LEARNING TO RECORRUPT: NOISE DISTRIBUTION AGNOSTIC SELF-SUPERVISED IMAGE DENOISING

Anonymous authors
Paper under double-blind review

ABSTRACT

Self-supervised image denoising methods have traditionally relied on architectural constraints, pseudo-pair constructions, or specialized loss functions to avoid the trivial identity mapping. Among these, approaches such as Noisier2Noise or R2R create training pairs by adding synthetic noise to noisy images. While effective, these recorruption-based approaches require precise knowledge of the noise distribution, which is often unavailable. We present Learning to Recorrupt (L2R), a self-supervised framework that does not require exact knowledge of the noise distribution. Our method introduces a learnable recorruptor jointly optimized with the denoiser through a min–max saddle-point objective. The proposed method achieves state-of-the-art performance among methods without prior knowledge of the noise distribution across unconventional and heavy-tailed noise distributions, such as log-gamma and Laplace, as well as spatially correlated noise, while obtaining competitive results on real-world denoising benchmarks characterized by signal-dependent noise.

1 INTRODUCTION

Image denoising remains an essential component of computer vision, serving as a critical preprocessing step in domains such as computational photography (Plötz & Roth, 2017), remote sensing (Rasti et al., 2021), and medical imaging (Kaur & Dong, 2023). It operates as a fundamental block of contemporary image restoration algorithms (Elad et al., 2023; Milanfar & Delbracio, 2025) and generative models (Ho et al., 2020; Li & He, 2025). While supervised image denoising has achieved remarkable success by exploiting large-scale datasets of clean–noisy pairs, its applicability is often restricted in real-world scenarios, where ground-truth images are either expensive or physically impossible to obtain (Lehtinen et al., 2018). This limitation has driven the development of self-supervised denoising, which seeks to learn the underlying image manifold directly from noisy observations (Zhang et al., 2026; Diaz-Delgado et al., 2026). The fundamental challenge in this paradigm is to define a training objective that avoids the trivial identity mapping (where the network simply learns to output the noisy input) while ensuring that the network achieves meaningful denoising.

Existing self-supervised strategies address this challenge through various structural or statistical priors. *Noise2Noise* (Lehtinen et al., 2018) utilizes independent noisy images of the same scene, though such pairs are rarely available in practice. *Blind-spot networks* (BSNs), such as those of Laine et al. (2019) and their variants (Krull et al., 2019; Wang et al., 2022; Zhang et al., 2023; Liu et al., 2025), impose architectural constraints that prevent a pixel from seeing itself, effectively leveraging spatial correlation to predict noise-free values. However, BSNs often suffer from reduced performance due to the missing central pixel information and are computationally expensive. Downsampling/masking-based strategies, such as Noise2Self (Batson & Royer, 2019) and Neighbor2Neighbor (NBR2NBR) (Huang et al., 2021), instead limit the information presented to the denoiser to construct self-supervised targets from a single noisy image, without imposing a blind-spot architecture. However, these approaches typically rely on a pixel-wise noise independence assumption, which limits their effectiveness in the presence of spatially correlated corruptions. Alternatively, methods based on *Stein's Unbiased Risk Estimator* (SURE) and its generalizations (Hudson, 1978; Eldar, 2008) utilize the divergence of the denoiser to estimate the supervised loss (Ramani et al., 2008). While powerful, these divergence-based methods typically require explicit knowledge of the noise likelihood.

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: **Non-Gaussian Image Denoising.** L2R is capable of handling complex non-Gaussian noise models without requiring prior knowledge of the noise distribution statistics, including log-gamma, Laplace, and correlated noise.

Recently, *UNSURE* (Tachella et al., 2025a) and other extensions (Herbreteau & Meunier, 2026) proposed a middle ground between BSNs and SURE by tackling the case where the noise distribution is known, but some of its parameters, such as the covariance, are unknown.

Recorruption-based methods, such as Noisier2Noise (Moran et al., 2020) and R2R (Recorrupted-to-Recorrupted) (Pang et al., 2021), construct training pairs by adding noise to a single noisy image. More recently, the *Generalized Recorrupted-to-Recorrupted framework* (GR2R) (Monroy et al., 2025) extended this mechanism to additive and natural exponential-family noise, providing a computationally efficient, architecture-agnostic, and unbiased proxy for the supervised risk. However, GR2R requires knowledge of the noise distribution to construct a valid recorruption; unknown noise statistics therefore require manual tuning. Related blind-denoising approaches, such as GAN2GAN (Cha et al., 2020) and DCD-Net (Zou et al., 2023), learn noise models from single noisy images and generate Noise2Noise pairs from progressively denoised estimates. Furthermore, complex noise can be modeled with normalizing flows (Abdelhamed et al., 2019), while recent noise-translation work learns to map real noise into Gaussian noise for a pretrained denoiser (Ha et al., 2026); however, these approaches typically require clean–noisy pairs, leaving a gap for noise distribution agnostic approaches that work with single noisy observations (Lee et al., 2022).

To overcome this limitation, we propose **Learning to Recorrupt (L2R)**, a unified framework based on a noise distribution agnostic objective that extends recorruption-based self-supervision to settings with unknown noise statistics. L2R casts denoising as a min–max saddle-point problem in which the recorruption mechanism is treated as a learnable model whose parameterization can handle i.i.d., spatially correlated, and signal-dependent corruption structures. By employing a monotonic neural network, the recorruptor learns to map a standard Gaussian distribution into an auxiliary distribution satisfying the GR2R-compatible statistics. Rather than constraining the restoration architecture to enable self-supervision, L2R shifts the statistical constraint to the learned recorruptor, avoiding blind-spot or masking constraints on the denoiser. This min–max optimization jointly adapts the auxiliary recorruption statistics and trains the denoiser to reduce the discrepancy between the self-supervised objective and its supervised counterpart, without requiring exact knowledge of the noise distribution. This enables robust denoising across noise types such as log-gamma, Laplace, and spatially correlated noise, as illustrated in Figure 1. The resulting framework retains the same objective across these corruption structures and the architectural neutrality of GR2R while providing a robust, data-driven approach for denoising across diverse and heavy-tailed noise regimes, with competitive performance on real-world signal-dependent noise.

2 Related Work

Given the rapid growth of self-supervised learning for denoising and imaging inverse problems, this section concentrates on the core methodological frameworks aligned with our formulation, namely supervised risk minimization, unbiased risk estimation, and recorruption-based training objectives. For a more comprehensive overview (including additional self-supervised paradigms for noisy and incomplete measurements, their theoretical foundations, and practical imaging applications), we refer the reader to the recent survey (Tachella & Davies, 2026).

2.1 SUPERVISED AND SELF-SUPERVISED LEARNING

The goal of supervised learning is to learn a deep denoiser operator $\hat{\boldsymbol{x}} = \hat{f}(\boldsymbol{y})$ given clean–noisy paired data $(\boldsymbol{y}, \boldsymbol{x})$. The population-optimal denoiser is defined by

$$f^* = \arg\min_f \ \mathbb{E}_{\boldsymbol{x},\boldsymbol{y}} \, \mathcal{L}_{\mathrm{SUP}}(\boldsymbol{x}, \boldsymbol{y}; f), \tag{1}$$

where $\mathcal{L}_{\mathrm{SUP}}(\boldsymbol{x}, \boldsymbol{y}; f) = \|f(\boldsymbol{y}) - \boldsymbol{x}\|_2^2$ is the MSE between the estimated and clean images[1]. Under the squared loss, the population optimum satisfies $f^*(\boldsymbol{y}) = \mathbb{E}\{\boldsymbol{x} \mid \boldsymbol{y}\}$. In practice, the population expectation is approximated with a finite dataset of N samples, yielding the empirical estimator $\hat{f}$ through $\frac{1}{N}\sum_{i=1}^{N} \mathcal{L}_{\mathrm{SUP}}(\boldsymbol{x}^{(i)}, \boldsymbol{y}^{(i)}; f)$. By simple algebraic manipulation, $\mathcal{L}_{\mathrm{SUP}}$ can be split into self-supervised and supervised terms as follows

$$\mathcal{L}_{\mathrm{SUP}}(\boldsymbol{x}, \boldsymbol{y}; f) = \|f(\boldsymbol{y}) - \boldsymbol{y}\|_2^2 + 2f(\boldsymbol{y})^\top(\boldsymbol{y} - \boldsymbol{x}) + \mathrm{const.} \tag{2}$$

When there is no access to the set of clean images $\boldsymbol{x}$, self-supervised methods aim to approximate or eliminate the term that depends on $\boldsymbol{x}$, namely $f(\boldsymbol{y})^\top(\boldsymbol{y} - \boldsymbol{x})$, by building a self-supervised loss $\mathcal{L}_{\mathrm{SELF}}$ such that

$$\mathbb{E}_{\boldsymbol{y}}\mathcal{L}_{\mathrm{SELF}}(\boldsymbol{y}; f) = \mathbb{E}_{\boldsymbol{y},\boldsymbol{x}}\mathcal{L}_{\mathrm{SUP}}(\boldsymbol{x}, \boldsymbol{y}; f) + \mathrm{const.} \tag{3}$$

2.2 STEIN'S UNBIASED RISK ESTIMATOR AND EXTENSIONS

A more realistic scenario is the self-supervised case, with access only to unpaired noisy measurements. Assuming Gaussian observations $\boldsymbol{y}|\boldsymbol{x} \sim \mathcal{N}(\boldsymbol{x}, \sigma^2\mathbf{I})$ and a weakly differentiable estimator $f : \mathbb{R}^n \to \mathbb{R}^n$, Stein's unbiased risk estimator (SURE) defines the self-supervised loss as

$$\mathcal{L}_{\mathrm{SURE}}(\boldsymbol{y}; f) = \|f(\boldsymbol{y}) - \boldsymbol{y}\|_2^2 + 2\sigma^2 \operatorname{div} f(\boldsymbol{y}), \tag{4}$$

with $\operatorname{div} f(\boldsymbol{y}) := \sum_{i=1}^{n} \frac{\partial f_i}{\partial y_i}(\boldsymbol{y})$ which relies on Stein's identity (Stein, 1972)

$$\mathbb{E}_{\boldsymbol{y}|\boldsymbol{x}}\Big[\big(f(\boldsymbol{y}) - \boldsymbol{x}\big)^\top(\boldsymbol{y} - \boldsymbol{x})\Big] = \sigma^2 \, \mathbb{E}_{\boldsymbol{y}|\boldsymbol{x}}[\operatorname{div} f(\boldsymbol{y})] \,, \tag{5}$$

so that $\mathbb{E}_{\boldsymbol{y}|\boldsymbol{x}}[\mathcal{L}_{\mathrm{SURE}}(\boldsymbol{y}; f)]$ is an unbiased estimator of the supervised risk $\mathbb{E}_{\boldsymbol{y}|\boldsymbol{x}}\mathcal{L}_{\mathrm{SUP}}$ up to f-independent constants, with the minimizer approaching the MMSE estimator. Noteworthy extensions beyond i.i.d. Gaussian noise include Hudson's lemma (Hudson, 1978) and GSURE (Eldar, 2008), which extends SURE to the exponential family.

For correlated Gaussian noise $\boldsymbol{y}|\boldsymbol{x} \sim \mathcal{N}(\boldsymbol{x}, \boldsymbol{\Sigma})$, the corresponding Stein identity replaces $\sigma^2 \operatorname{div} f(\boldsymbol{y})$ with a $\boldsymbol{\Sigma}$-weighted divergence term $\operatorname{tr}\Big(\boldsymbol{\Sigma}\, \frac{\partial f}{\partial \boldsymbol{y}}(\boldsymbol{y})\Big)$.

UNSURE (Tachella et al., 2025a) further addresses *unknown* noise statistics by replacing the explicit variance-dependent correction with a constraint on the expected divergence. In its basic form, UNSURE considers the denoising objective

$$\min_f \ \mathbb{E}_{\boldsymbol{y}} \, \|f(\boldsymbol{y}) - \boldsymbol{y}\|_2^2 \quad \text{s.t.} \quad \mathbb{E}_{\boldsymbol{y}}\big[\operatorname{div} f(\boldsymbol{y})\big] = 0, \tag{6}$$

referred to as the *zero expected divergence* constraint. Equivalently, a Lagrangian (saddle) form is

$$\min_f \ \max_{\eta \in \mathbb{R}} \ \mathbb{E}_{\boldsymbol{y}} \, \|f(\boldsymbol{y}) - \boldsymbol{y}\|_2^2 + 2\eta \, \mathbb{E}_{\boldsymbol{y}}\big[\operatorname{div} f(\boldsymbol{y})\big]. \tag{7}$$

2.3 GENERALIZED RECORRUPTED-TO-RECORRUPTED

The GR2R framework extends the R2R approach (Pang et al., 2021), originally proposed for Gaussian noise, to additive noise and the natural exponential family (including Poisson, Gamma, hypergeometric, among others), exhibiting a close connection with SURE. In the additive noise setting, given a noisy observation $\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon}$ denotes the measurement noise, GR2R generates a recorrupted pair as follows

$$\boldsymbol{y}_1 = \boldsymbol{y} + \tau\boldsymbol{\omega}, \quad \boldsymbol{y}_2 = \boldsymbol{y} - \boldsymbol{\omega}/\tau, \tag{8}$$

[1] It is important to highlight that other cost functions can be used, such as the mean absolute error.

where $\tau > 0$ and $\boldsymbol{\omega}$ is drawn independently of $\boldsymbol{\varepsilon}$. In the original R2R setting (Gaussian noise with known variance), one may take $\boldsymbol{\omega} \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I})$. In the general additive non-Gaussian case, $\boldsymbol{\omega}$ is instead chosen so that its low-order moments match those of $\boldsymbol{\varepsilon}$, e.g.,

$$\mathbb{E}_{\varepsilon_i,\omega_i}\left[\left(\varepsilon_i - \frac{\omega_i}{\tau}\right)(\varepsilon_i + \tau\omega_i)^k\right] = 0, \qquad i = 1,\dots,n, \quad k = 0,\dots,K. \tag{9}$$

Here, K denotes the highest moment order enforced by the matching conditions.

Then, the self-supervised denoising loss is defined as

$$\mathcal{L}_{\mathrm{GR2R}}(\boldsymbol{y}; f) = \mathbb{E}_{\boldsymbol{y}_1,\boldsymbol{y}_2|\boldsymbol{y}}\big\|f(\boldsymbol{y}_1) - \boldsymbol{y}_2\big\|_2^2. \tag{10}$$

Under moment-matching conditions such as those in Eq. 9 (for instance, in the case of Gaussian noise with variance-matching $\boldsymbol{\omega} \overset{d}{=} \boldsymbol{\varepsilon}$), the population GR2R objective becomes, up to an additive constant, equivalent to a supervised denoising objective on $\boldsymbol{y}_1$, that is, $\mathbb{E}_{\boldsymbol{y}}\mathcal{L}_{\mathrm{GR2R}}(\boldsymbol{y}; f) \propto \mathbb{E}_{\boldsymbol{y}_1|\boldsymbol{x}}\mathcal{L}_{\mathrm{SUP}}(\boldsymbol{x}, \boldsymbol{y}_1; f)$, with the denoiser model satisfying $f^*(\boldsymbol{y}_1) \approx \mathbb{E}\{\boldsymbol{x} \mid \boldsymbol{y}_1\}$.

3 Learning to Recorrupt

We consider the additive denoising model

$$\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{\varepsilon}, \qquad \boldsymbol{\varepsilon} \perp \boldsymbol{x}, \qquad \mathbb{E}[\boldsymbol{\varepsilon}] = 0, \tag{11}$$

where $\boldsymbol{y} \in \mathbb{R}^n$ denotes the noisy observation, $\boldsymbol{x} \in \mathbb{R}^n$ is the unknown clean image, and $\boldsymbol{\varepsilon} \sim p_{\boldsymbol{\varepsilon}}$ follows an *unknown* distribution. Specifically, let $\boldsymbol{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and define

$$\boldsymbol{\varepsilon} = g(\boldsymbol{w}), \tag{12}$$

for a map $g : \mathbb{R}^n \to \mathbb{R}^n$ that pushes $\mathcal{N}(\mathbf{0}, \mathbf{I})$ forward to the target noise distribution, i.e., $g_{\#}\mathcal{N}(\mathbf{0}, \mathbf{I}) = p_{\boldsymbol{\varepsilon}}$, where $g_{\#}$ denotes the pushforward measure induced by g. With this reparameterization, the observation model becomes

$$\boldsymbol{y} = \boldsymbol{x} + g(\boldsymbol{w}), \qquad \boldsymbol{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \tag{13}$$

so the unknown corruption is represented as the pushforward of a Gaussian distribution. We follow the intuition of GR2R and construct a recorrupted image $\boldsymbol{y}_1$ from the single noisy image $\boldsymbol{y}$. In classical GR2R, the choice of recorruption relies on explicit knowledge of the noise distribution. To overcome this limitation, we propose a learned alternative by introducing an auxiliary recorruption mapping $h : \mathbb{R}^n \to \mathbb{R}^n$ constrained to a class $\mathcal{H}$ of admissible vector-valued recorruption maps that are (a.e.) differentiable, and whose outputs have zero mean, i.e., $\mathbb{E}_{\boldsymbol{w}' \sim \mathcal{N}(\mathbf{0},\mathbf{I})}[h(\boldsymbol{w}')] = \mathbf{0}$. This mapping generates a recorrupted observation as

$$\boldsymbol{y}_1 = \boldsymbol{y} + \tau h(\boldsymbol{w}'), \tag{14}$$

where $\boldsymbol{w}' \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $\tau > 0$ controls the strength of the recorruption.

As in Eq. 2, we can expand the simple measurement consistency loss into a supervised term and a term measuring correlation with the noise:

$$\begin{aligned} &\mathbb{E}_{\boldsymbol{y}_1,\boldsymbol{y}}\|f(\boldsymbol{y}_1) - \boldsymbol{y}\|_2^2 \\ &= \mathbb{E}_{\boldsymbol{y}_1,\boldsymbol{x}}\|f(\boldsymbol{y}_1) - \boldsymbol{x}\|_2^2 - 2\,\mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'} f(\boldsymbol{x} + g(\boldsymbol{w}) + \tau h(\boldsymbol{w}'))^\top g(\boldsymbol{w}) + \text{const.} \end{aligned} \tag{15}$$

Since g is unknown, we seek to control this correlation through a family of constraints:

$$f^* = \arg\min_f \ \mathbb{E}_{\boldsymbol{y}_1,\boldsymbol{y}}\|f(\boldsymbol{y}_1) - \boldsymbol{y}\|_2^2 \quad \text{s.t. } \mathbb{E}_{\boldsymbol{y},\boldsymbol{w}'} f(\boldsymbol{y} + \tau h(\boldsymbol{w}'))^\top h(\boldsymbol{w}') = 0 \quad \forall h \in \mathcal{H}. \tag{16}$$

An auxiliary recorruption map is compatible when its correlation term is proportional to the unknown correlation with $g(\boldsymbol{w})$, according to Eq. 9. If $\mathcal{H}$ contains such a map h_c, then $\mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'} f^*(\boldsymbol{x} + g(\boldsymbol{w}) + \tau h_c(\boldsymbol{w}'))^\top g(\boldsymbol{w}) = 0$, and, for this fixed compatible map, Eq. 16 is equivalent to the *supervised* constrained minimization problem:

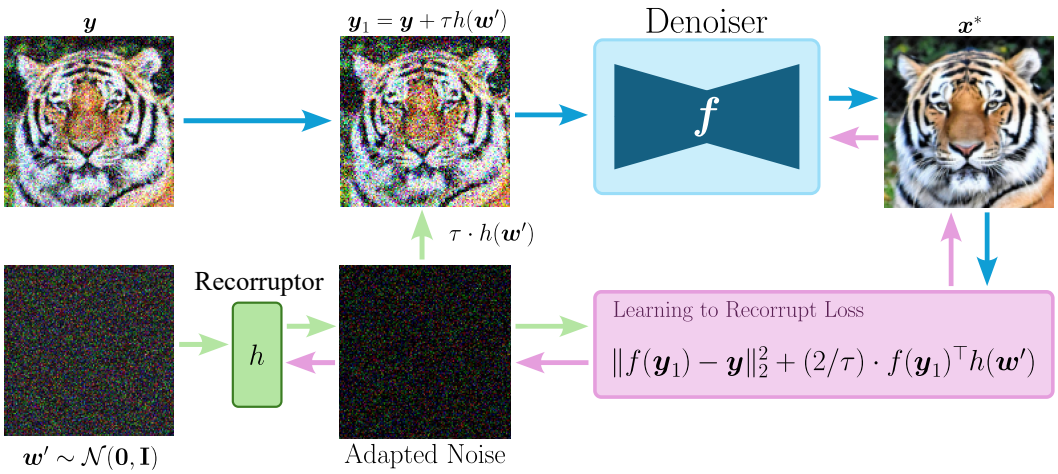

Figure 2: **Learning to Recorrupt.** Given a noisy image $\boldsymbol{y}$ with unknown noise, L2R generates a recorrupted version $\boldsymbol{y}_1$ through a learned recorruptor h to enable self-supervised denoiser training.

$$f^* = \arg\min_f \ \mathbb{E}_{\boldsymbol{y}_1,\boldsymbol{x}} \|f(\boldsymbol{y}_1) - \boldsymbol{x}\|_2^2 \quad \text{s.t. } \mathbb{E}_{\boldsymbol{y},\boldsymbol{w}'} f(\boldsymbol{y} + \tau h(\boldsymbol{w}'))^\top h(\boldsymbol{w}') = 0 \quad \forall h \in \mathcal{H}. \tag{17}$$

Further analysis is provided in Appendix A.1. Assuming that the output of f is bounded and that $\mathcal{H}$ is a conic set[2], the constraints in Eq. 16 motivate the following adversarial saddle-point objective:

$$\begin{gathered} \min_f \max_{h \in \mathcal{H}} \ \mathbb{E}_{\boldsymbol{y}}[\mathcal{L}_{\text{L2R}}(\boldsymbol{y}; f, h)], \\ \mathcal{L}_{\text{L2R}}(\boldsymbol{y}; f, h) := \mathbb{E}_{\boldsymbol{w}'} \left[\|f(\boldsymbol{y} + \tau h(\boldsymbol{w}')) - \boldsymbol{y}\|_2^2 + \frac{2}{\tau} f(\boldsymbol{y} + \tau h(\boldsymbol{w}'))^\top h(\boldsymbol{w}') \right]. \end{gathered} \tag{18}$$

In practice, we represent both f and h as neural networks and alternate between gradient descent with respect to f and gradient ascent with respect to h, as in adversarial learning (Arjovsky & Bottou, 2017). We use a stop-gradient surrogate for the recorruptor update. Unlike blind-spot networks (Laine et al., 2019), which impose architectural constraints to prevent access to the central pixel, the denoiser f in L2R can be parametrized using any architecture.

Designing the learnable recorruptor. The architecture and capacity of the recorruptor h determine the range of auxiliary noise statistics it can represent, providing a way to control the trade-off between robustness to unknown noise distributions, and the gap with respect to the optimal supervised loss Tachella et al. (2025a); Tachella & Davies (2026). Therefore, we propose four h variants for two scenarios encountered in practice: non-signal-dependent noise, with (1.a) an elementwise recorruptor and (1.b) a spatial recorruptor; and signal-dependent noise, with (2.a) a local recorruptor and (2.b) a global recorruptor.

For the **non-signal-dependent scenario**, we define

$$h(\boldsymbol{w}') = \boldsymbol{k}_h * \mathrm{N}(\mathrm{mMLP}(\boldsymbol{w}')) \,, \tag{19}$$

where mMLP is a monotonic network (Runje & Shankaranarayana, 2023), inspired by the inverse-transform principle that a monotonic map can transform a Gaussian variable into a target distribution. The operator $\mathrm{N}(\cdot)$ normalizes its output to zero mean and unit variance. Variant (1.a) uses a 1×1 kernel $\boldsymbol{k}_h$ to model non-correlated noise, while variant (1.b) uses a spatial kernel to model spatially correlated noise.

For the **signal-dependent scenario**, we define

$$h(\boldsymbol{w}', \boldsymbol{y}) = s(\boldsymbol{y}) \odot \tilde{h}_{\boldsymbol{y}}(\boldsymbol{w}'), \tag{20}$$

[2]For all $\alpha > 0$ and $h \in \mathcal{H}$, the scaled map αh also belongs to $\mathcal{H}$. For the additive parameterization in Eq. 19, this follows by scaling the learnable kernel $\boldsymbol{k}_h$ by α.

where $s(\boldsymbol{y})$ sets the recorruption scale as a function of the observation, and $\tilde{h}_{\boldsymbol{y}}$ follows the architecture in Eq. 19. In the local variant (2.a), the weights of $\tilde{h}_{\boldsymbol{y}}$ are shared across images; in the global variant (2.b), a hypernetwork predicts image-specific weights from $\boldsymbol{y}$. Further implementation details are provided in Appendix A.2.1.

Interpretation of h at convergence. At a saddle point (f^*, h^*), we can rewrite Eq. 18 as

$$\min_f \mathbb{E}_{\boldsymbol{y}_1,\boldsymbol{x}} \|f(\boldsymbol{y}_1) - \boldsymbol{x}\|_2^2 + 2\,\mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'} f(\boldsymbol{y} + \tau h^*(\boldsymbol{w}'))^\top \left(\frac{1}{\tau} h^*(\boldsymbol{w}') - g(\boldsymbol{w})\right).$$

The second term provides a measure of alignment with the supervised objective, which we examine empirically through the following relation:

$$\mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'} f^*(\boldsymbol{y} + \tau h^*(\boldsymbol{w}'))^\top \frac{1}{\tau} h^*(\boldsymbol{w}') \approx \mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'} f^*(\boldsymbol{y} + \tau h^*(\boldsymbol{w}'))^\top g(\boldsymbol{w}). \tag{21}$$

Thus, the learned auxiliary corruption tracks the true-noise contribution through the denoiser-weighted correlation. This relation is evaluated numerically in Appendix A.5.

Relationship to GR2R and UNSURE. The proposed self-supervised objective can be seen as a generalization of the UNSURE loss in Eq. 6. L2R uses a more flexible set of constraints parametrized by a (monotonic) neural network. The method can also be linked to GR2R in Eq. 10, as the loss in Eq. 18 can be rewritten as

$$\min_f \max_h \; \mathbb{E}_{\boldsymbol{y}_1,\boldsymbol{y}_2} \|f(\boldsymbol{y}_1) - \boldsymbol{y}_2\|_2^2 + \mathbb{E}_{\boldsymbol{y},\boldsymbol{w}'} (\boldsymbol{y} + \boldsymbol{y}_2)^\top h(\boldsymbol{w}')/\tau,$$

with $\boldsymbol{y}_2 = \boldsymbol{y} - \frac{1}{\tau} h(\boldsymbol{w}')$, where the recorruption process is learned instead of being fixed. In GR2R, the second term is omitted because it is not optimized. See Appendices A.1.1 and A.1.2 for details.

4 Simulations and Results

This section evaluates the proposed **L2R** framework in the distribution-agnostic self-supervised denoising setting, where the corruption follows an unknown noise distribution and clean targets are unavailable. We compare it with noise distribution-aware recorruption methods **GR2R** (Monroy et al., 2025) and **R2R** (Pang et al., 2021), divergence-based estimators **(PG)-SURE** (Ramani et al., 2008; Le Montagner et al., 2014) and **(PG)-UNSURE** (Tachella et al., 2025a), noise distribution agnostic baselines **NBR2NBR** (Huang et al., 2021) and **Blind2Sound** (Liu et al., 2025), and **Supervised** training as an upper bound. Note that GR2R can be seen as an oracle in our framework: our method reduces to GR2R when the learned recorruption operator matches the true noise-induced recorruption assumed in GR2R.

Datasets. Experiments are conducted on BSDS500 following the standard $200/100/200$ train/validation/test split, and all quantitative results are reported on the test subset. During training, images are randomly cropped into patches of spatial resolution 256×256 pixels. To assess robustness under distribution shift, the DIV2K dataset is employed exclusively for evaluation: performance is measured on the 100 validation images, using 512×512 pixel center crops.

Training setup. For synthetic-noise experiments, all methods employ a lightweight DRUNet (Zhang et al., 2021) variant, comprising approximately 2M trainable parameters. A single noisy realization is generated for each training image and kept fixed during training. The models are trained with a batch size of 32 for 2000 epochs using the AdamW optimizer with a weight decay of 0.01 and a cosine-annealed learning-rate schedule decreasing from 10^{-4} to 10^{-6}, with $\tau = 0.5$. For experiments on SIDD (Abdelhamed et al., 2018), we adopt the U-Net architecture and training schedule of Blind2Sound (Liu et al., 2025), using Adam for 400 epochs with $\tau = 1.0$ and an initial learning rate of 10^{-4}, halved every 80 epochs. The recorruptor learning rate is set to half that of the denoiser. To reduce training cost, stop-gradient is applied to the input of the denoiser $f(\mathrm{sg}[\boldsymbol{y}_1])$, preventing gradients through the denoiser with respect to $\boldsymbol{y}_1$ while keeping a similar performance, as shown in Appendix A.3.3. The implementation is based on the `DeepInverse` library (Tachella et al., 2025b), and the code will be released as open source.

Noise settings. We evaluate L2R under zero-mean log-gamma and Laplace noise, spatially correlated Gaussian noise, and signal-dependent Poisson–Gaussian noise. Across these settings, the L2R

Table 1: **Non-Gaussian Denoising.** Performance under different types of non-Gaussian noise. Best and second-best self-supervised results are shown in bold and underlined, respectively. Gray rows are reference methods excluded from the ranking.

Method	p_ϵ known?	Log-gamma Noise ℓ	Log-gamma Noise BSDS500	Log-gamma Noise DIV2K	Laplace Noise b	Laplace Noise BSDS500	Laplace Noise DIV2K	Correlated Noise σ	Correlated Noise BSDS500	Correlated Noise DIV2K
Supervised	no		31.44/0.898	31.96/0.894		29.48/0.856	30.21/0.859		30.56/0.897	30.75/0.887
GR2R	yes		31.14/0.889	31.98/0.893		28.10/0.803	28.74/0.807		30.43/0.893	30.61/0.884
SURE	yes		28.33/0.796	28.61/0.781		27.07/0.748	27.60/0.747		30.37/0.891	30.60/0.882
R2R	yes		28.63/0.809	29.03/0.799		27.85/0.797	28.52/0.804		24.52/0.653	24.97/0.634
UNSURE	no	1.0	27.83/0.756	28.11/0.737	0.1	22.93/0.615	23.42/0.639	0.2	29.56/0.841	29.56/0.821
NBR2NBR	no		29.45/0.830	29.78/0.818		27.39/0.746	27.86/0.740		25.16/0.680	25.53/0.657
Blind2Sound	no		29.61/0.857	30.11/0.849		27.86/0.802	28.43/0.801		24.20/0.643	24.53/0.620
L2R (Ours)	no		**30.77/0.871**	**31.23/0.865**		**28.40/0.823**	**29.12/0.830**		**29.85/0.863**	**29.75/0.836**
Supervised	no		33.40/0.934	33.57/0.926		32.46/0.915	32.84/0.909		34.48/0.949	34.25/0.938
GR2R	yes		32.34/0.915	33.01/0.915		31.90/0.901	32.29/0.896		34.16/0.945	34.03/0.936
SURE	yes		28.26/0.804	28.73/0.792		31.18/0.879	31.42/0.869		34.41/0.948	34.20/0.937
R2R	yes		27.93/0.781	28.33/0.768		31.30/0.885	31.61/0.878		29.64/0.853	30.00/0.837
UNSURE	no	0.1	26.95/0.745	27.22/0.719	0.05	30.10/0.843	30.27/0.827	0.1	33.16/0.915	32.58/0.891
NBR2NBR	no		30.63/0.864	30.90/0.850		31.49/**0.888**	**32.10/0.886**		30.58/0.863	30.57/0.836
Blind2Sound	no		30.39/0.878	30.81/0.867		30.89/0.878	31.44/0.879		28.33/0.781	28.70/0.756
L2R (Ours)	no		**32.29/0.901**	**32.40/0.887**		**31.59**/0.880	31.79/0.868		**33.31/0.921**	**32.87/0.900**

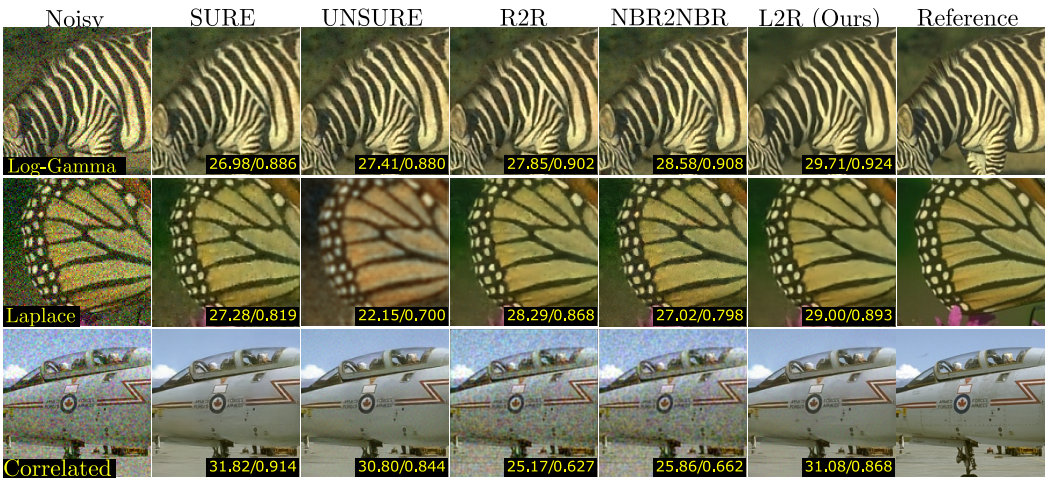

Figure 3: **Non-Gaussian Denoising.** Visual comparison on representative BSDS500 crops corrupted by log-gamma, Laplace, and spatially correlated noise. Columns show the noisy input, SURE, UNSURE, R2R, NBR2NBR, and the proposed L2R, followed by the ground truth. Numbers in each panel report PSNR (dB) / SSIM.

objective remains unchanged, while the recorruptor structure encodes only broad assumptions such as element-wise, spatially correlated, or signal-dependent corruption. The noise parameters and distribution statistics are not provided to L2R. Exact generation procedures and recorruptor settings are given in Appendix A.2.1. Additional visualizations of the learned recorruption maps and spatial kernel are provided in Appendix A.4.

Non-Gaussian Denoising.

Table 1 reports quantitative results on three representative non-Gaussian noise types, log-gamma (heavy-tailed), Laplace, and spatially correlated noise, evaluated on BSDS500 and DIV2K. In the case of log-gamma noise, L2R attains the highest PSNR among distribution-agnostic methods, reaching 30.77 dB (BSDS500) / 31.23 dB (DIV2K) in the first setting ($\ell = 1.0$) and improving to 32.29 dB / 32.40 dB in the second ($\ell = 0.1$). The strongest noise distribution-aware competitor (GR2R) remains slightly higher in PSNR and better in SSIM than L2R across both datasets, while UNSURE, NBR2NBR, and Blind2Sound exhibit larger gaps, particularly under the heavier-tailed regime. Overall, these results indicate that L2R preserves strong performance under heavy-tailed corruptions without requiring access to p_ϵ.

Regarding Laplace noise, L2R also provides the best PSNR in the first setting, achieving 28.40 dB on BSDS500 and 29.12 dB on DIV2K. In the second setting, L2R remains among the top-performing

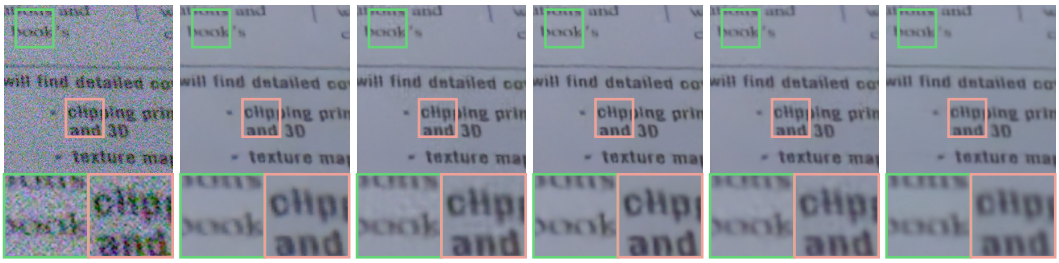

Figure 4: Visualization of qualitative results on a sample image from the SIDD benchmark.

methods, reaching 31.59 dB and 31.79 dB, with only small gaps to the best model-aware alternatives. Notably, the second-best method changes across Laplace settings (and across datasets), whereas L2R remains among the strongest distribution-agnostic methods across these settings.

For spatially correlated noise, SURE yields the highest PSNR when the correlation kernel is known, but L2R is the best self-supervised alternative without access to $p_{\boldsymbol{\varepsilon}}$. In the first setting, L2R reaches 29.85 dB (BSDS500) and 29.75 dB (DIV2K), and in the second it improves to 33.31 dB and 32.87 dB, substantially outperforming other distribution-agnostic baselines and narrowing the gap to SURE. Overall, while the strongest competing baseline varies with the noise distribution, L2R remains the most competitive self-supervised option across all non-Gaussian regimes considered.

The visual comparison in Figure 3 supports the quantitative trends. Under log-gamma and Laplace noise, L2R better suppresses granular artifacts while preserving high-frequency textures (e.g., zebra stripes and fine edges) compared to SURE/UNSURE and prior self-supervised baselines. Under correlated noise, L2R reduces spatially structured residuals more effectively than distribution-agnostic competitors, producing reconstructions that are visually closer to the ground truth while avoiding oversmoothing. Overall, L2R provides the best robustness–performance trade-off: it achieves state-of-the-art results among methods that do not assume access to the underlying noise distribution while remaining competitive with noise distribution-aware approaches in structured noise.

Real-World Denoising.

Table 2 reports results on SIDD, which contains real, spatially varying, and signal-dependent camera noise. The estimated training cost is computed as the ratio of the mean training image throughput of Supervised to that of each method, with Supervised normalized to $\times 1$. L2R with the global recorruptor from Eq. 20 achieves 51.35 dB and 0.992 SSIM on the validation set and 50.60 dB and 0.991 SSIM on the benchmark. Figure 4 supports these quantitative results: L2R removes strong chromatic noise while preserving text contours, producing a reconstruction visually comparable to Supervised and specialized blind-spot methods. L2R outperforms NBR2NBR by 0.29 dB on validation and 0.13 dB on the benchmark, while reducing the estimated training cost from $\times 2$ to $\times 1.4$. It remains within 0.15 dB of Blind2Sound on validation and 0.32 dB on the benchmark at approximately one-third of its estimated training cost. Compared with Blind2Unblind, L2R obtains similar accuracy with approximately $12\times$ lower training cost. Both L2R and Blind2Unblind attain the same benchmark SSIM at the reported precision.

Notably, L2R closely matches both the Supervised and Noise2Noise reference objectives. On the SIDD benchmark, L2R reaches 50.60 dB, compared with 50.61 dB for Supervised and 50.62 dB for Noise2Noise, while all three obtain an SSIM of 0.991. On validation, L2R slightly exceeds both references in PSNR and SSIM. Since Supervised and Noise2Noise share the same population MSE minimizer under standard zero-mean noise assumptions, these results indicate that the learned recorruption objective can recover performance comparable to clean-target and paired-noisy training without either form of supervision or a blind-spot architecture.

Real Cryo-EM Data Experiment.

We evaluate the proposed L2R method on real Cryo-EM data from the Topaz-EM open-source library (Bepler et al., 2020), for which the noise distribution is unknown. The dataset comprises five images with a resolution of 7676×7420 pixels. We employ the local signal-dependent recorruptor from Eq. 20. A single denoiser–recorruptor pair is trained jointly across all five images, with the

Table 2: Real-world denoising performance on SIDD. Training cost is reported relative to Supervised training. Best and second-best self-supervised results are highlighted.

Methods	Blind-spot	Train Cost	SIDD Val	SIDD Bench
Supervised	no	×1	51.19/0.991	50.61/0.991
Noise2Noise	no	×1	51.21/0.991	50.62/0.991
Noisy-as-Clean (NAC)	no	×1	43.24/0.961	-
R2R	no	×1	47.20/0.980	46.70/0.978
Noise2Void (N2V)	yes	×1	48.55/0.984	48.01/0.983
NBR2NBR	no	×2	51.06/0.991	50.47/0.990
DCD-Net	yes	×3	<u>51.40</u>/**0.992**	-
Blind2Sound (B2S)	yes	×4	**51.50/0.992**	**50.92/0.991**
Blind2Unblind (B2U)	yes	×17	51.36/**0.992**	<u>50.79</u>/**0.991**
L2R (Ours)	no	×1.4	51.35/**0.992**	50.60/**0.991**

monotonic mapping and learned noise statistics shared across images. These statistics are learned directly from the noisy observations without prior knowledge of the noise parameters. For comparison, we include results obtained with the PG-UNSURE approach introduced by Tachella et al. (2025a). Both methods are instantiated using a four-scale bias-free U-Net. Training is performed using random 512×512 crops and the Adam optimizer, with a learning rate of 1×10^{-4} and otherwise standard hyperparameters. Both methods produce visually plausible denoising results, but L2R produces smoother background textures than PG-UNSURE while preserving visible particle structures. Unlike PG-UNSURE, L2R does not require additional denoiser evaluations.

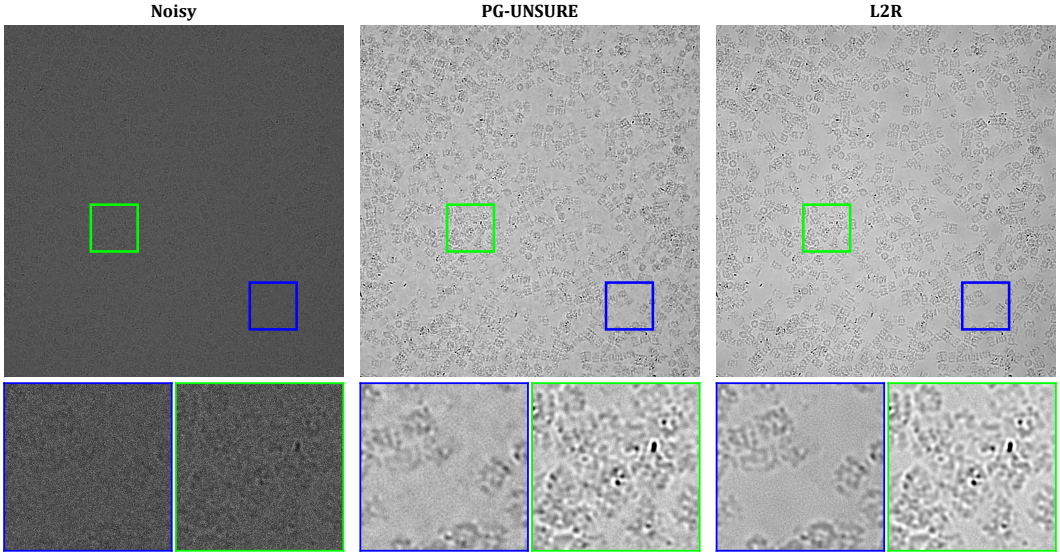

Figure 5: Blind denoising results of Cryo-EM data using the L2R method and PG-UNSURE.

5 Conclusions

This work introduced *Learning to Recorrupt* (L2R), a unified self-supervised framework combining a noise distribution agnostic objective with a learnable recorruptor. L2R learns the auxiliary GR2R statistics through a min–max objective without requiring explicit noise likelihoods, known statistics, or a moment-matched recorruption distribution a priori. Using the same objective and denoiser while adapting only the recorruptor structure, L2R achieves the highest PSNR among distribution-agnostic methods in 11 out of 12 non-Gaussian settings. On real SIDD data, it matches the Supervised and Noise2Noise references within 0.02 dB while remaining competitive with specialized blind-spot methods at substantially lower training cost. These results demonstrate the method's effectiveness across controlled and real-world noise. Beyond denoising, the learned recorruptor provides a proxy for the effective auxiliary noise statistics. Future work will explore more general recorruptor architectures for independent, correlated, and signal-dependent noise, as well as extensions of this approach to broader inverse problems with unknown measurement noise.

AI USE STATEMENT

In this work, we used generative AI tools to assist with method implementation and translation. We did not use generative AI tools to generate synthetic datasets, help develop theoretical models or conceptual frameworks, formulate mathematical claims, provide critical ingredients for proving mathematical claims, assist in writing proofs, propose or refine hypotheses, design or provide feedback on research methodology or experiments, clean or reformat datasets, support qualitative and thematic data analysis, or interpret results; these uses were not applicable to this work. Additionally, we used generative AI tools to create or modify scientific figures or images, create or edit software code, and summarize or analyze existing literature. We reviewed all AI-assisted work, and LLM-generated code was verified and tested for correctness by the authors. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

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, 2018.

Abdelrahman Abdelhamed, Marcus A Brubaker, and Michael S Brown. Noise flow: Noise modeling with conditional normalizing flows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 3165–3173, 2019.

Martin Arjovsky and Leon Bottou. Towards principled methods for training generative adversarial networks. In *International Conference on Learning Representations*, 2017.

Joshua Batson and Loïc Royer. Noise2self: Blind denoising by self-supervision. In *International Conference on Machine Learning*, pp. 524–533. PMLR, 2019.

Tristan Bepler, Kotaro Kelley, Alex J Noble, and Bonnie Berger. Topaz-denoise: general deep denoising models for cryoem and cryoet. *Nature communications*, 11(1):5208, 2020.

Sungmin Cha, Taeeon Park, Byeongjoon Kim, Jongduk Baek, and Taesup Moon. Gan2gan: Generative noise learning for blind denoising with single noisy images. In *International Conference on Learning Representations*, 2020.

Laura C Diaz-Delgado, Emmanuel Martinez, and Henry Arguello. Frozen clip priors for robust self-supervised poisson inverse problems. In *European Conference on Computer Vision*, pp. 588–604. Springer, 2026.

Michael Elad, Bahjat Kawar, and Gregory Vaksman. Image denoising: The deep learning revolution and beyond—a survey paper. *SIAM Journal on Imaging Sciences*, 16(3):1594–1654, 2023.

Yonina C Eldar. Generalized sure for exponential families: Applications to regularization. *IEEE Transactions on Signal Processing*, 57(2):471–481, 2008.

Inju Ha, Donghun Ryou, Seonguk Seo, and Bohyung Han. Learning to translate noise for robust image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5034–5043, 2026.

Sébastien Herbreteau and Étienne Meunier. Divergence-free neural networks with application to image denoising. In *The Fourteenth International Conference on Learning Representations*, 2026.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

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, 2021.

H Malcolm Hudson. A natural identity for exponential families with applications in multiparameter estimation. *The Annals of Statistics*, 6(3):473–484, 1978.

Amandeep Kaur and Guanfang Dong. A complete review on image denoising techniques for medical images. *Neural Processing Letters*, 55(6):7807–7850, 2023.

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, 2019.

Samuli Laine, Tero Karras, Jaakko Lehtinen, and Timo Aila. High-quality self-supervised deep image denoising. *Advances in neural information processing systems*, 32, 2019.

Yoann Le Montagner, Elsa D Angelini, and Jean-Christophe Olivo-Marin. An unbiased risk estimator for image denoising in the presence of mixed poisson–gaussian noise. *IEEE Transactions on Image processing*, 23(3):1255–1268, 2014.

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, 2022.

Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2Noise: Learning image restoration without clean data. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2965–2974. PMLR, 10–15 Jul 2018. URL https://proceedings.mlr.press/v80/lehtinen18a.html.

Tianhong Li and Kaiming He. Back to basics: Let denoising generative models denoise. *arXiv preprint arXiv:2511.13720*, 2025.

Jiazheng Liu, Zejin Wang, Bohao Chen, and Hua Han. Blind2sound: self-supervised image denoising without residual noise. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12937–12946, 2025.

Peyman Milanfar and Mauricio Delbracio. Denoising: a powerful building block for imaging, inverse problems and machine learning. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 383(2299), 2025.

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, 2025.

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, 2020.

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, 2021.

Tobias Plötz and Stefan Roth. Benchmarking denoising algorithms with real photographs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1586–1595, 2017.

Sathish Ramani, Thierry Blu, and Michael Unser. Monte-carlo sure: A black-box optimization of regularization parameters for general denoising algorithms. *IEEE Transactions on image processing*, 17(9):1540–1554, 2008.

Behnood Rasti, Yi Chang, Emanuele Dalsasso, Loïc Denis, and Pedram Ghamisi. Image restoration for remote sensing: Overview and toolbox. *IEEE Geoscience and Remote Sensing Magazine*, 10 (2):201–230, 2021.

Davor Runje and Sharath M Shankaranarayana. Constrained monotonic neural networks. In *International Conference on Machine Learning*, pp. 29338–29353. PMLR, 2023.

Charles Stein. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. In *Proceedings of the sixth Berkeley symposium on mathematical statistics and probability, volume 2: Probability theory*, volume 6, pp. 583–603. University of California Press, 1972.

Julián Tachella and Mike Davies. Self-supervised learning from noisy and incomplete data. *Foundations and Trends® in Signal Processing*, 20(2):85–184, 2026.

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*, 2025a.

Julián Tachella, Matthieu Terris, Samuel Hurault, Andrew Wang, Léo Davy, Jérémy Scanvic, Victor Séchaud, Romain Vo, Thomas Moreau, Thomas Davies, et al. Deepinverse: A python package for solving imaging inverse problems with deep learning. *The Journal of Open Source Software*, 10(115):8923, 2025b.

Zejin Wang, Jiazheng Liu, Guoqing Li, and Hua Han. Blind2unblind: Self-supervised image denoising with visible blind spots. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2027–2036, 2022.

Dan Zhang, Fangfang Zhou, Yuwen Jiang, and Zhengming Fu. Mm-bsn: Self-supervised image denoising for real-world with multi-mask based on blind-spot network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4189–4198, 2023.

Dan Zhang, Fangfang Zhou, Yuwen Jiang, Felix Albu, Yuanzhou Wei, Xiao Yang, and Qiang Li. Unleashing the power of self-supervised image denoising: A comprehensive review. *Computer Vision and Image Understanding*, pp. 104786, 2026.

Kai Zhang, Yawei Li, Wangmeng Zuo, Lei Zhang, Luc Van Gool, and Radu Timofte. Plug-and-play image restoration with deep denoiser prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):6360–6376, 2021.

Yunhao Zou, Chenggang Yan, and Ying Fu. Iterative denoiser and noise estimator for self-supervised image denoising. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 13219–13228. IEEE, 2023.

A Appendix

A.1 Connection to Supervised Denoising

This section examines the connection between the proposed self-supervised constrained optimization and supervised denoising. We focus on the role of an auxiliary recorruption operator h, which generates the perturbation $h(\boldsymbol{w}')$ used to construct the recorrupted observation $\boldsymbol{y}_1 = \boldsymbol{y} + \tau h(\boldsymbol{w}')$. Following Proposition 1 of Monroy et al. (2025), we consider a candidate operator h_c whose auxiliary noise $\boldsymbol{\omega} = h_c(\boldsymbol{w}')$ satisfies the mixed-moment conditions relevant to the denoiser class. For i.i.d. additive noise, these conditions are

$$\mathbb{E}_{\varepsilon_i,\omega_i}\left[\left(\varepsilon_i - \frac{\omega_i}{\tau}\right)(\varepsilon_i + \tau\omega_i)^k\right] = 0, \qquad i = 1,\dots,n, \quad k = 0,\dots,K, \tag{22}$$

where ε_i is the ith component of the measurement noise and ω_i is the ith component of $h_c(\boldsymbol{w}')$. For spatially correlated noise, the corresponding conditions are expressed through the appropriate mixed moments. These conditions characterize compatibility at the level of the required moments; they do not require the auxiliary perturbation to reproduce the measurement noise sample by sample. Consider the noisy image $\boldsymbol{y}$ and recorrupted image $\boldsymbol{y}_1$

$$\boldsymbol{y} = \boldsymbol{x} + g(\boldsymbol{w}), \qquad \boldsymbol{y}_1 = \boldsymbol{y} + \tau h(\boldsymbol{w}'), \tag{23}$$

where $\boldsymbol{w}, \boldsymbol{w}' \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ are independent and $\boldsymbol{\varepsilon} = g(\boldsymbol{w})$. Then

$$\boldsymbol{y}_1 = \boldsymbol{x} + g(\boldsymbol{w}) + \tau h(\boldsymbol{w}'). \tag{24}$$

Using $\boldsymbol{y} = \boldsymbol{x} + g(\boldsymbol{w})$, the self-supervised loss can be expanded as

$$\begin{aligned}\mathbb{E}_{\boldsymbol{y}_1,\boldsymbol{y}}\|f(\boldsymbol{y}_1) - \boldsymbol{y}\|_2^2 &= \mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'}\|f(\boldsymbol{x} + g(\boldsymbol{w}) + \tau h(\boldsymbol{w}')) - (\boldsymbol{x} + g(\boldsymbol{w}))\|_2^2 \\ &= \mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'}\|f(\boldsymbol{x} + g(\boldsymbol{w}) + \tau h(\boldsymbol{w}')) - \boldsymbol{x}\|_2^2 \\ &\quad - 2\,\mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'} f(\boldsymbol{x} + g(\boldsymbol{w}) + \tau h(\boldsymbol{w}'))^\top g(\boldsymbol{w}) + \text{const.}\end{aligned} \tag{25}$$

The term denoted by `const.` is independent of f. On the other hand, the proposed method imposes that, for every admissible recorruption map $h \in \mathcal{H}$,

$$\mathbb{E}_{\boldsymbol{y},\boldsymbol{w}'} f(\boldsymbol{y} + \tau h(\boldsymbol{w}'))^\top h(\boldsymbol{w}') = 0. \tag{26}$$

Since $\boldsymbol{y} = \boldsymbol{x} + g(\boldsymbol{w})$, this can be rewritten as

$$\mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'} f(\boldsymbol{x} + g(\boldsymbol{w}) + \tau h(\boldsymbol{w}'))^\top h(\boldsymbol{w}') = 0, \qquad \forall h \in \mathcal{H}. \tag{27}$$

Thus, the feasible set consists of denoisers whose outputs are orthogonal in expectation to every admissible recorruption direction generated by $\mathcal{H}$. Now assume that $\mathcal{H}$ contains a compatible auxiliary operator h_c satisfying Eq. 22 for the selected τ and the considered denoiser class. Since the constraint in Eq. 27 must hold for all $h \in \mathcal{H}$, it must in particular hold for the specific choice $h = h_c$, which yields

$$\mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'} f(\boldsymbol{x} + g(\boldsymbol{w}) + \tau h_c(\boldsymbol{w}'))^\top h_c(\boldsymbol{w}') = 0. \tag{28}$$

For this compatible operator, the moment relations imply the following correlation identity (Pang et al., 2021):

$$\frac{1}{\tau}\,\mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'} f(\boldsymbol{x} + g(\boldsymbol{w}) + \tau h_c(\boldsymbol{w}'))^\top h_c(\boldsymbol{w}') = \mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'} f(\boldsymbol{x} + g(\boldsymbol{w}) + \tau h_c(\boldsymbol{w}'))^\top g(\boldsymbol{w}). \tag{29}$$

Hence, the constraint associated with h_c also makes the correlation term with the true noise vanish. For any feasible f, the decomposition in Eq. 25 together with the previous argument then yields the following supervised-risk correspondence. Over the feasible set,

$$\mathbb{E}_{\boldsymbol{y}_1,\boldsymbol{y}}\|f(\boldsymbol{y}_1) - \boldsymbol{y}\|_2^2 = \mathbb{E}_{\boldsymbol{y}_1,\boldsymbol{x}}\|f(\boldsymbol{y}_1) - \boldsymbol{x}\|_2^2 + \text{const.} \tag{30}$$

Since the additive constant does not depend on f, minimizing the self-supervised objective over the constraint set

$$\arg\min_f \; \mathbb{E}_{\boldsymbol{y}_1,\boldsymbol{y}}\|f(\boldsymbol{y}_1) - \boldsymbol{y}\|_2^2 \quad \text{s.t.} \quad \mathbb{E}_{\boldsymbol{y},\boldsymbol{w}'} f(\boldsymbol{y} + \tau h(\boldsymbol{w}'))^\top h(\boldsymbol{w}') = 0 \quad \forall h \in \mathcal{H}, \tag{31}$$

can be written in the following supervised-risk form over the same constraint set:

$$\arg\min_f \; \mathbb{E}_{\boldsymbol{y}_1,\boldsymbol{x}}\|f(\boldsymbol{y}_1) - \boldsymbol{x}\|_2^2 \quad \text{s.t.} \quad \mathbb{E}_{\boldsymbol{y},\boldsymbol{w}'} f(\boldsymbol{y} + \tau h(\boldsymbol{w}'))^\top h(\boldsymbol{w}') = 0 \quad \forall h \in \mathcal{H}. \tag{32}$$

Interpretation. The central idea is that the constraint family indexed by $h \in \mathcal{H}$ should be expressive enough to contain a compatible auxiliary operator satisfying the moment relations in Eq. 22. In that case, the feasible set excludes denoisers whose outputs remain correlated with the true corruption, even though g is never directly observed. The preceding derivation therefore provides a formal connection between the self-supervised objective and the supervised denoising risk, up to a term that does not depend on the denoiser parameters. During training, the maximization over the operator h searches within $\mathcal{H}$ for auxiliary directions that enforce these constraints, while the denoiser minimizes the resulting objective. If $\mathcal{H}$ contains such a compatible operator and the saddle-point optimization is successful, the self-supervised objective can serve as a supervised-risk surrogate under the compatibility assumption; when the required moment relations cannot be represented exactly, this correspondence should be understood as approximate, which motivates choosing a sufficiently expressive class $\mathcal{H}$ for the learned recorruption operator.

A.1.1 Relationship to GR2R

The connection with GR2R becomes clearer by introducing, as in the main text, the paired observation

$$\boldsymbol{y}_1 = \boldsymbol{y} + \tau h(\boldsymbol{w}'), \qquad \boldsymbol{y}_2 = \boldsymbol{y} - \frac{1}{\tau} h(\boldsymbol{w}'), \tag{33}$$

where in the present setting the recorruption map h is *learned* rather than fixed from a known noise model. Using this notation, the GR2R-type loss takes the form

$$\mathcal{L}_{\mathrm{GR2R}}(\boldsymbol{y}; f) = \mathbb{E}_{\boldsymbol{y}_1, \boldsymbol{y}_2 | \boldsymbol{y}} \|f(\boldsymbol{y}_1) - \boldsymbol{y}_2\|_2^2. \tag{34}$$

Expanding $\boldsymbol{y}_2 = \boldsymbol{y} - \frac{1}{\tau} h(\boldsymbol{w}')$ yields

$$\begin{aligned} \|f(\boldsymbol{y}_1) - \boldsymbol{y}_2\|_2^2 &= \left\| f(\boldsymbol{y}_1) - \boldsymbol{y} + \frac{1}{\tau} h(\boldsymbol{w}') \right\|_2^2 \\ &= \|f(\boldsymbol{y}_1) - \boldsymbol{y}\|_2^2 + \frac{2}{\tau} \big(f(\boldsymbol{y}_1) - \boldsymbol{y}\big)^\top h(\boldsymbol{w}') + \frac{1}{\tau^2} \|h(\boldsymbol{w}')\|_2^2 \\ &= \|f(\boldsymbol{y}_1) - \boldsymbol{y}\|_2^2 + \frac{2}{\tau} f(\boldsymbol{y}_1)^\top h(\boldsymbol{w}') - \frac{2}{\tau} \boldsymbol{y}^\top h(\boldsymbol{w}') + \frac{1}{\tau^2} h(\boldsymbol{w}')^\top h(\boldsymbol{w}'). \end{aligned} \tag{35}$$

Rearranging the last two terms and using the definition of $\boldsymbol{y}_2$, one obtains

$$\|f(\boldsymbol{y}_1) - \boldsymbol{y}_2\|_2^2 = \|f(\boldsymbol{y}_1) - \boldsymbol{y}\|_2^2 + \frac{2}{\tau} f(\boldsymbol{y}_1)^\top h(\boldsymbol{w}') - \frac{1}{\tau} (\boldsymbol{y} + \boldsymbol{y}_2)^\top h(\boldsymbol{w}'). \tag{36}$$

Taking expectations, and recalling the proposed min–max objective

$$\min_f \max_h \; \mathbb{E}_{\boldsymbol{y}, \boldsymbol{w}'} \Big[\|f(\boldsymbol{y}_1) - \boldsymbol{y}\|_2^2 + \frac{2}{\tau} f(\boldsymbol{y}_1)^\top h(\boldsymbol{w}') \Big], \tag{37}$$

the method can be interpreted as a GR2R-type criterion in which the paired recorruption process is learned adversarially:

$$\begin{aligned} &\min_f \max_h \; \mathbb{E}_{\boldsymbol{y}, \boldsymbol{w}'} \Big[\|f(\boldsymbol{y}_1) - \boldsymbol{y}\|_2^2 + \frac{2}{\tau} f(\boldsymbol{y}_1)^\top h(\boldsymbol{w}') \Big] \\ = &\min_f \max_h \; \mathbb{E}_{\boldsymbol{y}_1, \boldsymbol{y}_2} \|f(\boldsymbol{y}_1) - \boldsymbol{y}_2\|_2^2 + \mathbb{E}_{\boldsymbol{y}, \boldsymbol{w}'} \frac{1}{\tau} (\boldsymbol{y} + \boldsymbol{y}_2)^\top h(\boldsymbol{w}'). \end{aligned} \tag{38}$$

This expression highlights the main distinction with standard GR2R. In GR2R, the paired variables $(\boldsymbol{y}_1, \boldsymbol{y}_2)$ are generated from a *known* oracle recorruption mechanism, so the loss $\mathbb{E}\|f(\boldsymbol{y}_1) - \boldsymbol{y}_2\|_2^2$ is optimized only with respect to f, and all remaining terms are fixed by construction. In the proposed setting, however, h is itself an optimization variable. As a consequence, the additional term

$$-\frac{1}{\tau} (\boldsymbol{y} + \boldsymbol{y}_2)^\top h(\boldsymbol{w}') \tag{39}$$

cannot be ignored when optimizing with respect to h, because it induces an extra bias on the learned recorruption map. This term is irrelevant for the update of f, since it does not depend on the denoiser parameters, but it does affect the update of h and may drive the recorruption mechanism toward degenerate solutions unrelated to the desired orthogonality condition.

Hence, the proposed approach can be understood as a learned extension of GR2R in which the oracle recorruption process is no longer assumed known and is instead approximated through the min–max optimization; see Appendix A.3.3. In this view, GR2R corresponds to the ideal case where the correct recorruption map is available from the outset, whereas the present method gradually guides h toward an oracle GR2R-like mechanism while simultaneously training the denoiser. The optimization encourages h to capture auxiliary statistics compatible with GR2R.

A.1.2 Relationship to UNSURE

The proposed objective can also be interpreted as a learned generalization of UNSURE. Recall the L2R saddle objective

$$\min_{f}\max_{h}\ \mathbb{E}_{\boldsymbol{y},\boldsymbol{w}'}\Big[\|f(\boldsymbol{y}+\tau h(\boldsymbol{w}'))-\boldsymbol{y}\|_2^2+\frac{2}{\tau}f(\boldsymbol{y}+\tau h(\boldsymbol{w}'))^\top h(\boldsymbol{w}')\Big]. \tag{40}$$

Since, by construction, we ensure that $\mathbb{E}_{\boldsymbol{w}'}[h(\boldsymbol{w}')]=\mathbf{0}$, the correlation term can be equivalently rewritten as

$$\frac{2}{\tau}\mathbb{E}_{\boldsymbol{y},\boldsymbol{w}'}\left[f(\boldsymbol{y}+\tau h(\boldsymbol{w}'))^\top h(\boldsymbol{w}')\right]=\frac{2}{\tau}\mathbb{E}_{\boldsymbol{y},\boldsymbol{w}'}\left[\big(f(\boldsymbol{y}+\tau h(\boldsymbol{w}'))-f(\boldsymbol{y})\big)^\top h(\boldsymbol{w}')\right], \tag{41}$$

where the subtraction of $f(\boldsymbol{y})$ is introduced only to center the perturbation term and expose a finite-difference form. This rewriting is useful because the quantity

$$\frac{f(\boldsymbol{y}+\tau h(\boldsymbol{w}'))-f(\boldsymbol{y})}{\tau} \tag{42}$$

is precisely the first-order increment of f along the perturbation direction $h(\boldsymbol{w}')$. Assuming that f is differentiable and τ is small, a first-order expansion gives

$$f(\boldsymbol{y}+\tau h(\boldsymbol{w}'))=f(\boldsymbol{y})+\tau\frac{\partial f}{\partial \boldsymbol{y}}(\boldsymbol{y})\,h(\boldsymbol{w}')+\mathcal{O}(\tau^2), \tag{43}$$

and therefore

$$\frac{2}{\tau}\mathbb{E}_{\boldsymbol{y},\boldsymbol{w}'}\left[f(\boldsymbol{y}+\tau h(\boldsymbol{w}'))^\top h(\boldsymbol{w}')\right]=2\,\mathbb{E}_{\boldsymbol{y}}\left[\mathrm{tr}\left(\mathbf{C}_h\,\frac{\partial f}{\partial \boldsymbol{y}}(\boldsymbol{y})\right)\right]+\mathcal{O}(\tau), \tag{44}$$

where

$$\mathbf{C}_h\triangleq\mathbb{E}_{\boldsymbol{w}'}\left[h(\boldsymbol{w}')h(\boldsymbol{w}')^\top\right]. \tag{45}$$

For the quadratic term, the same first-order expansion implies

$$\mathbb{E}_{\boldsymbol{y},\boldsymbol{w}'}\|f(\boldsymbol{y}+\tau h(\boldsymbol{w}'))-\boldsymbol{y}\|_2^2=\mathbb{E}_{\boldsymbol{y}}\|f(\boldsymbol{y})-\boldsymbol{y}\|_2^2+\mathcal{O}(\tau). \tag{46}$$

Hence, to first order in τ, the proposed loss becomes

$$\mathbb{E}_{\boldsymbol{y}}\|f(\boldsymbol{y})-\boldsymbol{y}\|_2^2+2\,\mathbb{E}_{\boldsymbol{y}}\left[\mathrm{tr}\left(\mathbf{C}_h\,\frac{\partial f}{\partial \boldsymbol{y}}(\boldsymbol{y})\right)\right]+\mathcal{O}(\tau), \tag{47}$$

which recovers several UNSURE-type objectives as special cases. From a recorruption-based perspective, L2R involves fewer approximations than UNSURE: it optimizes the recorrupted-input $\boldsymbol{y}_1$ objective directly, whereas the UNSURE connection emerges after a first-order reduction in which the correlation term is converted into a divergence penalty and the quadratic term is replaced by its small-τ surrogate. In addition, UNSURE requires a Monte Carlo finite-difference approximation of the divergence in practice, while L2R avoids divergence estimation altogether.

Three cases of the L2R–UNSURE connection. We consider three noise settings to show how the parameterization of the L2R recorruptor h shapes its connection to UNSURE. For isotropic and correlated Gaussian noise, linear recorruptors lead to the corresponding UNSURE divergence terms in the small-τ limit. For heavy-tailed noise, a nonlinear recorruptor retains higher-order contributions beyond that first-order limit.

1) Isotropic Gaussian noise. For the basic Gaussian denoising setting, choose a linear scaling recorruptor

$$h(\boldsymbol{w}')=\sqrt{\alpha}\,\boldsymbol{w}',\qquad \boldsymbol{w}'\sim\mathcal{N}(\mathbf{0},\mathbf{I}), \tag{48}$$

so that $\mathbf{C}_h=\alpha\mathbf{I}$. Then Eq. 47 reduces to

$$\mathbb{E}_{\boldsymbol{y}}\|f(\boldsymbol{y})-\boldsymbol{y}\|_2^2+2\alpha\,\mathbb{E}_{\boldsymbol{y}}[\mathrm{div}\,f(\boldsymbol{y})]+\mathcal{O}(\tau), \tag{49}$$

For a scalar linear-Gaussian recorruptor, this has the UNSURE form after identifying the UNSURE multiplier with the auxiliary covariance parameter, $\eta=\alpha$; the remaining discrepancy is $\mathcal{O}(\tau)$. Thus, as $\tau\to 0$, L2R recovers the UNSURE objective.

2) Correlated Gaussian noise. For correlated Gaussian noise, let the recorruptor be a linear operator applied to white Gaussian noise,

$$h(\boldsymbol{w}') = \mathbf{H}_\alpha \boldsymbol{w}', \tag{50}$$

or, equivalently in the spatially stationary case,

$$h(\boldsymbol{w}') = k_\alpha * \boldsymbol{w}', \tag{51}$$

with $*$ denoting convolution. In this case,

$$\mathbf{C}_h = \mathbb{E}_{\boldsymbol{w}'}[\mathbf{H}_\alpha \boldsymbol{w}' \boldsymbol{w}'^\top \mathbf{H}_\alpha^\top] = \mathbf{H}_\alpha \mathbf{H}_\alpha^\top \triangleq \mathbf{\Sigma}_\alpha, \tag{52}$$

and Eq. 47 becomes

$$\mathbb{E}_{\boldsymbol{y}} \|f(\boldsymbol{y}) - \boldsymbol{y}\|_2^2 + 2\,\mathbb{E}_{\boldsymbol{y}}\left[\operatorname{tr}\left(\mathbf{\Sigma}_\alpha \frac{\partial f}{\partial \boldsymbol{y}}(\boldsymbol{y})\right)\right] + \mathcal{O}(\tau). \tag{53}$$

This is the correlated-Gaussian version of UNSURE. In particular, when $\mathbf{H}_\alpha$ is chosen as a circulant convolution operator, the induced covariance $\mathbf{\Sigma}_\alpha$ coincides with the convolutional parameterization used by correlated-noise UNSURE (C-UNSURE) (Tachella et al., 2025a). Hence, restricting h to a linear convolution map recovers the same family of covariance-weighted divergence constraints.

3) Heavy-tailed noise. The previous two cases show that UNSURE is recovered when h is restricted to linear transformations of Gaussian noise. However, for heavier-tailed corruptions such as log-gamma and Laplace, the proposed formulation goes beyond the basic Gaussian and correlated-Gaussian UNSURE objectives. Indeed, the heavy-tailed UNSURE formulation still imposes a structured constraint of the form

$$\min_f \max_\eta \; \mathbb{E}_{\boldsymbol{y}} \|f(\boldsymbol{y}) - \boldsymbol{y}\|_2^2 + 2\eta \sum_{i=1}^n \mathbb{E}_{\boldsymbol{y}}\left[a(y_i)\,\frac{\partial f_i}{\partial y_i}(\boldsymbol{y})\right], \tag{54}$$

where the weighting function $a(\cdot)$ must be specified analytically. In contrast, L2R replaces this hand-crafted scalar weighting by a learned nonlinear recorruption map $h \in \mathcal{H}$ and enforces the wider family of constraints

$$\mathbb{E}_{\boldsymbol{y},\boldsymbol{w}'}\left[f(\boldsymbol{y} + \tau h(\boldsymbol{w}'))^\top h(\boldsymbol{w}')\right] = 0, \qquad \forall h \in \mathcal{H}. \tag{55}$$

The additional flexibility appears beyond the first-order UNSURE limit. In scalar notation, with $h = h(\boldsymbol{w}')$, retaining higher-order terms gives

$$\begin{aligned}\frac{2}{\tau}\mathbb{E}_{\boldsymbol{y},\boldsymbol{w}'}[f(\boldsymbol{y} + \tau h)h] = 2\mathbb{E}_{\boldsymbol{y}}[f'(\boldsymbol{y})]\mathbb{E}_{\boldsymbol{w}'}[h^2] + \tau\mathbb{E}_{\boldsymbol{y}}[f''(\boldsymbol{y})]\mathbb{E}_{\boldsymbol{w}'}[h^3] \\ + \frac{\tau^2}{3}\mathbb{E}_{\boldsymbol{y}}[f'''(\boldsymbol{y})]\mathbb{E}_{\boldsymbol{w}'}[h^4] + \mathcal{O}(\tau^3),\end{aligned} \tag{56}$$

showing that finite-τ nonlinear recorruption incorporates variance, skewness, and higher-order moments. Therefore, UNSURE is recovered under a linear first-order parameterization, whereas the learned monotonic recorruptor defines a richer family that is particularly relevant for heavy-tailed noise. L2R also avoids explicit divergence estimation and the additional denoiser evaluations used by Monte Carlo finite differences. This interpretation agrees with the empirical moment matching and the improvements over UNSURE reported for log-gamma and Laplace noise.

A.2 Implementation Details

A.2.1 Noise models and recorruptor parameterization

We evaluate performance under a diverse set of non-Gaussian noise regimes that capture statistical properties observed in real data, including heavy-tailed log-gamma, Laplace, spatially correlated Gaussian, and Poisson–Gaussian noise, in order to stress-test robustness beyond the standard i.i.d. Gaussian assumption.

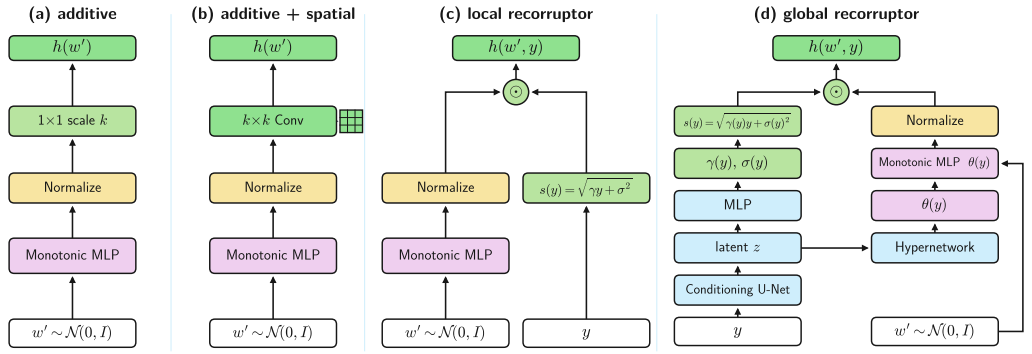

Figure 6: **Overview of the four recorruptor parameterizations.** (a) Additive i.i.d. noise; (b) additive noise with spatial correlations; (c) local recorruptor, with parameters shared across images and pixelwise scaling from $\boldsymbol{y}$; and (d) global recorruptor, which uses image-level information from $\boldsymbol{y}$ to estimate image-specific noise statistics and monotonic-map parameters.

Log-gamma noise. We model zero-mean additive log-gamma noise with fixed variance to capture different long-tailed distributions. Specifically, we parametrize the noise variable as $\boldsymbol{\varepsilon} = \tilde{\sigma}(\ln(\boldsymbol{z}) - c_\ell)$, where $\boldsymbol{z}$ follows a Gamma distribution $z_i \sim \mathcal{G}(\ell, \ell)$, with ℓ controlling the shape of the distribution. Here, c_ℓ removes the constant bias and $\tilde{\sigma}$ rescales to the target standard deviation. They can be computed as $c_\ell = \psi(\ell) - \ln(\ell)$ and $\tilde{\sigma} = \sigma/\sqrt{\psi_1(\ell)}$, where ψ and ψ_1 denote the digamma and trigamma functions, respectively. The standard deviation is fixed at $\sigma = 0.1$, and the recorruptor uses a 1×1 kernel for element-wise recorruption.

Laplace noise. In the case of additive Laplace noise, the noise variable is modeled as $\varepsilon_i \sim \text{Laplace}(0, b)$, where the location parameter is set to zero to ensure a zero mean, and the scale parameter is b. The corresponding standard deviation of the noise is given by $\sigma = b\sqrt{2}$. In this setting, the same element-wise recorruptor h is employed.

Correlated noise. In the correlated case, the noise variable is reparameterized as $\boldsymbol{\varepsilon} = \boldsymbol{z} * \boldsymbol{k}_{\text{gt}}$, where $\boldsymbol{z} \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I})$ and $\boldsymbol{k}_{\text{gt}}$ denotes the ground-truth Gaussian blur kernel. In these experiments, $\boldsymbol{k}_{\text{gt}}$ is instantiated with a spatial size 3×3 and a fixed standard deviation. The ground-truth kernel is presented in Figure 8. The recorruptor also uses a 3×3 kernel.

Poisson–Gaussian noise. Under the Poisson–Gaussian model, $\boldsymbol{y} = \gamma \boldsymbol{z}_p + \boldsymbol{z}$, with $\boldsymbol{z}_p \sim \mathcal{P}(\boldsymbol{x}/\gamma)$ and $\boldsymbol{z} \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I})$. Hence, $\mathbb{E}[\boldsymbol{y} \mid \boldsymbol{x}] = \boldsymbol{x}$ and $\text{Var}(\boldsymbol{y} \mid \boldsymbol{x}) = \gamma\boldsymbol{x} + \sigma^2$. Equivalently, $\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{\varepsilon}_{\boldsymbol{x}}$ with $\mathbb{E}[\boldsymbol{\varepsilon}_{\boldsymbol{x}} \mid \boldsymbol{x}] = \mathbf{0}$, so the GR2R moment relations extend by conditioning on $\boldsymbol{x}$. Following the conditional recorruptor defined above, we set $s(\boldsymbol{y}) = \sqrt{\gamma(\boldsymbol{y})\boldsymbol{y} + \sigma(\boldsymbol{y})^2}$, so that

$$h(\boldsymbol{w}', \boldsymbol{y}) = s(\boldsymbol{y}) \odot \tilde{h}_{\boldsymbol{y}}(\boldsymbol{w}').$$

Here, $s(\boldsymbol{y})$ sets the signal-dependent recorruption scale, and $\tilde{h}_{\boldsymbol{y}}$ follows the architecture in Eq. 19.

We consider two configurations.

Local recorruptor (Fig. 6(c)). For the synthetic Poisson–Gaussian and Cryo-EM experiments, γ, σ, and the weights of $\tilde{h}_{\boldsymbol{y}}$ are learned and shared across images. The scale $s(\boldsymbol{y}) = \sqrt{\gamma\boldsymbol{y} + \sigma^2}$ varies pixelwise with $\boldsymbol{y}$, while the learned statistics and transport remain shared.

Global recorruptor (Fig. 6(d)). For the real-world SIDD experiments, a conditioning U-Net extracts an image-level latent representation from $\boldsymbol{y}$. One part is passed through an MLP to predict image-specific parameters $\gamma(\boldsymbol{y})$ and $\sigma(\boldsymbol{y})$; the other is passed to a hypernetwork that predicts the image-specific weights of the monotonic network defining $\tilde{h}_{\boldsymbol{y}}$. An element-wise absolute-value activation constrains these weights to be nonnegative and preserves the monotonicity of $\tilde{h}_{\boldsymbol{y}}$. Thus, information from the whole image determines its noise statistics and recorruptor, while $s(\boldsymbol{y})$ provides pixelwise scaling.

Finally, the choice of recorruptor class reflects a trade-off between generality and specialization. A more expressive, general-purpose recorruptor could represent multiple noise distribution families within a single model, but it also enlarges the search space and may make learning the appropriate recorruption for each distribution more difficult. In contrast, a specialized recorruptor that incorporates known structural properties of the corruption process can reduce the learning burden and provide a more accurate parameterization. For example, prior knowledge about whether the noise is additive, signal-dependent, or spatially correlated can guide the design of the recorruptor class. Thus, the choice of recorruptor architecture acts as both a user-specified and structural prior: greater knowledge of the real noise process enables more targeted modeling, whereas limited knowledge motivates more general but potentially harder-to-optimize classes. Developing recorruptors that balance broad distributional coverage with reliable optimization is an important direction for future work.

A.3 Additional Simulations

A.3.1 Poisson–Gaussian Denoising

We evaluate the proposed L2R under two Poisson–Gaussian denoising setups: a Poisson-only case ($\gamma/\sigma = 0.1/0.0$) and a mixed Poisson–Gaussian case ($\gamma/\sigma = 0.05/0.05$). Here, γ controls the Poisson variance (gain) and σ denotes the Gaussian standard deviation; in both setups the dominant corruption is still Poisson. Comparisons are conducted against representative self-supervised baselines, including PG-SURE and PG-UNSURE (the Poisson–Gaussian variants of SURE and UNSURE, respectively), as well as GR2R and NBR2NBR.

Table 3: Denoising performance in removing Poisson–Gaussian noise.

Method	p_ε known?	Poisson–Gaussian Noise γ/σ	BSDS500	DIV2K
Supervised	no		28.44/0.833	29.45/0.845
PG-SURE	yes		27.21/<u>0.776</u>	27.94/<u>0.783</u>
PG-UNSURE	no	0.05/	<u>27.44</u>/0.774	<u>28.14</u>/0.771
NBR2NBR	no	0.05	27.12/0.751	27.64/0.744
Blind2Sound	no		27.30/0.793	27.89/0.789
L2R (Ours)	no		**27.80/0.791**	**28.70/0.797**
Supervised	no		27.16/0.792	28.27/0.815
GR2R	yes		27.03/0.786	28.15/0.812
PG-SURE	yes		25.91/0.729	26.69/0.744
PG-UNSURE	no	0.1/	25.80/0.710	26.57/0.724
NBR2NBR	no	0.0	<u>26.34</u>/<u>0.750</u>	<u>27.20</u>/<u>0.768</u>
Blind2Sound	no		26.12/0.747	26.78/0.754
L2R (Ours)	no		**26.63/0.756**	**27.68/0.779**

Figure 7 shows that L2R yields cleaner reconstructions with improved edge fidelity and fewer residual artifacts across representative BSDS500 crops. Compared to the strongest and most methodologically related baseline, PG-UNSURE, L2R produces sharper structures and reduced noise remnants.

Table 3 confirms these observations: L2R achieves the best PSNR among methods that do not assume knowledge of the noise distribution. At $\gamma/\sigma = 0.05/0.05$, L2R improves over PG-UNSURE from 27.44 to 27.80 dB on BSDS500 and from 28.14 to 28.70 dB on DIV2K. Notably, PG-UNSURE relies on a Monte Carlo divergence approximation to build its objective, whereas L2R avoids divergence estimation, removing this additional approximation while still improving PSNR. In contrast, GR2R performs the best for Poisson-only denoising ($\gamma = 0.1$, $\sigma = 0.0$), given that it has exact prior knowledge of the noise distribution.

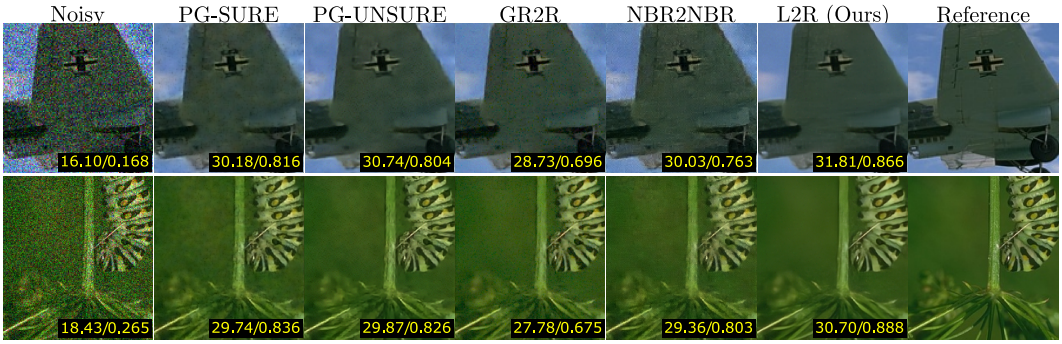

Figure 7: **Poisson–Gaussian Denoising.** Visual comparison on representative BSDS500 crops corrupted by Poisson–Gaussian Noise. Columns show the noisy input, PG-SURE, PG-UNSURE, GR2R, NBR2NBR, and the proposed L2R, followed by the ground truth.

A.3.2 MONOTONIC LEARNED RECORRUPTOR h

Table 4: Ablation of the design of model h.

network setup $h(\cdot)$	**PSNR**	**SSIM**
Supervised	31.44	0.898
Oracle (GR2R)	31.14	0.889
No learning	20.00	0.415
MLP	30.45	0.860
MLP + Residual	30.36	0.858
MLP + Id-pretrain	30.42	0.856
Monotonic MLP	30.45	0.864
Monotonic MLP + Id-pretrain	**30.77**	**0.871**

Table 4 reports an ablation on the design of h for self-supervised denoising of log-gamma noise with $\ell = 1.0$. *No learning* corresponds to evaluating the metrics directly on the noisy observation. Unconstrained MLP variants yield a large improvement and approach the GR2R oracle bound, but remain sensitive to optimization and do not explicitly enforce the intended order-preserving behavior. Adding a residual connection does not provide consistent gains. Enforcing monotonicity improves robustness, and combining it with identity initialization (*Id-pretrain*) yields the best overall self-supervised performance. In particular, *Monotonic MLP + Id-pretrain* achieves the top results (30.77 dB / 0.871), narrowing the gap to supervised training (31.44 dB / 0.898) while being competitive with GR2R in PSNR.

A.3.3 INFLUENCE OF TRAINING OPTIMIZATION SETUP

Table 5 analyzes the effect of different optimization strategies for the denoiser f and the learned recorruption map h in the correlated Gaussian denoising setting with $\sigma = 0.2$. The experiments preserve the same network, dataset, and training configuration reported in the main text. In the *joint* setting, both variables are optimized from a single loss evaluation by applying gradient descent on f and gradient ascent on h using the same forward pass. In contrast, the *non-joint* setting alternates the updates: the loss is first evaluated to update f, and then re-evaluated to update h. The results show that the two strategies perform similarly; the stop-gradient joint configuration gives the best result. This is relevant in practice, as the joint formulation is simpler and avoids repeated loss evaluations while preserving essentially the same denoising quality.

The table also compares GR2R-style and UNSURE-style formulations for the min–max objective. In particular, the GR2R-inspired versions $\|f(\boldsymbol{y}_1) - \boldsymbol{y}_2\|_2^2$ and their bias-corrected variant produce results that are close to those of the proposed UNSURE-style formulation, suggesting that both views are consistent in practice and that the empirical gap between them is minimal in this setting. However, the scaling of the correlation term is important: replacing the factor $\frac{2}{\tau}$ with an unscaled inner product $h(\boldsymbol{w}')^\top f(\boldsymbol{y}_1)$ consistently degrades performance, both in the joint and non-joint settings.

Finally, in the practical implementation we use a stop-gradient operation on $\boldsymbol{y}_1$ when updating the recorruptor, i.e., $f(\mathrm{sg}[\boldsymbol{y}_1])$. This yields the best overall results and prevents gradients from propagating through the denoiser with respect to the dependence of $\boldsymbol{y}_1$ on h. Consequently, the update of h is driven directly by the correlation term, providing an efficient surrogate ascent step while avoiding an additional backward pass through f and reducing computational overhead.

Table 5: **Ablation f-loss and h-loss.** Correlated denoising performance for different types of f-loss and h-loss with $\sigma = 0.2$. $\mathrm{sg}[\cdot]$ denotes the stop-gradient function.

Joint?	f-loss (min)	h-loss (max)	PSNR	SSIM
no	$\|f(\boldsymbol{y}_1) - \boldsymbol{y}_2\|_2^2$	$h(\boldsymbol{w}')^\top f(\boldsymbol{y}_1)$	29.72	0.854
no	$\|f(\boldsymbol{y}_1) - \boldsymbol{y}_2\|_2^2 + \frac{1}{\tau} \cdot (\boldsymbol{y}_2 + \boldsymbol{y})^\top h(\boldsymbol{w}')$		29.74	0.854
no	$\|f(\boldsymbol{y}_1) - \boldsymbol{y}\|_2^2 + 2 \cdot h(\boldsymbol{w}')^\top f(\boldsymbol{y}_1)$		29.38	0.841
no	$\|f(\boldsymbol{y}_1) - \boldsymbol{y}\|_2^2 + \frac{2}{\tau} \cdot h(\boldsymbol{w}')^\top f(\boldsymbol{y}_1)$		29.83	0.859
yes	$\|f(\boldsymbol{y}_1) - \boldsymbol{y}\|_2^2 + \frac{2}{\tau} \cdot h(\boldsymbol{w}')^\top f(\boldsymbol{y}_1)$		29.78	0.860
yes	$\|f(\mathrm{sg}[\boldsymbol{y}_1]) - \boldsymbol{y}\|_2^2 + \frac{2}{\tau} \cdot h(\boldsymbol{w}')^\top f(\mathrm{sg}[\boldsymbol{y}_1])$		29.85	0.863
yes	$\|f(\mathrm{sg}[\boldsymbol{y}_1]) - \boldsymbol{y}\|_2^2 + h(\boldsymbol{w}')^\top f(\mathrm{sg}[\boldsymbol{y}_1])$		29.25	0.847
yes	$\|f(\boldsymbol{y}_1) - \boldsymbol{y}\|_2^2 + h(\boldsymbol{w}')^\top f(\boldsymbol{y}_1)$		28.85	0.825

A.3.4 EXPRESSIVITY–ROBUSTNESS TRADE-OFF

This section analyzes the expressivity–robustness trade-off of the learned recorruptor h through a denoising evaluation on log-gamma noise with $\ell = 1.0$. In particular, the goal is to assess how the architectural capacity of h, controlled by its depth and number of neurons per layer, affects the final reconstruction performance. Since h is trained adversarially to generate informative recorruptions, this experiment helps identify the model size that best balances representational flexibility and optimization robustness.

Table 6: Performance for different model sizes of h on log-gamma noise.

Depth	#Neurons	PSNR	SSIM
2	16	30.57	0.871
2	8	30.60	0.874
3	8	30.56	0.870
4	8	30.53	0.871
2	4	30.50	0.866
3	4	30.60	0.872
4	4	<u>30.64</u>	<u>0.875</u>
5	4	**30.79**	**0.879**
6	4	30.57	0.870

Table 6 shows how the design of h affects performance. The most relevant trend appears for the narrow configuration with 4 neurons per layer: increasing the depth from 2 to 5 improves the PSNR from 30.50 dB to 30.79 dB, with depth 4 already reaching 30.64 dB. This indicates that additional depth is beneficial because it increases the nonlinear expressivity of the recorruptor while keeping the number of parameters per layer controlled. In this case, depth allows h to better approximate the required recorruption mechanism without making each layer excessively flexible.

At the same time, the results show that increasing the width alone does not necessarily improve performance. For example, shallow but wider models such as depth 2 with 8 and 16 neurons achieve 30.60 dB and 30.57 dB, respectively, which remain below the deeper narrow model with depth 5 and width 4. Moreover, increasing depth too aggressively is also detrimental, since depth 6 with 4 neurons drops back to 30.57 dB.

Table 7: **Influence of the architecture design h.** The table reports the ablation results for different kernel sizes and $s(\boldsymbol{y})$-scaling configurations, evaluated under three noise models: element-wise log-gamma noise, spatially correlated noise, and signal-dependent Poisson–Gaussian noise.

Noise Model	k-size	$s(\boldsymbol{y})$	PSNR	SSIM
Log-gamma	1×1	no	**30.77**	**0.871**
$\ell = 1.0$	3×3	no	<u>29.93</u>	<u>0.837</u>
	1×1	yes	26.98	0.721
	3×3	yes	24.02	0.657
Correlated	1×1	no	24.21	0.660
$\sigma = 0.2$	3×3	no	**29.85**	**0.863**
	1×1	yes	24.02	0.657
	3×3	yes	<u>27.05</u>	<u>0.754</u>
Poisson–Gaussian	1×1	no	24.99	0.628
$\gamma = 0.05, \sigma = 0.05$	3×3	no	23.77	0.581
	1×1	yes	**27.80**	**0.791**
	3×3	yes	<u>26.03</u>	<u>0.713</u>

A.3.5 INFLUENCE OF THE ARCHITECTURE DESIGN h

Table 7 shows how the design of the recorruption mapping h affects performance across noise models, indicating that the best architecture depends on the corruption structure. For element-wise log-gamma noise, the pointwise 1×1 design without $s(\boldsymbol{y})$ scaling achieves the best performance, reaching 30.77 dB PSNR and 0.871 SSIM, outperforming both the 3×3 variant and the scaled 1×1 version. A similar tendency is observed for Poisson–Gaussian noise, where the 1×1 architecture combined with $s(\boldsymbol{y})$-scaling attains the best results (27.80 dB, 0.791 SSIM), while removing the scaling leads to a noticeable degradation to 24.99 dB and 0.628 SSIM.

In contrast, the correlated noise setting benefits from incorporating the spatial convolution operator in h. In this case, the 3×3 configuration without $s(\boldsymbol{y})$ scaling reaches 29.85 dB and 0.863 SSIM, whereas the scaled 3×3 version degrades to 27.05 dB and 0.754 SSIM, and the pointwise 1×1 alternative further drops to 24.21 dB and 0.660 SSIM. This is expected, as employing a non-correlated recorruption scheme can induce substantial errors when the underlying noise exhibits spatial correlations. Overall, the ablation indicates that the design of h should remain aligned with the main characteristics of the corruption process: element-wise mappings appear adequate for independent noise, spatial kernels become beneficial for correlated perturbations, and $s(\boldsymbol{y})$ scaling is helpful in signal-dependent Poisson–Gaussian regimes.

A.3.6 TRAINING STABILITY

We further evaluate the sensitivity of L2R to its main optimization parameters by varying one parameter at a time while holding the other settings fixed. The random seed is fixed within each experiment.

Table 8 reports an additional sensitivity analysis of the recorruption parameter τ under heavy-tailed log-gamma noise. This experiment uses a separate random seed and 1,000 training epochs. L2R exhibits a broad effective operating region for moderate perturbation strengths, with the strongest performance observed around $\tau \in [0.3, 0.5]$. In particular, $\tau = 0.5$ attains the highest PSNR of 30.34 dB while essentially matching the best SSIM. Performance remains close throughout $\tau \in [0.3, 0.7]$, whereas more extreme values lead to larger degradation. This indicates that τ controls a meaningful recorruption trade-off.

Table 8: Sensitivity to τ on log-gamma noise $\ell = 1.0$

τ	0.1	0.3	0.4	0.5	0.6	0.7	0.9
PSNR	28.83	<u>30.21</u>	29.95	**30.34**	29.78	29.59	29.26
SSIM	0.821	**0.869**	0.865	<u>0.868</u>	0.860	0.854	0.843

We additionally report a separate SIDD experiment for sensitivity analysis, using independent training runs; these values are not the principal benchmark result in Table 9. Consistent with the scenario-dependent role of τ, the favorable operating region on SIDD is centered around $\tau = 1$. Performance remains stable in its neighborhood, varying by only 0.28 dB for $\tau \in [0.8, 1.25]$, while more extreme perturbation strengths lead to moderate degradation.

L2R is similarly robust to the relative learning rate between the denoiser and recorruptor. Although $R_{lr} = 0.5$ achieves the highest PSNR of 51.32 dB, values in the broad range $R_{lr} \in [0.3, 0.9]$ remain within 0.23 dB of one another, with $R_{lr} = 0.9$ reaching 51.30 dB. Finally, the method exhibits very low sensitivity to stochastic initialization: five independent random seeds yield 51.25 ± 0.06 dB, with all runs contained within a 0.17 dB interval. Overall, these results indicate that the min–max optimization is reproducible across initializations and does not depend on narrowly tuned optimization parameters.

Table 9: Training stability on SIDD under variations of the main optimization parameters.

Parameter	Values / PSNR (dB)					avg PSNR (mean ± std.)
Recorruption (τ)	0.6	0.8	**1.0**	1.25	1.67	51.04 ± 0.24
	50.68	51.04	**51.32**	51.20	50.98	
Learning Ratio (R_{lr})	0.1	0.3	**0.5**	0.7	0.9	51.12 ± 0.22
	50.76	51.09	**51.32**	51.14	51.30	
RNG Seed	0	1	2	3	4	51.25 ± 0.06
	51.29	51.15	**51.32**	51.27	51.25	

A.4 Learned recorruption maps and spatial kernel

Figure 8 compares the learned recorruption maps and spatial kernel with their oracle counterparts. For the log-gamma and Laplace settings (left and center), the recorruptor h is compared with (i) an *oracle* recorruption obtained by manually matching low-order moments of the target distribution, and (ii) the identity mapping, representing Gaussian noise. For correlated noise (right), the true 3×3 Gaussian blur kernel inducing spatial correlation is compared with the kernel recovered by the proposed min–max training.

For log-gamma noise, the learned mapping closely follows the oracle curve over most of the support, exhibiting the expected nonlinearity required to reproduce the heavy-tailed behavior. For Laplace noise, the learned mapping remains a reasonable approximation but shows larger deviations from the oracle, particularly in the tails. The learned map is optimized to capture auxiliary statistics compatible with the GR2R objective.

In the correlated-noise setting, each coefficient of the 3×3 kernel is treated as a free parameter within h, and the learned kernel approaches the overall structure of the true Gaussian blur kernel while not perfectly reproducing its symmetry. Such symmetry can, in principle, be imposed via a structured parameterization; however, empirical evidence indicates that this is not necessary to achieve satisfactory denoising performance. Overall, the same min–max objective adapts the recorruptor through its monotonic transport and learnable spatial kernel. The learned h provides an interpretable proxy for the effective auxiliary noise statistics.

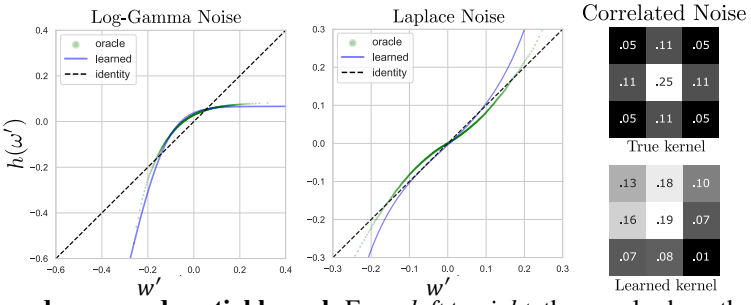

Figure 8: **Learned maps and spatial kernel.** From *left* to *right*, the panels show the learned recorruption maps for log-gamma and Laplace noise and the learned kernel for correlated noise.

A.5 Implicit noise-correlation tracking

To assess the learned recorruption and the effect of the maximization in Eq. 18, an additional evaluation is conducted on the log-gamma denoising setting with $\ell = 1.0$. During training, the following quantities are monitored:

$$\mathcal{C}_\varepsilon \triangleq \mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'}\left[f(\boldsymbol{y}+\tau h(\boldsymbol{w}'))^\top g(\boldsymbol{w})\right],\ \mathcal{C}_h \triangleq \mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'}\left[\frac{1}{\tau} f(\boldsymbol{y}+\tau h(\boldsymbol{w}'))^\top h(\boldsymbol{w}')\right], \tag{57}$$

$$\mathcal{C}_\Delta \triangleq \left|\mathbb{E}_{\boldsymbol{x},\boldsymbol{w},\boldsymbol{w}'}\left[f(\boldsymbol{y}+\tau h(\boldsymbol{w}'))^\top g(\boldsymbol{w}) - \frac{1}{\tau} f(\boldsymbol{y}+\tau h(\boldsymbol{w}'))^\top h(\boldsymbol{w}')\right]\right|. \tag{58}$$

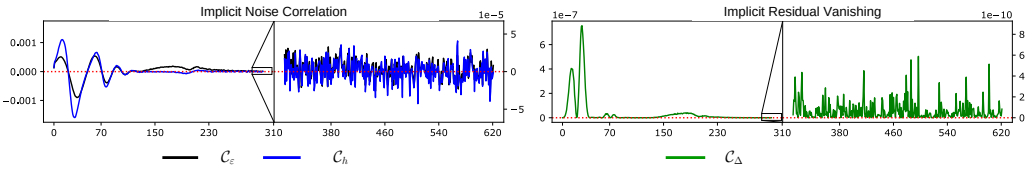

Figure 9: **Training dynamics of the implicit correlation terms.** *Left:* noise correlation $\mathcal{C}_\varepsilon$ (**black**) and auxiliary correlation $\mathcal{C}_h$ (**blue**) versus epoch. *Right:* residual gap $\mathcal{C}_\Delta$ (**green**). Both inner-product terms settle near zero. The residual gap decreases and stabilizes near zero, supporting the empirical alignment described in Eq. 21.

Here, $\mathcal{C}_\varepsilon$ is the *noise-correlation term*, $\mathcal{C}_h$ is the *auxiliary correlation term*, and $\mathcal{C}_\Delta$ measures the *residual gap* between both terms.

Figure 9 reports these statistics across epochs. Two regimes are observed. Early in training, both $\mathcal{C}_\varepsilon$ and $\mathcal{C}_h$ exhibit oscillatory behavior as the min–max optimization evolves. The auxiliary correlation term (blue) settles near zero as training progresses, consistent with the intended effect of the saddle formulation, i.e., enforcing $\mathbb{E}\big[\frac{1}{\tau} f(\boldsymbol{y}_1)^\top h(\boldsymbol{w}')\big] \approx 0$ without access to clean targets.

The noise-correlation term (black) also settles near zero and closely tracks $\mathcal{C}_h$. The residual gap $\mathcal{C}_\Delta$ (green) decreases and stabilizes near zero, supporting the empirical alignment described in Eq. 21.