Developing an Image Processing Pipeline for DSLR Astrophotography Using Classical and Deep Learning Methods

Raina Song, and Yifei Deng

Abstract—Amateur astrophotography using DSLR cameras often suffers from star trails, non-uniform backgrounds, and high noise levels, resulting in lower image quality compared to professional setups. To address these challenges, we developed an automated image processing pipeline that compares classical signal processing methods with neural network-based approaches. The classical method includes background gradient removal, PSF motion blur kernel estimation, Richardson-Lucy deblurring, and non-local means denoising, while the neural network method employs a multi-stage model for denoising, background removal, and deblurring. Our experiments show that classical methods excel in precise background gradient extraction and PSF estimation, whereas neural networks generalize better across diverse conditions. This work provides a robust solution for enhancing amateur astrophotography, making it more accessible and achieving results closer to professional standards.

Index Terms—Astrophotography, Richardson-Lucy Deblurring, Neural Networks, Deep Learning

1 Introduction

Amateur astrophotography using DSLR cameras has grown in popularity due to its affordability and accessibility. However, capturing high-quality images of celestial objects presents significant challenges, including motion blur from camera movement, non-uniform background illumination caused by light pollution, and high noise levels due to low-light conditions. Unlike professional observatories, which employ high-precision tracking mounts, adaptive optics, and cooled CCD sensors, amateur setups rely on consumer-grade equipment that lacks these advanced features. As a result, DSLR astrophotographers often struggle with star trailing, poor contrast, and degraded image clarity.

In this project, we develop and compare two approaches for improving DSLR astrophotography: a classical image processing pipeline and a multistage neural network. The classical pipeline includes steps such as background gradient removal, Point Spread Function (PSF) motion blur kernel estimation, Richardson-Lucy deblurring, and non-local means denoising. The neural network approach, on the other hand, employs a multi-stage model to perform denoising, background removal, and deblurring in an end-to-end manner. By evaluating the strengths and limitations of each method, we aim to provide a robust solution that bridges the gap between amateur and professional astrophotography.

This work focuses on enhancing images of star fields and deep-sky objects, addressing common issues such as star trailing, non-uniform background, and noise. Our goal is to make high-quality astrophotography more accessible to enthusiasts, enabling them to achieve results closer to professional standards using affordable DSLR equipment. Through this project, we explore the potential of combining classical and neural network-based methods to create a more effective and realistic image processing pipeline.

Fig. 1. Example Star Images Captured by DSLR Cameras with Star Streaks, Background Gradient, and Noises

2 Related Work

Image restoration, particularly in astrophotography and low-light conditions, has seen significant advancements in both classical and deep-learning-based approaches.

Many researchers have explored classical methods for low-light and astrophotography image restoration. Wiener and Richardson-Lucy deconvolution can effectively remove light trails [1]. However, in astrophotography, variations in the blur kernel caused by exposure time, tripod motion, and atmospheric conditions make determining the Point Spread Function (PSF) challenging. Traditional methods often rely on blind deconvolution and PSF estimation to correct motion blur and enhance image clarity. Hu et al. introduced a non-linear blur model that explicitly utilizes light streaks as constraints for estimating motion blur kernels [2]. They used a power-spectrum-based metric for selecting the best light streak, approximating the blur kernel's power spectrum directly from the input image. This approach proved effective where strong light streaks are present,

providing rich blur information. However, the method struggles with scenes lacking sufficient light streaks and can be adversely affected by saturation artifacts. Expanding on this idea, Su et al. incorporated a thinning algorithm to estimate motion blur in astronomical star images [3]. Their method refines the Richardson-Lucy deblurring approach by extracting more accurate motion blur kernels based on the trajectory of blurred stars, enabling better restoration of astrophotography images. Although effective for linear motion blur, their approach faces difficulties with complex rotational blur and spatially varying distortions common in astrophotography.

To address similar issues, other fields have used ADMM-based blind deconvolution and denoising techniques from low-photon fluorescence imaging [4]. Background removal in astrophotography has been approached through various techniques, including background subtraction, wavelet-based denoising, and filtering. For instance, [5] combines Gaussian filtering, background subtraction, histogram equalization, and Otsu's thresholding to address non-uniform backgrounds in astronomical images.

Recently, neural network-based approaches have gained attention in low-light and astrophotography image enhancement due to the availability of larger datasets and advancements in deep learning. [6] introduced MPRNet, a multi-stage progressive image restoration network designed for denoising, deblurring, and dehazing. Each stage refines an intermediate output passed to subsequent stages, effectively managing the complexity of restoration tasks. Building upon MPRNet, Restormer [7] further improved image restoration by employing Transformer blocks, enabling better global context capturing. Restormer excels in multiple tasks such as deblurring, motion deblurring, denoising, and deraining, demonstrating superior performance and flexibility. Nevertheless, these general-purpose methods are not specifically tailored for astrophotography, which has unique sparsity and noise characteristics. To address astrophotography-specific challenges, [8] proposed a deep learning approach integrating band-dependent PSF information into the deconvolution process, applying Wiener deconvolution separately to each color band. For motion blur correction, [9] estimated motion blur kernel parameters using hyper-Laplacian priors and an ensemble neural network, enabling robust deblurring of star images.

Although these approaches show promising results, there is no dedicated imaging pipeline designed specifically for DSLR astrophotography, particularly at the amateur level. Developing such a pipeline could make astrophotography more accessible to a broader audience.

3 Proposed Method

To address the issues DSLR cameras face when capturing astrophotography, we proposed two sets of approaches: the classical approach and the deep learning-based approach, and doing comparison between them.

3.1 Classical Approach

Our image restoration pipeline consists of several key steps aimed at improving the quality of astrophotography images by correcting motion blur, reducing noise, and enhancing contrast while preserving star details. Each step is carefully designed to address the specific challenges posed by DSLR astrophotography, particularly star trailing, background inconsistencies, and sensor noise. The workflow includes background gradient removal, PSF estimation, Richardson-Lucy deblurring, non-local means denoising, and selective star enhancement.

3.1.1 Background Gradient Removal

Firstly, astrophotography images often exhibit non-uniform background illumination due to light pollution, sensor inconsistencies, or natural sky brightness variations. Direct contrast enhancement without handling these gradients can lead to loss of faint celestial objects and inaccurate star reconstruction. To mitigate this, we use morphological opening with an elliptical kernel to estimate and subtract the large-scale background gradient while preserving star details. That is, given an input image I and a structuring element (kernel) B, morphological opening is defined as:

$$I \circ B = (I \ominus B) \oplus B$$

where: $\ominus$ represents erosion (shrinking objects by removing pixels from boundaries). $\oplus$ represents dilation (expanding the eroded image back).

3.1.2 PSF Blur Kernel Estimation

Then, to correct star trailing and motion blur, we must first estimate the Point Spread Function (PSF), which represents the blurring pattern caused by the Earth's rotation or camera movement during long exposures. Since the PSF varies across images, we dynamically estimate it from the brightest star in the scene.

3.1.3 Richardson-Lucy Deblurring

Once the PSF is estimated, we apply the Richardson-Lucy (RL) deconvolution algorithm to recover sharp star details. RL deconvolution is based on Bayesian probability and Maximum Likelihood Estimation (MLE). It uses an iterative approach to refine an estimate of the sharp image I, given the blurred image B and the known PSF. In each iteration, we do the following steps:
(1) Estimate the Blurred Image

$$B' = I' * PSF$$

where I' is the current estimate of the sharp image.

(2) Compute the Ratio of Actual Blurred Image to Estimated Blurred Image:

$$R = \frac{B}{B' + \epsilon}$$

where ϵ is a small constant to prevent division by zero.
(3) Apply the Correction:

$$I'_{\text{new}} = I' \times (R * PSF_{\text{flipped}})$$

This updates I' using the flipped PSF to refine the estimate.
(4) Iterative Refinement: this process is repeated for multiple iterations until the image is sufficiently sharpened.

3.1.4 Non-local Means Denoising

Even after deblurring, the image may contain high-frequency noise from the camera sensor, low-light conditions, or deconvolution artifacts. To mitigate this, we apply Non-Local Means (NLM) denoising, which is effective at preserving fine details while suppressing noise.

3.1.5 Selective Star Enhancement

While deblurring restores structural details, stars may appear less prominent due to overall contrast adjustments. To address this, we apply a selective star enhancement technique using a binary star mask. That is, the star mask is 1 for bright stars and is 0 for background. So we keep the deblurred foreground and the denoised background by blending them smoothly.

3.2 Deep Learning-Based Approach

While the classical approach is effective in many scenarios, it faces several limitations. Classical denoising techniques such as Non-Local Means (NLM) struggle to differentiate between fine celestial details and sensor noise, often leading to excessive smoothing and loss of faint objects. Background removal techniques that rely on morphological operations can introduce overcorrections, particularly in images containing extended nebulae or diffuse astronomical structures. Motion blur correction using deconvolution-based methods, such as Richardson-Lucy deconvolution, requires an accurately estimated Point Spread Function (PSF), which is difficult to determine due to variations in optical distortions. Moreover, the classical approach often require manual parameter tuning, making it less adaptable to diverse astrophotographic conditions, and often could take long computational time.

To overcome these limitations, we propose **StarEnhancementNet**, a deep learning-based multi-stage pipeline designed for progressive astrophotography image restoration. The network consists of four key stages: denoising, background removal, star streak correction, and end-to-end fine-tuning. Each stage is trained independently with task-specific objectives before a joint optimization step ensures optimal restoration performance.

StarEnhancementNet is inspired by MPRNet [10], and is designed as a progressive pipeline where each stage is responsible for a specific restoration task. The first three stages operate as independent modules, handling denoising, background removal, and structured artifact suppression. The final stage integrates these modules through end-to-end optimization, allowing the network to learn joint feature representations for improved restoration quality.

3.3 Stage 1: Denoising via DnCNN

The first stage removes sensor noise and low-light artifacts using **DnCNN**, a deep residual convolutional network optimized for blind Gaussian denoising. Unlike direct denoising models, DnCNN learns to predict the noise component, subsequently subtracting it from the input image. Given an input image X, the noise component $\hat{N}$ is estimated as:

$$\hat{N} = f(X; \theta) \tag{1}$$

where f represents the DnCNN network parameterized by weights θ. The denoised image $\hat{I}$ is reconstructed as:

$$\hat{I} = X - \hat{N} \tag{2}$$

The DnCNN architecture consists of:

- **Input convolutional layer**: Extracts initial low-level features.
- **15 intermediate convolutional layers**: Each followed by batch normalization and ReLU activations for deep feature learning.
- **Final convolutional layer**: Outputs the noise residual map to subtract from the input image.

Training utilizes a combined L1 loss with additional brightness and color consistency constraints, defined as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_1 + \alpha \cdot \mathcal{L}_{\text{brightness}} + \beta \cdot \mathcal{L}_{\text{color}} \tag{3}$$

where α and β can be tuned, ensuring accurate denoising while preserving fine astronomical features.

3.4 Stage 2: Background Removal via CNN

The second stage employs a convolutional neural network trained to adaptively remove large-scale background illumination gradients without manual parameter tuning. Unlike traditional handcrafted methods, this CNN directly learns background gradient patterns from paired datasets.

The architecture comprises:

- **Sequential convolutional layers** (3×3 kernels), each followed by ReLU activations, designed to learn and subtract background illumination patterns adaptively.

Training data for Stage 2 directly uses the denoised output from Stage 1 as input. The loss function remains consistent (L1 loss combined with brightness and color consistency) to ensure preservation of astronomical structures.

3.5 Stage 3: Star Streak Correction via Simplified MPRNet

Stage 3 addresses star streak artifacts resulting from camera movement or Earth rotation using a simplified **MPRNet**:

- **Encoder**: Two convolutional layers capture local and global multi-scale features, compressing information into deeper feature maps.
- **Decoder**: Upsampling layers restore spatial resolution, incorporating skip connections to retain essential spatial features.
- **Output layer**: Final 3×3 convolution reconstructs the corrected image, eliminating structured blur artifacts.

Stage 3 training uses synthetic training data with star streaks generated by applying randomized linear and curved motion blur kernels to background-removed images from Stage 2.

Fig. 2. Classical Approach Image Processing Pipeline

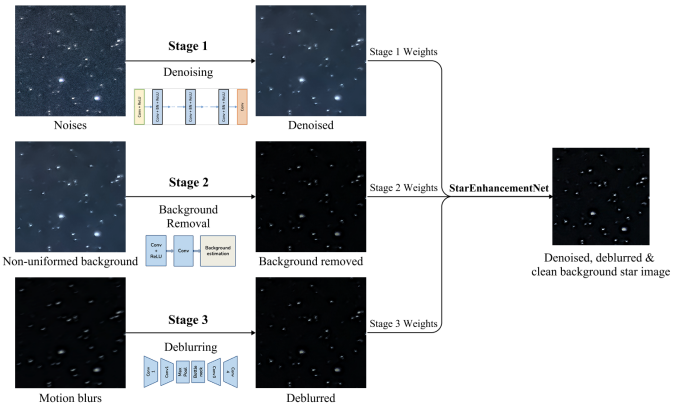

Fig. 3. Deep Learning-Based Image Approach Processing Pipeline and Model Architectures

3.6 Stage 4: End-to-End Fine-Tuning

The final stage jointly optimizes the pipeline by fine-tuning all previous stages simultaneously. It employs the same combined loss function (L1 loss with brightness and color consistency) to maintain training consistency across the pipeline. Optimization uses the Adam optimizer with a constant learning rate of 1×10^{-4}. This fine-tuning stage ensures seamless integration and enhanced generalization capability.

3.7 Training and Data Generation

Training datasets are explicitly generated for each stage as follows:

- **Stage 1 (Denoising)**: Randomly extracted 256×256 patches from raw astrophotography images with no star streaks but with noises and non-uniformed background, mildly denoised using the non-local means algorithm.
- **Stage 2 (Background Removal)**: Background removal performed using Gaussian blur (31×31 kernel) applied to Stage 1 outputs, followed by partial subtraction:

$$I' = \text{clip}(I - \alpha \cdot G_{\sigma=31}(I), 0, 255)$$

  where $\alpha = 0.9$, controlling background retention.
- **Stage 3 (Deblurring)**: Synthetic star streak artifacts created using linear and curved motion blur kernels with randomized angles and curvature parameters, applied to Stage 2 background-removed images:

$$I_{\text{streaked}} = I_{\text{Stage2}} * K_{\text{motion-blur}}$$

- **Stage 4 (All-in-One Degradation)**: Composite degradation applied simultaneously (random Gaussian noise $\sigma \in [5, 25]$, partial background reintroduction with Gaussian blur, and synthetic star streaks) to Stage 2 images:

$$I_{\text{final}} = \text{Streak}\big(\text{PartialBG}\big(\text{Noise}(I_{\text{Stage2}})\big)\big)$$

The entire dataset consists of around 200 pairs of images for training and around 80 pairs of images for validation.

All networks within the pipeline were trained on a system running Ubuntu Linux, equipped with an NVIDIA GeForce RTX 3080 GPU. Training utilized PyTorch with CUDA acceleration for efficient computation.

Each pipeline stage was independently trained as follows:

- **Stage 1 (DnCNN Denoising)**: 800 epochs with a batch size of 8, learning rate of 1×10^{-4}.
- **Stage 2 (Background Removal CNN)**: 800 epochs, batch size of 8, learning rate of 1×10^{-4}.
- **Stage 3 (Simplified MPRNet Deblurring)**: 1000 epochs, batch size of 8, learning rate of 1×10^{-5}.
- **Stage 4 (End-to-End Fine-Tuning)**: 300 epochs, batch size of 8, learning rate of 1×10^{-4}.

The Adam optimizer was consistently employed across all stages to optimize the models' parameters. The training performance was monitored through validation loss and PSNR metrics computed after each epoch to detect convergence and potential overfitting.

Model checkpoints were periodically saved (every 50 epochs and at training completion) to enable monitoring of intermediate restoration performance and facilitate fine-tuning. Also, visual validation examples were generated at intervals to qualitatively assess improvement over the training epochs. Additionally, all training images were normalized to a [0,1] range to ensure numerical stability and faster convergence.

4 Experiments & Results

To evaluate the effectiveness of our proposed methods, we conducted a series of experiments using both single blurred astrophotography images (without ground truth)

and paired datasets captured specifically for quantitative analysis. We compare our methods against established classical and deep learning-based methods, namely Hu et al. [2] and Restormer [7], respectively.

We first ran experiments using single blurred images without corresponding ground truth data. These blurred images were processed through our proposed classical pipeline and the deep learning-based StarEnhancementNet. Results from our methods were then visually compared side-by-side against outputs from Hu et al.'s classical approach and the Restormer neural network. This initial qualitative assessment allowed us to observe differences in deblurring capabilities, noise reduction, and background uniformity among the methods.

Subsequently, to enable quantitative evaluation, we captured our own ground truth and blurred image pairs using a Canon 90D DSLR camera, equipped with an APS-C CMOS sensor (32.5 Megapixels, sensor size: 22.3 × 14.9 mm). The images were captured with exposure times ranging between 10 to 20 seconds and ISO settings between 400 to 800. These pairs provided a reliable dataset for calculating image quality metrics.

4.1 Evaluation Metrics

To objectively evaluate the restoration quality, we employed two standard image quality metrics: Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM).

Peak Signal-to-Noise Ratio (PSNR) quantifies the pixel-level reconstruction accuracy between the restored image and the ground truth. PSNR is defined mathematically as:

$$\text{PSNR} = 10 \cdot \log_{10} \left(\frac{MAX^2}{MSE} \right) \tag{4}$$

where MAX represents the maximum possible pixel intensity value of the image (255 for 8-bit images), and MSE (Mean Squared Error) is computed as:

$$MSE = \frac{1}{mn} \sum_{i=1}^{m} \sum_{j=1}^{n} [I(i,j) - K(i,j)]^2 \tag{5}$$

Here, I is the restored image, K is the ground-truth image, and m, n denote image dimensions. Higher PSNR values indicate better restoration performance, typically above 30 dB for visually pleasing results.

Structural Similarity Index Measure (SSIM) evaluates the perceptual similarity between two images by comparing luminance, contrast, and structural information. SSIM is calculated as:

$$\text{SSIM}(I, K) = \frac{(2\mu_I \mu_K + C_1)(2\sigma_{IK} + C_2)}{(\mu_I^2 + \mu_K^2 + C_1)(\sigma_I^2 + \sigma_K^2 + C_2)} \tag{6}$$

where:

- μ_I, μ_K are the mean pixel values of images I and K, respectively.
- σ_I^2, σ_K^2 denote the variances of images I and K.
- σ_{IK} is the covariance between I and K.
- $C_1 = (k_1 L)^2$, $C_2 = (k_2 L)^2$ are stabilization constants, with L being the dynamic range of pixel values (255 for 8-bit images), and $k_1 = 0.01, k_2 = 0.03$ are constants.

SSIM values range between -1 and 1, with values closer to 1 indicating higher perceptual similarity and superior visual quality.

4.2 Results Analysis

The experimental results are presented in Figures 5, 6, and 4. Quantitative metrics (PSNR and SSIM) are indicated below each corresponding restored image in Figures 5, 6.

Upon analyzing the quantitative metrics, our proposed methods achieved relatively high PSNR and SSIM scores comparable to those obtained by Restormer, which is a complex, transformer-based deep neural network trained on more than 5,000 image samples. Notably, our deep learning-based method was trained on significantly fewer images (approximately 200), yet achieved competitive results. Specifically, our deep learning method resulted in PSNR values around 30 dB, demonstrating its ability to accurately reconstruct sharp celestial details and remove star streak artifacts.

A detailed visual analysis supports these quantitative findings. In the qualitative comparisons, our deep learning-based approach effectively removed motion-induced star streaks, suppressed sensor noise, and corrected non-uniform background illumination. Restormer, despite its extensive training dataset, demonstrated limitations in adequately correcting motion blur in astronomical images, leaving residual streaks visible. Furthermore, our classical method successfully performed both deblurring and denoising tasks, yielding visually clear and sharp star fields when compared with the Hu et al.'s [2] classical approach. However, our deep learning-based method occasionally oversmoothed very faint stars, resulting in a loss of visibility for smaller celestial objects in the final output.

Moreover, we evaluated and compared the computational efficiency of both the classical and deep learning-based restoration methods. For an astrophotography image of resolution $6960 \times 4640 \times 3$ (approximately 39MB JPEG file size), the classical approach pipeline requires approximately 90–100 seconds to process a single image. In contrast, our proposed deep learning-based method, StarEnhancementNet, significantly reduces the processing time, taking approximately 50 seconds per image on an NVIDIA GeForce RTX 3080 GPU with CUDA acceleration.

This computational improvement highlights the efficiency of the deep learning approach, making it preferable for large-scale astrophotographic restoration tasks where processing time is critical.

In summary, our experimental evaluation demonstrates the effectiveness of both the classical and the proposed deep learning-based approaches presented in this work. The classical pipeline utilizes traditional image processing methods, offering competitive restoration performance. Our proposed deep learning-based method, StarEnhancementNet, achieves restoration performance comparable to state-of-the-art transformer-based architectures despite being trained on a significantly smaller dataset. This highlights its practical potential and computational efficiency, making it suitable for real-world astrophotographic image restoration tasks.

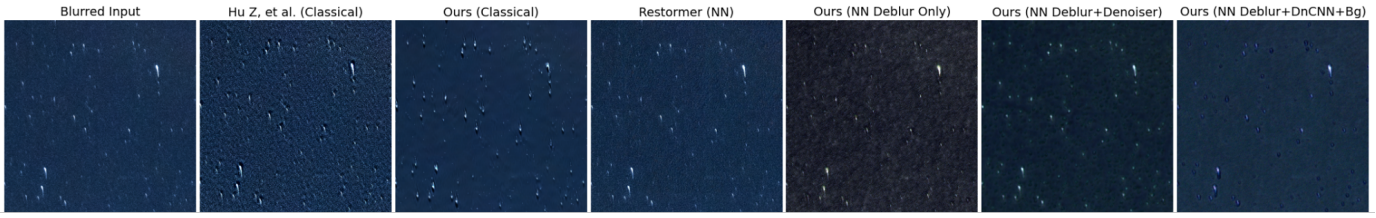

Fig. 4. Qualitative comparison of our proposed classical and deep learning-based methods against Hu et al. [2] and Restormer [7]

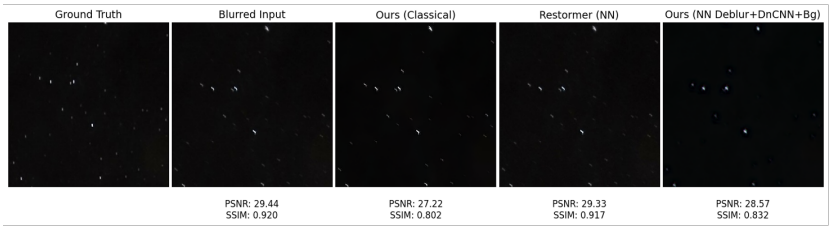

Fig. 5. Visual and quantitative comparisons of our proposed methods with Restormer [7] and our classical approach on real image pairs.

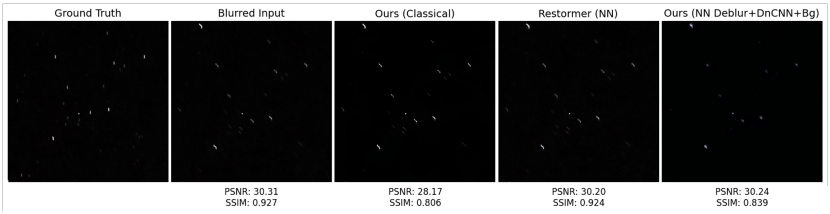

Fig. 6. Visual and quantitative comparisons of our proposed methods with Restormer [7] and our classical approach on real image pairs.

5 Discussion & Conclusion

In this project, we developed two automated image processing pipelines to address key challenges in DSLR astrophotography, such as star streaks (motion blur), non-uniform backgrounds, and high noise levels. By exploring both classical and neural network-based methods, we evaluated their strengths, weaknesses, and overall effectiveness in enhancing astrophotography images.

Classical image restoration techniques demonstrated strong performance in tasks such as background gradient extraction and Point Spread Function (PSF) kernel estimation. These methods rely on well-established mathematical models and signal processing techniques, making them highly effective at accurately identifying motion blur patterns and compensating for light pollution effects. However, their computational complexity, particularly for high-resolution images, remains a significant drawback. Processing a single image using classical deblurring algorithms can take several minutes, making them impractical for real-time or large-scale applications.

In contrast, deep learning-based neural network approaches offered greater efficiency and generalization across diverse astrophotographic conditions. Our experiments showed that neural networks achieved higher PSNR and SSIM scores, indicating superior image clarity and noise reduction. Additionally, neural networks proved more adaptable to varying imaging conditions, requiring minimal parameter tuning to perform well across a wide range of inputs. However, we observed that neural networks sometimes over-deblurred images, inadvertently removing small stars and faint celestial details, which reduced the realism of the final output.

To address these limitations, we propose a hybrid approach that combines the strengths of both methods. By integrating precise gradient extraction and PSF estimation from classical processing as priors for the neural network, we aim to refine the model's ability to differentiate between noise and fine astrophotographic details. This hybrid approach has the potential to preserve faint celestial structures while enhancing overall image clarity.

Despite these improvements, our current implementation has several limitations. The neural network was trained on a relatively small dataset of 200 paird of images, potentially restricting its robustness and generalization. Expanding the training dataset with a larger and more diverse collection of real astrophotography image pairs will be a critical next step. Furthermore, our pipeline currently processes JPEG images, which inherently contain compression artifacts and reduced dynamic range. Transitioning to RAW image processing will allow us to leverage higher bit-depth information and richer sensor data, enabling more precise restoration and enhanced professional image quality.

In conclusion, this project demonstrates the potential of combining classical and neural network-based approaches to achieve efficient, high-quality astrophotographic image restoration. By leveraging the precision of classical methods for PSF estimation and background normalization alongside the adaptability of deep neural networks, we can develop a more effective and realistic image processing pipeline. Future work will include training the neural network on

a larger dataset, incorporating RAW image processing for greater detail preservation, and optimizing the hybrid approach to balance computational efficiency, image realism, and restoration accuracy. These advancements will further empower amateur astrophotographers to capture and restore professional-quality astrophotography images using accessible DSLR equipment.

REFERENCES

[1] W. Wallace, L. H. Schaefer, and J. R. Swedlow, "A workingperson's guide to deconvolution in light microscopy," *Biotechniques*, vol. 31, no. 5, pp. 1076–1082, 2001.

[2] Z. Hu, S. Cho, J. Wang, and M.-H. Yang, "Deblurring low-light images with light streaks," in *2014 IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 3382–3389.

[3] L. Su, X. Shao, L. Wang, H. Wang, and Y. Huang, "Richardson-lucy deblurring for the star scene under a thinning motion path," in *Commercial + Scientific Sensing and Imaging*, 2015. [Online]. Available: https://api.semanticscholar.org/CorpusID:119081822

[4] H. Ikoma, M. Broxton, T. Kudo, and G. Wetzstein, "A convex 3d deconvolution algorithm for low photon count fluorescence imaging," *Scientific Reports*, vol. 8, no. 1, p. 11489, 2018.

[5] C. Zheng, J. Pulido, P. Thorman, and B. Hamann, "An improved method for object detection in astronomical images," *Monthly Notices of the Royal Astronomical Society*, vol. 451, no. 4, pp. 4445–4459, 2015.

[6] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, M.-H. Yang, and L. Shao, "Multi-stage progressive image restoration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 14 821–14 831.

[7] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, and M.-H. Yang, "Restormer: Efficient transformer for high-resolution image restoration," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5728–5739.

[8] H. Wang, S. Sreejith, Y. Lin, N. S. Ramachandra, A. Slosar, and S. Yoo, "Deconvolution of astronomical images with deep neural networks," in *NeurIPS 2022 AI for Science: Progress and Promises*, 2022.

[9] X. Chen, D. Liu, Y. Zhang, X. Liu, Y. Xu, and C. Shi, "Robust motion blur kernel parameter estimation for star image deblurring," *Optik*, vol. 230, p. 166288, 2021.

[10] S. W. Zamir, A. Arora, S. H. Khan, M. Hayat, F. S. Khan, M. Yang, and L. Shao, "Multi-stage progressive image restoration," *CoRR*, vol. abs/2102.02808, 2021. [Online]. Available: https://arxiv.org/abs/2102.02808