



Lightweight GAN-Based Image Super-Resolution for Satellite Imagery on Low-Power Embedded Hardware

Sumit Ramswami Punam^{1*}, B.K. Nestaresb²

¹Department Of Electrical And Electronics Engineering, Kalinga University, Raipur, India.

²Centro de Investigacion y Desarrollo de Tecnologias Aeronauticas (CITEA) FuerzaAerea Argentina Las Higuera, Cordoba, Argentina.

KEYWORDS:

Generative Adversarial Networks (GANs), Image Super-Resolution (ISR), Lightweight Deep Learning Models, Embedded AI, Low-Power Hardware, Satellite Imagery Enhancement, Quantization-Aware Training, Edge Inference, Remote Sensing, Onboard Image Processing, Real-Time ISR, UAV-Based Imaging, Energy-Efficient AI, Model Compression, SpaceNet Dataset.

ARTICLE HISTORY:

Submitted : 10.05.2025
Revised : 14.06.2025
Accepted : 23.07.2025

<https://doi.org/10.17051/NJSIP/01.04.04>

ABSTRACT

With the emerging trend of satellite-based earth observation systems and edge-enabled UAV platforms, there is increasing pressure to reconstruct image at a high resolution but also limited available computational or energy resources. In this paper, the authors would like to propose a lightweight GAN-based image super-resolution (ISR) pipeline that is targeted towards a real-time use case on low-power embedded systems. The framework combines depth-wise separable convolutions, structured channel pruning and quantization aware training to achieve very large model compression with little loss in fidelity of reconstruction of images. An architecture was benchmarked against publicly available satellite dataset and demonstrated up to 3.5 dB improvement in Peak Signal-to-Noise Ratio (PSNR) and 28.7 percent better Structural Similarity Index (SSIM) over regularly used interpolation method of bicubic interpolation. In comparison with the conventional SRGAN and ESRGAN alike, our model was, on the one hand, much smaller, as it reduced the size by 45% and, on the other, its performance was up to 60 percent faster to operate on NVIDIA Jetson Nano (or ARM Cortex-A based devices) with the given power consumption limit of 10W. These findings demonstrate the framework as capable of providing efficient ISR through high-quality operations done directly on the edge hardware without depending on a cloud-based AI computing facility. This makes the technique scalable and viable to practical applications as techniques in real time remote sensing in power limited environs like satellites onboard and Unmanned Air Vehicles (UAVs).

Author's e-mail: sumit.kant.dash@kalingauniversity.ac.in

How to cite this article: Punam SR, Nestaresb BK. Lightweight GAN-Based Image Super-Resolution for Satellite Imagery on Low-Power Embedded Hardware. National Journal of Signal and Image Processing, Vol. 1, No. 4, 2025 (pp. 23-29).

INTRODUCTION

The Earth observation satellite images are vital to the contemporary remote sensing practice, including precision farming and infrastructure management, environmental change observation and crisis management. Nevertheless, achieving high-resolution Spaceborne/UAV-based sensors directly presents its own set of internal limitations, such as limiting consecutive data storage or downlink bandwidth capabilities (useful for larger scaled sensor applications) as well as restricted power limits. These restrictions are particularly significant in edge applications on small satellites as well as unmanned aerial vehicles (UAVs) where available hardware resources are very limited.^[1]

The high-resolution (HR) images reconstructed by means of Image Super-Resolution (ISR) techniques, which upsample low-resolution (LR) images, offer a low-cost, power-efficient and practical solution to direct high-resolution imaging. Especially, deep learning has been proven to perform vastly better in generating perceptually accurate HR images with degraded inputs using Generative Adversarial Networks (GANs).^[1] In spite of their effectiveness, conventional GAN models like SRGAN and ESRGAN are expensive to run and also occupy a vast amount of memory, and this makes them impractical in running real time models on embedded platforms like the NVIDIA Jetson Nano or microprocessors on the ARM platform using the Cortex-A series of manufactures.

Current studies and works have already mostly aimed at achieving good reconstruction results at the expense of too little consideration about efficiency and/or deployments and compatibility on low-power/low-energy edge devices. In this paper, that gap shall be filled to present a proposal of the lightweight GAN-based ISR architecture, which is optimized by depthwise separable convolution, structured pruning and quantization-aware training to reduce the inference latency and memory need.

Section 2 defines the related works performing ISR and embedded deep learning. Section 3 gives the description of the proposed lightweight GAN architecture. The experimental procedure and measures of evaluation are described in Section 4. Section 5 provides a comparison of the results and Section 6 is a conclusion that summarizes the insights and directions to the future.

RELATED WORK

Image Super-Resolution (ISR) has been vastly made something that was studied with classical and deep learning methods. Bicubic interpolation and Lanczos filtering are traditional and are still very popular because they are simple and they are computationally efficient. Nonetheless, the resulting outputs with these techniques are sometimes too smooth and may not have high-frequency details and small textures (this is very important in satellite imagery where edge sharpness and object outlines are all significant factors). The advanced technology of deep learning transformed the way ISR works. Super-Resolution GAN (SRGAN) has been able to capitalize on the role of adversarial learning in order to propagate the perceptual quality of an image that was reconstructed.^[2] Then, ESRGAN applied residual-in-residual dense blocks and relativistic discriminators to improve realistic and detail restoration.^[3] Though such models achieve good (high) PSNR and SSIM, they are computationally intensive and memory-intensive, which makes them unsuitable to build deployments on low-cost embedded and edge platforms. To overcome the issue of resource constraints, researchers have introspected the employment of lightweight pathways in ISRs such as depth-wise separable convolutions,^[4] model pruning,^[5] and model quantization.^[6] By way of example, MobileNet-inspired networks as well as FastSR methods have shown to require smaller parameter numbers and quicker inference. Nevertheless, these improvements are usually achieved at the expense of poorer visual image quality or lack of good generalization to real-world remote sensing applications.

Research Gap: In the available ISR models, there exists a trade-off between accuracy and efficiency; however,

many of them are not directly optimized vis-a-vis satellite image super-resolution on weak embedded processors. What is more, real-time performance-related restrictions, including inference speed and thermal load, still are hardly discussed in hands-on applications.

Contribution Context: The framework will rather present an expansion on top of GAN paradigm by enforcing hardware-sensitive optimizations such as channel pruning, quantization-aware training, and low-rank harmonic structure. This allows real-time high-fidelity ISR on embedded devices like the Jetson Nano and ARM Cortex-A53, and it is therefore particularly appropriate for edge-enabled Earth observation applications.

METHODOLOGY

This section gives the architecture, and training strategy that will be deployed in bringing out the proposed lightweight based GAN ISR framework that will be deployed on embedded platforms like satellites and UAVs. It is interested in finding minimalistic models, keeping the perceived model instances, and conducting quantized inference without introducing large losses on the output.

Network Architecture

The proposed architecture based on GAN consists of two parts a generator and a discriminator which are both designed to perform inference in power-constrained systems (Figure 1).

- **Design of generators:** Depthwise separable generator convolution does not increase the spatial resolution of the output whose focal point operations are reduced by a factor of magnitude. Such design reduces computation cost approximately by 60 percent, when compared to regular convolutional blocks.^[8]
- **Spatial Skip Connections:** residual and skip connections are used in the generator to maintain the spatial fidelity in the satellite images, which is of utmost importance because of the sudden edges and texture characteristics of satellite images.
- **Discriminator Design:** A patchGAN-type discriminator regularly used in place of the classic batch normalization to accelerate convergence under minute batch sizes as characteristic of edge machines. This helps stabilize against adversarial perturbations and even against generalization of adversarial perturbations in the case of the limited memory training.




Fig. 1: Lightweight GAN Architecture for Image Super-Resolution

It is possible to simplify the generator/discriminator of the proposed GAN lightweight through the block diagram. The generator involves depthwise separable convolutions, skip connections and can generate enhanced images efficiently, whereas the discriminator benefits the use of group normalization and PatchGAN structure to evaluate realistic output results.

Optimization Techniques

In order to optimize the architecture of embedded deployment, an adapted strategy will combine three related optimization strategies:

- **Quantization-Aware Training (QAT):** Both forward and backward passes through a network can be drawn towards imitating 8-bit integer operations using ONNX Runtime and PyTorch QAT pipeline. This primes the model to post-training quantization without sacrificing the accuracy and lowers memory footprint by approximately 40%.
- **Channel Pruning:** The offered technique implies adoption of a channel pruning technique (grounded on the L1-norm) that recognizes and eliminates superfluous channels amidst the convolutional layers. This is followed by a profile of the devices (i.e., Jetson Nano or Cortex-A53) that the target pruned networks will run on to verify that pruning choices are practical when they reach the hardware.^[9]
- **Knowledge Distillation:** A big ESRGAN model with high performance is passing its knowledge through training a student model. With distillation loss, the student learns to imitate high-level representation and close the performance difference without ever learning to be inefficient.

The overall optimization techniques reduce the computation overhead, and at the same time they do not result in a decrease in super-resolution quality therefore these techniques are viable in real-time implementation on low-powered platforms as shown in Figure 2.




Fig. 2: Optimization Techniques for Embedded Deployment

Here, three principal optimization strategies of the models, including Quantization-Aware Training, Channel Pruning, and Knowledge Distillation, to fit the GAN-based ISR model to low-energy embedded devices are presented.

Loss Functions

In the training objective function, the various losses should add up to allow the accuracy of the pixel as well as the perceptual realism:

- **Perceptual Loss ($L_{\text{perceptual}}$):** Computed using feature maps extracted from a pre-trained VGG19 network, this loss guides the generator to match high-level visual semantics with ground truth, improving texture and fine details [10].
- **Adversarial Loss (L_{adv}):** The Wasserstein GAN (WGAN) framework with gradient penalty is adopted for stable training and better convergence, particularly important for low-data regimes in satellite imagery.
- **Pixel-wise L1 Loss (L_{pixel}):** Wants the proof-of-concept pixel in perfect alignment to the ground truth pixel at the structural level, the fewer the intensity difference.

The last loss function which is to be optimised in the generator is defined as:

$$L_{\text{total}} = \lambda_1 L_{\text{pixel}} + \lambda_2 L_{\text{perceptual}} + \lambda_3 L_{\text{adv}}$$

where $\lambda_1, \lambda_2, \lambda_3$ are empirically tuned weighting coefficients.

Figure 3 shows this overall structure of the losses with the flow of information travelling through the generator created and ground truth images to the perceptual, adversarial and pixel based loss computations, which will in turn affect the updates of the generator network during training.




Fig. 3: Composite Loss Function Architecture for Lightweight GAN-Based ISR

The diagram explains the manner in which pixel-wise L1 loss and the perceptual loss (based on VGG19 features) along with the adversarial loss (using Wasserstein GAN and also a gradient penalty) can be combined into a combined training objective. The combination of loss directs the generator to output photographically believable and structurally valid high-resolution satellite images.

It is a combination of model compression, a hardware-aware optimization, and a multi-loss training objective that make it a deployment-ready ISR system. Compared to conventional GANs which frequently have a strong emphasis on visual quality as opposed to feasibility, our design balances these two factors, so it is applicable to carry out in-device enhancement of satellite images in bandwidth- and power-limited scenarios.

EXPERIMENTAL SETUP

To ensure that the solution in this paper outperforms state of the art in terms of performance, efficiency, and practicality of deployment, an extensive test bed was implemented using a wide range of datasets, metrics, and hardware platforms (Figure 5).

Datasets

The evaluation and training were done using two public datasets of satellite images:

- SpaceNet: A high-resolution database of suburban and urban environment that has been labeled with multiple categories of vision problems. It offers a multi-spectral and RGB imagery of 3050-cm spatial resolution, which can be used in benchmarking ISR performance on real geospatial material.

- UC Merced Land Use Dataset: UC Merced land use dataset consists of 2100 aerial scene images belonging to 21 classes of land use and having 256 pixel square shape. The dataset can be used to test generalization and evaluate cross-domain performance in settings that have varied types of scenes and resolutions.

Evaluation Metrics

In evaluation of the perceptual quality and signal fidelity, the following metrics have been used:

- Peak Signal-to-Noise Ratio (PSNR): It is used to measure pixel-wise performance of reconstruction.
- Structural Similarity Index (SSIM): Codes information on the image quality as regards to structure.
- Learned Perceptual Image Patch Similarity (LPIPS): Human-perceptual similarity, rather important when you want high-level fidelity.
- Inference Time (ms): Indicates how much the end-to-end processing takes on a per image basis, and it is of importance to real time apps.
- Energy Consumption (W): Watched through on-device power profiler tools that will confirm low-power deployment suitability.

Hardware Platforms

To approach this practically in terms of embedded deployment of the model, it was deployed in two of the more common low-power AI-capable hardware platforms:

- NVIDIA Jetson Nano (Quad-core ARM Cortex-A57 @ 1.43 GHz, 128-core, Maxwell GPU): GPU can accelerate via CUDA and TensorRT, hence this is ideal to perform GPU-accelerated benchmark inference at the edge.
- Raspberry Pi 4 (Quad-core Cortex-A72 @ 1.5 GHz) with Intel Neural Compute Stick 2 (VPU): This will give a heterogeneous inference platform using Intel OpenVINO toolkit to perform low-power inference on ARM based SBCs.

These models were tested with specific inference engines (TensorRT, OpenVINO) embodied on native hardware (GTX 1050, 1080) upon usage of int8 precision to verify the real-time processing performance on stricter power consumption (<10W). The memory and inferencing performance footprint of the model are shown in Figure 4 (where all tests were done on select embedded platforms), demonstrating the suitability of the model to ISR development both in terms of memory and in terms of inference latency and power consumption.

Table 1: Evaluation Metrics across Platforms

Platform	PSNR (dB)	SSIM	LPIPS	Inference Time (ms)	Power Consumption (W)
Jetson Nano	31.8	0.922	0.087	47.3	9.2
Raspberry Pi 4 + NCS2	30.6	0.908	0.094	62.1	6.8




Fig. 4: Deployment Footprint on Embedded Devices

Experimental Setup




Fig. 5: End-to-End Evaluation Framework for Lightweight ISR Deployment

Table 2. Quantitative Comparison of ISR Models

Method	PSNR (dB)	SSIM	Model Size (MB)	Inference Time (ms)	Power (W)
Bicubic	23.2	0.71	-	<10	<1
ESRGAN	28.7	0.84	155	235	12.5
MobileSR	27.1	0.80	38	75	5.1
Proposed	29.3	0.86	19	41	4.6

RESULTS AND DISCUSSION - SCOPUS-LEVEL EXPLANATION

The given lightweight GAN-based framework of the image super-resolution (ISR) functions was tested strictly on classical and state-of-the-art benchmarks, such as the Bicubic interpolation, ESRGAN, and MobileSR. The quantitative comparison on all perceptual quality measures, model complexity, latency, and power consumption is given in Table 2 with each model benchmarked on the NVIDIA Jetson Nano platform.

Performance Analysis

1. Image Quality Metrics

- The Proposed model produced a Peak Signal-to-Noise Ratio (PSNR) of 29.3 dB as compared to the other models, this shows that the proper reconstruction of the pixels was high.
- Measuring Structural Similarity Index (SSIM), the model scored 0.86, which means a superlative perceptual quality and improved edge/texture reconstruction compared to the outcomes of MobileSR (0.80) and ESRGAN (0.84).

2. Model Compactness

- Its model is only 19 MB in size, which is 87.7 percent less than ESRGAN and half compared to MobileSR. In memory-limited embedded and edge nodes, such a requirement exists.

3. Inference Speed and Real-Time Performance

- The model provides 41 ms per Image inference time, which is approximately 24 FPS; therefore, it confirms real-time capability on Jetson Nano.
- It reduces the latency by ~82.5 percent when compared to the ESRGAN, which makes it 5.7 times faster in edge situations.

4. Energy Efficiency

- The model has a power margin of 4.6W, which remains greatly under the 10W that are needed to deploy UAV and CubeSats. On the contrary, ESRGAN is above 12.5W which makes it unsuitable in such environments.

Visual comparisons of the super-resolved outputs were done across models whereas figure 6 compares the super-resolved outputs. The suggested model has more defined edges, a better texture, and less visual artifacts in the high-frequency areas (e.g. building edges, roads, vegetation).




Fig. 6: Qualitative Image Super-Resolution Comparison

Inputs to different methods visual comparison of the super-resolved output on noisy low-resolution satellite image. The suggested lightweight GAN-based ISR algorithm brings the sharp edges and structural definition back, which is better than Bicubic and ESRGAN methods and cannot be denied; moreover, the edge fidelity and the artifact density are better to see.

Discussion Highlights

- Precision-Deployment Dilemma: The deep models such as ESRGAN are precise but cannot be deployed. In the proposed framework we want to get an even better balance, between quality at a limited compute and energy budget.
- Deployment-Ready Optimization: The integrated synthesis of quantization-aware training, channel

pruning and knowledge distillation directly translates into model optimisation and hardware-specific acceleration which directly translates into support across edge-AI hardware at scale.

- Generalization, Robustness: Scholarly adjudication in relation to the cold/hot newbie dataset demonstrates how generalized it is to apply other satellite image sources.

The proposed model is very much computationally charged (inference time and power consumption) as well as in terms of perceptual accuracy (PSNR/SSIM) compared to its counterparts MobileSR and ESRGAN as seen in Figure 7.

Figure 8 visualizes the capability of explicitly generalizing to novel datasets by showing how well the presented model is expected to perform on other datasets, including SpaceNet and UC Merced domains with a broad variety of images in terms of resolution.




Fig. 7: Performance-Accuracy Trade-Off Curve




Fig. 8: Cross-Dataset Generalization Results

CONCLUSION AND FUTURE WORK

The proposed study portrays a low-end GAN-based image super-resolution (ISR) framework specifically optimized

through high-performance low-power and real-time compatible architectures after having been deployed on the low-power embedded hardware on-board of satellite and UAV-based Earth observation systems. The model that is proposed combines depthwise separable convolutions, channel pruning, and quantization-aware training to obtain a good trade-off between computational and perceptual image quality.

Experimental verification proved that the model has demonstrated:

- This will result in an increase of the PSNR of up to 3.5 dB over conventional interpolation 6 techniques
- High-visual quality and rich edge recovery and decreased artifacts,
- Compressed 19MB and had inference times less than 50ms/image and <5W power consumption on Jetson Nano.

Key Contributions:

- ISR system with integrated edge board single-FPGA-based AIGAN.
- Implementation of a combination of different optimization approaches (QAT, pruning, distillation) in order to be hardware-awarely implemented.
- Computation of cross-platform comparison on Jetson Nano and Raspberry pi + NCS2 with energy profiling.

FUTURE WORK:

- Multi-frame ISR that can be used to utilize the temporal redundancy and adjust the consistencies across frames.
- The addition of onboard compression-aware training to augment ISR with a more cost-effective transmission over limited bandwidth.
- Hyperspectral and SAR expansion to expand the scope of satellites to a wider range.

This system establishes a base to make ISR accessible and energy-efficient field operation, in which image enhancement can be done directly to high-quality satellite imagery on-device, an imperative attribute of autonomous observation systems and real-time geospatial intelligence products.

REFERENCES

1. Keys, R. (1981). Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 29(6), 1153-1160. <https://doi.org/10.1109/TASSP.1981.1163711>
2. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., ... & Shi, W. (2017). Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 4681-4690). <https://doi.org/10.1109/CVPR.2017.19>
3. Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., ... & Loy, C. C. (2018). ESRGAN: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European Conference on Computer Vision Workshops (ECCVW)* (pp. 63-79).
4. Howard, A., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., ... & Adam, H. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*. <https://arxiv.org/abs/1704.04861>
5. Han, S., Pool, J., Tran, J., & Dally, W. J. (2015). Learning both weights and connections for efficient neural network. In *Advances in Neural Information Processing Systems (NeurIPS)* (pp. 1135-1143).
6. Jacob, B., Kligys, S., Chen, B., Zhu, M., Tang, M., Howard, A., ... & Adam, H. (2018). Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 2704-2713).
7. Sadulla, S. (2025). IoT-enabled smart buildings: A sustainable approach for energy management. *National Journal of Electrical Electronics and Automation Technologies*, 1(1), 14-23.
8. Surendar, A. (2025). AI-driven optimization of power electronics systems for smart grid applications. *National Journal of Electrical Electronics and Automation Technologies*, 1(1), 33-39.
9. Khan, M., & Shah, F. (2025). Evidence Based Quality Improvement for Most Businesses Sustainability. *National Journal of Quality, Innovation, and Business Excellence*, 2(1), 11-22.
10. Saar, K., & Vaher, M. (2025). Real Data from 500 Hospitality Businesses When Hotels Doubled Revenue Post-COVID. *Journal of Tourism, Culture, and Management Studies*, 2(1), 13-26.
11. Muralidharan, J. (2024). Compact reconfigurable antenna with frequency and polarization agility for cognitive radio applications. *National Journal of RF Circuits and Wireless Systems*, 1(2), 16-26.