 percentage points than Aggregation-based GroupFace, however, Aggregation-based GroupFace shows a much bet-

Method	#Image	LFW	YTF
DeepID [29]	0.2M	99.47	93.2
DeepFace [30]	4.4M	97.35	91.4
VGGFace [25]	2.6M	98.95	97.3
FaceNet [26]	200M	99.64	95.1
CenterLoss [35]	0.7M	99.28	94.9
RangeLoss [43]	5M	99.52	93.7
MarginalLoss [8]	3.8M	99.48	95.9
SphereFace [19]	0.5M	99.42	95.0
RegularFace [44]	3.1M	99.61	96.7
CosFace [34]	5M	99.81	97.6
UniformFace [9]	6.1M	99.80	97.7
AFRN [14]	3.1M	99.85	97.1
ArcFace [7]	5.8M	99.83	97.7
GroupFace	5.8M	99.85	97.8

Table 2. Verification accuracy (%) on LFW and YTF.

ter TAR@FAR= $1e-5$ by 1.16 percentage points. We chose the Aggregation-based GroupFace that is generally better performing with fewer feature dimensions.

Group-aware Similarity. The recognition-performance is once again improved significantly by evaluating the group-aware similarity (Table 1e). Even though the group-aware similarity increases the feature dimension for calculating a similarity, it is easy to extract the required feature because the feature is the intermediate output of the recognition network. Especially, this experiment shows that the group-based information is distinct from the conventional identity-based information enough to improve performance in practical usages. We show more detailed experiments in Table 5.

Lightweight Model. GroupFace is also effective for a lightweight model such as ResNet-34 [12] that requires only 8.9 GFLOPS less than ResNet-100 [12], which requires 24.2 GFLOPS. ResNet-34 based GroupFace shows a similar performance of ResNet-100 based ArcFace [7] and greatly outperforms ResNet-100 in a most difficult criterion (FAR= $1e-6$). In addition, the group-aware similarity significantly exceed the basic performance of ResNet-34 model (Table 1f).

4.3. Evaluations

LFW, YTF, CALFW, CPLFW, CFP-FP and AgeDB-30. We compare the verification-accuracy on LFW [13] and YTF [37] with the unrestricted with labelled outside data protocol (Table 2). On YTF, we evaluate all the images without the exclusion of noisy images from image sequences. Even though both datasets are highly-saturated, Our GroupFace surpasses the other recent methods. We also report the verification accuracy on the variant of LFW (CALFW [46], CPLFW [45]), CFP-FP [27] and AgeDB-30

Method	CALFW	CPLFW	CFP-FP	AgeDB-30
CenterLoss [35]	85.48	77.48	-	-
SphereFace [19]	90.30	81.40	-	-
VGGFace2 [3]	90.57	84.00	-	-
CosFace [34]	95.76	92.28	98.12	98.11
ArcFace [7]	95.45	92.08	98.27	98.28
GroupFace	96.20	93.17	98.63	98.28

Table 3. Verification accuracy (%) on CALFW, CPLFW, CFP-FP and AgeDB-30.

Method	Protocol	Ident	Verif
RegularFace [44]	Large	75.61	91.13
UniformFace [9]	Large	79.98	95.36
CosFace [34]	Large	80.56	96.56
ArcFace [7]	Large	81.03	96.98
GroupFace	Large	81.31	97.35
SphereFace [19]	Large / R	97.91	97.91
AdaptiveFace [18]	Large / R	95.02	95.61
CosFace [34]	Large / R	97.91	97.91
ArcFace [7]	Large / R	98.35	98.49
GroupFace	Large / R	98.74	98.79

Table 4. Identification and verification evaluation on MegaFace. Ident means rank-1 identification accuracy (%) and Verif means TAR@FAR=1e-6 (%). R denotes the evaluation procedure on the refined version [7] of MegaFace dataset.

[22] (Table 3). Our GroupFace shows the better accuracy on all of the above datasets.

MegaFace. We evaluate our GroupFace under the large-training-set protocol, in which models are trained by using the training set containing more than 0.5M images, on MegaFace [16] (Table 4). GroupFace is the top-ranked face recognition model among the recent published state-of-the-art methods. On the refined MegaFace [7], our GroupFace also outperforms the other models.

IJB-B and IJB-C. We compare the proposed method with other methods on IJB-B [36] and IJB-C [21] datasets (Table 5). Recent angular-margin-softmax based methods [7, 34] show great performance in the datasets. We reports the improvement of GroupFace in the verification accuracy based on both CosFace [34] and ArcFace [7] without any test-time augmentations such as horizontal flipping. Our GroupFace shows significant improvements on all FAR criteria by 8.5 percentage points on FAR=1e-6, 1.8 percentage points on FAR=1e-5 and 0.2 percentage points on FAR=1e-4 than the ArcFace [7] on IJB-B and by 4.3 percentage points on FAR=1e-6, 1.2 percentage points on FAR=1e-5 and 0.4 percentage points on FAR=1e-4 than the ArcFace [7] on IJB-C. The recognition-performance is once again improved significantly by applying the group-aware similarity (Eq. 5), especially on the most difficult criterion (TAR@FAR=1e-

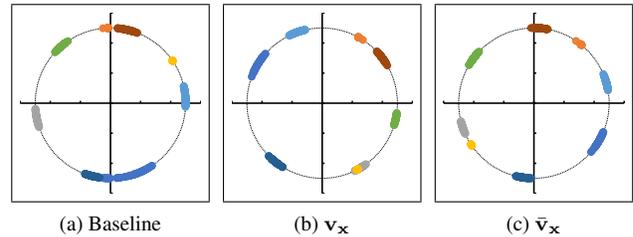


Figure 4. A quantitative comparison of the baseline network (ArcFace [7]), v_x that denotes an instance-based representation of GroupFace and \bar{v}_x that denotes a final representation of GroupFace for the first eight identities in the refined MSCeleb-1M [10] dataset on a 2D space. The eight colored-circles represents the eight identities.

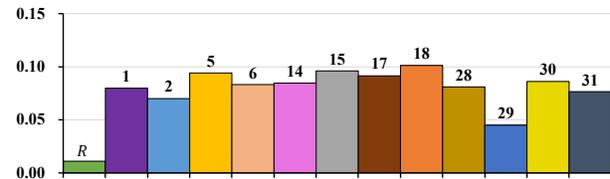


Figure 5. Average probabilities of 32 groups on the refined MSCeleb-1M [10] dataset. 12 groups can mainly activate and the rest groups, R , softly affect.

6) on IJB-B by 5.3 percentage points.

4.4. Visualization

To show the effectiveness of the proposed method, we visualize the feature representation, the average activation of groups and the visual interpretation of groups.

2D Projection of Representation. Figure 4 shows a quantitative comparison among (a) the final representation of the baseline network (ArcFace [7]), (b) the instance-based representation of GroupFace and the final representation of GroupFace on a 2D space. We select the first eight identities in the refined MSCeleb-1M dataset [10] and map the extracted features onto the angular space by using t-SNE [20]. The quantitative comparison shows that the proposed model generates more distinctive feature representations rather than the baseline model and also the proposed model enhances the instance-based representation.

Activation Distribution of Groups. The proposed Self-Grouping tries to make the samples evenly spread throughout the all groups, and at the same time, the softmax-based loss also simultaneously propagates gradients into GDN so that the identification works best. Thus, the probability distribution is not exactly uniform (Figure 5). Some probabilities of the groups are low and the others are high (e.g., 1, 2, 5, 6, 14, 15, 17, 18, 28, 29, 30, 31th groups). The overall

Method	TAR on IJB-B			TAR on IJB-C		
	FAR=1e-6	FAR=1e-5	FAR=1e-4	FAR=1e-6	FAR=1e-5	FAR=1e-4
VGGFace2 [3]	-	0.671	0.800	-	0.747	0.840
CenterFace [35]	-	-	-	-	0.781	0.853
ComparatorNet [39]	-	-	0.849	-	-	0.885
PRN [15]	-	0.721	0.845	-	-	-
AFRN [14]	-	0.771	0.885	-	0.884	0.931
CosFace [34]	0.3649	0.8811	0.9480	0.8801	0.9370	0.9615
ArcFace [7]	0.3828	0.8933	0.9425	0.8625	0.9315	0.9565
GroupFace ⁻	0.4166	0.8983	0.9453	0.8858	0.9399	0.9606
GroupFace	0.4678	0.9115	0.9445	0.9053	0.9437	0.9602
GroupFace [†]	0.5212	0.9124	0.9493	0.8928	0.9453	0.9626

Table 5. Verification evaluation according to different FARs on IJB-B and IJB-C. GroupFace is trained by ArcFace [7]. ⁻ denotes that a model is trained by CosFace [34] and [†] denotes that a model is evaluated by using the group-aware similarity.

distribution is not uniform as we expected, but we see that there is no dominant one among the high activated group.

Interpretation of Groups. The trained latent groups are not always visually distinguishable because they are categorized by a non-linear function of GDN using a latent feature, not a facial attribute (*e.g.*, hair, glasses, and mustache). However, there are two cases of groups (Group 5 and 20 in Figure 6) that we can clearly see their visual properties; 95 of randomly-selected 100 images are men in Group 5 and 94 of randomly-selected 100 images are bald men in Group 20. Others are not described as an one visual property, however, they seems to be described as multiple visual properties such as smile women, right-profile people and scared people in Group 1.

5. Conclusion

We introduce a new face-recognition-specialized architecture that consists of a group-aware network structure and a self-distributed grouping method to effectively manipulate multiple latent group-aware representations. By extensively conducting the ablation studies and experiments, we prove the effectiveness of our GroupFace. The visualization also shows that GroupFace fundamentally enhances the feature representations rather than the existing methods and the latent groups have some meaningful visual descriptions. Our GroupFace provides a significant improvement in the recognition-performance and is practically applicable to existing recognition systems. The rationale behind the effectiveness of GroupFace is summarized in two main ways: (1) It is well known that additional supervisions from different objectives can bring an improvement of the given task by sharing a network for feature extraction, *e.g.*, a segmentation head can improve accuracy in object detection [2, 11]. Likewise, learning the groups can be a helpful cue to train a more generalized feature extractor for face recognition. (2) GroupFace proposes a novel structure that fuses instance-

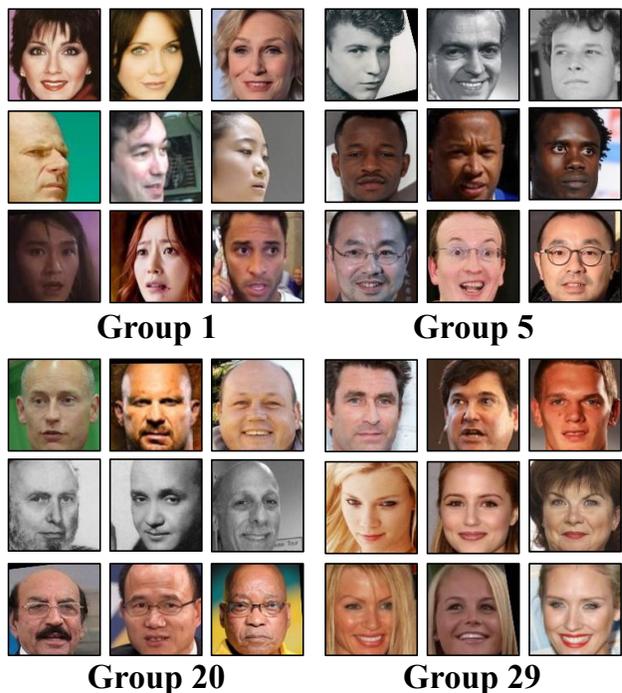


Figure 6. Example images belonging to each groups. As enormous identities (80k~) of large scale dataset cannot be mapped to a few groups (32), each group contains identities of multiple characteristics. Some groups have one common visual description (Group 5: Some Men, Group 20: Bald Men) with some variations while others have multi-mode visual descriptions.

based representation and group-based representation, which is empirically proved its effectiveness.

Acknowledgement

We thank AI team of Kakao Enterprise, especially Won-jae Kim and Yoonho Lee for their helpful feedback.

References

- [1] Timo Ahonen, Abdenour Hadid, and Matti Pietikäinen. Face recognition with local binary patterns. In Tomás Pajdla and Jiří Matas, editors, *European Conference on Computer Vision Workshops*, 2004. 2
- [2] Sean Bell, C Lawrence Zitnick, Kavita Bala, and Ross Girshick. Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 8
- [3] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *International Conference on Automatic Face and Gesture Recognition*, 2018. 1, 2, 7, 8
- [4] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *European Conference on Computer Vision*, 2018. 2
- [5] D. Chen, X. Cao, F. Wen, and J. Sun. Blessing of dimensionality: High-dimensional feature and its efficient compression for face verification. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2013. 2
- [6] D. Chen, X. Cao, D. Wipf, F. Wen, and J. Sun. An efficient joint formulation for bayesian face verification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. 2
- [7] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 1, 2, 3, 4, 5, 6, 7, 8
- [8] Jiankang Deng, Yuxiang Zhou, and Stefanos Zafeiriou. Marginal loss for deep face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017. 2, 6
- [9] Yueqi Duan, Jiwen Lu, and Jie Zhou. Uniformface: Learning deep equidistributed representation for face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 1, 2, 6, 7
- [10] Yandong Guo, Lei Zhang, Yuxiao Hu, Xiaodong He, and Jianfeng Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *European Conference on Computer Vision*, 2016. 1, 2, 5, 7
- [11] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *IEEE International Conference on Computer Vision*, 2017. 8
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 1, 5, 6
- [13] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical report, University of Massachusetts, Amherst, 2007. 5, 6
- [14] Bong-Nam Kang, Yonghyun Kim, Bongjin Jun, and Daijin Kim. Attentional feature-pair relation networks for accurate face recognition. In *IEEE International Conference on Computer Vision*, 2019. 1, 5, 6, 8
- [15] Bong-Nam Kang, Yonghyun Kim, and Daijin Kim. Pairwise relational networks for face recognition. In *European Conference on Computer Vision*, 2018. 8
- [16] Ira Kemelmacher-Shlizerman, Steven M Seitz, Daniel Miller, and Evan Brossard. The megaface benchmark: 1 million faces for recognition at scale. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 1, 2, 5, 7
- [17] N. Kumar, A. C. Berg, P. N. Belhumeur, and S. K. Nayar. Attribute and simile classifiers for face verification. In *IEEE International Conference on Computer Vision*, 2009. 2
- [18] Hao Liu, Xiangyu Zhu, Zhen Lei, and Stan Z Li. Adaptive-face: Adaptive margin and sampling for face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 7
- [19] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song. Sphereface: Deep hypersphere embedding for face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 1, 2, 5, 6, 7
- [20] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 2008. 7
- [21] B. Maze, J. Adams, J. A. Duncan, N. Kalka, T. Miller, C. Otto, A. K. Jain, W. T. Niggel, J. Anderson, J. Cheney, and P. Grother. Iarpa janus benchmark - c: Face dataset and protocol. In *International Conference on Biometrics*, 2018. 5, 7
- [22] Stylianos Moschoglou, Athanasios Papaioannou, Christos Sagonas, Jiankang Deng, Irene Kotsia, and Stefanos Zafeiriou. Agedb: the first manually collected, in-the-wild age database. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017. 5, 7
- [23] Hieu V. Nguyen and Li Bai. Cosine similarity metric learning for face verification. In Ron Kimmel, Reinhard Klette, and Akihiro Sugimoto, editors, *Asian Conference on Computer Vision*, 2011. 2
- [24] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, 2016. 2
- [25] Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. In *British Machine Vision Conference*, 2015. 1, 2, 6
- [26] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2, 6
- [27] S. Sengupta, J. Chen, C. Castillo, V. M. Patel, R. Chellappa, and D. W. Jacobs. Frontal to profile face verification in the wild. In *IEEE Winter Conference on Applications of Computer Vision*, 2016. 5, 6
- [28] Karen Simonyan, Omkar Parkhi, Andrea Vedaldi, and Andrew Zisserman. Fisher vector faces in the wild. In *British Machine Vision Conference*, 2013. 2
- [29] Yi Sun, Yuheng Chen, Xiaogang Wang, and Xiaoou Tang. Deep learning face representation by joint identification-

- verification. In *International Conference on Neural Information Processing Systems*, 2014. 6
- [30] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2014. 2, 6
- [31] M. A. Turk and A. P. Pentland. Face recognition using eigenfaces. In *IEEE Conference on Computer Vision and Pattern Recognition*, 1991. 2
- [32] Fei Wang, Liren Chen, Cheng Li, Shiyao Huang, Yanjie Chen, Chen Qian, and Chen Change Loy. The devil of face recognition is in the noise. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *European Conference on Computer Vision*, 2018. 1, 2
- [33] Feng Wang, Xiang Xiang, Jian Cheng, and Alan Loddon Yuille. Normface: L2 hypersphere embedding for face verification. In *ACM International Conference on Multimedia*, 2017. 2
- [34] H. Wang, Y. Wang, Z. Zhou, X. Ji, D. Gong, J. Zhou, Z. Li, and W. Liu. Cosface: Large margin cosine loss for deep face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 1, 2, 3, 5, 6, 7, 8
- [35] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *European Conference on Computer Vision*, 2016. 1, 2, 3, 6, 7, 8
- [36] C. Whitelam, E. Taborsky, A. Blanton, B. Maze, J. Adams, T. Miller, N. Kalka, A. K. Jain, J. A. Duncan, K. Allen, J. Cheney, and P. Grother. Iarpa janus benchmark-b face dataset. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017. 5, 6, 7
- [37] L. Wolf, T. Hassner, and I. Maoz. Face recognition in unconstrained videos with matched background similarity. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2011. 5, 6
- [38] L. Wolf, T. Hassner, and Y. Taigman. Descriptor based methods in the wild. In *European Conference on Computer Vision Workshops*, 2008. 2
- [39] Weidi Xie, Li Shen, and Andrew Zisserman. Pairwise relational networks for face recognition. In *European Conference on Computer Vision*, 2018. 8
- [40] Jianwei Yang, Devi Parikh, and Dhruv Batra. Joint unsupervised learning of deep representations and image clusters. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [41] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z. Li. Learning face representation from scratch. *CoRR*, 2014. 1, 2
- [42] Qi Yin, Xiaoou Tang, and Jian Sun. An associate-predict model for face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2011. 2
- [43] X. Zhang, Z. Fang, Y. Wen, Z. Li, and Y. Qiao. Range loss for deep face recognition with long-tailed training data. In *IEEE International Conference on Computer Vision*, 2017. 2, 6
- [44] Kai Zhao, Jingyi Xu, and Ming-Ming Cheng. Regularface: Deep face recognition via exclusive regularization. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 1, 2, 3, 6, 7
- [45] T. Zheng and W. Deng. Cross-pose lfw: A database for studying cross-pose face recognition in unconstrained environments. Technical Report 18-01, Beijing University of Posts and Telecommunications, February 2018. 5, 6
- [46] Tianyue Zheng, Weihong Deng, and Jiani Hu. Cross-age LFW: A database for studying cross-age face recognition in unconstrained environments. *CoRR*, abs/1708.08197, 2017. 5, 6
- [47] Y. Zheng, D. K. Pal, and M. Savvides. Ring loss: Convex feature normalization for face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2