



Hard-Mining Loss Based Convolutional Neural Network for Face Recognition

Yash Srivastava, Vaishnav Murali, and Shiv Ram Dubey^(✉)

Computer Vision Group, Indian Institute of Information Technology, Sri City,
Chittoor, Andhra Pradesh, India
{srivastava.y15,murali.v15,srdubey}@iiits.in

Abstract. Face Recognition is one of the prominent problems in the computer vision domain. Witnessing advances in deep learning, significant work has been observed in face recognition, which touched upon various parts of the recognition framework like Convolutional Neural Network (CNN), Layers, Loss functions, etc. Various loss functions such as Cross-Entropy, Angular-Softmax and ArcFace have been introduced to learn the weights of network for face recognition. However, these loss functions do not give high priority to the hard samples as compared to the easy samples. Moreover, their learning process is biased due to a number of easy examples compared to hard examples. In this paper, we address this issue by considering hard examples with more priority. In order to do so, We propose a Hard-Mining loss by increasing the loss for harder examples and decreasing the loss for easy examples. The proposed concept is generic and can be used with any existing loss function. We test the Hard-Mining loss with different losses such as Cross-Entropy, Angular-Softmax and ArcFace. The proposed Hard-Mining loss is tested over widely used Labeled Faces in the Wild (LFW) and YouTube Faces (YTF) datasets. The training is performed over CASIA-WebFace and MS-Celeb-1M datasets. We use the residual network (i.e., ResNet18) for the experimental analysis. The experimental results suggest that the performance of existing loss functions is boosted when used in the framework of the proposed Hard-Mining loss.

Keywords: Face Recognition · Deep learning · Loss functions · Sigmoid function · Hard-Mining loss

1 Introduction

In the past few years, the face recognition task has seen a tremendous growth in terms of the robust recognition and applications in various spheres of human lives. Face Recognition has been seen with a significant usage in multiple domains like biometric-based security tools and criminal identification system among many others. Such applications of the face recognition has lead to researchers and developers to work and design face recognition systems strongly built to work in

an unconstrained environment as its usage is expected to grow exponentially in the forthcoming years [35].

The advancements in deep learning have significantly accelerated the growth and performance of face recognition. AlexNet [14], proposed by Krizhevsky et al., is marked as the birth of the Convolutional Neural Networks (CNNs) which became a revolutionary architecture developed for the task of image classification and won the ImageNet Large Scale Challenge in 2012 [19]. Since then, many CNN based approaches have been introduced for face recognition such as, DeepFace [26], DeepID2 [24], FaceNet [20], SphereFace [17], and ArcFace [3]. The CNN based approaches [5, 10, 11, 14, 29] have shown a tremendous growth in the performance as compared to the hand-crafted features [1, 6–8, 13, 22]. The above growth was accompanied by the development of large-scale face datasets for training and testing the CNN based models, which majorly include CASIA-Webface [33], MS-Celeb-1M [9], Labeled Faces in the Wild (LFW) [12] and YouTube Faces (YTF) [31] among other face datasets. In this work, the CASIA-Webface and MS-Celeb-1M face datasets are used for training. However, the LFW and YTF face datasets are used for the testing.

The trend of CNN over time shows that the deep CNN architectures perform better as compared to the shallow networks. It was the motivation for the deeper architectures like GoogleNet [25] and ResNet [10]. The residual network shows that the performance of the deeper plain model is not improved because it is hard to optimize such model [10]. Thus, researchers also started exploring the relevance of loss functions in optimizing the deep networks. The Cross-Entropy (i.e., Softmax) loss is very widely used for optimizing the deep learning models. Recently, the work in loss functions has been quite significant with functions like SphereFace (i.e., Angular-Softmax) [17] and ArcFace [3], specially designed for the face recognition task and have shown very promising gain in the performance. Some other existing loss functions are Marginal loss [4], Soft-margin softmax loss [15], Large-margin softmax loss [18], Additive margin softmax [27], Minimum margin loss [30], Cosface: Large margin cosine loss [28], and Adaptive-Face: Adaptive margin loss [16]. Moreover, in another work, we have conducted a performance analysis of different loss functions and found that the ArcFace outperforms other losses [23].

A few attempts are also made to utilize the complexity of data in training such as the hardest positive pairs and hardest negative pairs are computed using margin sample mining loss by Xiao et al. [32]; an adaptive hard sample mining strategy it used by Chen et al. [2] to pick the hard examples in the training pair images; and an auxiliary embedding is used by Smirnov et al. [21] to pick the hard examples in mini-batches. Note that these methods try to find out the hard examples first and then use it for training. Whereas, the proposed method gives the high priority to hard examples inherently during training based on the performance of model in that iteration.

The main drawback of above mentioned loss functions is associated with its inefficiency while modelling the hard examples which lead to mis-classification. The loss due to the more number of easy examples dominates over the loss

due to the less number of hard examples. This is because while training is in progress, the number of hard examples decreases while the number of easy examples increases as network learns over iterations. In this paper, we address the above mentioned problem by giving more importance to hard examples through loss function in each iteration. We propose the Hard-Mining loss which increases the loss for the hard samples leading to high loss and decreases the loss for the easy samples leading to low loss. As a result, the average loss contains the significant contributions from the hard examples.

This paper is structured as follows: Sect. 2 proposes the Hard-Mining loss and existing losses in the Hard-Mining framework; Sect. 3 describes the experimental setup and details about the architecture and training and testing face datasets used. Section 4 presents the experimental results and comparisons; and finally, Sect. 5 concludes the paper with summarizing remarks.

2 Proposed Hard-Mining Loss

The loss functions are used in deep learning to judge the goodness of any model under given parameters. The stochastic gradient descent (SGD) optimization is widely adapted to train the Convolutional Neural Networks (CNNs). The SGD computes the gradient of loss function w.r.t. to the parameters which is used to update that parameter such that in the next iteration, the loss should decrease. Thus, the loss functions judge the performance of the designed architecture as well as guide the learning process. It is shown in introduction that most of the existing losses are not able to penalize the mis-classification efficiently caused by harder examples. In this paper, we propose the concept of Hard-Mining loss which increases the loss for harder examples and decreases the loss for easier examples such that the average loss should have the better representation of hard examples. A comparison between the Cross-Entropy loss and proposed Hard-Mining loss is presented in Fig. 1 as a function of probability of being classified in the correct class. In this section, first we present the Cross-Entropy loss, then we propose the idea of Hard-Mining loss, and finally we extend the existing losses such as Cross-Entropy, Angular-Softmax, and ArcFace in the proposed Hard-Mining framework.

2.1 Cross-Entropy Loss

The Cross-Entropy (or softmax) loss has been majorly used to judge the performance of CNN models for image classification task [10, 14]. Mathematically, the Cross-Entropy loss can be given as

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

where W is the weight matrix, b is the bias term, x_i is the i^{th} training sample, y_i is the class label for i^{th} training sample, N is the number of samples, W_j and




Fig. 1. Loss value vs Likelihood (i.e., probability for correct class) plot for the Cross-Entropy loss and Hard-Mining loss functions. Note that the Hard-Mining loss is computed on the output of Cross-Entropy loss.

W_{y_i} are the j^{th} and y_i^{th} columns of W , respectively. The Cross-Entropy loss is used as the baseline by the recent loss functions such as Angular-Softmax and ArcFace over the face recognition problem. Hence, we also use the Cross-Entropy loss as the baseline along with Angular-Softmax and ArcFace losses.

The behavior of the Cross-Entropy loss w.r.t. the probability of being classified in the correct class for an example is plotted in Fig. 1. It can be observed from this analysis that the Cross-Entropy loss gradually follows a downward slope and there is no big difference between easy and hard examples. We believe that if the probability is more than 0.5 then the loss should be minimum. Whereas, if the probability is less than 0.5 then the loss should be on higher side. This is our intuition to propose the Hard Mining Loss described next.

2.2 Hard-Mining Loss

Motivated from the fact that the loss for harder examples should be more, we propose the idea of Hard-Mining loss. The proposed Hard-Mining loss increases the loss if the probability is less than roughly 0.5, while at the same time it also decreases the loss if probability is more than 0.5 roughly. The Hard-Mining loss is defined as

$$\mathcal{L}_{HM} = \alpha \times \mathcal{L} \times \sigma(\beta \times \mathcal{L}) \quad (2)$$

where \mathcal{L} is the loss generated by any other loss function such as Cross-Entropy, Angular-Softmax, etc., α and β are the hyperparameters and σ is the sigmoid function given as:

$$\sigma(x) = \frac{1}{1 + e^{-A(x-B)}} \quad (3)$$

where A and B are the hyperparameters.

Note that the Hard-Mining operation is generic in nature, i.e., it can be used along with any existing loss function. In this paper, we use the Hard-Mining operation along with Cross-Entropy, Angular-Softmax, and ArcFace losses.

Algorithm 1. Hard-Mining Cross-Entropy Loss Algorithm

Input: Predicted class scores, Ground-truth class label, and hyper-parameters α , β , A and B .

Output: Loss generated.

1. $\mathcal{L}_{CE} \leftarrow \text{CrossEntropy}(\text{input}, \text{target}, W)$
 2. $x \leftarrow \beta \times \mathcal{L}_{CE}$
 3. $y \leftarrow A \times (x - B)$
 4. $z \leftarrow \text{Sigmoid}(y)$
 5. $\mathcal{L}_{HM_CE} \leftarrow \alpha \times x \times z$
 6. return \mathcal{L}_{HM_CE}
-

2.3 Hard-Mining Cross-Entropy Loss

As mentioned previously, the Hard-Mining concept is generic and can be used with existing losses. Primarily, we define the Hard-Mining loss with Cross-Entropy loss. The Hard-Mining Cross-Entropy loss (\mathcal{L}_{HM_CE}) is defined as

$$\mathcal{L}_{HM_CE} = \alpha * \mathcal{L}_{CE} * \sigma(\beta * \mathcal{L}_{CE}) \quad (4)$$

where α and β are the hyperparameters, σ is defined in (1), and \mathcal{L}_{CE} is the Cross-Entropy loss given in (1). Algorithm 1 shows the step-by-step instructions for the proposed Hard-Mining Cross-Entropy loss (\mathcal{L}_{HM_CE}).

The behavior of Hard-Mining operation on Cross-Entropy loss is depicted in Fig. 1. Note that the values of hyper-parameters α , β , A , and B are set to 1.5, 1.1, 35, and 0.75, respectively. It can be seen that the Hard-Mining operation increases the loss for hard examples (i.e., with less than half probability) while it decreases the loss for easy examples (i.e., with more than half probability). Our definition of hard/easy examples is relative to the probability of being classified in the correct class in a given iteration. Thus, the hard examples at the start of the training might become easy examples after training of some iterations.

Since, the Cross-Entropy is a very widely used loss function in various machine learning problems, it is paramount that we study the performance of Hard-Mining operation with loss functions specially designed for the face recognition problem. We consider two loss functions (i.e., Angular-Softmax [17] and ArcFace [3]) designed for the face recognition problem in the proposed Hard-Mining loss framework.

2.4 Hard-Mining Angular-Softmax Loss

The Hard-Mining Angular-Softmax loss (\mathcal{L}_{HM_AS}) is defined as follows:

$$\mathcal{L}_{HM_AS} = \alpha \times \mathcal{L}_{AS} \times \sigma(\beta \times \mathcal{L}_{AS}) \quad (5)$$

where α and β are the hyper-parameters, σ is given in (3), and \mathcal{L}_{AS} is the Angular-Softmax loss defined in the SphereFace model [17] and given as

$$\mathcal{L}_{AS} = -\frac{1}{N} \sum_{i=1}^N \log \left(\frac{e^{\|x_i\| \psi(\theta_{y_i,i})}}{e^{\|x_i\| \psi(\theta_{y_i,i})} + \sum_{j \neq y_i} e^{\|x_i\| \cos(\theta_{j,i})}} \right) \quad (6)$$

where x_i is the i^{th} training sample, $\psi(\theta_{y_i,i}) = (-1)^k \cos(m\theta_{y_i,i}) - 2k$ for $\theta_{y_i,i} \in [\frac{k\pi}{m}, \frac{(k+1)\pi}{m}]$, $k \in [0, m-1]$ and $m \geq 1$ is an integer controlling the size of angular margin.

2.5 Hard-Mining ArcFace Loss

ArcFace loss has been used in the recently developed ArcFace model for face recognition [3]. In a recent performance comparison study, ArcFace has been figured as the outstanding loss for face recognition [23]. The Hard-Mining ArcFace loss (\mathcal{L}_{HM_AF}) is defined as

$$\mathcal{L}_{HM_AF} = \alpha \times \mathcal{L}_{AF} \times \sigma(\beta \times \mathcal{L}_{AF}) \quad (7)$$

where α and β are the hyper-parameters, σ is given in (3), and \mathcal{L}_{AF} is the ArcFace loss [3] and given as

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

where s is the radius of the hypersphere, m is the additive angular margin penalty between x_i and W_{y_i} , and $\cos(\theta + m)$ is the margin which makes the class-separations more stringent.

3 Experimental Setup

In this section, we discuss the CNN architectures, training and testing datasets used for the experiments along with other settings like optimizers, learning rate, epochs, etc.

3.1 CNN Architectures

Several CNN architectures have been developed for different computer vision tasks. The recent trend is to utilize the power of residual learning. The ResNet model uses the residual blocks [10] which is very commonly used nowadays. In this paper, we consider ResNet architecture with 18 depth (i.e., ResNet18) for all the experiments.

3.2 Training Datasets

In our experiments, we primarily use two publicly available datasets such as CASIA-Webface [33] and MS-Celeb-1M [9] as the training datasets. The CASIA-Webface is one of the most widely adapted and available dataset used for the face recognition task. It contains 4,94,414 colored face images belonging to 10,575 different individuals. Second dataset used in our experiments is the MS-Celeb-1M dataset which consists of 1,00,000 face identities with each class containing 100 images leading to about 10M images, which are scraped from public sources. Being a humongous dataset, it contains a lot of noise and variations which impact the performance of the trained model. Hence, we use a cleaned and refined subset of the dataset as per the cleaned list provided by the ArcFace [10] authors.

3.3 Testing Datasets

We use the Labeled Faces in the Wild (LFW) [12] and Youtube Faces (YTF) [31] as the testing datasets in this paper. The LFW dataset contains 13,233 images of 5749 identities. The YTF dataset consists of 3,425 videos of 1,595 different people with images available in frame-by-frame format and retrieved through the provided meta data. Both the datasets use the standard LFW benchmark for face verification, which provide the verification accuracies over the testing dataset. These accuracies are used as the performance measure in the state-of-the-art face recognition works. Hence, we also use the accuracy as the performance measure in this paper.

3.4 Input Data and Network Settings

Following the recent trend [3,17], we use the MTCNN [34] to align the face images. The images are normalized by subtracting 127.5 from each pixel and then being divided by 128. The batch-size is kept at 64 with the initial learning rate as 0.01. The learning rate is multiplied by 0.1 at 8th, 12th and 16th epochs.

Table 1. Verification accuracies (%) using ResNet18 model over LFW and YTF face recognition testing datasets under different loss functions. The training is performed over CASIA-WebFace dataset.

Loss function	Accuracy on LFW dataset	Accuracy on YTF dataset
Cross-Entropy loss (\mathcal{L}_{CE})	95.35	91.8
Hard-Mining Cross-Entropy loss ($\mathcal{L}_{HM_{CE}}$)	96.75	93.1
Angular-Sofmax loss (\mathcal{L}_{AS})	97.12	93.9
Hard-Mining Angular-Sofmax loss ($\mathcal{L}_{HM_{AS}}$)	97.3	94.1
ArcFace loss (\mathcal{L}_{AF})	97.79	94.54
Hard-Mining ArcFace loss ($\mathcal{L}_{HM_{AF}}$)	97.9	94.67

The model is trained up to 20 epochs. The Stochastic Gradient Descent with Momentum (SGDM) is used as the optimizer to train the network. The values of hyper-parameters α , β , A , and B are empirically set to 1.5, 1.1, 35, and 0.75, respectively, in this paper.

4 Experimental Results and Observations

In order to show the effect of the proposed Hard-Mining loss, the face recognition experiments are conducted in this paper with ResNet18 model. Three existing loss functions, namely Cross-Entropy, Angular-Softmax and ArcFace, are used in the framework of the proposed Hard-Mining loss. The training is performed over the CASIA-WebFace and MS-Celeb-1M datasets and testing is performed over the LFW and YTF datasets.

The results in terms of the verification accuracies are reported in Table 1 using ResNet18 model for the CASIA-WebFace training dataset over the LFW and YTF testing datasets. It can be seen that an improvement is obtained by the Hard-Mining Cross-Entropy loss, Hard-Mining Angular-Softmax loss, and Hard-Mining ArcFace loss as compared to the Cross-Entropy loss, Angular-Softmax loss, and ArcFace loss, respectively, over both the LFW and YTF datasets.

The results in terms of the verification accuracies are reported in Table 2 using ResNet18 model for the MS-Celeb-1M training dataset over the LFW and YTF testing datasets. It is noticed from this result that the performance of Hard-Mining operation based losses is either better or comparable over LFW dataset w.r.t. the losses without Hard-Mining operation. Moreover, Hard-Mining operation is also suited with Cross-Entropy loss over YTF dataset when training is performed over MS-Celeb-1M datasets.

The experimental results suggest that increasing the loss for harder examples and decreasing the loss for easy examples in each iteration enforce the network to learn the characteristics of hard-examples as well. Overall, the proposed Hard-Mining loss is well suited for the face recognition problem along with the existing loss functions.

Table 2. Verification accuracies (%) using ResNet18 model over LFW and YTF face recognition testing datasets under different loss functions. The training is performed over MS-Celeb-1M dataset.

Loss function	Accuracy on LFW dataset	Accuracy on YTF dataset
Cross-Entropy loss (\mathcal{L}_{CE})	95.1	92.45
Hard-Mining Cross-Entropy loss ($\mathcal{L}_{HM_{CE}}$)	95.1	92.5
Angular-Softmax loss (\mathcal{L}_{AS})	96.9	94.1
Hard-Mining Angular-Softmax loss ($\mathcal{L}_{HM_{AS}}$)	97.05	93.8
ArcFace loss (\mathcal{L}_{AF})	97.6	95.1
Hard-Mining ArcFace loss ($\mathcal{L}_{HM_{AF}}$)	98	94.9

5 Conclusion

In this paper, a concept of Hard-Mining loss is proposed which increases the loss for hard examples being mis-classified and decreases the loss for easy examples. By doing so, we enforce the network to learn the characteristics of hard examples. The proposed concept is generic in nature and can be used with any existing loss function. We have tested the proposed Hard-Mining loss with Cross-Entropy, Angular-Softmax and ArcFace losses. The experiments are performed over CASIA-WebFace and MS-Celeb-1M training datasets and LFW and YTF testing datasets using ResNet18 model. It is observed from the experiments that the proposed Hard-Mining loss boosts the performance of existing losses in most of the cases.

References

1. Chakraborti, T., McCane, B., Mills, S., Pal, U.: Loop descriptor: local optimal-oriented pattern. *IEEE Signal Process. Lett.* **25**(5), 635–639 (2018)
2. Chen, K., Chen, Y., Han, C., Sang, N., Gao, C., Wang, R.: Improving person re-identification by adaptive hard sample mining. In: 2018 25th IEEE International Conference on Image Processing (ICIP), pp. 1638–1642. IEEE (2018)
3. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: additive angular margin loss for deep face recognition. arXiv preprint [arXiv:1801.07698](https://arxiv.org/abs/1801.07698) (2018)
4. 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)
5. Dubey, S.R., Roy, S.K., Chakraborty, S., Mukherjee, S., Chaudhuri, B.B.: Local bit-plane decoded convolutional neural network features for biomedical image retrieval. *Neural Comput. Appl.* **32**(11), 7539–7551 (2019). <https://doi.org/10.1007/s00521-019-04279-6>
6. Dubey, S.R., Singh, S.K., Singh, R.K.: Rotation and illumination invariant interleaved intensity order-based local descriptor. *IEEE Trans. Image Process.* **23**(12), 5323–5333 (2014)
7. Dubey, S.R., Singh, S.K., Singh, R.K.: Local wavelet pattern: a new feature descriptor for image retrieval in medical CT databases. *IEEE Trans. Image Process.* **24**(12), 5892–5903 (2015)
8. Dubey, S.R., Singh, S.K., Singh, R.K.: Multichannel decoded local binary patterns for content-based image retrieval. *IEEE Trans. Image Process.* **25**(9), 4018–4032 (2016)
9. Guo, Y., Zhang, L., Hu, Y., He, X., Gao, J.: MS-Celeb-1M: a dataset and benchmark for large-scale face recognition. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016*. LNCS, vol. 9907, pp. 87–102. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46487-9_6
10. 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)
11. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7132–7141 (2018)

12. Huang, G.B., Ramesh, M., Berg, T., Learned-Miller, E.: Labeled faces in the wild: a database for studying face recognition in unconstrained environments. Technical report, 07–49, University of Massachusetts, Amherst, October 2007
13. Kou, Q., Cheng, D., Zhuang, H., Gao, R.: Cross-complementary local binary pattern for robust texture classification. *IEEE Signal Process. Lett.* **26**(1), 129–133 (2018)
14. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems*, pp. 1097–1105 (2012)
15. Liang, X., Wang, X., Lei, Z., Liao, S., Li, S.Z.: Soft-margin softmax for deep classification. In: Liu, D., Xie, S., Li, Y., Zhao, D., El-Alfy, E.S. (eds.) *ICONIP 2017*. LNCS, vol. 10635, pp. 413–421. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70096-0_43
16. Liu, H., Zhu, X., Lei, Z., Li, S.Z.: AdaptiveFace: adaptive margin and sampling for face recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 11947–11956 (2019)
17. 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)
18. Liu, W., Wen, Y., Yu, Z., Yang, M.: Large-margin softmax loss for convolutional neural networks. In: *ICML*, p. 7 (2016)
19. Russakovsky, O., et al.: ImageNet large scale visual recognition challenge. *Int. J. Comput. Vis.* **115**(3), 211–252 (2015)
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. Smirnov, E., Melnikov, A., Oleinik, A., Ivanova, E., Kalinovskiy, I., Luckyanets, E.: Hard example mining with auxiliary embeddings. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 37–46 (2018)
22. Song, T., Xin, L., Gao, C., Zhang, G., Zhang, T.: Grayscale-inversion and rotation invariant texture description using sorted local gradient pattern. *IEEE Signal Process. Lett.* **25**(5), 625–629 (2018)
23. Srivastava, Y., Murali, V., Dubey, S.R.: A performance comparison of loss functions for deep face recognition. arXiv preprint [arXiv:1901.05903](https://arxiv.org/abs/1901.05903) (2019)
24. Sun, Y., Chen, Y., Wang, X., Tang, X.: Deep learning face representation by joint identification-verification. In: *Advances in Neural Information Processing Systems*, pp. 1988–1996 (2014)
25. Szegedy, C., et al.: Going deeper with convolutions. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–9 (2015)
26. Taigman, Y., Yang, M., Ranzato, M., Wolf, L.: DeepFace: closing the gap to human-level performance in face verification. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1701–1708 (2014)
27. Wang, F., Cheng, J., Liu, W., Liu, H.: Additive margin softmax for face verification. *IEEE Signal Process. Lett.* **25**(7), 926–930 (2018)
28. Wang, H., et al.: 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)
29. Wang, Y., Ward, R.K., Wang, Z.J.: Coarse-to-fine image dehashing using deep pyramidal residual learning. *IEEE Signal Process. Lett.* **26**, 1295–1299 (2019)

30. Wei, X., Wang, H., Scotney, B., Wan, H.: Minimum margin loss for deep face recognition. arXiv preprint [arXiv:1805.06741](https://arxiv.org/abs/1805.06741) (2018)
31. Wolf, L., Hassner, T., Maoz, I.: Face recognition in unconstrained videos with matched background similarity. IEEE (2011)
32. Xiao, Q., Luo, H., Zhang, C.: Margin sample mining loss: a deep learning based method for person re-identification. arXiv preprint [arXiv:1710.00478](https://arxiv.org/abs/1710.00478) (2017)
33. Yi, D., Lei, Z., Liao, S., Li, S.Z.: Learning face representation from scratch. arXiv preprint [arXiv:1411.7923](https://arxiv.org/abs/1411.7923) (2014)
34. Zhang, K., Zhang, Z., Li, Z., Qiao, Y.: Joint face detection and alignment using multitask cascaded convolutional networks. IEEE Signal Process. Lett. **23**(10), 1499–1503 (2016)
35. Zhou, Y., Liu, D., Huang, T.: Survey of face detection on low-quality images. In: 2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018), pp. 769–773. IEEE (2018)