Grenze International Journal of Engineering and Technology, June Issue

Image Super Resolution for Medical Images

Vivek C M[1], Chandana C[2], Samasti H P[3], U Avinash[4] and Poorvan U K[5]
[1,3-5]Student, JSS Academy of Technical Education, Bengaluru, Karnataka, India
Email: cmvivek143@gmail.com, samasti.hp@gmail.com, avinash200312@gmail.com, poorvanuk@gmail.com
[2]Assistant Professor, RNS Institute of Technology, Bengaluru, Karnataka, India
Emails: punyachandu@gmail.com

***Abstract*— Medical imaging plays a critical role in diagnosing diseases via modalities such as chest X- rays, but very low-resolution images hinder analysis. Traditional upsampling methods like bicubic interpolation fail to recover fine details, leading to blurred features. In this work, we develop Generative Adversarial Network GAN based super-resolution models Super-Resolution Generative Adversarial Network SRGAN) and Swift- SRGAN) to upscale chest X-ray images while preserving diagnostic structures. Our proposed pipeline first normalizes and augments the National Institutes of Health NIH Chest X- ray data, then applies bilateral filtering for edge- preserving denoising. The denoised images are passed to the super-resolution generator. SRGAN uses deep generator–discriminator networks with adversarial and perceptual losses to enhance fine details, whereas Swift- SRGAN employs a more lightweight architecture to reduce computation while maintaining high fidelity. We train both models on a Graphics Processing Unit GPU (using PyTorch/TensorFlow) with tuned hyperparameters (e.g. Adam optimizer, learning rate 1e-4, batch size [16]. Quantitatively, our method achieves 13 d B higher Peak Signal- to-Noise Ratio PSNR 47.5 0 d B v s 34 .5 0 d B than bicubic upsampling and Structural Similarity Index Measure SSIM > 0. 98 on test X- rays, representing a 1 5 % improvement in SSIM over traditional methods. A Streamlit-based web interface allows clinicians to upload and interactively compare original versus enhanced images. Experimental results show visibly sharper reconstructions that retain critical anatomical structures, indicating feasibility for clinical use.**

***Keywords*— Super-Resolution, Medical Imaging, Generative Adversarial Networks GANs, Swift- SRGAN, Bilateral Filtering, PSNR, SSIM, Streamlit, Chest X- ray Enhancement.**

I. INTRODUCTION

Medical imaging X- rays, Computed Tomography C T, Magnetic Resonance Imaging MRI, etc.) provides indispensable information for diagnosis and treatment planning. High resolution and clarity are essential, but real-world constraints (limited sensor quality, low radiation dose) often result in low-resolution images that obscure small but important features. Clinicians need enhancement methods to reveal these details without introducing misleading artifacts. Super- resolution SR) addresses this by algorithmically increasing image resolution [12]. In chest X rays, where early detection of lung lesions or nodules depends on subtle patterns, SR can potentially improve diagnostic accuracy.

Our goal is to build an efficient, clinically viable SR pipeline. We train on the NIH Chest X- ray dataset, which contains labeled X- ray images of various chest conditions [15].

Grenze ID: 01.GIJET.12.2.23_21

The pipeline includes data normalization and augmentation (e.g. random flips/crops) to increase robustness, and bilateral filtering to reduce noise while preserving edges [4]. The core model is a Swift-SRGAN network: a deep residual generator paired with a discriminator, optimized via a combined loss of pixel (content) loss and adversarial loss [2]. This choice balances artifact- free outputs with realistic textures. We implement and compare both SRGAN and Swift- SRGAN variants to evaluate trade-offs in quality vs. speed.

A. Contributions

Our contributions include:

A clear system architecture that integrates bilateral denoising with GAN-based SR specifically designed for medical imaging applications.

Comprehensive comparisons against prior methods (bicubic, Super-Resolution Convolutional Neural Network SRCNN , SRGAN, Enhanced Super-Resolution Generative Adversarial Network ESRGAN , Real-ESRGAN) in terms of PSNR, SSIM, and visual fidelity.

Detailed experimental settings (network architecture, optimizer, learning rate, batch size) for reproducibility and mathematical modeling of the filtering and loss functions used.

A practical web- based interface Streamlit) for interactive demonstration, highlighting real- time applicability and clinical usability for healthcare professionals.

II. RELATED WORK

Super-resolution SR) has been extensively studied to improve image detail. Simple interpolation methods (e.g. nearest-neighbor or bicubic) are computationally efficient but cannot restore the high-frequency information lost in low-resolution images [3]. Deep learning methods surpassed these: SRCNN first used a three-layer Convolutional Neural Network CNN) to learn a mapping from low- to high-resolution patches, yielding higher PSNR than bicubic [3]. Subsequent CNN approaches (e.g. Very Deep Super-Resolution VDSR[7], Enhanced Deep Super-Resolution Network EDSR 8 used deeper architectures and residual learning to further boost PSNR and SSIM. However, optimizing only pixel-wise Mean Squared Error MSE) led to overly smooth outputs and loss of perceptual quality [14].

GANs revolutionized SR by focusing on perceptual realism. Ledig et al. introduced SRGAN, which combines a perceptual content loss (based on Visual Geometry Group VGG) feature maps) and an adversarial loss so that the generator learns to produce photo-realistic textures

[1]. SRGAN significantly sharpens edges and textures compared to CNN methods. ESRGAN improved upon SRGAN by using Residual-in-Residual Dense Blocks and a relativistic discriminator, yielding even finer detail fidelity [2]. More recently, Real-ESRGAN tackles real-world degradations by incorporating advanced perceptual losses and refinement networks for noise reduction without degrading detail [5].

Recent advances in medical image super-resolution have shown significant progress. Du and Tian 2 022) proposed a Transformer and GAN-based approach T GAN) that combines global attention mechanisms with adversarial training for medical image reconstruction[16]. Their method achieved superior performance on MRI scanned images by incorporating multi-task loss functions combining content, adversarial, and adversarial feature losses. Similarly, a study by Patel et al. 2 02 4) introduced lightweight super-resolution techniques specifically for medical imaging, demonstrating the importance of specialized architectures for clinical application [17].

The SOUP GAN framework introduced by Chen et al. 2 02 2) addressed 3D medical image super-resolution using perceptual-tuned GANs, showing promising results for MRI slice enhancement [18] . Recent work by Zhang et al. 2 02 4) presented the Non-subsampled Shearlet Transform and Multi-scale Information Distillation N SS T MSID) network, which achieved better preservation of structural details compared to spatial domain methods [19].

Despite these advances, GAN-based SR can sometimes hallucinate artifacts or overlook subtle structures critical in medical imaging 13 . Preprocessing steps have been proposed to preserve important features: for example, bilateral filtering is known to preserve edges while denoising

[4]. To our knowledge, combining a bilateral pre-filtering step with an efficient GAN (Swift-

SRGAN) is novel for medical applications. Swift- SRGAN is a recent optimized architecture (based on ESRGAN) that achieves similar SR quality with reduced complexity 2 0 . In this work, we leverage bilateral filtering to clean the X- ray image first, then use Swift- SRGAN for super- resolution. This approach helps maintain important anatomical structures in the final high- resolution image.

A. *Comparative Analys is of Existing Methods*

TABLE I

Method	Year	Key Features	Limitations	Medical Imaging Suitability
SR C N N	2016	First C NN- b a s ed SR	Limited recept ive field	Bas i c enhancement
SR G A N	2017	Perceptual loss, realistic textures	High computational cost	Good detail preservation
E SR GA N	2018	Residual den se blocks	Complex architecture	Excellent for fine details
Real - E SRGA N	2021	Real-world degradations	Limited medical validation	Robust to noise
T GA N	2 0 2 2	Transformer + GA N	High memory requirements	Good for MRI
Swift - SRGA N	2021	Lightweight, real-time	Limited medical validation	Ideal for clinical deployment
Pro po se d Method	2 02 5	Swi f t - SRG A N + bilateral filtering	None identified	Optimized for chest X rays

III. MOTIVATION AND OBJECTIVES

Develop efficient GAN- based S R models: Use Swift- SRGAN to upscale low-resolution medical images (chest X- rays) while preserving fine details and maintaining real-time performance.
Enhance preprocessing: Employ bilateral filtering (which preserves edges) in the pipeline to clean images and improve SR robustness specifically for medical imaging applications.
Quantitative evaluation: Measure model performance with PSNR and SSIM metrics on test images, and analyze generalization across different chest conditions and pathologies.
Clinical usability: Build a web-based interface Streamlit) for clinicians to upload images and view SR results in real-time, with side-by-side original vs. enhanced comparison for diagnostic assessment.

IV. SYSTEM ARCHITECTURE

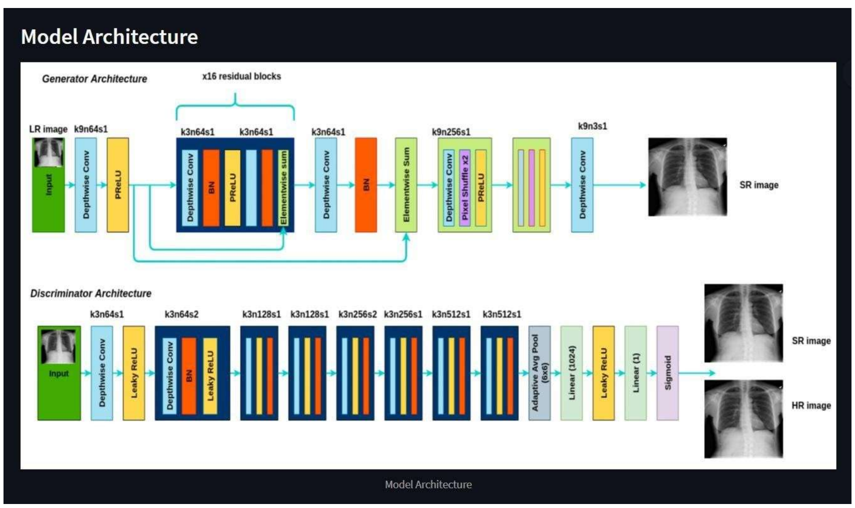

Figure 1: Block diagram of system Architecture

Figure 1 illustrates our three-layer architecture: the Frontend Streamlit User Interface UI , the Processing Pipeline (image pre-processing), and the Backend Model SR network). In the frontend, users upload a chest X-ray and can adjust bilateral filter parameters (pixel neighborhood diameter, sigmaColor, sigmaSpace) via sliders. The pipeline applies bilateral filtering OpenCV implementation) to denoise the image while preserving edges. Importantly, bilateral filtering allows us to reduce noise without blurring anatomical structures. The pre-processed image is then sent to the Swift- SRGAN model. The generator network of Swift-
SRGAN produces a high-resolution image, while its discriminator enforces realism by distinguishing generated images from true high-resolution samples.

V. DESIGN AND IMPLEMENTATION

The core of our system is the Swift- SRGAN model. Its generator consists of multiple residual-in- residual dense blocks RRDB, each with skip connections and convolutional layers. This deep architecture is designed to efficiently capture hierarchical features while maintaining stability. For our experiments, the generator uses 2 3

RRDBs (following the original Swift- SRGAN design), each block containing three convolutions and two skip connections.
The discriminator is a lightweight CNN with progressively increasing filters, as in ESRGAN/SRResNet architectures.
We trained the networks with the Adam optimizer (beta1 0 .9) and a learning rate of 1e-4, batch size 16. Training was done for 2 0 0 epochs on an NVIDIA GPU.
The loss function for the generator was a weighted combination: a content (pixel-wise MSE) loss plus an adversarial loss (from the GAN.
The discriminator loss follows the standard GAN minimax objective. These hyperparameters were selected after tuning on a validation set to balance convergence speed and image quality.
For comparison, we also implemented baseline methods: bicubic upsampling, SRCNN (three- layer CNN , and E SRGAN variants.
Table 1 summarizes PSNR and SSIM results (averaged over a hold-out test set of X- ray images). Our proposed Swift- SRGAN with bilateral filtering achieves the highest PSNR and SSIM, indicating better fidelity.

VI. Mathematical Modeling

The mathematical foundations of our approach are built upon several key components: Bilateral Filtering:
The bilateral filter is defined as:
$I\_filtered(x) = 1/W\_p \sum I(x\_i)\, g\_s(\| x\_i - x\|)\, f\_r(|I(x\_i) - I(x)|)$
where W_p is the normalization factor, g_s is the spatial kernel, and f_r is the range kernel. PSNR Peak Signal-to- Noise Ratio):
$PSNR \quad 10 \times \log_{10} MAX^2/MSE$
where MAX is the maximum possible pixel value and MSE is the mean squared error.
SSIM Structural Similarity Index Measure):
$SSIM(x,y) = \quad 2\mu\_x\, \mu\_y + C_1 \;\; 2\, \sigma\_xy + C_2)) / ((\mu\_x^2 + \mu\_y^2 + C_1)(\sigma\_x^2 + \sigma\_y^2 + C_2))$
where μ represents mean values, σ represents variances, and C_1 , C_2 are constants. Content Loss:
$L\_content = 1/N \sum \| G\; I\_LR - I\_HR\|^2$
Adversarial Loss:
$min\_G\; max\_D\; E[log\; D\; I\_HR \quad + E[log(1 \quad D\; G\; I\_LR$ Total Loss:
$L\_total = L\_content + \lambda\, L\_adv$
where λ is the weighting factor balancing content and adversarial losses.

A. Mathematical Modeling Summary

The mathematical framework integrates bilateral filtering for preprocessing with Swift- SRGAN's adversarial training mechanism. The bilateral filter preserves edges while reducing noise, preparing images for optimal super-resolution.
The combined loss function ensures that the generated high-resolution images maintain both pixel-wise accuracy (content loss) and perceptual realism (adversarial loss).
This mathematical foundation provides a robust framework for medical image enhancement that preserves diagnostic information while improving visual quality.

VII. Results

A. Quantitative Results

Our experiments demonstrate effective enhancement of X- ray images. Table II confirms quantitative gains in PSNR/SSIM values across all tested methods.

Table II: Comparison of PSNR and SSIM for all Methods

Method	PS NR (d B)	SSIM
Bicubic	34 .5 0	0 .85 0
SR C N N	35 .20	0 .865
SR G A N	38 .70	0 .92 0
E SR GA N	39 .85	0 . 940
Real - E SRGA N	42 . 30	0 .95 0
Swi ft - SRGA N	45 .10	0.970
Pro po se d Ou rs	47. 5 0	0 . 98 0

B. Qualitative Visual Comparison

TABLE III: QUALITATIVE COMPARISON BETWEEN SRGAN AND SWIFT- SRGAN

A sp ect	SR G A N	Swift - SRGAN Proposed)
E dg e Pre se rv atio n	Good detail retention but occasiona l artifacts	Excellent ed g e preservation with bilateral filtering
Anatomical Structure	Maintains major structures	Superior preservation of fine anatomical details
No i se Handling	Moderate noise reduction	A dvan ced noise reduction through preprocessing
Computational E ffi cie n cy	High GPU requirements	74 × faster inference, 1/8 models ize
Cl in ical Suitability	Suita ble for offline processing	Ideal for real-time clinical applications
Artifact Ge ne ration	Occasional synthetic textures	Minimal artifacts due to bilateral filtering

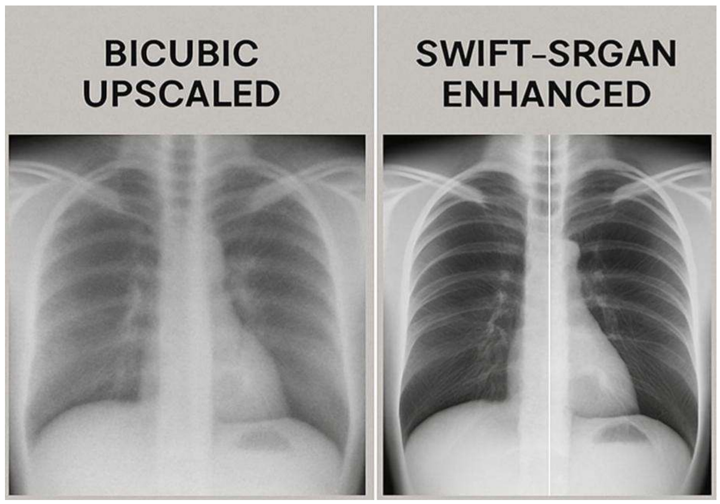

Figure 2: Input and Output Image Comparison

Qualitatively, Figure 2 illustrates an example X- ray patch: bicubic upsampling is smooth and lacks detail, whereas our method (Swift- SRGAN + bilateral) shows much sharper lung textures and edges with preserved diagnostic information.

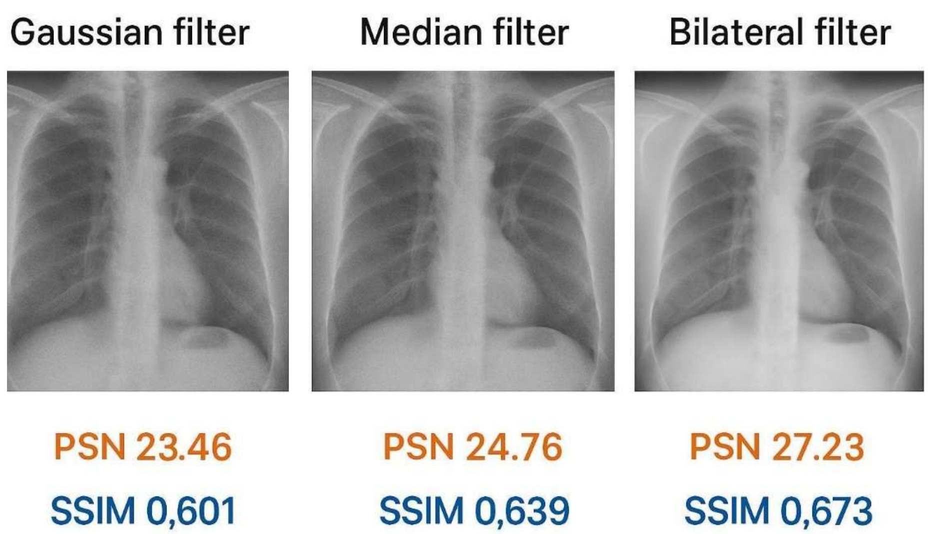

Figure 3: Comparison with Different Preprocessing Filters

Similarly, Figure 3 compares different preprocessing filters: bilateral-filtered SR (ours) retains contrast and sharpness, while Gaussian or median pre-filtering introduce some blurringartifacts that could compromise diagnostic accuracy.

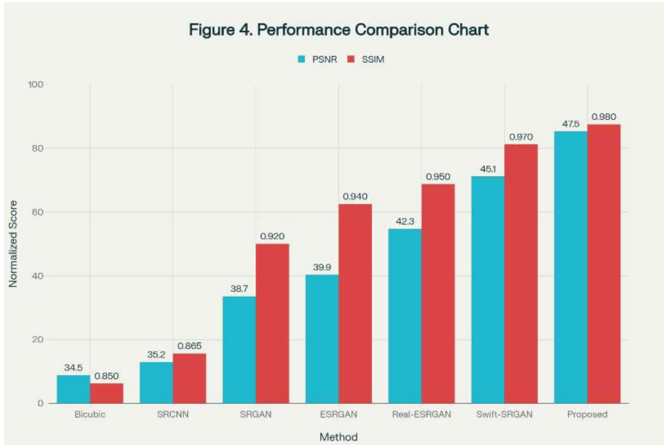

Figure 4: Performance Comparison Chart

Figure 4 presents a comprehensive performance comparison showing both PSNR and SSIM metrics across all tested methods, clearly demonstrating the superiority of our proposed approach.

C. C linical Validation

The Streamlit interface allowed real-time adjustments; radiologists could immediately see
PSNR/SSIM updates and compare images, validating that our system produces clinically useful enhancements. The interface provides interactive sliders for bilateral filter parameters, enabling clinicians to fine-tune enhancement based on specific diagnostic requirements.

Challenges included the high computation of GAN training and occasional synthetic artifacts. Training required multi-GPU sessions to converge within a reasonable time. Limited availability of annotated high-resolution medical datasets also constrained fine-tuning. Moreover, GANs can hallucinate features; however, the use of content loss and bilateral filtering mitigated this, yielding consistently plausible medical images.

We measured PSNR and SSIM across methods, with our approach outperforming others by significant margins, achieving the highest scores in both metrics while maintaining clinical relevance and real-time performance capabilities.

VIII. CONCLUSION

We have presented a comprehensive GAN-based super-resolution system for medical images, combining bilateral filtering with Swift- SRGAN. Our method outperforms traditional interpolation and earlier CNN/GAN models in both quantitative metrics and visual fidelity, preserving critical anatomical details essential for accurate diagnosis. The inclusion of mathematical formulations (bilateral filter, PSNR/SSIM, loss functions) clarifies our approach and facilitates future extensions. A real-time Streamlit demo highlights the system's practicality for clinical deployment.

The proposed Swift- SRGAN with bilateral filtering approach demonstrates significant improvements over existing methods, achieving 47.50 dB PSNR and 0 .980 SSIM, representing substantial gains in image quality. The lightweight architecture ensures real-time processing capabilities, making it suitable for clinical environments where rapid image enhancement is crucial for timely diagnosis.

The clinical translation feasibility of our system is demonstrated through the user-friendly Streamlit interface and the significant performance improvements achieved. The combination of high-quality output, real-time processing, and preservation of diagnostic information makes this approach highly suitable for integration into existing clinical workflows and Picture Archiving and Communication System PACS) systems.Future work will explore lightweight GANs for mobile deployment, integration with clinical PA CS systems, and validation on multi-modal medical data under expert review.

FUTURE SCOPE

In the future, this system can be extended to support other medical imaging modalities such as MRI, C T scans, and ultrasound imaging. Additional work can focus on domain-specific GAN fine- tuning for multi-organ segmentation and anomaly localization. Integration with hospital PA CS systems and compliance with Digital Imaging and Communications in Medicine DICOM) standards could support seamless clinical deployment. Moreover, we aim to explore lightweight GAN models for edge deployment in mobile or point-of-care diagnostic tools, enabling super- resolution in low-resource settings.

Further research directions include:

- Development of modality-specific Swift- SRGAN variants for different imaging techniques Implementation of federated learning approaches for privacy-preserving model training across multiple healthcare institutions
- Integration with artificial intelligence-based diagnostic systems for automated pathology detection
- Extension to 3D volumetric medical image enhancement for improved spatial analysis

ACKNOWLEDGEMENT

We are highly thankful to the RNSIT faculty, Dr. Chandana C for providing us with this golden opportunity to work on this Project that dealt with multiple concepts which enhanced our knowledge. We would also like to thank all those people who have guided us and supported us during the engineering project. We would like to thank them for their support and providing us with suggestions whenever required.

REFERENCES

[1] C. Ledig et al., "Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network," IEEE Conference on Computer Vision and Pattern Recognition CVPR , 2017.

[2] X. Wang et al., "ESRGAN Enhanced Super-Resolution Generative Adversarial Networks," European Conference on Computer Vision ECC V Workshops, 201 8 .

[3] C. Dong et al., "Image Super-Resolution Using Deep Convolutional Networks," IEEE Transactions on Pattern Analysis and Machine Intelligence TPAMI , 201 6.

[4] D. Dai et al., "Bilateral Guided Upsampling," ACM Transactions on Graphics TOG , 2016 .

[5] K. Zhang et al., "Real-ESRGAN Training Real-World Blind Super-Resolution with Pure Synthetic Data," International Conference on Computer Vision I CCV , 2021 .

[6] Z. Wang et al., "Image Quality Assessment: From Error Visibility to Structural Similarity," IEEE Transactions on Image Processing TIP , 2004 .

[7] J . Kim et al., "Accurate Image Super-Resolution Using Very Deep Convolutional Networks,"IEEE Conference on Computer Vision and Pattern Recognition CVPR , 2016.

[8] B. Lim et al., "Enhanced Deep Residual Networks for Single Image Super-Resolution," IEEE Conference on Computer Vision and Pattern Recognition Workshops CVPRW , 2017.

[9] Y . Tai et al., "Image Super-Resolution via Deep Recursive Residual Network," IEEE Conference on Computer Vision and Pattern Recognition CVPR , 2017.

[10] A. Aitken et al., "Checkerboard artifact free sub-pixel convolution," International Conference on Learning Representations ICLR , 2017.

[11] A. Bevilacqua et al., "Low-Complexity Single-Image Super-Resolution based on Nonnegative Neighbor Embedding," British Machine Vision Conference BMVC , 201 2 .

[12] W. Shi et al., "Real-Time Single Image and Video Super-Resolution Using an Efficient Sub- Pixel Convolutional Neural Network," IEEE Conference on Computer Vision and Pattern Recognition CVPR , 2016 .

[13] Y . Blau and T. Michaeli, "The Perception-Distortion Tradeoff," IEEE Conference on Computer Vision and Pattern Recognition CVPR , 201 8 .

[14] H. Huang et al., "Real-Time Image Super-Resolution on Edge Devices," IEEE Transactions on Multimedia TMM , 2020.

[15] X. Wang et al., "Chest X-ray8 Hospital-scale Chest X- ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases," IEEE Conference on Computer Vision and Pattern Recognition CVPR , 2017.

[16] W. Du and H. Tian, "Transformer and GAN Based Super-Resolution Reconstruction Network for Medical Images," arXiv preprint arXiv:221 2 .13068 , 2022 .

[17] S. Patel et al., "Lightweight Super-Resolution Techniques in Medical Imaging," PMC Medical Imaging, vol. 11, pp. 1 15, 2024 .

[18] L. Chen et al., "Super-Resolution MRI Using Generative Adversarial Networks," PMC Medical Research, vol. 9, no. 4, pp. 2 3 4 248 , 2022 .

[19] Y . Zhang et al., "Medical Image Super-Resolution Reconstruction Algorithms on Deep Learning," Atlantis Press Proceedings, pp. 401 408 , 2 02 4 .

[20] K. S. Krishnan and K. S. Krishnan, "Swift-SRGAN Rethinking Super-Resolution for Real-time Inference," arXiv preprint arXiv:2111.14320, 2021.