CONTROLLING RE-CORRUPTION BIAS FOR UNSUPERVISED IMAGE DENOISING WITH UNKNOWN NOISE

Anonymous authors
Paper under double-blind review

ABSTRACT

Unsupervised image denoising via re-corruption constructs noisy/noisy training pairs whose loss can reproduce supervised denoising under suitable noise assumptions. Existing guarantees, however, require sufficiently specified noise statistics, which are often unavailable in practice. We revisit re-corruption and identify a denoiser-independent criterion that characterizes the bias induced by a given re-corruption. This criterion depends only on a conditional moment of the re-corruption mismatch, yielding a weaker requirement than exact conditional noise matching and an equivalent variational formulation for learning re-corruption from noisy data alone. Based on this characterization, we first learn the re-corruption mechanism and then freeze it to fine-tune an initial denoiser, thereby decoupling re-corruption learning from denoiser optimization. By targeting the bias-relevant mismatch rather than estimating the full noise distribution, our method naturally accommodates unknown real-world noise. Experiments on smartphone and fluorescence-microscopy images show consistent denoising improvements and a clear reduction of the predicted bias.

1 INTRODUCTION

Image denoising is a fundamental problem in image processing and computer vision. It has been greatly advanced by deep learning (Zhang et al., 2017; 2018; Zamir et al., 2021; 2022). Most neural networks (NN) for denoising are trained from paired noisy/clean images, but reliable clean references can be expensive or impossible to acquire in real imaging systems. Noise2Noise (N2N) (Lehtinen et al., 2018) relaxes this requirement by using independent noisy observations of the same underlying image, yet such paired captures still require repeated acquisition and accurate alignment. These limitations have motivated learning denoisers directly from noisy observations alone.

Throughout this paper, we use *unsupervised denoising* to refer to dataset-level training from noisy images without clean targets, in contrast to zero-shot denoising, where a model is optimized separately for each test image. Existing methods mainly follow two paradigms. *Blind-spot methods* exclude the target observation from its own prediction and infer it from the remaining image context (Krull et al., 2019; Batson & Royer, 2019; Laine et al., 2019; Quan et al., 2020). *Data-augmentation-based methods* instead construct surrogate training pairs or losses directly from noisy observations (Moran et al., 2020; Pang et al., 2021; Huang et al., 2021; Zheng et al., 2023). Our work belongs to the latter class and focuses on re-corruption-based augmentation.

An important advantage of data-augmentation-based loss design is its flexibility as it does not require discarding observed measurements or imposing a particular blind-spot network architecture. More importantly, the same principle can extend to inverse problems by constructing self-supervision directly in the measurement domain to handle measurement noise (Quan et al., 2022; 2024). This flexibility motivates our focus on the re-corruption paradigm: improving its robustness to unknown noise can benefit not only denoising, but potentially a broader class of noisy inverse problems.

Motivation: What is required for valid re-corruption? Re-corruption methods retain the observed measurements and introduce additional randomness to construct noisy/noisy training pairs. R2R (Pang et al., 2021) established that the loss defined over re-corruption gives an unbiased estimator of the supervised loss, up to a constant. ER2R (Zheng et al., 2023) extends it to general random noise by requiring the injected noise z to be conditionally i.i.d. with the observation noise

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.

n given the clean image. This broader guarantee still requires sufficient knowledge of the noise distribution. L2R (Monroy et al., 2026) further learns re-corruption under unknown noise through a denoiser–re-corruptor saddle formulation.

These developments raise a more fundamental question:

> *What statistical property must a re-corruption satisfy so that its noisy/noisy training objective remains a faithful surrogate for supervised denoising?*

In this paper, we identify a weaker denoiser-independent criterion for re-corruption bias and show how to learn it from noisy observations alone.

Characterizing and controlling re-corruption bias. Let $y = x+n$ and consider the re-corruption pair $(y + z, y - z)$. For a denoiser F_θ, the corresponding loss satisfies

$$\mathcal{L}_{\mathrm{rec}}(\theta) = \mathbb{E}\left[\|F_\theta(y + z) - (y - z)\|_2^2\right] = \mathcal{R}(\theta) + \mathrm{const} + 2C_\theta, \tag{1}$$

where

$$\mathcal{R}(\theta) = \mathbb{E}\|F_\theta(x + n + z) - x\|_2^2, \qquad C_\theta = \mathbb{E}\Big[\big(F_\theta(x + n + z) - x\big)^\top (z - n)\Big]. \tag{2}$$

Thus, the gap from supervised training is determined by C_θ. To characterize this bias independently of a particular denoiser, we define

$$Q := \sup_{g\neq 0} \frac{\left|\mathbb{E}\big[g(x, n + z)^\top (z - n)\big]\right|^2}{\mathbb{E}\|g(x, n + z)\|_2^2}. \tag{3}$$

Since $g(x, n + z) = F_\theta(x + n + z) - x$ is an admissible choice,

$$C_\theta^2 \leq \mathcal{R}(\theta)Q, \tag{4}$$

so Q gives a denoiser-independent measure of re-corruption bias. A key result is that Q admits two equivalent representations:

$$Q = \mathbb{E}\,\|\mathbb{E}[z - n \mid x, n + z]\|_2^2 = \sup_{\psi\in L^2(x,n+z)} \left\{2\mathbb{E}\big[\psi(x, n + z)^\top (z - n)\big] - \mathbb{E}\|\psi(x, n + z)\|_2^2\right\}. \tag{5}$$

The first shows that full conditional noise matching is stronger than necessary: if z and n are conditionally i.i.d. given x, as in ER2R, then exchangeability gives $Q = 0$. Our criterion requires only this conditional-moment condition, rather than matching the full conditional noise distribution.

The second representation provides the basis for learning re-corruption. Since x and n are unavailable, we construct noisy-only proxies using an initial denoiser and parameterize ψ by a neural critic. The critic exposes the bias-relevant component of the re-corruption mismatch, while a learnable generator is optimized to suppress it. The generator is then frozen and used to fine-tune the denoiser, separating re-corruption learning from denoiser optimization.

Learning re-corruption bias control from noisy data. The oracle quantity Q in Eq. 5 cannot be evaluated directly, since neither the clean image x nor the observation noise n is available. We therefore use an initial denoiser f_0 to construct noisy-only proxies of both quantities, and define a corresponding surrogate of Q. To optimize this surrogate, we introduce a critic that estimates the conditional mean of the re-corruption mismatch. At the population level, the optimal critic recovers this conditional mean, yielding a variational formulation of the surrogate objective. A learnable re-corruption generator is then optimized against the critic to reduce the resulting bias measure.

Importantly, this optimization is used only to learn the re-corruption mechanism; the denoiser does not participate in the process. Once the generator is learned, we freeze it and use the resulting re-corrupted pairs to fine-tune the initial denoiser. This two-stage design separates re-corruption learning from denoiser optimization and targets only the statistical property needed to control the gap from supervised training, rather than attempting to recover the full observation-noise distribution.

The main contributions are summarized as follows:

- **Characterizing what makes re-corruption valid.** We identify the bias between re-corruption training and supervised denoising and derive a denoiser-independent sufficient condition to show that exact noise-distribution matching is stronger than necessary.
- **Learning re-corruption from noisy data alone.** Since noise realization and clean image are inaccessible, we construct a noisy-only surrogate and a critic-based variational formulation to learn a re-corruption mechanism without modeling the full noise distribution.
- **A decoupled two-stage denoising framework.** We first learn the re-corruption mechanism and then freeze it to fine-tune the denoiser. Experiments on smartphone and fluorescence-microscopy noise show consistent denoising gains.

2 Related Work

Risk estimation and implicit priors. One line of clean-target-free denoising estimates the supervised risk directly from noisy images. SURE-based methods (Soltanayev & Chun, 2018) use Stein-type identities to obtain unbiased risk estimates under suitable noise models, while later work relaxes some of the required noise parameters (Tachella et al., 2025). Another line exploits image priors induced by the model itself. Deep Image Prior (DIP) (Ulyanov et al., 2018) shows that an untrained network can act as an effective image prior when optimized directly on a noisy image.

Blind-spot denoising. Blind-spot denoising is a truth-free learning paradigm that can be used in both dataset-level and zero-shot settings. It predicts each target pixel without accessing its own noisy observation, thereby avoiding overfitting to the identity mapping. At the dataset level, this idea underlies Noise2Void (Krull et al., 2019), Noise2Self (Batson & Royer, 2019), and many architecture-based blind-spot networks Laine et al. (2019); Wu et al. (2020); Lee et al. (2022). The same principle can also be applied in the zero-shot setting, where the model is optimized on each test image, as in Self2Self (Quan et al., 2020) and others (Chihaoui & Favaro, 2024; Quan et al., 2025; Guo et al., 2025). Blind-spot methods avoid explicit modeling of the noise distribution, but their effectiveness depends on strong correlations of local image context and weak correlations of local noise.

Pseudo-pair construction via image redundancy. One data-augmentation approach constructs surrogate training pairs by exploiting redundancy of image structures. Neighbor2Neighbor (Huang et al., 2021) forms paired sub-images from neighboring pixels, while ZS-N2N (Mansour & Heckel, 2023) uses downsampled observations and Pixel2Pixel (Ma et al., 2025) exploits image-internal recurrence to obtain noisy-to-noisy supervision from a single image. These approaches require little or no explicit modeling of the observation noise. Their effectiveness, however, depends on the constructed pairs sharing sufficiently consistent underlying image content, which often does not hold true for scientific images.

Pseudo-pair construction via re-corruption. A second class of data-augmentation methods constructs noisy/noisy pairs by injecting random perturbations into noisy observations. Noisier2Noise (Moran et al., 2020) uses additional corruption, while R2R and its extensions (Pang et al., 2021; Zheng et al., 2023; Monroy et al., 2025) design the perturbation so that the resulting loss is statistically connected to supervised denoising. Although they progressively enlarge the admissible noise class, they still require the noise model to be specified sufficiently well. L2R (Monroy et al., 2026) instead learns the re-corruptor jointly with the denoiser under unknown noise. In contrast, we first characterize re-corruption bias through a denoiser-independent quantity Q, learn a re-corruption mechanism to reduce its noisy-only surrogate, and then freeze it for denoiser fine-tuning.

Table 1: Comparison of related re-corruption methods for unsupervised denoising.

Method	Main assumption / formulation
R2R	Gaussian noise with prescribed covariance structure
ER2R	Injected noise sampled from the same conditional noise law
GR2R	Broader but still distribution-specific re-corruption
L2R	Unknown noise; Denoiser–re-corruptor saddle formulation
Ours	Unknown noise; denoiser-independent bias criterion and decoupled learning

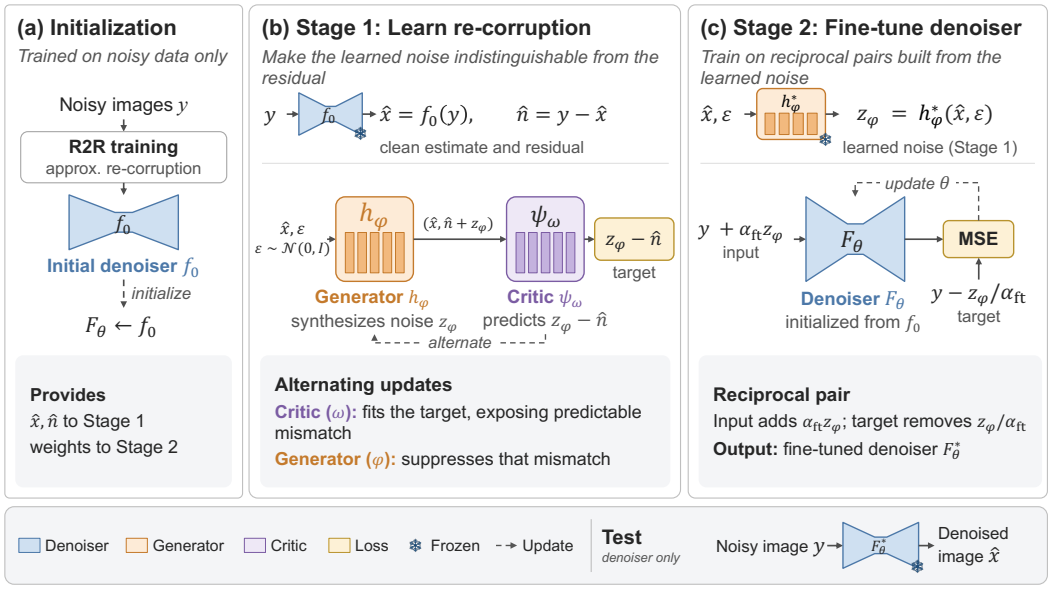

Figure 1: Overview of the workflow. **(a)** Train an initial denoiser using R2R with a standard re-corruption scheme. **(b)** Freeze the initial denoiser and learn the re-corruption mechanism through alternating critic and generator updates. **(c)** Freeze the learned generator and fine-tune the denoiser.

3 Main Results

Our framework follows a two-stage procedure (Figure 1): we first learn a re-corruption mechanism with reduced bias, and then freeze it to fine-tune the denoiser.

3.1 Equivalent Characterizations of Re-corruption Bias

Recall $y = x + n$ and the symmetric re-corruption pair $(y + z, y - z)$. As shown in equation 1,

$$\mathcal{L}_{\text{rec}}(\theta) = \mathbb{E}\left[\|F_\theta(y+z) - (y-z)\|_2^2\right] = \mathcal{R}(\theta) + \text{const} + 2C_\theta, \tag{6}$$

where

$$\mathcal{R}(\theta) = \mathbb{E}\|F_\theta(x+n+z) - x\|_2^2, \quad C_\theta = \mathbb{E}\Big[\big(F_\theta(x+n+z) - x\big)^\top (z-n)\Big]. \tag{7}$$

Thus, controlling C_θ is sufficient to make the re-corruption objective a low-bias surrogate of supervised denoising. For compact notation, let

$$U = (x, n+z), \qquad d = z - n. \tag{8}$$

Because C_θ depends on the particular denoiser, it does not by itself characterize the statistical quality of the re-corruption. We thus characterize the bias intrinsic to a given re-corruption by

$$Q := \sup_{g \in L^2(U), g \neq 0} \frac{\left|\mathbb{E}\left[g(U)^\top d\right]\right|^2}{\mathbb{E}\|g(U)\|_2^2}. \tag{9}$$

By construction, Q measures the largest normalized cross term that the re-corruption mismatch can induce. Since the residual of any denoiser,

$$g(x, n+z) = F_\theta(x+n+z) - x,$$

is an admissible choice in Eq. equation 9, we immediately obtain

$$C_\theta^2 \leq \mathcal{R}(\theta) Q. \tag{10}$$

Hence, unlike C_θ, Q depends only on the re-corruption mechanism and the underlying data distribution. A small Q controls the normalized cross-term bias for any denoiser, which is important for our two-stage framework: the re-corruption can be learned before the final denoiser is optimized.

The following result gives two equivalent characterizations of Q.

Proposition 1 (Equivalent characterizations of re-corruption bias)**.** *Assume that $z - n$ is square-integrable. Then*

$$Q = \mathbb{E}\,\|\mathbb{E}[z - n \mid x, n + z]\|_2^2\,, \tag{11}$$

and equivalently,

$$Q = \sup_{\psi \in L^2(x, n+z)} \left\{ 2\mathbb{E}\big[\psi(x, n + z)^\top (z - n)\big] - \mathbb{E}\|\psi(x, n + z)\|_2^2 \right\}. \tag{12}$$

The supremum in equation 12 is attained at

$$\psi^\star(x, n + z) = \mathbb{E}[z - n \mid x, n + z]. \tag{13}$$

The proof is given in Appendix A.1. The two forms in Proposition 1 have complementary interpretations. equation 11 shows that only the component of $z - n$ predictable from $(x, n + z)$ contributes to Q. Any remaining component is orthogonal to all square-integrable functions of $(x, n+z)$ and therefore cannot contribute to the cross term C_θ. equation 12, on the other hand, provides an optimization principle. For a fixed re-corruption, ψ searches for the component of $z - n$ that is predictable from $(x, n + z)$; its optimal value is exactly Q. This variational form will allow us to learn re-corruption without explicitly estimating the conditional expectation in equation 11.

Remark 1. *equation 11 also clarifies the relation between our criterion and classical exact re-corruption. If z and n are conditionally i.i.d. given x, then their exchangeability implies*

$$\mathbb{E}[z - n \mid x, n + z] = \mathbb{E}[n - z \mid x, z + n] = -\mathbb{E}[z - n \mid x, n + z] \Longrightarrow \mathbb{E}[z - n \mid x, n + z] = 0,$$

and therefore $Q = 0$. Hence, matching the full conditional distribution of the observation noise is a sufficient way to eliminate re-corruption bias, but it is stronger than necessary. Our formulation instead targets the weaker condition $Q = 0$ directly.

3.2 A TRUTH-FREE SURROGATE OF THE BIAS CRITERION

The bias measure Q cannot be evaluated from noisy observations alone, since it depends on the unavailable x and n. We therefore construct a noisy-only surrogate using an initial denoiser f_0, which remains fixed throughout re-corruption learning:

$$\hat{x} = f_0(y), \qquad \hat{n} = y - \hat{x}. \tag{14}$$

In practice, f_0 can be obtained using an R2R-style objective with approximate re-corruption.

We parameterize the re-corruption by a stochastic generator NN:

$$z_\phi = h_\phi(\hat{x}, \epsilon), \qquad \epsilon \sim \mathcal{N}(0, I), \tag{15}$$

where ϵ is independent of the noisy observation. We deliberately exclude $\hat{n}$ from the generator input. Otherwise, the generator could take the degenerate solution $z_\phi = \hat{n}$, for which the surrogate mismatch vanishes without producing a useful stochastic re-corruption.

Replacing the unavailable (x, n) in equation 11 by $(\hat{x}, \hat{n})$ gives the noisy-only surrogate

$$\widehat{Q}_\phi = \mathbb{E}\,\|\mathbb{E}[z_\phi - \hat{n} \mid \hat{x}, \hat{n} + z_\phi]\|_2^2\,. \tag{16}$$

This quantity has the same interpretation as the oracle Q: it measures the part of the surrogate mismatch $z_\phi - \hat{n}$ that is predictable from the corresponding re-corrupted variables.

Importantly, $\widehat{Q}_\phi$ is a surrogate rather than an unbiased estimator or upper bound of Q. Let $e_0 = \hat{x} - x$ denote the error of the initial denoiser. Since $\hat{n} = n - e_0$, we have

$$z_\phi - \hat{n} = (z_\phi - n) + e_0, \qquad (\hat{x}, \hat{n} + z_\phi) = (x + e_0, n + z_\phi - e_0). \tag{17}$$

Thus, the error of f_0 perturbs both the mismatch itself and the variables from which the mismatch is predicted. We therefore do not assume $\widehat{Q}_\phi = Q$ or $\widehat{Q}_\phi \geq Q$. Instead, equation 16 provides a noisy-data-only counterpart of the oracle criterion that can be optimized without clean references.

3.3 LEARNING RE-CORRUPTION WITH A NEURAL CRITIC

We now apply the variational representation equation 12 to the noisy-only surrogate. Parameterizing the predictor by a neural critic ψ_ω gives

$$\widehat{J}(\phi,\omega) = 2\mathbb{E}\big[\psi_\omega(\hat{x},\hat{n}+z_\phi)^\top(z_\phi-\hat{n})\big] - \mathbb{E}\,\|\psi_\omega(\hat{x},\hat{n}+z_\phi)\|_2^2\,. \tag{18}$$

With an unrestricted predictor class,

$$\sup_\psi \widehat{J} = \widehat{Q}_\phi, \tag{19}$$

and the optimal predictor is

$$\psi^\star(\hat{x},\hat{n}+z_\phi) = \mathbb{E}[z_\phi-\hat{n}\mid \hat{x},\hat{n}+z_\phi]\,. \tag{20}$$

Consequently, the noisy-only learning problem can be expressed as

$$\min_\phi \max_\omega \widehat{J}(\phi,\omega). \tag{21}$$

The inner optimization exposes the component of the surrogate mismatch that is predictable from the re-corrupted variables, while the generator changes the re-corruption so that this component becomes small. Unlike adversarial distribution matching, the critic is therefore not a real/fake classifier; it estimates the bias-related predictable component identified by our theoretical analysis.

Fixing ϕ, maximizing equation 18 for critic weights is equivalent to minimizing the following loss:

$$\mathcal{L}_{\text{critic}}(\phi,\omega) = \mathbb{E}\,\|\psi_\omega(\hat{x},\hat{n}+z_\phi)-(z_\phi-\hat{n})\|_2^2\,. \tag{22}$$

Thus, the critic has a direct interpretation: under population MSE regression, it estimates the conditional mean in equation 20. A finite neural critic may recover only part of this predictable component, but its maximized objective still provides the signal used to improve the generator.

Ideally, the generator would minimize the maximized objective in equation 21. Our implementation alternates several critic updates with one generator update. During the generator step, we treat the current critic prediction as a fixed estimate of the predictable mismatch and minimize

$$\mathcal{J}_{\text{gen}}(\phi) = 2\mathbb{E}\Big[\text{sg}(\psi_\omega(\hat{x},\hat{n}+z_\phi))^\top(z_\phi-\hat{n})\Big]\,, \tag{23}$$

where sg denotes stop-gradient. This update directly suppresses the component of the mismatch exposed by the critic. It is an approximation to the exact generator gradient of equation 18; the difference and its interpretation are given in Appendix A.2.

The variational criterion alone can admit undesirable changes in coarse marginal statistics, for example by shrinking the generated perturbation. We therefore weakly constrain its mean and energy:

$$\mathcal{L}_{\text{mean}} = \|\mathbb{E}[z_\phi]\|_2^2, \qquad \mathcal{L}_{\text{energy}} = \big(\mathbb{E}\|z_\phi\|_2^2 - \mathbb{E}\|\hat{n}\|_2^2\big)^2\,. \tag{24}$$

The generator objective adopted in the implementation is

$$\min_\phi \left[\mathcal{J}_{\text{gen}}(\phi) + \lambda_\mu\mathcal{L}_{\text{mean}} + \lambda_e\mathcal{L}_{\text{energy}}\right]. \tag{25}$$

These terms constrain only coarse first- and second-order statistics; the image-dependent structure of the re-corruption is learned from the predictability signal supplied by the critic.

3.4 DENOISER FINE-TUNING WITH LEARNED RE-CORRUPTION

The procedure above is used only to learn the re-corruption mechanism, while the initial denoiser f_0 remains fixed. Once we have learned the generator, h_{ϕ^*}, we freeze it and fine-tune f_0 to obtain an improved denoiser f_{θ^*} using the resulting noisy/noisy objective with reduced re-corruption bias.

Recall the generator is learned using the pair $(y+z_\phi, y-z_\phi)$. In this denoising fine-tuning stage, we use the reciprocal-scaling version of the R2R scheme:

$$\min_\theta \mathbb{E}\left\|F_\theta(y+\alpha_{\text{ft}}z_\phi) - \left(y-\frac{z_\phi}{\alpha_{\text{ft}}}\right)\right\|_2^2, \qquad z_\phi = h_\phi(f_0(y),\epsilon). \tag{26}$$

When $\alpha_{\text{ft}} = 1$, equation 26 uses exactly the same symmetric re-corruption pair as in generator training and the derivation of Q. For a general $\alpha_{\text{ft}} > 0$, the corresponding bias is instead governed by

$$\mathbb{E}\left[\frac{z_\phi}{\alpha_{\text{ft}}} - n \;\middle|\; x, n + \alpha_{\text{ft}} z_\phi\right]. \tag{27}$$

Nevertheless, in the denoiser fine-tuning stage, we allow α_{ft} to be a hyperparameter to be optimized.

Overall, the two stages are deliberately decoupled. Stage 1 uses the denoiser-independent criterion to learn a re-corruption with reduced bias-relevant mismatch while keeping f_0 fixed. Stage 2 freezes this re-corruption and optimizes the denoiser using the resulting noisy/noisy training pairs.

To conclude, the complete training procedure consists of three steps:

(1) We train an initial denoiser f_0 from scratch using R2R with a simple approximate re-corruption.

(2) Keeping f_0 fixed, we learn the re-corruption $z_\phi = h_\phi(\hat{x}, \epsilon)$ together with the critic.

(3) We freeze h_{ϕ^*} and fine-tune f_0 using the learned re-corruption to obtain the final denoiser f_{θ^*}.

At test time, only F_{θ^*} is applied to the noisy image; no re-corruption or sampling is required.

4 Experiments

Datasets. We consider two image domains: camera photography and fluorescence microscopy, with training based only on noisy images. For camera photography, we train on the noisy sRGB images of SIDD-Medium (Abdelhamed et al., 2018) and evaluate on SIDD validation and the SIDD benchmark , 1280 blocks of 256×256 each . For fluorescence microscopy, FMD (Zhang et al., 2019) contains 12 image types from three modalities; its noise is photon-limited and signal-dependent, a different regime from camera sRGB, and the reference of each FOV is the average of 50 captures. We train one model for all 12 types on FOV 1–18 and test on FOV 19, the mixed test set of 48 images. Checkpoints and hyper-parameters are selected on SIDD validation and FOV 20 of FMD.

Implementation details. For the initial R2R model, the approximate re-corruption is instantiated as per-channel Gaussian noise for camera images, and as signal-dependent Gaussian noise fitted from each noisy FMD image. We use a small Restormer (Zamir et al., 2022) (3.3M parameters) as the denoiser and, unless otherwise specified, six-layer 3×3 convolutional networks for the generator and critic. The fine-tuning scale α_{ft} is selected on validation for camera images and fixed to 1 on FMD. Full implementation details are given in Appendix B.

4.1 Denoising real-world camera images

In this experiment, we follow the protocol of R2R (Pang et al., 2021) to train an initial denoiser f_0. We then use our method to learn the re-corruption generator and use it to fine-tune f_0 on noisy images only. As shown in Table 2, the resulting model f_{θ^*} is the best performer. Among re-corruption methods, L2R (Monroy et al., 2026) is the closest prior work related to ours. L2R's lower performance on SIDD may partly reflect a mismatch between its target setting and real-camera noise: L2R aims at distribution-agnostic denoising across diverse synthetic non-Gaussian and structured noise models, rather than real-world camera noise. See Appendix C.3 for a more detailed discussion. In addition to quantitative performance gain, our approach also shows the improvement of visual quality of the results. See Figure 2 and Appendix D for visual comparison of different methods.

4.2 Denoising fluorescence microscopy images

In this experiment, we evaluate on the official mixed test set of FMD (Zhang et al., 2019) at the raw noise level. Table 3 compares three types of methods: models transferred from SIDD, zero-shot methods operating on individual test images, and methods trained on noisy FMD images. For the

[1]Running on the full validation blocks with a third-party implementation: `https://github.com/parsa21162/Zero-Shot-Image-Denoising-`.

[2]L2R does not report the results on SIDD; this number is obtained by applying its released configuration to SIDD-Medium under our evaluation protocol.

Table 2: Comparison on SIDD in PSNR (dB) and SSIM. Best results are shown in bold and second-best results are underlined.

	Method	SIDD validation	SIDD benchmark
Alternative approach	CVF-SID (Neshatavar et al., 2022)	34.51	35.03 / 0.856
	ZS-HPD (Zhao et al., 2026)[1]	27.10	31.01 / 0.511
Blind-spot approach	AP-BSN + R3 (Lee et al., 2022)	35.76	36.51 / 0.872
	C-BSN (Jang et al., 2023b)	36.22	36.85 / 0.882
	LG-BPN (Wang et al., 2023)	37.32	37.67 / 0.887
	PUCA (Jang et al., 2023a)	37.49	37.93 / 0.882
	SASL (Li et al., 2023)	37.39	37.85 / 0.879
	SDAP (Pan et al., 2023)	37.30	37.64 / 0.882
	AT-BSN (D) (Chen et al., 2024)	37.88	38.14 / 0.891
	APR-RD (Kim & Cho, 2025)	38.00	38.26 / 0.895
	TBSN (Li et al., 2025)	37.71	38.02 / 0.886
	NSP (Shan et al., 2026)	37.12	37.46 / 0.874
	TM-BSN ($s = 4$) (Park et al., 2026)	37.31	37.71 / 0.881
	TM-BSN (D) (Park et al., 2026)	<u>38.08</u>	<u>38.31</u> / <u>0.900</u>
Data augmentation approach	R2R (Pang et al., 2021)	35.04	35.50 / 0.8550
	L2R (Monroy et al., 2026)[2]	28.25	31.14 / 0.557
	Ours	**38.21**	**38.46 / 0.901**

Table 3: Comparison on the FMD mixed test set at the raw noise level. Results are reported in PSNR (dB) / SSIM; best and second-best results are shown in bold and underlined, respectively.

Method	Trained on	FMD `test_mix`
TransCLIP (Cheng et al., 2024)	SIDD	30.56 / 0.705
Restormer (Zamir et al., 2022)	SIDD	31.53 / 0.796
AP-BSN (Lee et al., 2022)	SIDD	31.99 / 0.836
Masked training (Chen et al., 2023)	SIDD	32.32 / 0.850
AT-BSN (Chen et al., 2024)	SIDD	32.40 / 0.838
CVF-SID (Neshatavar et al., 2022)	SIDD	32.73 / 0.843
Self2Self (Quan et al., 2020)	test image	30.76 / 0.695
BM3D (Dabov et al., 2007)	–	31.58 / 0.824
AMSNet (Liao et al., 2024)	test image	31.60 / 0.813
ZS-N2N (Mansour & Heckel, 2023)	test image	31.66 / 0.768
DIP (Ulyanov et al., 2018)	test image	32.07 / 0.811
Noise2Self (Batson & Royer, 2019)	test image	32.40 / 0.799
MPI (Ma et al., 2024)	test image	32.95 / 0.846
ScoreDVI (Cheng et al., 2023)	test image	33.10 / 0.865
MASH (Chihaoui & Favaro, 2024)	test image	33.69 / 0.881
FM2S (Liu et al., 2026)	test image	<u>34.15</u> / <u>0.886</u>
TM-BSN$^-$ (Park et al., 2026)	FMD (noisy)	32.14 / 0.840
Ours	FMD (noisy)	**34.16 / 0.888**

last category, we additionally train TM-BSN on FMD without its distillation stage. Our method trains a single model on noisy FMD images and shares it across all twelve imaging configurations.

As shown in Table 3, our model achieves 34.16 dB / 0.888 SSIM, the best result among all methods. In particular, it slightly outperforms FM2S (Liu et al., 2026), which is specifically designed for fluorescence microscopy and is per-image optimized, while our single model is shared across all twelve imaging configurations.

4.3 Ablation study

Validating the re-corruption bias criterion. We use an initial denoiser f_0, trained with a simple re-corruption via Gaussian white noise estimated from noisy images, and compare three re-corruptions for fine-tuning: white Gaussian, correlated Gaussian, and the learned z by our generator. Here,

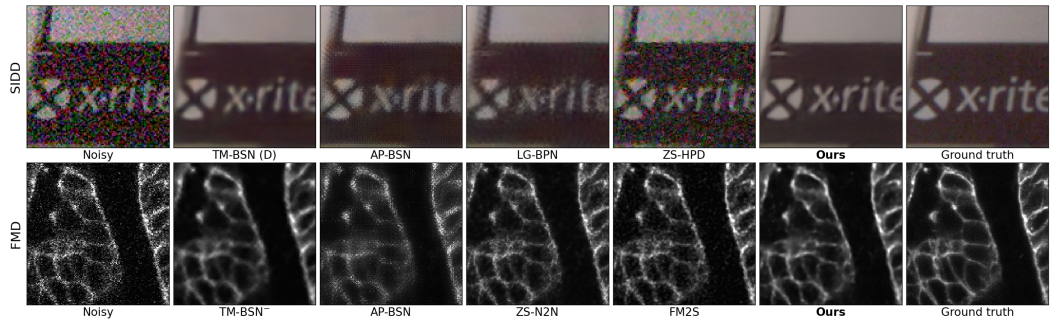

Figure 2: Visual comparison across the two domains. The first row shows SIDD validation blocks, and the second row shows FMD `test_mix` images. See more visual comparisons in Appendix D.

Table 4: Validation of the re-corruption bias criterion on SIDD. All methods start from the same f_0 and differ only in the re-corruption used for fine-tuning. $|2C_\theta|$ is measured on SIDD validation (10^{-4} per pixel), while $\hat{Q}$ is evaluated with a fresh critic and requires no ground truth. PSNR (dB) is reported on validation and benchmark sets for two fine-tuning scales.

Re-corruption z	$\lvert 2C_\theta\rvert$ at f_0	$\hat{Q}$	$\alpha_{\text{ft}} = 1$			$\alpha_{\text{ft}} = 0.5$		
			$\lvert 2C_\theta\rvert$	val.	bench.	$\lvert 2C_\theta\rvert$	val.	bench.
White Gaussian	17.75	0.89	18.08	34.96	35.69	17.90	35.01	35.73
Colored Gaussian	7.34	0.75	6.85	35.25	35.88	5.95	35.21	35.83
Learned z (ours)	**1.78**	**0.69**	**5.81**	**35.50**	**36.08**	**2.40**	**35.97**	**36.45**

correlated Gaussian denotes Gaussian perturbation with estimated spatial and cross-channel noise correlations (Yue et al., 2021). Table 4 reports the measured cross term $|2C_\theta|$, the noisy-only surrogate $\hat{Q}$, and the resulting denoising performance. The cross term is measured on SIDD validation using ground truth only for this diagnostic analysis, while $\hat{Q}$ is evaluated with a fresh critic trained independently of the critic used to learn z. All rows use the same f_0, critic architecture, training budget, and fine-tuning protocol.

The results show a consistent relation between re-corruption bias and denoising quality. Compared with white Gaussian noise, correlated Gaussian noise with estimated spatial and cross-channel correlations reduces $|2C_\theta|$ at f_0 by more than half, while the learned re-corruption reduces it to nearly one tenth of the white-Gaussian value. The noisy-only surrogate $\hat{Q}$ follows the same ordering without requiring ground truth, supporting its use as a practical bias-control criterion.

Such a reduction in re-corruption bias naturally translates into improved denoising performance. The learned z improves SIDD validation by $0.54\,\text{dB}$ at $\alpha_{\text{ft}} = 1$ and $1.01\,\text{dB}$ at $\alpha_{\text{ft}} = 0.5$, with corresponding benchmark gains of 0.40 and $0.77\,\text{dB}$. Although fine-tuning changes the denoiser and hence the cross term, the learned re-corruption still yields the smallest $|2C_\theta|$.

More ablation studies: Additional ablation studies are provided in Appendix C.

5 Conclusion and Future Work

We revisit re-corruption for unsupervised image denoising and show that full noise-distribution modeling is unnecessary: it is sufficient to control the bias-relevant re-corruption mismatch. This leads to a two-stage framework that learns re-corruption from noisy data and then uses it to improve the denoiser. Experiments on real camera and fluorescence-microscopy images consistently improve the initial R2R denoiser, while bias measurements support the proposed interpretation. Overall, controlling re-corruption bias provides a practical way to handle complex real-world noise without explicitly modeling its full distribution. Future work will explore adaptive re-corruption during denoiser training, stronger noisy-only surrogates for bias control, and extensions to inverse problems with unknown measurement noise.

AI USE STATEMENT

In this work, we used generative AI tools primarily for technical writing and editing, including improving clarity and grammar, reorganizing text, and compressing the presentation to meet the page limit. AI tools were also used to suggest alternative wording and to assist with the preparation of schematic figures and diagrams. All AI-assisted text and figures were reviewed and revised by the authors, who verified the accuracy of the technical content, experimental results, and scientific claims. The authors take full responsibility for the final content of the paper.

ETHICS STATEMENT

This work studies image denoising algorithms and does not introduce new data collection involving human subjects or sensitive personal information. The experiments use existing benchmark datasets for camera and fluorescence microscopy images. We do not identify any specific ethical concerns related to the methodology or experiments presented in this work. The authors have followed the applicable licenses and usage conditions of the datasets and take responsibility for the integrity and reporting of the results.

REPRODUCIBILITY STATEMENT

We provide detailed descriptions of the proposed method, training procedure, datasets, evaluation protocols, and implementation settings in the main paper and appendix. Complete proofs of the theoretical results and additional experimental details are also included in the appendix. Upon acceptance of the paper, we will release the source code and training configurations on GitHub to facilitate reproduction of the reported results.

REFERENCES

Abdelrahman Abdelhamed, Stephen Lin, and Michael S. Brown. A high-quality denoising dataset for smartphone cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1692–1700, June 2018.

Joshua Batson and Loic Royer. Noise2self: Blind denoising by self-supervision. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 524–533. PMLR, 2019.

Haoyu Chen, Jinjin Gu, Yihao Liu, Salma Abdel Magid, Chao Dong, Qiong Wang, Hanspeter Pfister, and Lei Zhu. Masked image training for generalizable deep image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1692–1703, 2023.

Shiyan Chen, Jiyuan Zhang, Zhaofei Yu, and Tiejun Huang. Exploring efficient asymmetric blind-spots for self-supervised denoising in real-world scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2814–2823, 2024.

Jun Cheng, Tao Liu, and Shan Tan. Score priors guided deep variational inference for unsupervised real-world single image denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 12937–12948, 2023.

Jun Cheng, Dong Liang, and Shan Tan. Transfer clip for generalizable image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 25974–25984, 2024.

Hamadi Chihaoui and Paolo Favaro. Masked and shuffled blind spot denoising for real-world images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3025–3034, June 2024.

Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16 (8):2080–2095, 2007.

Qilong Guo, Tianjing Zhang, Zhiyuan Ma, and Hui Ji. Zero-shot blind-spot image denoising via cross-scale non-local pixel refilling. In *Advances in Neural Information Processing Systems*, volume 38, 2025.

Tao Huang, Songjiang Li, Xu Jia, Huchuan Lu, and Jianzhuang Liu. Neighbor2neighbor: Self-supervised denoising from single noisy images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14781–14790, June 2021.

Hyemi Jang, Junsung Park, Dahuin Jung, Jaihyun Lew, Ho Bae, and Sungroh Yoon. Puca: Patch-unshuffle and channel attention for enhanced self-supervised image denoising. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, pp. 19217–19229, 2023a.

Yeong Il Jang, Keuntek Lee, Gu Yong Park, Seyun Kim, and Nam Ik Cho. Self-supervised image denoising with downsampled invariance loss and conditional blind-spot network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 12196–12205, 2023b.

Hyunjun Kim and Nam Ik Cho. Apr-rd: Complemental two steps for self-supervised real image denoising. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 4257–4265, 2025.

Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2void—learning denoising from single noisy images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2129–2137, June 2019.

Samuli Laine, Tero Karras, Jaakko Lehtinen, and Timo Aila. High-quality self-supervised deep image denoising. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Wooseok Lee, Sanghyun Son, and Kyoung Mu Lee. Ap-bsn: Self-supervised denoising for real-world images via asymmetric pd and blind-spot network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17725–17734, June 2022.

Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2noise: Learning image restoration without clean data. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2965–2974. PMLR, 2018.

Junyi Li, Zhilu Zhang, Xiaoyu Liu, Chaoyu Feng, Xiaotao Wang, Lei Lei, and Wangmeng Zuo. Spatially adaptive self-supervised learning for real-world image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9914–9924, 2023.

Junyi Li, Zhilu Zhang, and Wangmeng Zuo. Rethinking transformer-based blind-spot network for self-supervised image denoising. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 4788–4796, 2025.

Xiangyu Liao, Tianheng Zheng, Jiayu Zhong, Pingping Zhang, and Chao Ren. Asymmetric mask scheme for self-supervised real image denoising. In *European Conference on Computer Vision*, pp. 199–215. Springer, 2024.

Jizhihui Liu, Qixun Teng, Qing Ma, Junhui Hou, and Junjun Jiang. Fm2s: Towards spatially-correlated noise modeling in zero-shot fluorescence microscopy image denoising. *Machine Intelligence Research*, 23(1):200–213, 2026.

Qing Ma, Junjun Jiang, Xiong Zhou, Pengwei Liang, Xianming Liu, and Jiayi Ma. Pixel2pixel: A pixelwise approach for zero-shot single image denoising. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(6):4614–4629, 2025. doi: 10.1109/TPAMI.2025.3546870.

Xiaoxiao Ma, Zhixiang Wei, Yi Jin, Pengyang Ling, Tianle Liu, Ben Wang, Junkang Dai, and Huaian Chen. Masked pre-training enables universal zero-shot denoiser. *Advances in Neural Information Processing Systems*, 37:32570–32610, 2024.

Youssef Mansour and Reinhard Heckel. Zero-shot noise2noise: Efficient image denoising without any data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14018–14027, June 2023. doi: 10.1109/CVPR52729.2023.01347.

Brayan Monroy, Jorge Bacca, and Julián Tachella. Generalized recorrupted-to-recorrupted: Self-supervised learning beyond gaussian noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 28155–28164, June 2025.

Brayan Monroy, Jorge Bacca, and Julián Tachella. Learning to recorrupt: Noise distribution agnostic self-supervised image denoising. *arXiv preprint arXiv:2603.25869*, 2026.

Nick Moran, Dan Schmidt, Yu Zhong, and Patrick Coady. Noisier2noise: Learning to denoise from unpaired noisy data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12064–12072, June 2020. doi: 10.1109/CVPR42600.2020.01208.

Reyhaneh Neshatavar, Mohsen Yavartanoo, Sanghyun Son, and Kyoung Mu Lee. Cvf-sid: Cyclic multi-variate function for self-supervised image denoising by disentangling noise from image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 17583–17591, June 2022.

Yizhong Pan, Xiao Liu, Xiangyu Liao, Yuanzhouhan Cao, and Chao Ren. Random sub-samples generation for self-supervised real image denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 12150–12159, 2023.

Tongyao Pang, Huan Zheng, Yuhui Quan, and Hui Ji. Recorrupted-to-recorrupted: Unsupervised deep learning for image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2043–2052, June 2021.

Junyoung Park, Youngjin Oh, and Nam Ik Cho. Tm-bsn: Triangular-masked blind-spot network for real-world self-supervised image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 29877–29886, June 2026.

Yuhui Quan, Mingqin Chen, Tongyao Pang, and Hui Ji. Self2self with dropout: Learning self-supervised denoising from single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1890–1898, June 2020.

Yuhui Quan, Xinran Qin, Tongyao Pang, and Hui Ji. Dual-domain self-supervised learning and model adaption for deep compressive imaging. In *European Conference on Computer Vision (ECCV)*, pp. 409–426, 2022.

Yuhui Quan, Xinran Qin, Tongyao Pang, and Hui Ji. Siamese cooperative learning for unsupervised image reconstruction from incomplete measurements. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(7):4866–4879, 2024.

Yuhui Quan, Tianxiang Zheng, Zhiyuan Ma, and Hui Ji. Zero-shot blind-spot image denoising via implicit neural sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7502–7512, June 2025.

Yiwen Shan, Haiyu Zhao, Peng Hu, Xi Peng, and Yuanbiao Gou. Next-scale prediction: A self-supervised approach for real-world image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22669–22678, June 2026.

Shakarim Soltanayev and Se Young Chun. Training deep learning based denoisers without ground truth data. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Julián Tachella, Mike Davies, and Laurent Jacques. UNSURE: self-supervised learning with unknown noise level and stein's unbiased risk estimate. In *The Thirteenth International Conference on Learning Representations*, 2025. URL `https://openreview.net/forum?id=ScVnYBaSEw`.

Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9446–9454, June 2018.

Zichun Wang, Ying Fu, Ji Liu, and Yulun Zhang. Lg-bpn: Local and global blind-patch network for self-supervised real-world denoising. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 18156–18165. IEEE, 2023.

Xiaohe Wu, Ming Liu, Yue Cao, Dongwei Ren, and Wangmeng Zuo. Unpaired learning of deep image denoising. In *European Conference on Computer Vision*, pp. 352–368, 2020. doi: 10.1007/978-3-030-58548-8_21.

Huanjing Yue, Zhongyu Jiang, Shengdi Zhou, Jingyu Yang, Yonghong Hou, and Chunping Hou. Deep noise estimation and removal for real-world noisy images. *Signal Processing: Image Communication*, 94:116231, 2021.

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14821–14831, June 2021.

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*, pp. 5728–5739, June 2022.

Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26 (7):3142–3155, 2017. doi: 10.1109/TIP.2017.2662206.

Kai Zhang, Wangmeng Zuo, and Lei Zhang. Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *IEEE Transactions on Image Processing*, 27(9):4608–4622, 2018. doi: 10.1109/TIP.2018.2839891.

Yi Zhang, Dasong Li, Ka Lung Law, Xiaogang Wang, Hongwei Qin, and Hongsheng Li. Idr: Self-supervised image denoising via iterative data refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2098–2107, 2022.

Yide Zhang, Yinhao Zhu, Evan Nichols, Qingfei Wang, Siyuan Zhang, Cody Smith, and Scott Howard. A poisson-gaussian denoising dataset with real fluorescence microscopy images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11710–11718, June 2019.

Xiaole Zhao, Qingsong Pang, Xiaobo Zhang, Xun Xu, Xun Gong, Yan Yang, and Tianrui Li. Zero-shot image denoising via hybrid prior-guided pseudo sample generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 22648–22657, 2026.

Huan Zheng, Tongyao Pang, and Hui Ji. Unsupervised deep video denoising with untrained network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 3651–3659, 2023. doi: 10.1609/aaai.v37i3.25476.

A Theoretical Details

A.1 Proof of Proposition 1

Proof. Recall the shorthand $U = (x, n + z)$ and $d = z - n$, and write

$$m(U) = \mathbb{E}[d \mid U].$$

Square-integrability of d and the fact that conditional expectation is an L^2 contraction give $m \in L^2(U)$.

First claim. For any $g \in L^2(U)$, $g \neq 0$, the tower property gives

$$\mathbb{E}\big[g(U)^\top d\big] = \mathbb{E}\big[g(U)^\top m(U)\big]\,,$$

so by the Cauchy–Schwarz inequality in $L^2(U)$,

$$\frac{\big|\mathbb{E}\big[g(U)^\top d\big]\big|^2}{\mathbb{E}\|g(U)\|_2^2} = \frac{\big|\mathbb{E}\big[g(U)^\top m(U)\big]\big|^2}{\mathbb{E}\|g(U)\|_2^2} \leq \mathbb{E}\|m(U)\|_2^2.$$

If $\mathbb{E}\|m(U)\|_2^2 > 0$, the choice $g = m$ is admissible and attains the bound, so the supremum in Eq. 9 equals $\mathbb{E}\|m(U)\|_2^2$. If $m(U) = 0$ almost surely, the numerator vanishes for every admissible g and the supremum is $0 = \mathbb{E}\|m(U)\|_2^2$. In both cases

$$Q = \mathbb{E}\|m(U)\|_2^2 = \mathbb{E}\,\|\mathbb{E}[z - n \mid x, n + z]\|_2^2\,,$$

which is Eq. 11.

Second claim. For any $\psi \in L^2(U)$, the tower property and completing the square give

$$\begin{aligned} 2\mathbb{E}\big[\psi(U)^\top d\big] - \mathbb{E}\|\psi(U)\|_2^2 &= 2\mathbb{E}\big[\psi(U)^\top m(U)\big] - \mathbb{E}\|\psi(U)\|_2^2 \\ &= \mathbb{E}\|m(U)\|_2^2 - \mathbb{E}\|\psi(U) - m(U)\|_2^2. \end{aligned}$$

The second term is non-negative and vanishes exactly when $\psi = m$ almost surely. Since $m \in L^2(U)$, that choice is admissible, so

$$\sup_{\psi \in L^2(U)} \big\{2\mathbb{E}\big[\psi(U)^\top d\big] - \mathbb{E}\|\psi(U)\|_2^2\big\} = \mathbb{E}\|m(U)\|_2^2 = Q,$$

attained at $\psi^\star = m$. With the first claim this proves Eqs. 12 and 13. □

A.2 The stop-gradient generator update

We relate the practical update of Eq. (23) to the gradient of J in Eq. (18). Fix the critic parameters ω and write $c = \psi_\omega(U_\phi)$, $d = \hat{d}_\phi$, $G_\phi = \partial z_\phi / \partial \phi$ for the Jacobian of the generator output, and $A = \partial \psi_\omega(U_\phi)/\partial(\hat{n} + z_\phi)$ for the Jacobian of the critic with respect to its second input; $\hat{x}$ and $\hat{n}$ do not depend on ϕ since f_0 is frozen. Since $\partial d/\partial\phi = G_\phi$ and $\partial c/\partial\phi = A\,G_\phi$, the chain rule gives

$$\nabla_\phi J = 2\,\mathbb{E}\big[G_\phi^\top c\big] + 2\,\mathbb{E}\big[G_\phi^\top A^\top (d - c)\big], \qquad \nabla_\phi \mathcal{J}_{\mathrm{gen}} = 2\,\mathbb{E}\big[G_\phi^\top c\big], \tag{28}$$

so the stop-gradient update omits exactly the second term, which propagates through the critic input:

$$\nabla_\phi J - \nabla_\phi \mathcal{J}_{\mathrm{gen}} = 2\,\mathbb{E}\big[G_\phi^\top A^\top (d - c)\big]. \tag{29}$$

Optimal critic. Suppose the critic realizes the conditional mean, $c = \hat{m}_\phi(U_\phi) = \mathbb{E}[d \mid U_\phi]$, so that $\mathbb{E}[d - c \mid U_\phi] = 0$. Write the per-sample integrand of J as $j_\omega(U, d) = 2\,\psi_\omega(U)^\top d - \|\psi_\omega(U)\|^2$ and consider its gradient with respect to the generator output z_ϕ, which enters both U_ϕ and d: $g_z = \nabla_z j_\omega = 2c + 2A^\top(d - c)$, whereas the stop-gradient form gives $g_z^{\mathrm{sg}} = 2c$. Since A and c are functions of U_ϕ,

$$\mathbb{E}\big[g_z \mid U_\phi\big] = 2\,\hat{m}_\phi(U_\phi) = g_z^{\mathrm{sg}}. \tag{30}$$

In output space the stop-gradient direction is therefore the conditional expectation of the exact per-sample gradient given U_ϕ; the omitted term has zero conditional mean given U_ϕ and contributes only conditional variation. For the parameters ϕ, the omitted term in Eq. (29) is $2\,\mathbb{E}[G_\phi^\top A^\top (d - c)]$, which vanishes whenever $\mathbb{E}[G_\phi^\top A^\top (d - c) \mid U_\phi] = 0$, for instance when G_ϕ is a function of U_ϕ. A generic neural generator $h_\phi(\hat{x}, \epsilon)$ does not guarantee this, since G_ϕ depends on the seed ϵ, so we treat Eq. (23) as a practical approximation rather than the exact gradient. In practice five critic steps are taken per generator step.

B Implementation details

B.1 SIDD: R2R original setting

The initial denoiser of Table 2 is an R2R model with a small Restormer architecture trained with the protocol of Pang et al. (2021): the re-corruption is drawn in the raw domain with the NLF provided by SIDD and the pair is rendered to sRGB through the SIDD pipeline. It reaches 37.98 dB on validation and 38.26 dB / 0.889 on the benchmark. Its generator is the exact energy-matching variant, four layers with 32 channels, trained for 16k iterations on the sRGB noisy images against that model's residual. Fine-tuning uses $\alpha_{\mathrm{ft}} = 5$ with learning rate $1.2 \cdot 10^{-4}$; every cell with $\alpha_{\mathrm{ft}} \in \{4.5, 5, 5.5\}$ and learning rate $1.6 \cdot 10^{-4}$ or $2.4 \cdot 10^{-4}$ lands within 0.04 dB of it.

B.2 SIDD: BLIND SETTING

Here nothing about the noise is given. The patch is divided into 16×16 tiles and the noise level $\hat{\sigma}_c$ of each channel is the 20th percentile of the per-tile pixel standard deviation, which is dominated by noise in the flattest tiles and, being a marginal statistic, is less sensitive to spatial correlation than high-frequency estimators; then $z_c = \hat{\sigma}_c \, \epsilon_c$ with $\epsilon \sim \mathcal{N}(0, I)$. The colored Gaussian z additionally estimates, from the same tiles, the lag-one correlation of each channel and the 3×3 cross-channel correlation, and imposes them with a separable three-tap filter and a Cholesky factor. The initial denoiser is trained with $\alpha_0 = 5$, the generator for 40k iterations with $\lambda_e = 300$, and fine-tuning uses $\alpha_{\mathrm{ft}} = 0.5$ with learning rate 10^{-5} for 1200 iterations.

B.3 FMD

For FMD, we model the observation noise as Poisson–Gaussian, with $\mathrm{Var}(n \mid x) = a\,x + b$. Per image we take the mean and variance of 8×8 tiles, group them into 16 intensity bins, use a low-to-median order statistic of the variance in each bin to reject tiles with structure, and fit (a, b) by weighted least squares; then $z = \sqrt{a\,\bar{y} + b}\;\epsilon$ with $\bar{y}$ a local mean of y and ϵ standard Gaussian, so z reproduces the variance of that model but not its higher moments. Where the estimated lag-one correlation exceeds 0.15 (wide field), z is given that correlation. We use $\alpha_0 = 1$ and $\alpha_{\mathrm{ft}} = 1$, train the generator for 20k iterations with the energy matched exactly rather than penalized, and fine-tune for 4000 iterations. Inspired by IDR (Zhang et al., 2022), which refines a self-supervised denoiser by rebuilding its training data from the output of the previous round, we add one further round as a post-process. The reported FMD result includes this additional round, which is applied to both the R2R baseline and our method. Before the round, the baseline and our method achieve 33.53 and 33.70 dB, respectively; afterward, they achieve 34.06 and 34.16 dB. The gains over the baseline are therefore 0.17 dB before the additional round and 0.10 dB afterward.

B.4 ARCHITECTURES AND OPTIMIZATION

The denoiser is a Restormer with 3.3M parameters. The generator maps $(\hat{x}, \epsilon)$ to z and the critic maps $(\hat{x},\, z + \hat{n})$, $\hat{n} = y - \hat{x}$, to a prediction of $z - \hat{n}$; except where Sec. B states otherwise, both are six 3×3 convolutions with 64 channels and GELU, without normalization or skip connections. R2R pairs are $(y + \alpha_0 z,\; y - z/\alpha_0)$ and f_0 is trained with AdamW (learning rate 3×10^{-4} on camera images and 2×10^{-4} on FMD, weight decay 10^{-4}, gradient-norm clipping at 0.01) on 128×128 patches with batch size 8, enlarged to 256×256 on FMD, with random flips and rotations; the loss is the squared error. The generator and critic are trained alternately with Adam (learning rates 5×10^{-5} and 2×10^{-4}, five critic steps per generator step, the generator using the stop-gradient form) on 128×128 patches with batch size 16, under a zero-mean penalty with $\lambda_\mu = 1000$ and an energy penalty whose weight is given above, energies being per-pixel means of the squared channel norm. The generator is then frozen and f_0 is fine-tuned with AdamW. The initial training of f_0 and the fine-tuning both keep the checkpoint with the best validation PSNR; the generator is not selected on validation; we use its last iterate. Two of the settings replace the energy penalty by exact matching, rescaling z so that its mean squared norm over the batch equals that of $y - f_0(y)$; that variant therefore couples z to the residual through one scalar per batch.

B.5 BASELINES

TM-BSN$^-$ on FMD uses the authors' single-channel network trained by us for 100k iterations without the distillation stage; the checkpoint and the test-time hole size are selected on FOV 20. ZS-HPD is run with the third-party implementation cited in Table 2.

C MORE ABLATION STUDIES

C.1 EFFECT OF THE INITIAL RE-CORRUPTION:

Table 5 varies the re-corruption used to train f_0. The gain decreases as that re-corruption becomes more accurate, from $+1.01$ dB for the deliberately simple white Gaussian to $+0.23$ dB with the camera NLF and ISP, and remains positive throughout: the more closely the initial re-corruption

Table 5: Effect of the initial re-corruption. SIDD validation, PSNR (dB). Note that the first column is the re-corruption used to *train* f_0, not the one used for fine-tuning as in Table 4.

Initial re-corruption of f_0	f_0	Ours	Gain
White Gaussian	34.96	35.97	$+1.01$
Colored Gaussian	35.60	35.88	$+0.28$
NLF and ISP	37.98	38.21	$+0.23$

already matches the observation noise, the less bias is left to correct. The gain is not explained by perturbation energy alone — in the strongest setting an energy-matched Gaussian perturbation reaches only 38.01 dB, against 38.21 dB for the learned z. We use that strongest initialization for the comparison below.

C.2 Energy constraint and generator capacity.

Without a sufficiently strong energy constraint ($\lambda_e = 10$), the generator lowers the critic objective by shrinking z_ϕ and the fine-tuning gain vanishes; with a very strong one ($\lambda_e \geq 1000$) the constraint dominates and the gain drops to about $+0.05$ dB. Relative to the 6×64 generator, 8×96 is 0.02 dB and 8×128 is 0.11 dB lower on SIDD validation.

C.3 Analysis of the comparison with L2R.

L2R (Monroy et al., 2026) is the closest prior work on learning re-corruption under unknown noise. It learns a re-corruptor within a denoiser–re-corruptor saddle-point formulation and demonstrates strong performance across several controlled non-Gaussian and structured noise models. Our formulation addresses the same broad unknown-noise setting from a complementary perspective: we characterize re-corruption bias through a denoiser-independent criterion, learn the re-corruption with the initial denoiser fixed, and subsequently freeze it for denoiser fine-tuning.

Since L2R does not report results on SIDD, Table 6 studies how its released method transfers to real camera noise. In the first setting, we apply the released configuration[3] directly to noisy crops of SIDD-Medium, retaining its re-corruptor, loss, denoiser, optimizer, training schedule, and stopping rule. As shown in Table 6, this transfer achieves 28.25 dB, compared with 23.66 dB for the noisy input.

We further examine the L2R objective under our training recipe. In this setting, the mean absolute amplitude of the injected perturbation decreases from 0.22 to approximately 0.003, compared with 0.067 for the observed noise. This behavior makes the objective progressively closer to the unperturbed noisy-image fitting objective and may partly explain the limited improvement in this transfer setting. When warm-started from the same f_0, Table 6 shows that L2R reaches 34.78 dB, compared with 34.96 dB for f_0, whereas our decoupled procedure reaches 35.97 dB.

These results should be interpreted as a transfer study rather than a reproduction of the original L2R experiments. L2R is primarily designed and validated for robustness across diverse unknown synthetic noise models, where it reports strong performance; here we examine its transfer to real camera noise.

C.4 Qualitative effect of learned re-corruption.

Figure 3 isolates the effect of the learned re-corruption in the unknown-noise setting. The two models share the same initial denoiser and differ only in the re-corruption used for fine-tuning. Compared with the baseline re-corruption, the learned z better preserves fine structures and produces visibly smaller reconstruction errors. This is further confirmed by the residual maps in the last two columns, where both the residual magnitude and the full-block MSE are consistently reduced.

[3] https://github.com/bemc22/learning2recorrupt

Table 6: Transfer study of L2R on real camera noise. Results are reported on SIDD validation in PSNR (dB). The noisy input achieves 23.66 dB and the common initial denoiser f_0 achieves 34.96 dB.

Setting	Denoiser	SIDD validation
Released L2R configuration		
Released configuration	L2R default	28.25
Restormer, learning rate $3 \cdot 10^{-4}$	Restormer	27.57
Restormer, learning rate 10^{-3}	Restormer	26.27
$10\times$ more images, $5\times$ fewer epochs	L2R default	24.95
L2R objective under our training recipe		
From scratch	Restormer	25.74
without gradient clipping	Restormer	25.10
Warm-started from f_0	Restormer	34.78
Ours, warm-started from the same f_0	Restormer	**35.97**

D VISUAL COMPARISONS

This appendix provides qualitative comparisons with unsupervised denoisers on camera images (Fig. 4) and fluorescence microscopy (Fig. 5).

On camera images, our method better preserves fine structures that the compared methods tend to merge: small print remains legible, and periodic textures in fabric and paper stay resolved rather than being smoothed into a uniform tone. On fluorescence microscopy, the main difference appears in the background, where the compared methods leave more visible grain.

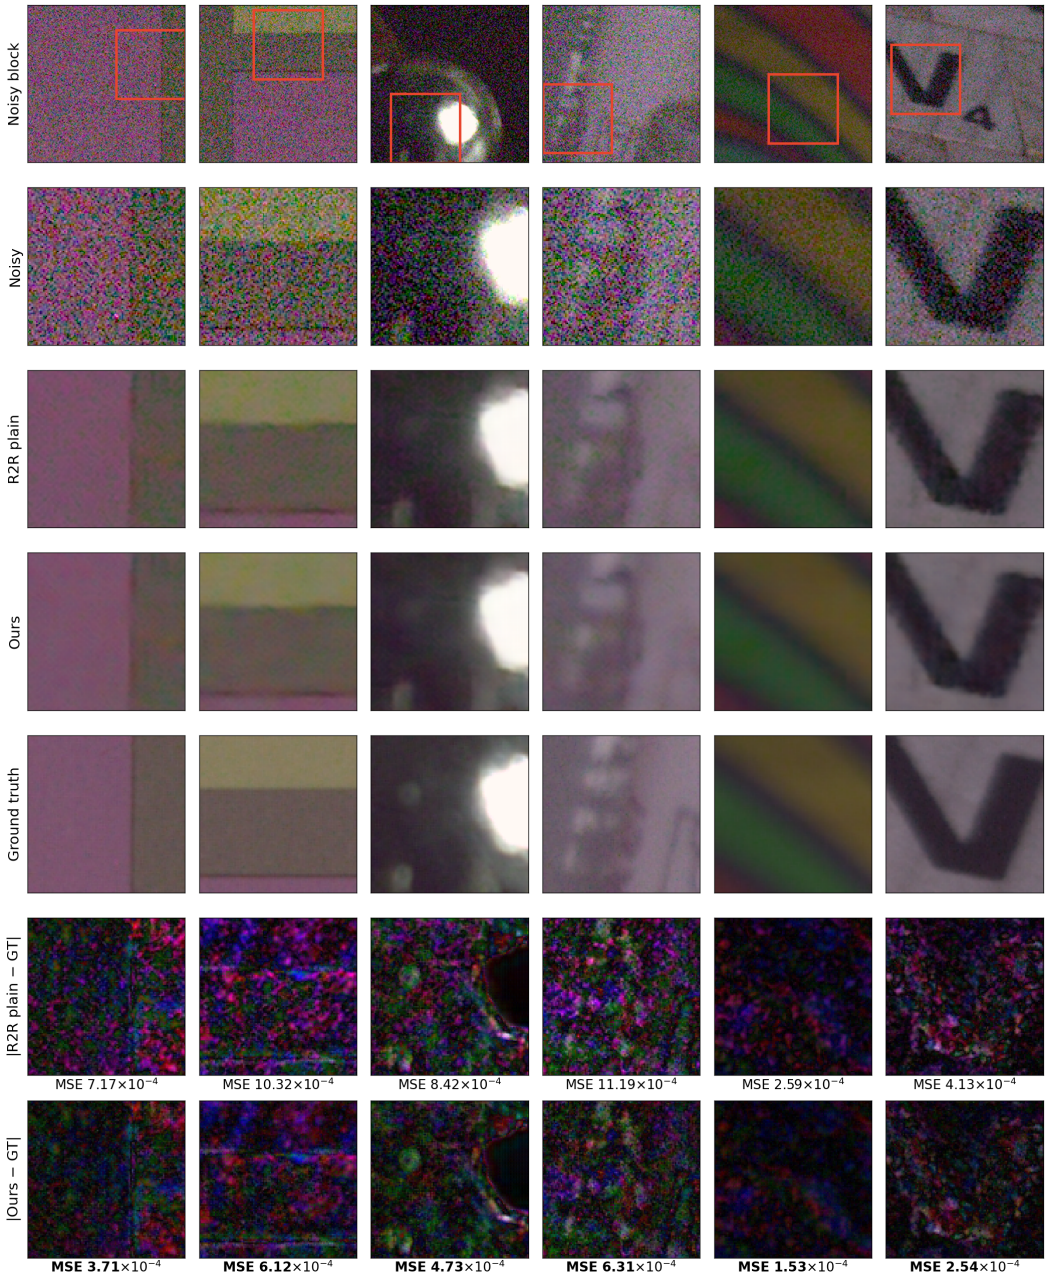

Figure 3: Effect of the learned re-corruption in the unknown-noise setting of Appendix B. The last two rows show residual maps with respect to the ground truth, amplified by $6\times$ for visualization; the corresponding full-block MSE is reported below each image.

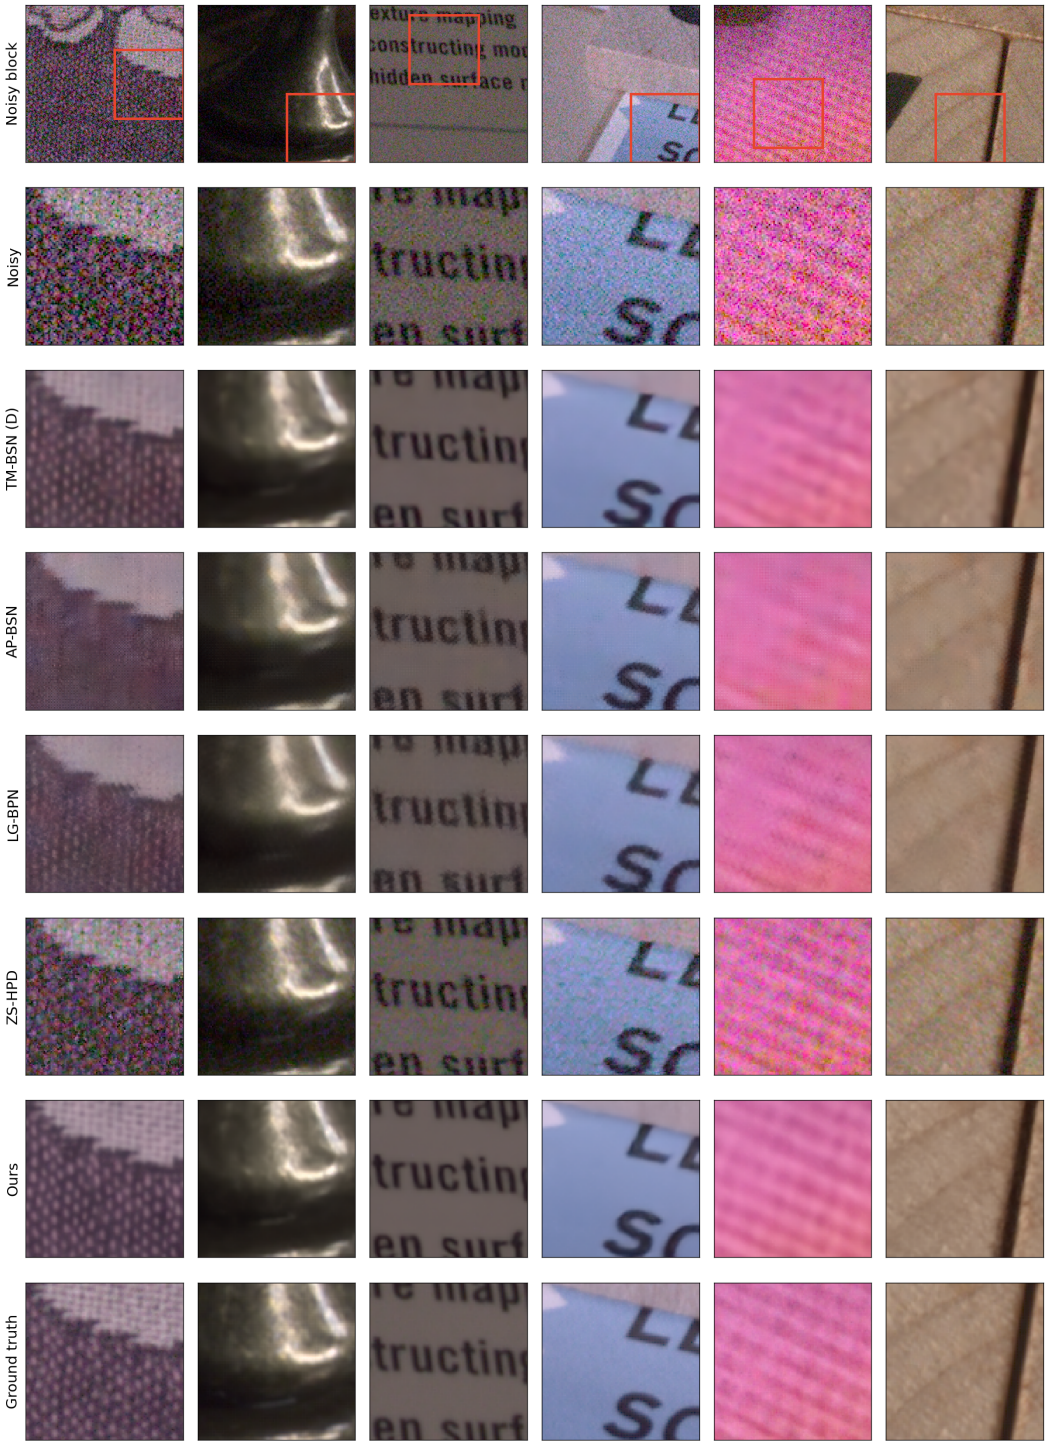

Figure 4: Visual comparison on SIDD validation, six blocks, one per column. Our method effectively suppresses noise while preserving color, texture, and fine details.

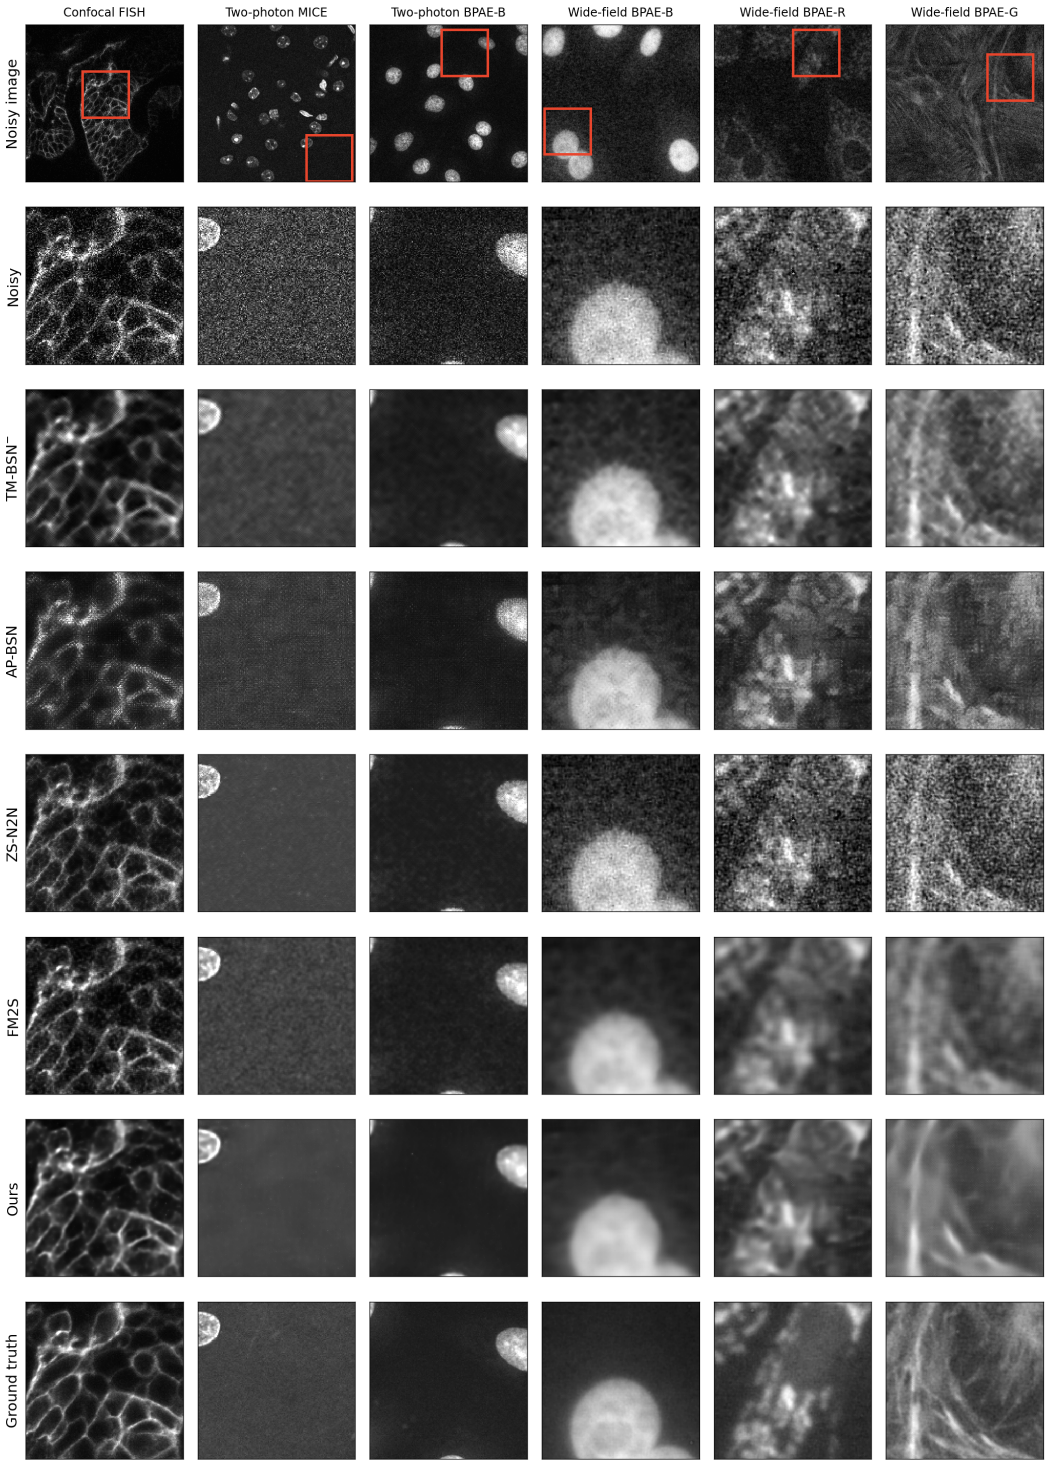

Figure 5: Visual comparison on FMD `test_mix`, six image types, one per column, covering all three modalities. Panels of a column share the display range. Our method exhibits better noise suppression, resulting in cleaner images.