

CURRICULARFACE: ADAPTIVE CURRICULUM LEARNING LOSS FOR DEEP FACE RECOGNITION

Anonymous authors

Paper under double-blind review

ABSTRACT

As an emerging topic in face recognition, designing margin-based loss functions can increase the feature margin between different classes for enhanced discriminability. More recently, absorbing the idea of mining-based strategies is adopted to emphasize the misclassified samples and achieve promising results. However, during the entire training process, the prior methods either do not explicitly emphasize the sample based on its importance that renders the hard samples not fully exploited; or explicitly emphasize the effects of semi-hard/hard samples even at the early training stage that may lead to convergence issue. In this work, we propose a novel Adaptive Curriculum Learning loss (CurricularFace) that embeds the idea of curriculum learning into the loss function to achieve a novel training strategy for deep face recognition, which mainly addresses easy samples in the early training stage and hard ones in the later stage. Specifically, our CurricularFace adaptively adjusts the relative importance of easy and hard samples during different training stages. In each stage, different samples are assigned with different importance according to their corresponding difficulty. Extensive experimental results on popular benchmarks demonstrate the superiority of our CurricularFace over the state-of-the-art competitors. Code will be available upon publication.

INTRODUCTION

The success of Convolutional Neural Networks (CNNs) on face recognition can be mainly credited to : enormous training data, network architectures, and loss functions. Recently, designing appropriate loss functions that enhance discriminative power is pivotal for training deep face CNNs.

Current state-of-the-art face recognition methods mainly adopt softmax-based classification loss. Since the learned features with the original softmax is not discriminative enough for the open-set face recognition problem, several margin-based variants have been proposed to enhance features' discriminative power. For example, explicit margin, *i.e.*, CosFace (Wang et al., 2018a), Sphereface (Li et al., 2017), ArcFace (Deng et al., 2019), and implicit margin, *i.e.*, Adacos (Zhang et al., 2019a), supplement the original softmax function to enforce greater intra-class compactness and inter-class discrepancy, which are shown to result in more discriminate features. However, these margin-based loss functions *do not explicitly emphasize each sample according to its importance*.

As demonstrated in Chen et al. (2019), hard sample mining is also a critical step to further improve the final accuracy. Recently, Triplet loss (Schroff et al., 2015) and SV-Arc-Softmax (Wang et al., 2018b) integrate the motivations of both margin and mining into one framework for deep face recognition. Triplet loss adopts a semi-hard mining strategy to obtain semi-hard triplets and enlarge the margin between triplet samples. SV-Arc-Softmax (Wang et al., 2018b) clearly defines hard samples as misclassified samples and emphasizes them by increasing the weights of their negative cosine similarities with a preset constant. In a nutshell, mining-based loss functions *explicitly emphasize the effects of semi-hard or hard samples*.

However, there are drawbacks in training strategies of both margin- and mining-based loss functions. For margin-based methods, mining strategy is ignored and thus the difficulty of each sample is not fully exploited, which may lead to convergence issues when using a large margin on small backbones, *e.g.*, MobileFaceNet (Chen et al., 2018). As shown in Fig. 1, the modulation coefficient for the negative cosine similarities $I(\cdot)$ is fixed as a constant 1 in ArcFace for all samples during the entire training process. For mining-based methods, over-emphasizing hard samples in early training




Figure 1: **Different training strategies** for modulating negative cosine similarities of hard samples (*i.e.*, the mis-classified sample) in ArcFace, SV-Arc-Softmax and our CurricularFace. **Left:** The modulation coefficients $I(t, \cos \theta_j)$ for negative cosine similarities of hard samples in different methods, where t is an adaptively estimated parameter and θ_j denotes the angle between the hard sample and the non-ground truth j -class center. **Right:** The corresponding hard samples' negative cosine similarities $N(t, \cos \theta_j) = I(t, \cos \theta_j) \cos \theta_j + c$ after modulation, where c indicates a constant. On one hand, during early training stage (*e.g.*, t is close to 0), hard sample's negative cosine similarities is usually reduced and thus leads to smaller hard sample loss than the original one. Therefore, easier samples are relatively emphasized; during later training stage (*e.g.*, t is close to 1), the hard sample's negative cosine similarities are enhanced and thus leads to larger hard sample loss. On the other hand, in the same training stage, we modulate the hard samples' negative cosine similarities with $\cos \theta_j$. Specifically, **the smaller the angle θ_j is, the larger the modulation coefficient should be.**

stage may hinder the model to converge. As SV-Arc-Softmax claimed, the manually defined constant t plays a key role in the model convergence property and a slight larger value (*e.g.*, >1.4) may cause the model difficult to converge. Thus t needs to be carefully tuned.

In this work, we propose a novel adaptive curriculum learning loss, termed CurricularFace, to achieve a novel training strategy for deep face recognition. Motivated by the nature of human learning that easy cases are learned first and then come the hard ones (Bengio et al., 2009), our CurricularFace incorporates the idea of Curriculum Learning (CL) into face recognition in an *adaptive* manner, which differs from the traditional CL in two aspects. First, **the curriculum construction is adaptive**. In traditional CL, the samples are *ordered* by the corresponding difficulty, which are often defined by a prior and then fixed to establish the curriculum. In CurricularFace, the samples are *randomly* selected in each mini-batch, while the curriculum is established adaptively via mining the hard samples online, which shows the *diversity* in samples with different importance. Second, **the importance of hard samples are adaptive**. On one hand, the relative importance between easy and hard samples is dynamic and could be adjusted in different training stages. On the other hand, the importance of each hard sample in current mini-batch depends on its own difficulty.

Specifically, the mis-classified samples in mini-batch are chosen as hard samples and weighted by adjusting the modulation coefficients $I(t, \cos \theta_j)$ of cosine similarities between the sample and the non-ground truth class center vectors, *i.e.*, negative cosine similarity $N(t, \cos \theta_j)$. To achieve the goal of adaptive curriculum learning in the entire training, we design a novel coefficient function $I(\cdot)$ that is determined by two factors: 1) the adaptively estimated parameter t that utilizes moving average of positive cosine similarities between samples and the corresponding ground-truth class center to unleash the burden of manually tuning; and 2) the angle θ_j that defines the difficulty of hard samples to achieve adaptive assignment. To sum up, the contributions of this work are:

- We propose an adaptive curriculum learning loss for face recognition, which automatically emphasizes easy samples first and hard samples later. To the best of our knowledge, it is the first work to introduce the idea of adaptive curriculum learning for face recognition.
- We design a novel modulation coefficient function $I(\cdot)$ to achieve adaptive curriculum learning during training, which connects positive and negative cosine similarity simultaneously without the need of manually tuning any additional hyper-parameter.
- We conduct extensive experiments on popular facial benchmarks, which demonstrate the superiority of our CurricularFace over the state-of-the-art competitors.

RELATED WORK

Margin-based loss function Loss design is pivotal for large-scale face recognition. Current state-of-the-art deep face recognition methods mostly adopt softmax-based classification loss. Since the learned features with the original softmax loss are not guaranteed to be discriminative enough for open-set face recognition problem, margin-based losses (Liu et al., 2016; Li et al., 2017; Deng et al., 2019) are proposed. Though the margin-based loss functions are verified to obtain good performance, they do not take the difficultness of each sample into consideration, while our CurricularFace emphasizes easy samples first and hard samples later, which is more reasonable and effectiveness.

Mining-based loss function Though some mining-based loss function such as Focal loss (Lin et al., 2017), Online Hard Sample Mining (OHEM) (Shrivastava et al., 2016) are prevalent in the field of object detection, they are rarely used in face recognition. OHEM focuses on the large-loss samples in one mini-batch, in which the percentage of the hard samples is empirically determined and easy samples are completely discarded. Focal loss is a soft mining variant that rectifies the loss function to an elaborately designed form, where two hyper-parameters should be tuned with a lot of efforts to decide the weights of each samples and hard samples are emphasized by reducing the weight of easy samples. The recent work, SV-Arc-Softmax (Wang et al., 2018b) fuses the motivations of both margin and mining into one framework for deep face recognition. They define hard samples as misclassified samples and enlarge the weight of hard samples with a preset constant. Our method differs from SV-Arc-Softmax in three aspects: 1) We do not always emphasize the hard samples, especially in the early training stages. 2) We assign different weights for hard samples according to their corresponding difficultness. 3) There’s no need in our method to manually tune the additional hyper-parameter t , which is estimated adaptively.

Curriculum Learning Learning from easier samples first and harder samples later is a common strategy in Curriculum Learning (CL) (Bengio et al., 2009), (Zhou & Bilmes, 2018). The key problem in CL is to define the difficultness of each sample. For example, Basu & Christensen (2013) takes the negative distance to the boundary as the indicator for easiness in classification. However, the ad-hoc curriculum design in CL turns out to be difficult to implement in different problems. To alleviate this issue, Kumar et al. (2010) designs a new formulation, called Self-Paced Learning (SPL), where examples with lower losses are considered to be easier and emphasized during training. The key differences between our CurricularFace with SPL are: 1) Our method focuses on easier samples in the early training stage and emphasizes hard samples in the later training stage. 2) Our method proposes a novel modulation function $N(\cdot)$ for negative cosine similarities, which achieves not only adaptive assignment on modulation coefficients $I(\cdot)$ for different samples in the same training stage, but also adaptive curriculum learning strategy in different training stages.

THE PROPOSED CURRICULARFACE

PRELIMINARY KNOWLEDGE ON LOSS FUNCTION

The original softmax loss is formulated as follows:

$$\mathcal{L} = -\log \frac{e^{W_{y_i} x_i + b_{y_i}}}{\sum_{j=1}^n e^{W_j x_i + b_j}}, \quad (1)$$

where $x_i \in R^d$ denotes the deep feature of i -th sample which belongs to the y_i class, $W_j \in R^d$ denotes the j -th column of the weight $W \in R^{d \times n}$ and b_j is the bias term. The class number and the embedding feature size are n and d , respectively. In practice, the bias is usually set to $b_j = 0$ and the individual weight is set to $\|W_j\| = 1$ by l_2 normalization. The deep feature is also normalized and re-scaled to s . Thus, the original softmax can be modified as follows:

$$\mathcal{L} = -\log \frac{e^{s(\cos \theta_{y_i})}}{e^{s(\cos \theta_{y_i})} + \sum_{j=1, j \neq y_i}^n e^{s(\cos \theta_j)}}. \quad (2)$$

Since the learned features with original softmax loss may not be discriminative enough for open-set face recognition problem, several variants are proposed and can be formulated in a general form:

$$\mathcal{L} = -G(p(x_i)) \log \frac{e^{sT(\cos \theta_{y_i})}}{e^{sT(\cos \theta_{y_i})} + \sum_{j=1, j \neq y_i}^n e^{sN(t, \cos \theta_j)}}, \quad (3)$$

where $p(x_i) = \frac{e^{sT(\cos \theta_{y_i})}}{e^{sT(\cos \theta_{y_i})} + \sum_{j=1, j \neq y_i}^n e^{sN(t, \cos \theta_j)}}$ is the predicted ground truth probability and $G(p(x_i))$ is an indicator function. $T(\cos \theta_{y_i})$ and $N(t, \cos \theta_j) = I(t, \cos \theta_j) \cos \theta_j + c$ are the functions to modulate the positive and negative cosine similarities, respectively, where c is a constant, and $I(t, \cos \theta_j)$ denotes the modulation coefficients of negative cosine similarities. In margin-based loss function, e.g., ArcFace, $G(p(x_i)) = 1$, $T(\cos \theta_{y_i}) = \cos(\theta_{y_i} + m)$, and $N(t, \cos \theta_j) = \cos \theta_j$. It only modifies the positive cosine similarity of each sample to enhance the feature discrimination. As shown in Fig. 1, the modulation coefficients of each sample’s negative cosine similarity $I(\cdot)$ is fixed as 1. The recent work, SV-Arc-Softmax emphasizes hard samples by increasing $I(t, \cos \theta_j)$ for hard samples. That is, $G(p(x_i)) = 1$ and $N(t, \cos \theta_j)$ is formulated as follows:

$$N(t, \cos \theta_j) = \begin{cases} \cos \theta_j, & T(\cos \theta_{y_i}) - \cos \theta_j \geq 0 \\ t \cos \theta_j + t - 1, & T(\cos \theta_{y_i}) - \cos \theta_j < 0. \end{cases} \quad (4)$$

If a sample is defined to be easy, its negative cosine similarity is kept the same as the original one, $\cos \theta_j$; if as a hard sample, its negative cosine similarity becomes $t \cos \theta_j + t - 1$. That is, as shown in Fig. 1, $I(\cdot)$ is a constant and determined by a preset hyper-parameter t . Meanwhile, since t is always larger than 1, $t \cos \theta_j + t - 1 > \cos \theta_j$ always holds true, which means the model always focuses on hard samples, even in the early training stage. However, the parameter t is sensitive that a large pre-defined value (e.g., > 1.4) may lead to convergence issue.

ADAPTIVE CURRICULAR LEARNING LOSS

Next, we present the details of our proposed adaptive curriculum learning loss, which is the first attempt to introduce adaptive curriculum learning into deep face recognition. The formulation of our loss function is also contained in the general form, where $G(p(x_i)) = 1$, positive and negative cosine similarity functions are defined as follows:

$$T(\cos \theta_{y_i}) = \cos(\theta_{y_i} + m), \quad (5)$$

$$N(t, \cos \theta_j) = \begin{cases} \cos \theta_j, & T(\cos \theta_{y_i}) - \cos \theta_j \geq 0 \\ \cos \theta_j(t + \cos \theta_j), & T(\cos \theta_{y_i}) - \cos \theta_j < 0. \end{cases} \quad (6)$$

It should be noted that the positive cosine similarity can adopt any margin-based loss functions and here we adopt ArcFace as the example. As shown in Fig. 1, the modulation coefficient of hard sample negative cosine similarity $I(t, \theta_j)$ depends on both the value of t and θ_j . In the early training stage, learning from easy samples is beneficial to model convergence. Thus, t should be close to zero and $I(\cdot)$ is smaller than 1. Therefore, the weights of hard samples are reduced and the easy samples are emphasized relatively. As training goes on, the model gradually focuses on the hard samples, i.e., the value of t shall increase and $I(\cdot)$ is larger than 1. Then, the weights of hard samples are enlarged, which are thus emphasized. Moreover, within the same training stage, $I(\cdot)$ is monotonically decreasing with θ_j so that harder sample can be assigned with larger coefficient according to its difficultness. The value of the parameter t is *automatically* estimated in our CurricularFace, otherwise it would require a lot of efforts for manually tuning.

Adaptive estimation of t It is critical to determine appropriate values of t in different training stages. Ideally the value of t can indicate the model training process. We empirically find the *average* of positive cosine similarities is a good indicator. However, mini-batch statistic-based methods usually face an issue: when many extreme data are sampled in one mini-batch, the statistics can be vastly noisy and the estimation will be unstable. Exponential Moving Average (EMA) is a common solution to address this issue (Li et al., 2019). Specifically, let $r^{(k)}$ be the *average of the positive cosine similarities* of the k -th batch and be formulated as $r^{(k)} = \sum_i \cos \theta_{y_i}$, we have:

$$t^{(k)} = \alpha r^{(k)} + (1 - \alpha)t^{(k-1)}, \quad (7)$$

where $t^0 = 0$, α is the momentum parameter and set to 0.99. As shown in Fig. 2, the parameter t increases with the model training, thus the gradient modulation coefficients’ range of hard sample, $M(\cdot) = 2 \cos \theta_j + t$, also increases. Therefore, hard samples are emphasized gradually. With the EMA, we avoid the hyper-parameter tuning and make the modulation coefficients of hard sample

Algorithm 1: CurricularFace

Input: The deep feature of i -th sample x_i with its corresponding label y_i , last fully-connected layer parameters W , cosine similarity $\cos \theta_j$ between two vectors, embedding network parameters Θ , learning rate λ , number of iteration k , parameter t , and margin m

$k \leftarrow 0, t \leftarrow 0, m \leftarrow 0.5;$

while not converged do

$k \leftarrow k + 1;$

if $\cos(\theta_{y_i} + m) > \cos \theta_j$ **then**

$N(t, \cos \theta_j) = \cos \theta_j;$

else

$N(t, \cos \theta_j) = (t^{(k)} + \cos \theta_j) \cos \theta_j;$

end

$T(\cos \theta_{y_i}) = \cos(\theta_{y_i} + m);$

 Compute the loss \mathcal{L} by Eq. 8;

 Compute the back-propagation error of x_i and W_j by Eq. 9;

 Update the parameters W and Θ by: $W^{(k+1)} = W^{(k)} - \lambda^{(k)} \frac{\partial \mathcal{L}}{\partial W}, \Theta^{(k+1)} = \Theta^{(k)} - \lambda^{(k)} \frac{\partial \mathcal{L}}{\partial x_i} \frac{\partial x_i}{\partial \Theta^{(k)}};$

 Update the parameter t by Eq. 7;

end

Output: W, Θ

negative cosine similarities $I(\cdot)$ adaptive to the current training stage. To sum up, the loss function of our CurricularFace is formulated as follows:

$$\mathcal{L} = -\log \frac{e^{s \cos(\theta_{y_i} + m)}}{e^{s \cos(\theta_{y_i} + m)} + \sum_{j=1, j \neq y_i}^n e^{s N(t^{(k)}, \cos \theta_j)}}, \quad (8)$$

where $N(t^{(k)}, \cos \theta_j)$ is defined in Eq. 6. The entire training process is summarized in Algorithm 1.

Fig. 3 illustrates how the loss changes from ArcFace to our CurricularFace during training. Here are some observations: 1) As we expected, hard samples are suppressed in early training stage but emphasized later. 2) The ratio is monotonically increasing with $\cos \theta_j$, since the larger $\cos \theta_j$ is, the harder the sample is. 3) The positive cosine similarity of a perceptual-well image is often large. However, during the early training stage, the negative cosine similarities of the perceptual-well image may also be large so that it could be classified as the hard one.

Optimization Next, we show our CurricularFace can be easily optimized by the conventional stochastic gradient descent. Assuming x_i denotes the deep feature of i -th sample which belongs to the y_i class, the input of the proposed function is the logit f_j , where j denotes the j -th class.

In the forwarding process, when $j = y_i$, it is the same as the ArcFace, *i.e.*, $f_j = sT(\cos \theta_{y_i})$, $T(\cos \theta_{y_i}) = \cos(\theta_{y_i} + m)$. When $j \neq y_i$, it has two cases, if x_i is an easy sample, it is the same as the original softmax, *i.e.*, $f_j = s \cos \theta_j$. Otherwise, it will be modulated as $f_j = sN(t, \cos \theta_j)$, where $N(t, \cos \theta_j) = (t + \cos \theta_j) \cos \theta_j$. In the backward propagation process, the gradient of x_i and W_j can also be divided into three cases and formulated as follows:

$$\frac{\partial \mathcal{L}}{\partial x_i} = \begin{cases} \frac{\partial \mathcal{L}}{\partial f_{y_i}} \left(s \frac{\sin(\theta_{y_i} + m)}{\sin \theta_{y_i}} \right) W_{y_i}, & j = y_i \\ \frac{\partial \mathcal{L}}{\partial f_j} s W_j, & j \neq y_i, \text{ easy} \\ \frac{\partial \mathcal{L}}{\partial f_j} s (2 \cos \theta_j + t) W_j & j \neq y_i, \text{ hard} \end{cases}, \quad \frac{\partial \mathcal{L}}{\partial W_j} = \begin{cases} \frac{\partial \mathcal{L}}{\partial f_{y_i}} \left(s \frac{\sin(\theta_{y_i} + m)}{\sin \theta_{y_i}} \right) x_i, & j = y_i \\ \frac{\partial \mathcal{L}}{\partial f_j} s x_i, & j \neq y_i, \text{ easy} \\ \frac{\partial \mathcal{L}}{\partial f_j} s (2 \cos \theta_j + t) x_i & j \neq y_i, \text{ hard} \end{cases} \quad (9)$$

Based on the above formulations, we can find the gradient magnitude of the hard sample is determined by two parts, the negative cosine similarity $N(\cdot)$ and the value of t .




Figure 2: **Illustrations on the adaptive parameter t (red line) and gradient modulation coefficients $M(\cdot) = 2 \cos \theta_j + t$ of hard samples (green area).** Since the number of mined hard samples reduces with the model training, the green area $M(\cdot)$ is relatively smooth in early stage and there are some burrs in later stage.




Figure 3: Illustrations on (ratio between our loss and ArcFace, maximum $\cos\theta_j$) from early (Top) to later (Bottom) training stages.

Table 1: Decision boundaries of popular loss functions.

Loss	Decision Boundary
Softmax	$\cos\theta_{y_i} = \cos\theta_j$
SphereFace	$\cos(m\theta_{y_i}) = \cos\theta_j$
CosFace	$\cos\theta_{y_i} - m = \cos\theta_j$
ArcFace	$\cos(\theta_{y_i} + m) = \cos\theta_j$
SV-Arc-Softmax	$\cos(\theta_{y_i} + m) = \cos\theta_j$ (easy) $\cos(\theta_{y_i} + m) = t \cos\theta_j + t - 1$ (hard)
CurricularFace (Ours)	$\cos(\theta_{y_i} + m) = \cos\theta_j$ (easy) $\cos(\theta_{y_i} + m) = (t + \cos\theta_j) \cos\theta_j$ (hard)

Table 2: Verification performance of different values of t .

Dataset (%)	$t = 0$	$t = 0.3$	$t = 0.7$	$t = 1$	Adaptive t
LFW	99.32	99.37	99.42	99.45	99.47
CFP-FP	95.90	96.47	96.66	93.94	96.96

DISCUSSIONS WITH SOTA LOSS FUNCTIONS

Comparison with ArcFace and SV-Arc-Softmax We first discuss the difference between our CurricularFace and the two competitors, ArcFace and SV-Arc-Softmax, from the perspective of the decision boundary in Tab. 1. ArcFace introduces a margin function $T(\cos\theta_{y_i}) = \cos(\theta_{y_i} + m)$ from the perspective of positive cosine similarity. As shown in Fig. 4, its decision condition changes from $\cos\theta_{y_i} = \cos\theta_j$ (i.e., blue line) to $\cos(\theta_{y_i} + m) = \cos\theta_j$ (i.e., red line) for each sample. SV-Arc-Softmax introduces additional margin from the perspective of negative cosine similarity for hard samples, and the decision boundary becomes $\cos(\theta_{y_i} + m) = t \cos\theta_j + t - 1$ (i.e., green line). Conversely, we adaptively adjust the weights of hard samples in different training stages. The decision condition becomes $\cos(\theta_{y_i} + m) = (t + \cos\theta_j) \cos\theta_j$ (i.e., purple line). During the training stage, the decision boundary for hard samples changes from one purple line (**early stage**) to another (**later stage**), which emphasizes easy samples first and hard samples later.

Comparison with Focal loss Focal loss is a soft mining-based loss, which is formulated as: $G(p(x)) = \alpha(1 - p(x_i))^\beta$, where α and β are modulating factors that need to be tuned manually. The definition of hard samples in Focal loss is ambiguous, since it always focuses on *relatively* hard samples by reducing the weight of easier samples during the entire training process. In contrast, the definition of hard samples in our CurricularFace is more clear, i.e., mis-classified samples. Meanwhile, the weights of hard samples are *adaptively* determined in different training stages.

EXPERIMENTS

IMPLEMENTATION DETAILS

Datasets We separately employ CASIA-WebFace (Yi et al., 2014) and refined MS1MV2 (Deng et al., 2019) as our training data for fair comparisons with other methods. We extensively test our method on several popular benchmarks, including LFW (Huang et al., 2007), CFP-FP (Sengupta et al., 2016), CPLFW (Zheng et al., 2018), AgeDB (Moschoglou et al., 2017), CALFW (Zheng et al., 2017), IJB-B (Whitelam et al., 2017), IJB-C (Maze et al., 2018), and MegaFace (Kemelmacher-Shlizerman et al., 2016).

Training Setting We follow Deng et al. (2019) to generate the normalised faces (112×112) with five landmarks (Zhang et al., 2016). For the embedding network, we adopt ResNet50 and ResNet100 as in Deng et al. (2019). Our framework is implemented in Pytorch (Paszke et al., 2017). We train models on 4 NVIDIA Tesla P40 (24GB) GPU with batch size 512. The models are trained with SGD algorithm, with momentum 0.9 and weight decay $5e - 4$. On CASIA-WebFace, the learning rate starts from 0.1 and is divided by 10 at 28, 38, 46 epochs. The training process is finished at 50 epochs. On MS1MV2, we divide the learning rate at 10, 18, 22 epochs and finish at 24 epochs. We follow the common setting as Deng et al. (2019) to set scale $s = 64$ and margin $m = 0.5$, respectively. Last but not least, since we only modify the loss function but use the same backbone as previous methods (e.g., ArcFace), **NO** additional time complexity is introduced for inference.




Figure 4: **From left to right**, decision boundaries of ArcFace, SV-Arc-Softmax, and ours. **Blue line, red line, green line and purple line** denote the decision boundary of Softmax, ArcFace, SV-Arc-Softmax, and ours, respectively. m denotes the angular margin added by ArcFace. d denotes the additional margin of SV-Arc-Softmax and ours. In SV-Arc-Softmax, $d = (t - 1) \cos \theta_j + t - 1$. In ours, $d = (t + \cos \theta_j - 1) \cos \theta_j$.

ABLATION STUDY

Effects on Fixed vs. Adaptive Parameter t We first investigate the effect of adaptive estimation of t . We choose four fixed values between 0 and 1 for comparison. Specifically, 0 means the modulation coefficient $I(\cdot)$ of each hard sample’s negative cosine similarity is always reduced based on its difficulty. In contrast, 1 means the hard samples are always emphasized. 0.3 and 0.7 are between the two cases. Tab. 2 shows that it is more effective to learn from easier samples first and hard samples later based on our adaptively estimated parameter t .

Effects on Different Statistics for Estimating t We now investigate the effects of several other statistics, *i.e.*, mode of positive cosine similarities in a mini-batch, or mean of the predicted ground truth probability for estimating t in our loss. As Tab. 3 shows, on one hand, the mean of positive cosine similarities is better than the mode. On the other hand, the positive cosine similarity is more accurate than the predicted ground truth probability to indicate the training stages.

Robustness on Training Convergence As claimed in Li (2019), ArcFace exists divergence issue when using small backbones like MobileFaceNet. As the result, softmax loss must be incorporated for pre-training. To illustrate the robustness of our loss function on convergence issue with small backbone, we use the MobileFaceNet as the network architecture and train it on CASIA-WebFace. As shown in Fig. 5, when the margin m is set to 0.5, the model trained with our loss achieves 99.25 accuracy on LFW, while the model trained with ArcFace does not converge and the loss is NAN at about 2,400-th step. When the margin m is set to 0.45, both losses can converge, but our loss achieves better performance (99.20% vs. 99.10%). Comparing the yellow and red curves, since the losses of hard samples are reduced in early training stages, our loss converges much faster in the beginning, leading to lower loss than ArcFace. Later on, the value of our loss is slightly larger than ArcFace, because we emphasize the hard samples in later stages. The results prove that learning from easy samples first and hard samples later is beneficial to model convergence.

COMPARISONS WITH SOTA METHODS

Results on LFW, CFP-FP, CPLFW, AgeDB and CALFW Next, we train our CurricularFace on dataset MS1MV2 with ResNet100, and compare with the SOTA competitors on various benchmarks, including LFW for unconstrained face verification, CFP-FP and CPLFW for large pose variations, AgeDB and CALFW for age variations. As reported in Tab. 4, our CurricularFace achieves comparable result (*i.e.*, 99.80%) with the competitors on LFW where the performance is near saturated. While for both CFP-FP and CPLFW, our method shows superiority over the baselines including general methods, *e.g.*, (Wen et al., 2016), (Cao et al., 2018b), and cross-pose methods, *e.g.*, (Tran et al., 2017), (Peng et al., 2017), (Cao et al., 2018a) and (Deng et al., 2018). As a recent face recognition method, SV-Arc-Softmax achieves better performance than ArcFace, but still worse than Our CurricularFace. Finally, for AgeDB and CALFW, as Tab. 4 shows, our CurricularFace again achieves the best performance than all of the other state-of-the-art methods.

Table 3: **Verification performance of different strategies for setting t .**

Methods (%)	LFW	CFP-FP
Mode($\cos \theta_{y_i}$)	99.42	96.49
Mean(p_{x_i})	99.42	95.39
Mean($\cos \theta_{y_i}$)	99.47	96.96




Figure 5: Illustration on **convergence issue with small backbone**.

Table 4: Verification comparison with SOTA methods on various small-scale benchmarks.


Methods (%)	LFW	CFP-FP	CPLFW	AgeDB	CALFW
Center Loss (ECCV'16)	98.75	—	77.48	—	85.48
SphereFace (CVPR'17)	99.27	—	81.40	—	90.30
DRGAN (CVPR'17)	—	93.41	—	—	—
Peng <i>et al.</i> (ICCV'17)	—	93.76	—	—	—
VGGFace2 (FG'18)	99.43	—	84.00	—	90.57
Dream (CVPR'18)	—	93.98	—	—	—
Deng <i>et al.</i> (CVPR'18)	99.60	94.05	—	—	—
ArcFace (CVPR'19)	99.77	98.27	92.08	98.15	95.45
SV-Arc-Softmax	99.78	98.28	92.83	97.95	96.10
CurricularFace (Ours)	99.80	98.37	93.13	98.32	96.20

Table 5: 1:1 verification TAR (@FAR = $1e-4$) on IJB-B and IJB-C.

Methods (%)	IJB-B	IJB-C
SENet50 (FG'18)	80.0	84.1
Multicolumn (BMVC'18)	83.1	86.2
DCN (ECCV'18)	84.9	88.5
ArcFace-R100 (CVPR'19)	94.2	95.6
Adacos (CVPR'19)	—	92.4
P2SGrad (CVPR'19)	—	92.3
PFE (ICCV'19)	—	93.3
SV-Arc-Softmax	93.6	95.2
CurricularFace (Ours)	94.8	96.1

Table 6: Verification comparison with SOTA methods on MegaFace Challenge 1 using FaceScrub as the probe set. **Left table:** 'Id' refers to the rank-1 face identification accuracy with 1M distractors, and 'Ver' refers to the face verification TAR at 10^{-6} FAR. 'R' refers to data refinement on both probe set and 1M distractors. **Right figure:** Rank-1 identification results of recent SOTA methods on probe set refined from ArcFace.

CASIA (%)	Id	Ver	MS1MV2 (%)	Id	Ver
Contrastive Loss (CVPR'14)	65.21	78.86	CosFace-MS1MV2-R100	80.56	96.56
Triplet (CVPR'15)	64.79	78.32	CosFace-MS1MV2-R100, R	97.91	97.91
Center Loss (ECCV'16)	65.49	80.14	ArcFace-MS1MV2-R100	81.03	96.98
SphereFace (CVPR'17)	72.73	85.56	ArcFace-MS1MV2-R100, R	98.35	98.48
CosFace (CVPR'18)	77.11	89.88	PFE (ICCV'19)	78.95	92.51
AM-Softmax (SPL'18)	72.47	84.44	Adacos, R (CVPR'19)	97.41	—
ArcFace-CASIA-R50 (CVPR'19)	77.50	92.34	P2SGrad, R (CVPR'19)	97.25	—
ArcFace-CASIA-R50, R	91.75	93.69	SV-Arc-Softmax, R	97.14	97.57
Ours-CASIA-R50	77.65	92.91	Ours-MS1MV2-R100	81.26	97.26
Ours-CASIA-R50, R	92.48	94.55	Ours-MS1MV2-R100, R	98.71	98.64



Results on IJB-B and IJB-C The IJB-B dataset contains 1,845 subjects with 21.8K still images and 55K frames from 7,011 videos. In the 1:1 verification, there are 10,270 positive matches and 8M negative matches. The IJB-C dataset is a further extension of IJB-B, which contains about 3,500 identities with a total of 31,334 images and 117,542 unconstrained video frames. In the 1:1 verification, there are 19,557 positive matches and 15,638,932 negative matches. On IJB-B and IJB-C datasets, we employ MS1MV2 and the ResNet100 for a fair comparison with recent methods. We follow the testing protocol in ArcFace and take the average of the image features as the corresponding template representation without bells and whistles. Tab. 5 exhibits the performance of different methods, *e.g.*, Multicolumn (Xie & Zisserman, 2018), DCN (Xie *et al.*, 2018), Adacos (Zhang *et al.*, 2019a), P2SGrad (Zhang *et al.*, 2019b), PFE (Shi *et al.*, 2019) and SV-Arc-Softmax (Wang *et al.*, 2018b) on IJB-B and IJB-C 1:1 verification, our method again achieves the best performance.

Results on MegaFace Finally, we evaluate the performance on the MegaFace Challenge. The gallery set of MegaFace includes 1M images of 690K subjects, and the probe set includes 100K photos of 530 unique individuals from FaceScrub. We report the two testing results under two protocols (large or small training set). Here, we use CASIA-WebFace and MS1MV2 under the small protocol and large protocol, respectively. In Tab. 6, our method achieves the best single-model identification and verification performance under both protocols, surpassing the recent strong competitors, *e.g.*, CosFace, ArcFace, Adacos, P2SGrad and PFE. We also report the results following the ArcFace testing protocol, which refines both the probe set and the gallery set. As shown from the figure in Tab. 6, our method still clearly outperforms the competitors and achieves the best performance on both verification and identification.

CONCLUSIONS

In this paper, we propose a novel Adaptive Curriculum Learning Loss that embeds the idea of adaptive curriculum learning into deep face recognition. Our key idea is to address easy samples in the early training stage and hard ones in the later stage. Our method is easy to implement and robust to converge. Extensive experiments on popular facial benchmarks demonstrate the effectiveness of our method compared to the state-of-the-art competitors. Following the main idea of this work, future research can be expanded in various aspects, including designing a better function $N(\cdot)$ for negative cosine similarity that shares similar *adaptive* characteristic during training, and investigating the effects of *noise* samples that could be optimized as hard samples.

REFERENCES

- Sumit Basu and Janara Christensen. Teaching classification boundaries to humans. In *AAAI*, 2013.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *ICML*, 2009.
- Kaidi Cao, Yu Rong, Cheng Li, Xiaoou Tang, and Chen Change Loy. Pose-robust face recognition via deep residual equivariant mapping. In *CVPR*, 2018a.
- Qiong Cao, Li Shen, Weidi Xie, Omkar M. Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *FG*, 2018b.
- Beidi Chen, Weiyang Liu, Animesh Garg, Zhiding Yu, Anshumali Shrivastava, and Anima Anandkumar. Angular visual hardness. In *ICML Workshop on Deep Phenomena*, 2019.
- Sheng Chen, Yang Liu, Xiang Gao, and Zhen Han. Mobilefacenet: Efficient cnns for accurate real-time face verification on mobile devices. In *CCBR*, 2018.
- Jiankang Deng, Shiyang Cheng, Niannan Xue, Yuxiang Zhou, and Stefanos Zafeiriou. Uv-gan: Adversarial facial uv map completion for pose-invariant face recognition. In *CVPR*, 2018.
- Jiankang Deng, Jia Guo, and Stefanos Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. In *CVPR*, 2019.
- 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 07-49, University of Massachusetts, Amherst, October 2007.
- Ira Kemelmacher-Shlizerman, Steven M Seitz, Daniel Miller, and Evan Brossard. The megaface benchmark: 1 million faces for recognition at scale. In *CVPR*, 2016.
- M Pawan Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. In *NIPS*, 2010.
- Buyu Li, Yu Liu, and Xiaogang Wang. Gradient harmonized single-stage detector. In *AAAI*, 2019.
- Weiyang Li, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphereface: Deep hypersphere embedding for face recognition. In *CVPR*, 2017.
- Xianyang Li. Airface: Lightweight and efficient model for face recognition. *arXiv:1907.12256*, 2019.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pp. 2980–2988, 2017.
- Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. Large-margin softmax loss for convolutional neural networks. In *ICML*, 2016.
- Brianna Maze, Jocelyn Adams, James A Duncan, Nathan Kalka, Tim Miller, Charles Otto, Anil K Jain, W Tyler Niggel, Janet Anderson, Jordan Cheney, et al. Iarpa janus benchmark-c: Face dataset and protocol. In *ICB*, 2018.
- Stylianos Moschoglou, Athanasios Papaioannou, Christos Sagonas, Jiankang Deng, Irene Kotsia, and Stefanos Zafeiriou. Agedb: the first manually collected, in-the-wild age database. In *CVPR Workshops*, 2017.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.
- Xi Peng, Xiang Yu, Kihyuk Sohn, Dimitris Metaxas, and Manmohan Chandraker. Reconstruction-based disentanglement for poseinvariant face recognition. In *ICCV*, 2017.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, 2015.

- S. Sengupta, J.-C. Chen, C. Castillo, V. M. Patel, R. Chellappa, and D.W. Jacobs. Frontal to profile face verification in the wild. In *WACV*, 2016.
- Yichun Shi, Anil K Jain, and Nathan D Kalka. Probabilistic face embeddings. In *ICCV*, 2019.
- Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, 2016.
- Luan Tran, Xi Yin, and Xiaoming Liu. Disentangled representation learning GAN for pose-invariant face recognition. In *CVPR*, 2017.
- Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *CVPR*, 2018a.
- Xiaobo Wang, Shuo Wang, Shifeng Zhang, Tianyu Fu, Hailin Shi, and Tao Mei. Support vector guided softmax loss for face recognition. *arXiv:1812.11317*, 2018b.
- Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *ECCV*, 2016.
- Cameron Whitelam, Emma Taborsky, Austin Blanton, Brianna Maze, Jocelyn Adams, Tim Miller, Nathan Kalka, Anil K Jain, James A Duncan, Kristen Allen, et al. Iarpa janus benchmark-b face dataset. In *CVPR Workshops*, 2017.
- Weidi Xie and Andrew Zisserman. Multicolumn networks for face recognition. In *BMVC*, 2018.
- Weidi Xie, Li Shen, and Andrew Zisserman. Comparator networks. In *ECCV*, 2018.
- Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z. Li. Learning face representation from scratch. *arXiv:1411.7923*, 2014.
- Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503, 2016.
- Xiao Zhang, Rui Zhao, Yu Qiao, Xiaogang Wang, and Hongsheng Li. Adacos: Adaptively scaling cosine logits for effectively learning deep face representations. In *CVPR*, 2019a.
- Xiao Zhang, Rui Zhao, Junjie Yan, Mengya Gao, Yu Qiao, Xiaogang Wang, and Hongsheng Li. P2sgrad: Refined gradients for optimizing deep face models. In *CVPR*, 2019b.
- Tianyue Zheng, Weihong Deng, and Jiani Hu. Cross-age lfw: A database for studying cross-age face recognition in unconstrained environments. *arXiv:1708.08197*, 2017.
- Tianyue Zheng, Weihong Deng, and Jiani Hu. 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.
- Tianyi Zhou and Jeff Bilmes. Minimax curriculum learning: Machine teaching with desirable difficulties and scheduled diversity. In *ICLR*, 2018. URL <https://openreview.net/forum?id=BywyFQ1AW>.