The International Journal of Engineering and Science (IJES)
|| Volume || 15 || Issue 1 || January || Pages || PP 21-29 || 2026 ||
ISSN (e): 2319 – 1813 ISSN (p): 2319 – 1805

Intelligent Facial Image Beautification System Using Modified CycleGAN

Peng-Chen Chang Chien,Chou-Chen Wang
Department of Electronic Engineering, I-Shou University, Kaohsiung, Taiwan
Corresponding Author: Chou-ChenWang

----------ABSTRACT----------
Facial beautification enhancement is a widely used function in image processing, often requiring manual parameter adjustment to achieve optimal results. To provide a more intelligent and automated solution, this study proposes an enhanced facial beautification system based on generative adversarial networks (GANs). While the Cycle-Consistent Generative Adversarial Network (CycleGAN) effectively performs unpaired image-to-image translation, it suffers from contour and detail loss. To address this limitation, we propose a Modified CycleGAN (MCycleGAN) that integrates the advantages of both CycleGAN and PairedCycleGAN while maintaining the unpaired training framework. The proposed MCycleGAN employs combinations of Sobel, Gaussian, and Bilateral filters to extract edge-aware image features, leading to improved preservation of facial contours and skin texture. Experimental results show that MCycleGAN improves the average intersection-over-union (IOU) by 7% compared to CycleGAN, and subjective human evaluation indicates a 16.5% higher preference for MCycleGAN-generated images.
KEYWORDS–*Facial enhancement, Generative adversarial networks (GAN), intersection-over-union (IOU)*

Date of Submission: 26-01-2026 Date of acceptance: 10-02-2026

I. INTRODUCTION

Facial image beautificationhas been an important topic in computer vision and image processing in recent years. Traditional methods typically rely on filters or pixel value manipulation to enhance facial appearance. However, such approaches often suffer from poor generalization when fixed parameters cannot adapt to different facial structures and lighting conditions.With the rapid advancement of artificial intelligence (AI), intelligent facial image beautificationsystems can now leverage machine learning to achieve more natural and adaptive enhancement. Since the introduction of generative adversarial networks (GANs) by Goodfellow et al. [1], GANs have become powerful tools for solving complex image translation tasks, including style transfer, super-resolution, and realistic image synthesis. Among them, the cycle-consistent generative adversarial network (CycleGAN) proposed by [2-3] enables unpaired image-to-image translation and has been successfully applied to facial beautification. However, CycleGAN often produces blurred edges and loss of fine details, as shown in Fig. 1, where freckles are removed but contours such as eyelids and teeth appear blurred.

To overcome these challenges without increasing computational complexity, we propose a modified CycleGAN (MCycleGAN) that incorporates Sobel, Gaussian, and Bilateral filtering into the loss function to improve edge preservation and maintain visual fidelity in beautified results.

The remainder of this paper is organized as follows. In Section II we briefly review background and related works. Section III elaborates the proposed Modified CycleGAN Method. The experimental results are presented in Section IV. Finally, Section V summarizes our conclusions.

(a) (b)

Fig.1Comparison of unsupervised learning (a) original image (b)facial enhancementusingCycleGAN.

II. BACKGROUND AND RELATED WORKS

Image-to-image translation is a fundamental problem in visual computing, and numerous methods have been developed to address the issue [2–6]. After the introduction of GANs in 2014, a wide range of algorithms emerged for style transfer, image synthesis, and even direct generation of images in specified styles. Among these methods, the image style transfer framework of CycleGAN proposed by Zhu et al. [2] represented a significant improvement, particularly in applications such as intelligent beauty enhancement.

2.1 Architectureof CycleGAN system

The generator architecture of CycleGAN is based on the network design proposed by Johnson et al. [5], which has demonstrated strong performance in image translation and super-resolution tasks. The generator consists of two convolutional layers with stride 2, several residual blocks, and two convolutional layers with stride 1/2. For input images with a resolution of 256×256, nine residual blocks are used, whereas six blocks are applied for 128×128 inputs.The purpose of the discriminator is to distinguish generated images from real ones, which essentially corresponds to a classification task. Therefore, CycleGAN adopts the 70×70 PatchGAN discriminator introduced in [6], which is constructed using a simple five-layer convolutional architecture. PatchGAN is designed to classify whether each 70×70 receptive field (RF) within an image is real or fake, and due to its fully convolutional structure, it can operate on images of arbitrary resolution.

The choice of the 70×70 receptive field is based on empirical findings from Isola et al. [6], showing that this receptive field size is most effective for capturing and distinguishing crucial information such as texture, color, and edge characteristics within the image.Image-to-image translation involves learning a mapping between two different domains. Suppose we have two domains, A and B, with training samples $\{a_i\}_{i=1}^{N}$ where each $a_i$in $a_i \in A$, and similarly $\{b_j\}_{j=1}^{N}$ where each $b_i \in B$.These data distributions for two domains can be expressed as $a{\sim}p_{data}(a)$and $b{\sim}p_{data}(b)$, respectively.The goal of training a CycleGAN model is to learn a bidirectional mapping between these two domains. This involves two generators,G_{AB} and G_{BA}, which learn to map samples from one domain to the other. In other words, these two mappings are $G_{AB} : A \rightarrow B$and $G_{BA} : B \rightarrow A$, respectively.

In addition, two discriminators,$D_{AB}$and D_{BA}, are trained simultaneously. The discriminator D_{AB} learns to classify real samples from domain {a}and generated samplesG_{AB} while $\{G_{AB}(a)\}$ learns to classify real samples from domain{b}and generated samples $\{G_{BA}(b)\}$. Their adversaries, the generators $G_{AB}$and G_{BA}, use the feedback from the discriminators to improve the realism of the generated outputs, making $\{G_{AB}(a)\}$and $\{G_{BA}(b)\}$ increasingly closer to the target domains.The overall training framework of CycleGAN is illustrated in Fig. 2, and the associated loss function can be expressed as follows:

$$L_{GAN}(G_{AB}, D_{AB}, A, B) = \mathbb{E}_{b\sim Pdata(b)}[\log D_{AB}(b)] + \mathbb{E}_{a\sim Pdata(a)}\left[1 - log D_{AB}\left(G_{AB}(a)\right)\right] \quad (1)$$

$$L_{GAN}(G_{BA}, D_{BA}, B, A) = \mathbb{E}_{b\sim Pdata(a)}[\log D_{BA}(a)] + \mathbb{E}_{a\sim Pdata(b)}[1 - \log D_{BA}(G_{BA}(b))] \quad (2)$$

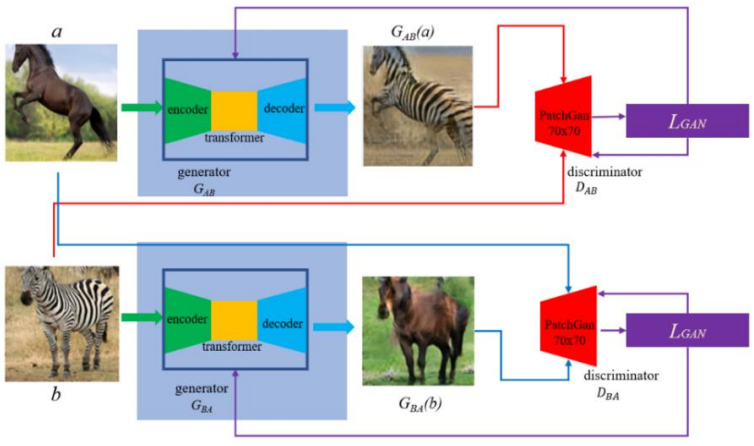

Fig.2Training frameworkof CycleGAN.

2.2 Cycle-consistency loss

From the previous discussion, the goal of the generators G_{AB} and G_{BA} is to transform input samples into their corresponding target-domain representations, namely $G_{AB}(a)$ and $G_{BA}(b)$. If the generated samples are fed back into the opposite generator to produce reconstructed samples, these reconstructions should ideally match the original inputs.For example, consider an input sample $a \in$ A . After being translated by generator G_{AB}, the output becomes $G_{AB}(a) \in$ B. When this generated output is further processed by the opposite generator G_{BA}, it produces

a reconstructed sample $G_{BA}(G_{AB}(a)) \in$ A. Theoretically, this reconstructed sample should be identical to the original input a,but discrepancies often exist in practice. The conceptual illustration of this process is shown in Fig. 3.

Based on this principle, the cycle-consistency loss (L_{cyc}) is introduced to perform the reconstruction to be as close as possible to the original input. Its formulation is given as follows:

$$L_{cyc}(G_{AB}, G_{BA}) = \mathbb{E}_{a\sim Pdata(a)}\left[||G_{BA}\left(G_{AB}(a)\right) - a||_1\right] + \mathbb{E}_{b\sim Pdata(b)}\left[||G_{AB}\left(G_{BA}(b)\right) - b||_1\right] \quad (3)$$

where $G_{BA}(G_{AB}(a))$ denotes the reconstructed image in domain A, and$a$ represents the real image from domain A. Similarly, $G_{AB}(G_{BA}(b))$ and b refer to the reconstructed image and the real image in domain B, respectively. The L_{cyc} is defined as the L_1norm difference between the reconstructed images and their corresponding real images.

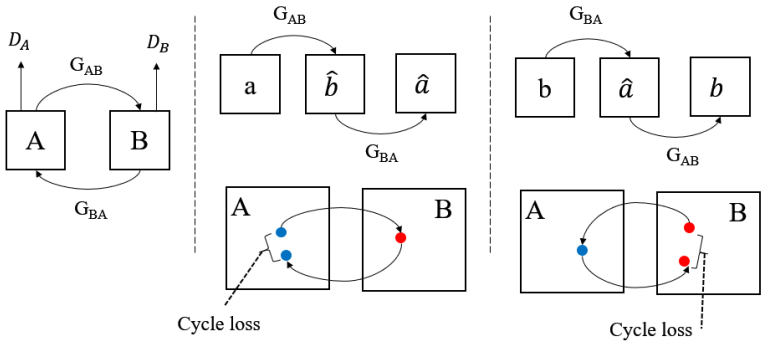

Fig. 3 The conceptual illustration of mapping process.

2.3 Objective function of CycleGAN

By combining the adversarial loss of the generative adversarial network with the cycle-consistency loss, the full objective function of CycleGAN can be expressed as follows:

$$L(G_{AB}, G_{BA}, D_{AB}, D_{BA}) = L_{GAN}(G_{AB}, D_{AB}, A, B) + L_{GAN}(G_{BA}, D_{BA}, B, A) + \lambda L_{cyc}(G_{AB}, G_{BA}) \quad (4)$$

whereλ controls the relative importance of two objects. With the full objective defined, the remaining task is to resolve the differing optimization goals of the generator and the discriminator. The corresponding training objectives can be expressed mathematically as follows:

$$G_{AB}^*, G_{BA}^* = arg \min_{G_{AB}, G_{BA}} \max_{D_{AB}, D_{BA}} L(G_{AB}, G_{BA}, D_{AB}, D_{BA}) \quad (5)$$

whereG_{AB}^* and G_{BA}^* denote the two GAN models corresponding to the bidirectional mappings. The generators $G_{AB}$and G_{BA} are trained to minimize the objective function, whereas the discriminators are trained to maximize it.

To ensure a more stable training process, produce higher-quality results, and enable each model to effectively achieve its respective objective. The overall objective function is decomposed into the following loss components:

$$loss\ G_{AB} = \mathbb{E}_{a\sim Pdata(a)}\left[\left(D_{AB}\left(G_{AB}(a)\right) - 1\right)^2\right] \quad (6)$$

$$loss\ D_{AB} = \mathbb{E}_{b\sim Pdata(b)}[(D_{AB}(b) - 1)^2] + \mathbb{E}_{a\sim Pdata(a)}[D_{AB}(G_{AB}(a))^2] \quad (7)$$

$$loss\ G_{BA} = \mathbb{E}_{b\sim Pdata(b)}\left[\left(D_{BA}\left(G_{BA}(b)\right) - 1\right)^2\right] \quad (8)$$

$$loss\ D_{BA} = \mathbb{E}_{a\sim Pdata(a)}[(D_{BA}(a) - 1)^2] + \mathbb{E}_{b\sim Pdata(b)}[D_{BA}(G_{BA}(a))^2] \quad (9)$$

With the individual loss functions defined, each component of the modelor each networkcan be trained to fulfill its corresponding objective.

To implement an intelligent enhancement system using the CycleGAN framework, it is essential to define the styles represented by domains A and B. In this work, we design domain A as facial images with freckles, while domain B corresponds to aesthetically enhanced facial images. Figure 4 illustrates the beautified results of domain-A images produced by CycleGAN under these definitions. As observed, the model successfully removes freckles,but it also introduces noticeable loss of edge details such as blurring around the teeth, double-eyelid folds, and other fine structural boundaries.

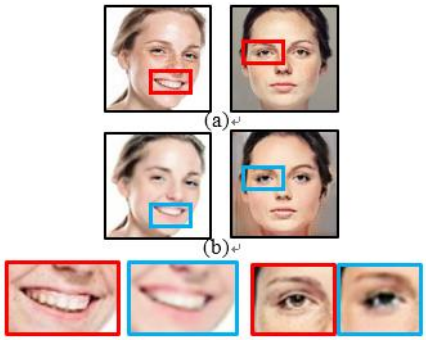

Fig. 4(a) Original facial image with freckles.(b)Beautified facial image generated by CycleGAN. (c) Comparison of fine image details between the original image and generated results.

III. PROPOSED METHOD

To address the limitations encountered when applying CycleGAN to intelligent facial enhancement, we propose a modified CycleGAN (MCycleGAN). In addition to preserving the unpaired training paradigm, the proposed MCycleGAN incorporates Sobel, Gaussian, and Bilateral filters [7] to extract image features and construct positive samples. These filtered representations are integrated into the loss function to alleviate the issue of edge-detail degradation commonly observed in conventional CycleGAN-based enhancement.

3.1 Incorporating SobelandGaussian filters

The Sobel operator extracts gradient-based edge features that emphasize intensity transitions. In contrast, the Gaussian filter is applied to obtain a smoothed version of the image, suppressing high-frequency noise while retaining essential structural information. By incorporating both filtered outputs into the loss function, the model is guided to maintain visually important contours and textures in the generated results. Therefore, the proposed MCycleGAN combined the Sobel and Gaussian filters (MCycle_SGGAN) to enhance edge preservation and reduce detail blurring during image translation. On the other hand, the MCycle_SGGAN introduces additional loss terms derived from Sobel and Gaussian filtered representations.

To preserve edge information while simultaneously suppressing freckle-related features, MCycle_SGGANcan enhance the freckle-removal capability of CycleGAN. Specifically, for each input image a, b and their corresponding generated outputs $G_{AB}(a)$and $G_{BA}(b)$, Sobel and Gaussian filters are applied to extract target features. The filtered results of (a, $G_{AB}(a)$) and (b,$G_{BA}(b)$) are then compared using the L_1 norm, and the resulting differences are included as additional losses during network training. The Sobel and Gaussian filtering operations are denoted as F_{Sobel} and $F_{Gaussian}$, respectively. Figures 5 and 6 show the processes of loss computation using Sobel filter and Gaussian filter in CycleGAN,respectively.The corresponding loss functions$L_{Sobel}$and $L_{Gaussian}$are defined as follows:

$$L_{Sobel}(G_{AB}, G_{BA}) = \mathbb{E}_{a\sim Pdata(a)}\left[||F_{Sobel}\left(G_{AB}(a)\right) - F_{Sobel}(a)||_1\right] + \mathbb{E}_{b\sim Pdata(b)}\left[||F_{Sobel}\left(G_{BA}(b)\right) - F_{Sobel}(b)||_1\right] \quad (10)$$

$$L_{Gaussian}(G_{AB}, G_{BA}) = \mathbb{E}_{a\sim Pdata(a)}\left[||F_{Gaussian}\left(G_{AB}(a)\right) - F_{Gaussian}(a)||_1\right] + \mathbb{E}_{b\sim Pdata(b)}\left[||F_{Gaussian}\left(G_{BA}(b)\right) - F_{Gaussian}(b)||_1\right] \quad (11)$$

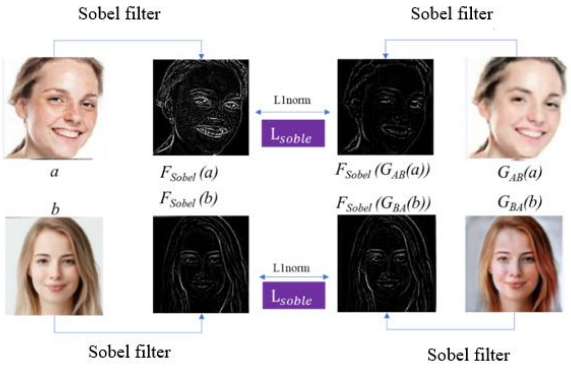

Fig. 5Illustration of edge feature extraction for loss computation.

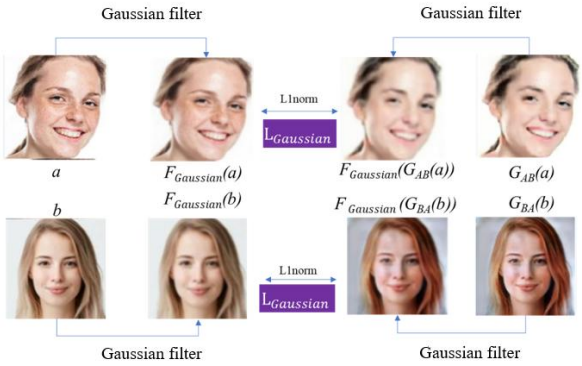

Fig. 6 Illustration of removing non-target features using Gaussian filtering for loss computation.

The loss function of proposed MCycle_SGGANcan be incorporated into the training process in parallel with the edge-based and smooth-based loss functions from Eqs. (4), (10) and (11). Therefore,the overall loss function is accordingly reformulated as follows:

$$L(G_{AB}, G_{BA}, D_{AB}, D_{BA}) = L_{GAN}(G_{AB}, D_{AB}, A, B) + L_{GAN}(G_{BA}, D_{BA}, B, A) + \lambda_1 L_{cyc}(G_{AB}, G_{BA}) + \lambda_2 L_{sobel}(G_{AB}, G_{BA}) + \lambda_3 L_{Gaussian}(G_{AB}, G_{BA}) \quad (12)$$

where$\lambda_1$controls the relative importance of two objects,$\lambda_2$controls the contribution of Sobel loss,and$\lambda_3$controls the contribution of Gaussian loss.

It can be observed when only incorporating the Sobel-based lossinto the CycleGAN framework significantly enhances its ability to preserve edge details. This improvement arises from the explicit reinforcement of high-frequency features introduced by the Sobel loss, which mitigates the tendency of unsupervised, unpaired training to bias the network toward learning non-target features. However, this modification also leads to a slight degradation in the freckle removal capability. Therefore, a Gaussian filter is further introduced to alleviate the interference between edge-detail features and the high-frequency characteristics of freckles.

3.2 Incorporating Sobeland Bilateralfilters

To preserve edge information while suppressing freckle-related features, we further take the replacement of the Gaussian filter with a Bilateral filter [8] in the MCycle_SGGAN, as called MCycle_SBGAN. However, directly incorporating the Bilateral filter into the loss function significantly increases computational complexity. This is because the Bilateral filter considers both spatial proximity and intensity similarity, requiring dynamic weight computation for every convolution operation. To avoid such computational overhead, the Bilateral filtering process is moved to the preprocessing stage.

Before training process, each freckled-face image in domain A is preprocessed using a Bilateral filter to generate a third domain C denoted as $\{c_i\}_{i=1}^{N}$ where $c_i \in C$. The filtered images in domain C serve as edge-preserving, freckle-reduced reference targets. During training, the difference between images $a$and $c$is computed using the L_1norm and incorporated as an additional loss term. This Bilateral filtering loss is denoted as $L_{\text{Bilateral}}$. The corresponding training objective can be mathematically expressed as follows:

$$L_{Bilateral}(G_{AB}, G_{BA}) = \mathbb{E}_{a\sim Pdata(a)}[||G_{AB}(a) - c||_1] \quad (13)$$

Therefore, the total loss function is modified as follows:

$$L(G_{AB}, G_{BA}, D_{AB}, D_{BA}) = L_{GAN}(G_{AB}, D_{AB}, A, B) + L_{GAN}(G_{BA}, D_{BA}, B, A) + \lambda_1 L_{cyc}(G_{AB}, G_{BA}) + \lambda_2 L_{sobel}(G_{AB}, G_{BA}) + \lambda_3 L_{Bilateral}(G_{AB}, G_{BA}) \quad (14)$$

where $\lambda_3$controls the relative contribution of Bilateral-filter-based loss within the overall objective loss. Using this enhanced loss function, we retrain the model and evaluate its performance.

IV. EXPERIMENTAL RESULTS

The dataset used in this work was collected from Google Images and Shutterstock. Using both Chinese and English keywords—*freckle face* and *pretty face*—we constructed a Female Face Dataset (FFD). The FFD consists of two training subsets representing two distinct domains, each containing 215 images with a resolution of 128 × 128. In addition, a test set of the same resolution was prepared, comprising 16 images belonging to Domain A.

All experiments were conducted using Python as the implementation environment [9-10]. Both CycleGAN and the proposed MCycleGAN were trained and evaluated on the FFD to ensure consistent benchmarking. Furthermore, we employed an objective semantic segmentation network to assess edge-detail preservation, complemented by subjective human perceptual evaluations to measure the overall beautification quality.

To quantitatively evaluate the preservation of facial structure and edge details after image translation, we employ a pretrained BiSeNet(Bilateral Segmentation Network) [11] to analyze both the original and the generated images. The segmentation results are assessed using three widely adopted metrics: intersection over union (IOU), per-pixel accuracy, and per-class accuracy. The corresponding metricsare defined as follows:

$$IOU = \frac{P_{ij}}{P_i \cup P_j} \quad (15)$$

$$\textit{per-pixel acc} = \frac{P_{ij}}{P_i} \quad (16)$$

$$\textit{per-pixel class} = \frac{1}{class} \times \frac{P_{ij}}{P_i} \quad (17)$$

where P_i denotes the semantic segmentation result of the original image produced by BiSeNet, P_j denotes the semantic segmentation result of the generated image, and P_{ij} represents the intersection of P_i and P_j.

Table 1 shows the experimental results on the FFD dataset under different configurations of the loss-function parameters $\lambda_2$and $\lambda_3$when$\lambda_1 = 10$. From the experimental results, we observe that setting $\lambda_2 = 1$ in MCycleGAN already leads to a significant improvement in the preservation of edge contours. Increasing λ_2 beyond this value does not yield further performance gains. Consequently, λ_2 is fixed at 1 in subsequent experiments, while λ_3 is varied to investigate its impact on the generated results. The results indicate that MCycle_SGGAN achieves an average IOU improvement of 0.07, while MCycle_SBGAN improves the average IOU by 0.06. These resultsin Table1 demonstrate that the proposed MCycleGAN produces facial image beautification results with lower distortion. Moreover, the IOU increases as $\lambda_3$becomes larger.

Table 1. Quantitative results of semantic segmentation.

method \ measurement	per-pixel acc	per-pixel class	IOU	(λ_2,λ_3)
CycleGAN	0.93	0.10	0.84	None
MCycle_SGGAN	0.95	0.11	0.90	(1,1)
MCycle_SGGAN	0.95	0.11	0.92	(1,20)
MCycle_SBGAN	0.95	0.11	0.89	(1,1)
MCycle_SBGAN	0.96	0.11	0.91	(1,20)

The corresponding visual resultsof MCycle_SGGANusingthe contributionparameters with (λ_1=1, λ_2=1, λ_3=10) are shown in Fig. 7. As observed in Fig. 7(b) and 7(c), we can find that MCycle_SGGANdemonstrate a strong capability to preserve edge details than the CycleGAN. Moreover, From Fig. 7(d), we observe that the MCycle_SGGAN freckle removal performance shows no significant degradation when compared with the CycleGAN.

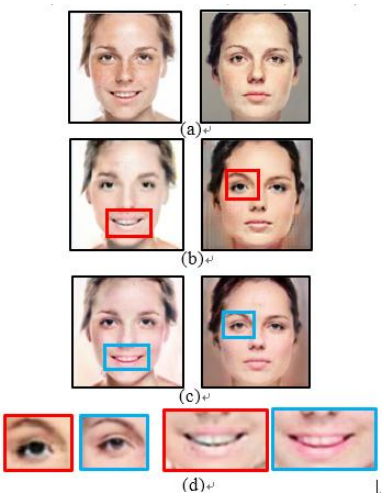

Fig. 7 (a) Original images. (b) Images byCycleGAN.(c) Images by MCycle_SGGAN. (d) Comparison of edge details and freckle removal.

On the other hand, the corresponding visual results of MCycle_SBGAN using the same contribution parameters as MCycle_SGGAN are shown in Fig. 8. As observed in Fig. 8(a), the edge contour features in the eye and nose regions partially overlap with freckle-related featuresin the left original image. From Fig. 8(b), we can find that slight loss of edge information can be observed when a basic Gaussian filter is applied to remove freckle features. Whereas such degradation does not occur when a Bilateral filter is employed, as shown in Fig. 8(c). Finally, we observe that the MCycle_SBGAN achieves superior beautification performance in these overlapping feature regions as compared with MCycle_SGGAN.

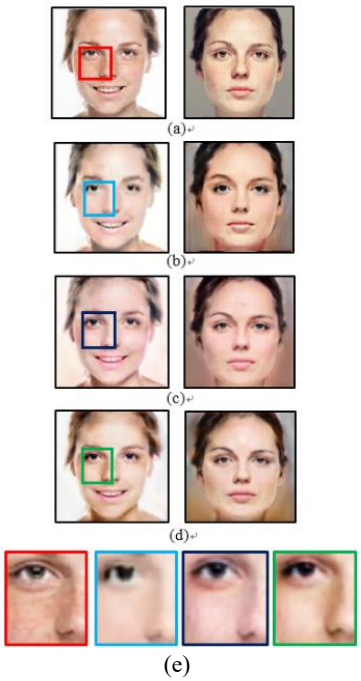

Fig. 8. (a) Original image. (b) Images by CycleGAN. (c) Images by MCycle_SGGAN. (d) Images by MCycle_SBGAN. (e) Comparison of results obtained by different methods.

We additionally invited a total of 34 image processing experts to participate in a preference-ranking survey conducted via the Survey Cake online platform [12]. Two subjective evaluation methods were designed as follows.

Method 1 assumes that users have access to the original image and subsequently apply an intelligent facial beautification system. As shown in Fig. 9, the statistical results indicate that the percentage of cases in which MCycleGAN was ranked first is 16.5% higher than that of CycleGAN on average. This demonstrates that, when the original image is available for reference, the facialimage beautification results produced by MCycleGAN are consistently preferred over those generated by CycleGAN.

Method 2 simulates a scenario in which MCycleGAN is embedded in a camera system, such that intelligent facial beautification is performed immediately after image capture. As shown in Fig. 10, under this setting, MCycleGAN achieves a first-place ranking rate that is 16% higher than CycleGANon average. These results further confirm that, even when only the beautified outputs are compared, MCycleGAN remains more visually appealing to human observers than CycleGAN.

Based on the comparative results of the objective and subjective evaluations, it is evident that different parameter configurations lead to variations in beautification performance. To better accommodate user preferences, we recommend the following parameter settings for the intelligent facial image beautification system.For Method 1, the optimal facial beautification results for MCycle_SBGAN and MCycle_SGGANare achieved using (λ_1 =10, λ_2 =1, λ_3 =20). For Method 2, the recommended parameter configurations for MCycle_SBGAN and MCycle_SGGANare(λ_1=10, λ_2=1, λ_3=1) and (λ_1=10, λ_2=1, λ_3=20), respectively. These configurations provide the most visually satisfactory facial image beautification outcomes under their respective evaluation scenarios.

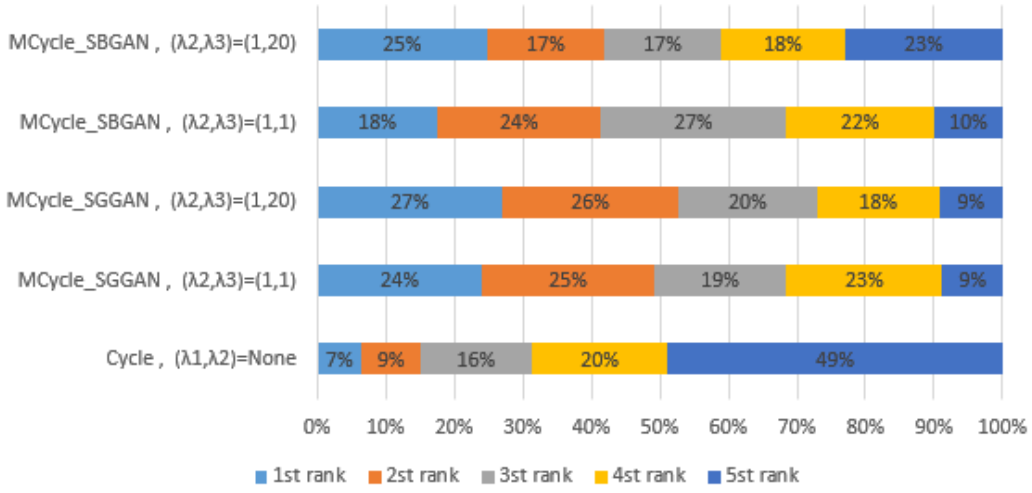

Fig.9 Results of human subjective perception evaluation using Method 1.

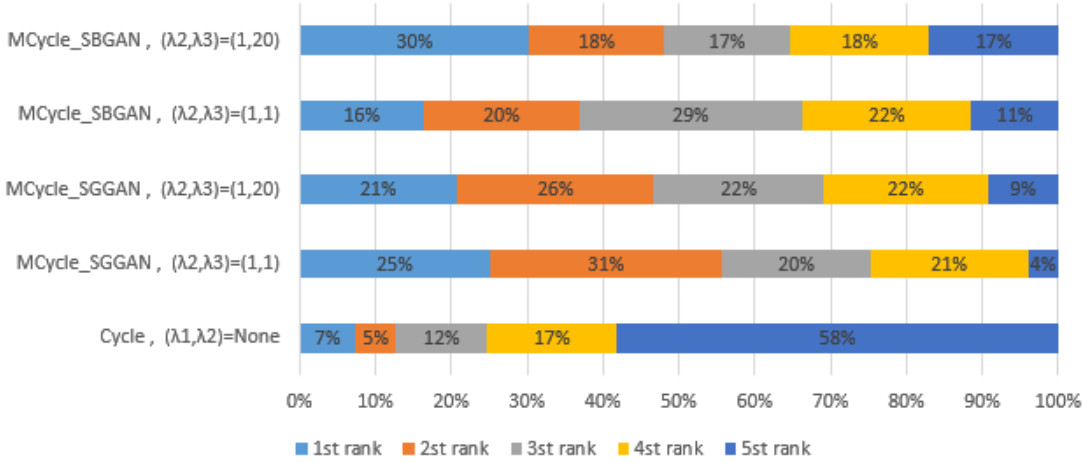

Fig.10Results of human subjective perception evaluation using Method 2.

V. CONCLUSION

Although CycleGAN has demonstrated remarkable success in image-to-image translation and has been widely applied across various domains, its reliance on unsupervised learning without explicit human annotation limits its ability to effectively distinguish between low-frequency and high-frequency features during feature mapping. As a result, the generated images often suffer from blurring artifacts. This limitation significantly affects the performance of CycleGAN when applied to facial image beautification tasks.

The main contributions of this work are twofold. First, edge features are explicitly extracted and incorporated into the loss function to enhance the network's capability to learn and preserve contour information. Second, two low-frequency filtering strategies are introduced to suppress non-target features. Therefore, these improvements enable the proposed method to reduce visual distortion in the translated images and achieve higher subjective preference scores in human perceptual evaluations.

Based on the semantic segmentation results obtained using BiSeNet, the proposed MCycleGAN consistently outperforms the original CycleGAN. Specifically, the average IOU is improved by 0.0575 and 0.065 on the two test sets, respectively. These objective improvements are further corroborated by subjective perceptual evaluations. When the original images are available for reference, MCycleGAN achieves a 16.5% higher first-place preference rate than CycleGAN. Even when only the generated results are compared without reference images, MCycleGAN still surpasses CycleGAN by an average margin of 16% in first-place preference.

REFERENCES

[1]. I. J. Goodfellow, J. P. Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozairy, A.Courville and Y. Bengio, "Generative Adversarial Nets," in *Advances in Neural Information Processing Systems 27 (NIPS 2014)*.

[2]. J. Zhu, T. Park, P. Isola and A. A. Efros, "Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks,"*2017 IEEE International Conference on Computer Vision (ICCV)*, Venice, 2017, pp. 2242-2251, doi: 10.1109/ICCV.2017.244.

[3]. J. Gan and J. Liu, "Applied Research on Face Image Beautification Based on a Generative Adversarial Network," Electronics 2024, Electronics 2024, 13, 4780. https://doi.org/10.3390/electronics13234780

[4]. H. Chang, J. Lu, F. Yu and A. Finkelstein, "PairedCycleGAN: Asymmetric Style Transfer for Applying and Removing Makeup," 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, 2018, pp. 40-48, doi: 10.1109/CVPR.2018.00012.

[5]. J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," In ECCV, 2016.

[6]. P. Isola, J. Zhu, T. Zhou and A. A. Efros, "Image-to-Image Translation with Conditional Adversarial Networks," 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, 2017, pp. 5967-5976, doi: 10.1109/CVPR.2017.632.

[7]. K. He, J. Sun and X. Tang, "Guided Image Filtering," in IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 35, no. 6, pp. 1397-1409, June 2013, doi: 10.1109/TPAMI.2012.213.

[8]. C. Tomasi and R. Manduchi,"Bilateral filtering for gray and color images," in Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271., 10.1109/ICCV.1998. 710815

[9]. https://github.com/eriklindernoren/PyTorch-GAN/tree/master/implementations/cyclegan

[10]. https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix

[11]. C. Yu, J. Wang, C. Peng, C. Gao, G. Yu, N. Sang, "BiSeNet: Bilateral Segmentation Network for Real-time Semantic Segmentation," *The European Conference on Computer Vision (ECCV)*, pp. 325-341,2018.

[12]. https://www.surveycake.com/