Denoising Pipeline for Deep-Sky Astrophotography

David Regordosa

AstroPics Lab

Abstract

This work presents a specialized denoising pipeline developed for *AstroPics Lab*, a software platform for advanced astrophotography post-processing. The system operates on paired noisy/clean 256×256 image tiles derived from community-provided deep-sky datasets and is enhanced through physically motivated synthetic noise modeling. A ResUNet architecture is trained using a multi-term loss that incorporates star-aware weighting, structural similarity, brightness consistency in non-stellar regions, and regularization in high-intensity zones. During inference, full-resolution images are processed using overlapping tile prediction with Gaussian feathering to eliminate seams. The resulting method achieves high-fidelity denoising while preserving the fine structure of celestial objects such as galaxies and emission or reflection regions, as well as maintaining the integrity of stellar profiles.

Keywords: astrophotography, denoising, deep learning, ResUNet, stars, deep-sky objects

1. Introduction

Astrophotographic data presents unique challenges compared to natural images: faint signals, high dynamic range, sensor-induced artifacts, and the requirement to preserve small-scale structures such as star fields and deep-sky objects. Traditional denoising models trained on generic image datasets often oversmooth these features or introduce distortions in high-intensity regions.

To address these limitations, we designed a domain-specific denoising module for *AstroPics Lab*, leveraging real user-generated astrophotography data from the AstroAnoia community. Each contributor provided both processed (clean) and lightly processed (noisy) captures, enabling the construction of a dataset of aligned noisy–clean patches suitable for supervised learning. Tiles are extracted at 256×256 resolution to balance detail preservation and computational efficiency.

A critical aspect of astrophotography denoising is the preservation of stars and the subtle structures of various astronomical objects, including galaxies, star-forming regions, and faint diffuse light. To this end, the training pipeline incorporates explicit star masks and custom losses that treat stellar and non-stellar regions differently. Furthermore, synthetic noise composed of realistic astrophotographic degradations—such as chroma blotching, banding, and Gaussian sensor noise—is injected during training to improve robustness under diverse imaging conditions.

For inference, a dedicated overlapping-tile algorithm processes full-resolution images with Gaussian feathering to avoid patch boundaries. This allows the model to operate seamlessly on images of arbitrary dimensions while maintaining prediction consistency across the frame.

2. Methods

2.1 Dataset Construction

High-resolution astrophotography images were sourced from the AstroAnoia community. For each scene, contributors supplied:

- a **noisy** version produced with minimal or no denoising;
- a **clean** version processed by the contributor.

All images were segmented into 256×256 non-overlapping tiles, generating a large corpus of paired examples suitable for supervised denoising.

2.2 Synthetic Noise Augmentation

To increase generalization and simulate real-world noise profiles, the training pipeline introduces synthetic perturbations on top of clean tiles. These perturbations include:

- **Random-variance Gaussian noise** applied independently per pixel;
- **Chroma block noise**, where low-resolution color fluctuations are upsampled to tile size, mimicking color blotches after extreme stretching;
- **Banding artifacts** (horizontal or vertical fixed-pattern variations), simulating sensor readout or calibration issues;
- **Identity and low-noise variants**, where the network is explicitly trained to leave already clean tiles unchanged.

This augmentation strategy ensures that the model encounters a wide range of astrophotographic noise conditions, improving robustness across different cameras, sensors, and acquisition pipelines.

2.3 Star Mask Generation

A luminance-based algorithm produces a soft star mask for every clean tile. The procedure is as follows:

- Convert the RGB tile to grayscale luminance;
- Apply a threshold to detect bright star cores;
- Dilate the result to include immediate halos around stars;
- Apply Gaussian blurring to create smooth, feathered transitions.

The resulting mask is stored as a single-channel float map in [0,1] and concatenated to the target RGB tile, forming a 4-channel ground truth:

$$y_{\mathrm{true}} = [\mathrm{clean}_{\mathrm{RGB}}, \mathrm{star\_mask}]$$

This representation enables the loss function to distinguish between stellar and non-stellar regions at training time.

2.4 Model Architecture

A tailored ResUNet architecture is employed for denoising. The network follows an encoder–bottleneck–decoder structure with residual blocks:

```
            Input (256×256×3)
                   │
                   ▼
          ┌─────────────────────┐
          │ Initial Conv 64     │
          └─────────────────────┘
                   │
                   ▼
          ┌─────────────────────┐
          │ Down Block 1        │
          │  - MaxPool          │
          │  - Residual Block 128│
          └─────────────────────┘
                   │
                   ▼
          ┌─────────────────────┐
          │ Down Block 2        │
          │  - MaxPool          │
          │  - Residual Block 256│
          └─────────────────────┘
                   │
                   ▼
```

```
┌──────────────────────┐
│ Bottleneck           │
│ - Residual Block 512 │
└──────────────────────┘
            │
            ▼
┌──────────────────────┐
│ Up Block 1           │
│  - UpSampling        │
│  - Concatenate skip3 │
│  - Residual Block 256│
└──────────────────────┘
            │
            ▼
┌──────────────────────┐
│ Up Block 2           │
│  - UpSampling        │
│  - Concatenate skip2 │
│  - Residual Block 128│
└──────────────────────┘
            │
            ▼
┌──────────────────────┐
│ Up Block 3           │
│  - UpSampling        │
│  - Concatenate skip1 │
│  - Residual Block 64 │
└──────────────────────┘
            │
            ▼
┌───────────────────────────┐
│ Output Conv (3 channels)  │
│ Activation: Sigmoid       │
└───────────────────────────┘
            │
            ▼
  Denoised Image (256×256×3)
```

Figure 1. Schematic diagram of the ResUNet architecture used in AstroPics Lab for denoising.

Residual connections improve gradient flow and help the network preserve fine astronomical structures, while skip connections allow the decoder to recover high-frequency details that might otherwise be lost in downsampling.

2.5 Loss Function

Training uses a composite loss tailored for astrophotography:

- **Star-weighted L1 loss**: a pixel-wise L1 term with reduced weight in stellar regions, preventing deformation of bright point sources;
- **SSIM loss**: encourages structural similarity between prediction and target, preserving detail and contrast;
- **Brightness consistency loss (object regions)**: applies a luminance constraint only to non-stellar regions, ensuring that deep-sky objects such as galaxies or emission regions are neither artificially darkened nor brightened;
- **Total variation regularization (star regions)**: applied inside stellar regions to smooth artifacts and reduce ringing around bright cores.

The combined loss can be expressed as:

$$\text{Loss} = (1 - \beta) \cdot \text{L1}_{\text{weighted}} + \beta \cdot \text{SSIM} + \gamma \cdot \text{Brightness} + \delta \cdot \text{TV}$$

2.6 Training Procedure

The model is trained on 256×256 tiles with a typical batch size of 8. The dataset is split into approximately 80% training and 20% validation subsets. The Adam optimizer with a learning rate of 1×10^{-4} is employed, alongside learning-rate scheduling and real-time visualization callbacks. Synthetic noise, star masks, and paired tile loading are integrated into the `tf.data` pipeline using `tf.py_function`, combining flexibility with efficient prefetching.

3. Inference on Full-Resolution Images

Inference on full-resolution astrophotography frames is performed using an overlapping tile strategy. Let *size* denote the tile size (256) and *overlap_fraction* the proportion of overlap (typically 0.2–0.3). The stride *step* and overlap are computed as:

```
step = size · (1 − overlap_fraction),    overlap = size − step.
```

This guarantees that each pixel in the original image is covered by multiple overlapping tiles.

Tile Extraction. The image is scanned in a grid pattern, extracting overlapping 256×256 patches. For edge regions, tiles are padded as needed so that every part of the image is processed.

Gaussian Weighting. For each tile, a two-dimensional Gaussian mask is generated, with maximum weight at the tile center and decreasing weights towards the edges. After the denoising network processes a tile, its output is multiplied by this mask.

Accumulation and Normalization. Two full-sized buffers are maintained: an *accumulator* for the sum of weighted predictions, and a corresponding *weight map* for the sum of Gaussian weights at each pixel. The final denoised image is obtained by:

```
final(x, y) = accumulator(x, y) / weight\_map(x, y).
```

This approach yields a seamless synthesis, eliminating visible seams and ensuring consistent predictions across the field.

4. Results

4.1 Visual Quality Improvements

The proposed denoising framework was evaluated qualitatively on a wide range of deep-sky images, including galaxies, star clusters, and emission or reflection regions. In all tested cases, the model successfully removed chroma blotching, banding patterns, sensor-induced noise, and small-scale random grain while preserving:

- stellar profiles and color integrity;
- fine structures in faint astronomical objects;
- global contrast and brightness relationships.

A side-by-side comparison between noisy inputs and denoised outputs reveals a clear reduction in random noise, especially in uniform areas such as backgrounds and low-signal regions. Furthermore, the method maintains small gradients and faint structures that are essential in astrophotography, avoiding the oversmoothing common in generic denoisers.

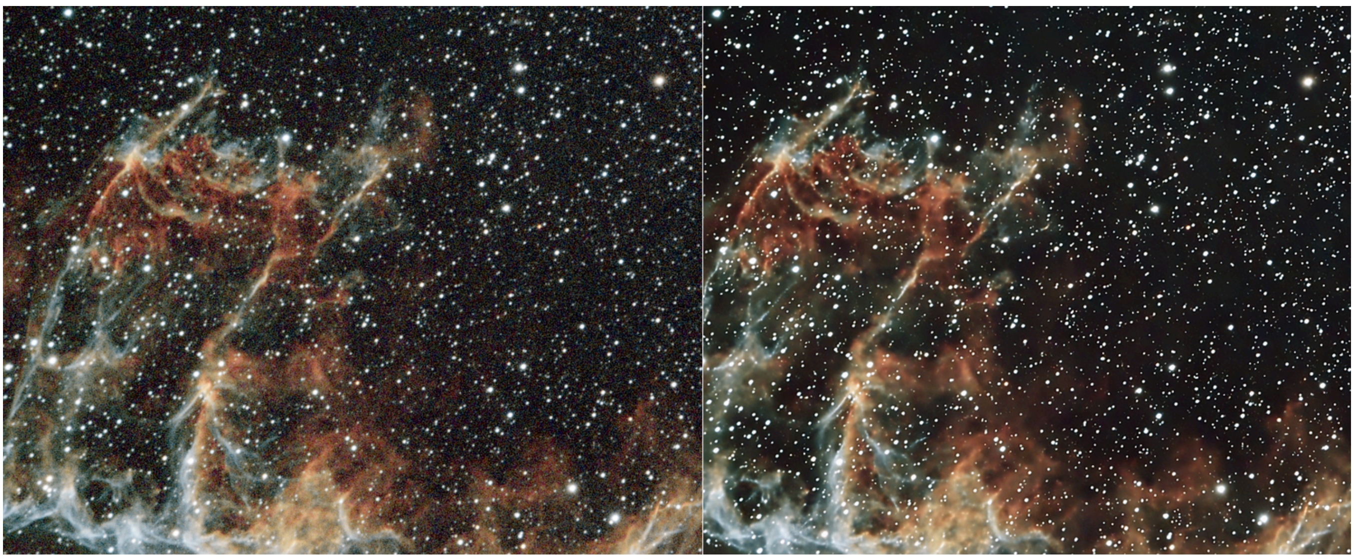

Figure 2. Comparison of noisy input and denoised output.

4.2 Tiling Strategies: Overlap + Gaussian Weighting vs. Naïve Non-Overlapping Tiles

Two inference strategies were compared:

Naïve tiling (no overlap). In this strategy, the image is divided into 256×256 blocks, each processed independently by the network and then placed back without blending. This approach leads to:

- visible block boundaries, especially in smooth backgrounds;
- regional inconsistencies due to local prediction biases in each tile;
- grid-like artifacts that are clearly perceptible in the final image.

Overlap + Gaussian weighting (AstroPics Lab default). With overlapping tiles and Gaussian feathering:

- no visible tile boundaries are present in the final image;
- predictions are consistent across tile borders;
- fine structures that cross tile boundaries are reconstructed smoothly.

When comparing both approaches on the same frame, the non-overlapping variant exhibits obvious tiling artifacts, while the overlapping method produces a uniform and coherent image. This confirms that the inference strategy plays a key role in achieving high-quality denoising results.

Figure 3. Comparison between noisy image (left), non-overlapping tiling (middle), and overlapping tiles with Gaussian weighting (right).

5. Conclusion

The denoising module developed for *AstroPics Lab* demonstrates strong performance in removing complex astrophotographic noise while preserving stellar integrity and fine object structure. The combination of real paired data from diverse imaging setups, synthetic astrophotography-inspired noise augmentation, a tailored ResUNet architecture, star-aware and brightness-preserving loss functions, and overlapping-tile inference with Gaussian weighting results in a robust denoising system suitable for galaxies, clusters, and a wide variety of deep-sky targets.

Although current results are highly encouraging, further improvements are planned. In particular, the training dataset will be expanded to incorporate images from additional telescopes, sensors, and optical configurations. Each camera model introduces characteristic noise patterns, and increasing dataset diversity will allow the system to generalize across the broad landscape of amateur and semi-professional astrophotography equipment.

Future work will explore additional domain-specific augmentations, adaptive inference strategies, and extended architectures. AstroPics Lab will continue to evolve with the goal of producing cleaner, more faithful, and scientifically meaningful deep-sky images while keeping star fields and object morphology intact.