

Denoising the Deep Sky: Physics-Based CCD Noise Formation for Astronomical Imaging

Shuhong Liu^{1,2}, Xining Ge², Ziyang Gu¹, Quanfeng Xu³, Lin Gu⁴
Ziteng Cui^{1,2}, Xuangeng Chu^{1,2}, Jun Liu², and Dong Li²
Tatsuya Harada^{1,5}

¹The University of Tokyo, Tokyo, Japan

²I2WM, Tokyo, Japan

³Shanghai Astronomical Observatory, Shanghai, China

⁴Tohoku University, Sendai, Japan

⁵RIKEN AIP, Tokyo, Japan

Abstract. Astronomical imaging remains noise-limited under practical observing conditions. Standard calibration pipelines remove structured artifacts but largely leave stochastic noise unresolved. Although learning-based denoising has shown strong potential, progress is constrained by scarce paired training data and the requirement for physically interpretable models in scientific workflows. We propose a physics-based noise synthesis framework tailored to CCD noise formation in the telescope. The pipeline models photon shot noise, photo-response non-uniformity, dark-current noise, readout effects, and localized outliers arising from cosmic-ray hits and hot pixels. To obtain low-noise inputs for synthesis, we stack multiple unregistered exposures to produce high-SNR bases. Realistic noisy counterparts synthesized from these bases using our noise model enable the construction of abundant paired datasets for supervised learning. Extensive experiments on our real-world multi-band dataset curated from two ground-based telescopes demonstrate the effectiveness of our framework in both photometric and scientific accuracy. Code is available at <https://github.com/ShuhongLL/Denoising-Deep-Sky>.

Keywords: Computational Imaging · Noise Calibration · CCD

1 Introduction

Modern imaging systems devote substantial effort to suppressing sensor noise through advances in both hardware design and computational methods. Nevertheless, digital images remain affected by fluctuations intrinsic to photon detection, electronic readout, and residual artifacts introduced by sensor architecture [49]. Reducing these noise sources while faithfully preserving signal fidelity has therefore become a central objective for computational imaging.

For consumer photography, particularly in extreme low-light conditions, learning based denoising has advanced substantially [8]. To mitigate the scarcity in acquiring strictly paired data, recent research has focused on developing sophisticated noise modeling and synthesis techniques [16, 30, 31, 45, 65, 71]. Within




Fig. 1: Overview of denoising pipeline. Left: a CMOS sensor performs in-pixel conversion and shows a near-uniform pixel responsivity gain, while a scientific CCD sensor transfers pixel charge and exhibits spatially structured gain. Right: paired training images are limited since clean signals cannot be recovered by extending exposure or traditional calibration. We stack raw inputs to form a high-SNR baseline, synthesize realistic noisy raws using a physics-based noise model, and train the denoising network.

this setting, physics-based models [7, 35, 64], implicit modeling [15, 68], and non-parametric approaches [9, 46] have been proposed to synthesize camera-dependent noisy RAW samples at scale, tailored to the ISO-aware noise characteristics of the modern Complementary Metal-Oxide-Semiconductor (CMOS) sensors that are prevalent in consumer cameras.

Scientific imaging, particularly in astronomy, presents a different landscape where Charge-Coupled Device (CCD) sensors remain the dominant technology [24]. The CCD architecture differs fundamentally from CMOS cameras. Instead of embedding amplifiers and converters at the pixel level, CCDs transfer charge serially to a small set of output amplifiers [28]. This readout process introduces distinctive noise signatures [14], including segment-specific bias and gain variation [33], charge transfer inefficiency [41], clock-induced charge [29], and bias drift and striping [21]. At longer wavelengths, back-illuminated CCDs exhibit interference fringes or etaloning [24], caused by internal reflections within the silicon substrate. Unlike consumer CMOS sensors, science CCDs record in a single spectral band defined by filters rather than through a mosaic color pattern, which alters the statistical structure of the data. Astronomical CCDs are also exposed to unique observational effects, including cosmic-ray strikes [61] and persistent hot pixels [43] that introduce localized outliers.

Unlike consumer CMOS imaging, where noise can often be reduced by simply extending the exposure time, astronomical observations are fundamentally constrained by atmospheric effects and target dynamics. Ground-based telescopes are subject to turbulence-induced wavefront aberrations [18], scintillation [53], intrinsic source variability [67], and Earth’s rotation [4], all of which evolve on timescales that make long integrations for obtaining high-SNR images impractical. As a result, CCD observations are typically acquired as a sequence of shorter

exposures that are individually calibrated and then registered and stacked [19]. While stacking improves SNR and enables the rejection of outliers, it also faces notable limitations. Residual offsets can arise from imperfect telescope tracking and co-registration. [5]. Moreover, the strict alignment required for stacking is time-consuming and not always feasible [58], especially for faint or time-variable sources [2], or for large-area survey observations where rapid coverage is prioritized [27]. Consequently, the scientific usability of individual exposures remains essential, while conventional calibration pipelines primarily suppress structured artifacts and leave stochastic noise largely unresolved [24].

To learn the denoising model for astronomical imaging and overcome the data sparsity commonly seen in scientific imaging, we propose a physics-based noise synthesis framework tailored to CCD noise formation. As illustrated in Figure 1, our pipeline models the principal sources of CCD noise, including photon shot noise, per-amplifier bias offsets and gain non-uniformity, dark-current noise, and readout effects such as read noise and charge transfer inefficiency at the output amplifier. It further accounts for observational phenomena, including transient cosmic-ray hits and persistent hot pixels, which can noticeably affect raw exposures. A natural question is how to obtain the “clean” bases for synthesis, since obtaining truly noise-free references is impractical [23]. We leverage common practice in modern telescopes that track targets with active guiding so the scene remains approximately fixed on the detector across sequential exposures. We directly stack multiple unregistered observation frames, which inevitably broadens the astrophysical point-spread function due to residual drift and seeing variation, but substantially suppresses stochastic noise without introducing interpolation artifacts. Such averaged images are not scientifically usable per se; however, they provide abundant high-SNR bases that preserve correct large-scale background structure and instrument signatures. Our synthesis pipeline can generate realistic noisy counterparts from these bases, enabling the construction of paired datasets for supervised learning. To validate the framework, we further construct a real astronomical image dataset that contains paired raw noisy frames and instrument-pipeline calibrated frames collected from two twin ground-based telescope observations for real-world evaluation.

Our main contributions can be summarized as follows:

- We propose a physics-based noise formulation model tailored to astronomical CCD imaging, capable of synthesizing realistic noisy RAW data that closely matches the statistical properties of real observations.
- We construct real-world multi-band astronomical imaging datasets from ground telescope observations, providing both raw and processed observations for noise synthesis and quantitative evaluation.
- Comprehensive in-domain and cross-instrument experiments demonstrate that our method consistently outperforms baseline approaches in both photometric fidelity and scientific accuracy.

2 Related Work

2.1 Denoising in Consumer CMOS Imaging

Recent efforts in CMOS image denoising have largely shifted from traditional model-based methods to deep learning approaches. While early learning-based methods trained on synthetic Gaussian noise showed promise [69], they often struggled to generalize to real-world scenarios where noise is more complex and signal-dependent [8, 36]. To bridge this gap, physically based noise models that account for signal-dependent and independent noise have been proposed [16, 31, 64, 71]. Further refinements to these models have incorporated ISO-correlated noise characteristics [7] and explored nonparametric approaches [46] to capture intricate noise distributions. To address the challenge of acquiring large-scale paired datasets, unsupervised and self-supervised approaches have been developed without requiring clean images [34, 48, 73]. Beyond single-image methods, burst-based methods [20, 44] exploit temporal redundancy across multiple exposures to further improve denoising quality. More recently, generative priors have been explored for zero-shot denoising without training [51, 56].

2.2 Denoising in Astronomical CCD Imaging

Astronomical imaging denoising is a fundamental challenge due to the extremely low photon counts associated with observing faint celestial objects. A cornerstone of conventional practice is the use of calibration frames, which correct for readout patterns, thermal signals, pixel-to-pixel nonuniformity, and optical vignetting [24, 28]. To mitigate transient artifacts and improve the SNR, multiple dithered exposures, sigma-clipping, and median filtering are commonly employed [12, 61], while advanced techniques such as drizzle have been developed to address undersampled images from the Hubble Space Telescope [19]. In addition, transform-domain methods, especially wavelet-based approaches, separate noise from astronomical structures by exploiting multi-scale representations [57]. More recently, deep learning has emerged as a promising approach, often bypassing the need for a multi-stage calibration pipeline by performing end-to-end denoising that simultaneously removes instrumental signatures and artifacts [37, 62, 70]. Self-supervised methods are also proposed to model noise directly from observations, offering an automated solution for the modern sky survey [23, 38, 72].

3 Method

To construct pair-wise training data for learning denoising models, our objective is to synthesize realistic CCD noise using only a limited number of stacked clean bases and calibration frames. We begin by introducing the CCD image formation process in Section 3.1, then describe the synthesis of signal-dependent noise and signal-independent noise in Section 3.2 and Section 3.3. The overview of the noise formation model is illustrated in Figure 2.




Fig. 2: Overview of noise formation in astronomical imaging. Unique telescope optics and CCD science cameras introduce noise characteristics that are distinct from commercial CMOS imaging. Zoom in the noise frames for better visibility.

3.1 Astronomical Image Formation Model

A CCD camera converts incident photons into photoelectrons and also accumulates thermally generated dark current during the exposure. The total charge collected in each pixel is then transferred through the serial register to one or more output amplifiers and digitized. At the frame level, this process can be separated into signal-dependent terms, which scale with illumination and exposure time, and signal-independent terms, which originate from the readout chain, sparse defects, and quantization. The system gain is G in electrons per Analog-to-Digital Unit (ADU). Let $S(x)$ denote the photon arrival rate at pixel x in photons per second, $\eta(x)$ the quantum efficiency in electrons per photon, and $\text{PRNU}(x)$ a multiplicative factor describing pixel-to-pixel Photo Response Non-uniformity (PRNU). The raw digital value $I(x)$ is then modeled as follows:

$$\begin{aligned}
 I(x) = & \underbrace{N_{\text{overscan}}(x)}_{\text{Constant}} + \underbrace{\frac{1}{G} (\eta(x)S(x)\text{PRNU}(x) + N_{\text{photon}}(x) + N_{\text{dark}}(x))}_{\text{Signal-dependent}} \\
 & + \underbrace{N_{\text{read}}(x) + N_{\text{digit}}(x) + N_{\text{hot}}(x)}_{\text{Signal-independent}}
 \end{aligned} \tag{1}$$

In this expression, $N_{\text{overscan}}(x)$ is the overscan pedestal, which provides a constant offset ensuring that all digitized pixel values remain positive. The signal-dependent component consists of three contributions, including the target signal describing the expected photoelectron count, the photon-shot noise term $N_{\text{photon}}(x)$, and the dark-current noise realization $N_{\text{dark}}(x)$. The factor $1/G$ provides the linear conversion from electrons to the ADU domain. The signal-independent component then comprises additive $N_{\text{read}}(x)$, $N_{\text{hot}}(x)$, and $N_{\text{digit}}(x)$. Here, $N_{\text{read}}(x)$ is read noise arising from the output amplifier and correlated-double-sampling chain, including contributions from reset and thermal noise and




Fig. 3: PRNU comparison between G' and R' bands. A captured sky flat frame (a) is typically expressed as the product of (b) the anti-fringing fixed pattern, (c) the low-frequency twilight sky background, and (d) dust-donut artifacts. PRNU is given by (b) $\times$ (d), where (c) is the external signal.

low-frequency amplifier fluctuations. $N_{\text{digit}}(x)$ is the digitization error. $N_{\text{hot}}(x)$ is a sparse defect term that aggregates hot-pixel outliers and cosmic-ray events.

3.2 Signal-Dependent Noise

Photo Response Non-Uniformity (PRNU) PRNU describes fixed, spatially varying differences in pixel responsivity, such that identical uniform illumination can yield different pixel outputs [24]. In commercial CMOS cameras, this effect is handled through *flat-field correction* or *shading correction*, which uses a uniform bright-field to compensate pixel-response variations [1, 55]. In consumer imaging pipelines, PRNU rarely dominates visual quality and is often neglected [3, 35, 65], as in-camera ISP steps typically mask subtle nonuniformity.

In astronomical CCD imaging, PRNU correction is essential for reliable photometry and background estimation: science targets are often faint and extended, and percent-level sensitivity variations can translate into systematic flux biases and large-scale background artifacts [24, 40, 50]. The observed PRNU pattern reflects both (i) intrinsic detector nonuniformity, such as pixel-to-pixel gain and quantum-efficiency variation, and (ii) multiplicative effects introduced by the optical path. For (i), thinned and back-illuminated CCDs can exhibit wavelength-dependent optical etaloning [63]. Fringe-suppression manufacturing processes therefore introduce controlled thickness variation to disrupt the Fabry–Pérot cavity [6] and reduce fringing, while leaving residual high-frequency structure in the flat-field response [24, 25] as shown in Figure 3b. For (ii), dust on the dewar window, filters, or nearby optics introduces out-of-focus shadow features that appear as ring-like *dust donuts*, shown in Figure 3d, in flat fields [24, 50, 59].

To correct these multiplicative response variations in real observations, flat-field nonuniformity is typically performed using twilight *skyflats* in the traditional calibration pipeline [42]. The telescope is pointed to a blank sky region and multiple exposures are acquired under near uniform illumination. Each skyflat frame is normalized by its median to remove exposure-to-exposure brightness differences, then combined to form a master skyflat frame. A science image is typically calibrated by dividing the raw measurement by the master skyflat map.

However, twilight skyflats contain both the intrinsic detector response and external illumination structure, so direct stacking can bias PRNU estimation. Large-scale twilight gradients, vignetting, and scattered light vary across frames, whereas $\text{PRNU}(x)$ is a fixed sensor property. To isolate the pixel-scale fixed pattern, we remove the low-frequency illumination component from each skyflat frame. Given a skyflat image $N_{\text{sky}}^i(x)$, we compute the median-normalized frame $\tilde{N}_{\text{sky}}^i(x)$, estimate an illumination background $N_{\text{bg}}^i(x)$ from block-wise large kernel median pooling, and form a high-pass residual $H_{\text{hp}}^i(x)$ as:

$$H_{\text{hp}}^i(x) = \frac{\tilde{N}_{\text{sky}}^i(x)}{N_{\text{bg}}^i(x)}, \quad \text{PRNU}_{\text{fp}}(x) = \text{norm}\left(\text{median}_{\sigma\text{-clip}}\left(\{H_{\text{hp}}^i(x)\}_i\right)\right). \quad (2)$$

Here $\text{median}_{\sigma\text{-clip}}(\cdot)$ denotes pixel-wise sigma-clipped median aggregation over multiple residuals, and $\text{norm}(\cdot)$ rescales the resulting map to unit median. The resulting $\text{PRNU}_{\text{fp}}(x)$ models the static sensor-intrinsic fixed pattern.

After estimating $\text{PRNU}_{\text{fp}}(x)$, we extract structured flat-field donut artifacts that remain after fixed-pattern removal. For each normalized skyflat $\tilde{N}_{\text{sky}}^i(x)$, we divide out $\text{PRNU}_{\text{fp}}(x)$ to obtain a donut map $\text{PRNU}_{\text{donut}}^i(x)$, which is randomly sampled during synthesis. The final PRNU model is defined as:

$$\text{PRNU}_{\text{donut}}^i(x) = \frac{\tilde{N}_{\text{sky}}^i(x)}{\text{PRNU}_{\text{fp}}(x)}, \quad \text{PRNU}(x) = \text{PRNU}_{\text{fp}}(x) \times \text{PRNU}_{\text{donut}}(x). \quad (3)$$

Photon-Shot Noise Unlike CMOS camera, where the effective gain varies with ISO and typically requires careful calibration [64] or hypothesized gain estimation [35] for each setting, the system gain G of an astronomical CCD is precisely characterized and known in advance. Under stationary illumination within the linear response regime, the photon-shot electron count E at pixel x over exposure time t is modeled as a Poisson random variable whose mean equals the expected photoelectron yield after quantum conversion and pixel sensitivity:

$$E(x) + N_{\text{photon}}(x) \sim \text{Poisson}(t\eta(x)S(x)\text{PRNU}(x)) \quad (4)$$

Dark-Current Noise Dark current arises from thermally generated electrons that accumulate in CCD pixels even without incident light. Since it is thermally activated, its level depends strongly on detector temperature and is substantially reduced by cooling [24, 66]. In typical astronomical observations, the detector




Fig. 4: Quantile–Quantile plots of fitting Gaussian models to dark frames residuals.

temperature is regulated and calibration frames are acquired under the same operating conditions as science images [32]. We therefore treat the dark current as constant for a given instrument setting and neglect explicit temperature dependence in our model.

Following the conventional procedure [42], we represent the dark current as a per-pixel rate in electrons per second, $D_{\text{dark}}(x)$ (e^-/s), which is retrieved from long-exposure dark frames (300 s) and normalized by the integration time [42]. For a science exposure time t , the accumulated dark signal can be estimated as:

$$N_{\text{dark}}(x) \sim t \times D_{\text{dark}}(x). \quad (5)$$

To model $D_{\text{dark}}(x)$, we first compute a reference map $D_{\text{ref}}(x)$ by averaging the empirically captured dark-current maps $\{D_{\text{dark}}^i(x)\}$. We then analyze the sample-to-sample residual $\epsilon_{\text{dark}}^i(x)$ and fit a Gaussian distribution to the residuals as:

$$\epsilon_{\text{dark}}^i(x) \sim \mathcal{N}(0, \text{Var}_{i,x}[\epsilon_{\text{dark}}^i(x)]), \text{ where } \epsilon_{\text{dark}}^i(x) = D_{\text{dark}}^i(x) - D_{\text{ref}}(x) \quad (6)$$

This formulation models the spatial pattern of dark-current noise N_{dark} using reference $D_{\text{ref}}(x)$, residual ϵ_{dark} , and target exposure duration t . The Quantile–Quantile plots for both the fitting set and testing set for ϵ_{dark} are shown in Figure 4, suggesting that using a simple Gaussian distribution is sufficient.

3.3 Signal-Independent Noise

Fixed-Pattern Noise and Read Noise A CCD bias frame is acquired with zero integration time and records the electronic offset added by the readout chain. Stacking multiple bias frames reduces stochastic read noise and makes deterministic spatial structures more apparent, including *column-wise banding* where certain columns exhibit a consistently elevated bias level [47, 50], as shown in the right columns of Figure 5a. This banding is a fixed-pattern artifact associated with CCD readout, where charge is transferred through a serial register and measured by one or a few shared output amplifiers, so small baseline drifts or electronic pickup can manifest as coherent column-dependent offsets [47, 50].




Fig. 5: Fixed-pattern noise and read noise within 97% sigma-clipping.

Given a set of bias frames $\{B_{\text{bias}}^i(x)\}$, we first estimate a fixed-pattern template $B_{\text{fp}}(x)$ by stacking. This template preserves the stable bias structure, including vertical banding and repeatable offset patterns. For each sample, we compute the residual by subtracting $B_{\text{fp}}(x)$, then extract a row-wise bias stripping by taking the mean residual within each image row. The remaining residual models pixel-wise readout fluctuations. We fit Gaussian mixture models to both components following [35] as:

$$N_{\text{row}}(\rho(x)) \sim \sum_{k=1}^{K_r} \pi_k^{(r)} \mathcal{N}_k, \quad \epsilon_{\text{pix}}^i(x) \sim \sum_{k=1}^{K_p} \pi_k^{(p)} \mathcal{N}_k, \quad (7)$$

where $\rho(x)$ denotes the row index of pixel x , K_r and K_p are the numbers of mixture components, and $\{\pi_k^{(\cdot)}\}$ are nonnegative mixture weights that sum to 1. Each $\mathcal{N}_k$ denotes a Gaussian model.

During synthesis, we sample row offset from the row mixture for each row and broadcast it to all pixels in that row via $N_{\text{row}}(\rho(x))$, and sample pixel-wise fluctuations $\epsilon_{\text{pix}}(x)$ independently from the pixel mixture. The resulting read-noise realization is constructed as:

$$N_{\text{read}}(x) = B_{\text{fp}}(x) + N_{\text{row}}(\rho(x)) + \epsilon_{\text{pix}}(x), \quad (8)$$

which preserves the observed fixed-pattern bias structure through $B_{\text{fp}}(x)$ while reproducing realistic row-level and pixel-level read-noise variability. Figure 6 shows the testing results of the Quantile-Quantile plot, with a close fit to both theoretical distributions.

Digitization Error After charge-to-voltage conversion, the CCD output is digitized by an analog-to-digital converter, mapping a continuous signal to a discrete digital number (ADU). This introduces a small quantization error $N_{\text{dig}}(x) \sim U(-\Delta/2, \Delta/2)$, where $\Delta = G$ in electron units. Because read noise is empirically estimated from bias frames, this digitization error is already absorbed into the measured read-noise term [24], so we do not model $N_{\text{dig}}(x)$ separately.




Fig. 6: Quantile–Quantile plots of fitting GMM models to row-pattern and pixel residuals.

Hot-Pixel and Cosmic Ray Hot pixels arise from detector defects with abnormally high dark current. Their intensities increase approximately linearly with exposure time due to dark-signal accumulation, and their locations remain fixed across frames [26]. Cosmic rays are transient high-energy events that deposit charge in the sensor and appear as impulsive bright outliers. In ground-based CCD imaging, cosmic-ray hits typically contaminate a small fraction of pixels per long exposure, affecting approximately 0.01%–0.1% of image pixels [13, 22, 39].

Rather than counting persistent hot pixels and transient cosmic rays separately, we model both effects using a single salt-and-pepper impulse noise model:

$$N_{\text{hot}}(x) = Z(x) A(x), \quad Z(x) \sim \text{Bernoulli}(p_{\text{hot}}), \quad (9)$$

where $A(x)$ is drawn from a uniform distribution [5000, 10000], and $p_{\text{hot}} = 0.01\%$ denotes the impulse probability that controls the fraction of corrupted pixels.

4 MuSCAT Observer Dataset

Instruments MuSCAT-3 (Haleakalā, Hawaii; LCO-FTN, 2020) and MuSCAT-4 (Siding Spring, Australia; LCO-FTS, 2023) are twin four-channel optical imagers installed on 2 m Las Cumbres Observatory telescopes. Each instrument uses a cascade of dichroic beam splitters to feed four independent CCD cameras in the G' , R' , I' , and Z_s bands. In this work, we focus on G' , R' , and I' , which follow the SDSS filter system and are the most consistently calibrated and frequently observed channels. Since MuSCAT-3 and MuSCAT-4 share nearly identical opto-mechanical designs and the same filter set, they enable cross-instrument zero-shot evaluation; we treat each band as an independent dataset to assess denoising performance under distinct, device-dependent noise characteristics.

Data Statistics We construct MuSCAT-3 for training and in-domain testing, and reserve MuSCAT-4 for zero-shot evaluation. For each observation, we provide (i) the raw science frame, (ii) the BANZAI “processed” product, and (iii) the

Table 1: Dataset statistics of MuSCAT-3 and MuSCAT-4 datasets.

Data	MuSCAT-3			MuSCAT-4		
	G'	I'	R'	G'	I'	R'
Training	215	435	395	-	-	-
Testing	54	73	61	46	96	72
Calibration	807	1524	1368	138	288	216
Processed	4843	9578	9677	661	1687	1566

Table 2: Comparison of background noise levels under different processing methods.

Methods	MuSCAT-3		MuSCAT-4	
	NMAD↓	STD↓	NMAD↓	STD↓
Raw	10.50	11.73	11.61	11.73
BANZAI	9.33	9.54	9.40	9.63
Mean	3.12	4.72	3.70	4.72
Drizzle	6.15	6.48	6.22	6.55

calibration frames, including bias, dark, and skyflats. BANZAI is the official LCO real-time reduction pipeline and performs standard instrumental signature removal [42]. Table 1 summarizes the dataset statistics per band.

While BANZAI removes deterministic detector and optical artifacts, the calibrated science images still largely contain photon-shot noise [42, 52]. To obtain high-SNR base images for noise synthesis, we additionally collect 10–20 consecutive BANZAI-processed exposures for each science target. We stack these exposures to form a high-SNR reference image [10, 17], denoted as **Mean**. Table 2 reports background noise levels measured on source-masked pixels using both normalized median absolute deviation (NMAD) and standard deviation (STD).

5 Experiment

Implementation Details Following previous studies [8, 35, 65], we construct a neural denoising pipeline using a standard U-Net architecture [54]. We use high-SNR stacking-mean images as clean baselines and synthesize noisy observations to form paired training data. We sample 100 calibration frames from the training set to build our physics-based noise model. We train the U-Net with 140 epochs using Raw-L1 loss and Adam optimizer with batch size 8 on a single H200 GPU. The learning rate is 1×10^{-4} . Detailed information is given in Appendix D.

Baseline Methods We evaluate both traditional and learning-based baselines. Traditional methods include BM3D [11] and A-BM3D [60] for single-image denoising, and Drizzle [19] for multi-image reconstruction from aligned exposures. Learning-based baselines include supervised methods trained with direct paired data or PMN [16], as well as self-supervised approaches including Noise2Noise (N2N) [34], ELD [7], SFRN [71], PNNP [15], and NMOH [35]. Implementation details of baseline approaches are outlined in the Appendix E.

Metric For synthetic data generated with our noise formation model, we report PSNR to evaluate photometric fidelity. For real observations, we report FLUX-SNR (dB) to assess scientific accuracy. We also compute NMAD and STD on source-masked background pixels, which are expected to be near constant, to quantify residual noise levels. We clip both prediction and GT at 5000,




Fig. 7: Visualization of baseline method comparisons on the MuSCAT-3 dataset, G' band, targeting TOI-4643 (2022-10-26 observation). Zoom in for better visibility.




Fig. 8: Visualization of zero-shot evaluation on the MuSCAT-4 dataset, R' band, targeting WASP-53 (2023-10-08 observation). Zoom in for better visibility.

corresponding to 99.99% of pixels, to suppress extreme outliers that could bias photometric evaluation. Best results are shaded as **first**, **second**, and **third**.

5.1 Evaluation on MuSCAT-3 Dataset

Table 3 reports quantitative results on the MuSCAT-3 dataset for both synthetic and real observations across three bands. For the synthetic setting, we use our noise formation model to generate realistic noisy raw observations from stacked clean proxies, while incorporating the corresponding per-observation calibration frames acquired on the same night to match real acquisition conditions. Under this synthetic evaluation, our method consistently outperforms all baselines and achieves the best photometric accuracy.

To evaluate scientific accuracy on real observations, where GT is typically inaccessible, we follow common practice in astronomy by injecting mock PSFs [23] with predefined photon flux into real observations and computing FLUX-SNR against the mocked GT. In addition, we report background noise statistics, including NMAD and STD, measured on source-masked regions, for a comple-

Table 3: Quantitative evaluation on MuSCAT-3 dataset.

	Method	PSNR(Syn) $\uparrow$				FLUX-SNR(Real) $\uparrow$				NMAD(Real) $\downarrow$				STD(Real) $\downarrow$			
		G'	I'	R'	Avg	G'	I'	R'	Avg	G'	I'	R'	Avg	G'	I'	R'	Avg
Trad	BM3D [11]	55.03	54.88	54.79	54.90	37.56	38.28	37.10	37.64	2.59	3.11	2.58	2.76	3.21	3.31	2.99	3.17
	ABM3D [60]	55.38	54.57	54.60	54.85	39.48	38.44	40.13	39.35	1.20	1.15	1.15	1.17	2.39	2.13	2.11	2.21
	Drizzle [19]	-	-	-	-	-	-	-	-	6.05	6.08	6.31	6.15	6.36	6.50	6.58	6.48
Learning	Paired	54.74	56.07	56.22	55.68	26.13	28.86	27.14	27.38	7.56	9.28	9.02	8.62	7.84	9.50	9.21	8.85
	PMN [16]	57.06	56.68	56.76	56.83	30.55	30.48	29.37	30.13	8.13	9.34	8.95	8.81	8.35	9.48	9.15	8.99
	N2N [34]	56.29	56.66	56.91	56.62	36.16	36.11	35.96	36.08	7.76	9.29	8.92	8.66	8.02	9.46	9.12	8.87
	ELD [7]	56.05	55.89	56.38	56.11	32.16	30.02	28.12	30.10	2.60	2.32	2.58	2.50	3.48	3.58	4.49	3.85
	SFRN [71]	51.23	56.13	55.70	54.35	28.87	29.47	28.39	28.91	4.20	4.98	4.82	4.67	5.27	5.75	5.55	5.52
	PNNP [15]	57.10	56.88	56.80	56.92	38.16	37.52	37.11	37.60	6.39	7.95	7.64	7.33	6.75	8.04	7.73	7.51
	NMOH [35]	56.03	56.28	56.86	56.39	37.51	36.26	36.83	36.87	2.58	2.35	2.45	2.46	3.46	3.21	3.45	3.37
	Ours	57.29	56.95	57.09	57.06	41.21	38.50	40.03	39.91	2.21	2.25	2.31	2.43	2.92	2.86	3.42	3.06

Table 4: Quantitative zero-shot evaluation on MuSCAT-4 dataset.

	Method	PSNR(Syn) $\uparrow$				FLUX-SNR(Real) $\uparrow$				NMAD(Real) $\downarrow$				STD(Real) $\downarrow$			
		G'	I'	R'	Avg	G'	I'	R'	Avg	G'	I'	R'	Avg	G'	I'	R'	Avg
Trad	BM3D [11]	55.07	55.55	56.61	55.74	39.00	37.56	39.13	38.56	2.99	2.34	2.26	2.53	4.04	3.96	3.62	3.87
	ABM3D [60]	57.75	57.25	56.92	57.31	40.51	37.70	40.04	39.42	2.42	2.45	2.18	2.35	3.69	3.51	3.37	3.52
	Drizzle [19]	-	-	-	-	-	-	-	-	6.17	6.51	5.98	6.22	6.24	6.90	6.51	6.55
Learning	Paired	57.11	56.52	56.63	56.75	25.33	28.36	29.12	27.60	8.11	10.06	9.12	9.10	8.49	10.57	9.36	9.47
	PMN [16]	59.13	56.39	56.62	57.38	29.88	30.97	29.64	30.16	8.77	10.12	9.05	9.31	8.99	10.63	9.38	9.67
	N2N [34]	58.21	56.75	57.25	57.40	37.03	36.51	36.06	36.53	8.42	10.10	9.10	9.21	8.70	10.59	9.31	9.53
	ELD [7]	59.13	57.29	56.81	57.74	32.56	30.65	31.04	31.42	2.87	2.59	2.53	2.66	3.49	4.33	3.81	3.88
	SFRN [71]	53.80	53.33	52.08	53.07	28.17	29.80	29.55	29.17	4.50	5.01	4.47	4.66	5.12	6.32	5.58	5.67
	PNNP [15]	59.30	57.36	57.72	58.13	37.76	36.27	37.07	37.03	7.15	7.60	6.85	7.20	7.56	8.27	7.40	7.74
	NMOH [35]	59.04	57.93	57.46	58.14	36.95	36.28	36.34	36.52	2.79	2.68	2.54	2.67	3.80	4.49	4.07	4.12
	Ours	60.22	58.34	57.98	58.85	42.07	37.98	39.56	39.87	2.35	2.51	2.53	2.49	3.48	3.83	3.62	3.64

mentary evaluation of background smoothness and residual noise suppression. Compared with learning-based approaches [7,15,16,34,35,71], our method consistently achieves superior flux accuracy and background smoothness, whereas baseline methods designed for commercial cameras often overlook PRNU and CCD-specific noise characteristics. Compared with traditional denoisers [11,19,60], although they may achieve stronger noise suppression, reflected by lower NMAD and STD, they tend to oversmooth stellar PSFs, attenuating peak intensities and fine PSF structures, which leads to lower FLUX-SNR.

5.2 Zero-shot Evaluation on MuSCAT-4 Dataset

We further conduct a cross-instrument evaluation by applying the model trained on MuSCAT-3 observations to denoise MuSCAT-4 in a zero-shot manner, as reported in Table 4. The results follow trends consistent with Table 3, validating the generalization ability of our physics-based training pipeline across instruments and observing conditions.

Moreover, Figure 7 and Figure 8 provide qualitative comparisons on real G' -band and I' -band observations from the MuSCAT-3 and 4 datasets, showing that




Fig. 9: Visualization of ablation studies on MuSCAT-3 dataset G' -band, targeting the quasar SDSS J0106 (2021-11-25 observation).

our method delivers optimal and photometrically faithful denoising performance among baseline approaches while preserving stellar PSF structures.

6 Ablation Study

Effect of Noise Formation Components

We conduct ablation studies by synthesizing training pairs with each component removed at a time and training the networks under the same settings. Quantitative results are reported in Table 5, with representative visual comparisons shown in Figure 9. Specifically, removing PRNU, N_{phot} , and N_{read} significantly reduces FLUX-SNR, leading to pronounced scientific artifacts in PSF flux recovery. These failure cases are consistent

with Figure 9, where stronger residual noise remains in the restored results. In particular, N_{read} models fixed-pattern noise, which can substantially affect flux estimation when the pattern overlaps with mocked PSFs. Moreover, removing N_{hot} can produce lower noise levels on source-masked backgrounds, reflected by lower NMAD and STD; however, modeling such artifacts is still essential for optimal flux estimation, as it improves robustness to real impulsive corruptions. Across the remaining ablations, removing any component consistently degrades performance, underscoring the necessity of the full noise formation model.

Table 5: Ablation study on each noise components in Equation (1) evaluated on *real* observations with mocked PSF. Left cell: MuSCAT-3 G' . Right cell: zero-shot on the MuSCAT-4 G' .

Methods		FLUX-SNR $\uparrow$		NMAD $\downarrow$		STD $\downarrow$	
w/o	PRNU	35.51	35.17	3.02	3.26	3.76	4.08
	N_{dark}	38.21	37.62	2.78	2.99	3.66	3.68
	N_{phot}	32.16	31.51	5.72	5.90	5.97	6.03
	N_{read}	31.69	32.03	2.38	2.98	3.81	3.83
	N_{hot}	40.88	41.64	2.46	2.82	3.32	3.36
	Ours	41.21	42.07	2.58	2.79	3.46	3.48

Sensitivity to the Number of Calibration Frames

We examine how the number of calibration frames affects learning noise characteristics by evaluating denoising performance on the G' -band, and compare against the direct sampling strategy used in NMOH [35]. Figures 10a and 10b show a consistent improvement as more calibration frames are available; in the low-data regime, direct sampling can achieve higher performance, whereas physics-based synthesis requires sufficient calibration data to reliably parameterize the statistical model. In contrast, Figures 10c and 10d indicate that direct sampling is prone to overfitting to the




Fig. 10: Ablation study on number of calibration frames. We compare our noise synthesis method with direct sampling of calibration data at different counts. We show results on MuSCAT-3 for supervised training and MuSCAT-4 on zero-shot inference.

observed noise realizations, resulting in elevated background noise relative to the physics-based approach. Given the practical scarcity of calibration frames, we use 100 frames throughout all experiments, while fitting more calibration frames can still consistently improve denoising performance.

7 Conclusion

In this study, we propose a physics-based CCD noise formation model and a practical synthesis pipeline that generates realistic noisy RAW observations for supervised denoising. Experiments of both in-domain and cross-instrument settings in our novel corrupted real-world multi-band dataset demonstrate that our method consistently outperforms baseline approaches in both photometric and scientific accuracy.

Limitation Strict noise-free ground-truth is unavailable for real astronomical observations, so our evaluation relies on stacked high-SNR bases and calibrated products as practical proxies. In addition, while we include cross-instrument evaluations, all experiments are conducted on two MuSCAT telescopes. Broader validation across other telescopes, sensors, and observing conditions remains an important direction for future studies.

Acknowledgements

This work was partially supported by JST Moonshot R&D Grant Number JPMJPS2011, CREST Grant Number JPMJCR2015 and Basic Research Grant (Super AI) of Institute for AI and Beyond of the University of Tokyo. In addition, this work was also partially supported by JST SPRING, Grant Number JPMJSP2108.

References

1. AG, B.: Enhanced shading and flat field correction. <https://docs.baslerweb.com/flat-field-correction> (2024), product documentation
2. Alard, C., Lupton, R.H.: A method for optimal image subtraction. *The Astrophysical Journal* **503**(1), 325 (1998)
3. Becker, G.S., Lovas, R.: Uniformity correction of cmos image sensor modules for machine vision cameras. *Sensors* **22**(24), 9733 (2022)
4. Benn, C., Ellison, S.: Brightness of the night sky over la palma. *New Astronomy Reviews* **42**(6-8), 503–507 (1998)
5. Bertin, E.: Swarp: resampling and co-adding fits images together. *Astrophysics Source Code Library* pp. ascl-1010 (2010)
6. Born, M., Wolf, E.: Principles of optics: electromagnetic theory of propagation, interference and diffraction of light. Elsevier (2013)
7. Cao, Y., Liu, M., Liu, S., Wang, X., Lei, L., Zuo, W.: Physics-guided iso-dependent sensor noise modeling for extreme low-light photography. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5744–5753 (2023)
8. Chen, C., Chen, Q., Xu, J., Koltun, V.: Learning to see in the dark. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3291–3300 (2018)
9. Colom, M., Buades, A., Morel, J.M.: Nonparametric noise estimation method for raw images. *Journal of the Optical Society of America A* **31**(4), 863–871 (2014)
10. Craig, M.: Image combination. *CCD Data Reduction Guide* (2023), <https://www.astropy.org/ccd-reduction-and-photometry-guide/v/dev/notebooks/01-06-Image-combination.html>
11. Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K.: Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on Image Processing* **16**(8), 2080–2095 (2007)
12. van Dokkum, P.G., Bloom, J., Tewes, M.: La cosmic: Laplacian cosmic ray identification. *Astrophysics Source Code Library* pp. ascl-1207 (2012)
13. Farage, C.L., Pimblet, K.A.: Evaluation of cosmic ray rejection algorithms on single-shot exposures. *Publications of the Astronomical Society of Australia* **22**(3), 249–256 (2005)
14. Faraji, H., MacLean, W.J.: Ccd noise removal in digital images. *IEEE Transactions on image processing* **15**(9), 2676–2685 (2006)
15. Feng, H., Wang, L., Huang, Y., Wang, Y., Zhu, L., Huang, H.: Physics-guided noise neural proxy for practical low-light raw image denoising. *arXiv preprint arXiv:2310.09126* (2023)
16. Feng, H., Wang, L., Wang, Y., Fan, H., Huang, H.: Learnability enhancement for low-light raw image denoising: A data perspective. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **46**(1), 370–387 (2023)
17. Fischer, P., Kochanski, G.P.: Optimal addition of images for detection and photometry. *arXiv preprint astro-ph/9310031* (1993)
18. Fried, D.L.: Optical resolution through a randomly inhomogeneous medium for very long and very short exposures. *Journal of the Optical Society of America* **56**(10), 1372–1379 (1966)
19. Fruchter, A.S., Hook, R.N.: Drizzle: A method for the linear reconstruction of undersampled images. *Publications of the Astronomical Society of the Pacific* **114**(792), 144 (2002)

20. Godard, C., Matzen, K., Uyttendaele, M.: Deep burst denoising. In: Proceedings of the European Conference on Computer Vision. pp. 538–554 (2018)
21. Grogin, N.A., Lim, P.L., Maybhate, A., Hook, R.N., Loose, M.: Post-sm4 ACS/WFC bias striping: Characterization and mitigation. Instrument Science Report ACS 2011-05, Space Telescope Science Institute (2011), https://www.stsci.edu/files/live/sites/www/files/home/hst/instrumentation/acs/documentation/instrument-science-reports-isrs/_documents/isr1105.pdf
22. Groom, D.: Cosmic rays and other nonsense in astronomical ccd imagers. *Experimental Astronomy* **14**(1), 45–55 (2002)
23. Guo, Y., Zhang, H., Li, M., Yu, F., Wu, Y., Hao, Y., Huang, S., Liang, Y., Lin, X., Li, X., Wu, J., Cai, Z., Dai, Q.: Deeper detection limits in astronomical imaging using self-supervised spatiotemporal denoising. *Science* (2026). <https://doi.org/10.1126/science.ady9404>, <https://www.science.org/doi/abs/10.1126/science.ady9404>
24. Howell, S.B.: Handbook of CCD Astronomy. Cambridge Observing Handbooks for Research Astronomers, Cambridge University Press, 2 edn. (2006)
25. Howell, S.B.: Fringe science: defringing ccd images with neon lamp flat fields. *Publications of the Astronomical Society of the Pacific* **124**(913), 263 (2012)
26. Institute, S.T.S.: Wfc3 data handbook, section 5.3: Dark current and hot pixels. HST Documentation (2024), <https://hst-docs.stsci.edu/wfc3dnhb/chapter-5-wfc3-uvis-sources-of-error/5-3-dark-current-and-hot-pixels>
27. Ivezić, Ž., Kahn, S.M., Tyson, J.A., Abel, B., Acosta, E., Allsman, R., Alonso, D., AlSayyad, Y., Anderson, S.F., Andrew, J., et al.: Lsst: from science drivers to reference design and anticipated data products. *The Astrophysical Journal* **873**(2), 111 (2019)
28. Janesick, J.R., Elliott, T., Collins, S., Blouke, M.M., Freeman, J.: Scientific charge-coupled devices. *Optical Engineering* **26**(8), 692–714 (1987)
29. Janesick, J., of Physics, A.I.: Photon Transfer: DN [λ]. SPIE (2007), <https://books.google.co.jp/books?id=RydDzQEACAAJ>
30. Jiang, H., Wazawa, T., Sato, I., Nagai, T., Zheng, Y.: Revolutionizing emccd denoising through a novel physics-based learning framework for noise modeling. In: The Thirteenth International Conference on Learning Representations (2025)
31. Jin, X., Xiao, J.W., Han, L.H., Guo, C., Zhang, R., Liu, X., Li, C.: Lighting every darkness in two pairs: A calibration-free pipeline for raw denoising. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13275–13284 (2023)
32. Kaaret, P.: Readout noise and dark current. https://homepage.physics.uiowa.edu/~pkaaret/2015f_a4850/Lab03_noise.html, university of Iowa Lab Notes
33. Lauer, T.R.: The reduction of wide field/planetary camera images. *Publications of the Astronomical Society of the Pacific* **101**(638), 445 (1989)
34. Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., Aila, T.: Noise2noise: Learning image restoration without clean data. In: Dy, J., Krause, A. (eds.) Proceedings of the International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 80, pp. 2965–2974. PMLR (10–15 Jul 2018)
35. Li, F., Jiang, H., Iso, D.: Noise modeling in one hour: Minimizing preparation efforts for self-supervised low-light raw image denoising. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5699–5708 (2025)

36. Liu, S., Chang, G., Liu, J., Chu, X., Zheng, Y., Harada, T., Cui, Z.: Rawild: Sensor-agnostic raw object detection via physics-guided curve and grid modeling. arXiv preprint arXiv:2605.05941 (2026)
37. Liu, S., Ge, X., Cui, Z., Li, L., Chang, G., Liu, J., Gu, Z., Li, D., Chu, X., Gu, L., et al.: Fluxflow: Conservative flow-matching for astronomical image super-resolution. arXiv preprint arXiv:2605.03749 (2026)
38. Liu, T., Quan, Y., Su, Y., Guo, Y., Liu, S., Ji, H., Hao, Q., Gao, Y., Liu, Y., Wang, Y., Sun, W., Ding, M.: Astronomical image denoising by self-supervised deep learning and restoration processes. *Nature Astronomy* **9**(4), 608–615 (2025)
39. Lupton, R.H., Ivezic, Z., Gunn, J.E., Knapp, G., Strauss, M.A., Richmond, M., Ellman, N., Newburg, H., Fan, X., Yasuda, N., et al.: Sdss image processing ii: The photo pipelines. Preprint, Princeton University. <https://www.astro.princeton.edu/~rhl/photo-lite.pdf> (2005)
40. Massey, P.: A user’s guide to ccd reductions with iraf. Tech. rep., National Optical Astronomy Observatory (NOAO/NOIRLab) (1997), <https://noirlab.edu/science/sites/default/files/media/archives/documents/scidoc0478.pdf>, nOIRLab document scidoc0478
41. Massey, R., Stoughton, C., Leauthaud, A., Rhodes, J., Koekemoer, A., Ellis, R., Shaghoulain, E.: Pixel-based correction for charge transfer inefficiency in the hubble space telescope advanced camera for surveys. *Monthly Notices of the Royal Astronomical Society* **401**(1), 371–384 (2010)
42. McCully, C., Volgenau, N.H., Harbeck, D.R., Lister, T.A., Saunders, E.S., Turner, M.L., Siiverd, R.J., Bowman, M.: Real-time processing of the imaging data from the network of las cumbres observatory telescopes using banzai. In: *Software and Cyberinfrastructure for Astronomy v. vol. 10707*, pp. 141–149. SPIE (2018)
43. McLean, I.S.: *Electronic imaging in astronomy: detectors and instrumentation*. Springer (2008)
44. Mildenhall, B., Barron, J.T., Chen, J., Sharlet, D., Ng, R., Carroll, R.: Burst denoising with kernel prediction networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2502–2510 (2018)
45. Monakhova, K., Richter, S.R., Waller, L., Koltun, V.: Dancing under the stars: video denoising in starlight. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16241–16251 (2022)
46. Mosleh, A., Zhao, L., Singh, A., Han, J., Punnappurath, A., Brubaker, M.A., Choe, J., Brown, M.S.: Non-parametric sensor noise modeling and synthesis. In: *Proceedings of the European Conference on Computer Vision*. pp. 73–89. Springer (2024)
47. Pan, H., Chen, B., Zhou, R.: Physics-based noise modeling and deep learning for denoising permanently shadowed lunar images. *Applied Sciences* **15**(5) (2025)
48. Pan, Y., Liu, X., Liao, X., Cao, Y., Ren, C.: Random sub-samples generation for self-supervised real image denoising. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 12150–12159 (2023)
49. Pratt, W.K.: *Digital Image Processing*, chap. 4, pp. 247–306. Wiley-Interscience (2007)
50. Project, A.: *Ccd reduction and photometry guide: Understanding an astronomical ccd image*. <https://www.astropy.org/ccd-reduction-and-photometry-guide> (2024), online documentation
51. Qin, J., Qin, P., Chai, R., Qin, J., Jin, Z.: El2nm: Extremely low-light noise modeling through diffusion iteration. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. pp. 1085–1094 (2024)

52. Robberto, M.: On the reference pixel correction of nircam detectors. Tech. Rep. JWST-STScI-003852, SM-12, Space Telescope Science Institute (2014), https://www.stsci.edu/files/live/sites/www/files/home/jwst/documentation/technical-documents/_documents/JWST-STScI-003852.pdf
53. Roddier, F.: V the effects of atmospheric turbulence in optical astronomy. In: Progress in optics, vol. 19, pp. 281–376. Elsevier (1981)
54. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer Assisted Intervention. pp. 234–241. Springer (2015)
55. Seibert, J.A., Boone, J.M., Lindfors, K.K.: Flat-field correction technique for digital detectors. In: Medical Imaging 1998: Physics of Medical Imaging. vol. 3336, pp. 348–354. SPIE (1998)
56. Shi, Y., Liu, D., Zhang, L., Tian, Y., Xia, X., Fu, X.: Zero-ig: Zero-shot illumination-guided joint denoising and adaptive enhancement for low-light images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3015–3024 (2024)
57. Starck, J.L., Murtagh, F.: Astronomical image and data analysis. Springer Science & Business Media (2007)
58. Stoughton, C., Lupton, R.H., Bernardi, M., Blanton, M.R., Burles, S., Castander, F.J., Connolly, A., Eisenstein, D.J., Frieman, J.A., Hennesy, G., et al.: Sloan digital sky survey: early data release. The Astronomical Journal **123**(1), 485 (2002)
59. (STScI), S.T.S.I.: Acs data handbook. <https://hst-docs.stsci.edu/acsdhb/chapter-4-ac-data-processing-considerations/4-4-flat-field-reference-files> (2024), online handbook
60. Tibbs, A.B., Daly, I.M., Roberts, N.W., Bull, D.R.: Denoising imaging polarimetry by adapted bm3d method. Journal of the Optical Society of America A **35**(4), 690–701 (2018)
61. Van Dokkum, P.G.: Cosmic-ray rejection by laplacian edge detection. Publications of the Astronomical Society of the Pacific **113**(789), 1420 (2001)
62. Vojtekova, A., Lieu, M., Valtchanov, I., Altieri, B., Old, L., Chen, Q., Hroch, F.: Learning to denoise astronomical images with u-nets. Monthly Notices of the Royal Astronomical Society **503**(3), 3204–3215 (2021)
63. Walsh, J., Pirzkal, N., Pasquali, A.: Modelling the fringing of the acs ccd detectors. In: HST Calibration Workshop: Hubble after the Installation of the ACS and the NICMOS Cooling System (2003), https://www.stsci.edu/largefiles/hst/HST_overview/documents/calworkshop/workshop2002/CW2002_Papers_for_indexing/walsh.pdf
64. Wei, K., Fu, Y., Yang, J., Huang, H.: A physics-based noise formation model for extreme low-light raw denoising. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2758–2767 (2020)
65. Wei, K., Fu, Y., Zheng, Y., Yang, J.: Physics-based noise modeling for extreme low-light photography. IEEE Transactions on Pattern Analysis and Machine Intelligence **44**(11), 8520–8537 (2021)
66. Widenhorn, R., Blouke, M.M., Weber, A., Rest, A., Bodegom, E.: Temperature dependence of dark current in a ccd. In: Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications III. vol. 4669, pp. 193–201. SPIE (2002)
67. Zhai, C., Shao, M., Nemati, B., Werne, T., Zhou, H., Turyshchev, S.G., Sandhu, J., Hallinan, G., Harding, L.K.: Detection of a faint fast-moving near-earth asteroid using the synthetic tracking technique. The Astrophysical Journal **792**(1), 60 (2014)

68. Zhang, F., Xu, B., Li, Z., Liu, X., Lu, Q., Gao, C., Sang, N.: Towards general low-light raw noise synthesis and modeling. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10820–10830 (2023)
69. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing* **26**(7), 3142–3155 (2017)
70. Zhang, K., Bloom, J.S.: deepcr: Cosmic ray rejection with deep learning. *The Astrophysical Journal* **889**(1), 24 (2020)
71. Zhang, Y., Qin, H., Wang, X., Li, H.: Rethinking noise synthesis and modeling in raw denoising. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4593–4601 (2021)
72. Zhang, Y., Nord, B., Pagul, A., Lepori, M.: Noise2astro: Astronomical image denoising with self-supervised neural networks. *Research Notes of the AAS* **6**(9), 187 (2022)
73. Zou, Y., Yan, C., Fu, Y.: Iterative denoiser and noise estimator for self-supervised image denoising. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13265–13274 (2023)