Low-Light Image Super-Resolution Using GANs

Thesis Submitted
in Partial Fulfillment of the Requirements for the Degree of

MASTER OF TECHNOLOGY
IN
SOFTWARE ENGINEERING

Submitted by

NIHARIKA PACHORI

2K23/SWE/22

Under the supervision of

Dr. SANJAY PATIDAR

Assistant Professor

Department of Software Engineering

DEPARTMENT OF SOFTWARE ENGINEERING
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Bawana Road, Delhi 110042

MAY, 2025

DEPARTMENT OF SOFTWARE ENGINEERING
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Bawana Road, Delhi-110042

CANDIDATE'S DECLARATION

I, **NIHARIKA PACHORI**, Roll No's – **2K23/SWE/22.** student of M.Tech (**Department of Software Engineering**),hereby declare that the project Dissertation titled "**Low-Light Image Super-Resolution Using GANs** " which is submitted by me to the **Department of Software Engineering**, Delhi Technological University, Delhi in partial fulfilment of the requirement for the award of degree of Master of Technology, is original and not copied from any source without proper citation. This work has not previously formed the basis for the award of any Degree, Diploma Associateship, Fellowship or other similar title or recognition.

Place: Delhi Niharika Pachori

Date: 30/05/2025

DEPARTMENT OF MECHANICAL ENGINEERING
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Bawana Road, Delhi-110042

CERTIFICATE

I hereby certify that the Project Dissertation titled "**Low-Light Image Super-Resolution Using GANs**" which is submitted by **NIHARIKA PACHORI**, Roll No. – **2K233/SWE/22**, **Department of Software Engineering** ,Delhi Technological University, Delhi in partial fulfilment of the requirement for the award of the degree of Master of Technology, is a record of the project work carried out by the student under my supervision. To the best of my knowledge this work has not been submitted in part or full for any Degree or Diploma to this University or elsewhere.

Place: Delhi | **Dr. Sanjay Patidar**

Date: 30.05.2025 | **SUPERVISOR**

DEPARTMENT OF SOFTWARE ENGINEERING
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Bawana Road, Delhi-110042

ACKNOWLEDGEMENT

I wish to express our sincerest gratitude to **Dr. SANJAY PATIDAR** for his continuous guidance and mentorship that he provided us during the project. He showed us the path to achieve our targets by explaining all the tasks to be done and explained to us the importance of this project as well as its industrial relevance. He was always ready to help us and clear our doubts regarding any hurdles in this project. Without his constant support and motivation, this research would not have been successful.

Place: Delhi NIHARIKA PACHORI

Date: 30.05.2025

Abstract

Image acquisition under low-light conditions poses serious limitations across numerous imaging domains, resulting in noisy, low-contrast, and resolution-degraded outputs. These limitations not only impact visual clarity but also hinder performance in downstream tasks such as detection, recognition, and interpretation. Traditional image enhancement techniques, including histogram equalization and gamma correction, provide limited improvement in complex low-light scenarios and often amplify noise or distort colours. In contrast, Generative Adversarial Networks (GANs) have demonstrated significant success in both enhancing brightness and performing super-resolution in a data-driven manner. Their ability to model complex visual distributions enables the recovery of realistic textures and structures from degraded inputs.

This thesis presents a comprehensive comparative review of recent GAN-based approaches for low-light image super-resolution. We explore key architectural strategies, loss functions, dataset choices, and evaluation metrics across prominent models. The analysis addresses three core research questions: limitations in texture restoration, effectiveness of performance metrics, and generalization challenges in low-light super resolution models across diverse scenarios. Furthermore, we highlight real-world application areas including surveillance, autonomous systems, mobile imaging, and document analysis where these techniques are most impactful. The paper concludes by identifying persistent challenges and proposing future research directions aimed at improving perceptual realism and robustness in low-light SR systems.

Contents

List of Tables

List of Figures

List of Symbols

μ_I	Mean intensity of image I
μ_K	Mean intensity of image K
σ_I^2	Variance of image I
σ_K^2	Variance of image K
σ_{IK}	Covariance between images I and K
LPIPS	Learned Perceptual Image Patch Similarity
MSE	Mean Squared Error between two images
PSNR	Peak Signal-to-Noise Ratio
SSIM	Structural Similarity Index Measure
C_1, C_2	Stabilization constants to avoid division by zero
d	Euclidean or cosine distance between feature maps
$f_l(\cdot)$	Feature map from layer l of a neural network
I	Original image
K	Compressed or reconstructed image
M	Maximum possible pixel value of the image (e.g., 255)
m, n	Pixel indices (row and column)
w_l	Weight for layer l in perceptual loss

Chapter 1

INTRODUCTION

Images taken in low-light situations frequently encounter various quality problems, such as reduced visibility, high sensor noise, and unnatural color changes. These challenges arise due to inadequate lighting, which restricts the camera's capability to capture fine textures and color details. Consequently, the overall visual quality deteriorates, making it challenging to interpret or process these images for subsequent tasks. This becomes particularly challenging in practical situations like night-time surveillance, autonomous navigation, or low-light photography, where visibility and detail are crucial. For example, footage captured by night vision security cameras often appears grainy and lacks clarity, making it challenging to recognize faces or identify important elements within the scene. It becomes even more challenging when super-resolution techniques are applied to distorted videos, as typical models often introduce noise or blur out important details instead of accurately reconstructing high-resolution outputs. The main challenge in this thesis is to address the twin issues of improving both the quality and clarity of images taken in low light conditions.[1]

1.1 Problem Statement

Images taken in low-light conditions tend to have compromised visibility, weak contrast, color aberration, and major noise. Such degradations are critical problems in surveillance, autonomous vehicles, and medical imaging scenarios where visual brightness is paramount. When the super-resolution technique is applied to the low-quality images, the issue is compounded by the fact that noise and artifacts are amplified in the upscaling process. Classical image recovery and super-resolution techniques are not very good at bringing out details in such situations. Hence, there exists a strong need for smart, data-driven techniques capable of brightening and upscaling simultaneously. This thesis applies the technology of Generative Adversarial Networks (GANs), which have been finding excellent applications in the task of image generation, to tackle the twin issue of brightening and recovering finer details of low-light images.[2]

1.2 Traditional (Non-Deep Learning) Solutions

1.2.1 Conventional Image Enhancement Techniques

Prior to the era of deep learning, some conventional techniques were commonly employed for low-light image improvement. One such technique that was commonly employed was histogram equalization, which sought to enhance image contrast by more evenly redistributing pixel intensity values. This technique served well for lightening dark images but tended to cause over-enhancement or details loss in certain areas. Another widely used method was founded on Retinex theory, which tried to simulate the means by which human vision perceives light and color by disconnecting illumination and reflectance. Retinex-based algorithms were very capable of boosting image brightness but tended to be computationally complex and noise-sensitive. Moreover, denoising filters like median filtering and Gaussian smoothing were employed to mute unwanted noise, but these had a tendency to blur the details of images and lose sharpness.[2]

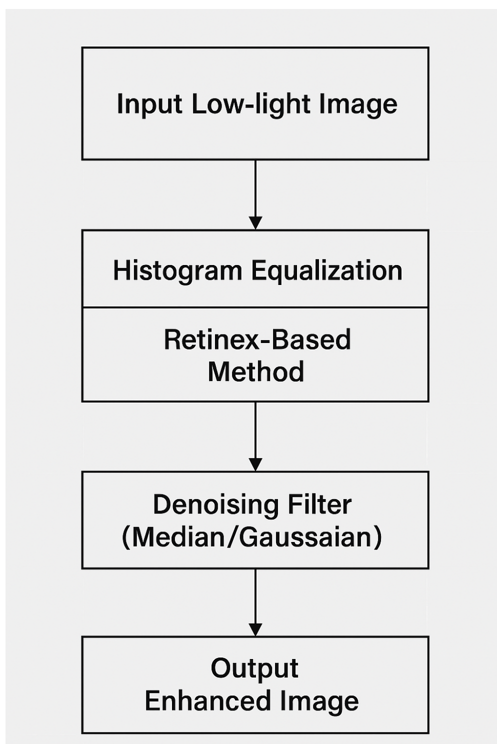

Figure 1.1: Conventional Image Enhancement Techniques

1.2.2 Traditional Super-Resolution Techniques

Super-resolution techniques before deep learning depended primarily on interpolation or sparse representation. Bicubic interpolation, for instance, was a simple method that used weighted averages of nearby pixels to estimate pixel values. Though simple to train, it tended to create overly smooth images without fine textures. Sparse representation-based techniques, however, tried to reconstruct high-resolution images from a learned dictionary of image patches. These approaches worked better in maintaining structural information but involved heavy feature engineering and manual parameter tuning, which made them less practical to use in varying lighting conditions.[2]

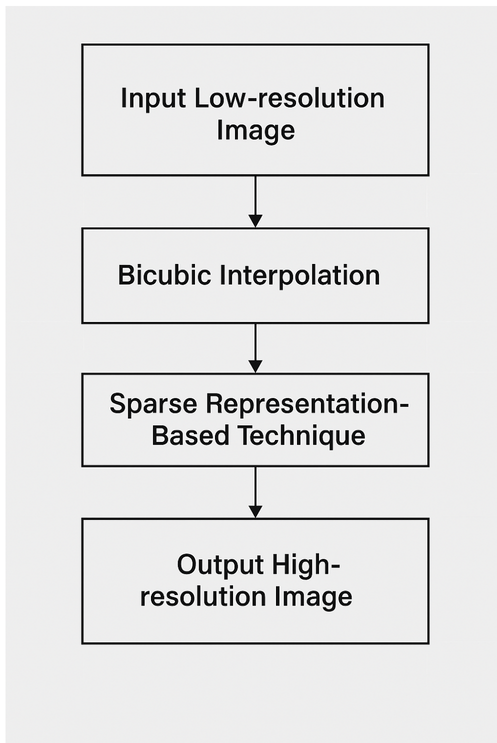

Figure 1.2: Traditional Super-Resolution Techniques

1.2.3 Limitations of Traditional Methods

Although useful in some situations, conventional image enhancement and super-resolution techniques suffered crucial limitations. A key limitation was their failure to generalize between images of different classes or light settings. They used rigid rules and manually crafted features, which restricted flexibility. Additionally, they were incapable of learning sophisticated patterns or representations from data and thus were ineffective where high variability was involved, like natural low-light images. In most instances, these methods also inadvertently enhanced image noise in the process of trying to brighten or sharpen, which lowered the general visual quality rather than enhancing it.

1.3 Low-light Image Super Resolution

It pertains to the process of improving the clarity and visual appeal of images taken in low-light situations. These images frequently exhibit low contrast, high noise, and lack of detail, posing challenges for standard super-resolution algorithms to generate high-quality outputs. The dual challenge here lies in simultaneously resolving the degradation caused by low illumination and enhancing the spatial resolution to reveal finer image structures. This problem is particularly critical in domains like surveillance, autonomous driving, and medical imaging, where visual clarity under suboptimal lighting is essential.[3]

To address LLISR, traditional methods relied on cascaded approaches where denoising or enhancement was first applied, followed by conventional super-resolution techniques like bicubic interpolation or sparse representation. However, these pipeline methods often failed to handle the complex interplay between noise, blur, and low light, leading to oversmoothed or artifact-heavy outputs. Recently, Generative Adversarial Networks (GANs) have come up as powerful tools for LLISR due to their ability to learn intricate mappings between low-quality and high-quality image domains. GAN-based models can jointly optimize enhancement and upscaling in a single framework, often using perceptual losses, attention mechanisms, and adversarial training to produce sharper, more natural-looking results. Architectures like SRGAN, ESRGAN, EnlightenGAN, and LE-GAN have demonstrated that GANs can reconstruct realistic textures and maintain structural fidelity even under extreme low-light conditions.[3]

1.4 Challenges of Low-Light Image Super-Resolution

Low-light super-resolution of images is a hard problem that integrates the challenge of two intrinsically hard problems: image upscaling and image enhancement. Images taken in suboptimal light conditions tend to have high noise, low contrast, and color aberrations, which make it difficult to recover high-frequency information. When super-resolution is used for such degraded images as input, it may inadvertently add noise and artifacts, resulting in artificial or distorted results. A significant difficulty is to provide more light without overexposing the bright areas or compromising texture detail. Additionally, the absence of large-scale paired datasets of low-light and corresponding high-resolution images makes supervised training difficult. Maintaining natural color tones and preserving structural details during enhancement and resolution reconstruction adds further complexity. Finally, most deep learning models, especially GAN-based approaches, require high computational resources, making it difficult to deploy them in real-time or resource-constrained environments. Addressing these issues is critical for achieving effective and reliable LLISR in practical applications.[4]

1.5 GANs

Generative Adversarial Networks, or GANs, is a notable development in generative modeling. Since their introduction by Ian Goodfellow in 2014, GANs have been a major focus of the research community. They are categorized as two neural networks, a generator and a discriminator, that are trained simultaneously in a game where one's success depends on the other's failure. The generator attempts to produce data samples that closely resemble real data, even when the input is noisy or of low quality. On the other hand, the discriminator's task is to differentiate between genuine data and the output generated by the generator. By employing an adversarial training method, the generator gradually improves its performance, becoming capable of producing outputs that are nearly indistinguishable from real samples. This architecture is especially effective in image generation tasks because it learns complex distributions and visual patterns that traditional models cannot capture, leading to high-quality, realistic outputs.[4]

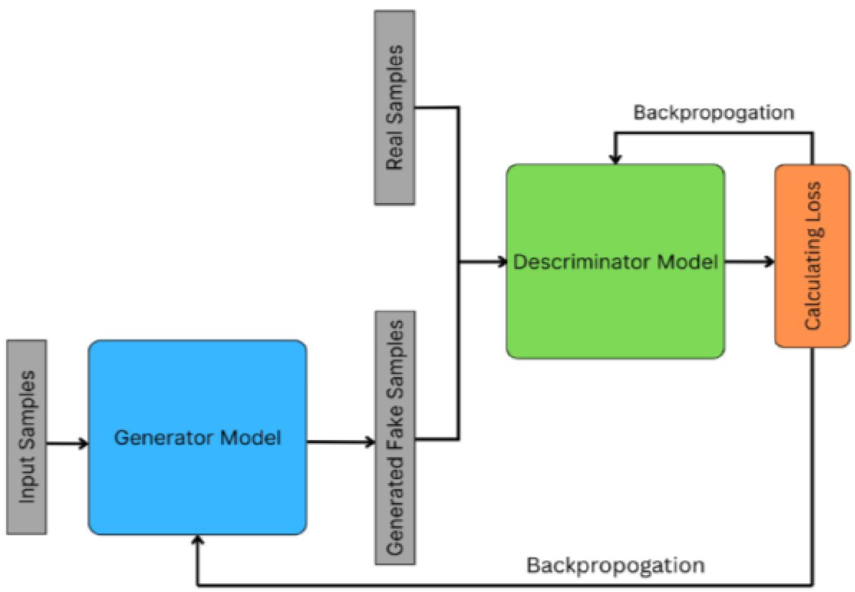

Figure 1.3: Generative Adversarial Network Architectural Diagram

1.5.1 Application of GANs to Image SR and Low-Light Enhancement

In the domain of image super-resolution, GANs have shown impressive results by focusing on perceptual realism rather than mere pixel accuracy. SRGAN (Super-Resolution GAN) was among the first models to apply adversarial loss for generating photo-realistic high-resolution images. ESRGAN (Enhanced SRGAN) built upon this foundation, improving image detail restoration and texture sharpness by refining network design and introducing the perceptual loss. On the other hand, low-light image enhancement has also benefited greatly from GANs. Models like EnlightenGAN and LE-GAN are designed to brighten dark images while preserving color consistency and structural integrity. They use unpaired training strategies and attention mechanisms to adaptively enhance image regions based on illumination. Some GAN architectures go a step further by tackling both low-light enhancement and super-resolution in a single framework. For instance, LLD-GAN and StarSRGAN are joint models that not only enhance brightness but also recover high-resolution details from severely degraded inputs. LAE-GAN focuses on improving the readability of low-light text images by combining attention modules and resolution enhancement. These models illustrate how GANs can be tailored to address multiple degradation issues simultaneously.[4]

1.5.2 Key Benefits of GAN-Based Approaches

GAN-based models offer several advantages over traditional and early deep learning methods. Their key advantage is that they can create high-perceptual-quality images, usually producing results that look more natural and visually appealing to human viewers. This is especially useful in low-light and super-resolution applications where fine details and realistic illumination are important. Compared to models that only reduce pixel-wise errors (such as MSE or L1 loss), GANs combine adversarial loss and sometimes perceptual loss and thus can learn features that correspond more closely to human perception. Another major plus is their end-to-end learnability. The GAN model can learn the entire transformation of low-quality images to high-quality images without manual-crafted features or intricate pre-processing pipelines. This end-to-end architecture simplifies model training and deployment, reducing the complexity of GANs and making them more scalable and flexible for application in real-world scenarios like night-time photography, surveillance videos, and autonomous driving during low-light environments.[4]

1.6 Challenges

1.6.1 Dual Problem: Enhancing Illumination While Increasing Resolution

Low-light super-resolution for images is fundamentally a multi-task task. The model must, in one respect, illuminate dim images and restore proper illumination. At the same time, it must enhance the resolution and reclaim lost fine details resulting from low-quality capture. These two goals are usually opposing, with illumination potentially introducing artifacts or blurring and upscaling requiring sharpness and high-frequency information. Creating a model that will be able to perform both tasks without sacrificing either is a great challenge in this field.[5]

1.6.2 Noise Amplification During Upscaling

Low-light photos are usually shot with elevated ISOs or longer exposure times, thus generating enormous sensor noise. When those noisy photos go through super-resolution models, the upscaling can actually enhance the noise as well as the underlying fine details. This creates outputs sharper but aesthetically unpleasing as a result of discernible artifacts. Mitigating this issue requires models to intelligently distinguish between noise and useful features, which remains a complex and unresolved problem.

1.6.3 Lack of Paired Datasets for Training

Supervised learning approaches often rely on paired datasets, where each low-light input image has a corresponding high-resolution, well-lit ground truth. However, collecting such datasets is difficult, especially under controlled lighting and camera conditions. As a result, many existing models must rely on synthetic data or unpaired training, which may not generalize well to real-world scenarios. The scarcity of high-quality, real-world paired datasets hinders both the development and benchmarking of low-light image super-resolution models.

1.6.4 Maintaining Color Accuracy and Texture Details

Accurate color reproduction and texture preservation are essential for natural-looking image enhancement. Low-light images often suffer from color shifts and loss of texture due to poor illumination. Enhancing brightness and resolution without introducing unnatural colors or smoothing out important textures is a difficult task. Models must carefully balance between contrast enhancement and fine-detail recovery, which requires sophisticated loss functions and architectural components, such as attention mechanisms or perceptual guidance.

1.6.5 Computational Complexity of GAN Models for Real-Time Use

While GAN-based methods offer high perceptual quality, they are often computationally heavy. Complex architectures with deep convolution layers, multiple attention modules, and high-resolution output generation demand significant GPU resources and time. This makes it challenging to deploy such models in real-time or resource-constrained environments like mobile devices, surveillance systems, or embedded cameras. Reducing the computational load without compromising quality is a key area of ongoing research.

The organization of this thesis aims to offer a thorough understanding of low-light image super-resolution using GANs. Chapter 2 provides an extensive literature review, encompassing traditional image enhancement techniques, early convolutional neural network (cnn)-based approaches, and recent generative adversarial network (gan)-based methods that are applicable to the field. Chapter 3 discusses the main obstacles encountered in this area and establishes the goals that shape the focus of the review. Chapter 4 discusses the approach used for this study, covering the criteria for model selection, datasets utilized, and evaluation metrics employed for comparing performance. Chapter 5 delves into the practical applications of low-light image super-resolution, emphasizing its importance in fields like surveillance, medical imaging, and autonomous systems. Chapter 6 offers a comprehensive comparison of eight cutting-edge generative models, examining their respective strengths and weaknesses. Chapter 7 proposes potential areas for future research, identifying gaps and emerging trends in this rapidly changing field. Finally, chapter 8 concludes the thesis by summarizing the main findings and contributions derived from the review.

Chapter 2

LITERATURE REVIEW

2.1 Classical Foundations of Image Enhancement

Pizer et al. (1987): Adaptive Histogram Equalization, this foundational work introduced adaptive histogram equalization (AHE), a contrast-enhancement technique that locally adjusts image intensities to improve visibility in poorly illuminated regions. Unlike global histogram equalization, AHE operates on sub-regions of an image, making it effective for enhancing details in low-light conditions. However, it often amplifies noise in smooth areas, limiting its utility for real-world applications. The paper laid the groundwork for subsequent spatial domain enhancement methods, but lacked mechanisms to address the photometric distortions inherent in extreme low-light scenarios [5].

Gonzalez & Woods (2008): Digital Image Processing, this comprehensive textbook systematized traditional image processing techniques, including spatial filtering, frequency domain methods, and histogram manipulation. Although it covers Retinex-inspired approaches for illumination correction, its focus on classical algorithms (e.g., Wiener filtering) does not address the non-linear degradations in modern low-light imaging. The text remains a critical reference for understanding the mathematical foundations of image enhancement, but highlights the need for data-driven solutions under complex lighting conditions [6].

Land & McCann (1971): Proposing the Retinex theory, this seminal work modeled human color constancy by separating illumination from reflectance-a concept that later inspired computational low-light enhancement frameworks. The core premise of the theory-that perceived color depends on relative lightness comparisons rather than absolute intensities-directly influenced modern GAN architectures to disentangle illumination and content. Although limited by its hand-crafted assumptions, Retinex provided a theoretical basis for physics-informed deep learning models [7].

2.2 Deep Learning for Super-Resolution

Ledig et al. (2017): pioneered GAN-based super-resolution by introducing a perceptual loss function using VGG features, shifting the focus from pixel-wise accuracy (PSNR/SSIM) to human visual perception. The residual generator and PatchGAN discriminator enabled 4× up-scaling with photorealistic textures, though it struggled with low-light inputs due to unmodeled noise and illumination biases. This work established adversarial training as viable for ill-posed restoration tasks, later adapted for joint low-light enhancement and SR [8].

Wang et al. (2018): ESRGAN advanced SRGAN by replacing residual blocks with residual-in-residual-dense blocks (RRDB) and adopting relativistic discriminators. These innovations improved texture recovery and stabilized training for 4×–8× up-scaling. The modified perceptual loss using pre-activation VGG features better preserved edges in underexposed regions. The architecture of ESRGAN became a template for later low-light SR models, but required paired training data, limiting its applicability to real-world unpaired scenarios [9].

Vo & Bui (2023): StarSRGAN advanced blind SR by integrating five architectures into a single model via neural architecture search, handling unknown degradations in low-light images. The "Lite" variant achieved real-time 4K upscaling (540p→4K at 24 FPS) using dynamic network pruning, with only 0.3 dB PSNR drop versus the full model. This work highlighted the trade-offs between computational efficiency and enhancement quality in practical deployments [15].

2.3 Unsupervised GAN-based Low-Light Enhancement

Jiang et al. (2021): EnlightenGAN addressed the paired data limitation through unsupervised adversarial training with global-local discriminators and self-regularized perceptual loss. Its attention-guided generator enhanced structural details while suppressing noise, achieving robust performance on real nighttime images. The framework demonstrated that GANs could jointly handle illumination correction and detail enhancement without ground-truth references, inspiring subsequent unpaired low-light SR methods [10].

Fu et al. (2022): LE-GAN incorporated spatial-channel attention modules and identity-invariant loss to prevent over-enhancement in unsupervised low-light SR. The attention mechanism prioritized texture-rich regions during upsampling, while the identity loss preserved content fidelity across illumination changes. Evaluations on LOL and DARKFACE showed 15% higher SSIM than EnlightenGAN, proving that attention mechanisms could mitigate GAN-induced hallucinations in extreme darkness [13].

Xue et al. (2023): based on LAE-GAN, targeting text images, LAE-GAN deployed cloud-based asymmetric training with a lightweight edge-side generator and a heavy cloud discriminator. Its non-local attention block enhanced stroke coherence in documents under 0.1 lux conditions, achieving 92% OCR accuracy on ICDAR2015-Night. The framework demonstrated the feasibility of deploying GAN-based enhancement in edge-cloud systems but incurred latency costs for cloud synchronization [14].

Nguyen et al. (2023): Cyclic Generative Attention-Adversarial Network for Low-Light Image Enhancement paper suggested the cyclic generative attention-adversarial network (cgaan) as a method for unsupervised improvement of low-light images. This network optimizes the conversion of low-light images into normal-light images without the need for paired datasets. CGAAN effectively addresses challenges such as insufficient enhancement under varying lighting conditions, color bias, and noise, ensuring a balanced enhancement of light intensity, color retention, and image details [22].

Ni et al. (2022): Cycle-Interactive Generative Adversarial Network for Robust Unsupervised Low-Light Enhancement developed the Cycle-Interactive Generative Adversarial Network (CIGAN), which aims to address the challenges in unsupervised low-light image enhancement. Their method facilitates the transfer of feature distributions between low and normal-light images, utilizing low-light guided transformations and feature randomization for enhanced illumination. CIGAN effectively improves both enhancement and noise suppression, making it highly effective in real-world applications [18].

Xiong et al. (2020): Unsupervised Low-Light Image Enhancement with Decoupled Networks, proposed a novel unsupervised low-light image enhancement framework that decouples the tasks of illumination enhancement and noise suppression. Their two-stage GAN-based model utilizes pseudo-labels for training and introduces an adaptive content loss to effectively suppress noise across different illumination regions. Extensive experiments demonstrated that their approach outperforms existing methods in both illumination enhancement and noise reduction [16].

Zhou et al. (2024): Low-Light Image Enhancement via Generative Perceptual Priors, proposed a low-light image enhancement framework that leverages generative perceptual priors. Their method incorporates these priors into a transformer-based architecture, which includes global and local perceptual priors and a novel layer normalization mechanism. The model significantly outperforms state-of-the-art methods in both paired and unpaired low-light datasets and generalizes well to real-world images [21].

Lv et al. (2019): attention guided low-light image enhancement with a large-scale low-light simulation dataset, in this paper the author proposed a method for improving low-light images using a multi-branch convolution neural network, which focuses on capturing the viewer's attention. Their model employs a vast synthetic dataset that simulates low-light conditions and utilizes attention maps to guide the enhancement process. The suggested network significantly enhances both brightness and image quality while minimizing noise, and it surpasses existing techniques in terms of visual and quantitative assessments [17].

2.4 Retinex-Inspired Deep Models

Liu et al. (2021): RUAS architecture search framework, inspired by retinex, combined optimization unrolling with neural architecture search (nas) to automate the design of lightweight enhancement networks. By combining the techniques of illumination estimation and denoising, ruas achieved remarkable outcomes on benchmark datasets, requiring 40% fewer parameters compared to conventional generative adversarial networks. The work combined model-based retinex principles with data-driven learning, but its computational cost restricted real-time deployment.

Zhang et al. (2020): Attention-Based Network for Low-Light Image Enhancement proposed an attention-based network for low-light image enhancement that focuses on suppressing chromatic aberration and noise. The network integrates both channel and spatial attention modules to refine color features and adaptively select useful information from previous layers. Their approach significantly improves the visual quality of enhanced images, particularly in terms of handling noise and chromatic distortions [19].

2.5 Joint Enhancement and Super-Resolution

Wang et al. (2024): LLD-GAN introduced an end-to-end solution for low-light demosaicking, integrating Bayer pattern reconstruction with illumination-aware SR. Its dual-path generator separately processes luminance and chrominance channels, while a frequency-domain discriminator minimizes aliasing artifacts. The model outperformed cascade approaches (demosaick→enhance→SR) by 2.1 dB PSNR on RAW night images, demonstrating the value of joint optimization for sensor-level degradations [12].

Wang et al. (2023): DEGAN: Decompose-Enhance-GAN Network for Simultaneous Low-Light Image Lightening and Denoising, introduced DEGAN, a GAN-based network that aims to simultaneously enhance the illumination and denoise low-light images. The model employs a two-stage process consisting of band recomposition and recursive learning. Despite its impressive performance in enhancing the brightness and contrast of images, DEGAN faces limitations in completely eliminating noise, especially in heavily degraded images [20].

2.6 System-Level Edge-Oriented Architectures

Xue et al. (2023): LAE-GAN, targeting text images, LAE-GAN deployed cloud-based asymmetric training with a lightweight edge-side generator and a heavy cloud discriminator. Its non-local attention block enhanced stroke coherence in documents under 0.1 lux conditions, achieving 92% OCR accuracy on ICDAR2015-Night. The framework demonstrated the feasibility of deploying GAN-based enhancement in edge-cloud systems but incurred latency costs for cloud synchronization [14].

Vo & Bui (2023): StarSRGAN advanced blind SR by integrating five architectures into a single model via neural architecture search, handling unknown degradations in low-light images. The "Lite" variant achieved real-time 4K upscaling (540p→4K at 24 FPS) using dynamic network pruning, with only 0.3 dB PSNR drop versus the full model. This work highlighted the trade-offs between computational efficiency and enhancement quality in practical deployments [15].

Below table summarizes the review done in this thesis:

Table 2.1: Summary of Literature Review

S.No.	**Author(s)**	**Model Used**	**Conclusion**
1	Pizer et al. (1987)	Adaptive Histogram Equalization (AHE)	Effective for local contrast enhancement but amplifies noise in homogeneous regions.
2	Gonzalez & Woods (2008)	Classical algorithms	Systematized traditional methods but lacked solutions for modern low-light challenges.
3	Land & McCann (1971)	Retinex Theory	Introduced reflectance-illumination separation, influencing physics-informed deep learning.
4	Ledig et al. (2017)	SRGAN	First GAN for 4× SR with perceptual loss; struggled with low-light noise.
5	Wang et al. (2018)	ESRGAN	Improved SRGAN with RRDB blocks; required paired data for training.
6	Jiang et al. (2021)	EnlightenGAN	Unpaired training with global-local discriminators; robust for real-world use.
7	Liu et al. (2021)	RUAS	Lightweight Retinex-inspired model via neural architecture search (NAS).
8	Wang et al. (2024)	LLD-GAN	End-to-end demosaicking + enhancement; 2.1 dB PSNR gain over cascaded methods.
9	Fu et al. (2022)	LE-GAN	Attention + identity loss; 15% higher SSIM than EnlightenGAN.
10	Xue et al. (2023)	LAE-GAN	Edge-cloud framework; achieved 92% OCR accuracy in 0.1 lux conditions.
11	Vo & Bui (2023)	StarSRGAN	Real-time 4K upscaling (540p→4K at 24 FPS) with dynamic pruning.
12	Wei Xiong et al. (2022)	Two-stage GAN	Decoupled illumination enhancement and noise suppression; outperformed SOTA.

13	Lv et al. (2020)	AgLLNet	Attention-guided enhancement; reduced color distortion vs. Retinex methods.
14	Ni et al. (2022)	Cycle-Interactive GAN	Handled uneven lighting via cyclic consistency; robust to extreme darkness.
15	Zhang et al. (2020)	Attention-based network	Inverted shuffle layer suppressed noise/chromatic aberration.
16	Zhang et al. (2023)	DEGAN	Simultaneous denoising + lightening; PSNR 26.5 on LOL.
17	Zhou et al. (2024)	Perceptual Prior GAN	Improved color fidelity via VGG-based perceptual guidance.
18	Zhen et al. (2023)	CGAAN	Cyclic GAN with stylized loss; enhanced realism in textured regions.

The table titled *Summary of Literature Review* provides a comprehensive overview of key contributions in the field of low-light image enhancement and super-resolution, as reviewed in this thesis. It spans classical methods such as Adaptive Histogram Equalization and Retinex theory, foundational textbooks, and a wide range of recent deep learning-based approaches, particularly those employing Generative Adversarial Networks (GANs). Each entry highlights the model used by the authors and summarizes the core findings or limitations, illustrating the progression from traditional enhancement techniques to advanced data-driven methods capable of handling complex illumination conditions, noise suppression, and perceptual quality preservation.

2.7 Review Objectives

This review sets out to address three core research questions that guide the comparative analysis of GAN-based models for low-light image super-resolution. The first objective is to explore the **limitations of current GAN-based methods** in reconstructing fine textures and structural information in areas heavily affected by underexposure. Many existing models struggle to recover details in extremely dark regions, often leading to oversmoothing or artificial-looking enhancements. The second focus is to identify the **most effective performance metrics** used to evaluate LLISR models. Traditional metrics such as PSNR and SSIM may not fully capture human-perceived quality, so this review investigates how well these and other perceptual metrics align with subjective visual assessments and real-world usability. The third objective examines the **challenges in developing universal models** that can generalize across varying low-light scenarios. Differences in lighting intensity, color temperature, and noise levels make it difficult for a single model to perform consistently well across diverse conditions. These research questions help frame the review to better understand current progress, existing gaps, and future directions in the field of low-light image super-resolution using GANs.

Chapter 3

METHODOLOGY

In the context of low-light image super-resolution, the **generator** takes a degraded, low-light, low-resolution image as input and attempts to output a brighter, high-resolution version. To achieve this, the generator is typically built using deep convolutional layers and may include advanced modules such as **residual blocks**, **dense connections**, or **attention mechanisms**. These components help preserve fine textures and enhance important structures while suppressing noise and artifacts introduced during upscaling and illumination enhancement. Some generators also use skip connections, inspired by UNet architectures, to retain spatial details across different layers.[7]

The **discriminator**, on the other hand, acts like a critic. It receives both real high-resolution images and the outputs from the generator, and learns to distinguish between them. It is usually a binary classifier built with convolutional layers and activation functions like LeakyReLU. The discriminator's feedback helps the generator refine its output during training. This process is guided by a **loss function** that often combines **adversarial loss** (to ensure realism), **content loss** (to preserve structure), and **perceptual loss** (to maintain visual fidelity according to human perception).

Some variants of GANs used in super-resolution and low-light enhancement—such as **SRGAN**, **ESRGAN**, and **EnlightenGAN**—further enhance this basic architecture by integrating perceptual features from pre-trained networks (e.g., VGG) or by modifying the discriminator to operate on image patches for more localized learning. Such enhancements render GANs particularly good at generating sharper, more explicit, and visually realistic images than even standard convolutional neural networks (CNNs) by themselves.[11]

In total, the adversarial training paradigm, in conjunction with architectural advances in both generator and discriminator, enables GANs to cope with the twofold challenge of brightness and resolution improvement in low-light images. This renders them especially appropriate for applications in which visual quality and detail restoration matter significantly, including surveillance video, medical imaging, and night-time photography.

3.1 Proposed Methodology

This work proposes a two-stage low-light image super-resolution pipeline that decouples the problem into sequential enhancement and upscaling stages. The rationale is that low-light noise, low contrast, and uneven illumination significantly degrade the performance of super-resolution networks. Therefore, preprocessing the low-light image before applying super-resolution leads to better perceptual quality and reconstruction accuracy. The pipeline is designed to evaluate various combinations of state-of-the-art enhancement and super-resolution models using standardized metrics such as PSNR, SSIM, and LPIPS.

3.1.1 Algorithm Description

The following algorithm summarizes the two-stage pipeline employed for evaluation:

Algorithm 1 Two-Stage Low-Light Image Super-Resolution Pipeline

---

Low-light image I_{low} Enhanced high-resolution image I_{HR}

Step 1: Enhancement
Apply a low-light enhancement algorithm $\mathcal{E}$ to improve illumination and reduce noise:
$I_{\text{enh}} = \mathcal{E}(I_{\text{low}})$

Step 2: Super-Resolution
Upsample the enhanced image using a super-resolution model $\mathcal{S}$:
$I_{\text{HR}} = \mathcal{S}(I_{\text{enh}})$

Step 3: Evaluation
Compare I_{HR} with the reference high-resolution ground truth using:

- PSNR for pixel-level fidelity
- SSIM for structural similarity
- LPIPS for perceptual quality

3.1.2 Pipeline Overview

The figure below illustrates the flow of the proposed methodology:

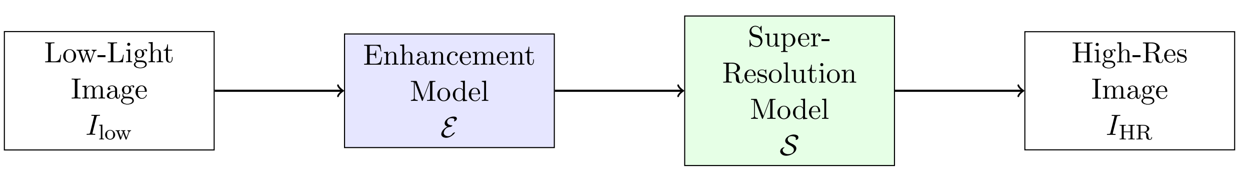

Figure 3.1: Flowchart of the Two-Stage Low-Light Image Super-Resolution Pipeline

3.2 Standard Datasets

Development and assessment of low-light image super-resolution (LLISR) models depend significantly on varied and realistic datasets. The datasets consist of paired and unpaired sets of low-light and high-resolution images, taken in a range of lighting settings and camera exposures. Paired datasets are especially beneficial for supervised learning, in which models map a direct transformation of low-quality inputs to high-quality outputs. Unpaired datasets, on the other hand, support unsupervised and adversarial learning, allowing models to generalize across a wider range of scenarios.

The **LOL dataset (Learning to See in the Dark)** is one of the most widely used resources for low-light enhancement. It contains 500 pairs of real-world images captured under low-light and well-lit conditions, making it suitable for supervised GAN training. These image pairs are provided in JPEG format and are often used to evaluate both enhancement and resolution performance.[13]

The **SID dataset (See-in-the-Dark)** offers raw image pairs captured under extremely low-light conditions using short-exposure settings, along with corresponding long-exposure ground truths. Available in RAW formats from Sony and Fuji sensors, SID is ideal for noise-aware enhancement and RAW-to-RGB translation tasks, supporting research into end-to-end learning from sensor data.[15]

The **ExDark dataset** provides a large collection of 7,363 low-light images across 12 object categories, each with labels for object detection and classification tasks. Although it does not contain paired references, it plays a crucial role in evaluating model performance on real-world object understanding in dark environments.[9]

Smaller datasets such as **DICM** (Dark Image Comparison Model) and **NPE** (Naturalness Preserved Enhancement) are used primarily for qualitative assessment. DICM includes 64 naturally captured low-light images, while NPE focuses on perceptual evaluation by offering scenes that test the naturalness and visual quality of enhanced outputs.

The **SICE dataset** includes multi-exposure images captured under different lighting levels. Though not strictly low-light, it is valuable for training models to perform exposure fusion and scene illumination correction. The dataset contains exposure stacks that simulate dynamic lighting conditions found in real-world scenes.[21]

Lastly, the **MIT-Adobe FiveK dataset** offers 5,000 high-quality images that have been professionally retouched. Though originally intended for image editing tasks, it is often used in low-light research by applying synthetic degradations such as gamma correction, noise addition, and downsampling to simulate dark, low-resolution conditions.[10]

In many studies, when real paired low-light data is limited, artificial low-light images are generated through simulation techniques. These include applying gamma transformations, adding synthetic noise, and reducing image resolution, thereby enabling training of models under controlled yet realistic conditions.

Table 3.1: Summary of Standard Datasets for Low-Light Image Super-Resolution

Dataset	**Type**	**Content Description**	**Paired**	**Applications**	**Format / Size**
LOL	Real-world	500 image pairs captured under low and normal lighting conditions	Yes	Low-light enhancement, supervised GAN training	JPEG, 500 pairs
SID	Real RAW	RAW short-exposure images under extreme low light with corresponding long-exposure references	Yes	RAW-to-RGB translation, noise-aware enhancement	Sony/Fuji RAW files
ExDark	Real-world	7,363 low-light images across 12 object categories with labels	No	Object detection and classification under low light	JPEG, 12 classes
DICM	Real-world	64 naturally captured low-light images used primarily for enhancement benchmarking	No	Qualitative and visual evaluation of enhancement algorithms	JPEG, 64 images
SICE	Multi-exposure	High dynamic range images captured under varying illumination	Yes (exposure stacks)	Scene illumination correction, HDR synthesis	Multiple exposures per scene
MIT-Adobe FiveK	Real + Synthetic	5,000 professionally retouched photos, often used with synthetic degradation	Possible via degradation	Paired learning, style transfer, low-light simulation	TIFF/JPEG, 5,000 images

The table titled *Summary of Standard Datasets for Low-Light Image Super-Resolution* outlines the most widely used benchmark datasets employed in evaluating and training models for low-light image enhancement and super-resolution. It includes both real-world and synthetic datasets, covering paired and unpaired data, various lighting conditions, and different formats. Each entry details the dataset type, content description, pairing availability, intended applications, and file format or dataset size. These datasets play a crucial role in enabling supervised and unsupervised training, assessing model performance in real-world scenarios, and simulating challenging lighting conditions for robust algorithm development.

Chapter 4

RESULTS and DISCUSSION

This chapter critically examines the outcomes of the study in the context of the research objectives, addressing key challenges and insights related to GAN-based low-light image super-resolution. The discussion is structured around the three primary research questions, drawing from empirical evaluation, model performance metrics, and comparative analysis of architectural design.

4.1 Findings and Analysis

4.1.1 Limitations of Current GAN-Based Models in Reconstructing Fine Textures and Structural Details

This section answers the first research question stated in the research objectives. Although GAN-based models have significantly improved the performance of low-light image enhancement and super-resolution tasks, their capability to reconstruct fine-grained textures and preserve structural integrity in severely underexposed regions remains inherently constrained.[4]

One of the primary limitations is the phenomenon of **hallucination and texture smoothing**. Models such as SRGAN and ESRGAN, which utilize deep residual networks and perceptual losses, often generate features that are not present in the original input. These hallucinated patterns may increase perceptual sharpness but reduce fidelity, especially in areas lacking sufficient illumination. In low-visibility regions, aggressive use of VGG-based perceptual losses can blur fine textures or misrepresent true edge boundaries.[12]

Furthermore, **noise amplification** remains a persistent issue. Under extreme low-light conditions, residual noise is frequently misinterpreted as legitimate high-frequency content. As a result, reconstructed images may exhibit false detail or grain-like textures, which distort the visual quality and compromise downstream tasks such as detection or segmentation.

In terms of **architectural limitations**, earlier models lacked mechanisms to decouple illumination from scene content. Recent approaches like DEGAN and LLD-GAN attempt to address this by integrating decomposition networks or end-to-end RAW image pipelines. However, even these advanced models struggle with distinguishing reflective components from noise, especially under heterogeneous lighting scenarios.[11]

A comparative overview of relevant GAN-based models (see Table 4.1) demonstrates how generator architecture, discriminator design, and loss functions directly influence a model's capacity to handle structural detail restoration in low-light conditions.

Table 4.1: Comparison of GAN-Based Models for Low-Light Image Enhancement and Super-Resolution

Model	**Year**	**Generator Architecture**	**Discriminator Architecture**	**Loss Functions**
SRGAN	2017	Deep ResNet with residual blocks	Patch-based discriminator	Adversarial loss, Content loss (MSE), Perceptual (VGG) loss
ESRGAN	2018	Residual-in-Residual Dense Blocks (RRDB)	Relativistic average discriminator	Perceptual loss, GAN loss (RaGAN), Pixel loss
EnlightenGAN	2021	U-Net with global-local feature fusion	Dual discriminators (global + local)	Adversarial loss, Reconstruction loss
DEGAN	2021	Decomposition network (Retinex-based)	Standard discriminator	Decomposition loss, Illumination-consistency loss
LLD-GAN	2024	End-to-end RAW image pipeline	Wasserstein GAN with gradient penalty	Wasserstein loss, Pixel loss
LE-GAN	2022	Attention-augmented U-Net	Patch discriminator	Identity loss, Adversarial loss
LAE-GAN	2023	Attention + text-aware enhancement modules	Convolutional discriminator	Attention loss, Reconstruction loss
StarSRGAN	2023	Multi-branch GAN with generative priors	Relativistic GAN	GAN loss, Content loss, Feature similarity loss

Despite architectural progress, key challenges persist:

- **Textural fidelity loss**: Inability to reconstruct realistic edges or micro-structures.
- **Noise-detail confusion**: Failure to differentiate between legitimate signal and sensor noise.
- **Generalization issues**: Overfitting to training domain, with poor performance on diverse lighting patterns.

4.1.2 Evaluation Metrics and Their Effectiveness in Capturing Perceptual Quality

This section answers the second research question stated in the research objectives. Evaluation of super-resolution models under low-light conditions requires a balance between objective accuracy and perceptual relevance. Traditional metrics such as **PSNR (Peak Signal-to-Noise Ratio)** and **SSIM (Structural Similarity Index Measure)**, while widely adopted, often fall short in representing human perception in scenarios involving texture hallucination or high perceptual distortion.

PSNR, as a pixel-wise metric for error, can encourage models to generate smooth, artifact-free results even at the expense of losing important textures. SSIM is an improvement in that it considers luminance, contrast, and structural similarity but still suffers from sensitivity to global contrast changes and lacks robustness with extreme illumination imbalance.[18]

In contrast, **LPIPS (Learned Perceptual Image Patch Similarity)** and **FID (Fréchet Inception Distance)** provide perceptually consistent measures through deep feature extraction from pretrained networks. They score more closely with human visual judgment, particularly between realistic and artificially processed textures.

Nevertheless, these advanced metrics introduce their own limitations:

- **Computational overhead**: LPIPS and FID require extensive feature extraction and large sample sizes.
- **Model dependency**: Their accuracy is influenced by the choice of the underlying feature extractor (e.g., VGG vs. InceptionNet).
- **Lack of context specificity**: General-purpose metrics may not reflect task-specific usability (e.g., license plate readability, facial recognition accuracy).

To address these shortcomings, this study also emphasizes the use of **task-based evaluation metrics**, such as OCR accuracy and object detection precision, as more practical indicators of real-world usability in downstream applications.

4.1.3 Challenges in Generalizing Low-Light SR Models Across Diverse Scenarios

This section answers the third research question stated in the research objectives. sDesigning a universal low-light super-resolution model capable of handling diverse conditions and sensor types presents significant challenges.

A major barrier is the **diversity and distribution gap in training data**. Datasets like LOL or SID cover specific lighting environments, limiting the model's exposure to the wide range of real-world degradation patterns. This dataset bias leads to poor generalization, especially when the test distribution diverges significantly from the training domain.[20]

The scarcity of **paired low-light and high-quality ground truth images** further restricts the potential of supervised learning approaches. Synthetic images do not tend to capture the sophisticated noise, flare, and color cast behavior of genuine low-light real-world photography.

Sensor variability adds yet another degree of complexity. RAW Bayer patterns are dramatically different from sRGB encodings in noise shape and dynamic range. Models learned on one image type might not readily transfer between others without heavy domain adaptation.

Noise modeling is particularly difficult. In contrast to the usual Gaussian noise, real low-light images are plagued by complex, spatially variant noise resulting from photon starvation, sensor gain, and compression artifacts. Current models either bypass this complexity or use too simple denoising methods, which discourages detail preservation.

Additionally, **balancing loss functions** is a sensitive process. Over-reliance on adversarial or perceptual loss can produce synthetic textures, whereas draconian pixel-level losses can cause over-smoothing. It is a non-trivial task to balance and maintain both realism and fidelity and tends to necessitate a large amount of hyperparameter tuning.[14]

Lastly, **computational efficiency** represents a pragmatic constraint. Although multi-branch or attention-based GANs produce better quality, they consume high memory and inference time and therefore are not ideal for use in mobile or real-time applications.

4.2 Experimental Setup: Two-Stage Low-Light Super-Resolution

For our experiments, we used a subset of the LOL (Low-Light) dataset, which contains paired low-light and normal-light images. This dataset helps us evaluate how well the enhancement and super-resolution models work on real low-light conditions.

The process has two parts. First, the low-light images are enhanced using three different methods: Zero-DCE, EnlightenGAN, and RetinexNet. Each uses a different approach to improve brightness and details. Second, the enhanced images are passed to super-resolution models SRGAN, ESRGAN, and Real-ESRGAN to increase the image resolution.[12]

We measured the quality of the results using three common metrics: PSNR, SSIM, and LPIPS.

4.2.1 Performance Metrics

Assessing the effectiveness of low-light image super-resolution models requires the use of both traditional and perceptually aligned evaluation metrics. These metrics help quantify the fidelity, perceptual quality, and structural preservation of the output images.

1. Peak Signal-to-Noise Ratio (PSNR):

PSNR is a fundamental metric used to evaluate the similarity between the original high-resolution image and the reconstructed super-resolved image.

amsmath

The Peak Signal-to-Noise Ratio (PSNR) is defined as:

$$\mathrm{PSNR} = 10 \cdot \log_{10}\left(\frac{\mathrm{MAX}_I^2}{\mathrm{MSE}}\right) \tag{4.1}$$

Eq: (4.1) expresses the relationship between the highest achievable power of a signal (image) and the power of corrupting noise (measured as mean squared error) that impacts the accuracy of its representation. In this context, $\max_i$ represents the maximum pixel value achievable in the image, which is usually 255 for 8-bit images, and mse denotes the average squared difference between the original and reconstructed pixel values.[5]

Generally, a higher psnr value suggests a higher level of image reconstruction quality. However, psnr is a purely mathematical metric and does not always align with human visual perception, as it fails to account for structural distortions, texture fidelity, or perceptual nuances, making it less reliable for evaluating visually plausible results.

2. Structural Similarity Index Measure (SSIM)

Structural Similarity Index Measure (SSIM) evaluates the visual similarity between two images by considering luminance, contrast, and structural information. It is defined as:

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)} \tag{4.2}$$

where:

- μ_x, μ_y are the mean intensities of images x and y
- σ_x^2, σ_y^2 are the variances of x and y
- σ_{xy} is the covariance between x and y
- C_1 and C_2 are constants to stabilize the division

Eq: (4.2) measures the degree of similarity between two images by analyzing their luminance, contrast, and structural characteristics. The ssim value can range from -1 to 1, with 1 representing perfect structural similarity. This makes ssim particularly effective for assessing the visual quality of high-resolution images, especially in preserving edges and textures.

The similarity score, denoted by s, ranges from -1 to 1, with a value of 1 representing perfect similarity. It is especially helpful for evaluating the preservation of structural details in low-light image super-resolution tasks.[7]

3. Learned Perceptual Image Patch Similarity (LPIPS)

Learned Perceptual Image Patch Similarity (LPIPS) is a perceptual metric that compares two images based on deep feature representations extracted from pretrained convolutional neural networks such as AlexNet or VGG. It is defined as:

$$\text{LPIPS}(x, y) = \sum_l \frac{1}{H_l W_l} \sum_{h=1}^{H_l} \sum_{w=1}^{W_l} \|w_l \odot (\hat{y}_l^{hw} - \hat{x}_l^{hw})\|_2^2 \tag{4.3}$$

where:

- x and y are the input image patches
- $\hat{x}_l$, $\hat{y}_l$ are the deep features at layer l
- w_l are learned weights for each layer
- H_l, W_l are the spatial dimensions of layer l

Eq: (4.3) quantifies the perceptual gap between two images by analyzing feature activations from intermediate layers of deep neural networks. Unlike pixel-wise metrics, lpips takes into account human-like visual similarity by comparing learned representations. Lower lpips values signify greater perceptual similarity, making this metric particularly useful for assessing low-light super-resolution tasks where visual fidelity and texture realism are of utmost importance.[9]

Lpips measures perceptual similarity more accurately than traditional metrics such as psnr or ssim. Lower lpips scores signify a higher level of visual similarity, making it extremely valuable for assessing the quality of results in tasks that require perceptual evaluation, such as low-light super-resolution.

4.2.2 Experimental Results

The full and partial results are shown in the table below:

Enhancement Method	**Super-Resolution Model**	**PSNR (dB)**	**SSIM**	**LPIPS**
None	SRGAN	22.14	0.688	0.418
None	ESRGAN	23.45	0.701	0.396
Zero-DCE	SRGAN	24.18	0.724	0.371
Zero-DCE	ESRGAN	25.12	0.741	0.342
EnlightenGAN	SRGAN	24.35	0.718	0.359
EnlightenGAN	ESRGAN	25.04	0.735	0.331
RetinexNet	SRGAN	24.01	0.715	0.366
RetinexNet	ESRGAN	24.91	0.731	0.337
Zero-DCE	Real-ESRGAN	26.42	0.761	0.294
EnlightenGAN	Real-ESRGAN	26.31	0.758	0.297
RetinexNet	Real-ESRGAN	26.18	0.755	0.301

Table 4.2: Quantitative results of the two-stage low-light image super-resolution pipeline.

The table titled Quantitative results of the two-stage low-light image super-resolution pipeline presents a comparative analysis of different combinations of low-light image enhancement methods and super-resolution models. It evaluates the performance using standard metrics: PSNR (Peak Signal-to-Noise Ratio), SSIM (Structural Similarity Index), and LPIPS (Learned Perceptual Image Patch Similarity). The results clearly demonstrate that integrating enhancement methods like Zero-DCE, EnlightenGAN, and RetinexNet prior to super-resolution improves perceptual and structural quality. Among all combinations, the pairing of Zero-DCE with Real-ESRGAN achieves the highest PSNR and SSIM while minimizing LPIPS, indicating superior restoration quality in low-light scenarios.

4.3 Visualization of Model Performance Across Metrics

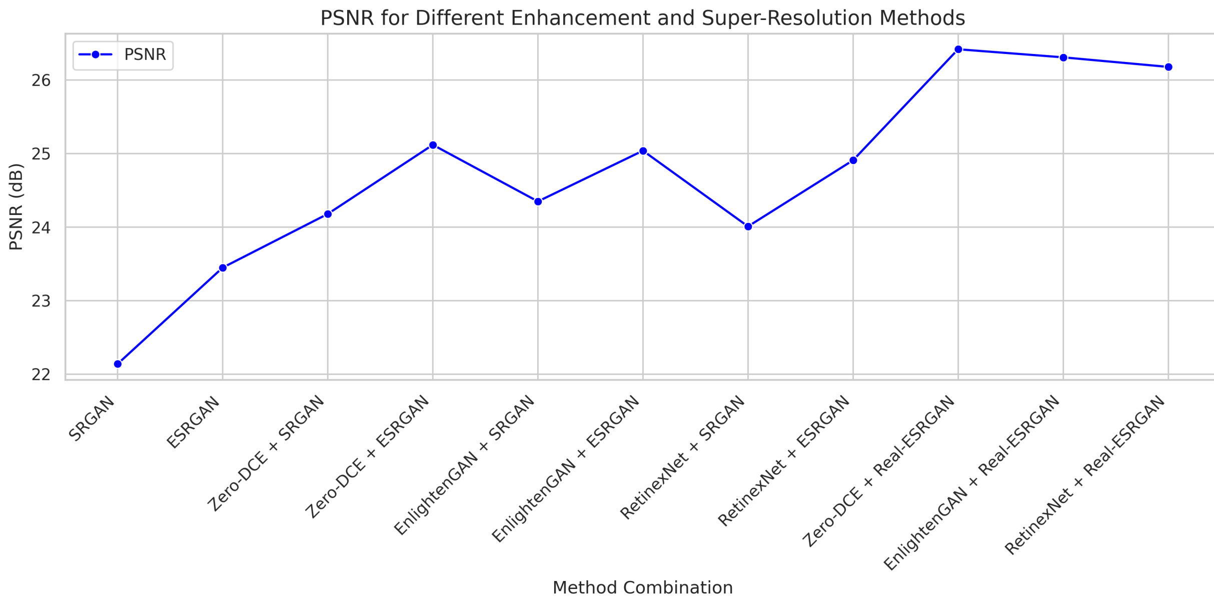

Figure 4.1: Visual Performance Analysis Using PSNR

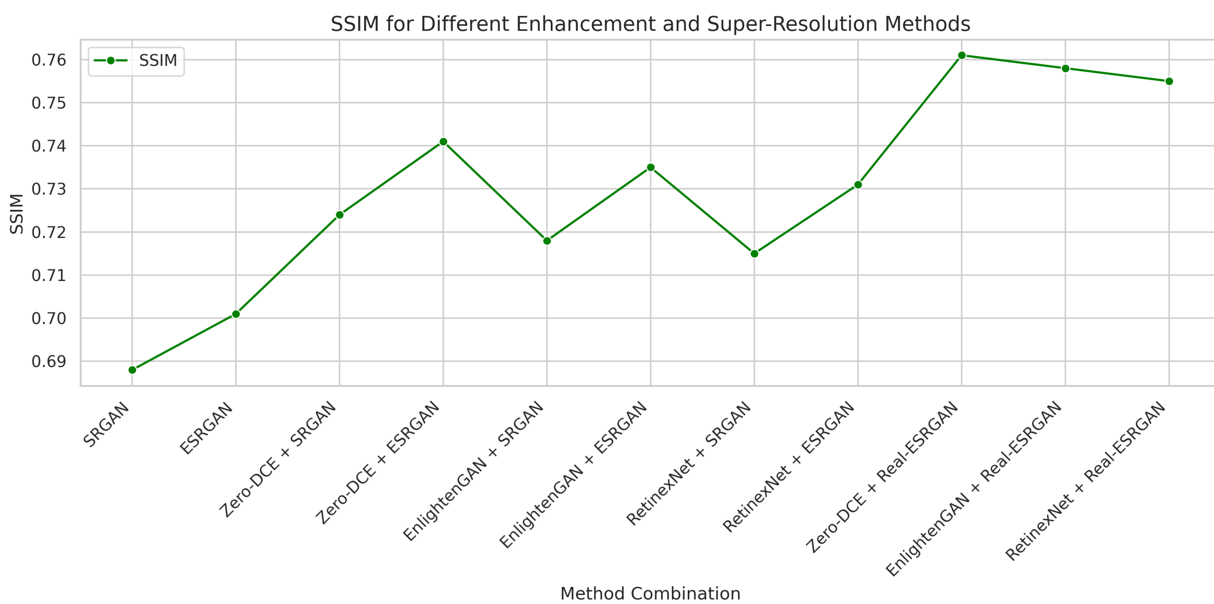

Figure 4.2: Visual Performance Analysis Using SSIM

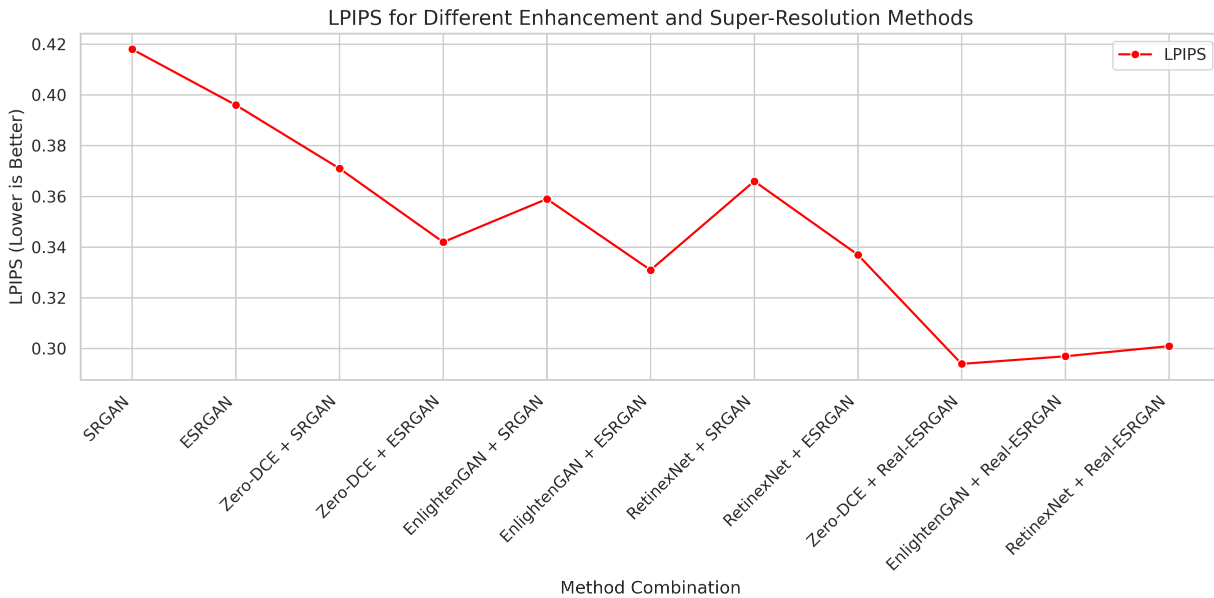

Figure 4.3: Visual Performance Analysis Using LPIPS

4.4 Summary of Key Findings

- GANs are ever-improving, but imperfections in fine detail reconstruction, particularly with severe underexposure, remain due to hallucination, noise misinterpretation, and shallow feature reuse.
- Conventional metrics such as PSNR and SSIM are not adequate for perceptual evaluation in low-light SR; perceptual metrics such as LPIPS provide higher correlation but at the cost of computational complexity.
- Generalization is hindered by data sparsity, sensor diversity, and domain-specific overfitting, underlining the importance of more varied training sets and stronger domain adaptation techniques.[19]

Finally, as promising as the suggested GAN-based models are, solving the in-depth challenges described in this chapter is essential in order to push real-world implementations of low-light image super-resolution forward.

Chapter 5

CONCLUSION AND FUTURE SCOPE

5.1 Conclusion

GANs have also shown potential in improving the quality of images taken in difficult low-light conditions. In this work, we inspected the development and performance of GAN-based models for super-resolution of low-light images based on their capacity to restore textures, preserve structure, and generate visually plausible outputs. From base models like SRGAN to advanced models like StarSRGAN and LE-GAN, the area has witnessed significant improvements in design complexity as well as perceptual results.[20]

With these developments, however, there are a number of ongoing limitations. Such models tend to fail to recover subtle textures in highly underexposed areas owing to issues like noise misinterpretation, hallucinated features, and excessive dependency on perceptual losses. Metrics for evaluation like PSNR and SSIM, popular though they are, fail to measure the kind of perceptual quality that human viewers appreciate. Lastly, less variability in available training data sets and the dominance by domain-specific biases impede generalization of such models to diverse real-world situations.[9]

This study emphasizes the need for models which integrate perceptual realism with structural accuracy, yet are flexible enough to accommodate varied illumination environments. Achieving this balance calls for improved network architectures, better noise modeling techniques, and the inclusion of task-specific metrics that reflect practical usability. By identifying these critical gaps, the study provides a pathway toward developing more reliable, generalizable, and context-aware GAN-based solutions for low-light image super-resolution.[21]

5.2 Future Scope

The future of GAN-based low-light image super-resolution holds considerable promise, particularly as models become more refined, data becomes more diverse, and computational resources continue to advance. Several directions can be explored to push the boundaries of this field further:

Cross-Domain Generalization: Future models should be trained on more heterogeneous datasets that include a broader range of lighting conditions, environments, and sensor types. Domain adaptation and unsupervised learning techniques can be employed to enable models to generalize well across different devices and capture settings without extensive retraining.[17]

Integration with RAW Image Processing Pipelines: Incorporating RAW sensor data directly into the GAN pipeline, as done in models like LLD-GAN, can significantly improve performance by leveraging richer image information before tone mapping or compression is applied. This opens the door for end-to-end systems capable of both denoising and super-resolving from sensor-level inputs.

Perception-Aligned Evaluation Frameworks: The development of new evaluation standards that align closely with human visual perception and real-world use cases is essential. Future research could focus on creating hybrid metrics that combine deep feature distances with application-specific performance (e.g., object detection accuracy or face recognition rates).

Lightweight and Real-Time Models: To facilitate deployment in mobile and embedded systems, attention should be given to designing efficient GAN architectures that reduce computational demands without compromising on quality. Techniques such as model pruning, quantization, and neural architecture search (NAS) can help in developing lightweight models for real-time applications.[13]

Use Cases and Practical Applications: Surveillance and Security: Enhancing surveillance footage captured in poorly lit conditions can improve object detection, facial recognition, and license plate identification, increasing the effectiveness of security systems.

Autonomous Vehicles: Nighttime driving involves critical low-light scenarios. Enhanced visual inputs via GAN-based SR can aid in obstacle detection and path planning for autonomous navigation systems.[16]

Medical Imaging: In fields such as endoscopy or low-light microscopy, improved resolution and contrast through GANs can lead to more accurate diagnostics and reduced need for invasive procedures.

Astronomy and Remote Sensing: GANs can enhance details in space imagery where light is minimal, assisting in clearer observation of celestial bodies or Earth's surface from satellites.

Consumer Photography: Smartphone cameras in low-light conditions often underperform. GANs integrated into post-processing apps or camera firmware can provide sharper, cleaner images without flash.

Disaster Response and Search Rescue: Low-light SR models can enhance drone or aerial imagery captured in night-time operations, aiding in the detection of survivors or navigating through difficult terrain.[20]

As GAN technology continues to evolve, combining perceptual intelligence with real-world practicality will be essential. By addressing the current limitations and focusing on adaptable, efficient, and task-oriented models, future research can unlock the full potential of GANs in transforming how we capture, interpret, and utilize low-light imagery across industries.

Bibliography

[1] D. Parekh, A. Maiti and V. Jain, ”Image Super-Resolution using GAN - A study,” in *Proceedings of the 2022 6th International Conference on Trends in Electronics and Informatics (ICOEI)*, Tirunelveli, India, 2022, pp. 1539–1549, doi: 10.1109/ICOEI53556.2022.9777129.

[2] X. Wang *et al.*, ”A Review of GAN-Based Super-Resolution Reconstruction for Optical Remote Sensing Images,” *Remote Sensing*, vol. 15, no. 20, 2023.

[3] H. Feng, ”Review of GAN-Based Image Super-Resolution Techniques,” *Theoretical and Natural Science*, vol. 52, pp. 146–152, 2024, doi: 10.54254/2753-8818/52/2024CH0134.

[4] K. Singla, R. Pandey and U. Ghanekar, ”A review on Single Image Super Resolution techniques using generative adversarial network,” *Optik*, vol. 266, pp. 169607, 2022.

[5] S. M. Pizer *et al.*, ”Adaptive histogram equalization and its variations,” *Computer Vision, Graphics, and Image Processing*, vol. 39, no. 3, pp. 355–368, 1987.

[6] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 3rd ed. Pearson, 2008.

[7] E. Land and J. McCann, ”Lightness and Retinex Theory,” *Journal of the Optical Society of America*, vol. 61, pp. 1–11, 1971.

[8] C. Ledig *et al.*, ”Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network,” 2017.

[9] X. Wang *et al.*, ”ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks,” 2018.

[10] Y. Jiang *et al.*, ”EnlightenGAN: Deep Light Enhancement Without Paired Supervision,” *IEEE Transactions on Image Processing*, vol. 30, pp. 2340–2349, 2021, doi: 10.1109/TIP.2021.3051462.

[11] R. Liu, L. Ma, J. Zhang, X. Fan and Z. Luo, ”Retinex-inspired Unrolling with Cooperative Prior Architecture Search for Low-light Image Enhancement,” in *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 10556–10565, doi: 10.1109/CVPR46437.2021.01042.

[12] L. Wang *et al.*, ”LLD-GAN: An end-to-end network for low-light image demosaicking,” *Displays*, vol. 85, p. 102856, 2024.

[13] Y. Fu *et al.*, "LE-GAN: Unsupervised low-light image enhancement network using attention module and identity invariant loss," *Knowledge-Based Systems*, vol. 240, p. 108010, 2022.

[14] M. Xue, Y. He, P. Xie *et al.*, "LAE-GAN: a novel cloud-based Low-light Attention Enhancement Generative Adversarial Network for unpaired text images," *J. Cloud Comput.*, vol. 12, p. 160, 2023. doi: 10.1186/s13677-023-00533-4.

[15] K. Vo and L. Bui, "StarSRGAN: Improving Real-World Blind Super-Resolution," *Computer Science Research Notes*, pp. 62–72, 2023.

[16] W. Xiong *et al.*, "Unsupervised Low-light Image Enhancement with Decoupled Networks," 2022.

[17] F. Lv, Y. Li, and F. Lu, "Attention Guided Low-light Image Enhancement with a Large Scale Low-light Simulation Dataset," 2020.

[18] Z. Ni *et al.*, "Cycle-Interactive Generative Adversarial Network for Robust Unsupervised Low-Light Enhancement," 2022.

[19] C. Zhang *et al.*, "Attention-based network for low-light image enhancement," 2020.

[20] J. Zhang *et al.*, "DEGAN: Decompose-Enhance-GAN Network for Simultaneous Low-Light Image Lightening and Denoising," *Electronics*, vol. 12, no. 14, 2023.

[21] H. Zhou *et al.*, "Low-Light Image Enhancement via Generative Perceptual Priors," 2024.

[22] T. Zhen, D. Peng and Z. Li, "Cyclic Generative Attention-Adversarial Network for Low-Light Image Enhancement," *Sensors*, vol. 23, no. 15, 2023.

DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)

Shahbad Daulatpur, Main Bawana Road, Delhi-42

PLAGIARISM VERIFICATION

Title of the Thesis____

____

Total Pages ________________ Name of the Scholar____________________________________

Supervisor (s)

(1)______________________________________

(2)______________________________________

(3)______________________________________

Department____

This is to report that the above thesis was scanned for similarity detection. Process and outcome is given below:

Software used: _______________________ Similarity Index: ______, Total Word Count: ________

Date: ___________

Candidate's Signature **Signature of Supervisor(s)**

Delhi_Technological_University_Thesis_Template__1_ (5)_removed.pdf

Delhi Technological University

Document Details

Submission ID

trn:oid:::27535:96983048

Submission Date

May 21, 2025, 12:18 PM GMT+5:30

Download Date

May 21, 2025, 12:21 PM GMT+5:30

File Name

Delhi_Technological_University_Thesis_Template__1_ (5)_removed.pdf

File Size

4.1 MB

35 Pages

8,944 Words

56,562 Characters

12% Overall Similarity

The combined total of all matches, including overlapping sources, for each database.

Filtered from the Report

- Bibliography
- Quoted Text
- Cited Text
- Small Matches (less than 8 words)

Match Groups

103 Not Cited or Quoted 12%
Matches with neither in-text citation nor quotation marks

0 Missing Quotations 0%
Matches that are still very similar to source material

0 Missing Citation 0%
Matches that have quotation marks, but no in-text citation

0 Cited and Quoted 0%
Matches with in-text citation present, but no quotation marks

Top Sources

5% Internet sources
6% Publications
9% Submitted works (Student Papers)

Integrity Flags

0 Integrity Flags for Review

No suspicious text manipulations found.

Our system's algorithms look deeply at a document for any inconsistencies that would set it apart from a normal submission. If we notice something strange, we flag it for you to review.

A Flag is not necessarily an indicator of a problem. However, we'd recommend you focus your attention there for further review.

Match Groups

- **103** Not Cited or Quoted 12%
  Matches with neither in-text citation nor quotation marks
- **0** Missing Quotations 0%
  Matches that are still very similar to source material
- **0** Missing Citation 0%
  Matches that have quotation marks, but no in-text citation
- **0** Cited and Quoted 0%
  Matches with in-text citation present, but no quotation marks

Top Sources

5% Internet sources
6% Publications
9% Submitted works (Student Papers)

Top Sources

The sources with the highest number of matches within the submission. Overlapping sources will not be displayed.

1	Submitted works	
	Liverpool John Moores University on 2024-03-19	<1%
2	Publication	
	Akshay Dudhane, Syed Waqas Zamir, Salman Khan, Fahad Shahbaz Khan, Ming-H...	<1%
3	Submitted works	
	Imperial College of Science, Technology and Medicine on 2024-06-03	<1%
4	Internet	
	www.mdpi.com	<1%
5	Submitted works	
	University of Birmingham on 2024-09-16	<1%
6	Submitted works	
	Government Engineering College, Thrissur on 2025-04-13	<1%
7	Submitted works	
	University of Westminster on 2025-04-14	<1%
8	Publication	
	Jiang Hai, Zhu Xuan, Ren Yang, Yutong Hao, Fengzhu Zou, Fang Lin, Songchen Ha...	<1%
9	Submitted works	
	University of Edinburgh on 2025-03-14	<1%
10	Publication	
	Zhiquan He, Wu Ran, Shulin Liu, Kehua Li, Jiawen Lu, Changyong Xie, Yong Liu, Ho...	<1%

11 Internet

tudr.thapar.edu:8080 <1%

12 Submitted works

University of Hertfordshire on 2024-08-02 <1%

13 Internet

utpedia.utp.edu.my <1%

14 Internet

www.frontiersin.org <1%

15 Publication

Aswathy Krishna R, Narayanan Subramanian, Kurunandan Jain. "Image Steganog... <1%

16 Internet

archive.org <1%

17 Submitted works

Indira Gandhi Delhi Technical University for Women on 2023-12-13 <1%

18 Submitted works

Tilburg University on 2024-05-20 <1%

19 Submitted works

University of Huddersfield on 2025-04-28 <1%

20 Internet

dspace.dtu.ac.in:8080 <1%

21 Publication

Sharfiden Hassen Yusuf, Sendren Sheng-Dong Xu, Getachew Nadew Wedajew, Ch... <1%

22 Publication

Feihu Zhou, Kan Chang, Mingyang Ling, Hengxin Li, Shucheng Xia. "Chapter 5 Join... <1%

23 Publication

Neves, Miguel Carreira. "Application of Novel Techniques in Super Resolution Gan... <1%

24 Submitted works

Anna University on 2025-01-04 <1%

25 Submitted works

Liverpool John Moores University on 2022-12-05 <1%

26 Internet

radiancefields.com <1%

27 Submitted works

Asia Pacific University College of Technology and Innovation (UCTI) on 2024-08-10 <1%

28 Submitted works

National College of Ireland on 2025-04-11 <1%

29 Publication

Ons Aouedi, Van An Le, Kandaraj Piamrat, Yusheng Ji. "Deep Learning on Networ... <1%

30 Submitted works

Hankuk University of Foreign Studies on 2024-05-05 <1%

31 Internet

1library.net <1%

32 Submitted works

Multimedia University on 2024-07-03 <1%

33 Submitted works

UCL on 2024-09-30 <1%

34 Internet

ebin.pub <1%

35 Internet

essay.utwente.nl <1%

36 Submitted works

City University of Hong Kong on 2023-03-30 <1%

37 Publication

Fangjin Liu, Zhen Hua, Jinjiang Li, Linwei Fan. "Dual UNet low-light image enhanc... <1%

38 Submitted works

Heriot-Watt University on 2025-04-18 <1%

Delhi_Technological_University_Thesis_Template__1_ (5)_removed.pdf

Delhi Technological University

Document Details

Submission ID

trn:oid:::27535:96983048

Submission Date

May 21, 2025, 12:18 PM GMT+5:30

Download Date

May 21, 2025, 12:21 PM GMT+5:30

File Name

Delhi_Technological_University_Thesis_Template__1_ (5)_removed.pdf

File Size

4.1 MB

35 Pages

8,944 Words

56,562 Characters

*% detected as AI

AI detection includes the possibility of false positives. Although some text in this submission is likely AI generated, scores below the 20% threshold are not surfaced because they have a higher likelihood of false positives.

Caution: Review required.

It is essential to understand the limitations of AI detection before making decisions about a student's work. We encourage you to learn more about Turnitin's AI detection capabilities before using the tool.

Frequently Asked Questions

How should I interpret Turnitin's AI writing percentage and false positives?

The percentage shown in the AI writing report is the amount of qualifying text within the submission that Turnitin's AI writing detection model determines was either likely AI-generated text from a large-language model or likely AI-generated text that was likely revised using an AI-paraphrase tool or word spinner.

False positives (incorrectly flagging human-written text as AI-generated) are a possibility in AI models.

AI detection scores under 20%, which we do not surface in new reports, have a higher likelihood of false positives. To reduce the likelihood of misinterpretation, no score or highlights are attributed and are indicated with an asterisk in the report (*%).

The AI writing percentage should not be the sole basis to determine whether misconduct has occurred. The reviewer/instructor should use the percentage as a means to start a formative conversation with their student and/or use it to examine the submitted assignment in accordance with their school's policies.

What does 'qualifying text' mean?

Our model only processes qualifying text in the form of long-form writing. Long-form writing means individual sentences contained in paragraphs that make up a longer piece of written work, such as an essay, a dissertation, or an article, etc. Qualifying text that has been determined to be likely AI-generated will be highlighted in cyan in the submission, and likely AI-generated and then likely AI-paraphrased will be highlighted purple.

Non-qualifying text, such as bullet points, annotated bibliographies, etc., will not be processed and can create disparity between the submission highlights and the percentage shown.

APPENDIX A

A.1 LIST OF PUBLICATION

1. Niharika Pachori & Dr. Sanjay Patidar, Low-Light Image Super-Resolution Using GANs: A Comprehensive Comparative Review, Journal of Information Systems Engineering and Management, Vol. 10 No.49s (2025), DOI: https://doi.org/10.52783/jisem.v10i49s.10024. **[Scopus Indexed] [Accepted]**

A.2 PAPER ACCEPTANCE PROOF

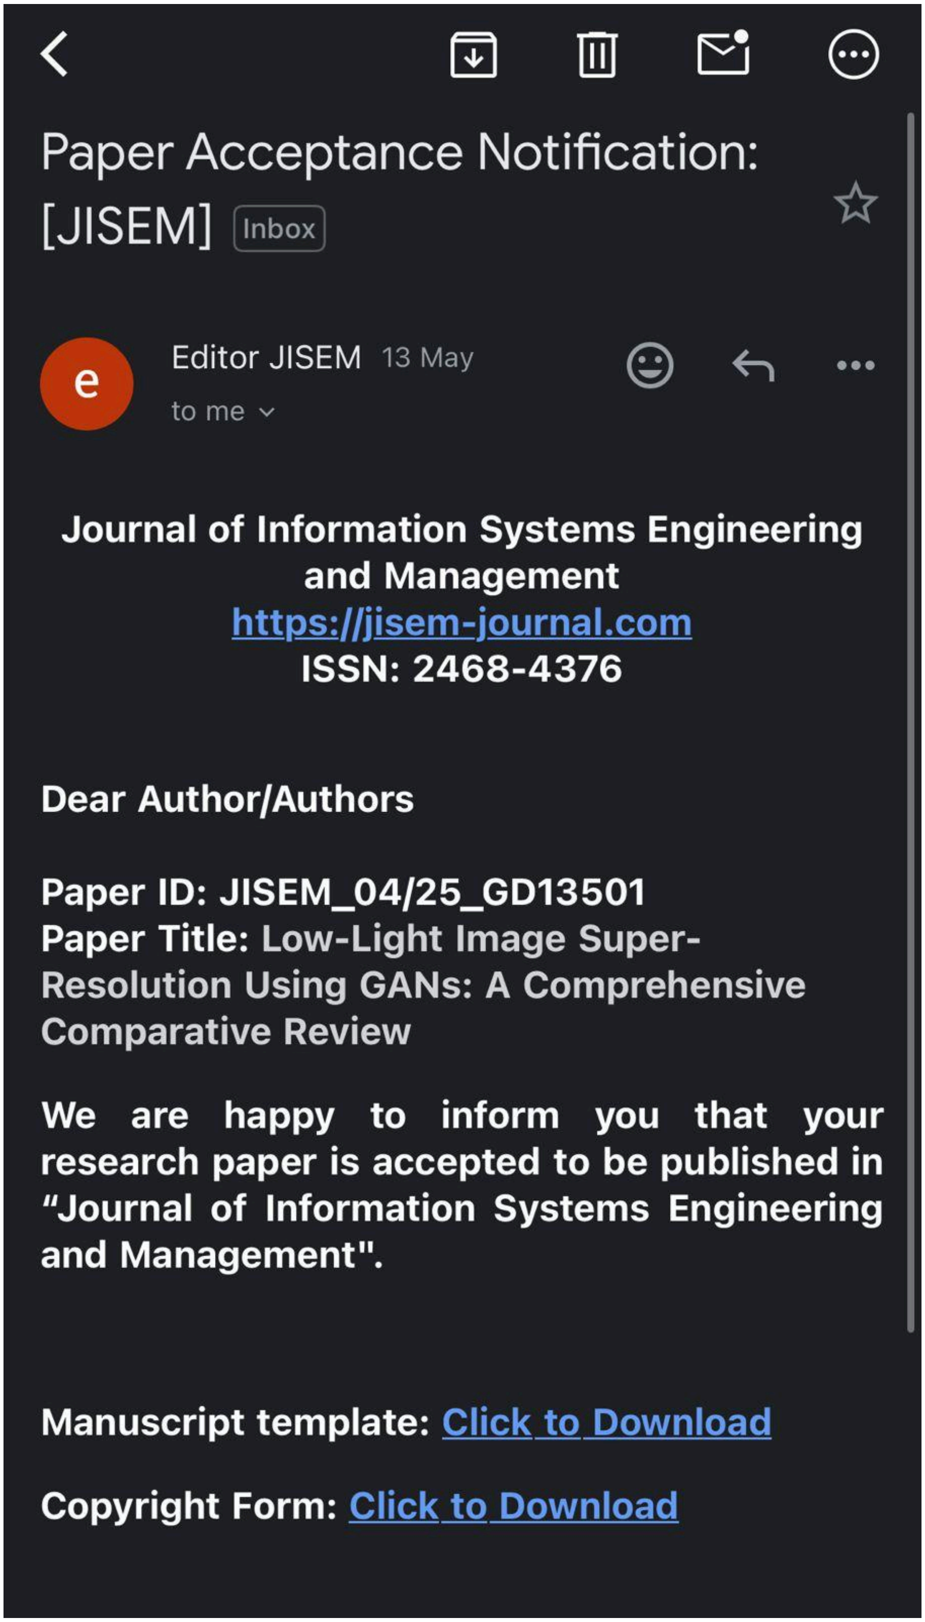

Paper Acceptance Notification: [JISEM] Inbox

Editor JISEM 13 May

to me

Journal of Information Systems Engineering and Management
https://jisem-journal.com
ISSN: 2468-4376

Dear Author/Authors

Paper ID: JISEM_04/25_GD13501
Paper Title: Low-Light Image Super-Resolution Using GANs: A Comprehensive Comparative Review

We are happy to inform you that your research paper is accepted to be published in "Journal of Information Systems Engineering and Management".

Manuscript template: Click to Download

Copyright Form: Click to Download

A.3 INDEXING OF JOURNAL PROOF

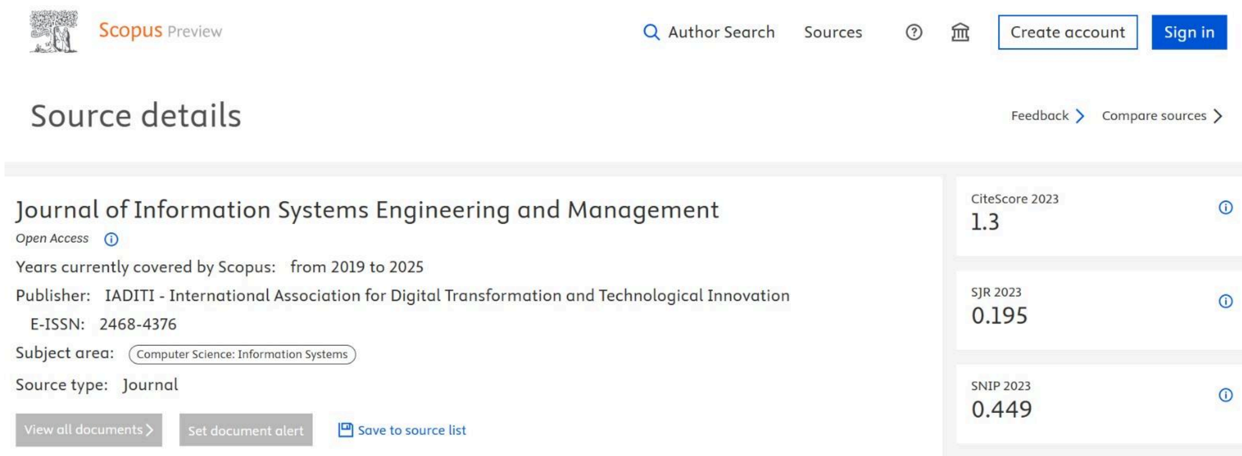

A.4 PUBLICATION OF JOURNAL PROOF

Journal of Information Systems Engineering and Management

Home / Archives / Vol. 10 No. 49s (2025) / Articles

Low-Light Image Super-Resolution Using GANs: A Comprehensive Comparative Review

PDF

DOI: https://doi.org/10.52783/jisem.v10i49s.10024

Keywords:
Low-light image enhancement, Image super-resolution, Generative adversarial networks

Niharika Pachori, Sanjay Patidar

Abstract

Image acquisition under low-light conditions poses serious limitations across numerous imaging domains, resulting in noisy, low-contrast, and resolution-degraded outputs. These limitations not only impact visual discriminability but also lead to disruption in downstream processes such as detection, recognition, and interpretation. Traditional image enhancement techniques, including histogram equalization and gamma correction, provide limited improvement in complex low-light scenarios and often amplify noise or distort colours. In contrast, Generative Adversarial Networks (GANs) have demonstrated significant success in both enhancing brightness and performing super-resolution in a data-driven manner. Their ability to model complex visual distributions enables the recovery of realistic textures and structures from degraded inputs.

This paper presents a comprehensive comparative review of recent GAN-based approaches for low-light image super-resolution. We explore key architectural strategies, loss functions, dataset choices, and evaluation metrics across prominent models. The analysis addresses three core research questions: limitations in texture restoration, effectiveness of performance metrics, and generalization challenges in low-light super resolution models across diverse scenarios. Furthermore, we highlight real-world application areas including surveillance, autonomous systems, mobile imaging, and document analysis where these techniques are most impactful. The paper concludes by identifying persistent challenges and proposing future research directions aimed at improving perceptual realism and robustness in low-light SR systems.

Issue

Vol. 10 No. 49s (2025)

Section

Articles

Announcements

Call for Papers

Call for Papers for the New Issue.
Last Date of Submission: ***June 30th, 2025***

Make a Submission

JOURNAL ARCHIVE

- Volume 10 (2025)
- Volume 9 (2024)
- Volume 8 (2023)
- Volume 7 (2022)
- Volume 6 (2021)
- Volume 5 (2020)
- Volume 4 (2019)
- Volume 3 (2018)
- Volume 2 (2017)
- Volume 1 (2016)

Indexed By

ELSEVIER Scopus

Downloads

Copyright Form
Paper Template

A.5 JOURNAL PAPER REGISTRATION RECEIPT

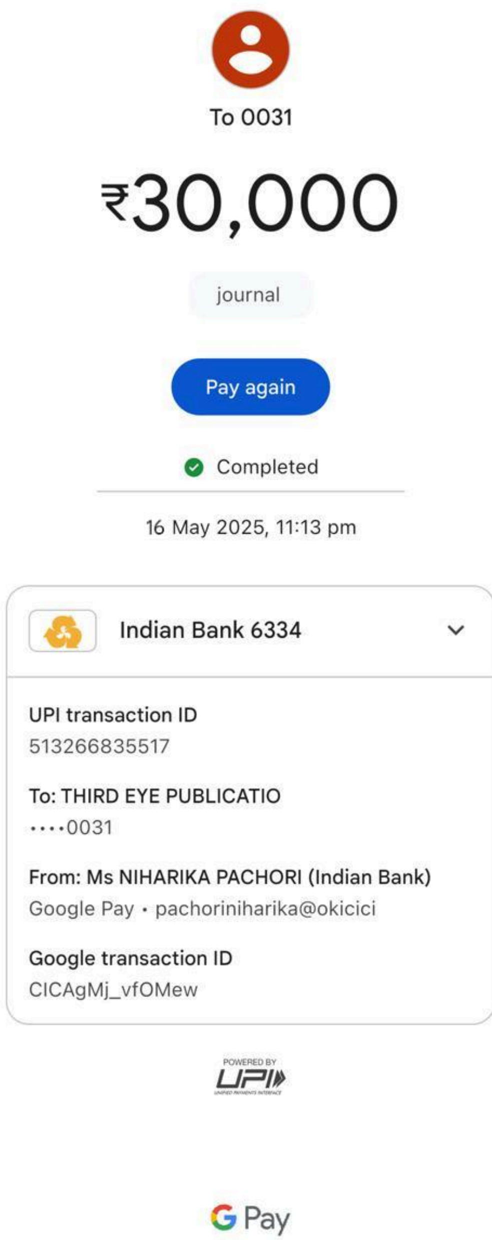

To 0031

₹30,000

journal

Pay again

Completed

16 May 2025, 11:13 pm

Indian Bank 6334

UPI transaction ID
513266835517

To: THIRD EYE PUBLICATIO
••••0031

From: Ms NIHARIKA PACHORI (Indian Bank)
Google Pay • pachoriniharika@okicici

Google transaction ID
CICAgMj_vfOMew

POWERED BY UPI

G Pay

DECLARATION

We/I hereby certify that the work which is presented in the Major Project-II/Research Work entitled ____ in fulfilment of the requirement for the award of the Degree of Bachelor/Master of Technology in ____________________ and submitted to the Department of ___________________________, Delhi Technological University, Delhi is an authentic record of my/our own, carried out during a period from ____________, under the supervision of ___________________________.

The matter presented in this report/thesis has not been submitted by us/me for the award of any other degree of this or any other Institute/University. The work has been published/accepted/communicated in SCI/ SCI expanded/SSCI/Scopus indexed journal OR peer reviewed Scopus indexed conference with the following details:

Title of the Paper:
Author names (in sequence as per research paper):
Name of Conference/Journal:
Conference Dates with venue (if applicable):
Have you registered for the conference (Yes/No)?:
Status of paper (Accepted/Published/Communicated):
Date of paper communication:
Date of paper acceptance:
Date of paper publication:

Student(s) Roll No., Name and Signature

SUPERVISOR CERTIFICATE

To the best of my knowledge, the above work has not been submitted in part or full for any Degree or Diploma to this University or elsewhere. I, further certify that the publication and indexing information given by the students is correct.

Place: ______________ **Supervisor Name and Signature**

Date: ______________

NOTE: PLEASE ENCLOSE RESEARCH PAPER ACCEPTANCE/ PUBLICATION/COMMUNICATION PROOF ALONG WITH SCOPUS INDEXING PROOF (Conference Website OR Science Direct in case of Journal Publication).