Physics-Aware Deblurring with Coded Exposure

Pedram Yazdinia and Seo Won Yi

Abstract—Motion blur represents a fundamental information bottleneck in photography. While coded-exposure cameras theoretically preserve high frequencies, practical reconstruction is limited by the "Spectral Gap"—the disconnect between idealized forward models and real-world signal-to-noise ratios. We bridge this gap with an Physics-aware deblurring framework that explicitly couples degradation physics with modern plug-and-play (PnP) reconstruction. We demonstrate that while classical baselines (e.g., ADMM with DRUNet) perform well in standard regimes, they collapse under severe degradation because they blindly amplify noise in spectral nulls. We introduce a frequency-domain "Trust Map" to gate updates based on the Modulation Transfer Function (MTF), maintaining reconstruction quality (≈ 24 dB) where baselines fail (≈ 14 dB). Furthermore, we propose a learnable Unrolled ADMM that predicts context-aware schedules from scene statistics, automating parameter tuning. Our framework runs on CPU, CUDA, or DirectML, providing a reproducible foundation for physics-aware computational photography.

Index Terms—Computational Photography, Coded Exposure, Image Restoration, Plug-and-Play Priors, Algorithm Unrolling, ADMM

1 Introduction

Motion blur acts as a low-pass filter, often introducing deep spectral nulls where spatial frequencies are irretrievably lost. Coded-exposure photography [1] mitigates this by fluttering the shutter to "broaden" the blur kernel and improve invertibility.

In practice, a gap remains. Under real-world low-light conditions, the interaction between the shutter code's Modulation Transfer Function (MTF) and sensor noise gives rise to a "Spectral Gap" in which frequencies that are theoretically preserved may have such low Signal-to-Noise Ratio (SNR) that inversion algorithms, even with deep priors, fail to recover them. This produces severe ringing or oversmoothing. Classical methods like Richardson Lucy struggle in these regions, and Plug-and-Play (PnP) approaches often treat the forward model as a black box, applying fixed schedules regardless of optical quality.

We argue that reconstruction algorithms must be *physics-aware*. Our framework explicitly conditions the optimization trajectory on the spectral characteristics of the degradation. We contribute:

1) **The "Spectral Gap" Insight:** We show that PnP methods fail in high-noise settings because they cannot "distrust" the data term in low-MTF bands. We introduce a frequency-domain **Trust Map** that gates updates in these regions, enforcing reliance on the prior.
2) **Adaptive Physics Scheduling:** We design a dynamic ADMM scheduler that adjusts the penalty parameter (ρ) and denoiser weight based on an "Optics Score" derived from the MTF, enabling aggressive updates for strong optics and conservative ones for weak optics, without manual tuning.
3) **Learnable Unrolled ADMM:** We implement an unrolled ADMM architecture with a lightweight context MLP that analyzes scene statistics (e.g., blur length, photon budget, MTF profile) to predict iteration schedules and spectral masks, achieving strong performance without hand-tuning.

Empirically, our physics-aware approach improves PSNR by ≈ 10 dB over strong baselines in severe conditions and matches or slightly exceeds them under standard settings. We release the full framework as open source, with a modular design to support future research at the intersection of optics and optimization.

2 Background and Related Work

2.1 Motion Deblurring and Coded Exposure

Motion blur renders blind deconvolution ill-posed [2]. To mitigate this, Raskar et al. [1] introduced *Coded Exposure*, modulating the shutter to preserve high-frequency information. Agrawal et al. [3] further demonstrated that code invertibility is paramount for noise robustness.

Among binary codes, Legendre sequences [4] are particularly effective. Defined by quadratic residues modulo a prime p, they exhibit a nearly flat power spectrum and peaky autocorrelation, minimizing noise amplification during inversion. Our work adopts Legendre codes as the high-quality reference, comparing them against random and box codes.

2.2 Deep Learning and Plug-and-Play Priors

Deep learning has revolutionized restoration, with models like DnCNN [5] and DRUNet [6] serving as powerful priors. Plug-and-Play (PnP) methods [7], [8] use such denoisers as implicit regularizers within iterative schemes such as ADMM [9]. PnP-ADMM allows these non-convex priors to be decoupled from the convex data fidelity term and has been successfully applied in domains such as MRI and remote sensing [7], [8], where partial observations create structured spectral nulls. Standard photographic reconstruction pipelines, however, often ignore the spectral characteristics of the forward model. Our work addresses this by explicitly conditioning the optimization schedule on the MTF.

2.3 Algorithm Unrolling

Algorithm unrolling [10] interprets optimization iterations as neural network layers. Diamond et al. [11] showed un-

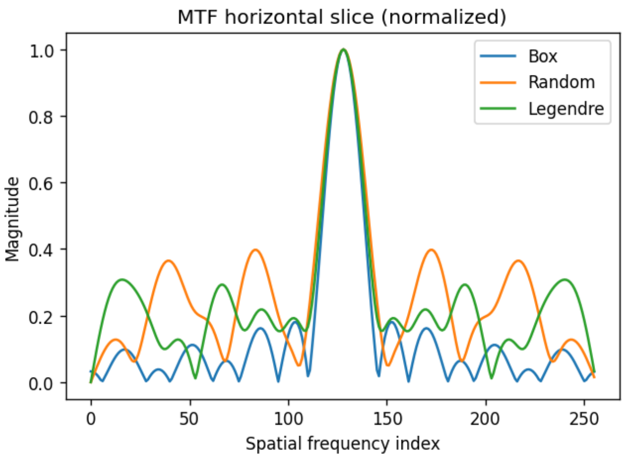

Fig. 1. Normalized MTF horizontal slice for Box, Random, and Legendre codes. Note the spectral zeros in the Box code (blue) which destroy information at specific frequencies. In contrast, the Random (orange) and Legendre (green) codes maintain broadband spectral coverage, preserving invertibility.

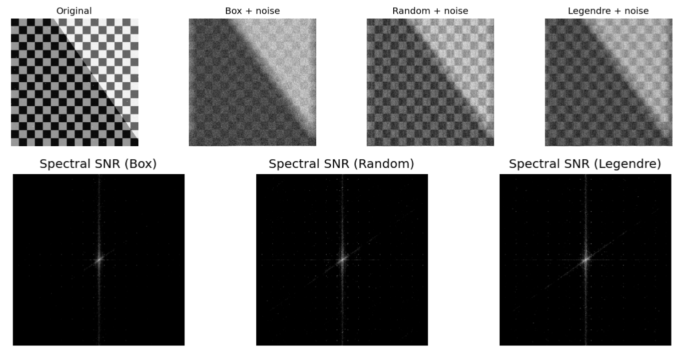

Fig. 2. Synthetic noisy and blurred observations for Box, Random, and Legendre codes. Despite the added noise, the coded exposures preserve more texture and high-frequency detail. The spectral SNR differences are present in the diagrams, but they are subtle; we recommend zooming in to observe them more clearly.

rolled methods could outperform generic CNNs by incorporating physical models. Our Unrolled ADMM extends this by learning not just scalar parameters, but also spectral trust masks adapted to specific degradation instances.

3 Methodology

Our framework introduces a "Physics-Aware PnP-ADMM" algorithm that explicitly couples the degradation physics with the reconstruction process. The pipeline consists of two main branches: a Physics-Aware Branch that pre-computes degradation statistics, and an Iterative PnP-ADMM Loop that solves the inverse problem using these statistics.

3.1 Forward Model

We model coded exposure as a linear convolution with additive noise, $y = k * x + n$. Our simulator generates exposure codes (Box, Random, Legendre) with configurable taps T, blur length (pixels), and duty cycle. It converts these codes to Point Spread Functions (PSFs) via a motion kernel and computes key diagnostics such as the Optical Transfer Function (OTF) and Modulation Transfer Function (MTF).

As shown in Figure 1, the choice of exposure code dramatically affects the conditioning of the inverse problem. The standard "Box" shutter (blue curve) exhibits characteristic spectral zeros where the MTF drops to zero. At these frequencies, signal information is irretrievably lost, making deconvolution ill-posed and sensitive to noise. In contrast, the "Random" (orange) and "Legendre" (green) codes are designed to be broadband, maintaining significant magnitude across the spectrum. This "spectral preservation" is the physical foundation that allows our method to recover high-frequency details that are otherwise lost.

To simulate realistic capture conditions, we inject Poisson–Gaussian noise parameterized by a fixed photon budget and read-noise standard deviation σ. Figure 2 visualizes these noisy observations, highlighting how the coded exposures retain more texture information than the box blur.

The pipeline supports both RGB and grayscale inputs from the DIV2K dataset, as well as synthetic scenes for validation.

3.2 Physics-Aware Branch and Trust Weights

The core of our method is the recognition that not all data is equally reliable. In the frequency domain, the utility of the observed signal $Y(\omega)$ is governed by the local SNR, which is strongly modulated by the MTF $|K(\omega)|$.

We construct a frequency-domain **Trust Map** $W(\omega)$ that acts as a spectral gate. It answers the question: *"Should we trust the observation at this frequency, or should we let the prior dominate?"* In practice, we combine two components:

- A **Wiener-like weight** $W_{\text{Wiener}}(\omega)$ derived from $|K(\omega)|^2$ and a scalar parameter τ inferred from the noisy observation.
- A **gamma-weighted MTF mask** that reshapes the normalized MTF and enforces both a hard cutoff and a soft floor.

We first compute a normalized MTF:

$$\widetilde{\text{MTF}}(\omega) = \frac{|K(\omega)|}{\max_{\omega'} |K(\omega')|} \in [0, 1]. \tag{1}$$

From this, we form a gamma mask

$$W_{\gamma}(\omega) = \widetilde{\text{MTF}}(\omega)^{\gamma}, \tag{2}$$

where γ controls the suppression of low-MTF frequencies. We apply two additional operations: a **hard cutoff** to zero out hopeless frequencies, and a **soft floor** to retain moderately weak bands preventing them from being completely ignored.

In parallel, we compute a Wiener-like weight $W_{\text{Wiener}}(\omega) = f(|K(\omega)|^2; \tau, \alpha)$, where α controls the trade-off between fidelity and robustness. In the **combined** mode used for our main results, the final Trust Map is:

$$W(\omega) \propto W_{\text{Wiener}}(\omega) \cdot W_{\gamma}(\omega), \tag{3}$$

followed by a global normalization to $[0, 1]$. The ADMM x-update then becomes a weighted frequency-domain least squares problem:

$$x^{(k+1)} = \mathcal{F}^{-1}\left(\frac{W \cdot K^* \cdot \mathcal{F}(y) + \rho\, \mathcal{F}(z^{(k)} - u^{(k)})}{W \cdot |K|^2 + \rho} \right). \tag{4}$$

When $W(\omega) \approx 0$ in spectral nulls, the data term vanishes and the update simplifies to $x^{(k+1)} \approx z^{(k)} - u^{(k)}$. In these regions, the solution is driven entirely by the prior.

3.3 Adaptive Physics Scheduler

Standard PnP-ADMM typically uses a fixed penalty parameter ρ, which assumes a stable relationship between the data term and the prior. This is suboptimal for motion deblurring, where the difficulty of the inverse problem can change dramatically with the exposure code and noise level.

We introduce an **Adaptive Physics Scheduler** that modulates the optimization trajectory based on simple optics and noise scores derived from the MTF and photon budget. Given a precomputed MTF map, we define an optics quality score

$$S_{\text{optics}} = \text{clip}\big(3 \cdot \text{mean}(\text{MTF}),\, 0,\, 1\big), \tag{5}$$

where the mean is taken over spatial frequencies. Well-conditioned optics produce large MTF values, hence S_{optics} near 1; box-like or diffuser-like patterns yield smaller values.

We then define a deep-dive start multiplier and a conservative end multiplier for the ADMM penalty. For good optics, we would ideally like to start at a very low penalty (around $0.02\rho_{\text{base}}$) to aggressively recover fine texture, and then ramp up to approximately $1.2\rho_{\text{base}}$ to lock in the prior. For poor optics, however, such a low starting value can cause instabilities. In practice we use

$$m_{\text{start}} = \begin{cases} 0.02, & S_{\text{optics}} \geq 0.15, \\ 0.15, & S_{\text{optics}} < 0.15, \end{cases} \qquad m_{\text{end}} = 1.2, \tag{6}$$

and ramp geometrically between them. At iteration k of K total steps, the scheduler uses

$$\rho_k = \rho_{\text{base}} \cdot m_{\text{start}} \left(\frac{m_{\text{end}}}{m_{\text{start}}}\right)^{\frac{k}{K-1}}. \tag{7}$$

This realizes the desired coarse-to-fine behavior: initially, the solver explores the solution space with a weak prior; later iterations enforce the prior more strongly.

The denoiser weight is also modulated by the optics score. We start from a base weight w_{base} and increase it for poor optics:

$$w_k = \text{clip}\Big(w_{\text{base}} \cdot \big(1 + 0.2(1 - S_{\text{optics}})\big),\, 0,\, 1\Big). \tag{8}$$

When the optics are good, w_k is close to w_{base} and the data term has more influence. When the optics are poor, the scheduler increases w_k so that the denoiser has greater control, which helps prevent catastrophic noise amplification. In our implementation the scheduler also exposes diagnostic quantities (such as the current multiplier and optics score) for analysis, but the core behavior is captured by the formulas above.

3.4 Learnable Unrolled ADMM

To reduce the reliance on manual hyperparameters, we introduce a learnable unrolled ADMM variant. This model unrolls K iterations of the PnP-ADMM solver into a neural module and learns physics-aware parameters directly from data, while reusing a fixed denoiser backbone.

At each iteration k, the model utilizes three scalar parameters: the ADMM penalty ρ_k, a denoiser blending weight w_k, and a sigma multiplier α_k. These are initialized from global values and refined by a lightweight **Context Network**. A feature vector containing blur and noise statistics (e.g., blur length, photon budget, MTF quality) is passed through a small MLP, which predicts per-step deltas for the parameters. When enabled, the network also predicts spectral mask parameters (gamma, floor, cutoff).

We train the unrolled model end-to-end on DIV2K, with dataset details provided in a later section. The denoiser backbone is kept frozen; thus, the network learns a *policy* for steering the optimization based on physics, rather than learning a new image reconstruction prior from scratch.

4 Software Framework and Reproducibility

To facilitate rigorous comparison and future research, we developed a modular, extensible software framework that serves as a sandbox for computational photography. The repository is architected to decouple the forward model (physics simulation) from the inverse solvers (reconstruction algorithms), allowing researchers to mix and match components seamlessly.

4.1 Modular Design

The codebase is organized into three primary modules:

- **Forward Pipeline:** A physics-faithful simulator that synthesizes motion blur, generates exposure codes (Box, Random, Legendre), computes PSF/OTF/MTF diagnostics, and injects realistic Poisson–Gaussian noise. It supports both synthetic scenes (checkerboards, rings) and real datasets (DIV2K).
- **Reconstruction Engines:** A collection of inverse solvers ranging from classical baselines to deep learning methods. This includes a flexible PnP-ADMM engine that supports custom schedules, a gradient-descent PnP-ADAM solver, and the learnable Unrolled ADMM. All solvers share a common interface for easy benchmarking.
- **Denoiser Backend:** An abstract interface for deep denoising priors, supporting TinyCNN (CPU-optimized), DnCNN, UNet, and DRUNet. This allows for easy integration of new architectures without modifying the solver logic.

4.2 Configuration and Parameters

The framework exposes a comprehensive set of parameters to fine-tune the simulation and reconstruction process. Users can control:

- **Physics Parameters:** Blur length, number of taps (T), duty cycle, photon budget, and read noise standard deviation (σ).
- **Solver Hyperparameters:** ADMM penalty (ρ), denoiser weight, number of iterations, and specific "physics hooks" like MTF weighting modes (Gamma, Wiener, Combined) and Sigma Adaptation.
- **Hardware Acceleration:** Seamless switching between CPU, CUDA, and DirectML backends for broad hardware compatibility.

4.3 CLI Usage Example

To facilitate reproducibility and experimentation, the framework provides a rich Command Line Interface (CLI). Below is an example command for running the Physics-Aware ADMM pipeline on the DIV2K dataset with specific noise and blur settings:

```
python -m \
  mtf_aware_deblurring.pipelines.reconstruct \
  --div2k-root data --subset train \
  --degradation bicubic --scale X2 \
  --image-mode rgb --limit 10 \
  --patterns box random legendre \
  --method admm \
  --admm-iters 60 \
  --admm-rho 1.04 \
  --admm-denoiser-weight 0.16 \
  --admm-mtf-weighting-mode combined \
  --admm-mtf-sigma-adapt \
  --use-physics-scheduler \
  --denoiser-type drunet_color \
  --photon-budget 1000 \
  --blur-length 15 \
  --read-noise 0.01 \
  --output-dir experiments/results
```

This command automatically handles data loading, forward simulation, reconstruction across multiple codes, and metric calculation (PSNR, SSIM, LPIPS), saving all artifacts to the specified output directory.

5 Experimental Setup

5.1 Datasets and Metrics

We utilize the **DIV2K** (Diverse 2K Resolution) dataset, a standard benchmark for image restoration tasks. The dataset consists of 800 high-resolution training images and 100 validation images.

- **Training:** For the learnable Unrolled ADMM, we train on the first 800 images using random 256×256 crops to ensure robustness to varied textures.
- **Evaluation:** All quantitative results (Table 1 and 2) are reported on the 100-image validation set. We use center crops of 256×256 to ensure consistent and reproducible benchmarking across all methods.

We report three key metrics:

- **PSNR (Peak Signal-to-Noise Ratio):** Measures pixel-wise fidelity.
- **SSIM (Structural Similarity Index):** Evaluates perceptual structural quality.
- **LPIPS (Learned Perceptual Image Patch Similarity):** Assesses high-level perceptual similarity using deep features, correlating better with human judgment for texture realism.

5.2 Noise Regimes

To test robustness, we simulate two distinct noise regimes using a Poisson–Gaussian model. The read noise standard deviation σ is applied relative to the normalized intensity range $[0, 1]$. Given our fixed photon budget of 1000 photons per pixel, these normalized values correspond to significant physical noise floors:

- **Noise1 (Standard):** Represents a typical low-light capture with a photon budget of 1000 and normalized read noise $\sigma = 0.01$. This corresponds to an effective read noise of ≈ 10 photons, typical of modern CMOS sensors. In this regime, the signal is weak but recoverable.
- **Noise2 (Severe):** Represents an extreme scenario (e.g., high analog gain). We maintain the photon budget at 1000 but increase normalized read noise to $\sigma = 0.05$, corresponding to an effective read noise of ≈ 50 photons. Combined with an increased Blur Length of 30 pixels, the SNR in the spectral nulls drops precipitously, rigorously testing the algorithm's ability to reject noise.

5.3 Baselines

We compare against:

- **Richardson–Lucy:** A classical iterative baseline.
- **PnP ADAM:** A gradient-descent PnP method with fixed steps.
- **Vanilla ADMM (DRUNet):** A standard ADMM implementation using the same strong DRUNet prior but with a fixed ρ and standard least-squares inversion (no Trust Map). This isolates the contribution of our physics-aware components.

6 Results

6.1 Standard Regime (Noise1)

In the standard noise regime (Blur 15px, $\sigma = 0.01$), the signal is degraded but not destroyed. Here, a strong prior can do the heavy lifting. Figure 3 provides a visual comparison, showing that our method reduces artifacts compared to the baseline, but results in a slightly oversmoothed image. As shown in Table 1, the "Vanilla" ADMM baseline with DRUNet achieves $\approx$ 26–28 dB, effectively matching our Physics-Aware method.

This result is crucial as it underscores that when the SNR is high enough, the "Spectral Gap" is narrow. The deep prior, trained on vast amounts of natural images, can successfully reconstruct the missing high frequencies even without explicit physics guidance. In this "easy" regime, the physics-aware weighting provides only marginal gains (≈ 0.3 dB), primarily by preventing minor artifacts.

TABLE 1
Performance Comparison - Standard Regime (Noise1)

Method	Box	Random	Legendre
Richardson–Lucy	19.01	19.00	19.02
PnP ADAM (Tiny)	22.16	22.78	22.69
Vanilla ADMM (DRUNet)	26.84	28.23	28.11
Physics-Aware ADMM	**27.10**	**28.50**	**29.55**
Unrolled ADMM	26.34	27.19	27.15

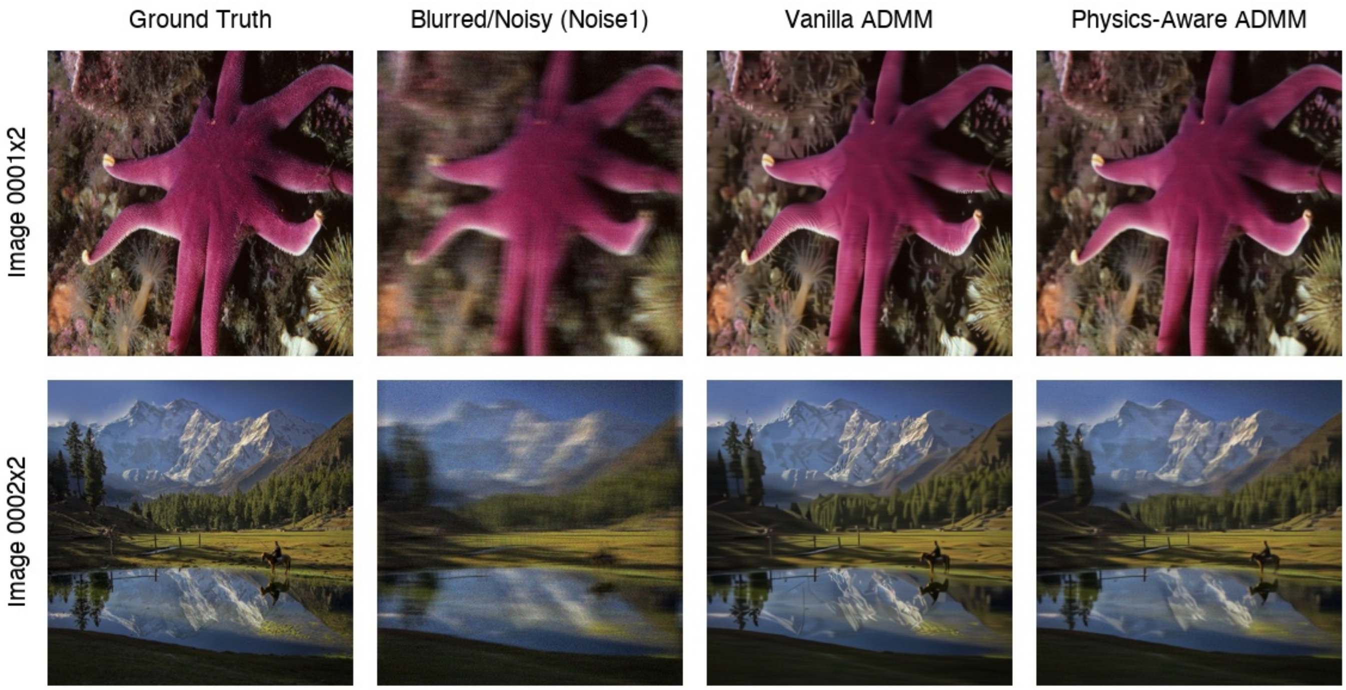

Fig. 3. Visual comparison for the Standard Regime (Noise1). From left to right: Ground Truth, Noisy Input, Vanilla ADMM, and Physics-Aware ADMM. Our method (right) produces fewer artifacts, although it results in a slightly oversmoothed image compared to the baseline.

6.2 Severe Regime (Noise2)

The performance divergence is most pronounced in the severe regime (Blur 30px, $\sigma = 0.05$). Figure 4 illustrates this: the Vanilla ADMM baseline produces output often worse than the input, while our Physics-Aware ADMM successfully reconstructs the image. As shown in Table 2, the baseline collapses to ≈ 10–14 dB.

This failure mode is instructive. The Vanilla ADMM blindly trusts the data term equally across all frequencies. In the spectral nulls of the blur kernel, the signal is negligible, yet the noise remains. The solver attempts to fit this noise, resulting in severe amplification and ringing. Our Physics-Aware method, protected by the Trust Map, effectively identifies these frequencies as noise-dominated and suppresses them, relying entirely on the prior. This empirically validates our central thesis: *robust deblurring requires explicit physics awareness.*

TABLE 2
Performance Comparison - Severe Regime (Noise2)

Method	Box	Random	Legendre
Richardson–Lucy	18.35	18.63	18.63
PnP ADAM (DnCNN)	12.45	17.05	17.08
Vanilla ADMM (DRUNet)	10.43	13.48	13.57
Physics-Aware ADMM	20.79	**23.59**	**24.04**
Unrolled ADMM	18.52	20.16	20.44

Table 2 also highlights the performance of the Unrolled ADMM. Despite lacking the heavy-weight DRUNet prior, it achieves approximately 20 dB and outperforms the Vanilla baseline by roughly 6 dB without requiring any manual tuning. This result indicates that under heavy degradation, explicit physics guidance, whether provided through hand-crafted schedules or through learned parameters, is more important than the raw representational power of the prior.

6.3 Unrolled ADMM: The Automation Trade-off

The Unrolled ADMM offers a compelling middle ground. While it trails the heavy-weight Physics-Aware ADMM (with DRUNet) by ≈ 2–4 dB, it outperforms standard baselines and, crucially, requires *zero manual tuning*. The learned context MLP successfully predicts effective schedules from the raw image statistics. The performance gap is largely due to the lighter backbone (TinyDenoiser vs. DRUNet) required for end-to-end training. This highlights a classic trade-off: the Unrolled model buys automation and speed at the cost of peak restoration quality.

6.4 Negative Results (Ablation Study)

We document alternative strategies that yielded negative results to aid future research:

Relay Denoising: We hypothesized that initializing the optimization with a lightweight denoiser (TinyCNN) and switching to a heavy prior (DRUNet) mid-stream would save computation. Experiments revealed negligible gains (± 0.1 dB) compared to using DRUNet exclusively, despite increased implementation complexity.

Spatial Sigma Maps: We attempted to feed a spatially varying noise map (derived from $\sqrt{I}$) into DRUNet to handle signal-dependent noise. This caused a performance collapse (-5 dB). We attribute this to a distribution shift: DRUNet is trained on uniform AWGN, and injecting structured sigma maps likely pushes the input out of the training distribution. This validates our choice of a scalar, optics-guided sigma multiplier.

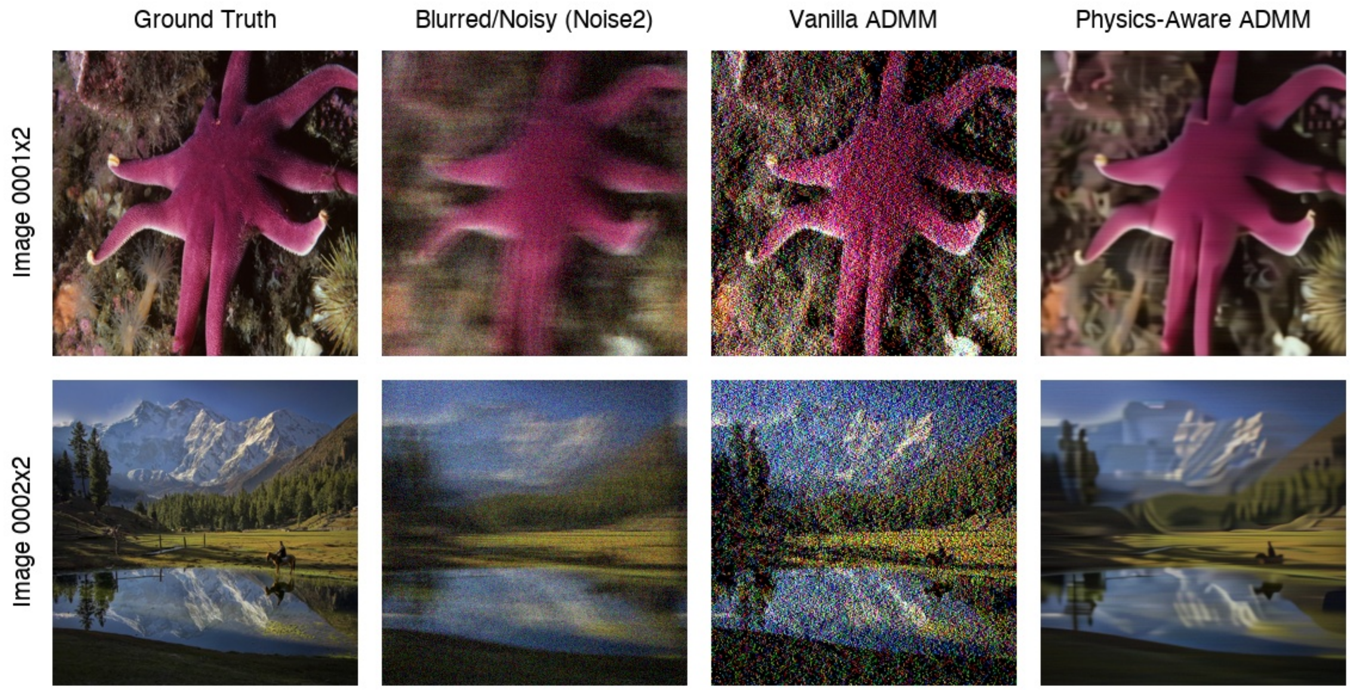

Fig. 4. Visual comparison for the Severe Regime (Noise2). The Vanilla ADMM baseline (third column) amplifies noise while trying to solve the deblur problem, while the Physics-Aware ADMM (right) successfully reconstructs the image, demonstrating robustness to heavy noise and blur.

7 Conclusion

We presented an Physics-aware deblurring framework that bridges the gap between clean theory and noisy reality. By explicitly modeling the "Spectral Gap" and introducing a frequency-domain Trust Map, we enable robust reconstruction in regimes where standard methods fail. Our results demonstrate that while deep learning is powerful, it is not a panacea; in the presence of severe physical degradation, the algorithm must be guided by the physics of the forward model.

Looking forward, we see significant potential in integrating *Diffusion Models* (score-based generative models) as priors within our physics-aware framework. Although the current DRUNet backbone is efficient, diffusion models provide stronger capabilities for texture hallucination and could help close the spectral gap in extreme low-light conditions. Their high inference cost remains a practical barrier, which suggests a promising direction for future work: distilling these powerful priors into lightweight and unrolled architectures similar to the one presented here. We hope this work, together with the accompanying open-source toolkit, encourages the community to move beyond black-box restoration and adopt physics-aware design principles.

References

[1] R. Raskar, A. Agrawal, and J. Tumblin, "Coded exposure photography: motion deblurring using fluttered shutter," *ACM Transactions on Graphics (TOG)*, vol. 25, no. 3, pp. 795–804, 2006.

[2] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman, "Understanding and evaluating blind deconvolution algorithms," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 1964–1971.

[3] A. Agrawal and Y. Xu, "Optimal single image capture for motion deblurring," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 2560–2567.

[4] H.-G. Jeon, J.-Y. Lee, Y. Han, S. Kim, and I. S. Kweon, "Fluttering pattern generation using modified legendre sequence for coded exposure," in *IEEE International Conference on Computational Photography (ICCP)*, 2013, pp. 1–8.

[5] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, 2017.

[6] K. Zhang, Y. Li, W. Zuo, L. Zhang, L. Van Gool, and R. Timofte, "Plug-and-play image restoration with deep denoiser prior," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 10, pp. 6360–6376, 2021.

[7] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *IEEE Global Conference on Signal and Information Processing*, 2013, pp. 945–948.

[8] S. H. Chan, X. Wang, and O. A. Elgendy, "Plug-and-play admm for image restoration: Fixed-point convergence and applications," *IEEE Transactions on Computational Imaging*, vol. 3, no. 1, pp. 84–98, 2016.

[9] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine learning*, vol. 3, no. 1, pp. 1–122, 2011.

[10] V. Monga, Y. Li, and Y. C. Eldar, "Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing," *IEEE Signal Processing Magazine*, vol. 38, no. 2, pp. 18–44, 2021.

[11] S. Diamond, V. Sitzmann, F. Heide, and G. Wetzstein, "Unrolled optimization with deep priors," *arXiv preprint arXiv:1705.08041*, 2017.