Adaptive Extensions of Unbiased Risk Estimators for Unsupervised Magnetic Resonance Image Denoising

Reeshad Khan[1], John Gauch[1], and Ukash Nakarmi[1]

Department of Computer Science and Computer Engineering,
University of Arkansas, Fayetteville, AR, USA
rk010@uark.edu, jgauch@uark.edu, unakarmi@uark.edu

Abstract. Deep neural networks (DNNs) have significantly advanced image denoising, particularly in challenging noise conditions encountered in medical imaging. However, both traditional and many learning-based approaches rely on high-quality, noiseless ground truth images, which are difficult to obtain in practice for modalities such as magnetic resonance imaging (MRI). This limitation has motivated unsupervised learning strategies based on unbiased risk estimation. In this work, we study estimator-driven denoising for MRI and benchmark Stein's Unbiased Risk Estimator (SURE), its paired-observation extension (eSURE), and an extended Poisson unbiased risk formulation (ePURE) within a unified experimental framework. We evaluate these methods on fully sampled 3T knee MRI data under synthetic Gaussian and Poisson noise, reflecting common additive and signal-dependent corruption models in medical imaging. Our results show that eSURE consistently improves over single-image SURE training and achieves performance comparable to Noise2Noise and supervised baselines, with gains of up to 2.4 dB in PSNR across axial and coronal views. Under Poisson noise, the proposed ePURE formulation attains PSNR values up to 32.45 dB while preserving structural fidelity. These findings demonstrate that unbiased risk estimation provides a practical and robust foundation for unsupervised MRI denoising when clean reference data is unavailable.

Keywords: MRI denoising, unsupervised learning, Stein's unbiased risk estimator, Noise2Noise, Poisson noise

1 Introduction

Magnetic Resonance Imaging (MRI) is a mainstay of clinical diagnosis because it provides rich soft-tissue contrast without ionizing radiation [1]. In practice, however, MRI acquisitions are inevitably corrupted by noise arising from the measurement process and the subsequent reconstruction pipeline, which can blur fine anatomical boundaries and reduce the reliability of downstream tasks such as segmentation, quantitative mapping, and longitudinal comparison [2, 3]. Classical denoisers, notably BM3D[8], remain strong baselines in many imaging

conditions [8, 9]. Yet, over the past decade, deep neural network (DNN) denoisers have repeatedly demonstrated superior restoration fidelity and a favorable accuracy-efficiency trade-off when trained appropriately for the noise statistics of interest [10].

The principal obstacle for supervised denoising in medical imaging is the absence of truly clean reference images. Collecting noiseless targets is rarely feasible because scan-time constraints, patient motion, and hardware limitations impose hard bounds on averaging and repeated acquisition. Consequently, "clean" targets are typically approximated using proxy procedures (e.g., averaging or vendor-specific processing), which may retain residual noise or introduce subtle bias [20]. These constraints have driven substantial interest in unsupervised and self-supervised learning paradigms that can train denoisers directly from noisy measurements. Classical noise modeling and restoration principles remain influential in medical imaging, and non-local methods such as Non-Local Means (NLM) and its variants have long served as important baselines for MRI denoising [5–7]. In the deep learning era, several strategies have emerged to relax the need for clean targets. Deep Image Prior (DIP) exploits the implicit bias of convolutional architectures to fit a single noisy image [11], while Noise2Noise (N2N) trains using pairs of noisy observations of the same underlying anatomy, eliminating the need for noiseless references [12]. Related masking-based approaches such as Noise2Void and Noise2Self further reduce data requirements by learning from single noisy images under conditional independence assumptions [13, 14]. Earlier work has also explored training deep denoisers without ground truth supervision under carefully specified statistical conditions [15].

Stein's Unbiased Risk Estimator (SURE) yields an unbiased estimate of MSE under additive i.i.d. Gaussian noise [16, 17], and Monte-Carlo SURE (MC-SURE) makes this principle practical for modern deep networks via stochastic approximation of the divergence term [18, 19]. When two noisy observations are available-even if they are correlated-extended SURE (eSURE) connects SURE-type objectives to N2N-style training under Gaussian assumptions [20]. For count-driven regimes where Poisson noise is more appropriate, Poisson unbiased risk estimators (PURE) and their extensions provide analogous training objectives under Poisson statistics [22–25].

SURE, eSURE, and PURE are not introduced by this work. The contribution of this paper is a careful adaptation and validation of these estimator-driven learning strategies in a realistic MRI denoising workflow, with an emphasis on clarifying when each estimator is preferable under Gaussian versus Poisson corruption. Concretely, we (i) implement SURE/MC-SURE and eSURE training for Gaussian MRI denoising, and (ii) operationalize an ePURE-style objective for Poisson-corrupted MRI using correlated noisy pairs, formed through the averaged observation $z = \frac{y_1+y_2}{2}$, to improve training stability when clean references are unavailable. All methods are evaluated on fully sampled 3T knee MRI data [28] using DnCNN as a common denoising backbone [10].

The remainder of the paper is organized as follows.

Contributions. While SURE, eSURE, and PURE have been previously introduced, this work makes the following concrete contributions in the context of MRI denoising:

- We provide a unified empirical comparison of SURE, eSURE, Noise2Noise, and supervised training under matched experimental conditions on fully sampled 3T knee MRI.
- We validate the practical benefit of paired-observation eSURE training for MRI, showing consistent improvements over single-image MC-SURE across views and noise levels.
- We operationalize a stabilized ePURE-style objective for Poisson-corrupted MRI using correlated noisy pairs, demonstrating stable training and high-fidelity reconstructions without clean references.
- We analyze estimator behavior under both additive (Gaussian) and signal-dependent (Poisson) noise to clarify when each unbiased risk formulation is preferable in practice.

Research Questions. This study is guided by the following research questions:

- **RQ1:** Does paired-observation eSURE training provide measurable advantages over single-image MC-SURE for MRI denoising under Gaussian noise?
- **RQ2:** How does eSURE compare empirically to Noise2Noise and supervised MSE training when applied to realistic MRI data?
- **RQ3:** Can a stabilized Poisson unbiased risk formulation enable effective unsupervised MRI denoising under signal-dependent noise?

Section 2 reviews SURE, MC-SURE, and related unbiased-risk estimators. Section 3 presents the estimator-driven training formulations used in our study for Gaussian and Poisson noise, including eSURE and an ePURE-style objective. Section 4 reports experimental results on 3T knee MRI, including quantitative PSNR comparisons and qualitative analyses. Section 5 presents Poisson unbiased risk estimation for MRI denoising, and Section 6 concludes with practical takeaways, limitations, and directions for future work. Throughout this paper, we focus on controlled Gaussian and Poisson corruptions to isolate estimator behavior under well-defined noise statistics; extension to Rician and mixed noise models that more directly reflect magnitude MRI is left for future work [4].

2 Background

We briefly summarize the theoretical foundations of unbiased risk estimation for image denoising, focusing on Stein's Unbiased Risk Estimator (SURE) and its Monte-Carlo approximation.

2.1 Stein's Unbiased Risk Estimator (SURE)

In the context of denoising Gaussian-contaminated signals or images, a standard observation model is

$$y = x + n, \tag{1}$$

where $x \in \mathbb{R}^N$ denotes an unknown deterministic signal or image, $y \in \mathbb{R}^N$ is the observed measurement, and $n \in \mathbb{R}^N$ represents additive i.i.d. Gaussian noise with distribution $n \sim \mathcal{N}(0, \sigma^2 I)$, where I is the identity matrix.

Stein's Unbiased Risk Estimator (SURE) yields an unbiased estimate of the mean squared error (MSE) associated with an estimator $h(y)$ of x. The SURE risk is defined as

$$\eta(h(y)) = \frac{\|y - h(y)\|^2}{N} - \sigma^2 + \frac{2\sigma^2}{N} \sum_{i=1}^{N} \frac{\partial h_i(y)}{\partial y_i}, \tag{2}$$

where the final term corresponds to the divergence of the estimator with respect to its input. As established by Stein [16] and later adapted to image denoising settings [17], the estimator in Eq. (2) satisfies

$$\mathbb{E}_{n \sim \mathcal{N}(0,\sigma^2 I)} \left\{ \frac{\|x - h(y)\|^2}{N} \right\} = \mathbb{E}_{n \sim \mathcal{N}(0,\sigma^2 I)} \left\{ \eta(h(y)) \right\}, \tag{3}$$

Equation (3) formalizes the unbiasedness property: the expected MSE equals the expected SURE risk under i.i.d. Gaussian noise. While Eq. (2) is attractive for parameter optimization, obtaining a closed-form expression for the divergence term is feasible only for a limited class of estimators, such as linear filters or non-local means [26, 27]. For more general estimators, including modern deep neural networks, direct evaluation of this term becomes intractable, motivating approximate strategies.

2.2 Monte-Carlo SURE (MC-SURE)

Monte-Carlo SURE (MC-SURE), proposed by Ramani et al. [18], offers a practical stochastic approximation of the divergence term in Eq. (2). Let $\tilde{\mathbf{b}} \sim \mathcal{N}(0, I) \in \mathbb{R}^N$ be an auxiliary Gaussian vector independent of both n and y. Ramani et al. [18] showed that the divergence term can be expressed as

$$\sum_{i=1}^{N} \frac{\partial h_i(y)}{\partial y_i} = \lim_{\epsilon \to 0} \mathbb{E}_{\tilde{\mathbf{b}}} \left\{ \tilde{\mathbf{b}}^T \left(\frac{h(y + \epsilon \tilde{\mathbf{b}}) - h(y)}{\epsilon} \right) \right\}. \tag{4}$$

Applying Eq. (4) to the divergence term in Eq. (2) yields the Monte-Carlo approximation

$$\frac{1}{N} \sum_{i=1}^{N} \frac{\partial h_i(y)}{\partial y_i} \approx \frac{1}{\epsilon N} \tilde{\mathbf{b}}^T \left(h(y + \epsilon \tilde{\mathbf{b}}) - h(y) \right), \tag{5}$$

where ϵ is a small positive scalar controlling the finite-difference approximation. This formulation enables the use of SURE-based objectives for estimators whose analytical divergence is unavailable.

2.3 Optimization of Deep Denoisers via SURE-Derived Losses

Building on the MC-SURE approximation in Eq. (5), SURE can be used as a surrogate loss for the unsupervised optimization of deep neural network (DNN)-based denoisers [19]. In this setting, only noisy observations are required during training, eliminating the dependence on clean ground truth data inherent to supervised approaches. Using a batch of M noisy images $\{\mathbf{y}^{(j)}\}_{j=1}^{M}$ and a denoiser parameterized by θ, the MC-SURE training objective can be written as

$$\mathcal{L}_{\mathrm{SURE}}(\theta) = \frac{1}{M}\sum_{j=1}^{M}\Bigg[\|\mathbf{y}^{(j)} - h_\theta(\mathbf{y}^{(j)})\|^2 - N\sigma^2 + \frac{2\sigma^2}{\epsilon}(\tilde{\mathbf{b}}^{(j)})^T\left(h_\theta(\mathbf{y}^{(j)} + \epsilon\tilde{\mathbf{b}}^{(j)}) - h_\theta(\mathbf{y}^{(j)})\right)\Bigg], \tag{6}$$

where N denotes the dimensionality of the data and $\tilde{\mathbf{b}}^{(j)}$ is an independent Gaussian perturbation for each training sample. The loss in Eq. (6) combines a data fidelity term, a noise-variance correction, and a Monte-Carlo estimate of the divergence, yielding a tractable surrogate for the MSE.

These estimators were originally developed and evaluated primarily in generic imaging settings, but they provide a flexible framework that can be adapted to domain-specific applications. In the remainder of this paper, we build upon these foundations to instantiate SURE-, eSURE-, and Poisson-based unbiased risk estimators for MRI denoising, and to examine their relative behavior under Gaussian and Poisson noise in realistic medical imaging scenarios.

3 Methods

This section presents the estimator-driven training formulations used in our study and clarifies their relationship to commonly used self-supervised denoising objectives. We first revisit the Noise2Noise (N2N) objective [12] to highlight the statistical condition under which training on pairs of noisy images is equivalent to supervised MSE training. We then summarize the SURE-derived training objectives and introduce the extended SURE (eSURE) setting for two noisy observations. These formulations are used directly in our Gaussian-noise experiments, while the Poisson-risk counterpart (PURE/ePURE) is presented in Section 5.

3.1 Revisiting Noise2Noise

Noise2Noise trains a denoiser using two noisy observations of the same underlying image, avoiding the need for clean targets [12]. Let x denote the unknown clean image and let y and z be two noisy measurements generated from the same x. The denoiser is represented by $h_\theta(\cdot)$ with parameters θ. The supervised objective minimizes the expected MSE $\mathbb{E}\{\|x - h_\theta(y)\|^2\}$, but x is unavailable in our setting.

To connect N2N to the supervised objective, we expand the MSE by adding and subtracting z:

$$\mathbb{E}_{(x,y)}\left\{\|x - h_\theta(y)\|^2\right\} = \mathbb{E}_x\left[\mathbb{E}_{(y,z)\,|\,x}\left\{\|x - z + z - h_\theta(y)\|^2 \,\middle|\, x\right\}\right]. \tag{7}$$

Expanding the square yields

$$\mathbb{E}_{(x,y)}\left\{\|x - h_\theta(y)\|^2\right\} = \mathbb{E}_x\left[\mathbb{E}_{(y,z)\,|\,x}\left\{\|z - h_\theta(y)\|^2 + 2(z-x)^T h_\theta(y) \,\middle|\, x\right\}\right] + \text{const.} \tag{8}$$

The key condition for equivalence is that the cross-term vanishes, i.e., $\mathbb{E}\{(z-x)^T h_\theta(y)\} = 0$. This holds when the residual noise $(z - x)$ has zero mean and is uncorrelated with $h_\theta(y)$ given x (for example, under independent noise realizations with zero mean under standard assumptions). Under this condition, minimizing the supervised MSE is equivalent (up to an additive constant) to minimizing the N2N objective

$$\mathbb{E}_{(x,y,z)}\left\{\|z - h_\theta(y)\|^2\right\}. \tag{9}$$

This perspective is useful for our study because it clarifies why, when two suitable noisy realizations are available, N2N can match supervised performance even without clean targets.

Finally, Eq. (8) also highlights a practical caveat. If the "target" image is itself a low-quality proxy for x (e.g., mildly denoised, averaged, or otherwise imperfect), then $(z-x)$ may not behave like a zero-mean term independent of the network input, and the cross-term may become non-negligible. This motivates risk-estimator-based objectives that do not rely on clean or proxy targets.

3.2 Extended SURE and Monte-Carlo SURE

SURE allows an unbiased estimate of the MSE under additive i.i.d. Gaussian noise using only a single noisy observation. However, in many imaging pipelines one can form *two* observations of the same anatomy (e.g., by repeated acquisition, split measurements, or controlled noise injection). eSURE extends SURE to exploit such paired observations, and it also provides a principled bridge between N2N-type training and SURE-derived losses [20].

Concretely, let $y_1 \sim \mathcal{N}(x, \sigma_{y_1}^2 I)$ denote a noisy observation of x, and let $z \sim \mathcal{N}(0, \sigma_z^2 I)$ be an additional Gaussian perturbation independent of y_1. We form a second observation $y_2 = y_1 + z$, so that $y_2 \sim \mathcal{N}(x, (\sigma_{y_1}^2 + \sigma_z^2) I)$. The eSURE identity states that the MSE of $h_\theta(y_2)$ can be estimated without x via

$$\mathbb{E}_{y_2}\left\{\frac{1}{N}\|x - h_\theta(y_2)\|^2\right\} = \mathbb{E}_{y_2}\left\{\gamma(h_\theta(y_2), y_1)\right\}, \tag{10}$$

where the eSURE risk $\gamma(\cdot)$ is given by

$$\gamma(h_\theta(y_2), y_1) = \frac{1}{N}\|y_1 - h_\theta(y_2)\|^2 - \sigma_{y_1}^2 + \frac{2\sigma_{y_1}^2}{N}\sum_{i=1}^{N}\frac{\partial h_i(y_2)}{\partial (y_2)_i}. \tag{11}$$

Equations (10)-(11) provide the training objective we use for Gaussian denoising when paired noisy observations are available.

As in standard SURE, the divergence term in Eq. (11) is generally intractable for deep networks. We therefore adopt the Monte-Carlo approximation (MC-SURE) [18, 19]. Let $\tilde{b} \sim \mathcal{N}(0, I)$ be an auxiliary Gaussian vector independent of y_2. The divergence can be estimated by the stochastic finite-difference identity

$$\mathbb{E}_{\tilde{b}} \left\{ \tilde{b}^T \left(\frac{h_\theta(y_2 + \epsilon\tilde{b}) - h_\theta(y_2)}{\epsilon} \right) \right\}, \tag{12}$$

where $\epsilon > 0$ controls the finite-difference approximation. In practice, we use a small ϵ and a single Monte-Carlo sample per training example, which yields an efficient and stable estimator for training.

4 Experiments and Results

This section evaluates estimator-driven training for MRI denoising under controlled Gaussian corruption and reports both quantitative and qualitative outcomes.

4.1 Dataset and Preprocessing

Experiments were conducted on fully sampled 3T knee MRI scans from 22 subjects (11 male, 11 female) [28]. The dataset facilitates high-resolution volumes of size $320 \times 320 \times 256$, acquired on a 3T whole-body scanner using a sagittal 3D FSE CUBE sequence with proton-density weighting and fat saturation. Each knee was positioned in an 8-channel HD knee coil and aligned from anterior to posterior with a tolerance of $\pm 10^\circ$ relative to the isocenter. We slice the reconstructed volumes into 320×256 planes and form two-view datasets (axial and coronal) for training and evaluation. For each view, 1000 images are used for training and 100 images are held out for testing.

To study estimator-driven training under known corruption levels, additive white Gaussian noise is synthetically applied to the clean reconstructions at test time. For paired-observation objectives, noisy pairs are generated from the same underlying slice according to the constructions described in Section 3; in particular, the conditions underlying the Noise2Noise equivalence in Eqs. (7)-(9) and the eSURE objective in Eqs. (10)-(11) guide our pairing protocol.

4.2 Experimental Setup

We consider two Gaussian noise regimes to evaluate robustness as corruption increases: (i) $\sigma = 25$, representing moderate corruption, and (ii) $\sigma = 50$, representing a substantially more challenging setting. We compare five approaches: BM3D [8], DnCNN trained with MC-SURE [18, 19], DnCNN trained with Noise2Noise

[12], DnCNN trained with eSURE [20], and a supervised DnCNN trained with MSE using clean targets (included as an upper-bound reference).

DnCNN is used as the denoising backbone [10]. The network consists of 20 convolutional layers with batch normalization and ReLU activations. Training uses an initial learning rate of 10^{-3}, reduced to 10^{-4} after 40 epochs, with a batch size of 128. All experiments are implemented in PyTorch and run on an NVIDIA GeForce RTX 3090 (24 GB). Performance is measured using peak signal-to-noise ratio (PSNR) and structural similarity (SSIM). The SURE-related objectives follow the formulations in Section 2: the single-observation risk is estimated using MC-SURE, using the finite-difference approximation in Eq. (5), and the paired-observation objective uses eSURE with a Monte-Carlo approximation of the divergence (Eq. (12)).

4.3 Quantitative Results

The following results address **RQ1** and **RQ2** by comparing single-image SURE, paired-observation eSURE, Noise2Noise, and supervised training under matched Gaussian noise conditions. Before reporting results, we clarify nomenclature in Table 1. *DnCNN-SURE* denotes DnCNN trained with the Monte-Carlo SURE objective in Eq. (6). *DnCNN-SURE** is a simplified baseline that omits the Monte-Carlo divergence correction term (the last term in Eq. (6)), i.e., it trains using only the data-fidelity component $\|y - h_\theta(y)\|^2$ together with the noise-variance offset. This variant is included to isolate the empirical contribution of the divergence term to optimization stability and final reconstruction quality. Table 1 reports PSNR for axial and coronal views across the two noise regimes. Several consistent patterns emerge. First, eSURE improves over single-observation SURE training in all settings, indicating that exploiting paired observations offers a stronger and more stable learning signal than relying on a single noisy measurement. Second, eSURE tracks Noise2Noise closely. In the moderate-noise setting ($\sigma = 25$), the two methods are essentially tied in axial views (33.96 vs. 33.98 dB) and differ only marginally in coronal views (32.46 vs. 32.56 dB). In the stronger-noise setting ($\sigma = 50$), eSURE yields the best PSNR among the unsupervised/self-supervised objectives in both views, with gains over N2N of 0.10 dB (axial) and 1.13 dB (coronal). These trends are consistent with the theoretical relationship between N2N and eSURE discussed in Section 3: when the cross-term in Eq. (8) is suppressed, N2N approximates supervised MSE training, and eSURE offers an estimator-driven alternative that remains effective even when the paired observations are not strictly independent. DnCNN-MSE denotes supervised training using fully sampled reconstructions as clean targets, evaluated under matched noise levels and normalized to the same intensity range as all unsupervised methods.

4.4 Qualitative Results

Quantitative metrics alone do not fully capture clinical utility, particularly when diagnostically relevant cues are subtle. Figures 1 and 2 therefore provide rep-

Table 1. PSNR (dB) on 3T knee MRI for axial and coronal views. Best results per row are in bold.

3T Knee MRI (Axial View)						
Methods	BM3D[8]	DnCNN-SURE	DnCNN-SURE*	DnCNN-N2N	DnCNN-eSURE	DnCNN-MSE
$\sigma = 25$	29.18	31.56	28.62	33.96	33.98	**34.12**
$\sigma = 50$	27.82	31.55	26.05	31.53	31.63	**31.88**
3T Knee MRI (Coronal View)						
Methods	BM3D[8]	DnCNN-SURE	DnCNN-SURE*	DnCNN-N2N	DnCNN-eSURE	DnCNN-MSE
$\sigma = 25$	28.96	32.55	28.41	32.46	32.56	**32.82**
$\sigma = 50$	27.61	28.75	25.72	28.86	29.99	**30.21**

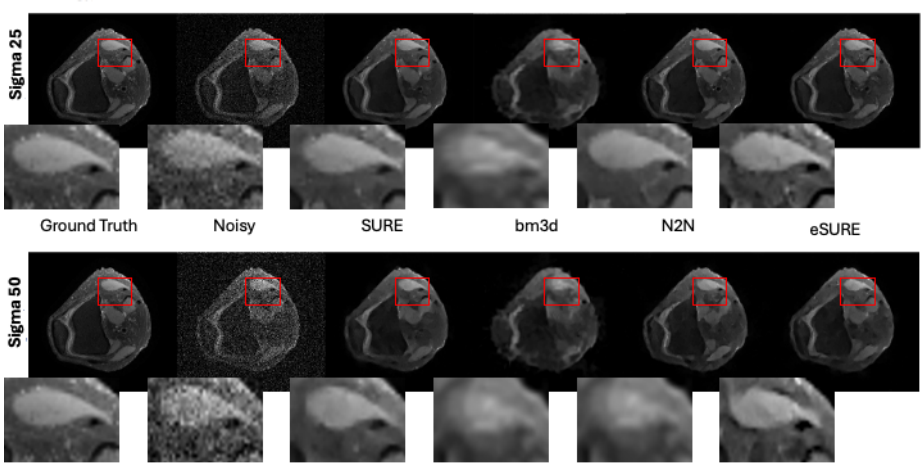

Fig. 1. Denoised results of BM3D[8] and DnCNN trained with different objectives on axial slices for Gaussian noise $\sigma = 25$ and $\sigma = 50$.

resentative denoising outputs for both views at $\sigma = 25$ and $\sigma = 50$. In these examples, BM3D[8] tends to suppress noise at the expense of fine texture, producing visibly smoother structures. Single-observation SURE training improves sharpness relative to BM3D[8], but residual grain and mild structural washout remain apparent in the stronger-noise regime. In contrast, eSURE consistently retains sharper boundaries and more coherent texture, with reconstructions that are visually close to N2N and, in the $\sigma = 50$ coronal setting, noticeably cleaner.

4.5 Discussion

The empirical findings reinforce two practical messages about estimator-driven learning in MRI denoising. First, paired-observation training is valuable even when the underlying estimator remains unbiased in expectation. MC-SURE replaces the intractable divergence with a stochastic approximation (Eq. (5)), which introduces variance into the training signal; in this setting, the additional

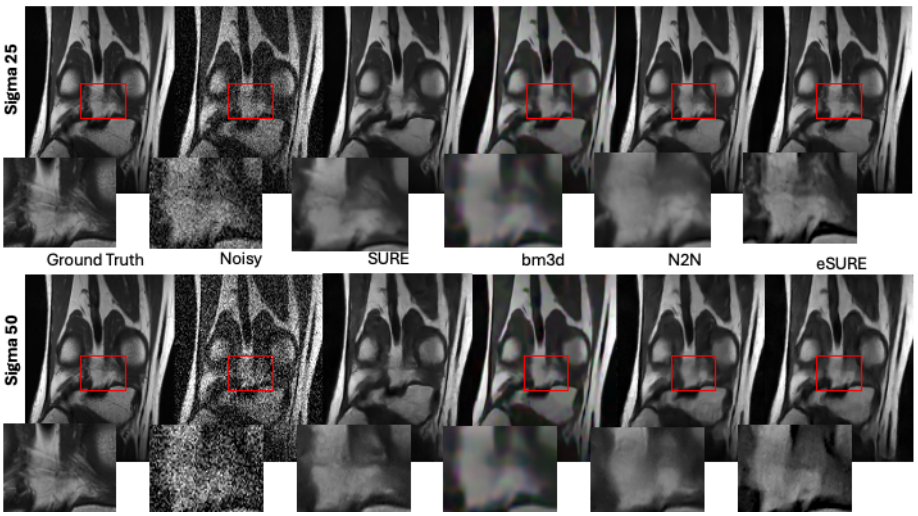

Fig. 2. Denoised results of BM3D[8] and DnCNN trained with different objectives on coronal slices for Gaussian noise $\sigma = 25$ and $\sigma = 50$.

structure provided by a paired observation in eSURE tends to stabilize optimization and improve final fidelity. This effect is reflected by consistent gains of eSURE over DnCNN-SURE across both views and both noise levels in Table 1.

Second, eSURE behaves in practice as a close counterpart to Noise2Noise under Gaussian corruption, while offering a more explicit risk-estimation interpretation. At $\sigma = 25$, eSURE essentially matches N2N, suggesting that both objectives recover similar solutions when the assumptions behind the N2N decomposition are well satisfied. As noise increases, the coronal-view results indicate that eSURE can provide an advantage, which is consistent with the intuition that coupling a paired-observation fidelity term with an explicit divergence correction can better control bias-variance trade-offs in strong corruption. Overall, these results suggest that eSURE serves as a reliable estimator-driven alternative when paired observations are available or can be constructed, and that it can approach N2N performance without requiring clean targets.

Finally, the observed differences in Figs. 1 and 2 underscore that improvements are not limited to PSNR. In knee MRI, clinically relevant cues often manifest as low-contrast boundaries and fine texture. Across both views, eSURE tends to preserve these cues more faithfully than BM3D[8] and single-observation SURE training, especially in the stronger-noise regime. Taken together, the quantitative and qualitative evidence supports the use of estimator-driven paired-observation training as a practical pathway for MRI denoising when clean references are unavailable.

5 Poisson Unbiased Risk Estimation for MRI Denoising

In this section, N denotes the number of pixels per image, consistent with the Gaussian-noise formulations in Sections 2 and 3. Gaussian noise is a convenient and widely used approximation for modeling MRI corruption, particularly when analyzing estimator behavior under additive perturbations. However, several imaging regimes exhibit *signal-dependent* variability, where uncertainty increases with intensity and low-signal regions become disproportionately unreliable. To study estimator-driven learning beyond the additive Gaussian setting, we consider a Poisson corruption model as a controlled proxy for intensity-dependent variance noise and adopt Poisson unbiased risk estimation (PURE) as the corresponding training principle.

5.1 Poisson Unbiased Risk Estimator (PURE)

Under a Poisson observation model, the variance of each measurement is tied to its mean, which invalidates Gaussian SURE assumptions and can bias training if Gaussian objectives are applied directly. Poisson Unbiased Risk Estimation (PURE) addresses this limitation by providing an unbiased surrogate of the mean squared error (MSE) that depends only on the noisy observation and discrete perturbations of the estimator [21, 24].

Let $x \in \mathbb{R}_+^N$ denote the clean image and let

$$y_i \sim \text{Poisson}(x_i), \quad i = 1, \dots, N, \tag{13}$$

denote the observed Poisson-corrupted image. For a denoising function $h(\cdot)$, the Poisson unbiased risk estimator can be written as

$$\text{PURE}(h, y) = \frac{1}{N} \sum_{i=1}^{N} \Big[(h_i(y) - y_i)^2 - y_i + 2y_i \big(h_i(y) - h_i(y - e_i) \big) \Big], \tag{14}$$

where e_i denotes the canonical basis vector corresponding to pixel i. This estimator is unbiased with respect to the Poisson distribution and depends only on the noisy observation y.

Direct evaluation of Eq. (14) requires computing $h(y - e_i)$ for every pixel, which is computationally prohibitive for high-resolution images and deep networks.

5.2 Monte-Carlo Approximation and Paired-Observation Stabilization

To make PURE tractable in deep learning settings, we adopt a Monte-Carlo approximation by randomly sampling a subset of pixel indices per image, following prior work on practical PURE implementations [24, 25]. Concretely, for each training image we sample K indices uniformly at random from $\{1, \dots, N\}$ (excluding indices with $y_{1,i} = 0$ to satisfy the nonnegativity constraint in $y_1 - e_i$).

The index set $\mathcal{S}$ is re-sampled every iteration, and we use a single Monte-Carlo sample per image (i.e., one sampled set $\mathcal{S}$) to keep training efficient while maintaining stable gradients. This approximation yields stable gradients while avoiding the cost of evaluating the estimator at every possible discrete perturbation.

In addition, unbiased risk estimators under Poisson noise can exhibit high gradient variance when only a single noisy realization is available. Motivated by the stabilization effect observed for eSURE in the Gaussian experiments, we further leverage *paired* Poisson-corrupted observations to reduce variance during training. Given two correlated Poisson observations y_1 and y_2 of the same underlying slice, we form an averaged observation

$$z = \frac{y_1 + y_2}{2}, \tag{15}$$

which serves as a lower-variance proxy while preserving anatomical structure. In practice, we retain the PURE-based term as the statistically grounded risk component, and introduce an additional paired-observation stabilization term based on z. Importantly, once this extra fidelity term is added, the overall training loss is no longer a strictly unbiased estimator of the MSE; rather, it is a variance-reduction modification that empirically stabilizes optimization and reduces gradient noise in low-count regimes, analogous in spirit to the stabilization gained by paired observations in eSURE under Gaussian noise. We therefore optimize an ePURE-style objective of the form

$$\begin{aligned} \mathcal{L}_{\mathrm{ePURE}}(\theta) = \frac{1}{K} \sum_{i \in \mathcal{S}} \Big[& (h_\theta(y_1)_i - y_{1,i})^2 - y_{1,i} \\ & + 2y_{1,i} \big(h_\theta(y_1)_i - h_\theta(y_1 - e_i)_i \big) \Big] \\ & + \ \beta \cdot \frac{1}{N} \| z - h_\theta(y_1) \|^2, \end{aligned} \tag{16}$$

where $\mathcal{S}$ is a randomly sampled set of K pixel indices, and β controls the strength of the stabilization term.

Nonnegativity constraint. Eq. (14) involves evaluating $h(y - e_i)$, which requires $y_i \geq 1$ to preserve nonnegativity. In practice, we restrict the sampled indices to those with $y_{1,i} \geq 1$; if a sampled index has $y_{1,i} = 0$, it is skipped (and replaced by another index) to avoid invalid negative inputs.

5.3 Poisson Noise Construction and Experimental Protocol

We apply synthetic Poisson corruption to magnitude MRI slices to create a controlled count-driven noise setting. Specifically, each clean slice x is first normalized to $[0, 1]$ and scaled to a photon-rate domain as $\tilde{x} = \lambda x$, where λ controls the overall count level. A noisy observation is then generated pixelwise as $y \sim \mathrm{Poisson}(\tilde{x})$, followed by rescaling back to the original intensity range by y/λ. We report results for $\lambda \in \{0.5, 1.0\}$ to span low- and moderate-count regimes.

Table 2. PSNR and SSIM for Poisson unbiased risk training on 3T knee MRI (coronal view).

Noise Level	PSNR (dB)	SSIM
$\lambda = 0.5$	31.68	0.934
$\lambda = 1.0$	32.45	0.947

All other training details follow the Gaussian experiments to ensure that observed differences are attributable to the learning objective rather than architecture or optimization. We train DnCNN using the Poisson unbiased risk objective described above and evaluate on coronal slices of the 3T knee MRI dataset described in Section 4 using PSNR and SSIM.

5.4 Results under Poisson Noise

These experiments address **RQ3** by evaluating whether a stabilized Poisson unbiased risk objective can support effective unsupervised MRI denoising under signal-dependent noise.

Table 2 reports quantitative results under Poisson corruption. The estimator-driven training remains stable across both intensity levels and achieves high structural fidelity, with SSIM above 0.93 in the low-count regime and above 0.94 in the higher-count regime. Figure 3 provides qualitative examples. As Poisson corruption increases, thin boundaries and low-contrast cartilage structures are easily obscured in the noisy inputs. The proposed estimator-driven approach suppresses count-driven noise while maintaining boundary sharpness, indicating that the objective preserves diagnostically relevant structure rather than simply smoothing low-intensity regions.

5.5 Discussion

The Poisson experiments complement the Gaussian findings by demonstrating that the estimator-driven perspective extends beyond additive noise when the loss formulation respects the underlying noise statistics. The paired-observation construction in Eq. (15) plays a role analogous to eSURE in the Gaussian setting by reducing estimator variance and improving training stability, while the discrete Poisson risk formulation ensures statistical consistency. Taken together with Section 4, these results support a unified message: when clean targets are unavailable, unbiased risk estimation provides a principled and practical foundation for training denoisers directly from noisy MRI across different noise regimes.

6 Conclusion

This work studied estimator-driven learning as a practical alternative to supervised denoising for magnetic resonance imaging, where clean ground truth data

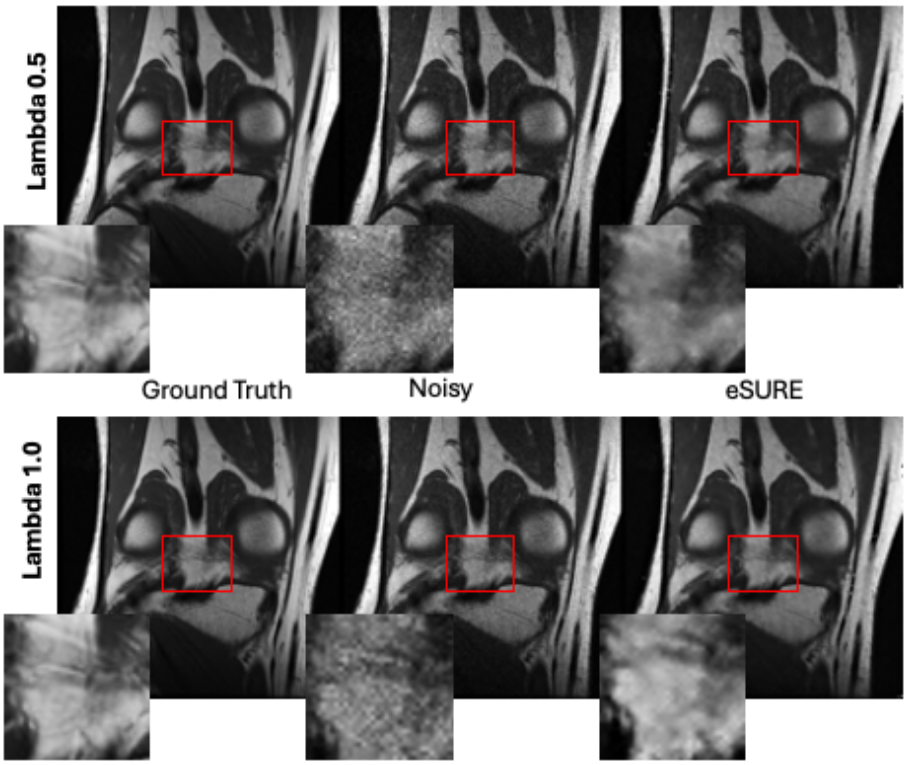

Fig. 3. Denoising results on coronal knee MRI under Poisson noise at different intensity levels using Poisson unbiased risk training.

is difficult to obtain. We focused on adapting and validating unbiased risk estimation frameworks—SURE, eSURE, and a Poisson-based extension—within a realistic MRI denoising pipeline, emphasizing stability, noise-model fidelity, and deployability. Under additive Gaussian noise, eSURE consistently outperformed single-image SURE/MC-SURE training and closely matched Noise2Noise performance across axial and coronal views. On 3T knee MRI, eSURE achieved gains of up to 2.4 dB in PSNR over standard SURE. These results indicate that exploiting paired noisy observations provides a more stable and effective learning signal, even when the observations are correlated and clean targets are unavailable. For signal-dependent noise, we extended the estimator-driven approach to the Poisson setting using an ePURE-style formulation based on correlated noisy pairs. On coronal knee MRI slices, this method achieved PSNR values exceeding 32 dB and SSIM above 0.94 at higher Poisson intensities, while preserving fine anatomical structure in low-signal regions. This demonstrates that unbiased risk estimation can be operationalized beyond Gaussian assumptions when the loss formulation reflects the underlying noise statistics. Overall, the results show that

unbiased risk objectives, particularly when combined with paired observations, provide a reliable foundation for unsupervised MRI denoising. The proposed methods offer performance comparable to supervised and Noise2Noise baselines without requiring clean or proxy references, making them well suited for clinical and research imaging workflows. Future work will explore extensions to mixed noise models, accelerated MRI reconstruction, and task-driven settings where denoising supports downstream clinical analysis.

Appendix

A. Implementation Details

This appendix summarizes implementation-specific details that support reproducibility and are not central to the methodological exposition in the main text.

Monte-Carlo approximation. For all SURE- and eSURE-based objectives, the divergence term was approximated using the Monte-Carlo finite-difference estimator proposed in [18, 19]. A single random perturbation vector was used per training sample, with a fixed step size of $\epsilon = 10^{-3}$. This configuration was found to provide stable optimization across all experiments, and increasing the number of Monte-Carlo samples did not yield observable performance gains.

DnCNN-SURE baseline.* The DnCNN-SURE* variant reported in Table 1 corresponds to a reduced SURE objective in which the Monte-Carlo divergence term is omitted. Specifically, training is performed using only the data-fidelity term and the noise-variance correction. This baseline is included to isolate the empirical contribution of the divergence term in SURE-based optimization.

Poisson noise generation. For experiments under Poisson corruption, magnitude MRI slices were normalized to the range $[0, 1]$ and scaled by a factor λ to define Poisson rate parameters. Noisy observations were generated via pixelwise Poisson sampling and rescaled by $1/\lambda$ to recover the original intensity range. The values $\lambda \in \{0.5, 1.0\}$ were selected to represent low- and moderate-count regimes, respectively.

Monte-Carlo PURE sampling and stabilization. For Poisson-risk training, we approximate PURE by sampling $K = 2048$ pixel indices per image uniformly at random (excluding indices with $y_{1,i} = 0$ to ensure $y_1 - e_i \geq 0$). The sampled index set is re-drawn each iteration. We optimize Eq. (16) with a small paired-observation stabilization weight β, using $\beta = 0.20$ for the low-count regime ($\lambda = 0.5$) and $\beta = 0.10$ for the moderate-count regime ($\lambda = 1.0$), which empirically reduces gradient variance without dominating the PURE term.

References

1. Ghadimi, M., Thomas, A.: Magnetic Resonance Imaging Contraindications. In: *StatPearls* [Internet]. StatPearls Publishing, Treasure Island (FL) (2025). Available from: https://www.ncbi.nlm.nih.gov/books/NBK551669/
2. Liu, R., Xiao, S., Liu, T., Jiang, F., Yuan, C., Chen, J.: Dual stage MRI image restoration based on blind spot denoising and hybrid attention. *BMC Medical Imaging* **24**(1), 259 (2024). doi:10.1186/s12880-024-01437-8
3. Khan, R., Nakarmi, U., Gauch, J.: From Noise Estimation to Restoration: A Unified Diffusion and Bayesian Risk Approach for Unsupervised Denoising. In: *Proceedings of the 20th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISAPP)*. SCITEPRESS (2025), pp. 547-555. doi: 10.5220/0013187300003912. http://dx.doi.org/10.5220/0013187300003912
4. Gudbjartsson, H., Patz, S.: The Rician distribution of noisy MRI data. *Magnetic Resonance in Medicine* **34**(6), 910-914 (1995)
5. Aja-Fernández, S., Tristán-Vega, A., Alberola-López, C.: *Noise and Signal Estimation in Magnetic Resonance Imaging*. Springer (2015)
6. Manjón, J.V., Coupé, P.: MRI denoising using non-local means. In: *Medical Image Processing* (tutorial/survey context) (2015)
7. Buades, A., Coll, B., Morel, J.-M.: A non-local algorithm for image denoising. In: *CVPR* (2005)
8. Dabov, K., Foi, A., Katkovnik, V., Egiazarian, K.: Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on Image Processing* **16**(8), 2080-2095 (2007)
9. Burger, H.C., Schuler, C.J., Harmeling, S.: Image denoising: Can plain neural networks compete with BM3D? In: *CVPR* (2012)
10. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing* **26**(7), 3142-3155 (2017)
11. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Deep Image Prior. In: *CVPR* (2018)
12. Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., Aila, T.: Noise2Noise: Learning image restoration without clean data. In: *ICML* (2018)
13. Krull, A., Buchholz, T.-O., Jug, F.: Noise2Void - Learning denoising from single noisy images. In: *CVPR* (2019)
14. Batson, J., Royer, L.: Noise2Self: Blind denoising by self-supervision. In: *ICML* (2019)
15. Soltanayev, S., Chun, S.Y.: Training deep learning based denoisers without ground truth data. In: *NeurIPS* (2018)
16. Stein, C.M.: Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics* **9**(6), 1135-1151 (1981)
17. Blu, T., Luisier, F.: The SURE-LET approach to image denoising. *IEEE Transactions on Image Processing* **16**(11), 2778-2786 (2007)
18. Ramani, S., Blu, T., Unser, M.: Monte-Carlo SURE: A black-box optimization of regularization parameters for general denoising algorithms. *IEEE Transactions on Image Processing* **17**(9), 1540-1554 (2008)
19. Shao, L., Li, H., Chang, H., Yang, M.-H.: Training deep convolutional neural networks for image denoising using SURE. *IEEE Transactions on Image Processing* **27**(10), 4755-4767 (2018)

20. Zhussip, M., Soltanayev, S., Chun, S.Y.: Extending Stein's unbiased risk estimator to train deep denoisers with correlated pairs of noisy images. In: *NeurIPS* (2019)
21. Luisier, F., Blu, T., Unser, M.: Image denoising in mixed Poisson-Gaussian noise. *IEEE Transactions on Image Processing* **20**(3), 696-708 (2011)
22. Le Montagner, Y., Angelini, E.D., Olivo-Marin, J.-C.: 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)
23. Cherry, S.R., Sorenson, J.A., Phelps, M.E.: *Physics in Nuclear Medicine*. Elsevier/Saunders (2012)
24. Kim, H., Chun, S.Y.: Unbiased risk estimator for image denoising in Poisson noise. *IEEE Signal Processing Letters* (2020)
25. Kim, H., Yie, S.Y., Chun, S.Y., Lee, J.S.: PureComb: Poisson Unbiased Risk Estimator based ensemble of self-supervised deep denoisers for clinical bone scan images. In: *IEEE ISBI* (2022)
26. Deledalle, C.-A., Denis, L., Tupin, F.: Iterative weighted maximum likelihood denoising with probabilistic patch-based weights. *IEEE Transactions on Image Processing* **18**(12), 2661-2672 (2009)
27. Deledalle, C.-A., Denis, L., Tupin, F.: Unbiased risk estimation for nonlocal means. *IEEE Transactions on Image Processing* **20**(12), 3309-3322 (2011)
28. Stanford University: 3D Knee MRI Dataset. *Stanford Digital Repository* (2013). https://mridata.org