



FIE-GAN: Illumination Enhancement Network for Face Recognition

Zhuo Wang, Weihong Deng^(✉), and Jiancheng Ge

Beijing University of Posts and Telecommunications, Beijing 100876, China
whdeng@bupt.edu.cn

Abstract. Low-light face images not only are difficult to be perceived by humans but also cause errors in automatic face recognition systems. Current methods of image illumination enhancement mainly focus on the improvement of the visual perception, but less study their applications in recognition systems. In this paper, we propose a novel generative adversarial network, called FIE-GAN to normalize the lighting of face images while try to retain face identity information during processing. Besides the perceptual loss ensuring the consistency of face identity, we optimize a novel histogram controlling loss to achieve an ideal lighting condition after illumination transformation. Furthermore, we integrate FIE-GAN as data preprocessing in unconstrained face recognition systems. Experiment results on IJB-B and IJB-C databases demonstrate the superiority and effectiveness of our method in enhancing both lighting quality and recognition accuracy.

Keywords: Illumination enhancement · Generative adversarial networks · Face recognition

1 Introduction

Face images in low-light conditions are poor visibility and low contrast, which may bring challenges not only for human visual perception but also for some intelligent systems such as face recognition. Recently, deep face recognition has shown excellent performance in the normal-light range because most face datasets [1, 8, 33] mainly contain normal-light face images for training. Therefore, models can learn to gain a powerful capability of identity feature extraction in a normal-light environment.

However, in real uncontrolled testing scenarios [23, 31], there also exist many low-light face images that may suffer mistakes during face recognition. A common strategy to solve this problem is brightening these low-light images to a normal range. To the best of our knowledge, most existing researches [13, 21, 26] on illumination enhancement mainly focus on promoting the visual perception of landscape images. The research on the application of illumination enhancement in face recognition has been less investigated for a long time.




Fig. 1. A face recognition system may be trapped in wrong judgment when encountering low-light face images. Our FIE-GAN can brighten low-light images to a normal range, which is beneficial for unconstrained face recognition.

Furthermore, there is a major obstacle to the training data: 1) It's difficult and inefficient to collect a paired training dataset of low/normal-light images under a regular situation. 2) Synthesizing corrupted images from clean images may usually not be realistic enough. Therefore, we quit seeking paired training data, but propose a more available method of achieving face illumination enhancement in unconstrained environments with unpaired supervision.

Inspired by [35] with unsupervised image-to-image translation, we adopt the generative adversarial network (GAN) to fetch close the gap between low and normal light distribution without paired images in the wild environment. GAN is first introduced by Goodfellow et al. [6]. The training strategy is like a min-max two-player game, where the generator tries to confuse the discriminator while the discriminator tries to distinguish between generated images and real images. In our method, we utilize this adversarial thought to achieve the conversion of images from low light toward a normal range.

Therefore, we propose FIE-GAN to achieve illumination enhancement and retain identity information during transformation. FIE-GAN disentangles illumination information from content information. Then, FIE-GAN can generate enhanced face images by fusing illumination information of normal-light images and content information of low-light images. As shown in Fig. 1, brightening low-light images to a normal range is beneficial for identity information extraction due to the powerful capability of the face recognition model in the normal-light range. Lastly, to investigate the effect of illumination enhancement on face recognition, we try deploying FIE-GAN in a face recognition system as a preprocessing to convert low-light images to normal-light ones. Overall, our contributions can be summarized as follows:

- A Face Illumination Enhancement based on Generative Adversarial Network (FIE-GAN) is proposed to bright low-light face images to a normal range. To avoid overfitting, a large unpaired training set is built up and the corresponding training strategy on unpaired data is adopted.
- Histogram vector and corresponding loss function are proposed to control image illumination transformation between different light conditions by a more flexible approach. Besides, Identity information is also appropriately

preserved during this transformation. The qualitative and quantitative analysis prove the superiority of our method commonly.

- We investigate the influence of face illumination enhancement in face recognition by employing FIE-GAN as an image preprocessing step in face recognition systems. Experiment results on IJB-B and IJB-C prove the rationality and effectiveness.

2 Related Works

We briefly review previous works from three aspects of image illumination enhancement, image-to-image translation and deep face recognition as follow.

Images Illumination Enhancement. Low-light face images in unconstrained environments are difficult to be distinguished by people and also become a challenge for face recognition systems. Gamma correction [24] and histogram equalization [11] are two representative traditional methods to achieve image illumination transformation. In deep learning time, most existing research works [13, 21, 26] on illumination enhancement focus on the visual improvement of landscape images. However, their training data is usually required to be paired, which is difficult or even impractical to be acquired for face images in the wild.

Image-to-Image Translation. Generative adversarial networks (GANs) [6] have demonstrated the powerful capabilities of adversarial learning in image synthesis [7, 17, 18] and translation [3, 12, 35]. The Image-to-image translation is a more general problem for image processing such as style transfer [15, 35], super resolution [14, 20] and face editing [2, 32]. Pix2pix[19] first defined this task and obtained impressive results on paired data. However, paired data are usually unavailable for most tasks. CycleGAN [2] proposed an unpaired image-to-image translation method to learn the mapping function between different domains.

Deep Face Recognition. Face recognition (FR) has been the prominent biometric technique for identity authentication [29]. This success is attributed to the development of loss function [4, 27, 28], network architectures [9, 30] and the availability of large-scale training datasets [1, 8, 33]. Data cleaning [34] and data augmentation [5] can also improve the performance of FR, which proves the importance of data preprocessing in FR and the advantages of deep methods in image operation.

Therefore, in this paper, we propose an unpaired image-to-image method for illumination enhancement. Furthermore, the effectiveness of this method is investigated in an FR system.

3 Methodology

Our goal is to train a generator G for converting low-light images to a normal-light range by using only unpaired images. Furthermore, to apply our method to face recognition, the identity information of face images should be maintained during the transformation. Therefore, the architecture and loss function will be described in detail below.

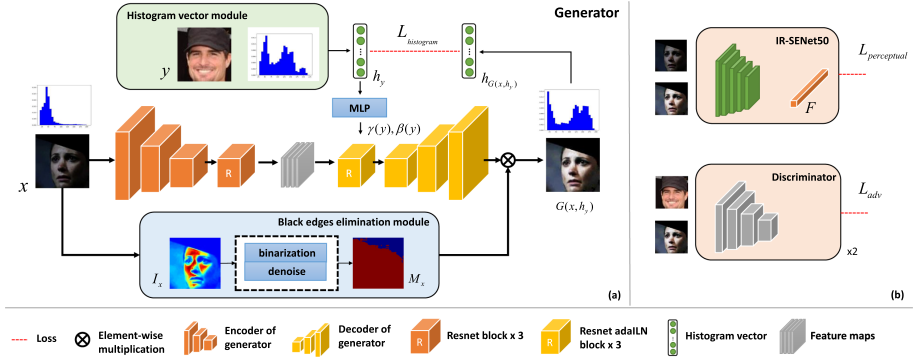


Fig. 2. The architecture of FIE-GAN. In part (a), the generator with integrated modules transforms low-light face images to normal ones. In part (b), the pre-trained expert network IR-SENet50 [4] is used to extract identity information, and two discriminators are applied to distinguish between the real image and generated image. All of the notations are listed at the bottom.

3.1 Architecture

Our framework consists of a generator G , two discriminators $\{D_1, D_2\}$ and an expert network. We integrate the black edges elimination module and the histogram vector module into the generator to ensure the effect of generated images.

Generator. The generator is shown in Fig. 2(a), where x from low-light face set and y from normal-light face set. Low-light face image x is fed into the backbone and generate corresponding normal-light face $G(x, h_y)$ under the control of h_y that is the histogram vector of normal-light image y .

Histogram vector is proposed to indicate the illumination distribution at pixel level. The illumination range is divided into m pieces with the same interval and the histogram vector of image y is defined as a m -dimensional vector $h_y = [h_y^{(1)} \dots h_y^{(m)}]$, $h_y^{(i)} = \frac{num_{I_p}}{\sum num_{I_p}}$, $I_p \leq I_{i+1}$, $0 \leq i \leq m$. I_p is the illumination of a pixel. num_{I_p} represents the number of pixels whose illumination is between I_i and I_{i+1} . Therefore, the histogram vector h_y represents the proportion of pixels in every piece interval.

In our method, content information of face images can be divided into face identity and illumination. Illumination is also regarded as style information. Inspired by [19], residual blocks with AdaLIN are equipped to implement style transfer while adaptively changing or keeping the content information to separate illumination information from identity information. Thus, the content information of low-light image x and the style information of normal-light image y are fused to generate corresponding normal-light image $G(x, h_y)$. Specifically, the histogram vector of normal-light face image y guide this style transfer to brighten low-light image x by controlling the affine transform parameters in AdaLIN: $\gamma(y), \beta(y) = MLP(h_y)$, and retaining face identity information.

Moreover, due to the defect of face alignment, low-light faces in public face datasets often contain black edges, which may affect the result of illumination enhancement. The generator will try to promote the illumination of pixels in black edges, but ignore the face region, because the former is largely darker than the latter. Therefore, we claim a mechanism of masking to ignore black edges in a pixel-wise multiplication way. As shown in Fig. 2(a), I_x is the heat map of x , representing the illumination distribution on pixels. M_x is a mask to ignore black edges in low-light face images by excluding consecutive zero-pixel regions. Furthermore, its impact is mitigated on the generated result by blocking the gradient propagation on black edges.

To sum up, our generator can convert input low-light face image x to a normal range in an image-to-image way. And the transformation process is under the control of h_y that is the histogram vector of normal-light face image y .

Discriminator. To better distinguish the real image y and the generated images $G(x, h_y)$, we adopt a multi-scale model as the other competitor in this adversarial game, which contains two discriminators with different sizes of the convolution kernel. Especially, multi-scale discriminators possess multiple receptive fields, which means they can cover more sufficient cues to split real images and generated images.

Expert Network. An expert network is a model that contains abundant prior knowledge. A pre-trained face feature extraction network IR-SENet50 [4] is employed as the expert network in our pipeline to ensure the consistency of identity information during the transformation. On the other hand, to obtain continuous prior guidance, we freeze the layers of the expert network in our training and evaluating period.

3.2 Loss Function

The aim of FIE-GAN is to brighten the low-light face image to a normal one. Thus, besides the basic adversarial loss, there also contain additional loss functions to preserve the identity information or assist network convergence.

Adversarial Loss. Adversarial loss is employed to minimize the gap between the distribution of source and target domain. The least-squares loss [22] is utilized for stable training and higher generated image quality,

$$\begin{aligned} L_{adv}^D &= \sum_{i=1}^2 E[D_i(G(x, h_y))^2 + (D_i(x) - 1)^2], \\ L_{adv}^G &= \sum_{i=1}^2 E[(D_i(G(x, h_y)) - 1)^2], \\ L_{adv} &= L_{adv}^D + L_{adv}^G, \end{aligned} \quad (1)$$

where x is the input low-light face image and y represents the real normal-light face image. h_y is the histogram vector of image y which tries to promote the generator G to transform x to normal-light range and confuse the discriminator D_1, D_2 simultaneously.

Perceptual Loss. Inspired by [16], the perceptual loss is introduced to retain identity information during illumination enhancement. High-dimension feature representation $F(\cdot)$ is extracted by a pre-trained expert network. So the loss can be formulated as:

$$L_{perceptual} = \|F(G(x, h_y)) - F(x)\|_1. \quad (2)$$

Histogram Loss. Histogram loss is an $L2$ loss in the space of the histogram vector for ensuring the similarity in histogram distribution between generated normal-light images and real ones. $H(\cdot)$ indicates the operation of computing histogram vector,

$$L_{histogram} = \|H(G(x, h_y)) - h_y\|_2^2. \quad (3)$$

Specially, $G(x, h_y)$ is the transformation result of low-light image x under the control of the histogram vector of normal-light image y . Thus, a similar brightness distribution at pixel level is expected between the generated normal-light image $G(x, h_y)$ and real one y .

Reconstruction Loss. Reconstruction loss is claimed to guarantee the input image can be reconstructed under the control of its histogram vector,

$$L_{reconstruction} = \frac{1}{W \times H \times C} \|G(y, h_y) - y\|_1, \quad (4)$$

which means that the generated result $G(y, h_y)$ under the control of histogram vector h_y should retain the consistency with the input normal-light image y . The above description represents a reconstruction process for a real normal-light image.

Full Objective. Finally, we acquire the overall loss function which is defined as follows:

$$\begin{cases} L_D = \lambda_1 L_{adv}^D, \\ L_G = \lambda_1 L_{adv}^G + \lambda_2 L_{perceptual} + \lambda_3 L_{histogram} + \lambda_4 L_{reconstruction}, \end{cases} \quad (5)$$

where λ_1 , λ_2 , λ_3 and λ_4 are hyper-parameters to balance the importance of different loss functions. We train our discriminator and generator in turn by optimizing L_D and L_G , respectively.

3.3 Deployment in Face Recognition

To explore the influence of face illumination enhancement toward face recognition (FR), we deploy FIE-GAN in the FR system, as shown in Fig. 3. FIE-GAN as an image preprocessing step is placed before the feature extraction step.

In the training phase, the FR model can acquire a more powerful capability of feature representation for images in the normal-light range. This judgment can be explained by the following two reasons: 1) FIE-GAN brightens low-light face images into a normal range, which is equal to augment training data in

normal-light range and promotes model by more learning on these adequate data; 2) This normalization operation also makes the brightness distribution of the dataset concentrated to a limited range, which means under this setting, FR model can focus its parameters to learn how to represent identity information instead of how to fit a few difficult cases such as low-light images.

Therefore, in the evaluation phase, the FR model can capture more sufficient features from brightened low-light face images. At the same time, the recognition accuracy is expected to be improved.




Fig. 3. The place of FIE-GAN in the face recognition system.

4 Experiments

4.1 Experimental Settings

In this section, we firstly train FIE-GAN to acquire a face illumination enhancement model. Then we integrate this pre-trained FIE-GAN as a preprocessing module into an unconstrained face recognition system to further evaluate its effect on face recognition (FR).

Databases. To train FIE-GAN in an unconstrained environment, a large-scale training set is built up, which contains unpaired low and normal light face images. Specifically, we collect 96,544 low-light face images from MS1M-RetinaFace [4] and 137,896 normal-light face images from CASIA-WebFace [33]. To exploit the effect of our method as a preprocess module in the FR system, we adopt CASIA-WebFace under different preprocessing methods as our training data to train multiple FR models. These FR models are tested on large-scale face image datasets such as IJB-B [31] and IJB-C [23], to evaluate their performance.

Implementation Details. We apply the Y value in YUV color space as the illumination of a single pixel. Image illumination is defined as the mean illumination of all pixels. Especially, images with light below 50 are regarded as low light while images with light between 80 and 120 are regarded as normal light. Therefore, illumination enhancement is equivalent to the transformation of low-light images to normal-light ones. Histogram vector guide this transformation and its dimension m is set to 25 in our experiments.

Our model training is composed of FIE-GAN training and FR model training.

When training FIE-GAN, we crop and resize all face images to 112×112 as input images. We keep the learning rate equal to 0.0001, set the batch size to 16, and finish training at 10K iteration by iteratively minimizing the discriminator and generator loss function with Adam optimization. According to experimental experience, we set hyper-parameters of the optimizer as follows: $\beta_1 = 0.5$, $\beta_2 = 0.999$. The values of λ_1 , λ_2 , λ_3 and λ_4 are 1, 100, 100 and 10, respectively.

When training the FR model, to compare the effectiveness of different illumination enhancement methods, we utilize the dataset whose low-light face images are preprocessed by these different methods as training sets. The batch size is equal to 512. The learning rate starts from 0.1 and is divided by 10 at 100,000, 160,000 and 220,000 iterations, respectively. We employ ResNet-34 as the embedding network and arface as the loss function. According to [4], the feature scale and angular margin m are set as 64 and 0.5 respectively. SGD optimization is adopted in the training procedure of FR models.

4.2 Visual Perception Results

To investigate the effectiveness of our method, we compare FIE-GAN with several commonly used algorithms of illumination enhancement, including Gamma Correction [24], Histogram Equalization [11] and CycleGAN [2]. Especially, we implement histogram equalization on the Y channel of YUV. Therefore, the histogram equalization method is also marked as YUV in the paper.

Qualitative Analysis. As shown in Fig. 4, the results of Gamma are generally darker compared with others. The YUV method generates some over-exposure artifacts and causes some color information distorted. CycleGAN generates unsatisfactory visual results in terms of both naturalness and definition. In contrast, FIE-GAN successfully not only realizes illumination enhancement in face regions but also preserves the identity information of face images. The results demonstrate the superiority of our method to realize image-to-image illumination enhancement in the wild environment.




Fig. 4. Comparison of illumination enhancement results. All source images are from unconstrained dataset IJB-C.

Quantitative Analysis. For measuring the quality of generated images, the Inception Score (IS) [25] and the Fréchet Inception Distance (FID) [10] are adapted in our experiments. The former is applying an Inception-v3 network pre-trained on ImageNet to compare the conditional distribution of generated samples. The latter is an improvement on the IS by comparing the statistics of generated samples to real samples. As shown in Table 1, our method acquires the highest Inception score and the lowest FID, which proves the generated images by FIE-GAN are high-quality and identity-preserved again. On the other hand, it is worth noting that the Inception score of source images is even lower than generated images, which indicates low-light images are difficult for feature extraction in not only face recognition networks but also the pre-trained Inception-v3 network there. Therefore, we can conclude that illumination enhancement before identity information extraction is necessary for low-light face images.

Table 1. The quantitative evaluation on the quality of the generated images.

Method	IS \uparrow	FID \downarrow
Source	2.62	–
Gamma	3.03	64.94
YUV	3.05	64.15
CycleGAN	2.82	74.83
FIE-GAN	3.18	47.74

4.3 Face Recognition Results

We further try employing our FIE-GAN in a complete face recognition (FR) system as a preprocess module and explore its effect on face recognition.

Table 2. The 1:1 verification TAR ($@FAR = 1e-5, 1e-4, 1e-3$) on the IJB-B and IJB-C dataset.

Method	IJB-B			IJB-C		
	1e-5	1e-4	1e-3	1e-5	1e-4	1e-3
Baseline	0.6432	0.7860	0.8771	0.7354	0.8273	0.9034
Gamma	0.6314	0.7862	0.8754	0.7211	0.8317	0.9044
YUV	0.6465	0.7835	0.8707	0.7290	0.8236	0.8988
CycleGAN	0.6369	0.7826	0.8692	0.7249	0.8248	0.8959
FIE-GAN	0.6608	0.7899	0.8844	0.7372	0.8360	0.9082

Results on IJB-B and IJB-C. To evaluate the performance of different FR models comprehensively, we apply IJB-B [31] and IJB-C [23] as test sets. They are significantly challenging benchmarks in unconstrained face recognition.

We employ CASIA-WebFace as training data and the ResNet34 as the embedding network. Different illumination enhancement methods, shown in Table 2 are applied as a pretreatment to brighten low-light images in training and testing data. The baseline model is trained by source images without any preprocessing. The experiment results shown in Table 2 prove the effectiveness of our method on illumination enhancement and identity information preservation. However, other methods fail to exceed the baseline model completely, which may be due to the loss of identity information during illumination enhancement.

4.4 Ablation Study

To verify the superiority of FIE-GAN as well as the contribution of each component, we train two incomplete models in terms of without the black edges elimination module and without the histogram vector module. Figure 5 illustrates their visual perception for a comparison. The second-row results prove black edges may cause the model to enhance the illumination of black background mistakenly, but ignore more significant face regions. From the third row results, we can conclude the importance of the histogram vector and corresponding loss on the guiding model to achieve illumination enhancement.




Fig. 5. Ablation study of the contribution of each component: (a) Source images, (b) Results w/o black edge elimination, (c) Results w/o histogram vector guidance, (d) Results with complete FIE-GAN.

5 Conclusion

In this paper, we propose a novel FIE-GAN model to achieve face image illumination enhancement in the condition of an unconstrained environment. The histogram vector we defined can guide image illumination transformation more flexibly and precisely. Multiple loss functions are utilized to ensure the effectiveness of illumination enhancement and the consistency of identity information during the transformation. After acquiring excellent visual perception results, we further investigate our FIE-GAN model as preprocessing in face recognition systems. As future work, illumination enhancement based on GAN can be assigned to other face-related missions to further test its validity.

References

1. Cao, Q., Shen, L., Xie, W., Parkhi, O.M., Zisserman, A.: Vggface2: a dataset for recognising faces across pose and age. In: 2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018), pp. 67–74. IEEE (2018)
2. Chang, H., Lu, J., Yu, F., Finkelstein, A.: Pairedcyclegan: asymmetric style transfer for applying and removing makeup. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 40–48 (2018)
3. Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S., Choo, J.: Stargan: unified generative adversarial networks for multi-domain image-to-image translation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 8789–8797 (2018)
4. Deng, J., Guo, J., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition (2018)
5. Fang, H., Deng, W., Zhong, Y., Hu, J.: Generate to adapt: resolution adaption network for surveillance face recognition. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) ECCV 2020. LNCS, vol. 12360, pp. 741–758. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58555-6_44
6. Goodfellow, I., et al.: Generative adversarial nets. In: Advances in Neural Information Processing Systems, pp. 2672–2680 (2014)
7. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.: Improved training of wasserstein gans. arXiv preprint [arXiv:1704.00028](https://arxiv.org/abs/1704.00028) (2017)
8. 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
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. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. arXiv preprint [arXiv:1706.08500](https://arxiv.org/abs/1706.08500) (2017)
11. Hum, Y.C., Lai, K.W., Mohamad Salim, M.I.: Multiobjectives bihistogram equalization for image contrast enhancement. *Complexity* **20**(2), 22–36 (2014)
12. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1125–1134 (2017)
13. Jenicek, T., Chum, O.: No fear of the dark: image retrieval under varying illumination conditions. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 9696–9704 (2019)
14. Jo, Y., Yang, S., Kim, S.J.: Investigating loss functions for extreme super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pp. 424–425 (2020)
15. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9906, pp. 694–711. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46475-6_43
16. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: European Conference on Computer Vision (2016)

17. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. arXiv preprint [arXiv:1710.10196](https://arxiv.org/abs/1710.10196) (2017)
18. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4401–4410 (2019)
19. Kim, J., Kim, M., Kang, H., Lee, K.: U-gat-it: unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation. arXiv preprint [arXiv:1907.10830](https://arxiv.org/abs/1907.10830) (2019)
20. Ledig, C., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4681–4690 (2017)
21. Lore, K.G., Akintayo, A., Sarkar, S.: Llnet: a deep autoencoder approach to natural low-light image enhancement. *Pattern Recogn.* **61**, 650–662 (2017)
22. Mao, X., Li, Q., Xie, H., Lau, R.Y.K., Wang, Z., Smolley, S.P.: Least squares generative adversarial networks (2016)
23. Maze, B., et al.: Iarpa janus benchmark-c: Face dataset and protocol. In: 2018 International Conference on Biometrics (ICB), pp. 158–165. IEEE (2018)
24. Poynton, C.: *Digital video and HD: Algorithms and Interfaces*. Elsevier (2012)
25. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. arXiv preprint [arXiv:1606.03498](https://arxiv.org/abs/1606.03498) (2016)
26. Shen, L., Yue, Z., Feng, F., Chen, Q., Liu, S., Ma, J.: Msr-net: Low-light image enhancement using deep convolutional network. arXiv preprint [arXiv:1711.02488](https://arxiv.org/abs/1711.02488) (2017)
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, M., Deng, W.: Deep face recognition: a survey. CoRR abs/1804.06655 (2018). <http://arxiv.org/abs/1804.06655>
30. Wang, Q., Guo, G.: Ls-cnn: characterizing local patches at multiple scales for face recognition. *IEEE Trans. Inf. Forensics Secur.* **15**, 1640–1653 (2019)
31. Whitelam, C., 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)
32. Wu, R., Zhang, G., Lu, S., Chen, T.: Cascade ef-gan: progressive facial expression editing with local focuses. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 5021–5030 (2020)
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, Y., et al.: Global-local gcn: Large-scale label noise cleansing for face recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7731–7740 (2020)
35. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2223–2232 (2017)