DORIAN: ZERO-SHOT DIRECTIONALLY STRUCTURED NOISE REMOVAL VIA GRAPH NEURAL PRIORS IN FOURIER SPACE

Anonymous authors
Paper under double-blind review

ABSTRACT

Various scientific imaging modalities, for instance, light-sheet fluorescence microscopy, atomic force microscopy, and remote sensing, suffer from degradations that persist along a dominant orientation over image-scale distances. We refer to them as Directionally Persistent Structured Noise (DPSN). Three characteristics render DPSN difficult to remove: it is non-local in the image domain, its spectrum overlaps with that of the underlying image, and its appearance varies across modalities in which clean references are rarely available. Motivated by its compact representation in the Fourier domain, we propose DORIAN, a zero-shot framework that poses DPSN removal as a Fourier-domain spectral reconstruction problem. Specifically, we construct a polar graph over Fourier coefficients, on which a graph neural prior reconstructs corrupted coefficients from spectrally related uncorrupted ones. The graph prior is optimized on each degraded observation alone, requiring neither external training data nor clean targets. Experiments on natural images, microscopy, and remote sensing data show that DORIAN removes DPSN of diverse physical origins and appearances within a single framework.

1 INTRODUCTION

In this work, we consider degradations whose intensity remains approximately constant, or varies only slowly, along a dominant orientation over an extended distance. We refer to them as Directionally Persistent Structured Noise (DPSN). Specifically, let $Y \in \mathbb{R}^{N_x \times N_y}$, $X \in \mathbb{R}^{N_x \times N_y}$, and $N_s \in \mathbb{R}^{N_x \times N_y}$ denote the degraded image, the clean image, and the DPSN component, respectively. The directional persistence of N_s is characterized as:

$$Y = X + N_s, \qquad \begin{cases} \left|N_s(\mathbf{p} + t\mathbf{d}_{\parallel}) - N_s(\mathbf{p})\right| \le t\,\epsilon_{\parallel}, & 0 \le t \le L, \\ \left|N_s(\mathbf{p} + s\mathbf{d}_{\perp}) - N_s(\mathbf{p})\right| \le s\,\epsilon_{\perp}, & [\mathbf{p}, \mathbf{p} + s\mathbf{d}_{\perp}] \cap \Gamma = \emptyset, \end{cases} \tag{1}$$

where $\mathbf{p}$ is a location at which DPSN is present, $\mathbf{d}_{\parallel}$ and $\mathbf{d}_{\perp}$ are its dominant persistence direction and the orthogonal direction, $\epsilon_{\parallel}$ and $\epsilon_{\perp}$ are small variation rates along $\mathbf{d}_{\parallel}$ and within a DPSN component, along $\mathbf{d}_{\perp}$, respectively, L is the persistence length, and Γ denotes the DPSN boundaries. The morphology of DPSN is thus characterized in two respects: along $\mathbf{d}_{\parallel}$, N_s varies slowly over the entire length L; across it, N_s varies slowly except at Γ, where it may change abruptly.

Degradations of this form arise across imaging systems (Fig. 1(b)), including light-sheet fluorescence microscopy (LSFM) (Ricci et al., 2022), atomic force microscopy (AFM) (Canale et al., 2011), and remote sensing (Chang et al., 2020). Although their physical effects may be additive or multiplicative, both can be represented by Eq. (1), with multiplicative degradations expressed in the log-intensity domain. Whereas independent noise is routinely handled by denoisers that transfer across modalities (Ma et al.; Lu et al.), DPSN has largely been treated modality by modality, since its appearance depends on the underlying physical mechanism and ranges from periodic to non-periodic and from globally distributed to spatially localized (Liu et al., 2022). Yet these manifestations share a common Fourier signature: compared with noise of weaker spatial correlation, DPSN concentrates in a directional frequency band (Fig. 1(a)) across all modalities (Fig. 1(b)), which motivates a unified restoration model.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

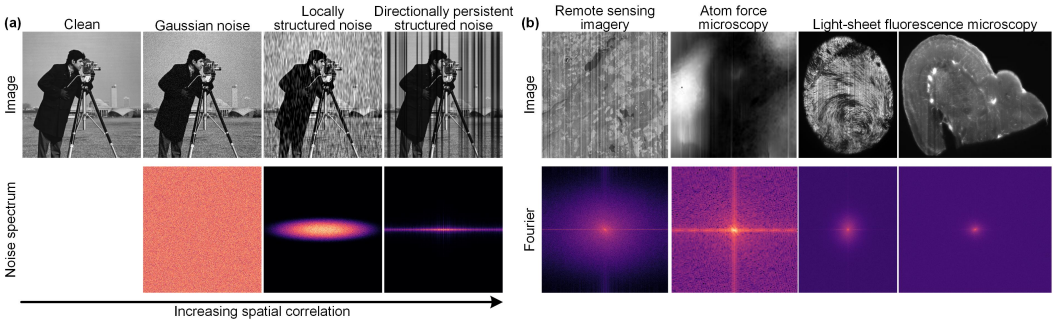


Figure 1: **Directionally Persistent Structured Noise (DPSN).** (a) Degradations of increasing spatial correlation and their noise spectra (bottom). (b) DPSN in remote sensing, AFM, and LSFM, with the Fourier magnitude of each observation (bottom).

A unified DPSN removal model must address three difficulties simultaneously. *(i) Non-local spatial support.* A single DPSN component extends over a large portion of the image, and its coherence is expressed only through long-range dependencies among large numbers of pixels. Restoration models defined at the pixel level act on individual pixels or bounded neighborhoods, and therefore capture it only in fragments. *(ii) Spectral overlap.* Although concentrated in orientation, the spectral energy of DPSN is not confined to isolated frequencies: its abrupt boundaries spread the energy along the band, and its finite extent broadens the band itself, so that it overlaps genuine image content. Coefficients therefore cannot be classified as corrupted or clean by their spectral position alone. *(iii) Cross-modality variability without references.* The orientation, extent, and intensity profile of DPSN vary substantially between modalities, while clean references are rarely obtainable in any of them.

To address these difficulties, we propose DORIAN, which exploits the spectral compact of DPSN to represent its non-local structure in the Fourier domain, addressing difficulty (*i*). DORIAN then reconstructs DPSN-dominated coefficients from spectrally related, comparatively reliable ones using a graph neural prior on a polar Fourier representation, recovering image content overlapping the DPSN band rather than attenuating it and thereby addressing difficulty (*ii*). Finally, the graph prior is optimized independently on each degraded observation, requiring neither clean targets nor external training data and thus addressing difficulty (*iii*). The main contributions are as follows:

- We recast DPSN removal as a spectral reconstruction problem, and introduce DORIAN, a graph neural prior on a polar Fourier representation that recovers DPSN-corrupted Fourier coefficients by aggregating information from radially matched reliable ones with observation-adaptive weights, optimized zero-shot on each degraded observation.
- We additionally mitigate the major limitation of zero-shot DPSN removal, the indistinguishability between DPSN and genuine structures extending along the same direction, through anisotropic resampling and residual transfer.
- DORIAN connects the two classical paradigms of DPSN removal, recovering spectral filtering as a degenerate case and image-domain variational restoration as a non-adaptive instance (see Appendix B).
- We evaluate DORIAN on synthetic DPSN of varying heterogeneity and on real data from three imaging modalities, where it matches or exceeds existing methods in nearly all settings with a single fixed configuration across all experiments.

2 Related Works

Spectral Filtering-Based Restoration Spectral filtering exploits the directional concentration of DPSN in the Fourier domain by attenuating predefined frequency bands, either directly, as in wavelet–FFT filtering (Münch et al., 2009), or in a transformed domain such as the contourlet (Liang et al., 2016). Although training-free and computationally efficient, these methods assume that degradation and image content are spectrally separable (Chang et al., 2013; Khalilian-Gourtani et al., 2019), thus failing on difficulty (*ii*). Consequently, wide-band suppression removes image structure, whereas narrow-band filtering leaves residual DPSN.

Optimization-Based Restoration Optimization-based methods formulate DPSN removal as an inverse problem in the image domain via an energy function that combines data fidelity with handcrafted regularization (Fehrenbach et al., 2012; Rottmayer et al., 2025). On the regularization side, the low-rank prior used in Chang et al. (2016) holds only when the degradation spans the entire field of view, while the anisotropic total variation (TV) prior (Rottmayer et al., 2025), although exploiting the directionality of DPSN, fails to capture the long-range persistence of DPSN due to its inherently local receptive field. Pixel-wise ℓ_1 or ℓ_2 fidelity (Chang et al., 2013; Yang et al., 2020) encourages the solution to reproduce the degraded intensities, which can be partially mitigated via adaptive or structure-aware fidelity terms, although their integration into a well-behaved energy function remains challenging (Naganuma & Ono, 2022). As a result, these locally defined energy functions are limited in modeling the long-range persistence of DPSN, corresponding to difficulty (*i*).

Learning-Based Parameterizations of Restoration Models Beyond handcrafted regularization, restoration models can be learned from data (Zamir et al., 2022; Chung et al., 2023; Cui et al., 2025). For DPSN, the main distinction lies in the supervision they require. For example, supervised methods learn degraded-to-clean mappings (Chang et al., 2020), but because clean references are rarely available in scientific imaging, training pairs are typically synthesized under assumed degradation models, limiting generalization when the physical corruption differs and thus failing on difficulty (*iii*).

Self-supervised methods avoid clean targets by exploiting assumptions about the noise. For instance, structured blind-spot methods (Broaddus et al., 2020; Qu et al., 2025) assume that the noise correlation is confined to a neighborhood small enough to be excluded from prediction. DPSN, however, spans a large portion of the image and cannot be masked in this way, failing on difficulty (*i*). Unpaired adversarial methods (Yang et al., 2024) likewise avoid paired data but still require a collection of DPSN-contaminated images, remaining tied to their training distribution and failing on difficulty (*iii*).

Zero-shot methods require only the degraded observation. S2DIP (Luo et al., 2021), following Deep Image Prior (Ulyanov et al., 2020), fits a network to a single hyperspectral observation with spatial-spectral TV, although its dependence on the spectral dimension limits transfer to single-channel modalities such as LSFM and AFM. Most closely related, DeStripe (Liu et al., 2022) reconstructs DPSN-affected Fourier coefficients using a graph-structured spectra representation. DORIAN differs substantially in its corruption modeling, graph construction, spectral readout, restoration objective, and structure-preserving recovery. A detailed comparison is provided in Appendix C.

3 METHODS

3.1 MASK-GUIDED FOURIER-DOMAIN SPECTRAL RESTORATION

As DORIAN in Fig. 2, we first transform the degradation model in Eq. (1) into the Fourier domain as $\mathcal{Y} = \mathcal{X} + \mathcal{N}_s$ by applying the centered 2D discrete Fourier transform (zero frequency shifted to the center, denoted as $\mathcal{F}(\cdot)$ with inverse $\mathcal{F}^{-1}(\cdot)$), where calligraphic letters denote the spectra of the corresponding signals. Inspired by the observation that the spectral energy of the DPSN is highly concentrated within a directional wedge in the Fourier domain, we first identify a DPSN-dominated spectral region Ω_s within which coefficients are most likely affected by DPSN. Specifically, for a coefficient i at frequency $\boldsymbol{\xi}_i = (u_i, v_i)$, its polar coordinate is defined as $(\rho_i = \|\boldsymbol{\xi}_i\|_2, \phi_i = \operatorname{atan2}(v_i, u_i))$. We then define the DPSN-dominated wedge and the spectral guidance mask as:

$$\Omega_s = \Big\{ i \;\Big|\; \min_{k\in\mathbb{Z}} |\phi_i - \phi_s + k\pi| \le \Delta_\phi \Big\} \cup \big\{ i \;\big|\; 0 < \rho_i \le \rho_0 \big\}, \qquad S_i = \mathbb{1}[\, i \in \Omega_s \,] \tag{2}$$

where ϕ_s denotes the dominant spectral orientation of the DPSN, and Δ_ϕ controls the angular half-width of the wedge. A low-frequency disk of radius $\rho_0 = 32$, irrespective of orientation, is also included, as DPSN components that are wide or vary slowly along $\mathbf{d}_\perp$ may concentrate their energy near DC, where an oriented wedge is too narrow to contain it.

Next, instead of attenuating the Fourier coefficients selected by S, DORIAN treats them as unreliable observations whose clean values are reconstructed from the remaining, comparatively reliable coefficients, which stay anchored to the observed spectrum. This avoids the irreversible loss of genuine image information that may overlap with the DPSN-dominated region. As a result, once the restored spectrum has been constructed, the restored image is obtained as $\widehat{X} = \mathcal{F}^{-1}(\widehat{\mathcal{X}})$, where $\widehat{\mathcal{X}}$ and $\widehat{X}$ denote the restored spectrum and the corresponding restored image, respectively.

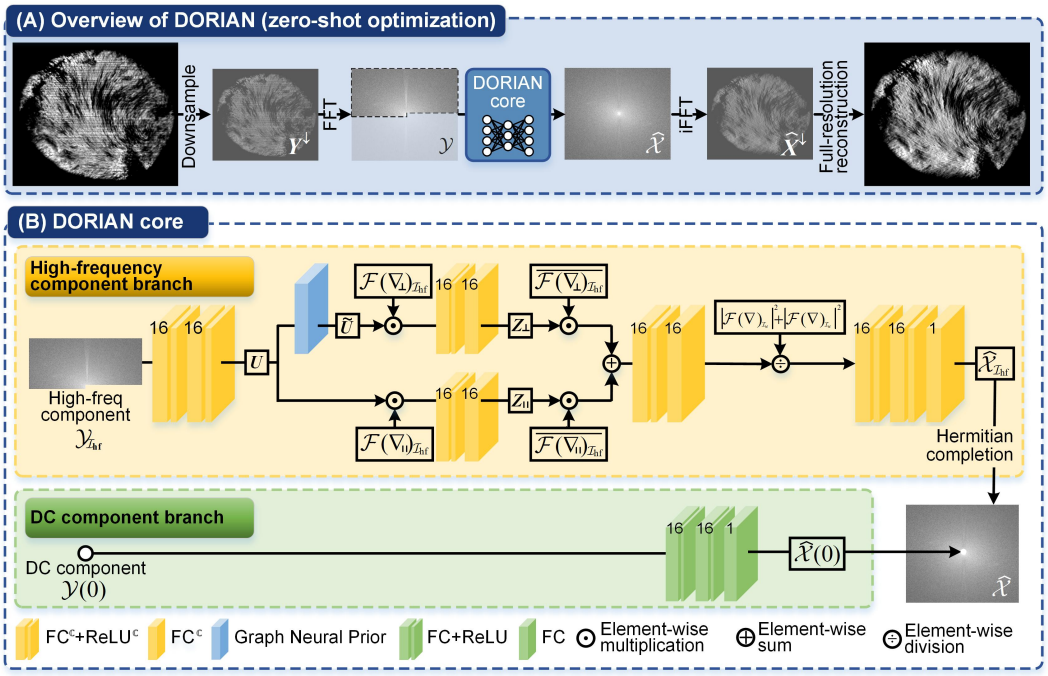


Figure 2: **Overview of DORIAN.** (A) Zero-shot restoration pipeline. (B) DORIAN core with the polar Fourier graph prior.

3.2 Polar Fourier Graph-Structured Neural Prior

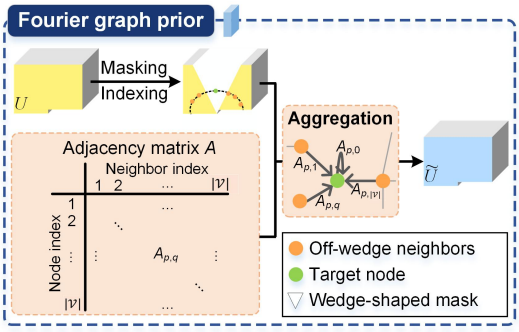


Figure 3: Architecture of the graph neural prior.

To reconstruct the highly corrupted spectral coefficients selected by S, we introduce a graph neural prior in Fourier space, as illustrated in Fig. 3. Since the discrete Fourier transform of a real-valued image exhibits conjugate symmetry, only the nonredundant half of the spectrum needs to be reconstructed, whereas the remaining half can subsequently be recovered by conjugate reflection. We thus denote the index set of that half, excluding the DC component, by $\mathcal{I}_{\mathrm{hf}}$, and restrict the observed spectrum and the spectral guidance mask accordingly as $\mathcal{Y}_{\mathcal{I}_{\mathrm{hf}}} \in \mathbb{C}^{|\mathcal{I}_{\mathrm{hf}}|}$ and $\mathbf{s} = S_{\mathcal{I}_{\mathrm{hf}}} \in \{0,1\}^{|\mathcal{I}_{\mathrm{hf}}|}$. Specifically, $\mathcal{Y}_{\mathcal{I}_{\mathrm{hf}}}$ is first embedded into a complex feature map $U \in \mathbb{C}^{|\mathcal{I}_{\mathrm{hf}}| \times d}$ by a point-wise complex-valued MLP[1] of alternating $\mathrm{FC}^{\mathbb{C}}(\cdot)$ and $\mathrm{ReLU}^{\mathbb{C}}(\cdot)$ layers.

We then construct an attributed graph $\mathcal{G} = (\mathcal{V}, U, A)$ with one node $v_p \in \mathcal{V}$ per frequency location $p \in \mathcal{I}_{\mathrm{hf}}$, where $U_p \in \mathbb{C}^d$ serves as the node attribute and A is a sparse adjacency matrix with $A_{p,q}$ representing the directed connection from node v_q to node v_p. Inspired by the observation that spectral coefficients at comparable radial frequencies carry scale-related spectral information, a polar structure is encoded in the design of A. Specifically, each DPSN-dominated node v_p ($s_p = 1$) draws N off-wedge references of similar radial frequency from:

$$\mathcal{C}(p) = \{q \in \mathcal{I}_{\mathrm{hf}} \mid s_q = 0,\ |\rho_q - \rho_p| \le \delta_\rho\}$$

where $\delta_\rho = 5$ is the radial bandwidth. We define $\mathcal{N}(p) = \mathrm{RandSample}_N(\mathcal{C}(p))$, with replacement when $|\mathcal{C}(p)| < N$, while off-wedge nodes retain only self-loops. The adjacency matrix is then defined

[1] $\mathrm{FC}^{\mathbb{C}}(Z) = ZW^\top + \mathbf{1}b^\top$ with complex W, b, and $\mathrm{ReLU}^{\mathbb{C}}(\cdot)$ applies ReLU to the real and imaginary parts separately.

as:

$$A_{p,q} = \big[(1 - s_p) + s_p \alpha_{p,p}\big]\mathbf{1}[p = q] + s_p\, \alpha_{p,q}\mathbf{1}[q \in \mathcal{N}(p)] \tag{3}$$

where $\alpha_{p,q} \in \mathbb{C}$ is a learnable signed aggregation weight optimized independently for each observation. Graph propagation is then performed using signed aggregation:

$$\widetilde{U}_p = (AU)_p \tag{4}$$

where the learnable signed weights allow both positive and negative contributions from the selected spectral references. As a result, off-wedge spectral coefficients remain unchanged after the graph neural prior, whereas DPSN-dominated coefficients are adaptively reconstructed from radially matched off-wedge references.

3.3 Graph-Guided Anisotropic Spectral Reconstruction

After graph propagation, the reconstructed $\widetilde{U} \in \mathbb{C}^{|\mathcal{V}|\times d}$, a high-dimensional spectral representation of the latent clean image, can be read out as the restored spectrum $\widehat{\mathcal{X}}_{\mathcal{I}_{\mathrm{hf}}} \in \mathbb{C}^{|\mathcal{V}|}$. Inspired by the high intensity variation orthogonal but low variation parallel to the DPSN direction, we propose to decompose the readout into directional gradients, which allows each direction to be reconstructed from the appropriate representation. Specifically, let $\nabla_{\parallel}$ and $\nabla_{\perp}$ denote the finite-difference operators along $\mathbf{d}_{\parallel}$ and $\mathbf{d}_{\perp}$, respectively, and $\mathcal{F}(\nabla_{\parallel})_{\mathcal{I}_{\mathrm{hf}}}$, $\mathcal{F}(\nabla_{\perp})_{\mathcal{I}_{\mathrm{hf}}}$ denote their Fourier responses, two directional branches predict the corresponding gradient spectra:

$$Z_{\perp} = \mathcal{R}_{\perp}^{\mathbb{C}}\big(\mathcal{F}(\nabla_{\perp})_{\mathcal{I}_{\mathrm{hf}}} \odot \widetilde{U}\big), \qquad Z_{\parallel} = \mathcal{R}_{\parallel}^{\mathbb{C}}\big(\mathcal{F}(\nabla_{\parallel})_{\mathcal{I}_{\mathrm{hf}}} \odot U\big) \tag{5}$$

where $\odot$ is element-wise multiplication broadcast across channels and $\mathcal{R}_{\perp}^{\mathbb{C}}(\cdot), \mathcal{R}_{\parallel}^{\mathbb{C}}(\cdot)$ alternate $\mathrm{FC}^{\mathbb{C}}(\cdot)$ layers and $\mathrm{ReLU}^{\mathbb{C}}(\cdot)$ activations. As a result, the gradient along $\mathbf{d}_{\parallel}$, which is essentially uncorrupted, is taken from U, whereas the gradient along $\mathbf{d}_{\perp}$, which must be recovered, is predicted through the graph-propagated $\widetilde{U}$. The two are then integrated into a gradient-consistent representation:

$$U^{\mathrm{grad}} = \frac{\overline{\mathcal{F}(\nabla_{\parallel})}_{\mathcal{I}_{\mathrm{hf}}} \odot Z_{\parallel} + \overline{\mathcal{F}(\nabla_{\perp})}_{\mathcal{I}_{\mathrm{hf}}} \odot Z_{\perp}}{|\mathcal{F}(\nabla_{\parallel})_{\mathcal{I}_{\mathrm{hf}}}|^2 + |\mathcal{F}(\nabla_{\perp})_{\mathcal{I}_{\mathrm{hf}}}|^2} \tag{6}$$

which is the closed-form least-squares solution of integrating $Z_{\parallel}$ and $Z_{\perp}$. The integrated representation is then projected onto $\widehat{\mathcal{X}}_{\mathcal{I}_{\mathrm{hf}}}$ by a further $\mathrm{FC}^{\mathbb{C}}(\cdot)$ stack. Additionally, as finite differences vanish at DC, Eq. (6) leaves the zero-frequency component undetermined. It is thus estimated by a separate real-valued branch from $\mathcal{Y}(0)$. Finally, the remaining half of the spectrum is inferred via conjugate symmetry, and the restored DPSN-free image is obtained by inverse Fourier transformation.

3.4 Zero-Shot Optimization Objective

DORIAN is optimized for each degraded observation alone by minimizing:

$$\mathcal{L}_{\mathrm{total}} = \mathcal{L}_{\mathrm{fid}} + \lambda_{\mathrm{TV}}\mathcal{L}_{\mathrm{TV}} + \lambda_{\mathrm{H}}\mathcal{L}_{\mathrm{H}} \tag{7}$$

where $\mathcal{L}_{\mathrm{fid}}$ enforces agreement between the reconstruction and the observation, $\mathcal{L}_{\mathrm{TV}}$ and $\mathcal{L}_{\mathrm{H}}$ are anisotropic total variation (TV) and Hessian terms, and λ_{TV}, λ_{H} are trade-off parameters.

Instead of pixel-wise fidelity, which would encourage the model to reproduce DPSN together with the underlying image details, we adopt a guided-filtering-based fidelity term that compares the two images only at a coarse scale:

$$\mathcal{L}_{\mathrm{fid}} = \big\|\mathrm{GF}_{\omega,\epsilon}(\widehat{X}) - \mathrm{GF}_{\omega,\epsilon}(Y)\big\|_2^2 \tag{8}$$

where $\mathrm{GF}_{\omega,\epsilon}(Z) \equiv \mathrm{GF}_{\omega,\epsilon}(Z; Z)$ is a self-guided filter whose large window ω and regularization parameter ϵ retain only coarse-scale intensity trends.

Next, fine-scale structure is constrained via the derivative terms, which are designed anisotropic due to the directionality of the DPSN. As a result, derivatives along $\mathbf{d}_{\perp}$, which DPSN dominates, are penalized on the restoration $\widehat{X}$, whereas derivatives involving $\mathbf{d}_{\parallel}$, which DPSN leaves largely intact, are penalized on the residual $(\widehat{X} - Y)$ so that image content along $\mathbf{d}_{\parallel}$ is preserved:

$$\mathcal{L}_{\mathrm{TV}} = \big\|m_{\mathrm{TV}} \odot \nabla_{\perp}\widehat{X}\big\|_1 + \overline{m}_{\mathrm{TV}}\,\big\|\nabla_{\parallel}(\widehat{X} - Y)\big\|_1, \tag{9}$$

$$\mathcal{L}_{\mathrm{H}} = \big\|m_{\mathrm{H}} \odot \nabla_{\perp\perp}\widehat{X}\big\|_1 + 2\overline{m}_{\mathrm{H}}\big\|\nabla_{\parallel\perp}(\widehat{X} - Y)\big\|_1 + \overline{m}_{\mathrm{H}}\big\|\nabla_{\parallel\parallel}(\widehat{X} - Y)\big\|_1 \tag{10}$$

where two adaptive weighting maps are derived as $m_{\mathrm{TV}} = \arctan\big(|\nabla_{\perp} Y|/(|\nabla_{\parallel} Y| + \varepsilon)\big)$ and $m_{\mathrm{H}} = \arctan\big(|\nabla_{\perp\perp} Y|/(|\nabla_{\parallel\parallel} Y| + \varepsilon)\big)$, so that smoothing along $\mathbf{d}_{\perp}$ is weakened wherever the observation also varies along $\mathbf{d}_{\parallel}$, which preserves structures that are curved or only approximately aligned with the DPSN direction while leaving strictly persistent components fully penalized. The arctangent is adopted to bound the ratio, and $\overline{m}_{\mathrm{TV}}, \overline{m}_{\mathrm{H}}$ denote their spatial averages.

3.5 Anisotropic Resampling and Residual Transfer for Structure Preservation

A major limitation of zero-shot DPSN removal is the indistinguishability, under Eq. (1), between DPSN and genuine image structures that extend along $\mathbf{d}_{\parallel}$. Downsampling the observation along $\mathbf{d}_{\parallel}$, however, amplifies the difference between the two, as the curvature and intensity fluctuation of genuine structures become steeper, while DPSN, which is nearly constant along $\mathbf{d}_{\parallel}$ by definition, is unaffected. Meanwhile, optimizing at the reduced resolution is correspondingly cheaper. DORIAN is therefore optimized on $Y^{\downarrow} = D_r^{\parallel}(Y)$, where $D_r^{\parallel}(\cdot)$ denotes downsampling by a factor r along $\mathbf{d}_{\parallel}$ only, yielding downsampled DPSN-free estimation $\widehat{X}^{\downarrow}$.

To recover $\widehat{X}$ at full-resolution, instead of upsampling $\widehat{X}^{\downarrow}$ directly, which would smooth structures that the downsampled observation resolves only partially, we transfer the correction through the residual $R^{\downarrow} = \widehat{X}^{\downarrow} - Y^{\downarrow}$:

$$\widehat{X}_i = Y_i + \underset{j \in W_i^{\parallel}}{\mathrm{median}} \big[U_r^{\parallel}(R^{\downarrow})\big]_j \tag{11}$$

where $U_r^{\parallel}(\cdot)$ denotes upsampling along $\mathbf{d}_{\parallel}$. A median filter, taken over a window $W_i^{\parallel}$ elongated along $\mathbf{d}_{\parallel}$, additionally smooths the residual, so that image structures leaking into it, being less persistent along $\mathbf{d}_{\parallel}$ than DPSN, are not subtracted from the observation.

4 Experiments

Experimental setup. We evaluate synthetic DPSN on the DIV2K validation set (100 images) (Agustsson & Timofte, 2017), BSD68 dataset (68 images) (Roth & Black, 2005), and a DPSN-free fluorescence volume of adipose tissue ($95 \times 2076 \times 1728$ voxels), under different degradation settings including periodic full-length, heterogeneous full-length, and heterogeneous finite-length. Real DPSN is evaluated on remote-sensing, LSFM, and AFM data. Note that the orientation of the DPSN is known throughout, by construction for synthetic data or from acquisition geometry for real data. We compare against six baselines, including WFF (Münch et al., 2009), GSR (Rottmayer et al., 2025), NR-CS (Schwartz et al., 2019), ASCNet (Yuan et al., 2025), Self-BSR (Qu et al., 2025), and DeStripe (Liu et al., 2022). Note that ASCNet and Self-BSR are retrained on the DIV2K training set. For synthetic data, we report $\mathrm{PSNR_D}$ over DPSN-affected pixels only and SSIM over the full image. Although the two effects are not fully decoupled, $\mathrm{PSNR_D}$ primarily reflects DPSN removal, whereas SSIM captures overall preservation of image content. DORIAN uses a single fixed configuration throughout; further details are provided in Appendix D. Runtime and memory analyses are provided in Appendix D.7.

4.1 Restoration of Synthetic DPSN

Table 1 reports quantitative results on DIV2K and BSD68 under three synthetic DPSN settings, with representative examples shown in Fig. 4. Complete qualitative comparisons are provided in Figs. 8–13. Specifically, under periodic full-length DPSN, all methods perform reasonably well. ASCNet achieves the highest $\mathrm{PSNR_D}$ on DIV2K (28.78 dB), but drops substantially on BSD68, whereas DORIAN remains competitive on both datasets. As DPSN becomes heterogeneous, DORIAN maintains strong performance, exceeding the strongest baseline by 4.43 and 2.47 dB on DIV2K and BSD68 under the finite-length setting. Consistent with the qualitative comparisons, competing methods either leave significant residual DPSN, for example, ASCNet (Yuan et al., 2025) and DeStripe (Liu et al., 2022), or remove image structures together with it, e.g., GSR (Rottmayer et al., 2025), whereas DORIAN consistently suppresses the degradation while preserving local details. Additionally, we evaluate DORIAN on a microscopy dataset, whose image statistics differ from natural images. DORIAN outperforms the strongest baseline by 3.59 dB in $\mathrm{PSNR_D}$ and 0.0296 in SSIM.

Table 1: Quantitative comparison on DIV2K, BSD68, and a microscopy dataset under different synthetic DPSN settings. Each setting reports $\mathrm{PSNR_D}$ (upper) and SSIM (lower, shaded). Red and blue indicate the best and the second-best result in each row.

Dataset	Setting	Input	WFF	GSR	NR-CS	ASCNet	Self-BSR	DeStripe	DORIAN
DIV2K	periodic	14.89	25.16	26.43	26.75	**28.78**	21.05	24.34	**27.93**
	full-length	0.8676	0.9664	0.9676	0.9696	**0.9722**	0.8966	0.9648	**0.9704**
	heterogeneous	11.13	20.52	**21.74**	16.91	16.14	12.24	18.26	**22.12**
	full-length	0.8302	0.9360	**0.9429**	0.9160	0.9255	0.7787	0.9193	**0.9524**
	heterogeneous	12.54	13.06	**18.03**	16.74	13.03	12.54	14.54	**22.46**
	finite-length	0.9473	0.9330	**0.9537**	0.9518	0.9493	0.8988	0.9460	**0.9619**
BSD68	periodic	14.25	19.83	**23.28**	22.14	20.03	20.13	20.00	**23.33**
	full-length	0.7000	0.9310	0.9352	**0.9573**	0.9204	0.8375	0.9156	**0.9588**
	heterogeneous	11.81	17.61	**18.42**	17.39	13.05	11.88	13.86	**20.12**
	full-length	0.9113	0.9413	**0.9451**	0.9424	0.9287	0.8621	0.9212	**0.9524**
	heterogeneous	11.93	16.26	**20.51**	15.30	13.30	13.30	18.64	**22.98**
	finite-length	0.8604	0.9042	**0.9394**	0.9075	0.8931	0.8098	0.9251	**0.9542**
Microsc.	heterogeneous	19.54	19.58	20.73	20.62	19.85	19.43	**22.43**	**26.02**
	finite-length	0.9397	0.9392	0.9326	**0.9510**	0.9449	0.9258	0.9467	**0.9763**

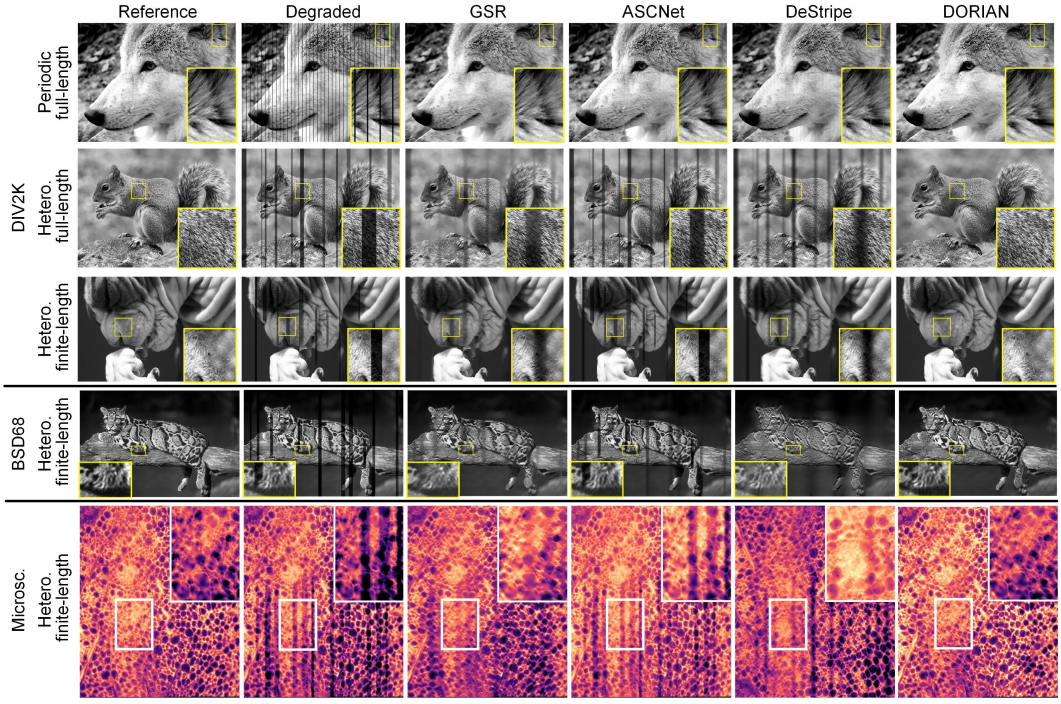


Figure 4: **Qualitative comparison under synthetic DPSN across DIV2K, BSD68, and microscopy data.** Insets magnify the boxed regions.

4.2 Restoration of Real DPSN

We evaluate DORIAN on real DPSN in remote sensing, AFM, and LSFM, where clean references are unavailable. We therefore report $|\rho|$, the absolute correlation between the removed component and the restored image, where lower values suggest less removal of image-correlated structure (Appendix D.5). Since $|\rho|$ does not measure residual DPSN, qualitative restoration results are shown alongside it in Fig. 5. Across the three modalities, DPSN patterns become increasingly heterogeneous in intensity and spatial extent. While competing methods either leave visible residual DPSN or remove image-correlated structures, DORIAN is the only method that consistently suppresses the

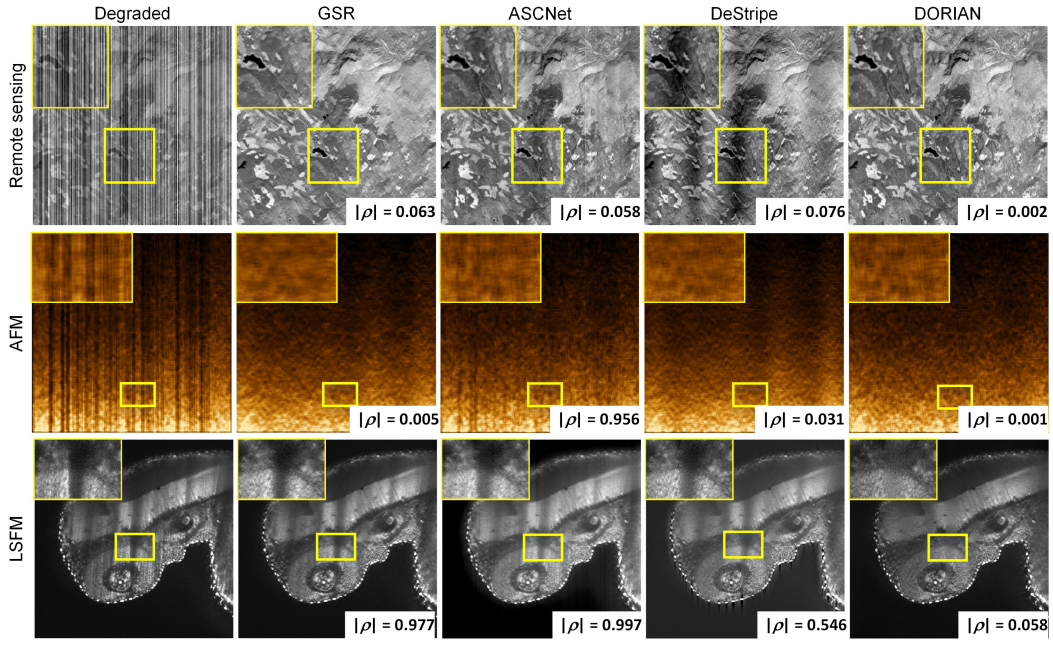


Figure 5: **Restoration of real DPSN across remote sensing, AFM, and LSFM.** Each panel reports $|\rho|$, the absolute correlation between the removed component and the restored image.

degradation across all three cases while maintaining low $|\rho|$. Complete comparisons are provided in Figs. 14–16.

4.3 Ablation Studies

All ablations are performed on the DIV2K validation set under the heterogeneous finite-length setting (Table 2). Additional analysis of the directional gradient readout is provided in Appendix G.

Table 2: Ablation studies on DIV2K.

	Variant	$\mathrm{PSNR_D}$	SSIM
Representation	Spatial U-Net	20.80	0.9523
	Fourier U-Net	15.40	0.8870
Graph	w/o propagation	12.82	0.9148
	Fixed polar	17.67	0.9382
	Learnable unstructured	18.62	0.9497
Objective	Pixel-wise fidelity	13.78	0.9400
	w/o Hessian	17.95	0.8501
	w/o TV	16.74	0.9472
	Uniform weighting	18.47	0.8886
Readout	Direct readout	18.97	0.9125
DORIAN (full)		**22.46**	**0.9619**

quality among these variants.

Necessity of the polar Fourier representation. In Fig. 17, we first replace the DORIAN core with a parameter-matched spatial U-Net (Appendix E). The residual DPSN left by the image-space U-Net reflects the difficulty of modeling the long-range persistence of DPSN with a bounded spatial receptive field. Moving the U-Net to the Fourier domain resolves this non-locality (Appendix F), although reconstructing the entire spectrum inevitably modifies reliable coefficients, degrading image texture. In contrast, DORIAN combines the compact Fourier representation with the selectivity of the graph neural prior, delivering the best reconstruction

Impact of the graph neural prior. Table 2 and Fig. 18 examine the effects of graph propagation, polar neighborhood structure, and observation-adaptive weighting. Without graph propagation, DPSN-contaminated coefficients are no longer effectively reconstructed and the residual DPSN remains nearly intact. Using fixed polar aggregation substantially suppresses the degradation but introduces local distortions, while adaptive aggregation over unstructured neighborhoods leaves

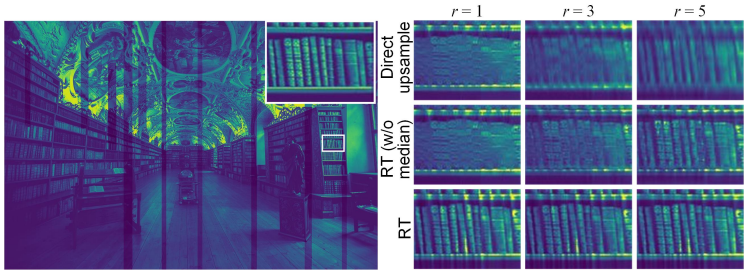


Figure 6: **Ablation of anisotropic resampling r and residual transfer (RT).**

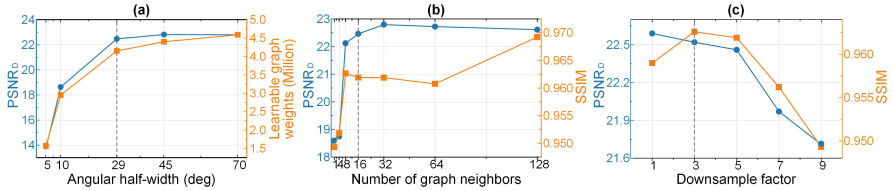


Figure 7: **Parameter sensitivity** to (a) angular half-width, (b) number of graph neighbors, and (c) downsampling factor. Dashed lines are the default settings.

residual DPSN. Combining the polar spectral structure with observation-adaptive aggregation, as in DORIAN, yields the best qualitative and quantitative performance.

Ablation of the restoration objective. Table 2 independently evaluates the fidelity term, the combination of first- and second-order regularization, and the adaptive weighting strategy in the regularization. Replacing the proposed guided-filtering-based fidelity with pixel-wise ℓ_2 preserves DPSN-contaminated local intensities and substantially degrades restoration (Fig. 19). Using either TV or Hessian alone is also suboptimal, where TV suppresses DPSN-induced directional variations more effectively but over-smooths fine image structures, whereas Hessian provides complementary second-order constraints that better preserve local continuity (Fig. 20). Their combination thus better balances artifact suppression and structural preservation. Finally, replacing the adaptive weighting maps m_{TV} and m_{H} with uniform regularization causes over-smoothing, confirming the importance of adapting the regularization strength to local directional variations in the observation (Fig. 21).

Effect of the anisotropic resampling. A key challenge in zero-shot DPSN removal is distinguishing DPSN from genuine image structures aligned with the same direction. As shown in Fig. 6, increasing the downsampling factor r progressively improves the preservation of aligned structures (zoom-in, first row). Residual transfer (RT) further preserves structures only partially resolved at low resolution by transferring the upsampled residual correction rather than the restoration itself, while the embedded directional median filter suppresses remaining structure leakage in the transferred residual.

Parameter sensitivity. Fig. 7 shows that the restoration is governed by lower bounds rather than by tuned optima: performance is unaffected once the spectral coverage and the number of graph references are sufficient, and remains stable under moderate downsampling. Detailed analyses are provided in Appendix H.

5 Conclusion

We introduce DORIAN, a zero-shot framework that recasts DPSN removal as Fourier-domain spectral reconstruction with a graph neural prior. Combined with anisotropic resampling and residual transfer, DORIAN handles DPSN with diverse spatial heterogeneity across natural, microscopy, and remote-sensing data using a single fixed configuration. DORIAN currently assumes a known dominant orientation of DPSN and single-channel input, and thus extending it to unknown or multiple orientations and more general degradation models remains future work.

AI USE STATEMENT

The DORIAN methodology, core algorithmic design, and primary implementation were developed without generative AI assistance. Generative AI tools were used to assist with ablation-study design, auxiliary experimental and baseline-code debugging, interpretation and presentation of experimental results, literature search, and manuscript editing. All AI-assisted code, analyses, claims, and references were reviewed and verified by the authors. All reported experimental results were obtained from the authors' own implementations and experiments. We take responsibility for the final content of this work.

REPRODUCIBILITY STATEMENT

Anonymous source code for reproducing the proposed method and experiments is available at `https://anonymous.4open.science/r/DORIAN-2E2D/`

REFERENCES

Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017.

Coleman Broaddus, Alexander Krull, Martin Weigert, Uwe Schmidt, and Gene Myers. Removing structured noise with self-supervised blind-spot networks. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pp. 159–163, 2020. doi: 10.1109/ISBI45749.2020.9098336.

Claudio Canale, Bruno Torre, Davide Ricci, and Pier Carlo Braga. *Recognizing and Avoiding Artifacts in Atomic Force Microscopy Imaging*, pp. 31–43. Humana Press, Totowa, NJ, 2011. ISBN 978-1-61779-105-5. doi: 10.1007/978-1-61779-105-5_3. URL `https://doi.org/10.1007/978-1-61779-105-5_3`.

Yi Chang, Houzhang Fang, Luxin Yan, and Hai Liu. Robust destriping method with unidirectional total variation and framelet regularization. *Opt. Express*, 21(20):23307–23323, Oct 2013. doi: 10.1364/OE.21.023307. URL `https://opg.optica.org/oe/abstract.cfm?URI=oe-21-20-23307`.

Yi Chang, Luxin Yan, Tao Wu, and Sheng Zhong. Remote sensing image stripe noise removal: From image decomposition perspective. *IEEE Transactions on Geoscience and Remote Sensing*, 54(12):7018–7031, 2016. doi: 10.1109/TGRS.2016.2594080.

Yi Chang, Meiya Chen, Luxin Yan, Xi-Le Zhao, Yi Li, and Sheng Zhong. Toward universal stripe removal via wavelet-based deep convolutional neural network. *IEEE Transactions on Geoscience and Remote Sensing*, 58(4):2880–2897, 2020. doi: 10.1109/TGRS.2019.2957153.

Keunwoo Choi, György Fazekas, Mark Sandler, and Kyunghyun Cho. A comparison of audio signal preprocessing methods for deep neural networks on music tagging. In *2018 26th European signal processing conference (EUSIPCO)*, pp. 1870–1874. IEEE, 2018.

Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, 2023. URL `https://openreview.net/forum?id=OnD9zGAGT0k`.

Yuning Cui, Syed Waqas Zamir, Salman Khan, Alois Knoll, Mubarak Shah, and Fahad Shahbaz Khan. AdaIR: Adaptive all-in-one image restoration via frequency mining and modulation. In *The Thirteenth International Conference on Learning Representations*, 2025. URL `https://openreview.net/forum?id=M5t0WvjfCg`.

Jérôme Fehrenbach, Pierre Weiss, and Corinne Lorenzo. Variational algorithms to remove stationary noise: Applications to microscopy imaging. *IEEE Transactions on Image Processing*, 21(10):4420–4430, 2012. doi: 10.1109/TIP.2012.2206037.

Amirhossein Khalilian-Gourtani, Mariano Tepper, Victor Minden, and Dmitri B. Chklovskii. Strip the stripes: Artifact detection and removal for scanning electron microscopy imaging. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1060–1064, 2019. doi: 10.1109/ICASSP.2019.8683119.

Mian Li, Jan Rieck, Beatriz Noheda, Jos B. T. M. Roerdink, and Michael H. F. Wilkinson. Stripe noise removal in conductive atomic force microscopy. 14(1):3931. ISSN 2045-2322. doi: 10.1038/s41598-024-54094-w. URL https://www.nature.com/articles/s41598-024-54094-w.

Xiao Liang, Yali Zang, Di Dong, Liwen Zhang, Mengjie Fang, Xin Yang, Alicia Arranz, Jorge Ripoll, Hui Hui, and Jie Tian. Stripe artifact elimination based on nonsubsampled contourlet transform for light sheet fluorescence microscopy. *Journal of Biomedical Optics*, 21(10):106005, 2016. doi: 10.1117/1.JBO.21.10.106005. URL https://doi.org/10.1117/1.JBO.21.10.106005.

Yu Liu, Kurt Weiss, Nassir Navab, Carsten Marr, Jan Huisken, and Tingying Peng. Destripe: A self2self spatio-spectral graph neural network with unfolded hessian for stripe artifact removal in light-sheet microscopy. In Linwei Wang, Qi Dou, P. Thomas Fletcher, Stefanie Speidel, and Shuo Li (eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, pp. 99–108, Cham, 2022. Springer Nature Switzerland. ISBN 978-3-031-16440-8.

Qiqi Lu, Xiuli Liu, Qianjin Feng, Shaoqun Zeng, and Shenghua Cheng. A foundation model for multi-task cross-distribution restoration of fluorescence microscopy images. 17(1):3729. ISSN 2041-1723. doi: 10.1038/s41467-026-70307-4. URL https://www.nature.com/articles/s41467-026-70307-4.

Yi-Si Luo, Xi-Le Zhao, Tai-Xiang Jiang, Yu-Bang Zheng, and Yi Chang. Hyperspectral mixed noise removal via spatial-spectral constrained unsupervised deep image prior. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14:9435–9449, 2021. doi: 10.1109/JSTARS.2021.3111404.

Chenxi Ma, Weimin Tan, Ruian He, and Bo Yan. Pretraining a foundation model for generalizable fluorescence microscopy-based image restoration. 21(8):1558–1567. ISSN 1548-7091, 1548-7105. doi: 10.1038/s41592-024-02244-3. URL https://www.nature.com/articles/s41592-024-02244-3.

Hongcheng Mai, Jie Luo, Luciano Hoeher, Rami Al-Maskari, Izabela Horvath, Ying Chen, Florian Kofler, Marie Piraud, Johannes C. Paetzold, Jennifer Modamio, Mihail Todorov, Markus Elsner, Farida Hellal, and Ali Ertürk. Whole-body cellular mapping in mouse using standard IgG antibodies. 42(4):617–627. ISSN 1087-0156, 1546-1696. doi: 10.1038/s41587-023-01846-0. URL https://www.nature.com/articles/s41587-023-01846-0.

Beat Münch, Pavel Trtik, Federica Marone, and Marco Stampanoni. Stripe and ring artifact removal with combined wavelet — fourier filtering. *Opt. Express*, 17(10):8567–8591, May 2009. doi: 10.1364/OE.17.008567. URL https://opg.optica.org/oe/abstract.cfm?URI=oe-17-10-8567.

Kazuki Naganuma and Shunsuke Ono. A general destriping framework for remote sensing images using flatness constraint. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–16, 2022. doi: 10.1109/TGRS.2022.3153995.

Chao Qu, Zewei Chen, Jingyuan Zhang, Xiaoyu Chen, and Jing Han. Self-bsr: Self-supervised image denoising and destriping based on blind-spot regularization. *IEEE Transactions on Circuits and Systems for Video Technology*, 35(9):8666–8678, 2025. doi: 10.1109/TCSVT.2025.3559214.

Pietro Ricci, Vladislav Gavryusev, Caroline Müllenbroich, Lapo Turrini, Giuseppe de Vito, Ludovico Silvestri, Giuseppe Sancataldo, and Francesco Saverio Pavone. Removing striping artifacts in light-sheet fluorescence microscopy: a review. *Progress in Biophysics and Molecular Biology*, 168:52–65, 2022. ISSN 0079-6107. doi: https://doi.org/10.1016/j.pbiomolbio.2021.07.003. URL https://www.sciencedirect.com/science/article/pii/S0079610721000821. The Resolution Revolution: Fluorescence Microscopy of Biological Samples from Micro to Meso.

Stefan Roth and Michael J. Black. Fields of experts: A framework for learning image priors. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pp. 860–867, 2005.

Niklas Rottmayer, Claudia Redenbach, and Florian O. Fahrbach. A universal and effective variational method for destriping: application to light-sheet microscopy, fib-sem, and remote sensing images. *Opt. Express*, 33(3):5800–5809, Feb 2025. doi: 10.1364/OE.542868. URL https://opg.optica.org/oe/abstract.cfm?URI=oe-33-3-5800.

Jonathan Schwartz, Yi Jiang, Yongjie Wang, Anthony Aiello, Pallab Bhattacharya, Hui Yuan, Zetian Mi, Nabil Bassim, and Robert Hovden. Removing stripes, scratches, and curtaining with nonrecoverable compressed sensing. *Microscopy and Microanalysis*, 25(3):705–710, 06 2019. ISSN 1431-9276. doi: 10.1017/S1431927619000254. URL https://doi.org/10.1017/S1431927619000254.

Chiheb Trabelsi, Olexa Bilaniuk, Ying Zhang, Dmitriy Serdyuk, Sandeep Subramanian, Joao Felipe Santos, Soroush Mehri, Negar Rostamzadeh, Yoshua Bengio, and Christopher J Pal. Deep complex networks. *arXiv preprint arXiv:1705.09792*, 2017.

Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. *International Journal of Computer Vision*, 128(7):1867–1888, March 2020. ISSN 1573-1405. doi: 10.1007/s11263-020-01303-4. URL http://dx.doi.org/10.1007/s11263-020-01303-4.

Jing-Hua Yang, Xi-Le Zhao, Tian-Hui Ma, Yong Chen, Ting-Zhu Huang, and Meng Ding. Remote sensing images destriping using unidirectional hybrid total variation and nonconvex low-rank regularization. *Journal of Computational and Applied Mathematics*, 363:124–144, 2020. ISSN 0377-0427. doi: https://doi.org/10.1016/j.cam.2019.06.004. URL https://www.sciencedirect.com/science/article/pii/S0377042719302766.

Shiqi Yang, Hanlin Qin, Shuai Yuan, Xiang Yan, and Hossein Rahmani. Destripecyclegan: Stripe simulation cyclegan for unsupervised infrared image destriping. *IEEE Transactions on Instrumentation and Measurement*, 73:1–14, 2024. doi: 10.1109/TIM.2024.3476560.

Shuai Yuan, Hanlin Qin, Xiang Yan, Shiqi Yang, Shuowen Yang, Naveed Akhtar, and Huixin Zhou. Ascnet: Asymmetric sampling correction network for infrared image destriping. *IEEE Transactions on Geoscience and Remote Sensing*, 63:1–15, 2025. doi: 10.1109/TGRS.2025.3534838.

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5728–5739, June 2022.

APPENDIX

A Physical Origins of DPSN Across Imaging Modalities

Although directionally persistent structured noise (DPSN) appears across different imaging modalities, its physical origin is modality-dependent. Here, we briefly describe representative mechanisms in light-sheet fluorescence microscopy (LSFM), remote sensing, and atomic force microscopy (AFM). Despite their different origins, these degradations share a common image-domain characteristic: they remain strongly correlated along a dominant direction over an extended spatial range, while their effect on the observed signal may be additive or multiplicative, and their intensity, width, spacing, and spatial extent may vary substantially.

Light-sheet fluorescence microscopy. LSFM illuminates the specimen from the side with a thin light sheet while detecting the emitted fluorescence along an orthogonal optical path (Ricci et al., 2022). When imaging large or optically inhomogeneous specimens, dense or absorbing structures along the excitation path may attenuate the light sheet, resulting in non-uniform illumination and stripy shadows. This attenuation acts multiplicatively on the underlying fluorescence signal. The resulting shadows extend along the light-propagation direction and can vary considerably in width and intensity depending on the local optical properties of the specimen. Nevertheless, they remain highly correlated along the illumination direction over an extended distance. Consequently, these absorption-induced stripe artifacts constitute a representative instance of DPSN.

Remote sensing. DPSN in remote-sensing images is mainly caused by differences in the response of adjacent detectors (Chang et al., 2020). Two common acquisition schemes are push-broom and whisk-broom imaging, and the appearance of the resulting DPSN depends strongly on the scanning mechanism of the imaging instrument. According to the detector response, DPSN in remote sensing can be either additive, where its intensity is largely independent of the underlying signal, or multiplicative, where the intensity varies with the image signal (Chang et al., 2020). It may extend over the entire image range or occur as broken stripes of arbitrary length. Their periodicity can also differ between acquisition systems: periodic stripes are commonly associated with whisk-broom imaging, whereas nonperiodic stripes mainly occur in push-broom systems (Chang et al., 2020).

Atomic force microscopy. DPSN is also frequently observed in conductive atomic force microscopy (c-AFM), particularly in lateral measurements (Li et al.). In such measurements, the generation of DPSN may be related to charge accumulation and drift, as its direction has been observed to coincide with the direction of the electrical current from the fixed electrode. Extremely strong potential differences may additionally contribute to its formation. In the observed c-AFM images, these artifacts can appear as closely spaced, parallel stripes extending over long spatial ranges, while exhibiting different intensities and thicknesses. They may also occur only in selected regions, with varying density, spacing, and spatial extent. Such stripe degradation is modeled as an additive component superimposed on the underlying clean image (Li et al.).

B Relation to Classical DPSN Restoration

DORIAN is closely related to the two major classical paradigms for DPSN removal: filtering-based methods in the Fourier domain and variational restoration in the image space. Below, we show that DORIAN retains the key restoration principles of both paradigms while generalizing their fixed assumptions through observation-adaptive spectral reconstruction.

B.1 DORIAN: From Frequency-Domain Filtering to Spectral Reconstruction

Frequency-domain DPSN removal methods, in general, consist of two components: a spectral region in which the degradation is assumed to be concentrated and an attenuation function applied within this region. This can be written as:

$$\widehat{\mathcal{X}} = H \odot \mathcal{Y} \tag{12}$$

where H denotes the frequency response prescribed by the method and $\odot$ denotes element-wise multiplication. Here, we denote the spectral region effectively modified by the filter as Ω_H. In conventional filtering methods, Ω_H is typically kept narrow to limit the removal of genuine image structures whose spectral components overlap with the degradation. In contrast, DORIAN defines a broader DPSN-dominated spectral region Ω_s, comprising a wedge of angular half-width Δ_ϕ together with a low-frequency disk in our implementation, which can encompass Ω_H, i.e.,

$$\Omega_H \subseteq \Omega_s \tag{13}$$

The coefficient-wise attenuation prescribed by H can also be recovered as a limiting configuration of the proposed graph-based reconstruction. Specifically, by suppressing all non-self connections for coefficients in Ω_H, the adjacency matrix A in Eq. (3) reduces graph propagation to independent coefficient-wise processing through the self-loop. The subsequent spectral reconstruction can then, in principle, reproduce the attenuation behavior prescribed by H.

As a result, DORIAN naturally contains conventional coefficient-wise frequency-domain filtering as a degenerate case, while extending fixed spectral attenuation to observation-adaptive spectral reconstruction. In the general case, moreover, Eq. (4) allows DPSN-dominated coefficients to additionally aggregate information from selected off-wedge reference nodes through observation-specific learnable signed weights.

B.2 Connecting DORIAN to Spatial-Domain Variational Restoration in the Fourier Domain

DORIAN can be viewed as an observation-adaptive generalization of classical directional variational restoration in the Fourier domain. To make this connection explicit, we first revisit a fundamental anisotropic variational model for DPSN removal as a representative example.

B.2.1 A Representative Directional Variational Model

We begin with a representative image-domain variational model that combines pixel-wise data fidelity with anisotropic directional regularization:

$$\begin{aligned} X^\star &= \arg\min_X \frac{1}{2}\|X-Y\|_2^2 + \lambda \mathcal{R}(X) \\ &= \arg\min_X \frac{1}{2}\|X-Y\|_2^2 + \lambda \left\{ \|\nabla_\perp X\|_1 + \|\nabla_\parallel (X-Y)\|_1 \right\} \\ &\xrightarrow{\text{split Bregman}} \arg\min_{X, G_\perp, G_\parallel} \Bigg[\frac{1}{2}\|X-Y\|_2^2 + \lambda \left\{ \|G_\perp\|_1 + \|G_\parallel - \nabla_\parallel Y\|_1 \right\} \\ &\qquad\qquad + \frac{\mu_\perp}{2}\|G_\perp - \nabla_\perp X - B_\perp\|_2^2 + \frac{\mu_\parallel}{2}\left\|G_\parallel - \nabla_\parallel X - B_\parallel\right\|_2^2 \Bigg]. \end{aligned} \tag{14}$$

where $\mathcal{R}(X)$ denotes an anisotropic total variation regularizer, and λ controls the trade-off between data fidelity and regularization. For DPSN removal, anisotropic total variation is particularly suitable because DPSN exhibits pronounced intensity variations across its dominant direction while remaining relatively persistent along it. Hence, $\|\nabla_\perp X\|_1$ suppresses DPSN-induced gradients across the dominant direction, whereas $\|\nabla_\parallel (X-Y)\|_1$ constrains residual variations along the dominant direction, thereby preserving image structures along this direction. To make the optimization in Eq. (14) tractable, split Bregman is commonly adopted, which introduces the auxiliary variables $G_\perp = \nabla_\perp X$ and $G_\parallel = \nabla_\parallel X$, together with the Bregman variables $B_\perp$ and $B_\parallel$. The resulting

problem can then be solved by alternating updates of $G_\perp$, $G_\parallel$, and X:

$$\begin{cases} G_\perp^{(k+1)} = \mathrm{shrink}\left(\nabla_\perp X^{(k)} + B_\perp^{(k)}, \dfrac{\lambda}{\mu_\perp}\right), \\ G_\parallel^{(k+1)} = \nabla_\parallel Y + \mathrm{shrink}\left(\nabla_\parallel X^{(k)} + B_\parallel^{(k)} - \nabla_\parallel Y, \dfrac{\lambda}{\mu_\parallel}\right), \\ X^{(k+1)} = \mathcal{F}^{-1}\left(\dfrac{\mathcal{Y} + \mu_\perp \overline{\mathcal{F}(\nabla_\perp)} \odot \mathcal{F}\left(G_\perp^{(k+1)} - B_\perp^{(k)}\right) + \mu_\parallel \overline{\mathcal{F}(\nabla_\parallel)} \odot \mathcal{F}\left(G_\parallel^{(k+1)} - B_\parallel^{(k)}\right)}{1 + \mu_\perp \left|\mathcal{F}(\nabla_\perp)\right|^2 + \mu_\parallel \left|\mathcal{F}(\nabla_\parallel)\right|^2}\right). \end{cases} \tag{15}$$

where k denotes the iteration index, $B_\perp^{(0)} = B_\parallel^{(0)} = 0$, and $\mathrm{shrink}(z, \tau) = \mathrm{sign}(z)\max(|z| - \tau, 0)$.

As a result, although the variational model is formulated in the image domain, the update of X is carried out in the Fourier domain, where the directional components are combined through the adjoint responses of $\nabla_\perp$ and $\nabla_\parallel$, followed by frequency-dependent normalization. Thus, classical spatial-domain variational restoration already contains an implicit directional spectral reconstruction step.

B.2.2 DORIAN: From Fixed Directional Updates to Adaptive Spectral Reconstruction

The connection to DORIAN becomes clear from the alternating solution above. If we interpret classical directional variational restoration as consisting of two key operations, that is, the estimation of directional quantities and their subsequent integration to recover the latent image, DORIAN retains this two-stage structure. Specifically, the three updates in Eq. (15) have the following counterparts in DORIAN:

$$\begin{cases} Z_\perp = R_\perp^{\mathbb{C}}\left(\mathcal{F}(\nabla_\perp)_{\mathcal{I}_{\mathrm{hf}}} \odot \widetilde{U}\right), \\ Z_\parallel = R_\parallel^{\mathbb{C}}\left(\mathcal{F}(\nabla_\parallel)_{\mathcal{I}_{\mathrm{hf}}} \odot U\right), \\ U^{\mathrm{grad}} = \dfrac{\overline{\mathcal{F}(\nabla_\parallel)_{\mathcal{I}_{\mathrm{hf}}}} \odot Z_\parallel + \overline{\mathcal{F}(\nabla_\perp)_{\mathcal{I}_{\mathrm{hf}}}} \odot Z_\perp}{\left|\mathcal{F}(\nabla_\parallel)_{\mathcal{I}_{\mathrm{hf}}}\right|^2 + \left|\mathcal{F}(\nabla_\perp)_{\mathcal{I}_{\mathrm{hf}}}\right|^2}. \end{cases} \tag{16}$$

where the first two lines replace the fixed shrinkage-based estimates $G_\perp$ and $G_\parallel$ with the observation-adaptive spectral representations $Z_\perp$ and $Z_\parallel$, respectively. The third line retains the directional spectral integration structure of the classical X-update, combining the two directional estimates through the adjoint Fourier responses of $\nabla_\perp$ and $\nabla_\parallel$.

As a result, DORIAN retains the directional spectral integration structure of classical variational restoration while generalizing its fixed directional updates into observation-adaptive spectral reconstruction.

B.2.3 DORIAN as a DPSN-Consistent Restoration Model

The connection discussed above further highlights two restrictive assumptions of the classical variational model in Eqs. (14) and (15) when applied to DPSN restoration:

- **Pixel-wise fidelity.** The fidelity term $\frac{1}{2}\|X - Y\|_2^2$ directly penalizes deviations from the degraded observation. For DPSN, however, the degradation is directionally persistent and spatially correlated, such that enforcing local agreement with Y can encourage the restoration to reproduce the structured degradation itself.
- **Fixed directional shrinkage.** The directional quantities are obtained through predefined shrinkage operators with fixed thresholds, which cannot adapt their restoration behavior to the spatially varying characteristics of an individual degraded observation.

DORIAN relaxes both assumptions while retaining the basic structure of variational restoration. For data fidelity, the direct pixel-wise term $\frac{1}{2}\|X - Y\|_2^2$ is removed from the directional spectral

reconstruction. Consequently, the gradient-based integration no longer determines the zero-frequency component, since the directional finite-difference operators have zero response at DC. DORIAN therefore estimates the DC component separately, while reconstructing the remaining frequency components from the directional spectral representations. Meanwhile, the predefined shrinkage operators are replaced by the observation-adaptive spectral prediction described above. The directional quantities are therefore inferred from the individual degraded observation rather than determined by fixed thresholds.

As a result, DORIAN retains the classical combination of data fidelity and directional priors, while adapting both components to the directionally persistent and spatially correlated characteristics of DPSN.

C Relation to DeStripe

The most closely related prior work is DeStripe (Liu et al., 2022), which introduced graph-based reconstruction of stripe-corrupted Fourier coefficients. DORIAN extends this idea through substantially different corruption modeling, graph construction, spectral readout, restoration objective, and full-resolution recovery:

- **Corruption localization.** DeStripe identifies corrupted Fourier coefficients by fitting a Rayleigh distribution to the Fourier-magnitude histogram, relying on a statistical assumption about the spectrum. DORIAN instead defines a directional wedge from the dominant DPSN orientation, together with a low-frequency disk, providing a deterministic and direction-aware localization of the corrupted spectral region.
- **Selective graph reconstruction.** DeStripe updates the spectral representation globally, including coefficients that are not identified as corrupted. DORIAN instead restricts graph-based aggregation to DPSN-dominated nodes only, which reconstruct from radially matched off-wedge references, while off-wedge nodes retain only self-loops. This preserves comparatively reliable spectral information during graph propagation.
- **Directional spectral readout.** DeStripe directly reconstructs the spectrum from the graph-updated representation. DORIAN instead decomposes reconstruction into directional gradients: the DPSN-affected direction is predicted from the graph-propagated representation, whereas the comparatively intact direction is read from the pre-propagation representation. The two are subsequently integrated to recover the restored spectrum, thereby reducing unnecessary modification of reliable directional content and improving structure preservation.
- **Structure-preserving regularization.** DeStripe preserves image structure through an unfolded Hessian prior with Bregman iterations, which increases model complexity and inference cost. DORIAN instead uses anisotropic total variation to constrain directional variations and retains the Hessian prior as an explicit loss term, avoiding deep unfolding while preserving fine image structures with lower computational cost.
- **Fidelity loss.** DeStripe uses a Self2Self-based denoising fidelity together with a handcrafted Fourier-domain isotropy prior. DORIAN instead adopts a guided-filtering-based fidelity that constrains coarse-scale intensity consistency while allowing DPSN-contaminated local variations to be corrected, providing a simpler optimization target.
- **Structure-preserving full-resolution recovery.** DeStripe operates directly at the target resolution. DORIAN instead performs optimization after anisotropic downsampling along the DPSN direction, which helps distinguish persistent degradation from similarly oriented image structures and reduces the optimization time. The estimated correction is then transferred back to full resolution through residual transfer with directional median filtering.
- **DC component modeling.** DeStripe does not explicitly model the DC component, whereas DORIAN introduces a separate branch to estimate it directly from the observation, explicitly constraining the global intensity level of the reconstruction.

D EXPERIMENTAL SETUP

D.1 DATASETS

For quantitative evaluation, we use two complementary natural-image benchmarks, DIV2K (Agustsson & Timofte, 2017) and BSD68 (Roth & Black, 2005). DIV2K contains high-resolution images with rich structural details and diverse scene content, while BSD68 provides a standard low-level vision benchmark with different textures and local image statistics. Together, they allow us to assess the robustness of DORIAN across heterogeneous image content and resolutions. We use the 100 images from the DIV2K validation set and all 68 images from BSD68 dataset. Starting from these clean images, we generate the corresponding degraded observations under three representative DPSN settings: periodic full-length, heterogeneous full-length, and heterogeneous finite-length DPSN. The detailed degradation settings are described below. We further evaluate DORIAN on a 3D fluorescence microscopy dataset of adipose tissue (Mai et al.), which extends the controlled evaluation from 2D natural images to volumetric microscopy data with substantially different image statistics and structural characteristics. This DPSN-free volume, consisting of $95 \times 2076 \times 1728$ voxels, is used as the clean reference, and synthetic DPSN are superimposed to generate the corresponding degraded observations.

In addition, we evaluate real DPSN across remote sensing, LSFM, and AFM. We use 15 remote-sensing images from the public ERSDNet dataset on IEEE DataPort[2], the `D2_I0.tif` zebrafish volume ($250 \times 1024 \times 1024$ voxels) from a public X-SPIM dataset[3], and five c-AFM images from the public dataset of Li et al. (Li et al.). These data contain naturally occurring DPSN with modality-dependent appearance, extent, and intensity variation. Since clean references are unavailable, they are used to assess the practical applicability of DORIAN across imaging modalities and real-world DPSN characteristics.

D.2 SYNTHETIC DEGRADATION.

We construct three synthetic DPSN settings with progressively increased spatial heterogeneity.

- **Periodic full-length.** We generate periodic DPSN by sampling a one-dimensional attenuation profile with regular spatial repetition for each 2D observation. The profile is extended along the orthogonal spatial dimension, yielding directional corruption that persists across the full image range. Within each observation, the spacing, width, and attenuation strength are kept uniform: the period is fixed to 20 pixels, the width of each attenuated region is set to 20% of the period, and the retained intensity is fixed at 50% of the original signal. A random phase shift varies the absolute DPSN locations across images, and the attenuation boundaries are slightly smoothed to avoid abrupt intensity transitions.
- **Heterogeneous full-length.** We next introduce spatial heterogeneity by allowing the spacing, width, and attenuation strength of individual DPSN components to vary within each observation. For each observation, the mean DPSN spacing is set according to the image width to produce approximately 12 components. Individual spacings are then randomly varied by up to 25% around this mean. The width of each component is sampled between 10% and 30% of its local spacing, while the retained intensity is independently sampled between 30% and 60%. All DPSN components extend across the full image range along the persistent direction.
- **Heterogeneous finite-length.** We further increase the spatial heterogeneity by allowing individual DPSN components to persist over only a finite extent along the directional axis, i.e., vertical. The spacing, width, and attenuation strength follow the same sampling strategy as in the heterogeneous full-range setting, while the longitudinal extent of each component is sampled between 30% and 80% of the image range. To avoid imposing artificial rectangular terminations, the start and end boundaries are anchored to strong local image gradients and tracked across the component width, allowing the finite-length DPSN to follow underlying image structures. The resulting boundaries are slightly smoothed while preserving their alignment with local gradients.

[2] `https://ieee-dataport.org/documents/ersdnet-multimodal-image-striping-noise-removal`
[3] `https://zenodo.org/records/14215090`

D.3 BASELINES

We compare DORIAN with six representative methods for DPSN removal: (1) model-based methods in the frequency domain, including WFF (Münch et al., 2009); (2) model-based approaches in the spatial domain, including GSR (Rottmayer et al., 2025); (3) hybrid spatial–spectral model-based methods, represented by NR-CS (Schwartz et al., 2019); and (4) learning-based methods, including ASCNet (Yuan et al., 2025), Self-BSR (Qu et al., 2025), and DeStripe (Liu et al., 2022). Among learning-based approaches, ASCNet uses supervised learning with paired clean-degraded data, while Self-BSR employs dataset-level self-supervised learning on synthetically degraded training images. In comparison, DeStripe is a zero-shot spatio-spectral model requiring no external training data. For ASCNet and Self-BSR, we retrain the models from scratch using 300 images randomly selected from the DIV2K training set, following the authors' official implementations, recommended training settings, and degradation models, so that each method is trained under the degradation model it was designed for.

D.4 METRICS

For synthetic experiments, since the applied DPSN attenuation field is known, we separately evaluate DPSN restoration and overall image fidelity. First, we report DPSN-region Peak Signal-to-Noise Ratio ($\mathrm{PSNR_D}$), computed only over pixels affected by the simulated DPSN, to measure the ability of each method to recover the corrupted signal. Meanwhile, we report global Structural Similarity Index Measure (SSIM) over the entire image as a complementary measure of overall structural fidelity, reflecting both residual DPSN and restoration-induced changes to the underlying image content.

D.5 METRIC FOR EVALUATION IN REAL DATA

For real observations, clean references are unavailable, preventing the use of full-reference metrics such as PSNR and SSIM. We therefore use the absolute correlation between the removed component and the restored image as a complementary measure of structure leakage. Specifically, given a degraded observation Y and restoration $\hat{X}$, we first define the removed component as:

$$R = Y - \hat{X} \tag{17}$$

based on which we compute:

$$|\rho| = \left| \frac{\sum_i (R_i - \bar{R})(\hat{X}_i - \bar{\hat{X}})}{\sqrt{\sum_i (R_i - \bar{R})^2 \sum_i (\hat{X}_i - \bar{\hat{X}})^2}} \right| \tag{18}$$

where i indexes all pixels, and $\bar{R}$ and $\bar{\hat{X}}$ denote their respective spatial means.

In general, a high $|\rho|$ indicates that the removed component is strongly correlated with structures retained in the restoration, suggesting potential removal of image content. Conversely, a low $|\rho|$ suggests less removal of image-correlated structure. Importantly, $|\rho|$ does not quantify residual DPSN and is therefore interpreted jointly with the qualitative restoration results.

D.6 IMPLEMENTATION DETAILS

In DORIAN, all hyperparameters are inherited unchanged from a prior LSFM destriping system and were fixed before any experiment in this paper was conducted. No parameter was selected or adjusted on DIV2K, BSD68, or any dataset used for evaluation. Specifically, we use a spectral angular half-width of $\Delta_\phi = 29^\circ$ and $N = 16$ graph neighbors. Regarding the restoration objective, the Hessian derivative kernels are generated using a Gaussian scale of $\sigma = 1$. Both the TV and Hessian derivative kernels are normalized to standardize their respective derivative scales beforehand, therefore $\lambda_{\mathrm{TV}} = \lambda_{\mathrm{H}} = 1$. Meanwhile, for the guided-filtering-based fidelity term, we use a window size of one twentieth of the image size and $\epsilon_{\mathrm{GF}} = 1$. Finally, high-resolution images are optimized with an anisotropic downsampling factor of $r = 3$, followed by residual-based full-resolution reconstruction using a directional median window of 49 pixels. As a zero-shot framework, DORIAN is optimized independently for each observation for 300 iterations using complex-valued Adam with a learning

rate of 10^{-2}, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$. Note that all parameters are kept fixed across experiments unless explicitly varied in the corresponding sensitivity analysis.

Following Eq. (1), multiplicative degradations are handled in the log-intensity domain, that is, the observation is transformed before optimization and the restoration is mapped back by the exact inverse. Whether a degradation is additive or multiplicative is known throughout, by construction for synthetic data and from the physical mechanism of the modality for real data (Appendix A), and DORIAN is applied accordingly.

D.7 Computational Efficiency

On a representative 1356×2040 image, DORIAN requires 12 s with the default $r = 3$, compared with 20 s for full-resolution optimization ($r = 1$), measured on a single NVIDIA Tesla T4 GPU. This confirms that anisotropic resampling reduces the optimization time of per-observation optimization. Runtime comparisons with baselines implemented in different languages and backends are omitted to avoid confounding algorithmic and implementation costs.

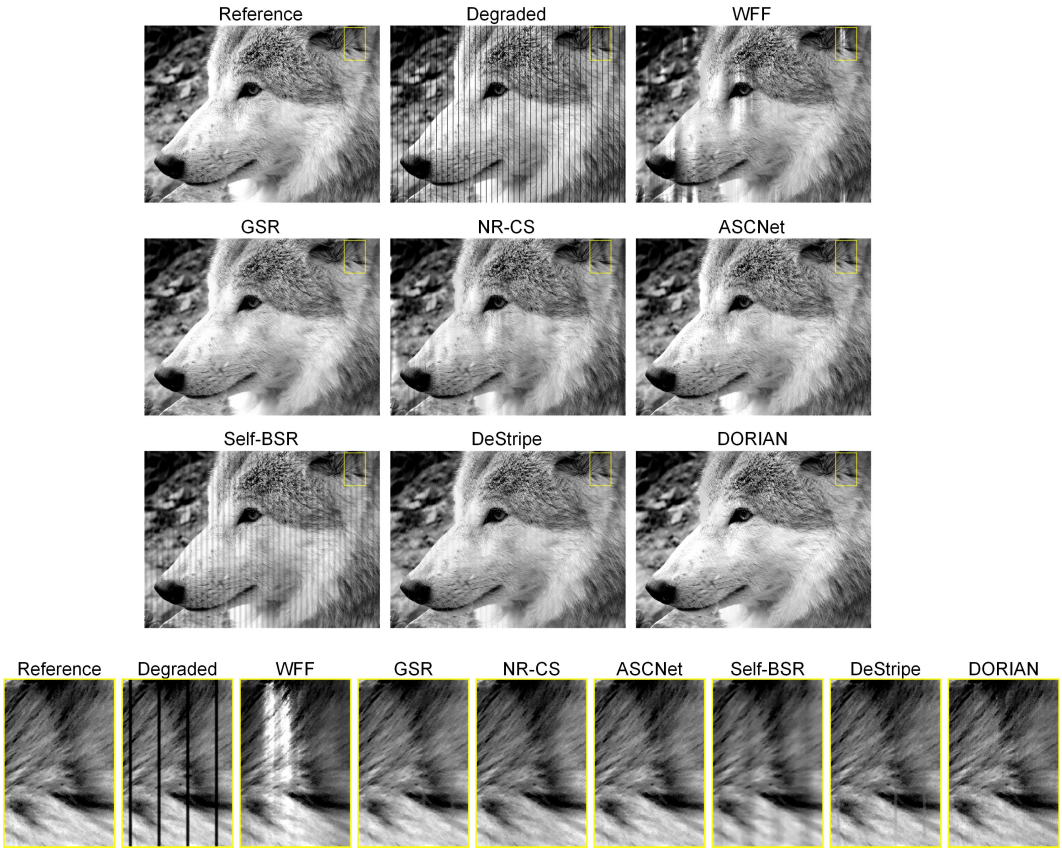


Figure 8: **Qualitative comparison on DIV2K under periodic full-length DPSN.** All methods suppress the degradation well at the image scale in this simplest setting, but the magnified region shows that the spectral filtering of WFF replaces the fur texture by a bright band. The yellow box indicates the region enlarged in the bottom row.

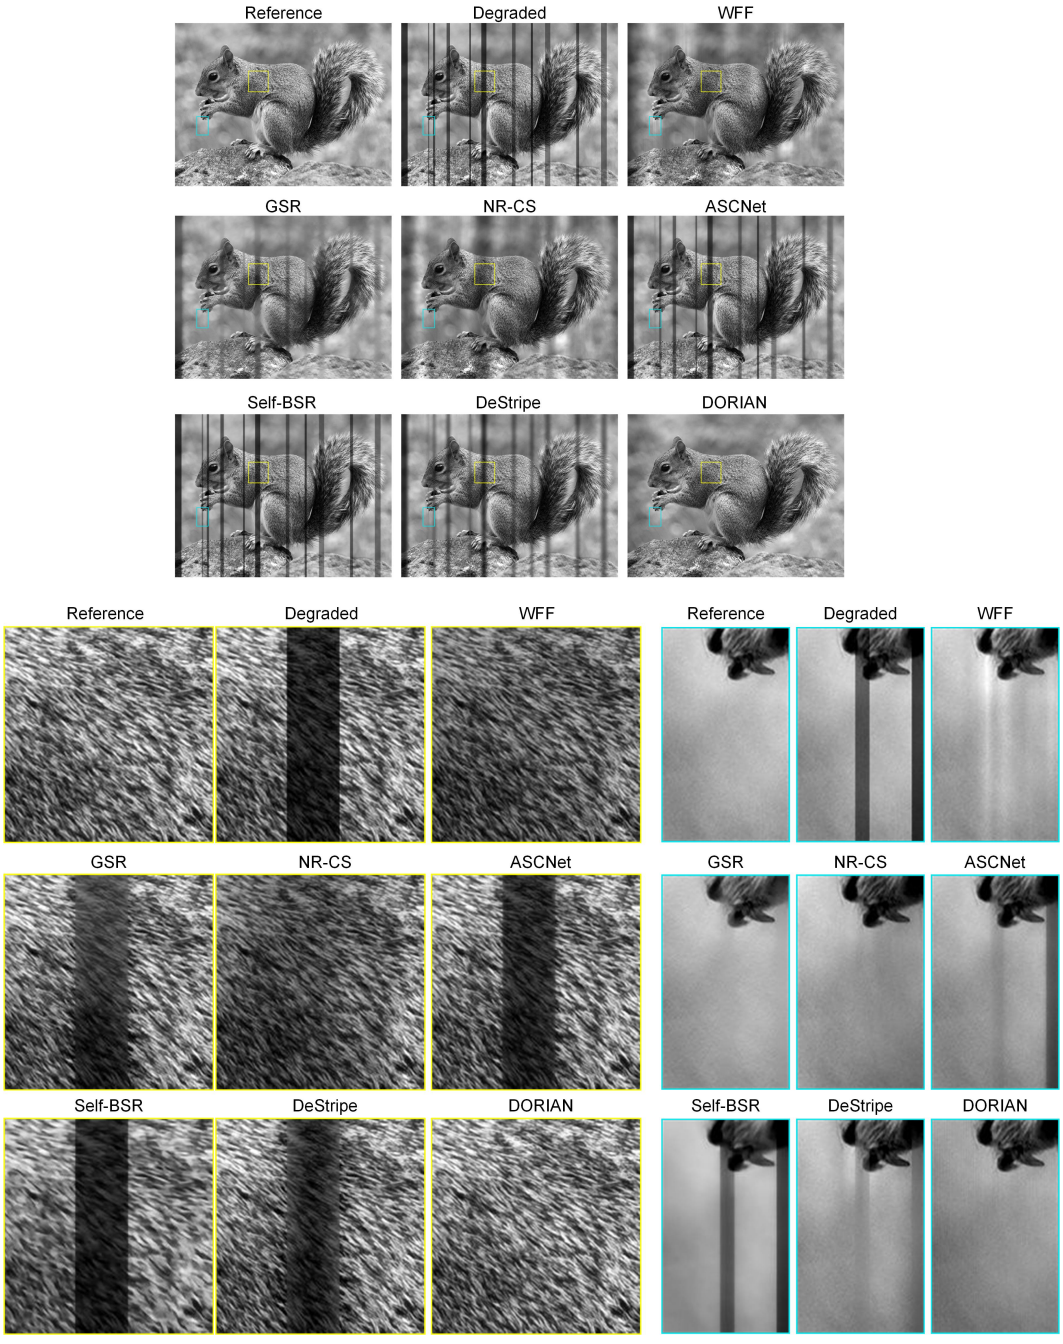


Figure 9: **Qualitative comparison on DIV2K under heterogeneous full-length DPSN.** Only DORIAN suppresses the degradation at the image scale (top rows). The yellow region shows that GSR, ASCNet and Self-BSR leave pronounced residuals, while the cyan region shows that the spectral filtering of WFF overshoots into a bright band in place of the removed DPSN. Yellow and cyan boxes indicate the two regions enlarged below.

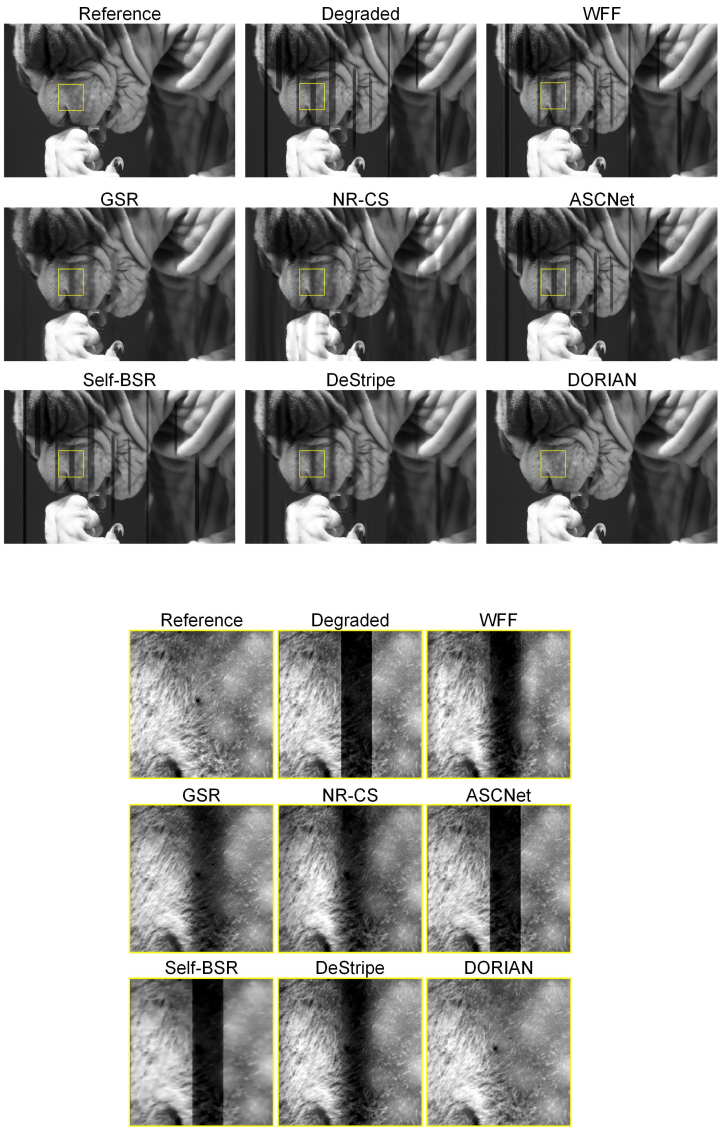


Figure 10: **Qualitative comparison on DIV2K under heterogeneous finite-length DPSN.** All competing methods leave a residual DPSN in the magnified region, whereas DORIAN restores the region successfully. The yellow box indicates the region enlarged below.

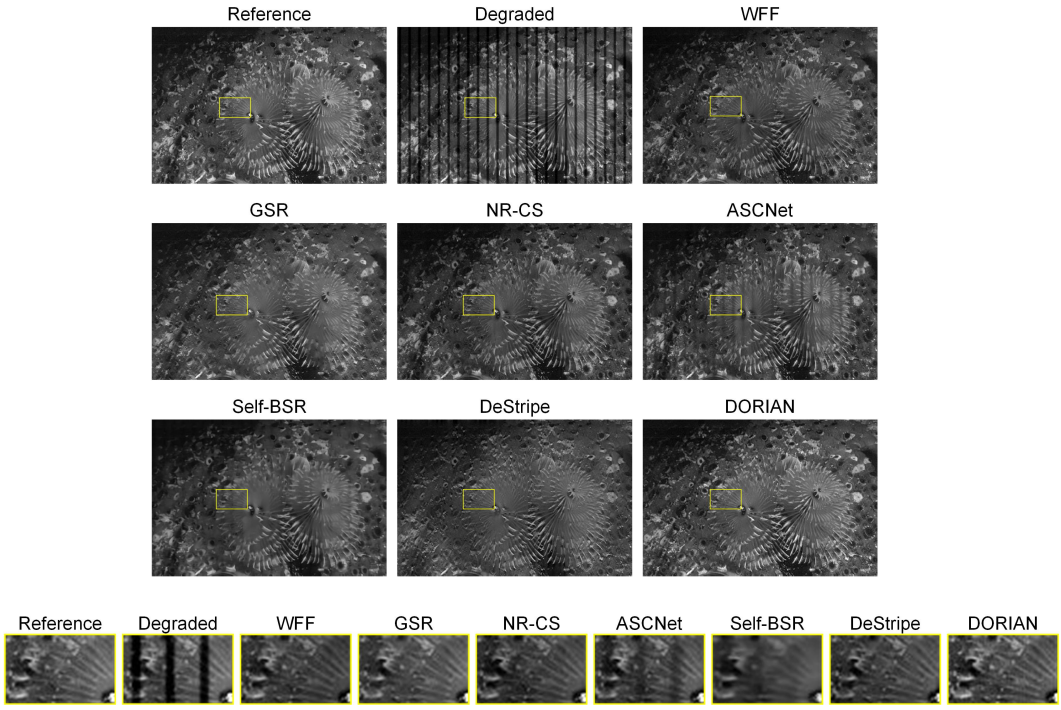


Figure 11: **Qualitative comparison on BSD68 under periodic full-length DPSN.** Although this is the simplest of the three settings, ASCNet, which is trained on DIV2K, degrades once applied to a different image distribution and leaves visible residuals, whereas the training-free methods are unaffected by the change of dataset. DORIAN removes the DPSN while retaining the fine radial structure of the object. The yellow box indicates the region enlarged in the bottom row.

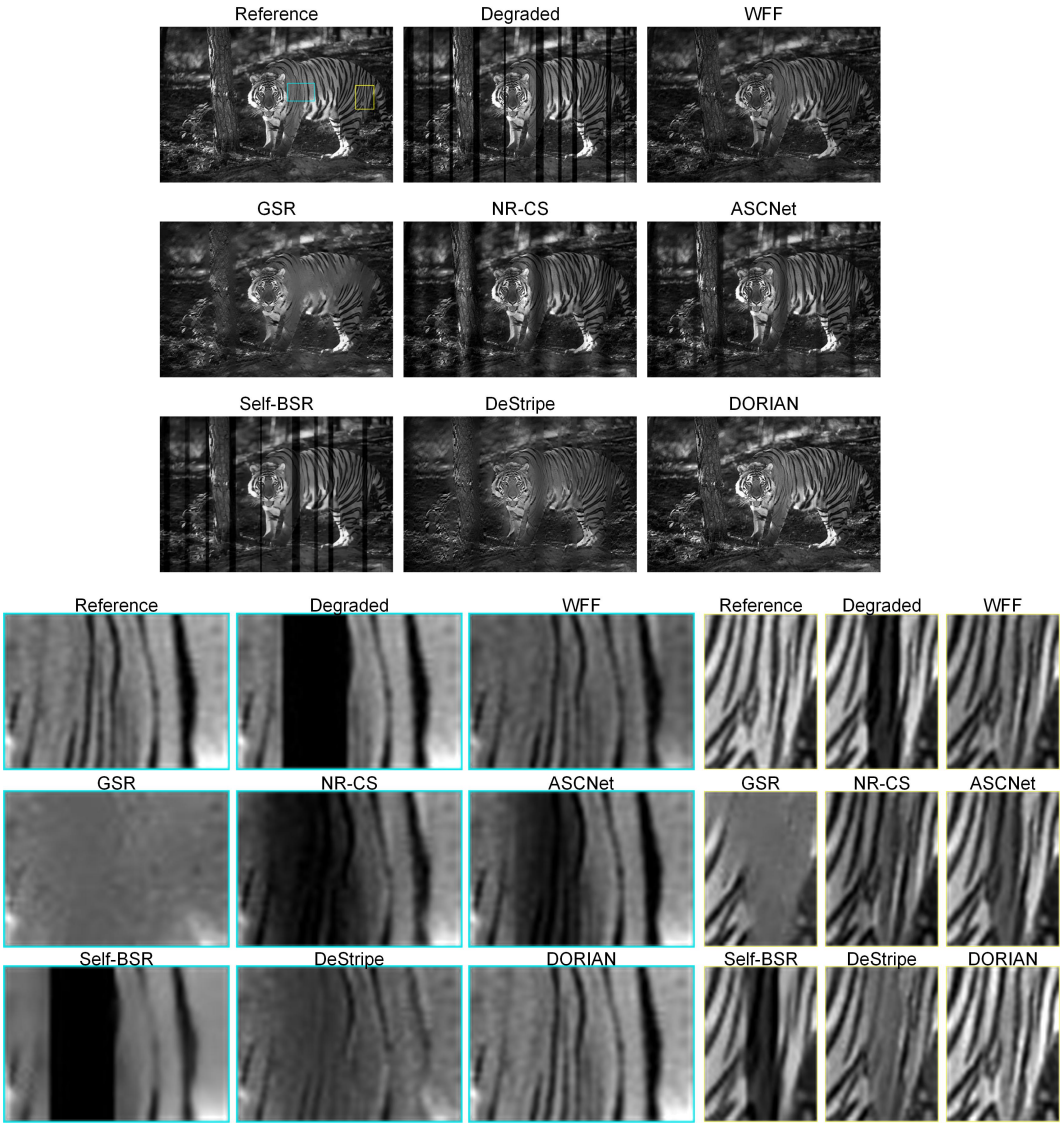


Figure 12: **Qualitative comparison on BSD68 under heterogeneous full-length DPSN.** GSR appears to suppress the degradation well at the image scale (top rows), yet both magnified regions reveal that it removes the tree bark and the fur pattern along with the DPSN. DORIAN suppresses the DPSN while retaining these structures. Cyan and yellow boxes indicate the two regions enlarged below.

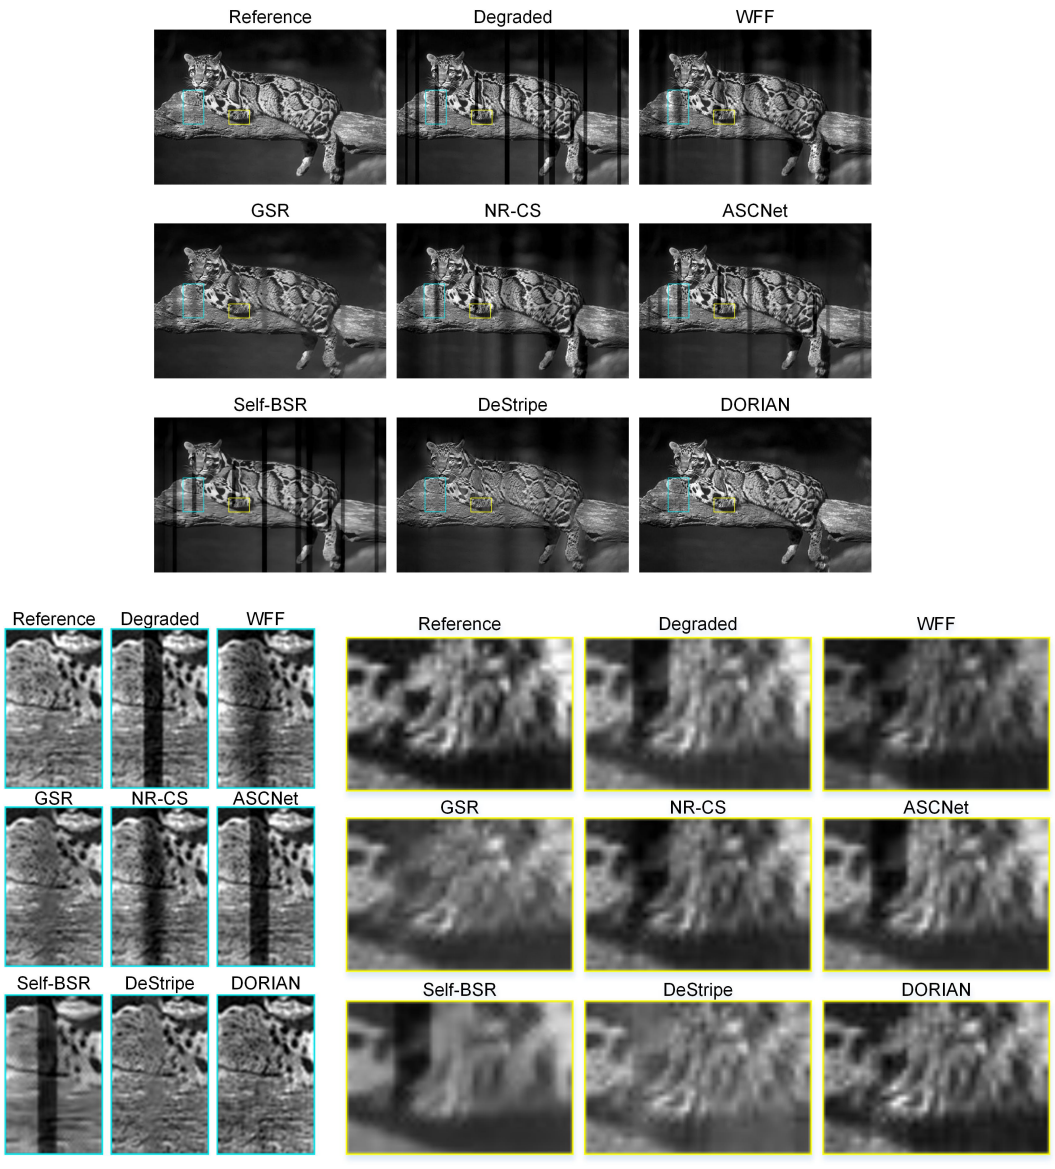


Figure 13: **Qualitative comparison on BSD68 under heterogeneous finite-length DPSN.** Only DORIAN removes the degradation without leaving visible residuals across the image (top rows). In the cyan region DeStripe suppresses the DPSN comparably well, but the yellow region shows that it does so at the cost of losing fine image detail, which DORIAN retains. Cyan and yellow boxes indicate the two regions enlarged below.

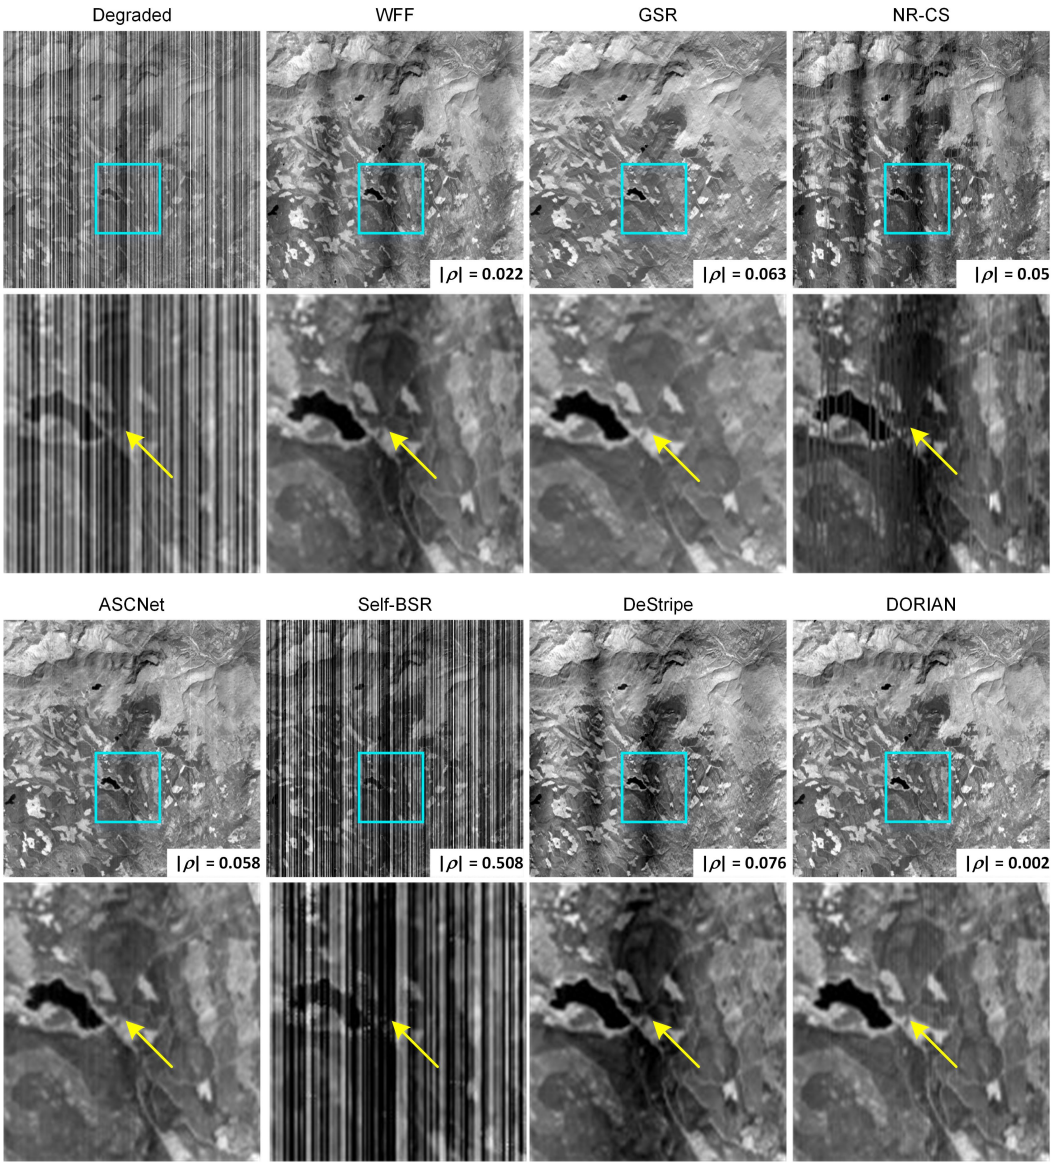


Figure 14: **Qualitative comparison on real remote sensing data.** Since no clean reference exists, each panel reports $|\rho|$, the absolute correlation between the removed component and the restored image; a value near zero indicates that the removed component carries no scene content. Self-BSR leaves the degradation almost intact and its $|\rho| = 0.508$ reflects a failure to separate the DPSN from the image, while NR-CS retains pronounced residuals despite a low $|\rho|$, since the measure is insensitive to under-removal. DORIAN suppresses the DPSN and preserves the road and the lake boundary indicated by the arrows. Cyan boxes indicate the regions enlarged below each image.

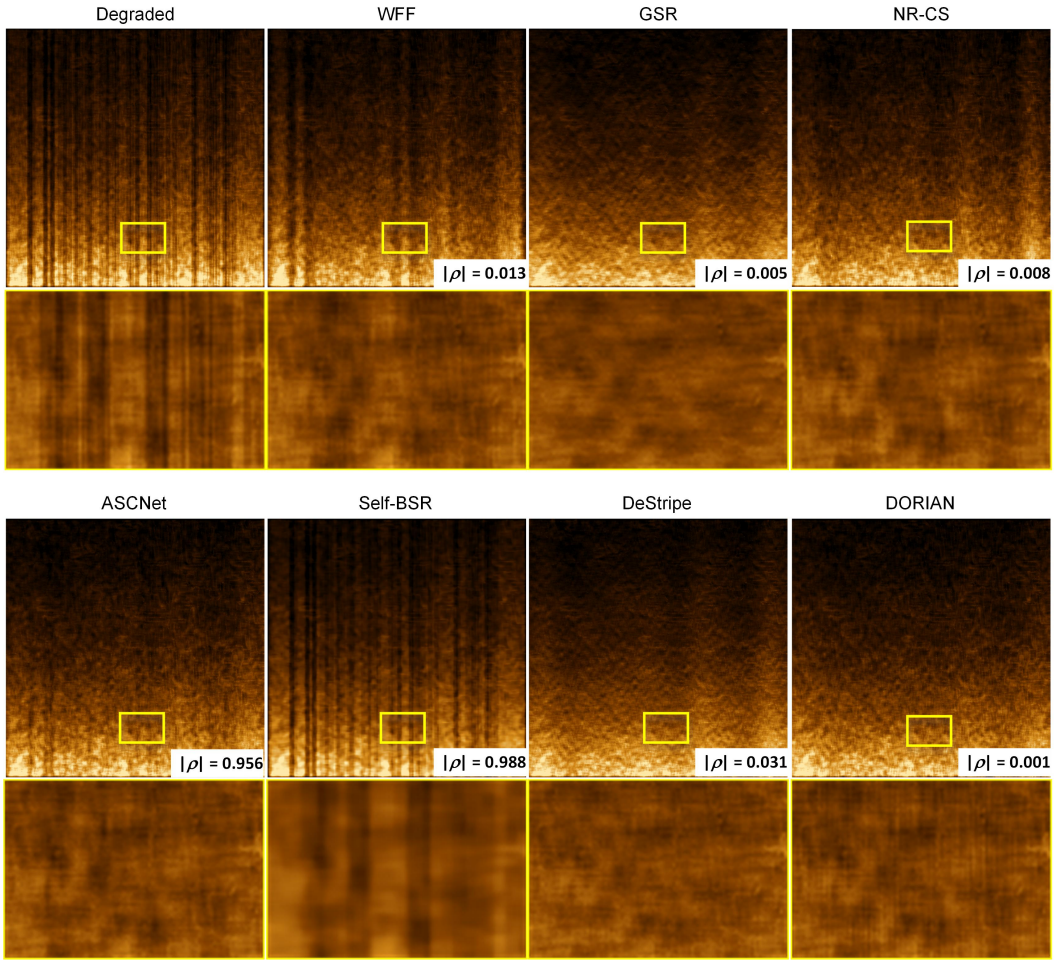


Figure 15: **Qualitative comparison on real AFM data.** Each panel reports $|\rho|$, the absolute correlation between the removed component and the restored image; a value near zero indicates that the removed component carries no sample structure. ASCNet and Self-BSR reach $|\rho| = 0.956$ and 0.988, meaning that what they subtract is almost entirely the sample itself, and both leave the DPSN largely intact. GSR and NR-CS attain a low $|\rho|$ but suppress the granular surface texture, illustrating that the measure is insensitive to over-smoothing that follows the degradation direction. DORIAN removes the DPSN at $|\rho| = 0.001$ while retaining the texture. Yellow boxes indicate the regions enlarged below each image.

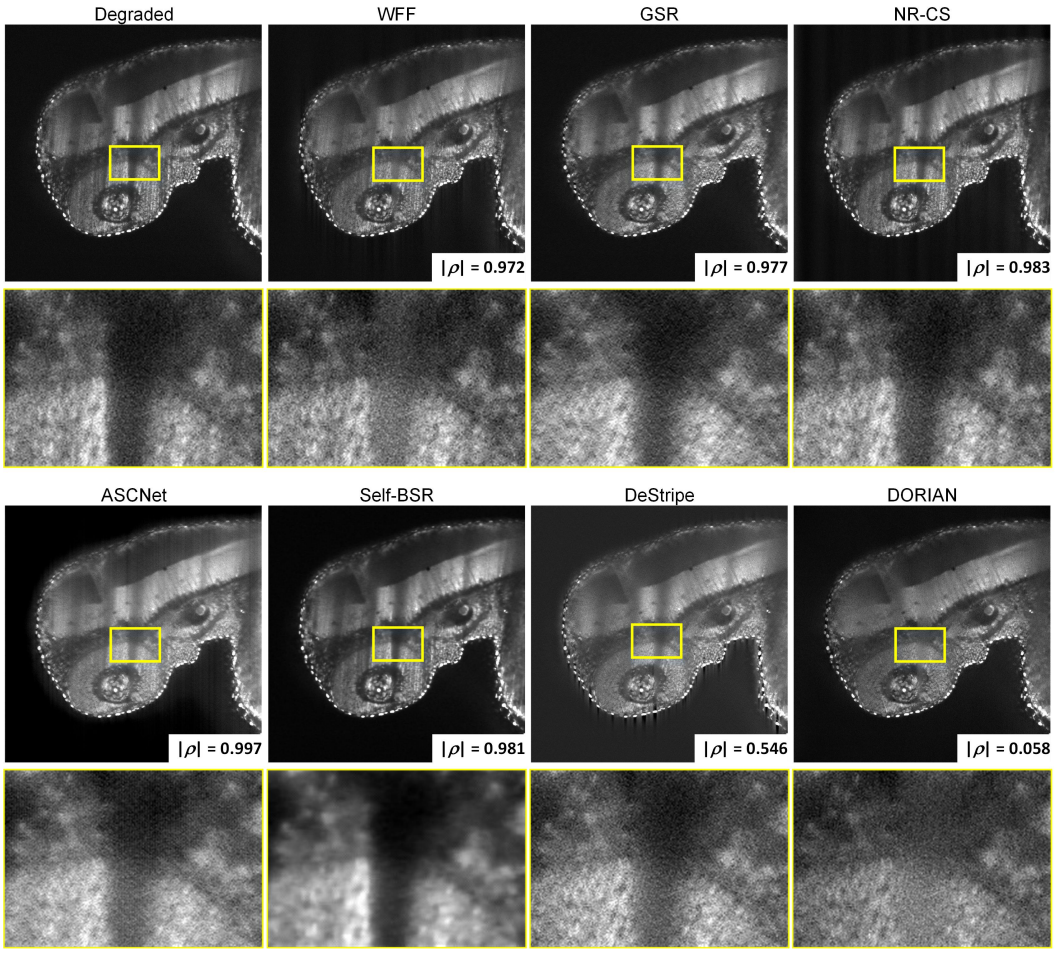


Figure 16: **Qualitative comparison on real LSFM data.** Each panel reports $|\rho|$, the absolute correlation between the removed component and the restored image; a value near zero indicates that the removed component carries no sample structure. The competing methods rescale the overall intensity, by factors between 0.07 and 0.54 relative to the observation, so that what they subtract is dominated by the sample itself and $|\rho|$ exceeds 0.97 for all of them except DeStripe. DeStripe additionally introduces structure that is absent from the observation: the background is lifted from black to grey and a row of dark vertical streaks appears below the lower boundary of the specimen. DORIAN preserves the intensity scale to within 1.2%, attains $|\rho| = 0.058$, and removes the absorption shadow in the magnified region while retaining the granular texture. Yellow boxes indicate the regions enlarged below each image.

E IMAGE-SPACE U-NET

The image-space U-Net in Table 2 replaces the DORIAN core by a real-valued U-Net acting directly on the observed image, so that the comparison isolates a non-local spectral representation against a spatial convolutional one.

Architecture. The network is a five-level U-Net with skip connections, two 3×3 convolutions per level, GELU activations, 2×2 max pooling and 2×2 transposed-convolution upsampling. Decoder features are resized to the exact encoder resolution before concatenation, so that the odd spatial sizes produced by the anisotropic downsampling remain valid. The network predicts a global residual, $\widehat{X}^{\downarrow} = Y^{\downarrow} + \mathrm{UNet}(Y^{\downarrow})$, so that optimization starts from the observation. The directional residual transfer of Eq. (11) is retained unchanged and carries no learnable parameters, so that the ablation isolates the restoration core alone.

Parameter matching. Since the graph prior holds complex weights, its 5,908,338 coefficients at 1356×2040 correspond to 11,816,676 real degrees of freedom. With a base width of 39 channels, doubling at every level up to 624 at the bottleneck, the U-Net holds 11,524,618 real parameters, comparable to that budget.

Optimization. Unlike the per-edge weights of the graph prior, each used once, a convolutional network shares its kernels across the image and does not converge under the default schedule. The learning rate and the number of iterations were therefore grid-searched over $\{10^{-2}, 3{\times}10^{-3}, 10^{-3}, 3{\times}10^{-4}, 10^{-4}\}$ and $\{300, 1000, 2000\}$, and the best configuration is reported, while DORIAN uses its default setting of learning rate of 10^{-2} for 300 iterations.

F FOURIER U-NET

The representation ablation in Table 2, i.e., Fourier U-Net, replaces the DORIAN core by a complex-valued U-Net acting on the Cartesian spectrum. The comparison is therefore between a polar graph neighborhood and a Cartesian convolutional neighborhood, both operating on the same complex spectrum with the same complex arithmetic and producing the restored spectrum from the same input.

Architecture. The network is a five-level complex U-Net with skip connections, two 3×3 complex convolutions per level, $\mathrm{ReLU}^{\mathbb{C}}(\cdot)$ activations, average pooling, and bilinear upsampling followed by a 1×1 channel reduction. Following Trabelsi et al. (2017), a complex convolution is realized as two real convolutions, $(W_r + iW_i) * (x_r + ix_i) = (W_r * x_r - W_i * x_i) + i\,(W_r * x_i + W_i * x_r)$, so that one complex weight costs two real parameters, matching the convention used for the complex edge weights $\alpha_{p,q}$ of the Fourier graph prior.

Three implementation details are necessary for convergence and ensure that the reported performance reflects the architecture rather than numerical issues. *(i) Log-magnitude companding.* Fourier magnitudes span many orders of magnitude, and a convolutional network fed the raw coefficients does not converge. Logarithmic compression of spectral magnitudes is therefore a standard practice when spectra are used as network inputs. In effect, although the compression is itself a simple function, it is difficult to learn implicitly when embedded in a larger task (Choi et al., 2018). The input is therefore compressed as $\log(1 + |z|/s)$ with the phase preserved, where s is the mean magnitude of the observed spectrum, and the output is expanded by the exact inverse. *(ii) Circular padding.* The fftshifted DFT grid is periodic, so circular padding is its natural continuation; reflection or zero padding would introduce artificial spectral edges. *(iii) Near-identity initialization.* The U-Net predicts a residual in the compressed domain and the final 1×1 convolution is scaled by 0.1 at initialization, so that optimization starts from the observed spectrum, which is the starting point the graph prior obtains through its self-loops.

Parameter matching. The graph prior stores one learnable complex weight per edge, so its size grows with the input shape, whereas that of a U-Net does not. A matched budget is therefore matched at one resolution only. Specifically, with $\Delta_\phi = 29^\circ$, $N = 16$ and $r = 3$, the Fourier graph prior holds 5,908,338 complex weights at 1356×2040, the average image size in DIV2K. The

channel configuration $(29, 58, 116, 232, 464)$ used throughout holds 5,944,218 complex parameters, comparable to that of the Fourier graph prior.

Optimization. The learning rate and the number of iterations of the baseline were grid-searched over $\{3\times10^{-2}, 10^{-2}, 3\times10^{-3}, 10^{-3}\}$ and $\{300, 1000, 2000\}$, and the best configuration is reported, while DORIAN uses its default setting of 10^{-2} for 300 iterations.

G Additional Ablation Studies

Effect of the directional gradient readout. Table 2 evaluates the decomposition of the spectral readout into two directional branches. By relying not only on the graph-completed representation $\widetilde{U}$ but also on the intact pre-propagation feature map U, DORIAN improves over direct readout, which relies on $\widetilde{U}$ alone, by 3.49 dB and 0.0494 SSIM. As shown in Fig. 22, both suppress the residual DPSN comparably, yet image texture is better preserved by the directional gradient readout (see zoom-in).

H Parameter Sensitivity

Fig. 7 varies three parameters that trade reconstruction capacity against model complexity. (i) An overly narrow spectral wedge restricts the number of Fourier coefficients available for adaptive reconstruction, and $\mathrm{PSNR_D}$ rises from 13.78 to 22.46 dB as Δ_ϕ increases from 5° to 29°, beyond which it saturates while the learnable weights keep growing from 4.15 to 4.59 million. (ii) A similar trend holds for the number of references N drawn from the annular candidate set $\mathcal{C}(p)$, which governs how many spectrally related coefficients each DPSN-dominated node aggregates: performance improves abruptly from 18.75 to 22.12 dB once eight references are considered, and beyond $N = 8$ saturates within 0.7 dB and 0.01 SSIM up to $N = 128$, whereas the learnable weights grow linearly with N. (iii) The restoration is likewise stable as the anisotropic downsampling factor r varies from 1 to 5, and degrades only once the resampling becomes coarse enough to remove genuine structures together with the DPSN, falling to 21.71 dB at $r = 9$. Note that all hyperparameters are inherited unchanged from a prior light-sheet destriping system and lie well inside the respective plateaus, so that a single fixed configuration suffices across observations and modalities.

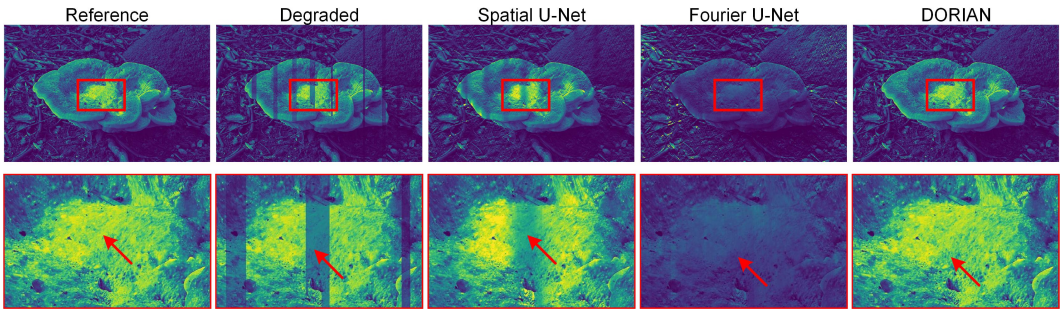


Figure 17: **Ablation of the polar Fourier representation.** The spatial U-Net leaves residual DPSN (arrows), whereas the Fourier U-Net deteriorates the image content (zoom-in). DORIAN suppresses the DPSN while preserving the fine structure of the reference. Red boxes indicate the regions enlarged in the second row; all panels share the same display range.

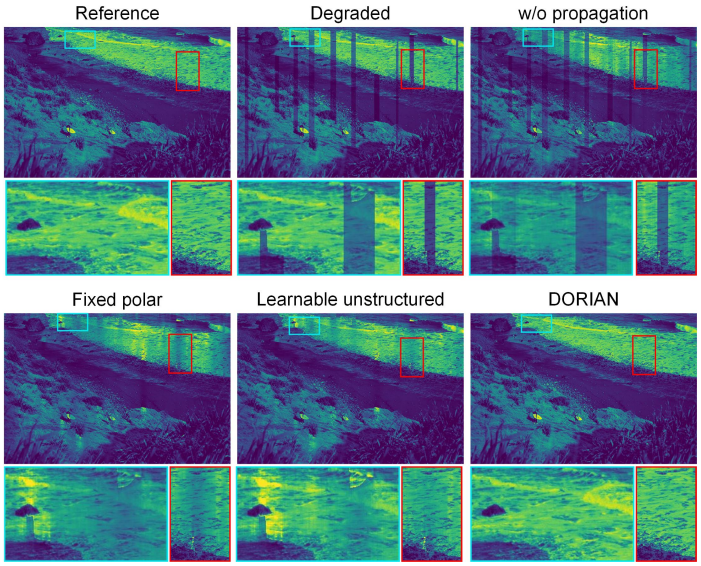


Figure 18: **Ablation study of the graph neural prior.** We compare three variants of DORIAN in which the graph neural prior is removed (*w/o propagation*), replaced by a fixed polar neighborhood (*Fixed polar*), or replaced by a learnable unstructured neighborhood definition (*Learnable unstructured*). DORIAN more effectively suppresses directional DPSN while preserving local image structure and texture. The cyan and red boxes indicate two representative regions enlarged below each image, highlighting distorted image contents and residual DPSN signals, respectively.

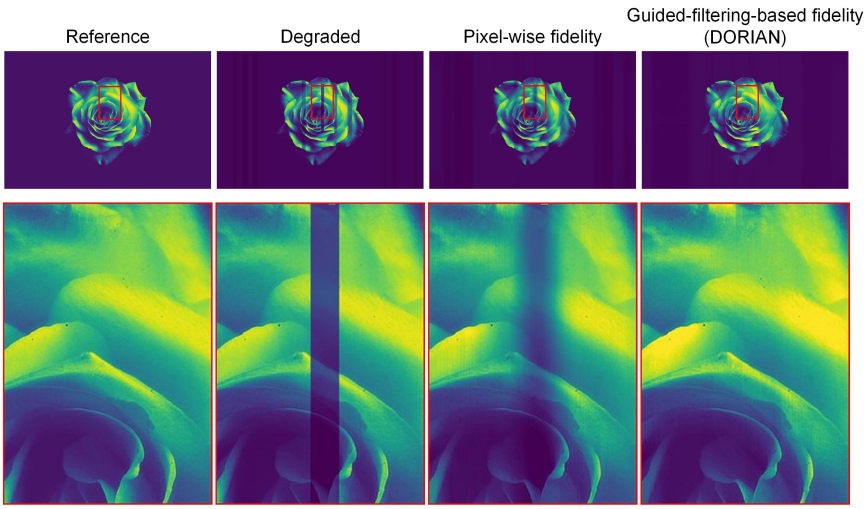


Figure 19: **Qualitative comparison of fidelity formulations.** Direct pixel-wise ℓ_2 fidelity retains visible residual DPSN by enforcing local agreement with the degraded observation. In contrast, the guided-filtering-based fidelity used in DORIAN constrains coarse-scale intensity consistency while allowing DPSN-contaminated local variations to be corrected. Red boxes indicate the regions enlarged in the second row.

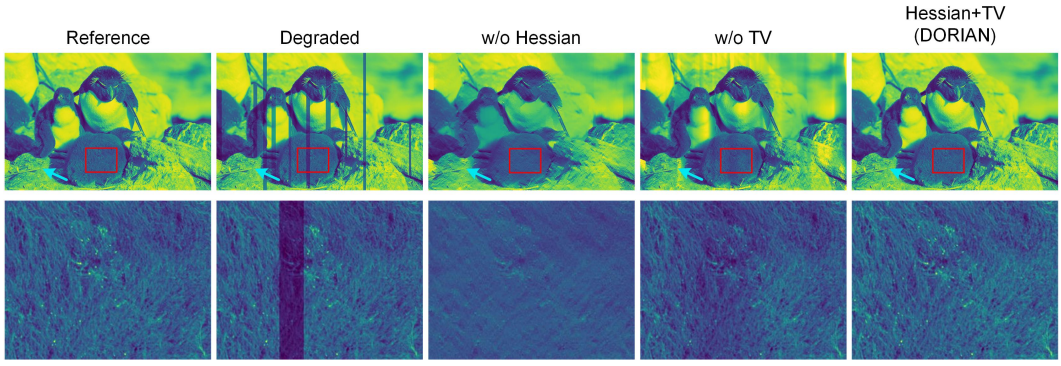


Figure 20: **Qualitative ablation of the TV and Hessian regularization terms.** Cyan arrows indicate residual DPSN, which remains visible when the TV term is removed. Red boxes mark the regions enlarged in the second row, where removing the Hessian term leads to substantial loss of fine structural details. Combining TV and Hessian regularization in DORIAN suppresses residual DPSN while preserving fine image structures.

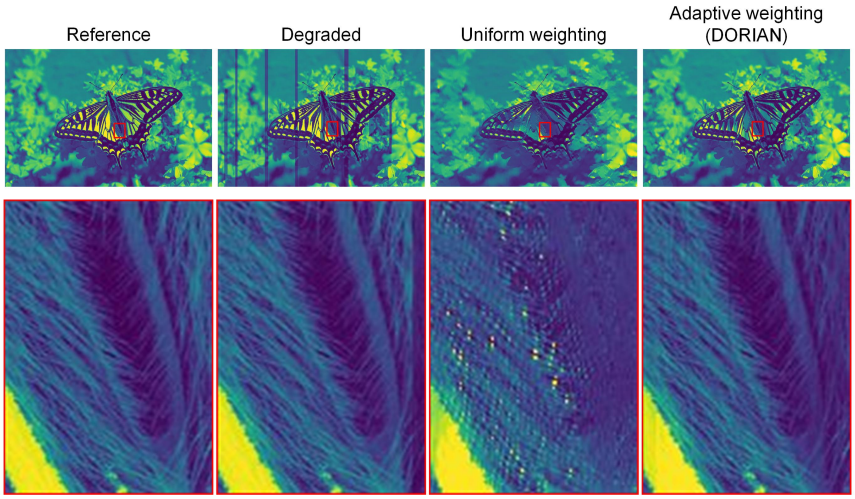


Figure 21: **Effect of adaptive regularization weighting.** Compared with uniform weighting, the adaptive weighting strategy used in DORIAN better balances DPSN suppression and structural preservation. Uniform regularization tends to over-smooth fine textures and locally distort image structures, whereas adaptive weighting preserves salient details while effectively removing DPSN.

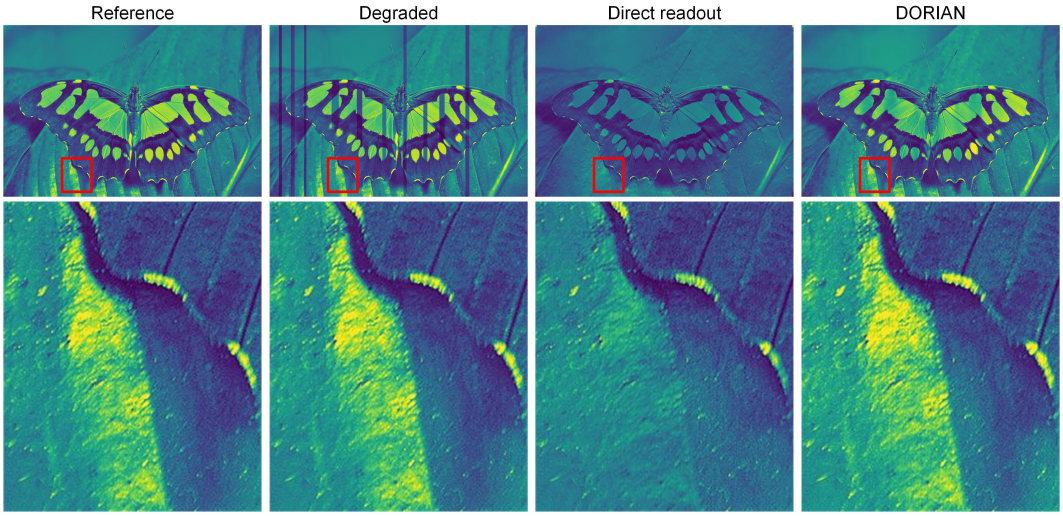


Figure 22: **Effect of the directional gradient readout.** *Direct readout* projects the restored spectrum entirely from the graph-propagated representation $\widetilde{U}$, whereas DORIAN predicts the gradient along $\mathbf{d}_{\|}$ from the pre-propagation representation U. Both variants share the same graph prior and suppress the DPSN well (top row), whereas, in the magnified region, direct readout attenuates image texture. All panels share the same display range.