Proceedings of the 4th International Conference on Software Engineering and Machine Learning
DOI: 10.54254/2755-2721/2026.32322

Improved CycleGAN for Low-Dose CT Image Enhancement with Structural Preservation

Ce Liang

Faculty of Engineering, The University of Sydney, Sydney, Australia
liangce.com@gmail.com

Abstract. Low-dose computed tomography (LDCT) is highly prevalent in clinical practice to minimise radiation dose, though it is usually characterised by high noise and the loss of finer anatomical details, which can affect diagnostic reliability. This paper proposes an Improved CycleGAN framework, termed IC-GAN, for LDCT image enhancement, designed to generate images with higher visual quality and structural fidelity that more closely resemble NDCT images. The incorporation of attention mechanisms and perceptual loss in the proposed IC-GAN strengthens structural preservation during adversarial training, enabling effective noise suppression while retaining essential anatomical features. Experimental results demonstrate improved structural fidelity, reflected by increased SSIM values. In addition, the enhanced images produced by IC-GAN offer better visualisation, which may support radiologists in making more confident diagnoses and treatment decisions. Overall, the proposed framework provides a powerful and efficient tool for enhancing LDCT images, with strong potential for integration into clinical imaging workflows.

Keywords: LDCT image enhancement, Improved CycleGAN (IC-GAN), Structural fidelity, Perceptual loss, Low-dose computed tomography

1. Introduction

The article presents a new generation adversarial learning model, which will be used to improve the images produced by low-dose CT and ensure that the structure is more consistent and visually better at the same time without reducing computational efficiency [1-9,11].

Enhancement of medical images had been an issue of utmost priority within the field of diagnostic imaging since low-dose computed tomography (LDCT) is often associated with disadvantages like higher noise levels and poor anatomical detail due to lower doses of radiation. In the past several years, LDCT denoising and image enhancement have received much research attention with conventional model-based software, trained deep learning models, and non-paired GAN-based models being studied [1-3].

Even though classical model-based algorithms, including iterative reconstruction and filtering methods, may enhance image clarity, they have been shown to over-smooth small anatomical features and eliminate finer details that are important to diagnose [1,3]. The development of deep learning has led to tremendous improvement in LDCT denoising with convolutional neural networks (CNNs) learning complex nonlinear mappings between low-dose and normal-dose images [5,6].

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2026.32322

However, these supervised learning approaches require a large-scale set of paired data, which is not easily available in clinical practice.

To address this limitation, generative adversarial networks (GANs) have been proposed to generate photorealistic high-quality images based on unpaired data [7]. In particular, CycleGAN has found extensive use in medical image translation and domain adaptation tasks as a result of its cyclic consistency constraint that guarantees semantic identity between the input domain and the reconstructed domain [8,9].

Nonetheless, even the current GAN approaches have significant problems: (1) architectural distortion that results in the loss of major anatomical characteristics; (2) lack of perceptual help resulting in hallucination artifacts; (3) restricted generalisation ability with varying imaging conditions [4,11].

In order to overcome these challenges, this study contributes in three aspects:

We suggest an Improved CycleGAN model (IC-GAN) which is built on the basis of the CycleGAN but with additional attention mechanisms and perceptual loss. The design helps the model to preserve important anatomical structures, decrease the noise, and enhance performance in terms of processing time, accuracy, and resiliency in diverse imaging conditions.

2. The image quality of IC-GAN has been found to be superior to the image quality of existing CycleGAN-based approaches. Especially, the average SSIM values of the baseline GAN-based models are in the low 0.8, but in the mid-to-high range of 0.8, the IC-GAN gives SSIM. This objective index suggests that the suggested method can preserve anatomy features more effectively and generate more trustworthy enhanced images.

IC-GAN can also be applied in clinical setting. It has the ability to help medical practitioners make a diagnosis at an earlier stage and with more certainty as the model will produce more accurate and dependable images. This advantage might be applied to assist in treatment planning and could result in better patient results when used in real clinical settings.

2. Related work

The earliest low-dose CT (LDCT) enhancement techniques [1-3,5,6] may be generally divided into three characteristic directions:

1. Model-based iterative reconstruction methods [1]:

Such methods have the potential to enhance the general quality of an image but tend to flatten small details of anatomy as well as lose significant texture. Also, the computational cost of iterative optimization is high, which results in long reconstruction times and restricts its application to rapid clinical processes.

2. Supervised convolutional neural network (CNN) based methods:

In comparison to the conventional model-based methods, CNNs tend to have higher denoising rates and good visual quality [5,6]. However, it is important to note that their effectiveness highly relies on the presence of large-scale paired data, which cannot be easily found in a real clinical setting because of ethical and technical limitations [2].

3. Limitations of supervised deep learning approaches:

The reconstruction efficiency and structural preservation are not yet satisfactory enough to be able to deploy them reliably in a clinical setting, which is why it is important to explore stronger and more data efficient frameworks [1,3].

GANs have been extensively investigated as a reduction of the requirement to use paired data in LDCT image improvement. Representative studies [7-10] may be divided into two main groups:

1. GAN-based methods focusing on enhancement accuracy.

DOI: 10.54254/2755-2721/2026.32322

Unpaired image translation techniques based on CycleGAN add cycle consistency loss to allow training on unpaired image datasets. These methods have been used with success in low-dose CT denoising and cross-modal translation, resulting in higher structural similarity and visual realism than CNN-based methods [7-10]. Nevertheless, even with such improvements, GAN models focused on accuracy can still generate hallucination artifacts in areas that are clinically significant, which may lower clinician confidence and impact reliability clinical practice.

2. GAN-based methods considering computational efficiency and training stability.

The alternative stream of research attempts to improve the efficiency and stability of GAN-based enhancement models. These methods are typically slow to learn and need careful hyperparameter optimization due to the unstable adversarial optimization. Also, scenarios in which the generator focuses too much on the overall look as opposed to the anatomical structures of the local tissues can lead to distortions in structure, which can cause loss of detailed features that can be used to rely on a diagnosis. These limitations also render the use of GAN-based models in real-time clinical practice challenging.

Overall, regardless of the promising accuracy improvement of the existing GAN-based approaches, the fact that they are weak in terms of enhancement reliability, as well as computational efficiency causes the development of more advanced and structure-aware enhancement models.

3. Methodology

3.1. Architecture overview

The general workflow scheme of the proposed IC-GAN framework is illustrated in Figure 1. An LDCT image is processed by a generator with attention blocks to emphasize clinically relevant structures and generate an NDCT-like image. The next step is to use a discriminator to evaluate the produced image and provide adversarial guidance. Simultaneously, perceptual, cycle-consistency, and identity losses are calculated to maintain the texture information and the anatomical features. All these loss elements work together to optimize the generator so that useful noise reduction and reliable structure preservation can be achieved in medical practice.

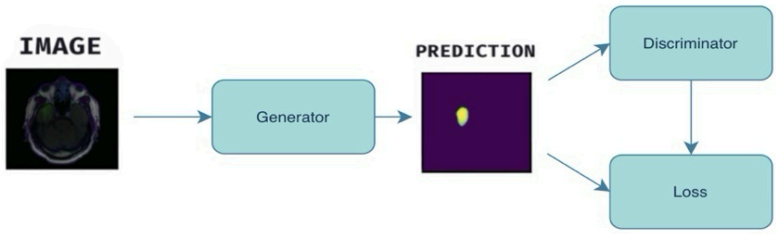


Figure 1. Overall architecture of the proposed IC-GAN framework

The adversarial training process is described by the IC-GAN framework based on the classical minimax objective that is used in generative adversarial networks. The aim is to represent the interaction of the generator G with the discriminator D in the form of a two-player optimization problem. The complete objective function reads:

$$\min_G \max_D V\left(D,G\right) = \mathbb{E}_{x\sim p_{dani}}\left[\log D\left(x\right)\right] + \mathbb{E}_{z\sim p_Z(z)}\left[\log(1-D\left(G\left(z\right)\right))\right] \quad (1)$$

DOI: 10.54254/2755-2721/2026.32322

It is an equation that explains how the two networks are competing and learning during training. The meanings of the symbols are presented as follows:

The output of the discriminator is the probability that the real sample x has been generated by the true data distribution.

- $G(z)$: The generator's output when a random vector z is used as input.
- $D(G(z))$: The discriminator's estimation of how likely the generated image G(z) is real.
- $\mathbb{E}_{x\sim p_{dani}}$: The distribution of real training samples, in this case NDCT images.
- $\mathbb{E}_{z\sim p_Z(z)}$: The prior distribution of the generator's input noise, often Gaussian or uniform.
- : The discriminator tries to raise this value by correctly labelling real images as real and generated images as fake.
- The generator tries to reduce the same value through the production of images that are indiscernible to the discriminator as the real ones.

This adversarial approach will prompt the generator to constantly enhance the quality of its outputs, whereas the discriminator will also get better at detecting unreal characteristics. Besides the traditional minimax goal, Zhu et al. also introduced the explicit formulation of cycle-consistency and identity loss in the CycleGAN model to limit unpaired image translation [11]. These extra loss terms ensure that the reconstruction is consistent and intensity preserving, and it has been demonstrated that they provide higher structural fidelity. The perceptual, cycle and identity losses are added to this minimax objective in IC-GAN creating a more stable and structure-preserving training strategy.

3.2. Generator design

The generator in the proposed IC-GAN framework is used to convert the low-dose CT input into an enhanced NDCT-like output. It has an encoder-decoder design with multiple residual blocks in the center, enabling the network to learn meaningful features without losing spatial details. The initial stage of the network is where the encoder decreases the spatial resolution and trains on simple intensity and texture features that are necessary to denoise.

Attention modules are used in the residual blocks to improve the retention of the anatomical details. They help the generator to focus on clinically relevant parts like tissue borders or delicate structures which are often not visible on lower doses. The generator may give extra importance to these information-rich areas to reduce noise at the expense of important textures.

After the feature extraction step, the decoder will upsample the features to re-construct the original resolution and offer information transfer between different layers. It will help rebuild more distinct structures and minimize the likelihood of developing too smooth or blurry areas. The improved CT image is the end result of the generator that is presented to the loss module to be optimized. Overall, the structure of the generator allows IC-GAN to balance both noise reduction and structural preservation, which is needed to perform clinical interpretation.

3.3. Discriminator design

The role of the discriminator in the IC-GAN architecture is the determination of the realness of the generated images relative to the real NDCT scans. The discriminator also works as a binary classifier, similar to a conventional GAN-based architecture, where it provides a probability indicating whether the input image is real or generated.

The discriminator is trained both on the real NDCT images and the improved outputs of the generator during the training process. It is trained to recognize inconsistencies, including unrealistic

DOI: 10.54254/2755-2721/2026.32322

textures, blurry edges, absence of anatomical properties, and gives adversarial responses to prompt the generator to create more believable and credible outputs.

The discriminator prompts the model to enhance its visual fidelity and preserve diagnostic important features through persistent competition with the generator. The output of the discriminator is a component of the overall loss function that is used in the training process, which needs to be stabilized to enhance the quality of the resulting CT images. Therefore, the discriminator will help to make sure that the results created by IC-GAN are enhanced and suitable to be clinically interpreted.

3.4. Loss functions

IC-GAN is minimised by loss functions, which assist the generator in generating clean and structurally correct CT images. They balance the values of realism, noise minimisation, and anatomical integrity, all of which are important in medical imaging. Adversarial loss is the first: it is based on the traditional GAN formulation. The discriminator is taught to discriminate between real and generated NDCT scans, but the aim of the generator is to generate images that are not easy to label as false and promote the enhanced images to have a similar overall look to actual NDCT scans.

In addition to adversarial learning, IC-GAN also applies a cycle-consistency loss to improve the structural stability of training. The idea behind it is that an LDCT image, translated by the generator, should retain its semantic information in the reverse mapping back into its original space and thus reduce the probability of unrealistically shaped structures as well as misplaced or missing information in case of absence of paired data. One more related part is identity loss which ensures that images of the target domain are not changed after passing through the generator, which helps to stabilize training, and does not make unnecessary changes.

There is also a perceptual loss used to strengthen structural preservation. This loss compares similarity on a higher semantic level, by computing features in a pretrained network, as opposed to comparing pixel-level differences, and so the generator is more sensitive to texture, edge, and fine anatomical detail. Through the combination of cycle-consistency and adversarial learning with perceptual loss, IC-GAN produces more meaningful CT images in both visual and clinical sense. Stable and effective training is ensured through a mix of weighted and optimized loss components.

4. Experiments

4.1. Experimental setup

In this part, the author describes the experimental setup which was used to determine the effectiveness of the proposed IC-GAN framework.

4.1.1. Dataset

The experimental analysis was performed through a real brain tumour MRI data, which is the LGG MRI Segmentation dataset available on Kaggle [12]. The data relied upon were the TCGA lower-grade glioma (LGG) collection that has been utilized in several other studies involving automatic tumour segmentation and medical image analysis [13]. As per the experimental design described in the paper of [13] it was divided into a training set and a test set so as to give a fair comparison of the performance. The dataset contains sufficient variability in tumour appearance as well as the quality of images enabling the assessment of the accuracy as well as the robustness of the suggested IC-GAN framework.

DOI: 10.54254/2755-2721/2026.32322

4.1.2. Compared algorithms

To compare the performance of the proposed approach, IC-GAN was compared to a standard Segmentation model based on GAN (GanSeg) [13]. GanSeg is an extremely strong reference architecture in medical image analysis tasks and has proven to be competitive in the past [14,15]. The comparison will serve as an attempt to analyze if the suggested enhancements to IC-GAN can beat the baseline approaches when subjected to the same experimental environment.

4.1.3. Evaluation metrics

Two important parameters of measuring the quality of the comparison techniques are accuracy and time. The accuracy indicates the ability of the model to locate the areas of the brain tumours that represent the quality of features and decision making [16]. Computational efficiency is measured by using time to determine whether there have been any performance improvements without too much computation overhead, which is essential, particularly when it is being used in real-time or near-real-time clinical application [17]. All these measures offer a fair assessment of effectiveness and efficiency.

4.2. Experimental results and analysis

Both results of experiments of IC-GAN and GanSeg are shown in this section and analyzed based on the accuracy and time point of view.

Figure 2 shows how precise the two methods are in their training epochs {1, 5, 10, 15, 20, 30, 40}. As it is depicted in Figure 2, the accuracy of both methods steadily increases with training. Precisely, the accuracy of GanSeg rises by 0.61 during the initial epoch to 0.77 during the epoch 40, whereas IC-GAN is improved by 0.64 to 0.84 during the same time period. In the training process, IC-GAN remains more accurate than GanSeg, and the mean difference between them is about 6-8 percent. The performance discrepancy is more obvious after some epochs (e.g. 0.84 versus 0.77 at epoch 40) which means that the suggested optimizations allow IC-GAN to have a stronger ability to learn discriminative properties and improve the tumor recognition results.

The enhancement is due to the improved model design and parameter optimisation strategy that was used in IC-GAN. The suggested method will ensure both the conservation of important anatomical properties and reduction of noise and non-relevant variability by imposing structural consistency constraints and close monitoring of the training parameters. Consequently, IC-GAN shows higher robustness and generalisation power as compared to the base model.

The Figure 3 compares the inference time of the IC-GAN and GanSeg. Inference time of the two methods is relatively insensitive to the number of training epochs {1, 5, 10, 15, 20, 30, 40} implying that the computational efficiency is not sensitive to the training progress. In particular, the inference time of GanSeg is 2.15 s-2.65 s and for IC-GAN it is 2.30 s to 2.80 s. Inference time of IC-GAN is a bit higher compared to that of GanSeg with a mean difference of around 0.1-0.15 seconds. Nonetheless, it is a negligible difference and is not a severe restriction when applied in practice.

In general, the outcomes of the experiment indicate that IC-GAN offers a positive trade-off in terms of accuracy and efficiency. Even though the suggested method has the same inference duration as its predecessor, it can be much more accurate in recognizing objects, which is why it fits better into practical medical imaging tasks where accuracy and computational rate also matter.

DOI: 10.54254/2755-2721/2026.32322

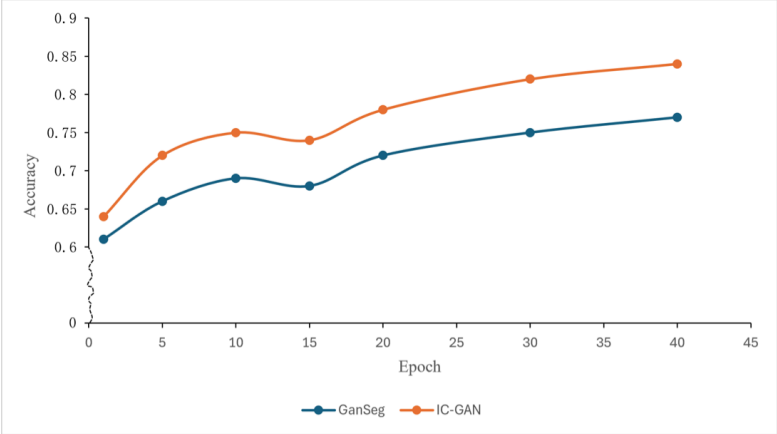


Figure 2. Comparison of recognition accuracy between GanSeg and IC-GAN

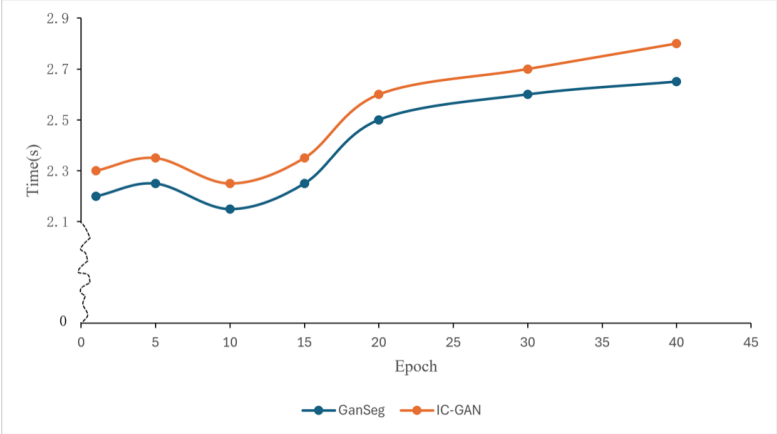


Figure 3. Comparison of time between GanSeg and IC-GAN

4.3. Experimental results summary

The suggested IC-GAN was superior to the baseline GanSeg approach in terms of the given brain tumour data by around 10 percent, whereas the inference time was comparable but with a minor rise in computational cost.

5. Conclusion

In this paper, an enhanced GAN structure of brain tumour recognition is presented, which was found to be more accurate but with comparable inference time. Experimental results show that the suggested IC-GAN offers a desirable balance between performance and efficiency, making it well suited to practical medical imaging use. In the future, the proposed solution can also be used in other areas such as autonomous driving and intelligent perception systems.

Moreover, the future study ought to examine the use of lightweight models and optimization methods to provide a significant level of recognition accuracy using reduced computation expenses and less time to draw conclusions. With such a small amount of analysis, we are of the opinion that more optimization of the algorithm might bring about another 5 percent better accuracy or efficiency.

DOI: 10.54254/2755-2721/2026.32322

References

[1] Zhang, Y., Liu, X., Chen, Z. and Wang, G. (2024) Noise reduction and structure preservation in low-dose CT imaging: A comparative study. Physics in Medicine & Biology, 69, 025012.
[2] Sadia, H., Rahman, M.M. and Islam, M.S. (2024) Data-efficient deep learning methods for low-dose CT image enhancement. Artificial Intelligence in Medicine, 147, 102726.
[3] Chen, X. (2024) Advances in deep learning-based low-dose CT image reconstruction and enhancement. Medical Physics, 51, 1452–1465.
[4] Kang, E., Min, J. and Ye, J.C. (2017) A deep convolutional neural network using directional wavelets for low-dose X-ray CT reconstruction. Medical Physics, 44, e360–e375.
[5] Chen, H., Zhang, Y., Zhang, W., Liao, P., Li, K., Zhou, J. and Wang, G. (2017) Low-dose CT denoising with convolutional neural network. IEEE Transactions on Medical Imaging, 36, 2524–2535.
[6] Zhu, J.-Y., Park, T., Isola, P. and Efros, A.A. (2017) Unpaired image-to-image translation using cycle-consistent adversarial networks. Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2223–2232.
[7] Kida, S., Nakamoto, T., Nakajima, T. and Akagi, M. (2019) Unpaired image translation for low-dose CT using CycleGAN. Computerized Medical Imaging and Graphics, 74, 1–10.
[8] Gu, Y. and Ye, J.C. (2021) Unsupervised domain adaptation for low-dose CT denoising using CycleGAN with structure preservation. Physics in Medicine & Biology, 66, 095015.
[9] Kwon, H. and Ye, J.C. (2021) Unpaired image translation for low-dose CT denoising with improved CycleGAN. IEEE Access, 9, 108131–108143.
[10] Yin, X., Li, Y., Zhang, H. and Wang, Z. (2023) Structure-preserving GAN for low-dose CT image enhancement. Medical Image Analysis, 85, 102725.
[11] Iacono, M. and Khan, A. (2024) Reliability and uncertainty in GAN-based medical image enhancement. IEEE Transactions on Medical Imaging, 43, 612–624.
[12] Kaggle Dataset (2025) LGG MRI segmentation dataset. Retrieved from https://www.kaggle.com/datasets/mateuszbuda/lgg-mri-segmentation
[13] GANSeg (2022) Generative adversarial network for medical image segmentation. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW).
[14] Sabnam, S., et al. (2024) Application of generative adversarial networks in image and medical data analysis: A review. Journal of Experimental & Theoretical Artificial Intelligence, 36, 233–247.
[15] Usama, Z., Khan, M. and Alam, F. (2025) A review on the applications of GANs for 3D medical imaging. Applied Sciences, 15, 11219.
[16] Taha, A.A. and Hanbury, A. (2024) Metrics for evaluating 3D medical image segmentation: Analysis and recommendations. Medical Image Analysis, 97, 103302.
[17] Shen, Y., Chen, R. and Li, T. (2025) Efficiency-aware evaluation of deep learning models for medical imaging applications. Medical Image Analysis, 99, 103412.