

Satellite Image Enhancement Using Generative Adversarial Networks

Gadi Mohan

Reg. No. 24Q71F0013

gadimohan@gmail.com

Department of Master of Computer Applications
Avanthi Institute of Engineering and Technology (Autonomous)
Vizianagaram, Andhra Pradesh, India

Under the guidance of Mr. P. Satyanarayana, M.Tech., (Ph.D.), Associate Professor
pukkallasatya84@gmail.com

Abstract—Satellite imagery is vital for applications such as environmental monitoring, urban planning, and disaster management, yet its quality is often degraded by atmospheric interference, sensor limitations, and transmission noise. Traditional enhancement techniques struggle to restore fine details, limiting the effectiveness of downstream analysis. Generative Adversarial Networks (GANs) offer a promising solution by learning complex data distributions to generate realistic, high-resolution outputs. In this work, a GAN-based framework for satellite image enhancement is presented that reconstructs sharper textures, reduces noise, and improves spatial resolution. The generator learns to map low-quality inputs to enhanced outputs, while the discriminator ensures fidelity to real high-quality references. Experimental evaluation demonstrates superior performance compared with conventional methods, with improvements in PSNR, SSIM, and visual clarity; enhanced images provide clearer terrain, vegetation, and urban structures, enabling more accurate geospatial analysis. The system is implemented in Python with a generator–discriminator architecture trained using adversarial and content losses and the Adam optimiser, and is deployed through a client–server web interface so users can upload images and receive enhanced outputs. The system was validated through nine functional test cases covering upload, preprocessing, GAN processing, output display, metric calculation, saving, and error handling, all of which passed. This study establishes GANs as a robust approach for advancing satellite image enhancement.

Keywords—Satellite Image Enhancement; Generative Adversarial Networks; Super-Resolution; Image Denoising; Deep Learning; PSNR; SSIM; Remote Sensing.

I. INTRODUCTION

Satellite imagery has become an essential resource in modern scientific, environmental, and technological applications. It plays a crucial role in monitoring the Earth's surface, analysing environmental changes, supporting disaster management, enabling urban planning, and assisting in agricultural assessment. However, despite advancements in satellite technology, the quality of captured images is often affected by unavoidable factors such as atmospheric disturbances, sensor noise, cloud cover, motion blur, and limitations in spatial and spectral resolution, which reduce clarity and make it difficult to extract accurate, meaningful information.

Traditional image-enhancement techniques such as histogram equalisation, filtering, and interpolation have been widely used to improve satellite images. While they can enhance contrast or reduce noise to

some extent, they are limited in reconstructing fine details and textures because they rely on predefined mathematical operations and fail to capture complex real-world patterns, making them insufficient for high-precision applications such as object detection, land-use classification, and change detection. Deep-learning models, particularly Convolutional Neural Networks (CNNs), perform well in image processing, but CNN-based enhancement sometimes produces overly smooth images and fails to generate realistic textures.

Generative Adversarial Networks (GANs) represent a significant breakthrough in deep learning and image generation. A GAN consists of two competing networks—a generator and a discriminator: the generator creates enhanced images from low-quality inputs, while the discriminator evaluates their authenticity against real high-quality images. Through this adversarial process, the generator gradually produces images that are visually indistinguishable from real data, achieving super-resolution, denoising, deblurring, and texture reconstruction simultaneously. By learning from large datasets of low- and high-quality image pairs, GANs improve spatial resolution, reduce noise, and preserve important features such as edges, patterns, and textures, while loss functions that consider both pixel-level accuracy and perceptual similarity preserve the semantic and structural integrity of geographical features. Advanced architectures such as conditional GANs, SRGAN, and ESRGAN further improve performance, though GAN training can be unstable and data-intensive.

II. LITERATURE SURVEY

Recent advancements in deep learning have significantly improved image-enhancement techniques, particularly for satellite imagery. Traditional methods such as interpolation, histogram equalisation, and filtering were widely used but often fail to reconstruct fine details and textures. With deep neural networks, especially CNNs, researchers achieved better performance in super-resolution, denoising, and deblurring; however, CNN-based approaches tend to optimise pixel-wise accuracy, which may result in overly smooth images lacking realistic details.

The emergence of Generative Adversarial Networks, introduced by Ian Goodfellow in 2014, marked a breakthrough in image-enhancement tasks. Models such as SRGAN by Ledig et al. and ESRGAN by Wang et al. demonstrated the ability to generate high-resolution images with sharper textures and perceptually convincing details by introducing perceptual loss and adversarial learning, significantly improving visual quality compared with traditional CNN-based methods. In the context of satellite imagery, GAN-based approaches have been increasingly adopted to overcome low resolution, atmospheric distortion, and noise, with enhancements including attention mechanisms, multi-scale architectures, and hybrid CNN-GAN models. Despite these advancements, training instability, high computational requirements, and generalisation across datasets remain challenges, motivating ongoing research into more robust and efficient GAN-based models.

TABLE I. SUMMARY OF REPRESENTATIVE PRIOR WORK

S.No	Author(s) / Year	Methodology	Contribution	Limitation
1	Goodfellow et al., 2014	GAN framework	Introduced adversarial learning	Training instability

S.No	Author(s) / Year	Methodology	Contribution	Limitation
2	Ledig et al., 2017	SRGAN (residual GAN)	Perceptual loss for super-resolution	Artifacts in some cases
3	Wang et al., 2018	ESRGAN (RRDB)	Sharper textures than SRGAN	Computationally expensive
4	He et al., 2016	Deep residual learning	Strong feature backbone	Not enhancement-specific
5	Zhang et al., 2018	Residual Dense Network	Improved super-resolution	High complexity
6	Liu et al., 2020	Multi-scale GAN	Better image enhancement	Resource demands

III. EXISTING SYSTEM AND PROPOSED SYSTEM

A. Existing System

The existing system for satellite image enhancement primarily relies on traditional image-processing techniques and basic machine-learning methods such as histogram equalisation, contrast stretching, interpolation, filtering, and noise reduction. While these provide basic improvements, they are limited in restoring fine details and textures because they are based on predefined mathematical operations and do not adapt well to complex variations such as atmospheric interference, sensor noise, cloud cover, and low spatial resolution. CNN-based approaches introduced later can learn features but often produce overly smooth images lacking realistic textures, and many traditional methods require manual parameter tuning, making them inefficient and poorly scalable.

Limitations of the existing system:

- Poor reconstruction of fine details and textures.
- Limited adaptability to complex degradations (haze, noise, blur).
- CNN-based methods produce overly smooth, unrealistic textures.
- Manual parameter tuning is inefficient and time-consuming.
- Lack of scalability for large volumes of satellite data.

B. Proposed System

The proposed system introduces a deep-learning approach using Generative Adversarial Networks for satellite image enhancement. A GAN architecture with a generator and a discriminator is used: the generator takes low-quality satellite images and produces enhanced high-resolution images, while the discriminator compares generated images with real high-quality images and judges their realism; through adversarial training the generator improves its output. The system performs super-resolution, denoising, deblurring, and texture reconstruction simultaneously, learns features automatically without manual parameter tuning, preserves important geographical features, supports evaluation using PSNR and SSIM, and is scalable and integrable with a web interface for real-time use.

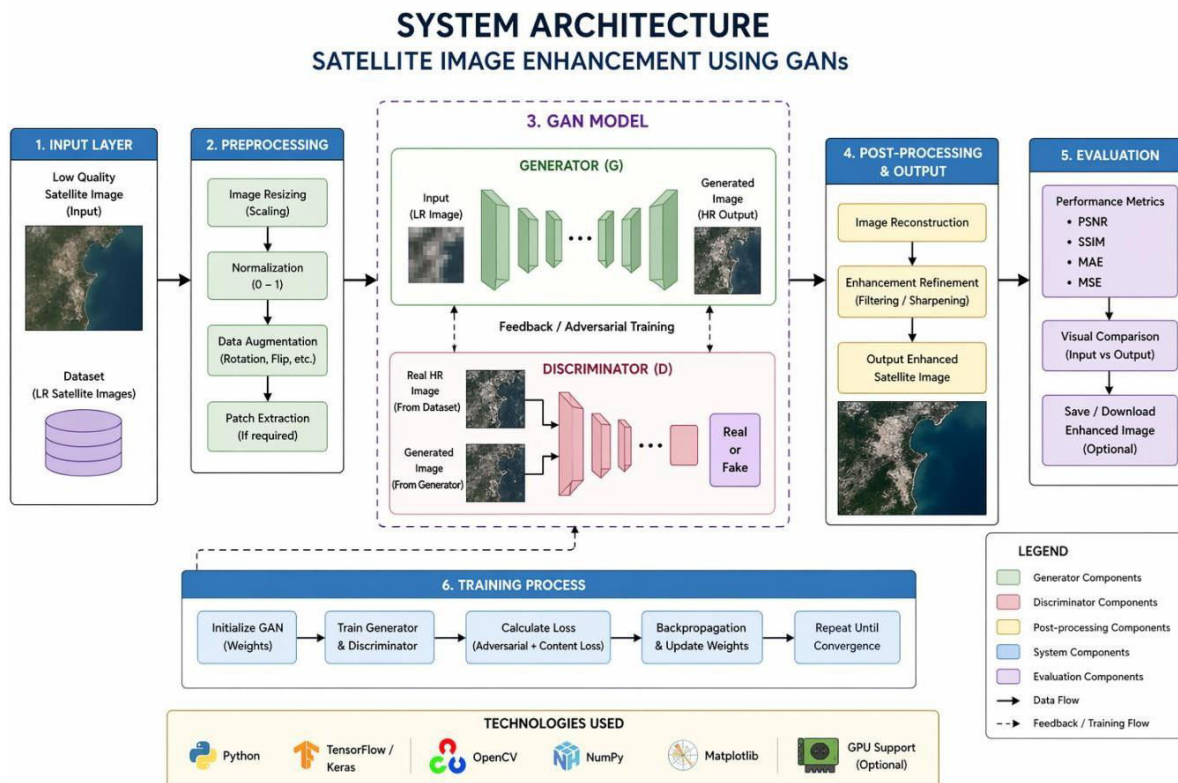
Advantages of the proposed system:

- Realistic, high-resolution enhancement via adversarial learning.
- Simultaneous super-resolution, denoising, deblurring, texture reconstruction.
- Automatic feature learning without manual parameter tuning.
- Preservation of edges, textures, and geographical features.
- Quality assessment using PSNR and SSIM metrics.
- Scalable and deployable through a web-based interface.

IV. SYSTEM DESIGN AND METHODOLOGY

A. GAN Architecture

The core of the system is the GAN architecture, consisting of a generator and a discriminator. The generator transforms low-quality satellite images into enhanced high-resolution outputs and is designed using convolutional layers, residual blocks, and upsampling techniques to capture both global and local features. The discriminator distinguishes between real high-quality images and generated images using convolutional layers followed by activation functions. Both networks are trained simultaneously in an adversarial manner: the generator tries to fool the discriminator while the discriminator learns to correctly identify real and generated images.



B. Training

During training, the generator receives low-quality images and produces enhanced images; these, along with real high-quality images, are fed to the discriminator, which estimates the probability that an image is real or fake, and the loss is computed accordingly. The generator is trained with a combination of adversarial loss and content loss (such as mean squared error or perceptual loss) to ensure both visual quality and accuracy, and the Adam optimiser updates the weights of both networks. Training is repeated over multiple epochs until the model converges and produces satisfactory results.

C. Deployment Architecture

Deployment uses a client-server architecture. The front-end interface is developed using web technologies so users interact through a browser, while the back-end server hosts the GAN model and handles image processing. Frameworks such as Flask or Django are used for the server-side application; the trained model is loaded on the server, and API endpoints handle image upload, processing, and output delivery. The system can also be deployed on cloud platforms for scalability and accessibility.

V. SYSTEM IMPLEMENTATION

A. Technology Stack

TABLE II. TECHNOLOGY STACK

Component	Technology / Tool
Programming Language	Python
Model	GAN (generator + discriminator; conv + residual + upsampling)
Losses	Adversarial loss + content loss (MSE / perceptual)
Optimiser	Adam
Backend Framework	Flask / Django
Frontend	Web technologies (browser-based client)
Evaluation Metrics	PSNR, SSIM
Deployment	Client-server; optional cloud

B. Implementation Pipeline

Low- and high-quality satellite image pairs are prepared and the input is preprocessed (resizing and normalisation). The generator is implemented with convolutional layers, residual blocks, and upsampling, and the discriminator with convolutional layers and activation functions. The two networks are trained adversarially using the combined loss and the Adam optimiser over multiple epochs. After training, the

generator is integrated into the back-end server, and API endpoints handle image upload, GAN-based enhancement, metric computation, and delivery of the enhanced output to the browser client.

C. Quality Evaluation

Enhanced images are evaluated using Peak Signal-to-Noise Ratio (PSNR) and the Structural Similarity Index Measure (SSIM), together with visual comparison against conventional methods. These metrics assess pixel-level fidelity and structural similarity to high-quality references; the source reports improvements qualitatively, and no specific numeric PSNR/SSIM values are asserted here.

VI. SYSTEM TESTING AND RESULTS

The system was validated through nine functional test cases covering valid image upload, invalid file-format handling, preprocessing, GAN model processing, output display, performance-metric calculation, saving the enhanced image, corrupted-image handling, and batch processing of multiple images. All test cases passed and behaved as expected.

TABLE III. REPRESENTATIVE TEST CASES

ID	Scenario	Input	Expected Output	Status
TC01	Upload valid image	JPG/PNG satellite image	Image uploaded successfully	Pass
TC02	Invalid file format	Text/PDF file	File rejected with error	Pass
TC03	Preprocessing	Low-resolution image	Image resized and normalised	Pass
TC04	GAN model processing	Preprocessed image	Enhanced HR image generated	Pass
TC06	Metrics calculation	Generated image	PSNR and SSIM displayed	Pass
TC07	Save enhanced image	Enhanced output	Image saved/downloaded	Pass
TC08	Corrupted image	Damaged image file	Error shown; no crash	Pass

A. Observed Results

The implementation shows that GAN-based approaches can effectively overcome common issues in satellite imagery such as noise, blur, and low resolution. The generator–discriminator framework produces realistic high-resolution images that are significantly better than those from traditional enhancement methods, and the evaluation metrics PSNR and SSIM, along with visual comparisons, confirm that the enhanced images preserve important spatial details and structural information. This makes the system useful for environmental monitoring, urban planning, agriculture, and defence analysis. The source reports these results qualitatively; high computational requirements and training instability remain challenges.

Representative screenshots from the prototype implementation:

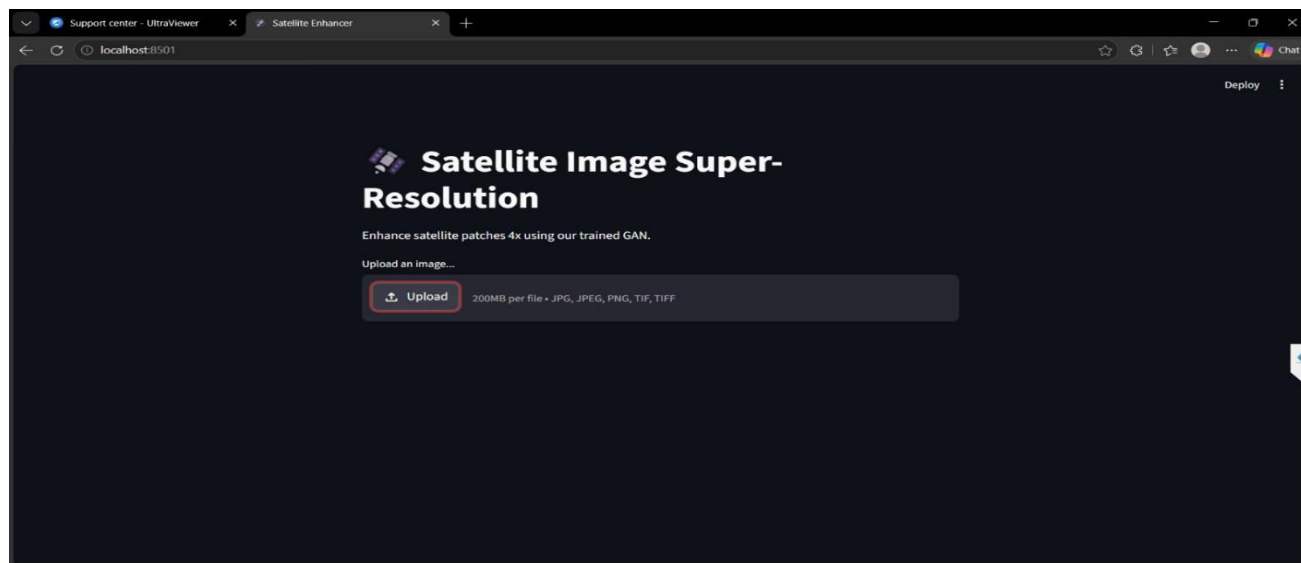


Fig. 1. Satellite-image upload interface.

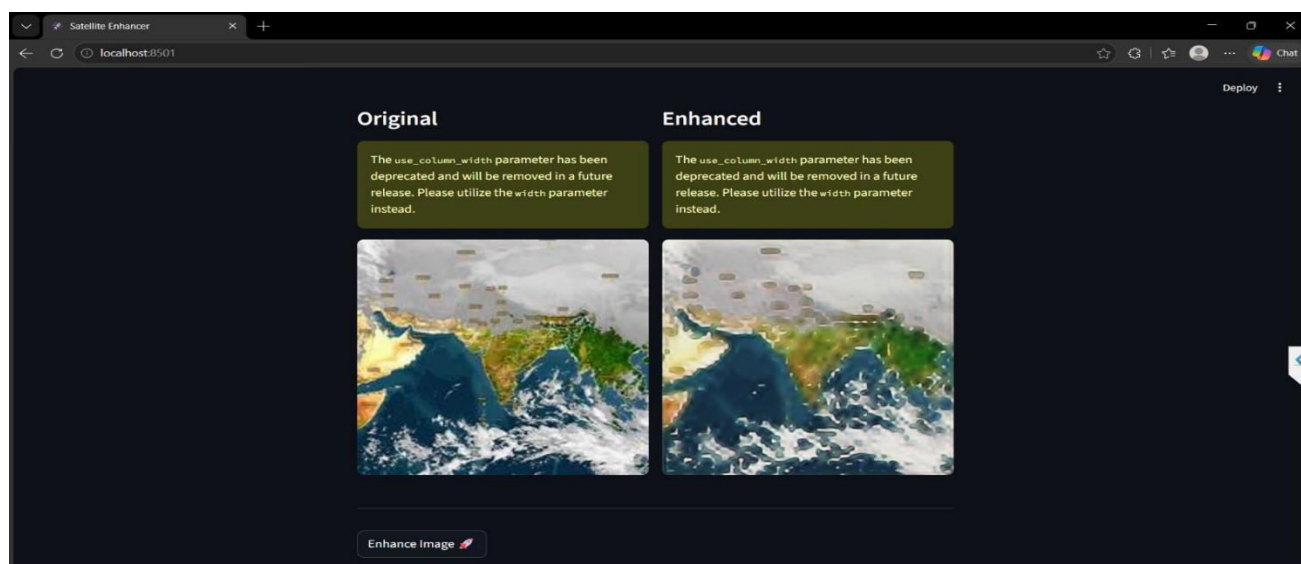


Fig. 3. GAN-enhanced output vs. input.

VII. CONCLUSION AND FUTURE SCOPE

The project successfully demonstrates the application of deep-learning techniques to improve the quality of satellite images. By utilising Generative Adversarial Networks, the system enhances low-resolution images by improving clarity, sharpness, and overall visual quality, and the generator–discriminator framework enables the model to produce realistic high-resolution images that are significantly better than those obtained through traditional enhancement methods. The implementation shows that GAN-based approaches can effectively overcome common issues such as noise, blur, and low resolution, and the

evaluation metrics PSNR and SSIM, along with visual comparisons, confirm that the enhanced images preserve important spatial details and structural information, making the system highly useful for environmental monitoring, urban planning, agriculture, and defence analysis. The project also highlights challenges including high computational requirements and training instability, but overall the results prove that GANs are a powerful tool for satellite image enhancement.

The project can be extended in several directions. Advanced GAN architectures such as ESRGAN, StyleGAN, or transformer-based GANs, and hybrid models combining CNNs and transformers, can produce more realistic, higher-quality images with better texture detail. Real-time image enhancement, optimisation of model efficiency for faster inference, training on larger and more diverse datasets, and deployment on scalable cloud infrastructure would further improve performance, usability, and real-world applicability.

REFERENCES

- [1] I. Goodfellow, J. Pouget-Abadie, M. Mirza, et al., “Generative Adversarial Nets,” in Proc. Advances in Neural Information Processing Systems (NIPS), 2014, pp. 2672–2680.
- [2] C. Ledig, L. Theis, F. Huszár, et al., “Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2017, pp. 4681–4690.
- [3] X. Wang, K. Yu, S. Wu, et al., “ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks,” in Proc. European Conf. Computer Vision (ECCV) Workshops, 2018.
- [4] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016.
- [5] Y. Zhang, Y. Tian, Y. Kong, et al., “Residual Dense Network for Image Super-Resolution,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2018.
- [6] J. Yu, Z. Lin, J. Yang, et al., “Generative Image Inpainting with Contextual Attention,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2018, pp. 5505–5514.
- [7] Z.-S. Liu, L. Li, S. Wu, and W. Zuo, “Multi-Scale Generative Adversarial Network for Image Enhancement,” IEEE Trans. Image Processing, vol. 29, pp. 8624–8636, 2020.
- [8] K. Zhang, W. Zuo, and L. Zhang, “Plug-and-Play Image Restoration with Deep Denoiser Prior,” IEEE Trans. Pattern Analysis and Machine Intelligence, vol. 43, no. 2, pp. 636–647, 2021.
- [9] J. Liang, J. Cao, G. Sun, et al., “SwinIR: Image Restoration Using Swin Transformer,” in Proc. IEEE Int. Conf. Computer Vision Workshops (ICCVW), 2021.
- [10] S. Ioffe and C. Szegedy, “Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift,” in Proc. Int. Conf. Machine Learning (ICML), 2015.