

AirFace: Lightweight and Efficient Model for Face Recognition

Xianyang Li Feng Wang Qinghao Hu Cong Leng

Nanjing Artificial Intelligence Chip Research, Institute of Automation, Chinese Academy of Sciences(AIRIA)
China

qq1052308104@gmail.com {wangfeng, huqinghao, lengcong}@airia.cn

Abstract

With the development of convolutional neural network, significant progress has been made in computer vision tasks. However, the commonly used loss function softmax loss and highly efficient network architectures for common visual tasks are not as effective for face recognition. In this paper, we propose a novel loss function named Li-ArcFace based on ArcFace. Li-ArcFace takes the value of the angle through a linear function as the target logit rather than through cosine function, which has better convergence and performance on low dimensional embedding feature learning for face recognition. In terms of network architecture, we improved the performance of MobileFaceNet by increasing the network depth, width and adding attention module. Besides, we found some useful training tricks for face recognition. Under all the above effects, we won the second place in the deeplint-light challenge of L-FR2019 [6].

1. Introduction

The development of deep convolutional networks(DCNN) has made remarkable progress in a series of computer vision tasks. However, it is not so effective while using the common method for face recognition. Softmax loss, which is commonly used in classification, can't maximize inter-class variance and minimize intra-class variance of embedding feature vectors. In order to obtain highly discriminative embedding features(See Figure 1), a series of novel loss functions have been proposed in recent years, such as A-Softmax [12], CosFace/AM-Softmax [21, 19], ArcFace [5]. Among them, ArcFace achieved state-of-the-art performance by adding additive margin between classes in the angle space. But, when ArcFace is used to train some efficient networks with small(128) embedding size, it's hard to train from scratch(for example, the embedding feature size of MobileFaceNet [4] is 128, it can't converge when trained from scratch with ArcFace). As a result, to guarantee the convergence of training, pre-training is always re-




Figure 1. Schematic diagram of discriminative embedding features. $W1$ refers to the center of the class 1, $f1$ refers to the embedding feature of class 1, $f2$ refers to the embedding feature of class 2. During training, the variance of the same person becomes smaller, and the variance of different people becomes larger.

quired with softmax loss. To overcome this problem, we carefully designed a novel loss function named Li-ArcFace based on ArcFace, which performs better in convergence.

Face recognition technology is now widely used on mobile devices, which requires that the computational cost of the model should not be too large. In recent years, some highly efficient neural network architectures have been proposed, such as MobileNetV1 [10], ShuffleNet [23], and MobileNetV2 [16], but they're all designed for common visual recognition tasks instead of face recognition. Their performance is really general when used for face recognition directly. MobileFaceNet is designed for face recognition based on MobileNetV2, achieved remarkable accuracy on LFW [11], AgeDB [15]. And it is even comparable to state-of-the-art big DCNN model on MegaFace [14] Challenge 1 under the much smaller computational resources. In this paper, with limited amount of computation, we carefully designed a higher performance network architecture based on MobileFaceNet.

2. Related Work

Loss function. Softmax loss is the most commonly used loss function in classification, which is presented as:

$$L_1 = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{W_{y_i}^T x_i + b_{y_i}}}{\sum_{j=1}^n e^{W_j^T x_i + b_j}} \quad (1)$$

where x_i denotes the embedding feature of the i -th sample belonging to the y_i -th class, and the dimension of the embedding feature (hereinafter referred to as embedding size) is set as d . $W_j \in \mathbb{R}^d$ denotes the j -th column of the weight $W \in \mathbb{R}^{d \times n}$ and $b_j \in \mathbb{R}$ is the bias term. The batch size is N and the class number of training data is n . However, the traditional softmax loss lacks the power to supervise the embedding feature to minimize inter-class similarity and maximize intra-class similarity. In SphereFace [12] and NormFace [20], the bias term is being removed at first and then fixing the $\|W_j\| = 1$, $\|x_i\| = s$ by l_2 normalisation, such that the logit is:

$$W_j^T x_i = \|W_j\| \times \|x_i\| \times \cos \theta_j = s \times \cos \theta_j \quad (2)$$

Where θ_j denotes the angle between x_i and W_j . Thus the logit is only depend on the cosine of the angle. The modified loss function can be formulated as :

$$L_2 = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{s \cos \theta_{y_i}}}{\sum_{j=1}^n e^{s \cos \theta_j}} \quad (3)$$

Although L_2 guarantees a high similarity of features of the same person, it can't separate different classes very well. In this paper, we use N-Softmax denotes L_2 . In ArcFace [5], the authors added an additive angular margin m within $\cos \theta_{y_i}$, which can enhance the intra-class compactness and inter-class discrepancy, the ArcFace is formulated as :

$$L_3 = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{s \cos(\theta_{y_i} + m)}}{e^{s \cos(\theta_{y_i} + m)} + \sum_{j=1, j \neq y_i}^n e^{s \cos \theta_j}} \quad (4)$$

When using ArcFace to train models with 512-dimensional embedding feature, it has well convergence and state-of-the-art performance. However, it will be difficult to converge while training some highly efficient networks with 128-dimensional embedding feature from scratch.

network architectures. Face recognition is being used more and more on mobile devices. So it's really important to optimize the trade-off between accuracy and computational cost when designing deep neural network architecture. In recent years, some highly efficient neural network architectures have been proposed for common visual tasks. MobileNetV1 [10] used depthwise separable convolution instead of traditional convolution to reduce




Figure 2. Target logit curves.




Figure 3. Decision margins of different loss functions under binary classification case. The yellow areas are the decision margins, the red areas are the overlap area of class1 and class2.

computational cost and improve network efficiency. MobileNetV2 [16] introduced inverted residuals and linear bottlenecks to further improve network efficiency. However, these lightweight network architectures are not so accurate when using these unchanged for face recognition. The author of MobileFaceNet [4] found the weakness of common mobile networks for Face recognition, and solved it by replacing global average pooling with global depthwise convolution(GDC). And the network architecture of MobileFaceNet is specifically designed for face recognition with smaller expansion factors in bottlenecks and more output channels at the beginning of the network architecture.

3. Proposed Approach

3.1. Li-ArcFace

In ArcFace, the authors added an angular margin m within $\cos \theta_{y_i}$, which takes $s \times \cos(\theta_{y_i} + m)$ as the target logit. The loss function proposed by us takes the angle after a linear function as the logit rather than cosine function. In the same way, we remove the bias term and then fix the $\|W_j\| = 1$, $\|x_i\| = s$ by l_2 normalisation, the θ_j denotes the angle be-

tween x_i and W_j , thus $\theta_j = \arccos(W_j^T x_i/s) \in [0, \pi]$. At first, we constructed a linear function $f(x) = (\pi - 2x)/\pi$, and we have $f(\theta_j) \in [-1, 1]$. Then we add an additive angular margin m in the target logit. In the end we have $s \times (\pi - 2(\theta_{y_i} + m))/\pi$ as the target logit. We call this novel loss function Li-ArcFace, the prefix Li refers to the linear function. The whole Li-ArcFace can be formulated as follows

$$L_4 = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{s(\pi - 2(\theta_{y_i} + m))/\pi}}{e^{s(\pi - 2(\theta_{y_i} + m))/\pi} + \sum_{j=1, j \neq y_i}^n e^{s(\pi - 2\theta_j)/\pi}} \quad (5)$$

There are two main advantages of using this linear function to replace cosine function. Firstly, it is monotonic decreasing when the angle is between 0 and $\pi + m$, which will have better convergence, especially for model with small embedding size. For example, when training MobileFaceNet with ArcFace, it will lead to divergence(NaN). Therefore, softmax loss must be used for pre-training before convergence. The proposed loss does not require these two-stage training. Secondly, the penalty of the proposed loss function increases linearly as the angle between embedding feature x_i and center W_{y_i} increasing, so that the target logit decreases linearly which is more intuitive(See Figure 2). It will not decline rapidly in some angles, but slowly in others (corresponding to the gradient value of the target logit curve), which makes the proposed loss function have better performance in the end. In terms of geometric decision margins, ArcFace will have a part of overlap area when the x_i deviates too much from the center W_{y_i} , because of its non-monotonicity of the target logit curve. This area can be distinguished as class 1 and class 2 with ArcFace. The proposed loss function would't have this overlap area(See Figure 3).

3.2. Network Architectures

In this section, we introduce our proposed network architecture. Our network architecture is based on a deeper MobileFaceNet(y2) [5], so the residual bottlenecks proposed in MobileNetV2 are used as our main building blocks. Table 1 shows the details of our network architecture. We follow the MobileFaceNet, expansion factors for bottlenecks in our architecture are much smaller than those in MobileNetV2 and using PReLU as the non-linearity rather than Relu. Nevertheless, we noticed the importance of network width in face recognition. MobileFaceNet is significantly wider than MobileNetV2 at the beginning of the network. But MobileFaceNet does not double the output channels of the network when sampling at last stage. We did that within limited amount of computation. At the same time, we carefully adjusted the depth of the network, and we introduced the attention module CBAM [22] into every bottleneck in the network. But we changed the last activation function in

Input	Operator	t	c	n	s
112x112x3	conv 3x3	-	64	1	2
56x56x64	depthwise conv3x3	-	64	1	1
56x56x64	bottleneck	2	64	1	2
28x28x64	bottleneck	2	64	9	1
28x28x64	bottleneck	4	128	1	2
14x14x128	bottleneck	2	128	16	1
14x14x128	bottleneck	8	256	1	2
7x7x256	bottleneck	2	256	6	1
7x7x256	conv1x1	-	1024	1	1
7x7x1024	linear GDCConv7x7	-	1024	1	1
1x1x1024	linear conv1x1	-	512	1	1

Table 1. The proposed network architecture, n refers to the number of repetitions, c refers to output channels, t refers to The expansion factor.

the attention module from sigmoid to 1+tanh. The range of 1+tanh is 0 to 2. When a channel or spatial position needs to be enhanced, it is multiplied by a weight between one and two; when it needs to be weakened, it is multiplied by a weight between zero and one. Using 1+tanh instead of sigmoid is more intuitive and makes training converge faster.

3.3. Training Tricks for Face Recognition

During the competition of LFR, we found some useful training tricks for face recognition. Firstly, using a variety of loss functions to fine-tune the model will make the features more robust and improve the accuracy to some extent. In the competition, we used Li-ArcFace, ArcFace, combined loss to fine-tune our model. Secondly, in 512-dimensional embedding feature space, it is difficult for the lightweight model to learn the distribution of the features. It is an effective method to use some large models to guide the feature distribution of lightweight models [9, 13].

4. Experiments

4.1. Ablation Experiment of Li-ArcFace

In this section, we mainly introduce some comparative experiments of different loss functions.

Implementation Details. We use the MobileFaceNet as the network architecture, in which the embedding size is set as 128. And batch size is set as 256 x 4. Training models on four NVIDIA TITIAN XP GPUs. We use SGD with momentum 0.9 to optimize models. The scale of the feature is set as 64. The training data set is CASIA-Webface, which contains 10K Identity 0.5M pictures. The learning rate starts at 0.1, divides by 10 at 18k, 26k and 29K iterations, and finally stops training at 30K iterations. At last, we compare the performance on Labelled Faces in the Wild (LFW) [11], Celebrities in Frontal Profile (CFP) [18] and Age Database (AgeDB) [15]. In this paper, ArcFace (m=0.5, NS) refers

Weight decay	LFW	CFP-FP	AgeDB
5e-4	99.17	94.09	92.95
5e-4 wdm10	99.18	93.86	93.18
4e-5 wdm10	99.05	92.07	92.00

Table 2. Verification performance(%) of different weight decay values. Wdm10 represents the weight decay multiplier (wd_mult) parameter of the last convolution layers is 10.

Loss	LFW	CFP-FP	AgeDB
N-Softmax	97.87	90.46	86.55
Li-ArcFace(m=0.35)	99.20	94.20	92.87
Li-ArcFace(m=0.4)	99.27	94.11	93.25
Li-ArcFace(m=0.45)	99.23	94.20	93.15
Li-ArcFace(m=0.5)	99.20	94.10	92.93
Li-ArcFace(m=0.4, NS)	99.23	94.51	92.90
Li-ArcFace(m=0.45, NS)	99.27	94.39	93.22
ArcFace(m=0.4, NS)	99.27	93.93	92.72
ArcFace(m=0.45, NS)	99.27	94.00	92.98
ArcFace(m=0.5, NS)	99.18	93.86	93.18
CosFace/AM-Softmax	99.20	93.39	92.55

Table 3. Verification performance (%) of different loss functions. ArcFace (m=0.5, NS) refers to the margin of ArcFace is set as 0.5, and pre-training via N-Softmax before using ArcFace.

to the margin of ArcFace is set as 0.5, and pre-training via N-Softmax before using ArcFace.

Weight decay. Before we start the comparison of different loss functions, we first do the numerical experiment of weight decay, and finally determine that the weight decay is set to 5e-4, except the weight decay parameter of the last layer to the embedding layer being 5e-3. According to the experimental results(See Table 2), the weight decay of the last layer is more demanding.

Effect of Hyper-parameter m. In Table 3, we firstly explored the optimal setting for m of Li-ArcFace, and then we found it was between 0.4 and 0.45. We tried training model from scratch with ArcFace, but diverged after about 1600 iterations. Therefore, all the experiments with ArcFace were pre-trained by N-Softmax. Since the embedding size is smaller than 512, we also explored the optimal setting for m of ArcFace, and we found it was between 0.45 and 0.5.

Comparison with state-of-the-art loss functions. In Table 3, the difference between Li-ArcFace, ArcFace and Cosface is tiny on LFW, but all of them are obviously better than N-Softmax; On CFP-FP and AgeDB, Li-ArcFace is slightly better than ArcFace and CosFace. We have drawn the accuracy On CFP-FP and AgeDB during the training, which makes the contrast more obvious(See Figure 4). We also compared Li-ArcFace and ArcFace in the same situation pre-trained by N-Softmax. On CFP-FP, Li-ArcFace




Figure 4. The accuracy of the verification sets during training. CFP-FP is on the left, AgeDB is on the right.

Methods	Id(%)	Ver(%)	Flops
FaceNet [17]	70.49	86.47	-
CosFace [21]	82.72	96.65	-
R100,MS1MV2,ArcFace [5]	81.03	96.98	27G
R100,MS1MV2,CosFace [5]	80.56	96.56	27G
R100,MS1MV2,ArcFace,R [5]	98.35	98.48	27G
R100,MS1MV2,CosFace,R [5]	97.91	97.91	27G
MobileFaceNet [4]	-	90.16	440M
MobileFaceNet,R [4]	-	92.59	440M
AirFace,MS1M-retina	80.80	96.52	1G
AirFace,MS1M-retina,R	98.04	97.93	1G

Table 4. Identification and verification evaluation on MegaFace Challenge1. Id refers to the rank 1 face identification accuracy with 1M distractors, and Ver refers to the face verification TAR under $1e - 6$ FAR. R refers to data refinement on both probe set and 1M distractors. The list of data cleansing and the code for calculating flops are from InsightFace [1]

achieves the highest accuracy, and Li-ArcFace is still slightly better than ArcFace on AgeDB. In general, there is little difference in accuracy on the verification sets, but Li-ArcFace has better convergence and does not need pre-training stage when training the model with small embedding size.

4.2. Evaluation Results of Network Architecture and Training Tricks

We name our model that contains our network architecture and training tricks as AirFace. Under the same training data set MS1M-retina [2] and model constraints, (MS1M-retina is cleaned from [8] by [5], and [7] is the face detector and alignment tool used to pre-process the data) the accuracy of AirFace reached 88.415%@FPR=1e-8 in deepglint-light challenge of LFR19 [6, 3]. Meanwhile, we verified the performance of AirFace in MegaFace Challenge 1 compared with the previous state-of-the-art models. In Table 4, AirFace has reached incredible efficiency and performance.

5. Conclusions

In this paper, first of all, we propose a novel additive margin loss function for deep face recognition based on ArcFace. The proposed loss function solves the problem that ArcFace does not converge in training model with small embedding feature size. And it achieves the state-of-the-art results on several face verification datasets. Second, we have carefully designed an efficient network architecture and explored some useful training tricks for face recognition, which makes our model AirFace extremely efficient at both deepglint-light challenge and MegaFace Challenge 1.

References

- [1] <https://github.com/deepinsight/insightface>. 4
- [2] <https://ibug.doc.ic.ac.uk/resources/lightweight-face-recognition-challenge-workshop>. 4
- [3] <http://www.insightface-challenge.com/overview>. 4
- [4] S. Chen, Y. Liu, X. Gao, and Z. Han. Mobilefacenets: Efficient cnns for accurate real-time face verification on mobile devices. In *Chinese Conference on Biometric Recognition*, pages 428–438. Springer, 2018. 1, 2, 4
- [5] J. Deng, J. Guo, N. Xue, and S. Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019. 1, 2, 3, 4
- [6] J. Deng, J. Guo, D. Zhang, Y. Deng, X. Lu, S. Shi, and S. Zafeiriou. Lightweight face recognition challenge. 2019. 1, 4
- [7] J. Deng, J. Guo, Y. Zhou, J. Yu, I. Kotsia, and S. Zafeiriou. Retinaface: Single-stage dense face localisation in the wild. *arXiv preprint arXiv:1905.00641*, 2019. 4
- [8] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *European Conference on Computer Vision*, pages 87–102. Springer, 2016. 4
- [9] G. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 3
- [10] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 1, 2
- [11] E. Learned-Miller, G. B. Huang, A. RoyChowdhury, H. Li, and G. Hua. Labeled faces in the wild: A survey. In *Advances in face detection and facial image analysis*, pages 189–248. Springer, 2016. 1, 3
- [12] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song. SpheroFace: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 212–220, 2017. 1, 2
- [13] P. Luo, Z. Zhu, Z. Liu, X. Wang, and X. Tang. Face model compression by distilling knowledge from neurons. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016. 3
- [14] D. Miller, E. Brossard, S. Seitz, and I. Kemelmacher-Shlizerman. Megaface: A million faces for recognition at scale. *arXiv preprint arXiv:1505.02108*, 2015. 1
- [15] S. Moschoglou, A. Papaioannou, C. Sagonas, J. Deng, I. Kotsia, and S. Zafeiriou. Agedb: the first manually collected, in-the-wild age database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 51–59, 2017. 1, 3
- [16] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018. 1, 2
- [17] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. 4
- [18] 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 *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9. IEEE, 2016. 3
- [19] F. Wang, J. Cheng, W. Liu, and H. Liu. Additive margin softmax for face verification. *IEEE Signal Processing Letters*, 25(7):926–930, 2018. 1
- [20] F. Wang, X. Xiang, J. Cheng, and A. L. Yuille. Normface: 1 2 hypersphere embedding for face verification. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 1041–1049. ACM, 2017. 2
- [21] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5265–5274, 2018. 1, 4
- [22] S. Woo, J. Park, J.-Y. Lee, and I. So Kweon. Cbam: Convolutional block attention module. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–19, 2018. 3
- [23] X. Zhang, X. Zhou, M. Lin, and J. Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6848–6856, 2018. 1