

Rethinking Robust Representation Learning Under Fine-grained Noisy Faces

Bingqi Ma^{1*}, Guanglu Song^{1*}, Boxiao Liu^{1,2}, and Yu Liu^{1†}

¹ Sensetime Research

² SKLP, Institute of Computing Technology, CAS

{mabingqi, songguanglu}@sensetime.com, liuboxiao@ict.ac.cn,
liuyuisanai@gmail.com

Abstract. Learning robust feature representation from large-scale noisy faces stands out as one of the key challenges in high-performance face recognition. Recent attempts have been made to cope with this challenge by alleviating the intra-class conflict and inter-class conflict. However, the unconstrained noise type in each conflict still makes it difficult for these algorithms to perform well. To better understand this, we reformulate the noise type of each class in a more fine-grained manner as **N-identities|K^C-clusters**. Different types of noisy faces can be generated by adjusting the values of N , K , and C . Based on this unified formulation, we found that the main barrier behind the noise-robust representation learning is the flexibility of the algorithm under different N , K , and C . For this potential problem, we propose a new method, named Evolving Sub-centers Learning (ESL), to find optimal hyperplanes to accurately describe the latent space of massive noisy faces. More specifically, we initialize M sub-centers for each class and ESL encourages it to be automatically aligned to **N-identities|K^C-clusters** faces via producing, merging, and dropping operations. Images belonging to the same identity in noisy faces can effectively converge to the same sub-center and samples with different identities will be pushed away. We inspect its effectiveness with an elaborate ablation study on the synthetic noisy dataset with different N , K , and C . Without any bells and whistles, ESL can achieve significant performance gains over state-of-the-art methods on large-scale noisy faces.

Keywords: Fine-grained Noisy Faces, Evolving Sub-centers Learning

1 Introduction

Owing to the rapid development of computer vision technology [24, 23, 22], face recognition [26, 4, 16, 30, 15] has made a remarkable improvement and has been widely applied in the industrial environment. Much of this progress was sparked by the collection of large-scale web faces as well as the robust learning strategies [4, 26] for representation learning. For instance, MS-Celeb-1M (MS1MV0) [8]

*Equal contributions.

†Corresponding author.

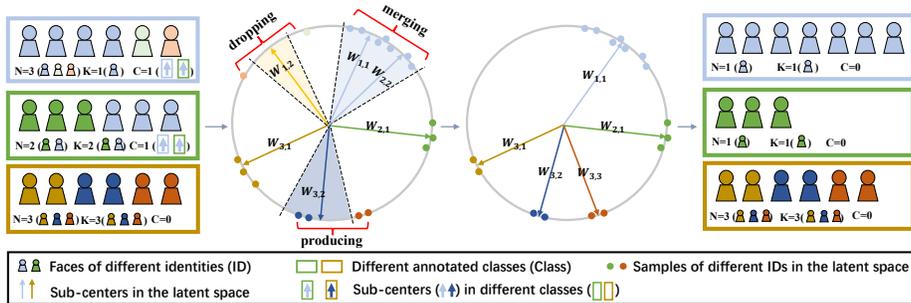


Fig. 1. Illustration of fine-grained noisy faces and ESL (Best viewed in color). For **ID2** and **ID3** in **Class1**, there is only one image for each ID, and they will be removed by the dropping operation. **ID1** appears in both **Class1** and **Class2**, so that the merging operation will merge images in **Class2** with **ID1** into **Class1**. In **Class3**, there are 3 IDs but only 2 sub-centers, so the producing operation will produce another valid sub-center. Our proposed ESL can be flexibly adapted to different combinations of NKC , and is more robust to unconstrained real-world noise.

provides more than 10 million face images with rough annotations. The growing scale of training datasets inevitably introduces unconstrained noisy faces and can easily weaken the performance of state-of-the-art methods. Learning robust feature representation from large-scale noisy faces has become an important challenge for high-performance face recognition. Conventional noisy data learning, such as recursive clustering, cleaning, and training process, suffers from high computational complexity and cumulative error. For this problem, Sub-center ArcFace [2] and SKH [13] are proposed to tackle the intra-class conflict or inter-class conflict by designing multiple sub-centers for each class. These algorithms demonstrate remarkable performance in the specific manual noise. However, they are still susceptible to the unconstrained types of real-world noisy faces. Naturally, we found that it is far from enough to just divide label noise in face recognition roughly into intra-class noise and inter-class noise. It greatly limits our understanding of the variant noise types and the exploration of noise-robust representation learning strategies.

To better understand this, we reformulate the noise data in a more fine-grained manner as N -identities $|K^C$ -clusters faces for each class. Faces sharing *identity* (ID) means these images come from the same person. Faces annotated with the same label construct a *class*, and there may be annotation errors in the class. If there are no less than two faces for an identity, these images build a meaningful *cluster* [7]. Please refer to the Sec. 1 in the appendix for the holistic description of terms and notations. Taking the Class1 in Fig. 1 as an example, there are 3 IDs marked with ID1, ID2 and ID3, so the N in Class1 is 3. However, only ID1 contains more than 2 images, so the K in Class1 is 1. Furthermore, ID1 appears in both Class1 and Class2, which indicates one inter-class conflict, so the C in Class1 is 1. As shown in Table 1, our proposed N -identities $|K^C$ -clusters formulation can clearly represent different fine-grained noisy data.

Table 1. Noise type in different combination of N , K , and C . Δ represents intra-class conflict in which there are multiple clusters in the class. \square represents intra-class conflict where there are outlier faces in the class. \diamond represents the inter-class conflict in which there are multiple clusters with the same identify in different classes.

	$N = K = 1$	$N = K > 1$	$N > K > 1$	$N > K = 1$	$N > K = 0$
$C = 0$	-	Δ	$\Delta \square$	\square	\square
$C > 0$	\diamond	$\Delta \diamond$	$\Delta \square \diamond$	$\square \diamond$	-

However, if N and K are larger than the predefined sub-center number in Sub-center ArcFace [2] and SKH [13], images without corresponding sub-center will lead to intra-class conflict. If C exceeds the sub-center number in SKH [13], extra conflicted clusters will bring inter-class conflict. Both intra-class conflict and inter-class conflict will lead to the wrong gradient, which would dramatically impair the representation learning process.

In this paper, we constructively propose a flexible method, named Evolving Sub-centers Learning (ESL), to solve this problem caused by unconstrained N , K , and C . More specifically, we initialize M sub-centers for each class first. Images belonging to the same identity will be pushed close to the corresponding positive sub-center and away from all other negative sub-centers. Owing to elaborate designed producing, dropping, and merging operations, ESL encourages the number of sub-centers to be automatically aligned to `N-identities|KC-clusters` faces. As shown in Fig. 1, our proposed ESL can be flexibly adapted to different combinations of NKC , and is more robust to unconstrained real-world noise. We inspect its effectiveness with elaborate ablation study on variant `N-identities|KC-clusters` faces. Without any bells and whistles, ESL can achieve significant performance gains over current state-of-the-art methods in large-scale noisy faces. To sum up, the key contributions of this paper are as follows:

- We reformulate the noise type of faces in each class into a more fine-grained manner as `N-identities|KC-clusters`. Based on this, we reveal that the key to robust representation learning strategies under real-world noise is the flexibility of the algorithm to the variation of N , K , and C .
- We introduce a general flexible method, named Evolving Sub-centers Learning (ESL), to improve the robustness of feature representation on noisy training data. The proposed ESL enjoys scalability to different combinations of N , K , and C , which is more robust to unconstrained real-world noise.
- Without relying on any annotation post-processing or iterative training, ESL can easily achieve significant performance gains over state-of-the-art methods on large-scale noisy faces.

2 Related work

Loss function for Face Recognition. Deep face recognition models rely heavily on the loss function to learn discriminate feature representation. Previous

works [4, 25, 26, 12, 20, 14, 28, 6, 31] usually leverage the margin penalty to optimize the intra-class distance and the inter-class distance. Facenet [20] uses the Triplet to force that faces in different classes have a large Euclidean distance than faces in the same class. However, the Triplet loss can only optimize a subset of all classes in each iteration, which would lead to an under-fitting phenomenon. It is still a challenging task to enumerate the positive pairs and negative pairs with a growing number of training data. Compared with the sample-to-sample optimization strategy, Liu et al. [14] proposes the angular softmax loss which enables convolutional neural networks to learn angularly discriminative features. Wang et al. [26] reformulates Softmax-base loss into a cosine loss and introduces a cosine margin term to further maximize the decision margin in the angular space. Deng et al. [4] directly introduces a fixed margin, maintaining the consistency of the margin in the angular space. Liu et al. [16] adopts the hard example mining strategy to re-weight temperature in the Softmax-base loss function for more effective representation learning.

Dataset for Face Recognition. Large-scale training data can significantly improve the performance of face recognition models. MS1MV0 [8], in which there are about 100K identities and 1M faces, is the most commonly used face recognition dataset. MS1MV3 is a cleaned version from MS1MV0 with a semi-automatic approach [5]. An et al. [1] cleans and merges existing public face recognition datasets, then obtains Glint360K with 17M faces and 360K IDs. Recently, Zhu et al. [35] proposes a large-scale face recognition dataset WebFace260M and a automatically cleaning pipeline. By iterative training and cleaning, they proposed well-cleaned subset with 42M images and 2M IDs.

Face Recognition under Noisy Data. Iterative training and cleaning is an effective data cleanup method. However, it is extremely inefficient as the face number increases. Recent works [27, 10, 32, 4, 13, 33, 34] focus on efficient noisy data cleanup methods. Zhong et al. [32] decouples head data and tail data of a long-tail distribution and designs a noise-robust loss function to learn the sample-to-center and sample-to-sample feature representation. Deng et al. [2] designs multiple centers for each class, splitting clean faces and noisy faces into different centers to deal with the inter-class noise. Liu et al. [13] leverages multiple hyper-planes with a greedy switching mechanism to alleviate both inter-class noise and intra-class noise. However, these methods are sensitive to hyper-parameter and can not tackle the complex noisy data distribution.

3 The Proposed Approach

In this section, we are committed to eliminating the unconstrained real-world noise via a flexible and scalable learning manner, named Evolving Sub-centers Learning, that can be easily plugged into any loss functions. The pipeline of ESL is as shown in Fig. 2. We will first introduce our proposed ESL and then give a deep analysis to better understand its effectiveness and flexibility under fine-grained noisy faces. Finally, we conduct a detailed comparison between ESL and the current state-of-the-art noise-robust learning strategies.

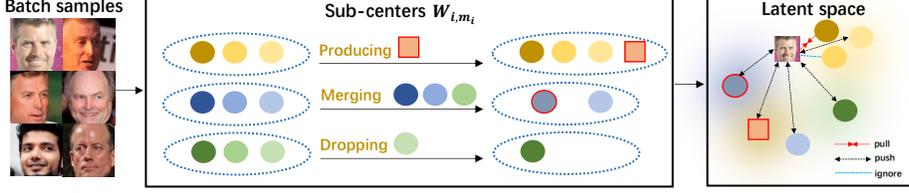


Fig. 2. The pipeline of the Evolving Sub-centers Learning. We initialize M sub-centers for each class and they will evolve adaptively to align the data distribution. It pushes the images belonging to an identity close to the specific sub-center and away from all other negative sub-centers. The sub-center with a confusing similarity to the current sample will be ignored in the latent space. This can effectively dispose of the label conflict caused by fine-grained noisy faces.

3.1 Evolving Sub-centers Learning

In face recognition tasks, the unified loss function can be formulated as:

$$\mathcal{L}(x_i) = -\log \frac{e^{\hat{f}_{i,y_i}}}{e^{\hat{f}_{i,y_i}} + \sum_{j=1, j \neq y_i}^S e^{f_{i,j}}}, \quad (1)$$

where i is the index of face images, y_i represents the label ID of image I_i and S indicates the total class number in the training data. Let the x_i and W_j denote the feature representation of face image I_i and the j -th class center, the logits \hat{f}_{i,y_i} and $f_{i,j}$ can be computed by:

$$\hat{f}_{i,y_i} = s \cdot [m_1 \cdot \cos(\theta_{i,y_i} + m_2) - m_3], \quad (2)$$

$$f_{i,j} = s \cdot \cos(\theta_{i,j}), \quad (3)$$

where s is the re-scale parameter and $\theta_{i,j}$ is the angle between the x_i and W_j normalized with \mathcal{L}_2 manner. For ArcFace with $m_1 = 1$ and $m_3 = 0$, we can compute the θ_{i,y_i} by:

$$\theta_{i,y_i} = \arccos\left(\frac{W_{y_i}^T x_i}{\|W_{y_i}^T\|_2 \|x_i\|_2}\right). \quad (4)$$

As shown in SKH [13], Eq. (1) will easily get wrong loss under $N > 1$ which indicates at least two different identities exist in the images currently labeled as the same identity. In this paper, we address this problem by proposing the idea of using class-specific sub-centers for each class, which can be directly adopted by any loss functions and will significantly increase its robustness. As illustrated in Fig 2, we init M_j sub-centers for j -th class where each center is dominated by a learnable vector W_{j,m_j} , $m_j \in [1, M_j]$. The original class weight $W_j \in \mathbb{R}^{1 \times D}$ can be replaced by all sub-centers $W_j \in \mathbb{R}^{1 \times M_j \times D}$. Based on this, Eq. (1) can be re-written as:

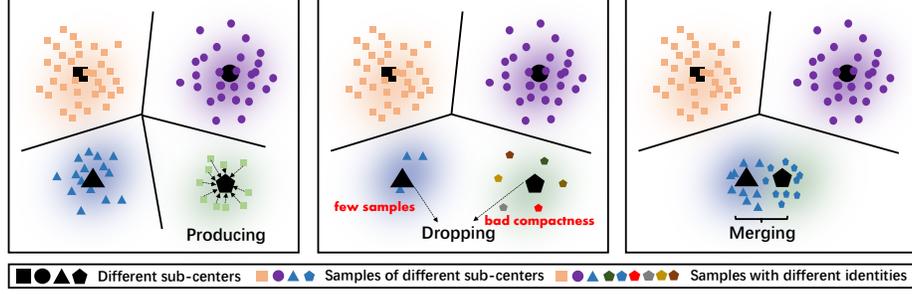


Fig. 3. Illustration of the sub-centers producing, dropping and merging. The instances with black color indicate the sub-centers. The instance belonging to each sub-center is represented by the same shape and different colors mean different identities.

$$\mathcal{L}(x_i) = -\log \frac{e^{\hat{f}_{i,y_i,m}}}{e^{\hat{f}_{i,y_i,m}} + \sum_{\substack{j \in [1,C], m_j \in [1,M_j] \\ (j,m_j) \neq (y_i,m)}} (1 - \mathbb{1}\{\cos(\theta_{i,j,m_j}) > \mathcal{D}_{j,m_j}\}) e^{f_{i,j,m_j}}}, \quad (5)$$

where indicator function $\mathbb{1}\{\cos(\theta_{i,j,m_j}) > \mathcal{D}_{j,m_j}\}$ returns 1 when $\cos(\theta_{i,j,m_j}) > \mathcal{D}_{j,m_j}$ and 0 otherwise. We calculate the mean μ_{j,m_j} and standard deviation σ_{j,m_j} of the cosine similarity between the sub-center W_{j,m_j} and the samples belonging to it. \mathcal{D}_{j,m_j} can be generated by:

$$\mathcal{D}_{j,m_j} = \mu_{j,m_j} + \lambda_1 \sigma_{j,m_j} \quad (6)$$

$\hat{f}_{i,y_i,m}$ and f_{i,j,m_j} are computed by:

$$\hat{f}_{i,y_i,m} = s \cdot [m_1 \cdot \cos(\theta_{i,y_i,m} + m_2) - m_3], \quad (7)$$

$$f_{i,j,m_j} = s \cdot \cos(\theta_{i,j,m_j}), \quad (8)$$

where θ_{i,j,m_j} is the angle between the feature representation x_i of the i -th face image and the m_j -th sub-center W_{j,m_j} in j -th class. We determine m for sample I_i by the nearest distance priority manner as:

$$m = \arg \max_{m_{y_i}} \cos(\theta_{i,y_i,m_{y_i}}). \quad m_{y_i} \in [1, M_{y_i}] \quad (9)$$

Given an initial M_j for each class, Eq. (5) can capture the unconstrained distribution of the whole training data with potential label noise. It pushes the images belonging to the same identity close to a specific sub-center and away from all other negative sub-centers. Meanwhile, the sub-center with a confusing similarity to the current sample will be ignored to dispose of the label conflict. To make it more flexible with unconstrained changes in N , K , and C , we further

introduce the producing, merging, and dropping operations as shown in Fig. 3. **Sub-centers Producing.** Based on the aforementioned design, it can effectively alleviate the conflict caused by label noise when $M_j > N > 1$. However, the unconstrained N makes it difficult to select the appropriate M_j for each class. To make it more flexible, we introduce the sub-centers producing operation to automatically align the sub-centers and the actual identity number in each class. Given \mathcal{N} images with label y and assigned to sub-center m with Eq. 9, a new sub-center W_{y, M_y+m} can be generated by:

$$W_{y, M_y+m} = \frac{1}{\mathcal{T}} \sum_{i=1}^{\mathcal{N}} \mathbb{1}\{\cos(\theta_{i,y,m}) < \mu_{y,m} - \lambda_2 \sigma_{y,m}\} x_i, \quad \text{if } \mathcal{T} > 0, \quad (10)$$

where $\mathcal{T} = \sum_{i=1}^{\mathcal{N}} \mathbb{1}\{\cos(\theta_{i,y,m}) < \mu_{y,m} - \lambda_2 \sigma_{y,m}\}$. If $\mathcal{T} = 0$, there is no new sub-center to be formed. After this, we can progressively produce new sub-center to house additional identities beyond M_j . It effectively improves intra-class compactness and reduces the conflict caused by unconstrained N and K .

Sub-centers Dropping. As demonstrated by [2, 13], many state-of-the-art methods are susceptible to the outlier faces (the image number belonging to an identity is less than 2, $N > K \geq 1$). These outlier images are hard to be pushed close to any corresponding positive sub-center. During the producing process, outlier images from each sub-center will generate a new sub-center. The dropping operation should remove the sub-center from outlier images but preserve the sub-center with a valid identity. Considering the standard deviation can not reflect the density of a distribution, we just leverage μ_{i, m_i} as the metric. The condition of dropping can be formulated as :

$$\mathcal{J}(W_{i, m_i}) = \mathbb{1}\{\mu_{i, m_i} \leq \lambda_3\}. \quad (11)$$

If the μ_{i, m_i} is less than λ_3 , we will ignore these images during the training process and then erase the specific sub-center.

Sub-centers Merging. Using sub-centers for each class can dramatically improve the robustness under noise. However, the inter-class discrepancy will be inevitably affected by inter-class conflict caused by the shared identity between different sub-centers. SKH [13] sets the same fixed number of sub-centers for each class, which can not handle complex inter-class conflict with unconstrained C . Meanwhile, the sub-center strategy undermines the intra-class compactness as the samples in a clean class also converge to different sub-centers. To deal with this potential problem, we employ the sub-centers merging operation to aggregate different $W_{*,*}$. The condition of merging can be formulated as:

$$\mathcal{J}(W_{i, m_i}, W_{j, m_j}) = \mathbb{1}\{W_{i, m_i}^T W_{j, m_j} \geq \max(\mu_{j, m_j} + \lambda_4 \sigma_{j, m_j}, \mu_{i, m_i} + \lambda_4 \sigma_{i, m_i})\}, \quad (12)$$

where W_{i, m_i} and W_{j, m_j} are normalized with L_2 manner. According to Eq. (12), we merge multiple sub-centers satisfying $\mathcal{J}(*, *) = 1$ into a group and combine

them into a single sub-center as following:

$$W_{*,new} = \frac{1}{|G|} \sum_{(p,m_p) \in G} W_{p,m_p}, \quad (13)$$

where G and $|G|$ indicate the merged group and its sub-center number. Furthermore, images belonging to G will be assigned to a new label (we directly select the minimum label ID in the G as the target).

3.2 Progressive Training Framework

At the training stage, we perform the sub-centers producing, dropping and merging operations progressively to effectively alleviate the label conflict cause by unconstrained N , K , and C . The training framework is summarized in Alg. 1.

Algorithm 1 Evolving Sub-centers Learning

Input: Training data set \mathcal{X} , label set \mathcal{Y} , total training epoch E , start epoch ε for ESL.

Initialize: Label number C , sub-centers number M_* and $W_{*,*}$ for each class.

$e \leftarrow 0$;

while $e < E$ **do**

 sample data X, Y from \mathcal{X}, \mathcal{Y} ;

 compute loss function $L(X, Y)$ by Eq. (5) and update model;

 generating μ_* and σ_* for each sub-center;

if $e > \varepsilon$ **then**

for $i = 1$ **to** C **do**

for $j = 1$ **to** M_i **do**

 // Producing

 computing \mathcal{T} via Eq. (10);

if $\mathcal{T} > 0$ **then**

 └ computing W_{i,M_i+j} via Eq. (10);

 // Dropping

 generating $\mathcal{J}(W_{i,j})$ via Eq. (11);

if $\mathcal{J}(W_{i,j})$ **then**

 dropping sub-center $W_{i,j}$;

 generating \mathcal{X}_i as images with label i ;

foreach *image* X *in* \mathcal{X}_i **do**

 └ computing m via Eq. (9);

if $j = m$ **then**

 └ dropping image X ;

 // Merging

 generating vertex set V with each sub-center W in $W_{*,*}$;

 generating edge set E with (W_i, W_j) if $\mathcal{J}(W_i, W_j) = 1$ via Eq. (12);

 generating graph $G = (V, E)$;

foreach *connected component* g *in* G **do**

 └ generating new sub-center via Eq. (13);

 └ relabeling images belonging to sub-centers in g ;

$e = e + 1$;

In this manner, ESL is able to capture the complex distribution of the whole training data with unconstrained label noise. It tends to automatically adjust the sub-centers to align the distribution of N , K , and C in the given datasets. This allows it to be flexible in solving the real-world noise while preventing the network from damaging the inter-class discrepancy on clean faces.

3.3 Robustness Analysis on Fine-grained Noisy Faces

When applied to the practical N -identities/ K -clusters faces, the key challenge is to process different combinations of N , K , and C . In Tab. 1, we have analyzed the noise type in different combinations of N , K , and C . Now we investigate the robustness of ESL on fine-grained noisy faces.

To simplify the analysis, we first only consider the circumstance when $\mathbf{C} = \mathbf{0}$. (1) $\mathbf{N} = \mathbf{K} = \mathbf{1}, \mathbf{C} = \mathbf{0}$. This phenomenon indicates the training dataset is absolutely clean. In this manner, most of the feature learning strategies can perform excellent accuracy. However, introducing the sub-centers for each class will damage the intra-class compactness and degrade the performance. The sub-centers merging allows ESL to progressively aggregate the sub-centers via Eq. (12) to maintain the intra-class compactness. (2) $\mathbf{N} = \mathbf{K} > \mathbf{1}, \mathbf{C} = \mathbf{0}$. This means there are several identities existing in a specific class. $N = K$ represents the images for each identity are enough to form a valid cluster in the latent space and there are no outlier images. Under this manner, with appropriate hyper-parameter, the state-of-the-art methods Sub-center ArcFace [2] and SKH [13] can effectively cope with this label conflict. However, the unconstrained N and K still make them ineffective even with some performance gains. In ESL, the sub-centers producing via Eq. (10) adaptively produce new sub-centers to accommodate the external identities beyond the initialized sub-center number if there are fewer sub-centers than identities in the class. If the number of sub-centers is larger than the identity number, the merging strategy will merge clusters with the same identity to keep the intra-class compactness. (3) $\mathbf{N} > \mathbf{K} > \mathbf{1}, \mathbf{C} = \mathbf{0}$. Besides the conflict clusters, several identities can not converge to valid clusters in the latent space. We find that this is caused by the few-shot samples in each identity. It lacks intra-class diversity, which prevents the network from effective optimization and leads to the collapse of the feature dimension. To deal with these undiscoverable outliers, we design the sub-centers dropping operation to discard these sub-centers with few samples or slack intra-class compactness based on Eq. 11 in ESL. This is based on our observation that these sub-centers are not dominated by any one identity. Multiple outliers try to compete for the dominance, leading to bad compactness. (4) $\mathbf{N} > \mathbf{K} = \mathbf{1}, \mathbf{C} = \mathbf{0}$. It indicates that there is one valid identity and several outlier images in this class. ESL will enable the dropping strategy to remove the noise images and keep valid faces. (5) $\mathbf{N} > \mathbf{K} = \mathbf{0}, \mathbf{C} = \mathbf{0}$. It indicates each identity in the class only owns few-shot samples. We proposed dropping operation will discard all the sub-centers in this class. For $\mathbf{C} > \mathbf{0}$, there are multiple clusters with the same identity but different labels. This introduces the inter-class conflict. The state-of-the-art method

Table 2. Comparison with other noise-robust learning strategies under different types of noise. + indicates the method can solve the noise under the specific setting. +++ indicates the method can solve the noise problem. – indicates the method can not handle the problem.

Method	C	$N = K = 1$	$N = K > 1$	$N > K > 1$	$N > 1 \geq K$
ArcFace [4]	$C = 0$	+++	-	-	-
	$C > 0$	-	-	-	-
NT [10]	$C = 0$	+++	-	+	+
	$C > 0$	-	-	-	-
NR [32]	$C = 0$	+++	-	+	+
	$C > 0$	-	-	-	-
Sub-center [2]	$C = 0$	+	+	+	+
	$C > 0$	-	-	-	-
SKH [13]	$C = 0$	+	+	+	+
	$C > 0$	+	+	+	+
ESL	$C = 0$	+++	+++	+++	+++
	$C > 0$	+++	+++	+++	+++

SKH [13] can not perform well under unconstrained C . For this potential conflict, Eq. (12) in ESL can also accurately alleviate this by dynamically adjusting the label of images belonging to the merged sub-centers.

3.4 Comparison with Other Noise-robust Learning Strategies

The main difference between the proposed ESL and other methods [2, 13, 10, 32, 4] is that ESL is less affected by the unconstrained N , K , and C from the real-world noise. It’s more flexible to face recognition under different types of noise while keeping extreme simplicity, only adding three sub-centers operation. To better demonstrate this, we make a detailed comparison with other methods under fine-grained noisy faces as shown in Tab. 2. The superiority of our method is mainly due to the flexible sub-center evolving strategy, which can handle variant intra-class noise and inter-class noise simultaneously.

4 Experiments

4.1 Experimental Settings

Datasets. MS1MV0 [8] and MS1MV3 [5] are popular academic face recognition datasets. MS1MV0 [8] is raw data that is collected from the search engine based on a name list, in which there is around 50% noise. MS1MV3 [5] is the cleaned version of MS1MV0 [8] by a semi-automatic pipeline. To further explore the effectiveness of our proposed ESL, we also elaborately construct synthetic noisy datasets. We establish intra-conflict, inter-class conflict, and mixture conflict noisy datasets, which will be detailed introduced in the supplementary material. As for the performance evaluation, we tackle the True Accept Rate (TAR) at a

Table 3. Experiments of different settings on MS1MV0 and synthetic mixture noisy dataset comparing with state-of-the-art methods.

Method	Dataset	IJB-B			IJB-C		
		$1e-3$	$1e-4$	$1e-5$	$1e-3$	$1e-4$	$1e-5$
ArcFace [4]	MS1MV0	93.27	87.87	74.74	94.59	90.27	81.11
Sub-center ArcFace M=3 [2]	MS1MV0	94.88	91.70	85.62	95.98	93.72	90.59
Co-ming [27]	MS1MV0	94.99	91.80	85.57	95.95	93.82	90.71
NT [10]	MS1MV0	94.79	91.57	85.56	95.86	93.65	90.48
NR [32]	MS1MV0	94.77	91.58	85.53	95.88	93.60	90.41
SKH + ArcFace M=3 [13]	MS1MV0	95.89	93.50	89.34	96.85	95.25	93.00
ESL + ArcFace	MS1MV0	96.61	94.60	91.15	97.58	96.23	94.24
ArcFace [4]	Mixture of Noises	93.17	87.54	74.02	94.99	90.03	82.40
Sub-center ArcFace M=3 [2]	Mixture of Noises	92.83	86.80	73.11	94.20	89.32	81.43
SKH + ArcFace M=4 [13]	Mixture of Noises	95.76	93.62	89.18	96.89	95.16	92.71
ESL + ArcFace	Mixture of Noises	96.48	94.51	90.95	97.62	96.22	93.60

specific False Accept Rate (FAR) as the metric. We mainly consider the performance on IJB-B [29] dataset and IJB-C [17] dataset. Moreover, we also report the results on LFW [11], CFP-FP [21] and AgeDB-30 [18].

Implementation Details. Following ArcFace [4], we generate aligned faces with RetinaFace [3] and resize images to (112×112) . We employ ResNet-50 [9] as backbone network to extract 512-D feature embedding. For the experiments in our paper, we initialize the learning rate with 0.1 and divide it by 10 at 100K, 160K, and 220K iteration. The total training iteration number is set as 240K. We adopt an SGD optimizer, then set momentum as 0.9 and weight decay as $5e-4$. The model is trained on 8 NVIDIA A100 GPUs with a total batch size of 512. The experiments are implemented with Pytorch [19] framework. For experiments on ESL, we set the initial number of sub-centers for each class as 3. The λ_1 , λ_2 , λ_3 and λ_4 is separately set as 2, 2, 0.25, and 3.

4.2 Comparison with State-of-the-art

We conduct extensive experiments to investigate our proposed Evolving Sub-centers Learning. In Table. 3, we compare ESL with state-of-the-art methods on both real-word noisy dataset MS1MV0 [8] and the synthetic mixture of noise dataset. Without special instructions, the noise ratio is 50%.

When training on noisy data, ArcFace has an obvious performance drop. It demonstrates that noise samples would do dramatically harm to the optimization process. ESL can easily outperform current methods by an obvious margin. To be specific, ESL can outperform Sub-center ArcFace [2] by 2.51% and SKH [13] by 0.98% on IJB-C dataset. On the synthetic mixture noisy dataset, we make a grid search of the sub-center number in Sub-center ArcFace [2] and SKH [13]. Sub-center ArcFace [2] achieves the best performance when M=3 and SKH [13] achieves the best performance when M=4. ESL can easily outperform Sub-center ArcFace [2] by 6.9% and SKH [13] by 1.06% on IJB-C dataset. Our proposed

Table 4. Ablation experiments to explore the hyperparameters.

λ_1 (Eq. (5))	λ_2 (Eq. (10))	λ_3 (Eq. (11))	λ_4 (Eq. (12))	M_j (Eq. (5))	TAR@FAR=-4
2	2	0.25	3	3	96.22
1	2	0.25	3	3	95.73
3	2	0.25	3	3	95.88
2	1	0.25	3	3	96.11
2	3	0.25	3	3	95.89
2	2	0.2	3	3	96.05
2	2	0.3	3	3	96.18
2	2	0.25	1	3	95.07
2	2	0.25	2	3	95.82
2	2	0.25	3	1	96.03
2	2	0.25	3	2	96.14
2	2	0.25	3	4	96.20
2	2	0.25	3	5	96.19

Table 5. Ablation experiments to verify the effectiveness of proposed operations.

ArcFace	class-specific sub-center	Merging	Producing	Dropping	Dataset	IJB-C		
						1e-3	1e-4	1e-5
✓	✗	✗	✗	✗	Mixture of Noises	94.99	90.03	82.40
✓	✓	✗	✗	✗	Mixture of Noises	95.35	93.76	90.88
✓	✓	✓	✗	✗	Mixture of Noises	96.02	94.51	92.14
✓	✓	✓	✓	✗	Mixture of Noises	97.23	95.54	92.98
✓	✓	✓	✓	✓	Mixture of Noises	97.62	96.22	93.60

ESL can handle the fine-grained intra-class conflict and inter-class conflict under unconstrained N , K , and C , which brings significant performance improvement.

4.3 Ablation Study

Exploration on Hyperparameters. The hyperparameters in our proposed ESL contain the initial sub-center number for each class and the λ in each proposed operation. In Tab. 4, we investigate the impact of each hyperparameter.

Effectiveness of Proposed Operations. To demonstrate the effectiveness of our proposed ESL, we decouple each operation to ablate each of them on the mixture noisy dataset in Tab. 5.

We take turns adding each component to the original ArcFace [4] baseline. Due to the gradient conflict from massive fine-grained intra-class noise and inter-class noise, ArcFace [4] only achieves limited performance. Sub-center loss introduces sub-centers for each identity to deal with the intra-class conflict. Meanwhile, the ignore strategy can ease part of conflict from inter-class noise. Sub-center loss brings 3.73% performance improvement. The merging operation aims to merge images that share the same identity but belong to different sub-centers. The merging operation boosts the performance by 0.75%. The producing operation can automatically align the sub-centers and the actual identity number in each class, which improves intra-class compactness effectively. It further brings 1.03% performance improvement. The dropping operation tends to drop

Table 6. Ablation experiments to compare ESL with posterior cleaning methods. The GPU hour is measured on NVIDIA A100 GPU. $M=n \downarrow 1$ indicates the posterior data clean strategy proposed in Sub-center Arcface [2].

Method	Dataset	Posterior Clean	GPU Hour	IJB-C		
				$1e-3$	$1e-4$	$1e-5$
Sub-center ArcFace $M=3 \downarrow 1$	MS1MV0	✓	128	97.40	95.92	94.03
SKH + ArcFace $M=3 \downarrow 1$	MS1MV0	✓	128	96.55	96.26	94.18
ESL + ArcFace	MS1MV0	✗	80	97.58	96.23	94.24
ArcFace	MS1MV3	✗	64	97.64	96.44	94.66
Sub-center ArcFace $M=3 \downarrow 1$	Mixture of Noises	✓	128	97.13	95.89	92.67
SKH + ArcFace $M=4 \downarrow 1$	Mixture of Noises	✓	128	97.46	96.14	92.87
ESL + ArcFace	Mixture of Noises	✗	80	97.62	96.22	93.60

the outlier faces without the specific positive sub-center. These faces are hard to optimize and would harm the optimization process. It can obtain a significant performance gain by 0.68% under the fine-grained noisy dataset.

Efficiency of ESL. Deng et al. [2] and Liu et al. [13] adopt a posterior data clean strategy to filter out noise samples in an offline manner. Deng et al. [2] searches for the intra-class margin to drop the outlier samples for each domain center. Liu et al. [13] further introduces inter-class margin to merge samples belonging to different centers. For each margin setting, they should train for 20 epochs to verify its effectiveness, which is extremely time and computation resources consuming. In Table. 6, we compare ESL with these posterior cleaning strategies. ESL can also achieve better performance on both MS1MV0 [8] and synthetic mixture noisy dataset. Meanwhile, there is only a slight gap between ESL and ArcFace [4] training on cleaned MS1MV3.

Robustness under Various Noise Ratio. To further investigate the effectiveness of our proposed ESL, we conduct sufficient experiments under various noise ratios. As shown in Fig. 4, we visualize the relationship between noise ratio and evaluation results. ESL can remain robustness under different noise ratio and surpass Sub-center ArcFace [2] and SKH [13] by a large margin.

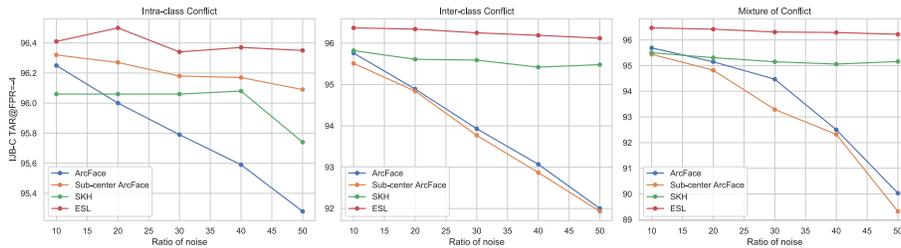


Fig. 4. Experiments of ArcFace, Sub-center ArcFace, SKH and ESL on different noise ratio. We tackle the TAR@FAR=-4 on IJB-C dataset as evaluation metric.

In Table. 7, we also compare our proposed ESL with other methods on the cleaned MS1MV3 dataset. Samples in a clean class would converge to different

Table 7. Experiments on cleaned MS1MV3 dataset. For IJB-B and IJB-C dataset, we adopt the $\text{TPR@FPR}=-4$ as evaluation metric.

Method	Dataset	IJB-B	IJB-C	LFW	CFP-FP	AgeDB-30
ArcFace	MS1MV3	95.04	96.44	99.83	98.57	98.12
Sub-center ArcFace M=3	MS1MV3	94.84	96.35	99.75	98.50	98.14
Sub-center ArcFace M=3 $\downarrow 1$	MS1MV3	94.87	96.43	99.78	98.52	98.19
SKH + ArcFace M=3	MS1MV3	93.50	95.25	99.78	98.59	98.23
SKH + ArcFace M=3 $\downarrow 1$	MS1MV3	94.98	96.48	99.77	98.70	98.25
ESL + ArcFace	MS1MV3	95.12	96.50	99.80	98.72	98.43

Table 8. Experiments on CosFace loss function.

Method	Dataset	IJB-B			IJB-C		
		$1e-3$	$1e-4$	$1e-5$	$1e-3$	$1e-4$	$1e-5$
CosFace	Mixture of Noises	93.44	86.87	74.20	95.15	90.56	83.01
Sub-center CosFace M=3	Mixture of Noises	91.85	84.40	69.88	94.25	89.19	80.25
SKH + CosFace M=4	Mixture of Noises	95.07	93.15	87.13	96.28	94.46	91.87
ESL + CosFace	Mixture of Noises	96.52	94.64	88.93	97.50	96.10	93.51

sub-centers so that the performance of Sub-center ArcFace [2] slightly drops. SKH [13] leads to a significant performance drop when directly training on the cleaned dataset. The restraint of SKH [13] forces each hyperplane to contain a subset of all IDs in the cleaned dataset, which does great harm to the inter-class representation learning. Compare with Sub-center ArcFace [2] and SKH [13], our proposed ESL can further boost the performance on the cleaned dataset, which further verifies the generalization of ESL.

Generalization on Other Loss Function. We also verify the generalization ability of proposed ESL on CosFace [26], which is another popular loss function for deep face recognition. In Tabel. 8, we can observe that ESL can significantly outperform Sub-center [2] and SKH [13] by a large margin.

5 Conclusions

In this paper, We reformulate the noise type of faces in each class into a more fine-grained manner as N -identities $|K^C$ -clusters. The key to robust representation learning strategies under real-world noise is the flexibility of the algorithm to the variation of N , K , and C . Furthermore, we introduce a general flexible method, named Evolving Sub-centers Learning (ESL), to improve the robustness of feature representation on noisy training data. The proposed ESL enjoys scalability to different combinations of N , K , and C , which is more robust to unconstrained real-world noise. Extensive experiments on noisy data training demonstrate the effectiveness of ESL and it provides a new state-of-the-art for noise-robust representation learning on large-scale noisy faces.

Acknowledgments The work was supported by the National Key R&D Program of China under Grant 2021ZD0201300.

References

1. An, X., Zhu, X., Gao, Y., Xiao, Y., Zhao, Y., Feng, Z., Wu, L., Qin, B., Zhang, M., Zhang, D., et al.: Partial fc: Training 10 million identities on a single machine. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1445–1449 (2021)
2. Deng, J., Guo, J., Liu, T., Gong, M., Zafeiriou, S.: Sub-center arcface: Boosting face recognition by large-scale noisy web faces. In: European Conference on Computer Vision. pp. 741–757. Springer (2020)
3. Deng, J., Guo, J., Ververas, E., Kotsia, I., Zafeiriou, S.: Retinaface: Single-shot multi-level face localisation in the wild. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5203–5212 (2020)
4. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4690–4699 (2019)
5. Deng, J., Guo, J., Zhang, D., Deng, Y., Lu, X., Shi, S.: Lightweight face recognition challenge. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 0–0 (2019)
6. Deng, J., Zhou, Y., Zafeiriou, S.: Marginal loss for deep face recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 60–68 (2017)
7. Du, H., Shi, H., Liu, Y., Wang, J., Lei, Z., Zeng, D., Mei, T.: Semi-siamese training for shallow face learning. In: European Conference on Computer Vision. pp. 36–53. Springer (2020)
8. Guo, Y., Zhang, L., Hu, Y., He, X., Gao, J.: Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In: European conference on computer vision. pp. 87–102. Springer (2016)
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
10. Hu, W., Huang, Y., Zhang, F., Li, R.: Noise-tolerant paradigm for training face recognition cnns. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11887–11896 (2019)
11. Huang, G.B., Mattar, M., Berg, T., Learned-Miller, E.: Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In: Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition (2008)
12. Huang, Y., Wang, Y., Tai, Y., Liu, X., Shen, P., Li, S., Li, J., Huang, F.: Curricular-face: adaptive curriculum learning loss for deep face recognition. In: proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5901–5910 (2020)
13. Liu, B., Song, G., Zhang, M., You, H., Liu, Y.: Switchable k-class hyperplanes for noise-robust representation learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3019–3028 (2021)
14. Liu, W., Wen, Y., Yu, Z., Li, M., Raj, B., Song, L.: Sphereface: Deep hypersphere embedding for face recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 212–220 (2017)
15. Liu, Y., Song, G., Shao, J., Jin, X., Wang, X.: Transductive centroid projection for semi-supervised large-scale recognition. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 70–86 (2018)

16. Liu, Y., et al.: Towards flops-constrained face recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 0–0 (2019)
17. Maze, B., Adams, J., Duncan, J.A., Kalka, N., Miller, T., Otto, C., Jain, A.K., Niggel, W.T., Anderson, J., Cheney, J., et al.: Iarpa janus benchmark-c: Face dataset and protocol. In: 2018 international conference on biometrics (ICB). pp. 158–165. IEEE (2018)
18. Moschoglou, S., Papaioannou, A., Sagonas, C., Deng, J., Kotsia, I., Zafeiriou, S.: Agedb: the first manually collected, in-the-wild age database. In: proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 51–59 (2017)
19. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32** (2019)
20. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 815–823 (2015)
21. Sengupta, S., Chen, J.C., Castillo, C., Patel, V.M., Chellappa, R., Jacobs, D.W.: Frontal to profile face verification in the wild. In: 2016 IEEE winter conference on applications of computer vision (WACV). pp. 1–9. IEEE (2016)
22. Song, G., Leng, B., Liu, Y., Hetang, C., Cai, S.: Region-based quality estimation network for large-scale person re-identification. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)
23. Song, G., Liu, Y., Jiang, M., Wang, Y., Yan, J., Leng, B.: Beyond trade-off: Accelerate fcn-based face detector with higher accuracy. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7756–7764 (2018)
24. Song, G., Liu, Y., Wang, X.: Revisiting the sibling head in object detector. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11563–11572 (2020)
25. Sun, Y., Wang, X., Tang, X.: Deep learning face representation from predicting 10,000 classes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1891–1898 (2014)
26. Wang, H., Wang, Y., Zhou, Z., Ji, X., Gong, D., Zhou, J., Li, Z., Liu, W.: Cosface: Large margin cosine loss for deep face recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5265–5274 (2018)
27. Wang, X., Wang, S., Wang, J., Shi, H., Mei, T.: Co-mining: Deep face recognition with noisy labels. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9358–9367 (2019)
28. Wang, X., Zhang, S., Wang, S., Fu, T., Shi, H., Mei, T.: Mis-classified vector guided softmax loss for face recognition. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 12241–12248 (2020)
29. Whitelam, C., Taborsky, E., Blanton, A., Maze, B., Adams, J., Miller, T., Kalka, N., Jain, A.K., Duncan, J.A., Allen, K., et al.: Iarpa janus benchmark-b face dataset. In: proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 90–98 (2017)
30. Zhang, M., Song, G., Zhou, H., Liu, Y.: Discriminability distillation in group representation learning. In: European Conference on Computer Vision. pp. 1–19. Springer (2020)

31. Zhang, X., Zhao, R., Qiao, Y., Wang, X., Li, H.: Adacos: Adaptively scaling cosine logits for effectively learning deep face representations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10823–10832 (2019)
32. Zhong, Y., Deng, W., Wang, M., Hu, J., Peng, J., Tao, X., Huang, Y.: Unequal-training for deep face recognition with long-tailed noisy data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7812–7821 (2019)
33. Zhu, M., Martínez, A.M.: Optimal subclass discovery for discriminant analysis. In: 2004 Conference on Computer Vision and Pattern Recognition Workshop. pp. 97–97. IEEE (2004)
34. Zhu, M., Martínez, A.M.: Subclass discriminant analysis. *IEEE transactions on pattern analysis and machine intelligence* **28**(8), 1274–1286 (2006)
35. Zhu, Z., Huang, G., Deng, J., Ye, Y., Huang, J., Chen, X., Zhu, J., Yang, T., Lu, J., Du, D., et al.: Webface260m: A benchmark unveiling the power of million-scale deep face recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10492–10502 (2021)