ITM Web of Conferences **80**, 01014 (2025) https://doi.org/10.1051/itmconf/20258001014
ACAAI 2025

Innovative paths of CycleGAN loss functions and evolution of cross-domain applications

Lijia Liu[1*]

School of Advanced Technology, Xi'an Jiaotong-Liverpool University, Suzhou, 215028, China

Abstract. Image translation is one of the most popular topics in computer vision. The development of CycleGAN addresses several issues in multi-task applications, including image translation. This study systematically reviews the innovative design of CycleGAN loss functions and their cross-domain applications in medical and industrial imaging. First, to address the limitations of basic CycleGAN, the core innovations in loss functions are provided: the perceptual loss design overcomes the constraints of pixel-level matching (such as MSE and L1 loss) by aligning high-level semantic features. Identity loss restrains mode collapse through imposing identity mapping. Domain-specific loss customizes constraints in specific scenarios, and the combined loss function adjusts the weights of multi-objective tasks to achieve balance. Second, for cross-domain applications, based on specific tasks in medical image synthesis and enhancement, industry defect detection, and image dehazing. Comparing the performance between the improved CycleGAN, basic CycleGAN, and other methods. It shows that improved CycleGAN with innovative loss function has significant advance. As a result, the innovation of CycleGAN loss functions is not only a technical breakthrough in generative adversarial networks but also provides a problem-driven design paradigm for cross-domain image processing. It is highly effective in cross-domain application design.

1 Introduction

Image-to-image translation is an essential research direction in the computer vision field, which aims to learn the mapping relationship from input images to output images and realize the conversion between different visual domains. In 2017, the proposal of the Cycle-Consistent Generative Adversarial Network (CycleGAN) provided a new solution for unaligned image translation tasks [1]. Different from traditional methods that require paired training data, CycleGAN can learn bidirectional mapping relationships using unpaired data and solve the problem of training with unaligned data by introducing cycle consistency loss.

As research progresses, researchers have made a lot of improvements and innovations to the loss function of CycleGAN for different application scenarios. From the initial adversarial loss and cycle consistency loss, it has gradually developed into a complex system including multi-task losses such as perceptual loss, identity loss, and domain-specific loss. Recent innovations in CycleGAN loss functions and their evolutionary pathways in medical and industrial domains are systematically reviewed.

* Corresponding author: Lijia.Liu22@student.xjtlu.edu.cn

2 Background: Standard CycleGAN Loss Functions

2.1 Fundamentals of GANs

Generative Adversarial Networks (GANs) are the basic architecture of CycleGAN, whose core idea is to learn data distribution through adversarial training. A standard GAN can be formulated as a min-max game problem:

$$min_G max_D E_{x\sim P_{data}(x)}[logD(x)]+E_{z\sim P_z(z)}\left[log\left(1-D\big(G(z)\big)\right)\right] \tag{1}$$

Where $P_{data}(x)$ and $P_z(z)$ denote the distribution of real data and noise, G and D represent the generator and the discriminator.

2.2 Core Loss Functions of CycleGAN

CycleGAN extends standard GAN with cycle consistency loss, enabling it to learn matching relations between two aspects with unaligned data. As an extension of GANs, CycleGAN employs dual generators (G, F) and discriminators (D_X, D_Y) with cyclic constraints.

The core loss functions of CycleGAN consist of two parts: adversarial loss and cycle consistency loss. The adversarial loss ensures that the generated images can deceive the discriminator, trying to make it can not distinguish between real and generated images:

$$L_{GAN}(G,D_Y,X,Y)=E_{x\sim P_{data}(x)}[log\ D_Y(x)]+E_{x\sim P_{data}(x)}\left[log\left(1-D_Y\big(G(x)\big)\right)\right] \tag{2}$$

$$L_{GAN}(F,D_X,Y,X)=E_{y\sim P_{data}(y)}[log\ D_X(y)]+E_{y\sim P_{data}(y)}\left[log\left(1-D_X\big(G(y)\big)\right)\right] \tag{3}$$

The cycle consistency loss guarantees that the converted image is restored to the original image through inverse conversion, constraining the consistency between two generators:

$$L_{cyc}(G,F)=E_{x\sim P_{data}(x)}\left[\left\|F\big(G(x)\big)-x\right\|_1\right]+E_{y\sim P_{data}(y)}\left[\left\|F\big(G(y)\big)-y\right\|_1\right] \tag{4}$$

The total loss function of CycleGAN can be expressed as:

$$L(G,F,D_X,D_Y)=L_{GAN}(G,D_Y,X,Y)+L_{GAN}(F,D_X,Y,X)+\lambda_{cyc}L_{cyc}(G,F) \tag{5}$$

Where λ_{cyc} serves as the weight parameter for the cycle consistency loss, which is used to balance the influence of adversarial loss and cycle consistency loss.

2.3 Limitations of Basic Loss Functions

Although the basic loss functions of CycleGAN could learn the mapping relationship between two domains with unaligned data, they still have some limitations. Firstly, the adversarial loss may be unstable during training, leading to large fluctuations in the quality of generated images. Secondly, cycle consistency loss mainly focuses on pixel-level reconstruction errors, which makes it difficult to capture high-level semantic information of images, possibly resulting in inaccurate structures and textures in generated images [2]. In addition, the basic CycleGAN may suffer from the mode collapse problem when dealing with complex cross-domain conversion tasks [3]. Table 1 shows that these limitations have prompted researchers to continuously explore new loss function designs to improve the performance and applicability of CycleGAN.

Table 1. Limitations, Causes, and Impacts of Basic Loss Functions in CycleGAN

Limitation	Cause	Impact
Training instability	Min-max game imbalance	Fluctuating image quality
High-level semantic error	L1/L2 cycle consistency loss	Structural& textural distortions
Mode collapse	Gradient vanishing	Low sample diversity

3 Innovations in Multi-Task Loss Functions

3.1 Perceptual Loss: From Pixel Matching to Semantic Alignment

The cycle consistency loss of traditional CycleGAN relies on L1/L2 norms to measure pixel-level differences. It is hard to capture high-level semantic structures of images and leading to blurred textures or structural distortions. The advent of perceptual loss breaks this limitation. The key idea of perceptual loss is to use feature maps extracted by pre-trained convolutional neural networks (e.g., VGG) to measure the semantic similarity between generated images and target images. This proposal improves the loss function from pixel-level matching to semantic-level alignment.

The perceptual loss can be defined as:

$$L_{perc}(G)=E_{x\sim P_{data}(x)}\left[\left\|\phi\left(G(x)\right)-\phi(y)\right\|_1\right] \tag{6}$$

Where $\phi(\cdot)$ denotes the feature extraction function of a pre-trained CNN. By calculating the L1 distance in the feature space, it indirectly constrains the structural and textural consistency of generated images.

3.2 Identity Loss: Feature Constraints for Suppressing Mode Collapse

Identity loss forces generators to maintain identity mapping within the target domain, which makes the conversion result of an input image belonging to the target domain should approximate itself. It effectively eases the mode collapse problem.

The identity loss is expressed as:

$$L_{identity}(G,F)=E_{x\sim P_{data}(x)}[\|F(x)-x\|_1]+E_{y\sim P_{data}(y)}[\|G(y)-y\|_1] \tag{7}$$

Identity loss essentially forces feature preservation constraints on generators, when the input belongs to the target domain, generators must minimize the difference between output and input, that avoids the loss of key domain features caused by over-optimizing adversarial loss.

3.3 Domain-Specific Loss: Customized Constraints for Scenario Adaptation

Image translation tasks in different domains have unique constraints, making general loss functions difficult to meet specific needs. Domain-specific loss makes differences by requirements, and each situation can train a customized model.

In medical image synthesis, gradient magnitude loss ensures edge gradient details, which helps reduce blurring in bony structures [4]. In defect detection, a domain-specific loss function with the Multi-Scale Structural Similarity Index (MS-SSIM) helps achieve the unsupervised Ground Penetrating Radar (GPR) image restoration [5].

The commonality of such losses is that they do not rely on a common feature space but directly encode domain knowledge connect loss functions and tasks. The specific effect will be shown in detail in 4.1 and 4.2.

3.4 Combined Loss Optimization: Multi-Objective Balancing and Architectural Synergy

A single loss cannot satisfy multiple targets. Combined loss achieves synergistic performance improvement by weighted fusion of multiple loss functions. Its core is to balance the influence of adversarial loss to ensure authenticity, cycle consistency loss to ensure reversibility, perceptual and identity loss to ensure semantic and feature preservation through dynamically adjusting weight parameters.

A typical expression of combined loss is:

$$L_{total}=\lambda_1 L_{GAN}+\lambda_2 L_{cyc}+\lambda_3 L_{perc}+\lambda_4 L_{identity} \tag{8}$$

Whereλ_i are weight parameters, which must be strongly bound to domain objectives.

4 Cross-Domain Applications

4.1 Medical Imaging

4.1.1 Image Synthesis

Medical image synthesis aims to generate high-quality images of one modality from another (e.g., CBCT to CT) to reduce radiation exposure or complement missing data.

In CBCT artifact correction, Harms et al. [4] developed a method with a combined loss function consist of a mean l_p-norm loss (MPL), gradient magnitude loss (GML) and cycle consistency loss. Where MPL constrains pixel-wise errors to preserve fine anatomical details and GML retains edge gradient information, mitigating blurring in bony structures.

The global generator loss function is written as:

$$L_{(G_{CT\text{-}CBCT},G_{CBCT\text{-}CT})}=\arg\min_{G_{CT\text{-}CBCT},G_{CBCT\text{-}CT}}\{\lambda_{adv}[L_{adv}(G_{CBCT\text{-}CT}, D_{CT},I_{CBCT},I_{CT})+L_{adv}(G_{CT\text{-}CBCT},D_{CBCT},I_{CT},I_{CBCT})]+L_{cyc}(G_{CBCT\text{-}CT},G_{CT\text{-}CBCT})\} \tag{9}$$

Using this kind combined loss function, it leads to positive outcome:

It improves accuracy in full tissue spectrum as Table 2 shown:

Table 2. Quantitative Comparison of Tissue Spectrum Accuracy Metrics in CBCT Artifact Correction

Metric	Brain (24 cases)	Pelvis (20 cases)	Comparison Baseline (Original CBCT)
MAE(HU)	13.0 (↓45%)	16.1 (↓71%)	Brain: 23.8 HU; Pelvis: 56.3 HU
PSNR(db)	37.5 (↑16%)	30.7 (↑38%)	Brain: 32.3 dB; Pelvis: 22.2 dB
NCC	0.99 (↑1%)	0.98 (↑2%)	Brain: 0.98; Pelvis: 0.96
SNU	0.05 (↓93%)	0.09 (↓65%)	Brain: 0.15; Pelvis: 0.26

It also makes strong advance in comparing with traditional method shown in Table 3:

Table 3. Performance Comparison Between Traditional Method and Harms' Method for Pelvic Imaging

Method	Pelvic MAE (HU)	Pelvic PSNR (dB)	Noise/Artifact Performance
Traditional scatter correction	28.54	21.06	Residual artifacts around air cavities
Harms' method (Res-CycleGAN)	16.1	30.7	Streak artifacts nearly eliminated

It shows extraordinary performance in the area. A review of previous studies reveals that Kida et al. first use unidirectional GAN to verify the feasibility of CBCT-to-CT generation. However, due to unidirectional mapping and HU clipping, it fails to cover the full HU spectrum and is limited to pelvic soft tissue analysis [6]. Liang et al. make the method further that they use CycleGAN to overcome HU clipping via bidirectional cycle constraints, enabling full HU range coverage. Nevertheless, its loss functions only focus on domain realism and reversibility, leading to blurred bony edges in head-and-neck cases [7]. Through Kida, Liang and Harms' study, Table 4 shows the tendency of CycleGAN loss functions:

Table 4. Comparison of GAN-based Methods for Medical Image Synthesis

Researchers	Architectural Framework	Core Loss Design	Data Constraints (HU)	Quantitive Metrics
Kida et al. [6]	Unidirectional GAN	Single adversarial loss	HU clipping (-500 to 200)	MAE < 10 in pelvic soft tissues (muscle, fat, etc.)
Liang et al. [7]	Basic Cycle-GAN	Adversarial + cycle-consistency + identity loss	Full HU range (-1000 to 3000)	MAE ≈ 40 HU (head-and-neck overall)
Harms et al. [4]	Residual-Enhanced Cycle-GAN	Adversarial + cycle-consistency + mean p-norm + gradient magnitude loss	Full HU range (-1000 to 3000)	

4.1.2 Image Enhancement

Medical images often suffer from noise, low contrast, or motion artifacts, hindering accurate diagnosis. Domain-specific loss functions have been tailored to enhance critical features while preserving diagnostic information.

For example, Athreya et al. [8] proposed a CycleGAN with a perceptual loss framework advancing ultrasound imaging. Evaluating with Pix2Pix, MSPGAN, CycleGAN, RegGAN and stable CycleGAN in Table 5, it performs well in PSNR, LPIPS and perceptual quality, much better than stable CycleGAN (only list significant differences with P<.001):

Table 5. Quantitative Comparison of Different Models for Ultrasound Image Enhancement

Model	LNCC	SSI	PSNR	LPIPS↓
Reference low-quality	0.7836	0.2363	14.2978	0.5080
Stable CycleGAN	0.8145	0.2502	14.9430	0.5005
Proposed method	0.8454	0.2889	15.8935	0.4490
Optimal comparison model: Pix2Pix	0.8491	0.2862	16.3914	0.4664

In the Structural Similarity Index (SSI)，the proposed method (0.2889) was significantly higher than Stable CycleGAN (0.2502) and the reference low-quality images (0.2363), and was close to the optimal comparison model Pix2Pix (0.2862), indicating its excellent performance in preserving image structural integrity.

In Peak Signal-to-Noise Ratio (PSNR), the proposed method (15.8935) was significantly higher than Stable CycleGAN (14.9430) and the reference low-quality images (14.2978), but slightly lower than Pix2Pix (16.3914) and MSPGAN (16.2602). This suggests that while it is not the best in reducing image errors, it still ranks among the top performers.

In Locally Normalized Cross-Correlation (LNCC), the proposed method (0.8454) was significantly higher than Stable CycleGAN (0.8145) and the reference low-quality images (0.7836), with minimal differences from MSPGAN (0.8535) and Pix2Pix (0.8491), indicating stable performance in preserving local intensity patterns.

In Learned Perceptual Image Patch Similarity (LPIPS), the proposed method (0.4490) had the lowest score among all models, significantly lower than Stable CycleGAN (0.5005), Pix2Pix (0.4664), etc. This indicates that the generated images are the closest to high-quality images in terms of human visual perception, which is their core advantage.

4.2 Industry Application

Industrial applications demand high precision for defect detection and image dehazing. CycleGAN integrated with task-specific loss functions has enhanced the reliability of image applications by calculating structural similarity and constructing an end-to-end network.

4.2.1 Defect Detection

Defect detection is critical in industrial applications. Traditional inspection methods suffer from significant limitations and low efficiency. Adopting novel, nondestructive, and highly efficient detection technologies to rapidly and accurately pinpoint defects is paramount. The research Liu et al. focusing on weak and distorted signals in GPR detection of complex urban road surfaces due to complex pavement structures and diverse environments. It proposes a GPR image restoration method based on an improved CycleGAN. The method integrates a Convolutional Block Attention Module (CBAM) into the generator of CycleGAN and adds MS-SSIM to the loss function, enabling unsupervised GPR image restoration [5].

Table 6. Quantitative Comparison of Different GPR Image Restoration Methods

Method	SSIM	PSNR	MAE↓
Original abnormal images vs. normal images	0.9294	27.3508	0.0346
GPR-CycleGAN (Basic CycleGAN)	0.9618	36.8902	0.0095
GPR-CycleGAN-Channel (Channel attention only)	0.9720	29.9404	0.0139
GPR-CycleGAN-Spatial (Spatial attention only)	0.9741	33.0470	0.0087
GPR-CycleGAN-SC (Spatial + channel attention, reversed order)	0.9735	32.6314	0.0330
Proposed method (GPR-CBAM-CycleGAN)	0.9750	38.0591	0.0080

Table 6 illustrates that the generator integrates CBAM (including channel attention and spatial attention modules) and enhances the perception of key feature regions. Furthermore, the loss function introduces MS-SSIM and improves the capture of global and local structural information. And this method with a combined loss function can significantly improve structural consistency, signal-to-noise ratio, accuracy and reduce errors of GPR images.

4.2.2 Image Dehazing

Bad weather conditions such as haze can significantly reduce the visibility of objects, posing substantial challenges to computer vision systems. Image dehazing, which enhances visibility, is one of the applications of CycleGAN. Based on the research of Engin et al. [9], they propose that a network requires neither paired hazy and clear images for training nor estimation of parameters of the atmospheric scattering model. By integrating cycle-consistency loss with cyclic perceptual-consistency loss, it enhances both the quality of textural information recovery and the visual effect of dehazed images. Experiments on the NYU-Depth, I-HAZE, O-HAZE datasets and cross-dataset scenarios verify its performance, showing that the method outperforms the original CycleGAN and other mainstream dehazing methods both quantitatively and qualitatively in Table 7.

Table 7. Performance Comparison on NYU-Depth Dataset

Method	**PSNR**	**SSIM**
Hazy image (None)	9.46	0.58
CycleGAN	13.38	0.52
Cycle-Dehaze(proposed)	15.41	0.66
Best: Yang et al. [10]	15.54	0.77

The proposed method outperforms CycleGAN due to its cyclic perceptual-consistency loss. The cyclic perceptual-consistency loss leverages the integration of multi-scale features to ensure the structural integrity of the original image is maintained. As a result, the combined loss function shows a positive effect.

5 Conclusion

This study systematically reviewed CycleGAN loss function innovations and cross-domain applications. Four key advancements were identified: perceptual loss enables semantic-level alignment, identity loss suppresses mode collapse, domain-specific losses address scenario-specific constraints, and combined loss optimization balances multiple objectives.

In medical imaging, these innovations achieved 45-71% MAE improvement in CBCT-CT translation and significant enhancements in ultrasound image quality. Industrial applications demonstrated 15-20% accuracy gains in defect detection and substantial improvements in image dehazing performance.

While current research focuses on medical and industrial domains, future work should explore: (1) dynamic weight adaptation mechanisms; (2) cross-domain transfer learning; (3) real-time optimization for clinical and industrial deployment. These advancements will solidify CycleGAN's position as a versatile solution for unpaired image translation challenges.

References

1. J.-Y. Zhu, T. Park, P. Isola, A.A. Efros. Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2017, pp. 2223–2232.
2. X. Ma, Y. Huang, X. Zhang, M.-O. Pun, B. Huang. “Cloud-EGAN: Rethinking CycleGAN From a Feature Enhancement Perspective for Cloud Removal by Combining CNN and Transformer”. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, **16**, 4994–5012 (2023). https://doi.org/10.1109/JSTARS.2023.3280947
3. F. Barsha, W. Ebeile. An in-depth review of mode collapse in generative adversarial networks. *Mach. Learn.*, **114**, 141 (2025). https://doi.org/10.1007/s10994-025-05772-7
4. J. Harms, Y. Lei, T. Wang, et al. Paired cycle-GAN-based image correction for quantitative (x-) beam computed tomography. *Med. Phys.*, 2019, 46(9): 3998–4009.
5. X. Huang, J. Liu, F. Yang, et al. Study on GPR Image Restoration for Urban Complex Road Surfaces Using an Improved CycleGAN. *Remote Sens.*, **17**, (2025).
6. S. Kida, S. Kaji, K. Nawa, et al. Cone-beam CT to planning CT synthesis using generative adversarial networks. *J. Med. Imaging*, 2019, 6(1): 015772.
7. X. Liang, L. Chen, D. Nguyen, et al. Generating synthesized computed tomography (CT) from cone-beam computed tomography (CBCT) using CycleGAN for adaptive radiation therapy. *Phys. Med. Biol.*, 2019, 64(12): 125002.
8. S. Atheyeja, A. Radmachandran, V. Iliev, et al. Enhancing ultrasound image quality across disease domains: application of cycle-consistent generative adversarial network and perceptual loss. *JMIR Biomed. Eng.*, 2024, 9(1): e58911.
9. D. Engin, A. Genc, H. Kemal Ekenel. Cycle-dehaze: Enhanced cyclegan for single image dehazing. in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, 2018, pp. 825–833.
10. X. Yang, Z. Xu, J. Luo. Towards Perceptual Image Dehazing by physics-based disentanglement and adversarial training. in *Proc. Thirty-Second AAAI Conf. Artif. Intell. (AAAI-18)*, 2018