

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

Niharika Pachori¹, Dr. Sanjay Patidar²

¹ MTech candidate, Department of Software Engineering, Delhi Technological University, Delhi 110042, India, niharikapachori09@gmail.com

² Asst. Professor, Department of Software Engineering, Delhi Technological University, Delhi 110042, India, sanjaypatidar@dtu.ac.in

ARTICLE INFO

Received: 29 Dec 2024

Revised: 12 Feb 2025

Accepted: 27 Feb 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 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.

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

INTRODUCTION

Low-light image recording with high quality is still a lingering problem in consumer and industrial image applications. Images recorded under low light conditions leads to low contrast, too much noise, and loss of detail and are difficult to process and analyze. Such degradations not only restrict visual quality but also the accuracy of critical downstream operations such as object recognition, surveillance, and autonomous navigation.

Super-resolution (SR) algorithms seek to create images with high-resolution from their low-resolution version. When implemented for low-light inputs, this problem becomes even more challenging due to cumulative issues like underexposure, noise, and texture loss. Conventional enhancement methods only provide partial relief, tending to exaggerate artifacts or not generalize well over varied lighting conditions.

Deep learning in the form of Generative Adversarial Networks (GANs) has exhibited strong capabilities in enhancing and enhancing low-light images simultaneously over the past few years. GAN models based on perceptual learning and adversarial training use visually coherent textures and retrieve structural detail to produce. Present techniques are still wanting in areas such as detail retention in very dark regions, correct evaluation, and generalization across domains. This paper presents a comparative review of recent GAN-based techniques for low-light image super-resolution. It critically examines architectural innovations, loss functions, and dataset usage, while also evaluating

key performance metrics and practical deployment challenges. The objective is to provide a consolidated picture of the current state of the art, uncover research gaps, and recommend directions for future research.

The following paper ahead is structured in the way such as Section 2 presents related work and model progressions. Section 3 outlines key challenges and research objectives. Section 4 explains the methodology used to analyse and compare models. Section 5 discusses results in the context of our research questions. Section 6 highlights real-world use cases, and Section 7 concludes with insights and future directions. References are listed in Section 8.

RELATED WORK

Low-light image enhancement has evolved significantly over the past two decades from traditional image processing methods. Histogram equalization, which is one of the simplest approaches, relies on redistributing the pixel intensities to constitute an image with enhanced contrast. Although it is successful in certain situations, it also leads to over-saturation and amplifies the noise in the comparatively uniform dark regions [1]. To adjust brightness in a more controlled manner, gamma correction applies a nonlinear transformation to pixel values to improve visibility in dark areas without perturbing luminance equilibrium [2]. Another important class of techniques stems from Retinex theory, which models human visual perception by separating an image into reflectance and illumination [3]. These conventional approaches are largely interpretable and speedy but compromise the performance under heavy noise or non-uniform illumination conditions. Deep learning, and especially Generative Adversarial Networks (GANs), has shown remarkable competence in solving the challenging task of enhancement of low-light images in recent years.

Methods based on GANs are naturally suited for tasks requiring perceptual realism, e.g., low-light super-resolution. SRGAN is one of the earliest useful models that introduced perceptual loss to guide the generation of high-resolution images with improved preservation of texture and fine details [4]. ESRGAN, a later refinement, extended SRGAN by employing deeper residual blocks and a relativistic discriminator to enhance perceptual fidelity and maintain structural integrity [5]. Targeting low-light photographs, EnlightenGAN proposed a novel unsupervised fix to enhancement of low-light images through the assistance of a global and a local discriminator for realistic brightness and contrast learning from unpaired data [6]. Inspired by Retinex theory, DEGAN incorporated decomposition techniques into a GAN framework to facilitate the separation of illumination and reflectance and thereby suppress colour distortion and noise amplification common in low-light images [7].

Another direction involves integrating enhancement with tasks like demosaicking reconstructing full-colour images from raw sensor data. LLD-GAN employed a Wasserstein-based GAN model with gradient penalty, demonstrating robust performance on raw Bayer data from low-light sources [8]. LE-GAN advanced this by introducing illumination-aware attention mechanisms that allowed the network to dynamically adjust focus in underexposed regions, improving both naturalness and detail preservation [9]. For domain-specific use, LAE-GAN was designed to enhance legibility in low-light text documents, leveraging spatial attention to emphasize character boundaries and suppress background noise [10].

Most recently, StarSRGAN tackled the challenge of blind super-resolution under unknown degradation models. By combining multi-branch architectures and generative priors, it adapts to real-world degradation—often present in low-light environments—without requiring explicit noise modelling [11].

These developments illustrate a transition from handcrafted enhancement rules to data-driven, end-to-end learning systems. GANs have proven particularly effective at generating photorealistic, high-resolution images from degraded low-light inputs, offering flexibility and adaptability that traditional techniques lack.

CHALLENGES AND RESEARCH OBJECTIVES

1. Challenges of Low-Light Image Super-resolution

Below, we outline several critical challenges identified in recent studies concerning low-light image super-resolution. These insights highlight the shortcomings of current methods and the complexity of enhancing image quality under poor illumination conditions.

Loss of Texture and Fine Detail: GAN-based models such as SRGAN and ESRGAN have shown promise in recovering perceptual details. However, under severe low-light conditions, these models may hallucinate details or fail to accurately reconstruct textures and edges.

Color Distortion and Inconsistent Illumination: Uneven lighting and low signal-to-noise ratios can cause significant color shifts. Methods that rely on Retinex decomposition (e.g., DEGAN) attempt to address this, but often face challenges in generalizing across diverse lighting conditions.

Lack of Paired Training Data: Capturing paired set of low-light and well-lit high-resolution image datasets is difficult, especially in real-world settings. Unsupervised approaches like EnlightenGAN attempt to mitigate this, but training remains complex.

Evaluation Difficulties: Standard metrics like SSIM and PSNR often fail to match with human perception, especially in assessing perceptual quality in very dark regions. This raises questions about how best to benchmark models.

2. Research Questions

RQ 1: What are the limitations of current GAN-based models in accurately reconstructing fine textures and structural details in severely underexposed regions of low-light images?

RQ 2: What key performance metrics are most effective in evaluating low-light image super-resolution models, and how well do they correlate with perceptual quality and real-world usability?

RQ 3: What are the challenges in training super-resolution models that generalize across diverse low-light conditions?

METHODOLOGY

1. Low-Light Image Super-Resolution

Low-light image super-resolution (SR) involves generating a high-resolution (HR) image from a input that is a low-resolution (LR) image degraded by both underexposure and noise. In such conditions, traditional SR models fail to reconstruct fine details and often enhance noise or artifacts. This methodology focuses on leveraging deep learning, specifically Generative Adversarial Networks (GANs) to improve visual quality, perceptual realism, and structural accuracy in low-light conditions.

The process typically follows a two-stage pipeline:

- **Enhancement:** Improve brightness and reduce noise.
- **Super-Resolution:** Up sample the enhanced image to a higher resolution while preserving perceptual details.

Rather than treating these as separate stages, modern GAN-based approaches perform enhancement and SR in a single end-to-end trainable network, learning directly from data.

2. Generative Adversarial Network (GANs)




Figure 1: Generative Adversarial Network Architecture Diagram

Generative Adversarial Networks (GANs) consist of two neural networks, Generator (G) and Discriminator (D), and are trained in competition. The generator aims to create real high-resolution images from low-light LR inputs and the discriminator aims to classify real high-resolution images and generated images.

Gan Objective Function:

$$\min_G \max_D V(D, G) = E_{X \sim P_{data}(x)} [\text{Log} D(X)] + E_{Z \sim P_Z(Z)} [\text{Log}(1 - D(G(Z)))]$$

Perceptual and Content Losses:

To improve perceptual quality, many models incorporate content loss and perceptual loss.

- Content Loss is typically the Mean Squared Error (MSE) between the generated image and the ground truth HR image.

$$L_{content} = ||G(z) - x||_2^2$$

- Perceptual Loss uses features from a pretrained VGG network to measure high-level similarity:

$$L_{perc} = ||\Phi_{VGG}(x) - \Phi_{VGG}(\hat{x})||_2^2$$

- Total Loss Function:

$$L_{total} = \lambda_1 L_{content} + \lambda_2 L_{perc} + \lambda_3 L_{GAN}$$

The weights $\lambda_1, \lambda_2, \lambda_3$ balance the contributions of each component.

3. Standard Datasets

To simulate low-light degradation, training data may be artificially darkened using gamma correction, noise addition, and down sampling. To train and evaluate low-light SR models, publicly available datasets are used. These include paired and unpaired low-light and high-resolution image sets captured under various lighting and resolution conditions:

Dataset	Type	Content Description	Paired Data	Applications	Format / Size
LOL [12]	Real-world	500 image pairs captured under low and normal lighting conditions	Yes	Low-light enhancement, supervised GAN training	JPEG, 500 pairs
SID [13]	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 [14]	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 [15]	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 [16]	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 [17]	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

Table 1. Low-Light Image Benchmark Datasets

4. Performance Metrics

PSNR (Peak Signal-to-Noise):

$$\text{PSNR} = 10 \log_{10} \left(\frac{\text{MAX}_I^2}{\text{MSE}} \right)$$

Peak Signal-to-Noise ratio measures the pixel-level difference between images, mainly capturing noise and compression artifacts.

SSIM (Structural Similarity Index Measure):

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)}$$

Structural Similarity Index Measure evaluates structural similarity, focusing on how image luminance, contrast, and structure align.

LPIPS (Learned Perceptual Image Patch Similarity):

$$\text{LPIPS}(x, y) = \sum_i \left(1/(H_i W_i) \right) \sum_{h=1}^{H_i} \sum_{w=1}^{W_i} \|w_i \odot (\hat{y}_i^{hw} - \hat{x}_i^{hw})\|^2$$

Learned Perceptual Image Patch Similarity compares perceptual similarity using deep neural network features, reflecting human visual judgments better.

RESULTS

RQ1: What are the limitations of current GAN-based models in accurately reconstructing fine textures and structural details in severely underexposed regions of low-light images?

While GAN-based models have demonstrated notable improvements in low-light image enhancement and super-resolution, their ability to faithfully reconstruct fine textures and structural details in severely underexposed regions remains limited. Several factors contribute to this:

- **Hallucination and texture smoothing:** Models like SRGAN and ESRGAN generate perceptual features but sometimes hallucinate textures not present in the input or blur fine details due to aggressive perceptual loss optimization.
- **Noise amplification and detail confusion:** Under extreme low-light conditions, residual noise is often mistaken for high-frequency image components, leading to inaccurate texture restoration.
- **Architecture depth and feature reuse limitations:** Earlier architectures lack the adaptive attention or decomposition required to distinguish between reflectance, noise, and illumination in low-visibility scenarios.

The table below outlines key GAN-based models, highlighting their architectural choices generator architecture, discriminator architecture, and loss functions, which directly influence their performance in restoring structural integrity under low-light conditions:

Model	Year	Generator Architecture	Discriminator Architecture	Loss Functions
SRGAN [4]	2017	Deep ResNet with residual blocks	Patch-based discriminator	Adversarial loss, Content loss (MSE), Perceptual (VGG) loss
ESRGAN [5]	2018	Residual-in-Residual Dense Blocks (RRDB)	Relativistic average discriminator	Perceptual loss, GAN loss (RaGAN), Pixel loss
EnlightenGAN [6]	2021	U-Net with global-local feature fusion	Dual discriminators (global + local)	Adversarial loss, Reconstruction loss
DEGAN [7]	2021	Decomposition network (Retinex-based)	Standard discriminator	Decomposition loss, Illumination-consistency loss
LLD-GAN [8]	2024	End-to-end RAW image pipeline	Wasserstein GAN with gradient penalty	Wasserstein loss, Pixel loss
LE-GAN [9]	2022	Attention-augmented U-Net	Patch discriminator	Identity loss, Adversarial loss
LAE-GAN [10]	2023	Attention + text-aware enhancement modules	Convolutional discriminator	Attention loss, Reconstruction loss
StarSRGAN [11]	2023	Multi-branch GAN with generative priors	Relativistic GAN	GAN loss, Content loss, Feature similarity loss

Table 2. Summary GAN-based Models for Low-Light Super-Resolution

RQ 2: What key performance metrics are most effective in evaluating low-light image super-resolution models, and how well do they correlate with perceptual quality and real-world usability?

The evaluation of low-light SR models hinges on metrics that balance fidelity, structural consistency, and perceptual quality. Traditional pixel-based metrics often fail to reflect human visual perception in low-light contexts, while newer, feature-based methods offer more realistic assessments.

- PSNR and SSIM, while standard, are inadequate in capturing perceptual realism in enhanced low-light images, particularly when hallucinated textures are involved.
- LPIPS has emerged as better metrics for perceptual evaluation but depend on pretrained models and large batch sizes.

Metric	Formula / Basis	Primary Use Case	Strengths	Limitations
PSNR	$PSNR = 10 \log_{10} \left(\frac{L^2}{MSE} \right)$	Measures pixel-wise accuracy	Easy to compute, widely used	Poor correlation with visual quality, especially under perceptual distortions
SSIM	Based on luminance, contrast, structure comparison	Captures structural similarity	Better than PSNR for structure	Sensitive to global contrast shifts
LPIPS	Deep feature distance from VGG or AlexNet	Evaluates perceptual closeness	Aligns with human perception	Computationally expensive, requires pretrained nets

Table 2. Summary of Performance Metrics for Low-Light Super-Resolution

RQ 3: What are the challenges in training super-resolution models that generalize across diverse low-light conditions?

Training a low-light SR model involves a multitude of challenges stemming from the nature of low-light data and the diversity of real-world degradation patterns. Key challenges include:

1. **Data Diversity and Distribution Gap:** Many datasets represent specific lighting conditions or sensor types. A model trained on LOL or SID may not generalize to nighttime street scenes or indoor environments without domain adaptation.
2. **Noise Modeling Complexity:** Real low-light images often contain non-Gaussian noise, compression artifacts, and lens flares. Most models assume simplified noise distributions, which hinders performance under actual capture conditions.
3. **Balancing Loss Functions:** Designing an effective training objective is non-trivial. Overweight perceptual loss can lead to hallucinated textures, while focusing too much on pixel accuracy may blur details.
4. **Overfitting to Dataset Bias:** GANs, particularly when trained on narrow distributions, can memorize artifacts or lighting patterns specific to the training set, thus failing to generalize to unseen lighting geometries or exposure levels.
5. **Real-Time Inference Constraints:** Models with complex architectures (e.g., multi-branch generators or attention-heavy networks) may be unsuitable for real-time or embedded use due to high computational demands.

USECASE

Low-light SRCNN is of practical value as it can be applied to numerous real-world cases where vision is limited. GAN-based architectures aid in improving detail and definition, and have applications in various areas like surveillance, medical imaging, autonomous driving, and document analysis. This section discusses some important pieces of evidence that these approaches are beneficial.

1. Surveillance and Security Systems

In surveillance, improving clarity of an image, especially under low illumination conditions, is important for identification of faces and licence plates. Superresolution methods, along with techniques such as exposure fusing and color correction, can have a remarkable impact on the quality of a video in these conditions [19]. The quality of these improvements, however, depends on the quality of the reference data used for evaluation, as ground truth of low quality may also result in the biased performance [18].

2. Medical Imaging

In radiological imaging, such as CT and endoscopy, image resolution enhancement with minimal exposure to radiation is crucial for successful diagnoses. Super-resolution techniques enabled by deep learning, like GANs, can reconstitute high-resolution images from low-resolution inputs without compromising patient safety [20].

3. Autonomous Driving and Robotics

Autonomous cars and drones often perceive the world with limited ambient illumination (e.g. night or fog scenarios). Super-resolved images provide better image quality and subsequently stronger object detection, path planning, and obstacle avoidance. Recent developments, such as second-order attention network (SAN), allow more expressive feature representation and detail recovery even for the case of low resolution, which further boosts the detection accuracy of RT systems [21]. Besides, RefSR [22] could adaptivelearns the textures from the similar references to generalize the high resolution, when the reference is also irrelevant to the input.

4. Satellite and Aerial Imaging

In remote sensing, satellite and drone images are often captured in low-light or poor weather conditions. GAN-based super-resolution (SR) methods can improve terrain recognition and resource monitoring by enhancing spatial detail. For instance, a CNN model used for the Sentinel-2 mission super-resolves lower-resolution bands, improving spatial accuracy from 20 m and 60 m to 10 m Ground Sampling Distance (GSD), with improved performance in both quantitative metrics and visual quality [23].

CONCLUSION

Generative Adversarial Networks have come forth as powerful tools for handling the difficulties of low-light image super-resolution. Through this review, we have examined various GAN-based models, from early frameworks like SRGAN to advanced approaches such as StarSRGAN and LE-GAN. While progress has been substantial, critical issues remain, including poor texture reconstruction in severely underexposed regions, unreliable evaluation metrics that poorly reflect perceptual quality, and limited model generalization due to dataset bias.

This study emphasizes the need for models that are not only perceptually accurate but also robust across diverse real-world lighting conditions. Addressing these challenges requires a combination of better architectural designs, more representative datasets, and task-aware evaluation strategies. By highlighting these gaps, we aim to support the development of next-generation models capable of delivering consistent performance in practical low-light scenarios.

REFERENCES

- [1] Stephen M. Pizer, undefined., et al. "Adaptive histogram equalization and its variations," in *Computer Vision, Graphics, and Image Processing*, vol. 39, no. 3, pp. 355-368, 1987.
- [2] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 3rd ed. Pearson, 2008.
- [3] E. Land, J. McCann. "Lightness and Retinex Theory," in *Journal of the Optical Society of America*, vol. 61, pp. 1-11, 1971.
- [4] Christian Ledig, undefined., et al, "Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network," 2017.
- [5] Xintao Wang, undefined., et al, "ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks," 2018.
- [6] Y. Jiang *et al.*, "EnlightenGAN: Deep Light Enhancement Without Paired Supervision," in *IEEE Transactions on Image Processing*, vol. 30, pp. 2340-2349, 2021, doi: 10.1109/TIP.2021.3051462.
- [7] R. Liu, L. Ma, J. Zhang, X. Fan and Z. Luo, "Retinex-inspired Unrolling with Cooperative Prior Architecture Search for Low-light Image Enhancement," *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Nashville, TN, USA, 2021, pp. 10556-10565, doi: 10.1109/CVPR46437.2021.01042.
- [8] Li Wang, undefined., et al. "LLD-GAN: An end-to-end network for low-light image demosaicking," in *Displays*, vol. 85, pp. 102856, 2024.
- [9] Ying Fu, undefined., et al. "LE-GAN: Unsupervised low-light image enhancement network using attention module and identity invariant loss," in *Knowledge-Based Systems*, vol. 240, pp. 108010, 2022.
- [10] Xue, M., He, Y., Xie, P. *et al.* LAE-GAN: a novel cloud-based Low-light Attention Enhancement Generative Adversarial Network for unpaired text images. *J Cloud Comp* **12**, 160 (2023). <https://doi.org/10.1186/s13677-023-00533-4>
- [11] K. Vo, L. Bui, "StarSRGAN: Improving Real-World Blind Super-Resolution," in *Computer Science Research Notes*, 2023, pp. 62-72.
- [12] C. Wei, W. Wang, W. Yang, and J. Liu, 'Deep Retinex Decomposition for Low-Light Enhancement', *arXiv [cs.CV]*. 2018.
- [13] C. Chen, Q. Chen, J. Xu, and V. Koltun, 'Learning to See in the Dark', *arXiv [cs.CV]*. 2018.
- [14] Loh, Y., & Chan, C. (2019). Getting to Know Low-light Images with The Exclusively Dark Dataset. *Computer Vision and Image Understanding*, 178, 30-42.
- [15] C. Lee, C. Lee and C. -S. Kim, "Contrast enhancement based on layered difference representation," *2012 19th IEEE International Conference on Image Processing*, Orlando, FL, USA, 2012, pp. 965-968, doi: 10.1109/ICIP.2012.6467022.
- [16] Cai, J., Gu, S., & Zhang, L. (2018). Learning a Deep Single Image Contrast Enhancer from Multi-Exposure Images. *IEEE Transactions on Image Processing*, 27(4), 2049-2062.
- [17] Vladimir Bychkovsky, Sylvain Paris, Eric Chan, & Fredo Durand (2011). Learning Photographic Global Tonal Adjustment with a Database of Input / Output Image Pairs. In *The Twenty-Fourth IEEE Conference on Computer Vision and Pattern Recognition*.
- [18] S. Su, J. M. Rocafort, D. Xue, D. Serrano-Lozano, L. Sun, and J. Vazquez-Corral, 'Rethinking Image Evaluation in Super-Resolution', *arXiv [cs.CV]*. 2025.

- [19] Keerthi, Halaharvi & B, Sreepathi. (2023). Enhancement of Low Light Videos using Artificial Intelligence. *International Journal for Research in Applied Science and Engineering Technology*. 11. 5933-5938. 10.22214/ijraset.2023.52903.
- [20] C. You *et al.*, 'CT Super-Resolution GAN Constrained by the Identical, Residual, and Cycle Learning Ensemble (GAN-CIRCLE)', *IEEE Transactions on Medical Imaging*, vol. 39, no. 1, pp. 188–203, Jan. 2020.
- [21] T. Dai, J. Cai, Y. Zhang, S. -T. Xia and L. Zhang, "Second-Order Attention Network for Single Image Super-Resolution," *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, USA, 2019, pp. 11057-11066, doi: 10.1109/CVPR.2019.01132.
- [22] Z. Zhang, Z. Wang, Z. Lin, and H. Qi, 'Image Super-Resolution by Neural Texture Transfer', *arXiv [cs.CV]*. 2019.
- [23] C. Lanaras, J. Bioucas-Dias, S. Galliani, E. Baltsavias, and K. Schindler, 'Super-resolution of Sentinel-2 images: Learning a globally applicable deep neural network', *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 146, pp. 305–319, Dec. 2018.