
Zero-shot Denoising via Neural Compression: Theoretical and algorithmic framework

Ali Zafari*
ali.zafari@rutgers.edu

Xi Chen*
xi.chen15@rutgers.edu

Shirin Jalali
shirin.jalali@rutgers.edu

Rutgers University
New Brunswick, NJ, USA

Abstract

Zero-shot denoising aims to denoise observations without access to training samples or clean reference images. This setting is particularly relevant in practical imaging scenarios involving specialized domains such as medical imaging or biology. In this work, we propose the *Zero-Shot Neural Compression Denoiser* (ZS-NCD), a novel denoising framework based on neural compression. ZS-NCD treats a neural compression network as an untrained model, optimized directly on patches extracted from a single noisy image. The final reconstruction is then obtained by aggregating the outputs of the trained model over overlapping patches. Thanks to the built-in entropy constraints of compression architectures, our method naturally avoids overfitting and does not require manual regularization or early stopping. Through extensive experiments, we show that ZS-NCD achieves state-of-the-art performance among zero-shot denoisers for both Gaussian and Poisson noise, and generalizes well to both natural and non-natural images. Additionally, we provide new finite-sample theoretical results that characterize upper bounds on the achievable reconstruction error of general maximum-likelihood compression-based denoisers. These results further establish the theoretical foundations of compression-based denoising. Our code is available at: <https://github.com/Computational-Imaging-RU/ZS-NCDdenoiser>.

1 Introduction

Background and motivation Denoising is a fundamental problem in classical signal processing and has recently gained renewed attention from the machine learning community. Let $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}_+^n$ denote a non-negative signal of length n , where signal \mathbf{x} is not observable in many systems. Instead, we observe a noisy version $\mathbf{y} = (y_1, \dots, y_n)$, where the observations are conditionally independent given \mathbf{x} , and each entry is distributed according to a common conditional distribution:

$$\mathbf{y} \sim \prod_{i=1}^n p(y_i \mid x_i).$$

We assume that the noise mechanism is memoryless (independent across coordinates) and homogeneous (identical across entries). The goal of a denoising algorithm is to estimate \mathbf{x} from the noisy observations \mathbf{y} . Given its prevalence in imaging and data acquisition systems, denoising has been a central topic in signal processing for decades. Classical denoising methods rely on explicit structural assumptions about the underlying signal \mathbf{x} , often hand-crafted by domain experts [1, 2, 3, 4, 5, 6, 7, 8]. In contrast, recent advances in machine learning have enabled a new class of data-driven denoising

*Equal contribution.

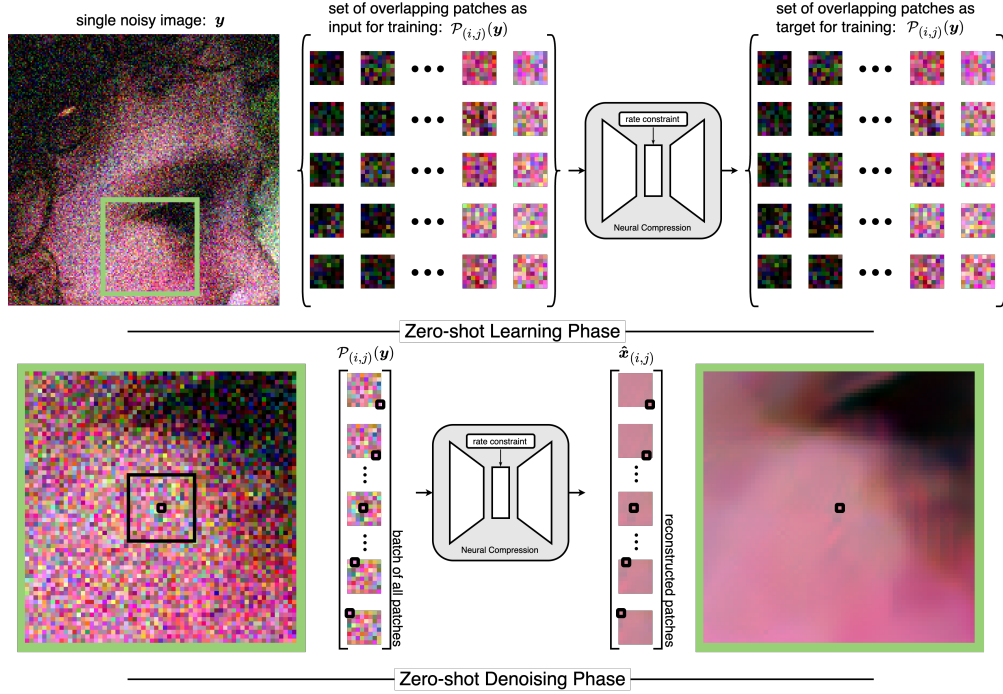


Figure 1: Zero-Shot Neural Compression Denoiser (ZS-NCD). Learning phase: a neural compression model (architecture shown in Fig. 5 of the supplementary material) is trained on overlapping patches extracted from a single noisy image. Denoising phase: each pixel is reconstructed by averaging predictions across neighboring patches processed by the trained model.

algorithms. These methods learn the optimal denoising function from data, leveraging statistical patterns directly from signal and noise distributions.

While learning-based approaches achieve state-of-the-art performance and often outperform classical methods in controlled settings, they face significant challenges in practice:

1. Supervision requirement: Most learning-based methods require training set of paired samples $\{(y_i, x_i)\}_{i=1}^m$, where x_i is clean signal and y_i is its noisy counterpart. In practical scenarios such as medical imaging, $\{x_i\}_{i=1}^m$ are unavailable or prohibitively expensive to obtain.
2. Data efficiency: These methods usually need lots of training data. Acquiring sufficient samples is difficult or costly, particularly in domains with strict data acquisition constraints.

To mitigate the reliance on paired clean and noisy samples, several self-supervised denoising methods have been developed that learn directly from noisy observations, without access to clean ground truth signals [9, 10, 11, 12, 13]. While these approaches alleviate the supervision requirement, they typically depend on access to large collections of noisy data and often yield suboptimal performance compared to methods trained with clean targets. Moreover, the absence of clean supervision necessitates the use of complex neural architectures and training schemes, which can make these methods computationally demanding and difficult to optimize in practice.

These challenges have sparked growing interest in *zero-shot denoisers*, which aim to recover clean signals from noisy observations without access to paired data or extensive noisy training data. Such methods are particularly appealing in domains where acquiring clean data is infeasible, and they offer the potential for deployable denoisers that adapt to individual inputs with general purpose.

From neural compression to zero-shot denoising Denoising algorithms—ranging from classical signal processing techniques to deep learning methods—fundamentally rely on the assumption that real-world signals are highly structured. Compression-based denoising leverages this same principle, but rather than directly solving the inverse problem, it instead performs lossy compression on the

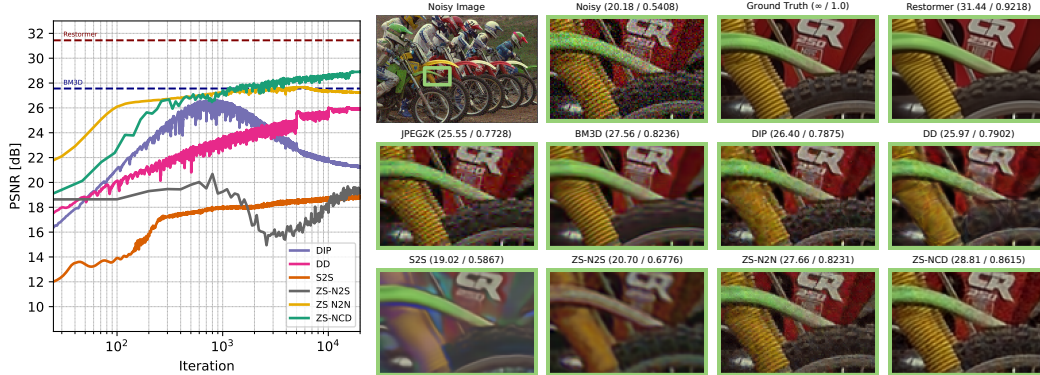


Figure 2: Zero-shot denoising of *Kodim05* with AWGN ($\sigma = 25$). **Left:** PSNR versus training iterations for zero-shot denoisers. Performance of BM3D [18] and Restormer [19] are included as a classical baseline and as a supervised empirical upper bound, respectively. **Right:** Visual reconstructions with PSNR above each image. Compression-based denoising based on JPEG-2K [20] achieves inferior performance. Learning-based zero-shot denoisers often struggle with either overfitting or high bias. DIP [21] and DD [22] require early stopping to avoid overfitting. ZS-N2S [12] and S2S [23] struggle with high-resolution color images, and ZS-N2N [24] often produces noisy outputs with potential overfitting. BM3D tends to oversmooth the denoised image. In contrast, ZS-NCD avoids these issues.

noisy observation \mathbf{y} , under the hypothesis that the clean signal \mathbf{x} lies in a lower-complexity subspace and is therefore more compressible.

In lossy compression, the goal is to represent signals from a target class using discrete encodings with minimal distortion. When applied to noisy data, the intuition is that a lossy compressor—operating at a distortion level matched to the noise—will favor reconstructions close to the original clean signal. While this approach has a strong theoretical foundation [14, 15], classical compression-based denoisers have shown limited empirical success, particularly for natural image denoising.

In this work, we revisit this idea in light of recent progress in *neural compression*, where learned encoders and decoders have demonstrated strong rate-distortion performance across a variety of image domains [16, 17]. Building on this foundation, we propose a *zero-shot* denoising method that we call the **Zero-Shot Neural Compression Denoiser (ZS-NCD)**. Unlike traditional neural compression models that are trained on large corpora of clean high-resolution images, ZS-NCD learns directly from a single noisy input image. Specifically, we extract overlapping patches from the noisy image, and train a neural compression network on those patches alone—without any clean supervision or prior dataset. Once trained, the denoiser is applied to all patches from the same image, and the final output is obtained by averaging the predictions in overlapping regions. This approach is illustrated in Figure 1.

Despite relying solely on the noisy input and operating without supervision, ZS-NCD achieves state-of-the-art performance among zero-shot denoising methods across diverse noise models, and remains robust even on inputs that lie outside the natural image distribution. We compare it against the baselines in Figure 2, ZS-NCD shows superior performance in denoising and training stability.

Paper contributions This paper introduces a zero-shot image denoising framework based on neural compression. Our main contributions are:

- **A zero-shot denoising algorithm using neural compression.** We propose a fully unsupervised method that trains a neural compression network on image patches of the noisy input. It does not rely on clean images, paired datasets, or prior training on the target distribution. It is architecture-agnostic and leverages only the structure present in the observed noisy image.
- **Theoretical results connecting denoising and compression performance.** We establish finite-sample upper bounds on the reconstruction error of the proposed compression-based maximum likelihood denoisers, for both Gaussian and Poisson noise models.

- **Extensive empirical validation.** We demonstrate that our method achieves state-of-the-art performance among zero-shot denoising techniques across a range of noise models and datasets.

2 Related work

Self-supervised and zero-shot denoising Supervised learning-based denoisers such as DnCNN [25] and Restormer [19] achieve state-of-the-art performance across various noise models, but require large datasets of paired clean and noisy images—often impractical in real-world settings. To avoid clean images, self-supervised methods have been proposed, including Noise2Noise [10], Noise2Self [12], Noise2Void [11], Noise2Same [26] and Noise2Score [13], which only use noisy images for training. However, their reliance on large noisy datasets remains a limitation. Zero-shot denoisers address this by training on a single noisy image. These include (i) untrained networks like DIP [27] and Deep Decoder [22], and (ii) single-image adaptations of self-supervised methods, *e.g.*, N2F [28], ZS-N2N [24], ZS-N2S [12] and its augmented variant with ensembling, S2S [23]. More recently, DS-N2N [29] improves ZS-N2N by further upsampling the downsampled paired noisy images. Pixel2Pixel [30] boosts the performance by using non-local similarity approach. DIP-based models avoid masking and leverage full-image context, but require early stopping or under-parameterization to avoid overfitting. Self-supervised variants suffer from masking-induced information loss. Hybrid approaches, such as masked pretraining-based method [31], uses external datasets for training and perform zero-shot inference, thus falling outside the zero-shot setting studied here.

Neural compression Learning-based lossy compression, often referred to as neural compression, uses an autoencoder architecture combined with an entropy model to estimate and constrain the bitrate at the bottleneck [32, 33]. These methods have significantly outperformed traditional codecs, particularly in image [32, 34, 35, 36, 37] and video [38, 39, 40] compression. In addition to these standard settings, several works have also explored applying neural compression to noisy data, either by adapting neural compression models for more efficient encoding of noisy images [41, 42, 43, 44].

Compression-based denoising Compression-based denoising leverages the insight that structured signals are inherently more compressible than their noisy counterparts. This connection was formalized by Donoho [14], who introduced the minimum Kolmogorov complexity estimator, and further refined by Weissman et al. [15], showing that, under certain conditions on both the signal and the noise, *optimal* lossy compression of a noisy signal—followed by suitable post-processing—can asymptotically achieve optimal denoising performance. Prior to these theoretical developments, early empirical methods such as wavelet-based schemes [45, 46, 47] and MDL-inspired heuristics [48] had explored this principle. Nevertheless, traditional compression-based denoisers have generally underperformed in high-dimensional settings such as image denoising.

Learning-based joint compression and denoising using neural compression has been explored in recent works [49, 50], where the goal is to achieve lower rate in compression. The empirical application of training neural compression for AWGN denoising was also proposed in [51]. Training the neural compression models in these works requires a dataset of images. In contrast, our proposed ZS-NCD is a two-step denoiser based on neural compression, trained on a single noisy image. It achieves state-of-the-art performance across both AWGN and Poisson noise models. Moreover, we contribute new theoretical results that advance the foundations of compression-based denoising.

3 Compression-based denoising: Theoretical foundations

Lossy compression Let $\mathcal{Q} \subset \mathbb{R}^n$ denote the signal class of interest, such as vectorized natural images of a fixed size. A lossy compression code for \mathcal{Q} is defined by an encoder-decoder pair (f, g) , $f : \mathcal{Q} \rightarrow \{1, \dots, 2^R\}$, and $g : \{1, \dots, 2^R\} \rightarrow \mathbb{R}^n$. The performance of a lossy code is characterized by: i) *Rate* R , indicating the number of distinct codewords; ii) *Distortion* δ , defined as the worst-case per-symbol mean squared error (MSE) over the signal class:

$$\delta = \sup_{\mathbf{x} \in \mathcal{Q}} \frac{1}{n} \|\mathbf{x} - g(f(\mathbf{x}))\|_2^2.$$

The set of reconstructions produced by the decoder forms the codebook:

$$\mathcal{C} = \{g(i) : i = 1, \dots, 2^R\} \subset \mathbb{R}^n.$$

Compression-based denoising We propose compression-based denoising as a structured maximum likelihood (ML) estimation. Given a noisy observation $\mathbf{y} \sim \prod_{i=1}^n p(y_i | x_i)$ and a lossy compression code (f, g) for \mathcal{Q} , the compression-based ML denoiser solves

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{c} \in \mathcal{C}} \mathcal{L}(\mathbf{c}; \mathbf{y}), \quad \text{where} \quad \mathcal{L}(\mathbf{c}; \mathbf{y}) := -\sum_{i=1}^n \log p(y_i | c_i).$$

This formulation leverages the fact that clean signals, by virtue of their structure, are more compressible than their noisy counterparts. Therefore, the most likely codeword under the noise model, when selected from a codebook designed to represent clean signals, serves as a natural denoising estimate. This ML-based view unifies denoising across noise models and provides a principled way to select reconstructions from a discrete, structure-aware prior.

In the case of AWGN: $\mathbf{y} = \mathbf{x} + \mathbf{z}$, where $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \sigma_z^2 \mathbf{I}_n)$, the described denoiser simplifies to:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{c} \in \mathcal{C}} \|\mathbf{y} - \mathbf{c}\|_2^2. \quad (1)$$

That is, denoising corresponds to projecting the noisy observation onto the nearest codeword.

Poisson noise commonly arises in low-light and photon-limited imaging scenarios. In this setting, each y_i is modeled as a Poisson random variable with mean αx_i : $y_i \sim \text{Poisson}(\alpha x_i)$. Under this model, the compression-based ML denoiser simplifies to

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{c} \in \mathcal{C}} \sum_{i=1}^n (\alpha c_i - y_i \log c_i). \quad (2)$$

While the loss function in (2) is statistically well-motivated, it is more sensitive to optimization issues than its Gaussian counterpart due to the curvature and nonlinearity of the log term. To improve robustness and simplify optimization, we also consider an alternative loss based on a normalized squared error between \mathbf{c} and the rescaled observations:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{c} \in \mathcal{C}} \left\| \mathbf{c} - \frac{1}{\alpha} \mathbf{y} \right\|_2^2. \quad (3)$$

Theoretical analysis We begin by analyzing the performance of compression-based ML denoising under AWGN. The following result provides a non-asymptotic upper bound on the reconstruction error in terms of the compression rate and distortion. All proofs can be found in Appendix A.

Theorem 1. Assume that $\mathbf{x} \in \mathcal{Q}$ and let (f, g) denote a lossy compression for \mathcal{Q} that operates at rate R and distortion δ . Consider $\mathbf{y} = \mathbf{x} + \mathbf{z}$, where $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \sigma_z^2 \mathbf{I}_n)$. Let $\hat{\mathbf{x}}$ denote the output of the compression-based denoiser defined by (f, g) as in (1). Then,

$$\frac{1}{\sqrt{n}} \|\mathbf{x} - \hat{\mathbf{x}}\|_2 \leq \sqrt{\delta} + 2\sigma_z \sqrt{\frac{(2 \ln 2)R}{n}} (1 + 2\sqrt{\eta}), \quad (4)$$

with a probability larger than $1 - 2^{-\eta R + 2}$.

This bound decomposes the denoising error into two terms: a *distortion term* $\sqrt{\delta}$, which reflects the approximation quality of the compression code, and a *rate-dependent term* that scales with the square root of the code rate R . The latter captures the likelihood concentration around the clean signal in high-probability regions of the noise distribution. Notably, the result holds non-asymptotically and does not assume the code is optimal, only that it provides a distortion- δ covering of \mathcal{Q} . This highlights that even non-ideal compression codes can enable effective denoising, provided the rate-distortion tradeoff is well-calibrated.

To better understand the implications of Theorem 1, in the following corollary, we focus on the special case of k -sparse signals.

Corollary 1 (AWGN, sparse signals). Let \mathcal{Q}_n denote the set of k -sparse vectors in \mathbb{R}^n satisfying $\|\mathbf{x}\|_\infty \leq 1$. Fix a parameter $\eta \in (0, 1)$, and suppose $\mathbf{y} = \mathbf{x} + \mathbf{z}$ where $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \sigma_z^2 \mathbf{I}_n)$. Then, there exists a family of compression codes such that, when used with the denoiser defined in (1), with a probability larger than $1 - e^{-\eta k \log(n/k)}$, the estimate $\hat{\mathbf{x}}$ satisfies

$$\frac{1}{\sqrt{n}} \|\mathbf{x} - \hat{\mathbf{x}}\|_2 \leq \sigma_z C \sqrt{\frac{k}{n} \log_2 \left(\frac{n}{k} \right) + \gamma_n} + \frac{1}{\sqrt{n}}, \quad (5)$$

where $C = 2(1 + 2\sqrt{\eta})\sqrt{2 \ln 2}$ and $\gamma_n = \frac{k \log_2 k}{2n} + \frac{\log_2 k + k(\log_2 e + 1)}{n}$.

Corollary 1 provides a high-probability bound on the normalized error $\frac{1}{\sqrt{n}}\|\hat{\mathbf{x}} - \mathbf{x}\|_2$. Squaring both sides and having $(\sqrt{a} + \sqrt{b})^2 \leq 2(a + b)$, it follows that, with high probability

$$\frac{1}{n}\|\hat{\mathbf{x}} - \mathbf{x}\|_2^2 \leq 2\sigma_z^2 C^2 \left(\frac{k}{n} \log_2 \left(\frac{n}{k} \right) + \gamma_n \right) + \frac{2}{n}.$$

Thus, up to universal constants, the dominant term in the upper bound scales as $\sigma_z^2 \frac{k}{n} \log(\frac{n}{k})$. This matches the known minimax rate for estimating k -sparse signals in Gaussian noise when $k/n \rightarrow 0$; see [52]. Determining whether the residual term $\frac{2}{n}$ is an artifact of our proof or a real barrier to optimality is an interesting problem. Finally, while the comparison above is high-probability rather than in expectation, one can integrate the tail bound from the proof to obtain an expected-risk bound.

We next extend our analysis to signal-dependent noise model. Poisson noise is particularly relevant in imaging applications such as microscopy and astronomy, where photon counts vary with signal intensity. Unlike Gaussian noise, Poisson observations induce a non-linear likelihood surface, making analysis more delicate. Theorem 2 and 3 establish performance guarantees for compression-based Poisson denoising, using both exact ML formulation and a practical squared-error surrogate.

Theorem 2. Consider the same setup of lossy compression as in Theorem 1. Assume that for any $\mathbf{x} \in \mathcal{Q}$, $x_i \in (x_{\min}, x_{\max})$, where $0 < x_{\min} < x_{\max} < 1$. Assume that y_1, \dots, y_n are independent with $y_i \sim \text{Poisson}(\alpha x_i)$. Let $\hat{\mathbf{x}}$ denote the solution of (2). Let $C_1 = x_{\max}^5 / (x_{\min}^2)$ and $C_2 = \frac{x_{\max}^2}{x_{\min}^3} \beta \sqrt{(\frac{4}{\ln 2})(\sqrt{1+\eta} + \sqrt{\eta})}$. Then, with a probability larger than $1 - 2^{-\eta R+2}$,

$$\frac{1}{n}\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 \leq C_1 \delta + C_2 \sqrt{\frac{R}{n\alpha}}. \quad (6)$$

Theorem 3. Consider the same setup as in Theorem 2. Let $\hat{\mathbf{x}}$ denote the solution of (3). Let $C = 4\sqrt{\ln 2}(\sqrt{1+\eta} + \sqrt{\eta} + 1)$. Then, with a probability larger than $1 - 2^{-\eta R+2}$,

$$\frac{1}{n}\|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 \leq \delta + C \sqrt{\frac{R}{n\alpha}}. \quad (7)$$

Remark 1. Theorems 2 and 3 show that, in the case of Poisson noise, minimizing either the ML loss function or the computationally efficient MSE loss function can recover the signal. This result is also consistent with our simulations reported later in Section 5.

4 Zero-shot compression-based denoiser

We refer to a general class of learning-based denoisers that operate by compressing noisy images using neural compression as the Neural Compression Denoiser (NCD). In this framework, denoising is achieved by identifying a low-complexity reconstruction from the output of a neural compression model. In the previous section, we characterized the performance of such denoisers in a setting where the compression code is fixed in advance, either learned from external data or designed using classical methods, and applied independently of the noisy input. This setup is not zero-shot, as it relies on prior training or code design. Inspired by this idea, we now propose a fully unsupervised variant: the **Zero-Shot Neural Compression Denoiser (ZS-NCD)**. In ZS-NCD, a neural compression network is trained directly on patches extracted from a single noisy image, without access to clean targets or external data. This section describes the ZS-NCD architecture and optimization procedure in detail.

Proposed zero-shot denoiser: ZS-NCD Let $\mathcal{P}_{(i,j)} : \mathbb{R}^{h \times w} \rightarrow \mathbb{R}^{k \times k}$ denote the patch extraction operator, which returns a $k \times k$ patch whose top-left corner is at pixel (i, j) . Let f_{θ_1} and g_{θ_2} denote the encoder and decoder networks, parameterized by weights θ_1 and θ_2 , respectively. Define \mathcal{I} as the set of all coordinates $(i, j) \in \{1, \dots, h - k + 1\} \times \{1, \dots, w - k + 1\}$ from which a valid $k \times k$ patch can be extracted.

Given a single noisy image \mathbf{y} , the ZS-NCD is trained to minimize the following patchwise objective:

$$(\hat{\theta}_1, \hat{\theta}_2) = \arg \min_{(\theta_1, \theta_2)} \sum_{(i,j) \in \mathcal{I}} (\mathcal{L}_K(g_{\theta_2}(f_{\theta_1}(\mathcal{P}_{(i,j)}(\mathbf{y}))), \mathcal{P}_{(i,j)}(\mathbf{y})) - \lambda \log \mathbb{P}(f_{\theta_1}(\mathcal{P}_{(i,j)}(\mathbf{y})))), \quad (8)$$

where $\mathbb{P}(f_{\theta_1}(\mathcal{P}_{(i,j)}(\mathbf{y})))$ denotes the likelihood (or entropy model) of the latent code produced by the encoder, and $\lambda > 0$ is a hyperparameter controlling the trade-off between fidelity and compressibility.

Algorithm 1 Finding Lagrangian coefficient λ

```

1: Initialize:  $\lambda^{(0)} > 0, \text{tol} > 0, \zeta \in (0, 1)$ 
2: for  $k = 0, 1, \dots, K_{\max} - 1$  do
3:   Estimate  $\hat{\mathbf{x}}^{(k)}$  from (9).
4:   Compute residual  $r^{(k)} = \|\hat{\mathbf{x}}^{(k)} - \mathbf{y}\|^2$ .
5:   Compute scale factor  $\beta = (r^{(k)} - n\sigma_z^2)/n\sigma_z^2$ .
6:   if  $|\beta| \leq \text{tol}$  then
7:     break
8:   else
      $\lambda^{(k+1)} := \begin{cases} \lambda^{(k)}/(1 + \zeta|\beta|), & \beta > 0 \\ \lambda^{(k)}(1 + \zeta|\beta|), & \text{otherwise} \end{cases}$ 
9:   end if
10: end for
11: Return:  $\lambda^* = \lambda^{(k)}$ 

```




Figure 3: Effect of λ in denoising Mouse Nuclei image.

In (8), $K = k^2$, and the function $\mathcal{L}_K : \mathbb{R}^K \times \mathbb{R}^K \rightarrow \mathbb{R}_+$ is a distortion loss determined by the noise model, as defined in Section 3. For example, in the AWGN case, \mathcal{L}_K corresponds to the squared ℓ_2 norm of the distance between a noisy patch and its neural compression reconstruction. Note that f_{θ_1} maps the input into a discrete latent space, which is non-differentiable and thus incompatible with standard gradient-based optimization. To address this, we follow the neural compression framework of [32], using a continuous relaxation during training (e.g., uniform noise injection) and applying actual discretization only at test time. The entropy term \mathbb{P} is modeled using a factorized, non-parametric density [34].

After training, the denoised image is obtained by applying the encoder and decoder to each patch and averaging the overlapping outputs. For each pixel (i, j) , let $\mathcal{I}_{(i,j)} \subset \mathcal{I}$ denote the set of patch locations such that $\mathcal{P}_{(i',j')}$ includes the pixel (i, j) . The final estimate at location (i, j) is given by

$$\hat{x}_{(i,j)} = \frac{1}{|\mathcal{I}_{(i,j)}|} \sum_{(i',j') \in \mathcal{I}_{(i,j)}} g_{\theta_2}(f_{\theta_1}(\mathcal{P}_{(i',j')}(\mathbf{y})))|_{(i-i',j-j')}, \quad (9)$$

where $|\mathcal{I}_{(i,j)}|$ denotes the number of patches covering pixel (i, j) , and $\cdot|_{(a,b)}$ denotes the (a, b) -th pixel of the patch output. For interior pixels away from the boundary, $|\mathcal{I}_{(i,j)}| = K$. As shown later in Section 5, this aggregating of reconstructed patches significantly enhances denoising performance.

Setting the hyperparameter λ The ZS-NCD objective in (8) includes a hyperparameter, λ , which balances reconstruction fidelity and compressibility. Interpreted through the lens of lossy compression, varying λ allows the model to explore different rate-distortion trade-offs. However, in the context of denoising, our goal is not compression but accurate signal recovery from the noisy observation \mathbf{y} . This raises the central question: *how should λ be selected to optimize denoising performance?* In the following, we explain our approach for setting λ under both AWGN model and Poisson noise model.

Case I: Gaussian noise. Let $\hat{\mathbf{x}}$ denote the output of the ZS-NCD denoiser, and consider the AWGN model $\mathbf{y} = \mathbf{x} + \mathbf{z}$ with $\mathbf{z} \sim \mathcal{N}(0, \sigma_z^2 \mathbf{I}_n)$. Then,

$$\frac{1}{n} \mathbb{E}[\|\mathbf{y} - \hat{\mathbf{x}}\|_2^2] = \frac{1}{n} \mathbb{E}[\|\mathbf{y} - \mathbf{x} + \mathbf{x} - \hat{\mathbf{x}}\|_2^2] = \sigma_z^2 + \frac{1}{n} \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 + \frac{2}{n} \mathbb{E}[\mathbf{z}^T (\mathbf{x} - \hat{\mathbf{x}})], \quad (10)$$

the first term is the noise variance, and the second is the true denoising error. While \mathbf{z} and $\hat{\mathbf{x}}$ are not fully independent, they are intuitively weakly correlated in successful denoising regimes, where the estimate $\hat{\mathbf{x}}$ depends only indirectly on the noise. Thus, the cross term is expected to be small: $\frac{1}{n} \mathbb{E}[\mathbf{z}^T (\mathbf{x} - \hat{\mathbf{x}})] \approx 0$. This approximation suggests that, ideally, under low-noise regimes, $\frac{1}{n} \|\mathbf{y} - \hat{\mathbf{x}}\|_2^2$ is expected to be close to σ_z^2 . Based on this insight, we propose a simple and effective heuristic for choosing λ : select λ such that $\frac{1}{n} \|\mathbf{y} - \hat{\mathbf{x}}\|_2^2$ is closest to the known noise variance σ_z^2 . This procedure can be implemented efficiently via a tree-based search strategy, as described in Algorithm 1. To apply Algorithm 1, one needs an estimate of the noise power σ_z^2 . This is a well-studied problem and there exist robust algorithms for estimating the variance of noise [2, 53]. For example, in [2], it is shown that the noise power can be estimated from the median of the absolute differences of wavelet coefficients.

Finally, we observe that the performance of ZS-NCD is relatively robust to the choice of λ . For instance, on the Nuclei dataset (Figure 3), ZS-NCD outperforms the state-of-the-art zero-shot learning-

based denoiser, ZS-Noise2Noise, across a wide range of λ values, for both $\sigma_z = 10$ and $\sigma_z = 20$. A similar approach can also be applied to the case of Poisson noise, as well.

Case II: Poisson noise In the case of Poisson noise, in addition to estimating λ , we need to estimate α , which is used to normalize the measurements, in both the MSE-based and the MLE-based methods. Note that in the case of Poisson noise, $\mathbb{E}[y_i] = \alpha x_i$, and therefore, with high probability, $\frac{1}{n} \sum_{i=1}^n y_i \approx \alpha \frac{1}{n} \sum_{i=1}^n x_i$. Using this observation, and assuming that $\frac{1}{n} \sum_{i=1}^n x_i \approx 0.5$, we estimate α as $\hat{\alpha} = \frac{2}{n} \sum_{i=1}^n y_i$. We then use the estimated noise level $\hat{\alpha}$ to normalize both MSE and MLE based optimization for denoising Poisson noise. See Table 1 for the result of this noise parameter estimation on a sample image. Having an estimate of noise parameter α , we then follow a similar procedure we used in the case of AWGN to set the parameter λ . Specifically, we write the MSE between normalized \mathbf{y} and the denoised image $\hat{\mathbf{x}}$ as

$$\begin{aligned} \frac{1}{n} \mathbb{E}[\|\mathbf{y}/\alpha - \hat{\mathbf{x}}\|_2^2] &= \frac{1}{n} \mathbb{E}[\|\mathbf{y}/\alpha - \mathbf{x} + \mathbf{x} - \hat{\mathbf{x}}\|_2^2] \\ &= \frac{1}{n\alpha} \sum_{i=1}^n x_i + \frac{1}{n} \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 + \frac{2}{n} \mathbb{E}[(\mathbf{y} - \alpha \mathbf{x})^T (\mathbf{x} - \hat{\mathbf{x}})]. \end{aligned} \quad (11)$$

Again, assuming that we are in a low-noise regime, i.e., the second and the third terms in (11) are close to zero, and using $\frac{1}{n} \sum_{i=1}^n x_i \approx 0.5$ approximation, it follows that $\frac{1}{n} \mathbb{E}[\|\mathbf{y}/\alpha - \hat{\mathbf{x}}\|_2^2] \approx \frac{1}{2\alpha}$. This implies that, in the case of Poisson noise, we can still use Algorithm 1 to find λ , after updating $\|\hat{\mathbf{x}}^{(k)} - \mathbf{y}\|^2 > n\sigma_z^2$ to $\|\hat{\mathbf{x}}^{(k)} - \mathbf{y}\|^2 > \frac{1}{2\alpha}$. We empirically observe that $\text{MSE}(\mathbf{y}/\alpha, \hat{\mathbf{x}})$ in training our networks is close to $\frac{1}{2\alpha}$ as reported in Table 1, which indicates that $\frac{1}{2\alpha}$ is a good approximation of the MSE that can be used as a threshold for selecting λ . When α is not known, we obtain the estimate $\hat{\alpha}$, and use $\frac{1}{2\hat{\alpha}}$ as a valid threshold to set λ .

Table 1: Analyzing the estimation of Poisson noise parameter for *Barbara* image in Set11 (MSE values are reported in terms of PSNR).

true α	estimated $\hat{\alpha}$	empirical $\text{MSE}(\mathbf{y}/\alpha, \hat{\mathbf{x}})$ [dB]	$1/(2\alpha)$ [dB]	$1/(2\hat{\alpha})$ [dB]
25	23.02	17.12	16.98	16.63
50	46.05	19.66	20.00	19.64

5 Experiments

In this section, we evaluate the denoising performance of ZS-NCD on both synthetic and real-world noise, across natural and microscopy images. We compare against representative zero-shot denoisers, including both traditional and learning-based methods. All baselines are dataset-free, i.e., they operate solely on the noisy image to be denoised. For non-learning methods, we include JPEG-2K and BM3D. Although rarely used as a denoising baseline, JPEG-2K provides a useful point of comparison from the perspective of compression-based denoising, as it represents a fixed, pre-defined compression code. For learning-based methods, we evaluate Deep Image Prior (DIP) [21], Deep Decoder (DD) [22], Zero-Shot Noise2Self (ZS-N2S) [12], Self2Self (S2S) [23], and Zero-Shot Noise2Noise (ZS-N2N) [24].

Due to instability in training for several baselines, we report their **best achieved performance** (with early stopping or model selection), whereas ZS-NCD is evaluated at its **final training iteration**, without manual tuning or stopping criteria.

Natural images with synthetic noise We consider two synthetic noise models, AWGN $\mathcal{N}(0, \sigma_z^2)$, where σ_z is the standard deviation of the Gaussian distribution, and Poisson noise defined as $\text{Poisson}(\alpha x)$, where α is the scale factor. Note that Poisson noise is signal-dependent noise with $\mathbb{E}[\mathbf{y}] = \alpha \mathbf{x}$. To re-scale the noisy image to the range of clean image, we followed the literature by assuming that the scale α is known, and normalize the noisy image as \mathbf{y}/α in the experiments of this section. We evaluate on grayscale *Set11* [25], RGB *Set13* [54] (center-cropped to 192×192) and *Kodak24* [55] datasets. Table 2 presents the denoising performance of various methods. BM3D achieves the strongest results on grayscale images, though it relies on accurate knowledge of the noise power parameter. Existing learning-based zero-shot denoisers, in contrast, often exhibit inconsistent

Table 2: Denoising performance comparison under AWGN and Poisson Noise, average PSNR(dB) and SSIM are reported. Best results are in **bold**, second-best are underlined.

Noise Parameter		AWGN, $\mathcal{N}(0, \sigma^2)$			Poisson, $\text{Poisson}(\alpha x)/\alpha$		
σ or α	Method	Set11	Set13	Kodak24	Set11	Set13	Kodak24
15	JPEG2K	27.45 / 0.7699	26.69 / 0.7543	27.86 / 0.7457	22.35 / 0.5882	21.76 / 0.5494	22.56 / 0.5249
	BM3D	32.22 / 0.8991	<u>31.15 / 0.8808</u>	<u>32.37 / 0.8754</u>	26.66 / 0.7505	<u>25.64 / 0.6912</u>	<u>27.04 / 0.6900</u>
	DIP	29.11 / 0.7990	30.31 / 0.8570	31.42 / 0.8454	23.69 / 0.5863	25.14 / 0.6916	26.37 / 0.6761
	DD	28.83 / 0.8215	29.22 / 0.8371	28.71 / 0.8016	24.37 / 0.6629	24.96 / <u>0.7006</u>	25.59 / 0.6679
	S2S	26.81 / 0.8158	20.61 / 0.6879	23.08 / 0.7695	21.75 / 0.6872	19.23 / 0.6553	22.52 / 0.7418
	ZS-N2S	28.92 / 0.8495	18.18 / 0.5690	18.68 / 0.5540	25.06 / 0.7051	21.23 / 0.6066	22.24 / 0.6170
	ZS-N2N	30.01 / 0.8169	30.95 / 0.8701	32.30 / 0.8650	24.04 / 0.5766	25.37 / 0.6878	26.80 / 0.6757
	ZS-NCD	<u>31.35 / 0.8580</u>	31.93 / 0.8983	33.18 / 0.9026	<u>25.65 / 0.7132</u>	26.44 / 0.7434	27.64 / 0.7432
25	JPEG2K	24.91 / 0.6997	24.32 / 0.6676	25.43 / 0.6550	23.03 / 0.6108	22.65 / 0.5952	23.58 / 0.5680
	BM3D	29.79 / 0.8523	<u>28.81 / 0.8213</u>	<u>29.98 / 0.8092</u>	22.70 / 0.5741	22.17 / 0.5992	24.13 / 0.5931
	DIP	26.60 / 0.7128	27.85 / 0.7837	28.90 / 0.7738	24.94 / 0.6512	26.13 / 0.7289	27.49 / 0.7243
	DD	26.93 / 0.7530	27.40 / 0.7832	27.62 / 0.7496	25.48 / 0.7022	26.04 / 0.7373	26.56 / 0.7060
	S2S	23.32 / 0.7306	17.95 / 0.5998	20.69 / 0.6949	23.40 / 0.7355	20.18 / 0.6927	23.09 / <u>0.7674</u>
	ZS-N2S	27.30 / 0.7971	20.39 / 0.6200	20.89 / 0.6156	<u>26.01 / 0.7478</u>	21.19 / 0.6312	21.47 / 0.6277
	ZS-N2N	27.18 / 0.7173	28.36 / 0.8001	29.54 / 0.7798	25.40 / 0.6432	<u>26.75 / 0.7455</u>	<u>28.21 / 0.7374</u>
	ZS-NCD	<u>28.93 / 0.8079</u>	29.33 / 0.8351	30.60 / 0.8144	27.10 / 0.7431	27.60 / 0.7827	28.77 / 0.7677
50	JPEG2K	22.05 / 0.5794	21.43 / 0.5295	22.17 / 0.5055	24.77 / 0.6811	24.25 / 0.6696	25.52 / 0.6608
	BM3D	26.56 / 0.7619	<u>25.78 / 0.7134</u>	<u>27.06 / 0.7047</u>	23.09 / 0.5787	23.00 / 0.6281	24.49 / 0.6008
	DIP	23.46 / 0.5783	24.82 / 0.6748	25.90 / 0.6494	26.30 / 0.7004	27.72 / 0.7845	29.12 / 0.7845
	DD	24.01 / 0.6584	24.56 / 0.6779	24.98 / 0.6413	26.87 / 0.7455	27.43 / 0.7867	27.71 / 0.7543
	S2S	17.41 / 0.5200	14.21 / 0.3938	17.00 / 0.5325	25.70 / <u>0.7896</u>	21.75 / 0.7365	23.88 / 0.8014
	ZS-N2S	24.74 / 0.6883	20.62 / 0.5880	20.05 / 0.5774	27.08 / 0.7855	20.75 / 0.6033	20.25 / 0.5993
	ZS-N2N	23.52 / 0.5457	24.67 / 0.6444	25.82 / 0.6151	<u>27.26 / 0.7216</u>	<u>28.57 / 0.8112</u>	30.13 / 0.8076
	ZS-NCD	<u>25.58 / 0.7144</u>	25.87 / 0.7269	27.89 / 0.7464	28.44 / 0.7914	29.09 / 0.8223	30.60 / 0.8235

performance across noise levels and image resolutions. For example, ZS-N2S and Self2Self degrade on high-resolution images, likely due to the limitations of training with masked pixels. ZS-N2N performs well on high-resolution images from Kodak24 but suffers on lower-resolution images in Set13 (192 \times 192), as it is trained to map between two downsampled versions of the same noisy image. In comparison, ZS-NCD maintains robust performance across different noise levels and image sizes. The more realistic case of not having access to the noise parameter α was discussed in Section 4. In both noise regimes, we use MSE as the loss function. However, for Poisson noise, minimizing the negative log-likelihood is also a natural choice. We defer the results using this loss to Appendix B.3.

Table 3: First 2 rows: Denoising performance (average PSNR / SSIM of 6 images) under AWGN $\mathcal{N}(0, \sigma^2 I)$ on Mouse Nucle fluorescence microscopy images (image size 128 \times 128). Noise levels are 10 and 20. Last row: Real camera denoising performance on camera image dataset: PolyU. The images are cropped into size of 512 \times 512. We report the average PSNR / SSIM of 6 random images.

σ	JPEG2K	BM3D	DIP	DD	ZS-N2N	ZS-N2S	S2S	ZS-NCD
10	32.89 / 0.8294	38.65 / 0.9640	36.43 / 0.8789	37.33 / 0.9533	36.17 / 0.9319	31.26 / 0.8812	12.63 / 0.2966	<u>38.23 / 0.9508</u>
20	28.57 / 0.6986	34.96 / 0.9296	32.32 / 0.7889	33.50 / 0.9092	32.25 / 0.8532	30.41 / 0.8600	10.09 / 0.1559	<u>34.71 / 0.9093</u>
Unknown	32.89 / 0.8294	<u>35.71 / 0.9506</u>	35.43 / 0.9408	34.83 / 0.9395	34.07 / 0.9028	23.61 / 0.8344	35.66 / <u>0.9527</u>	35.84 / 0.9534

Fluorescence microscopy and real camera images To evaluate performance in low-data and domain-shift settings, we test ZS-NCD on Mouse Nuclei fluorescence microscopy images [56], which differ significantly from natural images in structure and texture. We also assess real-world denoising using the PolyU dataset [57], which contains high-resolution images captured by Canon, Nikon, and Sony cameras. Ground-truth images are obtained by averaging multiple captures, while the noisy inputs are single-shot acquisitions. Results are shown in Table 3. ZS-NCD consistently outperforms other learning-based zero-shot denoisers, demonstrating robustness to unknown noise models and non-natural image distributions.

Table 4: Comparison of AWGN denoising of Conv and MLP based ZS-NCD on Set11.

	ZS-NCD	$\sigma = 25$	$\sigma = 50$
Conv		28.93 / 0.8079	25.58 / 0.7144
MLP		29.52 / 0.8363	25.89 / 0.7306




Figure 4: Denoising *Parrot* using ZS-NCD, where only a single pixel from each overlapping patch (stride 1) is retained after compression. (AWGN, $\sigma_z = 25$.) Each heatmap value indicates the PSNR achieved when denoising is based solely on the pixel at that specific location within each patch.

Robustness to overfitting. Most learning-based zero-shot methods are prone to overfitting due to the lack of clean targets and the use of overparameterized networks. In contrast, ZS-NCD, grounded in compression-based denoising theory, overcomes this issue given the entropy constraint. To further highlight this key aspect of ZS-NCD, we replace the convolutional encoder-decoder ($\approx 0.4M$ params) with a fully connected MLP ($\approx 2.3M$ params) and observe that, instead of degradation, the performance improves using the same λ (see Table 4).

Effect of overlapping patch aggregation. As described in Section 4 and illustrated in Fig. 1, ZS-NCD denoises each pixel by aggregating outputs from overlapping patches, where each patch is first compressed and then decompressed using a learned neural compression model. Intuitively, one might expect the most accurate reconstruction for a given pixel to come from the patch in which it lies at the center, as this location benefits from the largest available spatial context, which has been observed in [58].

This observation leads to the question: *Does averaging over overlapping reconstructions improve denoising quality, or would it suffice to use only the patch where pixel appears at a fixed position (e.g., the center)?* From a computational perspective, both strategies are equivalent, since in both methods every patch is processed, but in averaging scheme, each patch contributes to all the pixels it covers.

To investigate this, we conducted an ablation in which, instead of averaging, each pixel (i, j) is reconstructed solely from one of the $k \times k$ patches in which it appears, using a fixed location in the patch (e.g., top-left, center, etc.). The results are shown in Fig. 4, where each heatmap entry reports the PSNR obtained by using only that specific location in the patch for reconstruction. As expected, performance is best when the pixel is centrally located, and degrades as it moves toward the patch boundaries.

However, the key observation is that averaging across all overlapping reconstructions yields a substantial performance gain. For instance, in denoising *Parrot* (from Set11 dataset), the best single-location reconstruction achieves 25.90 dB (center), while averaging achieves 28.14 dB, a gain of over 2 dB. This highlights the denoising benefit of combining multiple noisy views of each pixel, consistent with principles from ensembling and variance reduction.

6 Conclusions

We have studied maximum likelihood compression-based denoising, and provided theoretical characterization of its performance under both AWGN and Poisson noise. Furthermore, we introduced ZS-NCD, a new zero-shot neural-compression-based denoising and demonstrated that it achieves state-of-the-art performance among zero-shot methods, in both AWGN and Poisson denoising.

The presented theoretical results are derived by assuming a fixed (e.g., pre-trained/defined) compression code. Extending these results to the case of zero-shot learned compression codes is an interesting direction for future research.

Acknowledgment

A.Z., X.C., S.J. were supported by NSF CCF-2237538.

References

- [1] Norbert Wiener. *Extrapolation, interpolation, and smoothing of stationary time series*. The MIT press, 1964.
- [2] David L Donoho and Iain M Johnstone. Ideal spatial adaptation by wavelet shrinkage. *biometrika*, 81(3):425–455, 1994.
- [3] Stéphane Mallat. *A wavelet tour of signal processing*. Elsevier, 1999.
- [4] D. L. Donoho. De-noising by soft-thresholding. *IEEE transactions on information theory*, 41(3):613–627, 2002.
- [5] Javier Portilla, Vasily Strela, Martin J Wainwright, and Eero P Simoncelli. Image denoising using scale mixtures of gaussians in the wavelet domain. *IEEE Transactions on Image processing*, 12(11):1338–1351, 2003.
- [6] Michael Elad and Michal Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image processing*, 15(12):3736–3745, 2006.
- [7] Stefan Roth and Michael J Black. Fields of experts. *International Journal of Computer Vision*, 82:205–229, 2009.
- [8] Shuhang Gu, Lei Zhang, Wangmeng Zuo, and Xiangchu Feng. Weighted nuclear norm minimization with application to image denoising. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2862–2869, 2014.
- [9] Shakarim Soltanayev and Se Young Chun. Training deep learning based denoisers without ground truth data. *Advances in neural information processing systems*, 31, 2018.
- [10] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2noise: Learning image restoration without clean data. *arXiv preprint arXiv:1803.04189*, 2018.
- [11] 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*, pages 2129–2137, 2019.
- [12] Joshua Batson and Loic Royer. Noise2self: Blind denoising by self-supervision. In *International conference on machine learning*, pages 524–533. PMLR, 2019.
- [13] Kwanyoung Kim and Jong Chul Ye. Noise2score: tweedie’s approach to self-supervised image denoising without clean images. *Advances in Neural Information Processing Systems*, 34:864–874, 2021.
- [14] David Leigh Donoho. *The kolmogorov sampler*. Department of Statistics, Stanford University, 2002.
- [15] Tsachy Weissman and Erik Ordentlich. The empirical distribution of rate-constrained source codes. *IEEE transactions on information theory*, 51(11):3718–3733, 2005.
- [16] Johannes Ballé, Philip A Chou, David Minnen, Saurabh Singh, Nick Johnston, Eirikur Agustsson, Sung Jin Hwang, and George Toderici. Nonlinear transform coding. *IEEE Journal of Selected Topics in Signal Processing*, 15(2):339–353, 2020.
- [17] Yibo Yang, Stephan Mandt, Lucas Theis, et al. An introduction to neural data compression. *Foundations and Trends® in Computer Graphics and Vision*, 15(2):113–200, 2023.
- [18] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8):2080–2095, 2007.
- [19] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5728–5739, 2022.

- [20] David S Taubman, Michael W Marcellin, and Majid Rabbani. Jpeg2000: Image compression fundamentals, standards and practice. *Journal of Electronic Imaging*, 11(2):286–287, 2002.
- [21] Ulyanov Dmitry, Andrea Vedaldi, and Lempitsky Victor. Deep image prior. *International Journal of Computer Vision*, 128(7):1867–1888, 2020.
- [22] Reinhard Heckel and Paul Hand. Deep decoder: Concise image representations from untrained non-convolutional networks. In *International Conference on Learning Representations*, 2019.
- [23] Yuhui Quan, Mingqin Chen, Tongyao Pang, and Hui Ji. Self2self with dropout: Learning self-supervised denoising from single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1890–1898, 2020.
- [24] Youssef Mansour and Reinhard Heckel. Zero-shot noise2noise: Efficient image denoising without any data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14018–14027, 2023.
- [25] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017.
- [26] Yaochen Xie, Zhengyang Wang, and Shuiwang Ji. Noise2same: Optimizing a self-supervised bound for image denoising. *Advances in neural information processing systems*, 33:20320–20330, 2020.
- [27] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9446–9454, 2018.
- [28] Jason Lequyer, Reuben Philip, Amit Sharma, Wen-Hsin Hsu, and Laurence Pelletier. A fast blind zero-shot denoiser. *Nature Machine Intelligence*, 4(11):953–963, 2022.
- [29] Jibo Bai, Daqi Zhu, and Mingzhi Chen. Dual-sampling noise2noise: Efficient single image denoising. *IEEE Transactions on Instrumentation and Measurement*, 2025.
- [30] Qing Ma, Junjun Jiang, Xiong Zhou, Pengwei Liang, Xianming Liu, and Jiayi Ma. Pixel2pixel: A pixelwise approach for zero-shot single image denoising. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.
- [31] Xiaoxiao Ma, Zhixiang Wei, Yi Jin, Pengyang Ling, Tianle Liu, Ben Wang, Junkang Dai, and Huaian Chen. Masked pre-training enables universal zero-shot denoiser. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [32] Johannes Ballé, Valero Laparra, and Eero P. Simoncelli. End-to-end optimized image compression. In *International Conference on Learning Representations*, 2017.
- [33] Lucas Theis, Wenzhe Shi, Andrew Cunningham, and Ferenc Huszár. Lossy image compression with compressive autoencoders. In *International Conference on Learning Representations*, 2017.
- [34] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. In *International Conference on Learning Representations*, 2018.
- [35] David Minnen, Johannes Ballé, and George D Toderici. Joint autoregressive and hierarchical priors for learned image compression. In *Advances in neural information processing systems*, volume 31, 2018.
- [36] Yin hao Zhu, Yang Yang, and Taco Cohen. Transformer-based transform coding. In *International conference on learning representations*, 2022.
- [37] Jinming Liu, Heming Sun, and Jiro Katto. Learned image compression with mixed transformer-cnn architectures. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14388–14397, 2023.

- [38] Eirikur Agustsson, David Minnen, Nick Johnston, Johannes Balle, Sung Jin Hwang, and George Toderici. Scale-space flow for end-to-end optimized video compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8503–8512, 2020.
- [39] Fabian Mentzer, George Toderici, David Minnen, Sergi Caelles, Sung Jin Hwang, Mario Lucic, and Eirikur Agustsson. VCT: A video compression transformer. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [40] Jiahao Li, Bin Li, and Yan Lu. Neural video compression with feature modulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26099–26108, 2024.
- [41] Michela Testolina, Evgeniy Upenik, and Touradj Ebrahimi. Towards image denoising in the latent space of learning-based compression. In *Applications of Digital Image Processing XLIV*, volume 11842, pages 412–422. SPIE, 2021.
- [42] Saeed Ranjbar Alvar, Mateen Ulhaq, Hyomin Choi, and Ivan V Bajić. Joint image compression and denoising via latent-space scalability. *Frontiers in Signal Processing*, 2:932873, 2022.
- [43] Benoit Brummer and Christophe De Vleeschouwer. On the importance of denoising when learning to compress images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2440–2448, 2023.
- [44] Yuxin Xie, Li Yu, Farhad Pakdaman, and Moncef Gabbouj. Joint end-to-end image compression and denoising: Leveraging contrastive learning and multi-scale self-onns. *arXiv preprint arXiv:2402.05582*, 2024.
- [45] Naoki Saito. Simultaneous noise suppression and signal compression using a library of orthonormal bases and the minimum description length criterion. In *Wavelet Analysis and Its Applications*, volume 4, pages 299–324. Elsevier, 1994.
- [46] S Grace Chang, Bin Yu, and Martin Vetterli. Image denoising via lossy compression and wavelet thresholding. In *Proceedings of International Conference on Image Processing*, volume 1, pages 604–607. IEEE, 1997.
- [47] S Grace Chang, Bin Yu, and Martin Vetterli. Adaptive wavelet thresholding for image denoising and compression. *IEEE transactions on image processing*, 9(9):1532–1546, 2000.
- [48] Balas K Natarajan. Filtering random noise from deterministic signals via data compression. *IEEE transactions on signal processing*, 43(11):2595–2605, 1995.
- [49] Léo Larigauderie, Michela Testolina, and Touradj Ebrahimi. On combining denoising with learning-based image decoding. In *Applications of Digital Image Processing XLV*, volume 12226, pages 193–206. SPIE, 2022.
- [50] Zhihao Li, Yufei Wang, Alex Kot, and Bihan Wen. Compress clean signal from noisy raw image: A self-supervised approach. In *Forty-first International Conference on Machine Learning*, 2024.
- [51] Ali Zafari, Xi Chen, and Shirin Jalali. Decompress: Denoising via neural compression. *arXiv preprint arXiv:2503.22015*, 2025.
- [52] Iain M. Johnstone. *Gaussian estimation: Sequence and wavelet models*. Unpublished Book, 2017.
- [53] Guangyong Chen, Fengyuan Zhu, and Pheng Ann Heng. An efficient statistical method for image noise level estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 477–485, 2015.
- [54] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *International conference on curves and surfaces*, pages 711–730. Springer, 2010.
- [55] Eastman Kodak. Kodak lossless true color image suite (PhotoCD PCD0992). URL <http://r0k.us/graphics/kodak>, 6:2, 1993.

- [56] Tim-Oliver Buchholz, Mangal Prakash, Deborah Schmidt, Alexander Krull, and Florian Jug. Denoiseg: joint denoising and segmentation. In *European Conference on Computer Vision*, pages 324–337. Springer, 2020.
- [57] Jun Xu, Hui Li, Zhetong Liang, David Zhang, and Lei Zhang. Real-world noisy image denoising: A new benchmark. *arXiv preprint arXiv:1804.02603*, 2018.
- [58] Dailan He, Yaoyan Zheng, Baocheng Sun, Yan Wang, and Hongwei Qin. Checkerboard context model for efficient learned image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14771–14780, 2021.

A Proofs

A.1 Auxiliary lemmas

Before stating the proofs of the mains theorems, here we state some lemmas that will be used later in the proofs.

Lemma 1. Assume that $0 < \alpha_m \leq \alpha_1, \alpha_2 \leq \alpha_M < \infty$. Then,

$$\left(\frac{\alpha_m}{\alpha_M}\right)^2 \frac{(\alpha_2 - \alpha_1)^2}{2\alpha_1} \leq D_{\text{KL}}(\text{Poisson}(\alpha_1) \parallel \text{Poisson}(\alpha_2)) \leq \left(\frac{\alpha_M}{\alpha_m}\right)^2 \frac{(\alpha_2 - \alpha_1)^2}{2\alpha_1}.$$

Lemma 2. Consider independent Poisson random variables Y_1, \dots, Y_n , where $Y_i \sim \text{Poisson}(\alpha_i)$. Consider $w_1, \dots, w_n \in \mathbb{R}$. Let $\sigma_n^2 = \sum_{i=1}^n w_i^2 \alpha_i$ and $w_M \triangleq \max_{i \in \{1, \dots, n\}} |w_i|$. Then, for any $t \in [0, \frac{3\sigma_n^2}{2w_M}]$,

$$\mathbb{P}\left(\sum_{i=1}^n w_i Y_i \geq \sum_{i=1}^n w_i \alpha_i + t\right) \leq \exp\left(-\frac{t^2}{2(\sigma_n^2 + w_M t/3)}\right). \quad (12)$$

and

$$\mathbb{P}\left(\sum_{i=1}^n w_i Y_i \leq \sum_{i=1}^n w_i \alpha_i - t\right) \leq \exp\left(-\frac{t^2}{2(\sigma_n^2 + w_M t/3)}\right). \quad (13)$$

A.1.1 Proof of Lemma 1

Proof.

$$D_{\text{KL}}(\text{Poisson}(\alpha_1) \parallel \text{Poisson}(\alpha_2)) = \alpha_2 - \alpha_1 + \alpha_1 \log \frac{\alpha_1}{\alpha_2} = \alpha_2 - \alpha_1 - \alpha_1 \log\left(1 + \frac{\alpha_2 - \alpha_1}{\alpha_1}\right). \quad (14)$$

Using the Taylor's theorem,

$$\log(1 + u) = u - f''(\alpha) \frac{u^2}{2}, \quad (15)$$

where $f(u) = \log(1 + u)$ and $\alpha \in (0, u)$. Note that

$$f''(u) = -\frac{1}{(1 + u)^2}.$$

Letting $u = \frac{\alpha_2 - \alpha_1}{\alpha_1}$, for $\alpha \in (0, u)$,

$$-\left(\frac{\alpha_M}{\alpha_m}\right)^2 \leq f''(\alpha) \leq -\left(\frac{\alpha_m}{\alpha_M}\right)^2. \quad (16)$$

Combining (14), (15) and (16) yields the desired result. \square

A.1.2 Proof of Lemma 2

Proof. Define

$$\mu_n = \mathbb{E}\left[\sum_{i=1}^n w_i Y_i\right] = \sum_{i=1}^n w_i \alpha_i,$$

and

$$\sigma_n^2 = \mathbb{E}\left[\left(\sum_{i=1}^n w_i (Y_i - \alpha_i)\right)^2\right] = \sum_{i=1}^n w_i^2 \alpha_i.$$

Consider $s > 0$, then using the Chernoff bound, we have

$$\begin{aligned} \mathbb{P}\left(\sum_{i=1}^n w_i Y_i \geq \sum_{i=1}^n w_i \alpha_i + t\right) &\leq \frac{\prod_{i=1}^n \mathbb{E}[\exp(sw_i(Y_i - \alpha_i))]}{\exp(st)} \\ &= \exp\left(\sum_{i=1}^n (\alpha_i(e^{sw_i} - 1 - sw_i)) - st\right). \end{aligned} \quad (17)$$

Note that for $u \in (-1, 1)$, $e^u - 1 - u \leq \frac{u^2}{2(1-u/3)}$. Assuming that $s \leq \frac{1}{w_M}$, then $|sw_i| < 1$, for all i . Therefore,

$$\begin{aligned} \exp\left(\sum_{i=1}^n (\alpha_i(e^{sw_i} - 1 - sw_i)) - st\right) &\leq \exp\left(\sum_{i=1}^n \alpha_i \left(\frac{(sw_i)^2}{2(1-sw_i/3)}\right) - st\right) \\ &\leq \exp\left(\sum_{i=1}^n \alpha_i \left(\frac{(sw_i)^2}{2(1-sw_M/3)}\right) - st\right) \\ &= \exp\left(\frac{s^2 \sigma_n^2}{2(1-sw_M/3)} - st\right). \end{aligned} \quad (18)$$

Evaluating this bound at $s = \frac{t}{\sigma_n^2 + w_M t/3}$, since $1 - sw_M/3 = \frac{\sigma_n^2}{\sigma_n^2 + w_M t/3}$, it follows that

$$\frac{s^2 \sigma_n^2}{2(1-sw_M/3)} - st = -\frac{st}{2} = -\frac{t^2}{2(\sigma_n^2 + w_M t/3)}. \quad (19)$$

To derive the other bound, we can follow the same steps and apply Chernoff bound as done in (17) to get

$$\mathbb{P}\left(\sum_{i=1}^n w_i Y_i \leq \mu_n - t\right) \leq \exp\left(\sum_{i=1}^n \alpha_i (e^{-sw_i} - 1 + sw_i) - st\right). \quad (20)$$

We now use the inequality for $u \in (-1, 1)$:

$$e^{-u} - 1 + u \leq \frac{u^2}{2(1+u/3)}.$$

Assume $s \leq \frac{1}{w_M}$ so that $s|w_i| \leq 1$ for all i . Then:

$$\begin{aligned} \sum_{i=1}^n \alpha_i (e^{-sw_i} - 1 + sw_i) &\leq \sum_{i=1}^n \alpha_i \cdot \frac{(sw_i)^2}{2(1+sw_i/3)} \\ &\leq \sum_{i=1}^n \alpha_i \cdot \frac{(sw_i)^2}{4/3} \\ &= \frac{s^2 \sigma_n^2}{4/3}. \end{aligned} \quad (21)$$

Hence,

$$\mathbb{P}\left(\sum_{i=1}^n w_i Y_i \leq \mu_n - t\right) \leq \exp\left(\frac{s^2 \sigma_n^2}{4/3} - st\right). \quad (22)$$

Setting $s = \frac{3t}{2\sigma_n^2}$, which satisfies $s \leq \frac{1}{w_M}$,

$$\frac{s^2 \sigma_n^2}{4/3} - st = -\frac{3t^2}{4\sigma_n^2} \leq -\frac{t^2}{2(\sigma_n^2 + w_M t/3)}. \quad (23)$$

□

A.2 Proof of Theorem 1

Proof. Recall that $\mathbf{y} = \mathbf{x} + \mathbf{z}$, with \mathbf{z} is i.i.d. $\mathcal{N}(0, \sigma_z^2)$, and

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{c} \in \mathcal{C}} \|\mathbf{c} - \mathbf{y}\|_2^2, \quad \tilde{\mathbf{x}} = \arg \min_{\mathbf{c} \in \mathcal{C}} \|\mathbf{c} - \mathbf{x}\|_2^2.$$

Since both $\hat{\mathbf{x}}, \tilde{\mathbf{x}}$ are in \mathcal{C} ,

$$\|\hat{\mathbf{x}} - \mathbf{y}\|_2^2 \leq \|\tilde{\mathbf{x}} - \mathbf{y}\|_2^2 \quad (24)$$

$$\|(\hat{\mathbf{x}} - \mathbf{x}) - \mathbf{z}\|_2^2 \leq \|(\tilde{\mathbf{x}} - \mathbf{x}) - \mathbf{z}\|_2^2 \quad (25)$$

$$\|\hat{\mathbf{x}} - \mathbf{x}\|^2 - 2\langle \mathbf{z}, \hat{\mathbf{x}} - \mathbf{x} \rangle + \|\mathbf{z}\|^2 \leq \|\tilde{\mathbf{x}} - \mathbf{x}\|^2 - 2\langle \mathbf{z}, \tilde{\mathbf{x}} - \mathbf{x} \rangle + \|\mathbf{z}\|^2 \quad (26)$$

$$\|\hat{\mathbf{x}} - \mathbf{x}\|^2 \leq \|\tilde{\mathbf{x}} - \mathbf{x}\|^2 - 2\langle \mathbf{z}, \tilde{\mathbf{x}} - \mathbf{x} \rangle + 2\langle \mathbf{z}, \hat{\mathbf{x}} - \mathbf{x} \rangle \quad (27)$$

$$\|\hat{\mathbf{x}} - \mathbf{x}\|^2 \leq \|\tilde{\mathbf{x}} - \mathbf{x}\|^2 + 2|\langle \mathbf{z}, \tilde{\mathbf{x}} - \mathbf{x} \rangle| + 2|\langle \mathbf{z}, \hat{\mathbf{x}} - \mathbf{x} \rangle|. \quad (28)$$

Let $\mathbf{e} = \hat{\mathbf{x}} - \mathbf{x}$ denote the error of the compression-based estimate of ground truth \mathbf{x} from its noisy version \mathbf{y} , and $\mathbf{d} = \tilde{\mathbf{x}} - \mathbf{x}$ denote the distortion from the compressing the ground truth \mathbf{x} with the compression code \mathcal{C} , then we have

$$\|\mathbf{e}\|^2 \leq \|\mathbf{d}\|^2 + 2|\langle \mathbf{z}, \mathbf{e} \rangle| + 2|\langle \mathbf{z}, \mathbf{d} \rangle| \quad (29)$$

$$= \|\mathbf{d}\|^2 + 2\|\mathbf{e}\| \left| \left\langle \mathbf{z}, \frac{\mathbf{e}}{\|\mathbf{e}\|} \right\rangle \right| + 2\|\mathbf{d}\| \left| \left\langle \mathbf{z}, \frac{\mathbf{d}}{\|\mathbf{d}\|} \right\rangle \right|. \quad (30)$$

For any possible reconstruction $\mathbf{c} \in \mathcal{C}$, we define error vector $\mathbf{e}^{(\mathbf{c})} = \mathbf{c} - \mathbf{x}$. Given $t_1, t_2 > 0$, define event \mathcal{E}_1 and \mathcal{E}_2 as

$$\mathcal{E}_1 = \left\{ \left| \sum_{i=1}^n z_i \frac{e_i^{(\mathbf{c})}}{\|\mathbf{e}^{(\mathbf{c})}\|} \right| \leq t_1 : \mathbf{c} \in \mathcal{C} \right\}. \quad (31)$$

and

$$\mathcal{E}_2 = \left\{ \left| \sum_{i=1}^n z_i \frac{e_i^{(\tilde{\mathbf{x}})}}{\|\mathbf{e}^{(\tilde{\mathbf{x}})}\|} \right| \leq t_2 \right\}, \quad (32)$$

respectively. Conditioned on $\mathcal{E}_1 \cap \mathcal{E}_2$, it follows from (30) that

$$\|\mathbf{e}\|^2 \leq \|\mathbf{d}\|^2 + 2t_1\|\mathbf{e}\| + 2t_2\|\mathbf{d}\|. \quad (33)$$

Therefore,

$$\|\mathbf{e}\|^2 - 2t_1\|\mathbf{e}\| + t_1^2 \leq \|\mathbf{d}\|^2 + 2t_2\|\mathbf{d}\| + t_2^2 + (t_1^2 - t_2^2), \quad (34)$$

$$|\|\mathbf{e}\| - t_1| \leq \sqrt{(\|\mathbf{d}\| + t_2)^2 + (t_1^2 - t_2^2)}, \quad (35)$$

and finally,

$$\|\mathbf{e}\|_2 \leq \|\mathbf{d}\|_2 + t_1 + t_2 + \sqrt{t_1^2 - t_2^2}, \quad (36)$$

where the last line follows because $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$, for all $a, b > 0$. To finish the proof we need to bound $\mathbb{P}((\mathcal{E}_1 \cap \mathcal{E}_2)^c)$ and set parameters t_1 and t_2 .

Note that for each \mathbf{c} , $\frac{\mathbf{e}^{(\mathbf{c})}}{\|\mathbf{e}^{(\mathbf{c})}\|}$ is a unit vector in \mathbb{R}^n . Therefore, $\sum_{i=1}^n z_i \frac{e_i^{(\mathbf{c})}}{\|\mathbf{e}^{(\mathbf{c})}\|} \sim \mathcal{N}(0, \sigma_z^2)$. Hence,

$$\mathbb{P}\left(\left| \sum_{i=1}^n z_i \frac{e_i^{(\mathbf{c})}}{\|\mathbf{e}^{(\mathbf{c})}\|} \right| \geq t\right) \leq 2 \exp\left(-\frac{t^2}{2\sigma_z^2}\right). \quad (37)$$

Therefore, applying the union bound and noting that $|\mathcal{C}| \leq 2^R$,

$$\mathbb{P}(\mathcal{E}_1^c) \leq 2^{R+1} \exp\left(-\frac{t_1^2}{2\sigma_z^2}\right), \quad (38)$$

and

$$\mathbb{P}(\mathcal{E}_2^c) \leq 2 \exp\left(-\frac{t_2^2}{2\sigma_z^2}\right). \quad (39)$$

For $\eta \in (0, 1)$, set

$$t_1 = \sigma_z \sqrt{2 \ln 2R(1 + \eta)},$$

and

$$t_2 = \sigma_z \sqrt{2 \ln 2R\eta}.$$

Then,

$$\mathbb{P}(\mathcal{E}_1^c \cup \mathcal{E}_2^c) \leq 2^{-\eta R+2}. \quad (40)$$

Using the selected values of t_1 and t_2 in (36) yields the desired result, i.e.,

$$\frac{1}{\sqrt{n}} \|\mathbf{x} - \hat{\mathbf{x}}\|_2 \leq \sqrt{\delta} + 2\sigma_z \sqrt{\frac{(2 \ln 2)R}{n}} (1 + 2\sqrt{\eta}),$$

where we have used the fact that $\frac{1}{n} \|\mathbf{d}\|_2^2 \leq \delta$. \square

A.3 Proof of Corollary 1

Proof. We start by designing a lossy compression code for the set of signals in \mathcal{Q}_n , defined as

$$\mathcal{Q}_n = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_0 \leq k, \|\mathbf{x}\|_\infty \leq 1\}. \quad (41)$$

For a k -sparse $\mathbf{x} \in \mathcal{Q}_n$, let $\mathbf{x}^{(k)} \in \mathbb{R}^k$ denote the k -dimensional vector derived from the non-zero coordinates of \mathbf{x} . We define a lossy compression code (f, g) that achieves distortion δ . Specifically, given a k -sparse $\mathbf{x} \in \mathcal{Q}_n$, the encoder operates as follows:

1. Encode the number of non-zero entries and their locations. This requires at most

$$\log_2 k + \log_2 \binom{n}{k} \leq \log_2 k + k \log_2 \left(\frac{ne}{k}\right)$$

bits.

2. Quantize the values of the non-zero coordinates such that overall distortion δ is achieved. To achieve this goal it quantizes each non-zero coordinate of \mathbf{x} into b bits. Let $[x_i]_b$ denote the b -bit quantized version of x_i . Then, $|x_i - [x_i]_b| \leq 2^{-b}$. Therefore, the overall ℓ_2 distortion can be bounded as

$$\|\mathbf{x} - g(f(\mathbf{x}))\|_2^2 \leq k 2^{-2b}.$$

Choosing

$$b = \left\lceil \frac{1}{2} \log_2 \frac{k}{n\delta} \right\rceil$$

ensures that $\|\mathbf{x} - g(f(\mathbf{x}))\|_2^2 \leq n\delta$.

It can be observed that overall the number of bits required for describing the signals in \mathcal{Q}_n within distortion δ can be bounded as

$$R \leq \frac{k}{2} \log_2 \left(\frac{k}{n\delta}\right) + k \log_2 \frac{n}{k} + c_k,$$

where $c_k = \log_2 k + k(\log_2 e + 1)$. Using the defined lossy compression code to solve (1) and applying Theorem 1, it follows that, with a probability larger than $1 - 2^{-\eta R+2}$,

$$\begin{aligned} \frac{1}{\sqrt{n}} \|\mathbf{x} - \hat{\mathbf{x}}\|_2 &\leq \sqrt{\delta} + \sigma_z C \sqrt{\frac{k}{2n} \log_2 \left(\frac{k}{n\delta}\right) + \frac{k}{n} \log_2 \left(\frac{n}{k}\right) + \frac{c_k}{n}} \\ &= \sqrt{\delta} + \sigma_z C \sqrt{\frac{k}{2n} \log_2 \left(\frac{n}{k\delta}\right) + \frac{c_k}{n}}, \end{aligned}$$

where $C = 2(1 + 2\sqrt{\eta})\sqrt{2\ln 2}$. Let

$$\delta = \frac{1}{n}.$$

Then,

$$\begin{aligned} \frac{1}{\sqrt{n}}\|\mathbf{x} - \hat{\mathbf{x}}\|_2 &\leq \frac{1}{\sqrt{n}} + \sigma_z C \sqrt{\frac{k \log_2 k}{2n} + \frac{k}{n} \log_2\left(\frac{n}{k}\right) + \frac{c_k}{n}} \\ &= \frac{1}{\sqrt{n}} + \sigma_z C \sqrt{\frac{k}{n} \log_2\left(\frac{n}{k}\right) + \gamma_n}, \end{aligned}$$

where

$$\gamma_n = \frac{k \log_2 k}{2n} + \frac{c_k}{n}.$$

Finally, since $R \geq k \log_2(\frac{n}{k})$,

$$1 - 2^{-\eta R + 2} \geq 1 - e^{-\eta k \log(\frac{n}{k})}.$$

□

A.4 Proof of Theorem 2

Proof. Recall that $\arg \min_{\mathbf{c} \in \mathcal{C}} \mathcal{L}(\mathbf{c}; \mathbf{y})$. Let

$$\tilde{\mathbf{x}} = \arg \min_{\mathbf{c} \in \mathcal{C}} \|\mathbf{x} - \mathbf{c}\|_2.$$

Since both $\hat{\mathbf{x}}$ and $\tilde{\mathbf{x}}$ are in \mathcal{C} , we have $\mathcal{L}(\hat{\mathbf{x}}; \mathbf{y}) \leq \mathcal{L}(\tilde{\mathbf{x}}; \mathbf{y})$, or

$$\sum_{i=1}^n (\alpha \hat{x}_i - y_i \log \hat{x}_i) \leq \sum_{i=1}^n (\alpha \tilde{x}_i - y_i \log \tilde{x}_i). \quad (42)$$

Given the input signal $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{c} \in \mathcal{C}$, let $\text{Poisson}(\alpha \mathbf{x})$ and $\text{Poisson}(\alpha \mathbf{c})$ denote the distributions corresponding to independent Poisson random variables with respective means αx_i and αc_i . Note that

$$D_{\text{KL}}(\text{Poisson}(\alpha \mathbf{x}) \parallel \text{Poisson}(\alpha \mathbf{c})) = \sum_{i=1}^n \left(\alpha(c_i - x_i) + \alpha x_i \log \frac{x_i}{c_i} \right). \quad (43)$$

Adding $\sum_i (-\alpha x_i + \alpha x_i \log x_i)$ to the both sides of (42), it follows that

$$\sum_{i=1}^n (\alpha(\hat{x}_i - x_i) - y_i \log \hat{x}_i + \alpha x_i \log x_i) \leq \sum_{i=1}^n (\alpha(\tilde{x}_i - x_i) - y_i \log \tilde{x}_i + \alpha x_i \log x_i), \quad (44)$$

or

$$\begin{aligned} &D_{\text{KL}}(\text{Poisson}(\alpha \mathbf{x}) \parallel \text{Poisson}(\alpha \hat{\mathbf{x}})) + \sum_{i=1}^n (\alpha x_i - y_i) \log \hat{x}_i \\ &\leq D_{\text{KL}}(\text{Poisson}(\alpha \mathbf{x}) \parallel \text{Poisson}(\alpha \tilde{\mathbf{x}})) + \sum_{i=1}^n (\alpha x_i - y_i) \log \tilde{x}_i. \end{aligned} \quad (45)$$

Given $t_1, t_2 > 0$, define events \mathcal{E}_1 and \mathcal{E}_2 as

$$\mathcal{E}_1 = \left\{ \sum_{i=1}^n (y_i - \alpha x_i) \log c_i \leq t_1 : \forall \mathbf{c} \in \mathcal{C} \right\} \quad (46)$$

and

$$\mathcal{E}_2 = \left\{ \sum_{i=1}^n (y_i - \alpha x_i) \log \tilde{x}_i \geq -t_2 \right\}, \quad (47)$$

respectively. Conditioned on $\mathcal{E}_1 \cap \mathcal{E}_2$,

$$D_{\text{KL}}(\text{Poisson}(\alpha \mathbf{x}) \parallel \text{Poisson}(\alpha \hat{\mathbf{x}})) \leq D_{\text{KL}}(\text{Poisson}(\alpha \mathbf{x}) \parallel \text{Poisson}(\alpha \tilde{\mathbf{x}})) + t_1 + t_2, \quad (48)$$

and consequently from Lemma 1,

$$\frac{x_{\min}^2}{x_{\max}^3} \alpha \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 \leq \frac{x_{\max}^2}{x_{\min}^3} \alpha \|\mathbf{x} - \tilde{\mathbf{x}}\|_2^2 + (t_1 + t_2). \quad (49)$$

To finish the proof, we bound $\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2)$ and set t_1 and t_2 .

To bound $\mathbb{P}(\mathcal{E}_1^c)$, we apply Lemma 2, where for each \mathbf{c} , we set $w_i(\mathbf{c}) = \log \frac{1}{c_i}$. Then,

$$\sigma^n(\mathbf{c}) = \sum_{i=1}^n (\log \frac{1}{c_i})^2 \alpha x_i \leq \alpha \beta^2 \sum_{i=1}^n x_i,$$

and

$$w_M = \max_i |w_i| \leq \beta,$$

where

$$\beta = \log\left(\frac{1}{x_{\min}}\right).$$

Therefore, using the union bound, it follows that

$$\begin{aligned} \mathbb{P}(\mathcal{E}_1^c) &\leq 2^R \exp\left(-\frac{t_1^2}{2(\alpha \beta^2 \sum_{i=1}^n x_i + \beta t_1/3)}\right) \\ &\leq 2^R \exp\left(-\frac{t_1^2}{2(n\alpha \beta^2 x_{\max} + \beta t_1/3)}\right). \end{aligned} \quad (50)$$

To bound $\mathbb{P}(\mathcal{E}_2^c)$, we again apply Lemma 2, with $w_i = \log \frac{1}{x_i}$, and derive

$$\begin{aligned} \mathbb{P}(\mathcal{E}_2^c) &\leq \exp\left(-\frac{t_2^2}{2(\alpha \beta^2 \sum_{i=1}^n x_i + \beta t_2/3)}\right) \\ &\leq \exp\left(-\frac{t_2^2}{2(n\alpha \beta^2 x_{\max} + \beta t_2/3)}\right). \end{aligned} \quad (51)$$

Setting t_1 and t_2 such that they are both smaller than $3n\alpha\beta$, and noting that $x_{\max} < 1$, we have

$$\mathbb{P}(\mathcal{E}_1^c \cap \mathcal{E}_2^c) \leq 2^R \exp\left(-\frac{t_1^2}{4n\alpha\beta^2}\right) + \exp\left(-\frac{t_2^2}{4n\alpha\beta^2}\right). \quad (52)$$

Choosing $t_1 = \beta \sqrt{\frac{4}{\ln 2} n R (1 + \eta) \alpha}$ and $t_2 = \beta \sqrt{\frac{4}{\ln 2} n R \eta \alpha}$, it follows that

$$\mathbb{P}(\mathcal{E}_1^c \cap \mathcal{E}_2^c) \leq 2^{-\eta R}. \quad (53)$$

□

A.5 Proof of Theorem 3

Proof. Recall that

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{c} \in \mathcal{C}} \|\mathbf{c} - \mathbf{y}/\alpha\|_2^2, \quad \tilde{\mathbf{x}} = \arg \min_{\mathbf{c} \in \mathcal{C}} \|\mathbf{c} - \mathbf{x}\|_2^2.$$

Following the similar setup as in Section A.2, we get

$$\begin{aligned} \|\hat{\mathbf{x}} - \mathbf{y}/\alpha\|^2 &\leq \|\tilde{\mathbf{x}} - \mathbf{y}/\alpha\|^2 \\ \|\hat{\mathbf{x}} - \mathbf{x} + \mathbf{x} - \mathbf{y}/\alpha\|^2 &\leq \|\tilde{\mathbf{x}} - \mathbf{x} + \mathbf{x} - \mathbf{y}/\alpha\|^2 \\ \|\hat{\mathbf{x}} - \mathbf{x}\|^2 - 2\langle \mathbf{y}/\alpha - \mathbf{x}, \hat{\mathbf{x}} - \mathbf{x} \rangle + \|\mathbf{y}/\alpha - \mathbf{x}\|^2 &\leq \|\tilde{\mathbf{x}} - \mathbf{x}\|^2 - 2\langle \mathbf{y}/\alpha - \mathbf{x}, \tilde{\mathbf{x}} - \mathbf{x} \rangle + \|\mathbf{y}/\alpha - \mathbf{x}\|^2 \\ \|\hat{\mathbf{x}} - \mathbf{x}\|^2 &\leq \|\tilde{\mathbf{x}} - \mathbf{x}\|^2 - 2\langle \mathbf{y}/\alpha - \mathbf{x}, \tilde{\mathbf{x}} - \mathbf{x} \rangle + 2\langle \mathbf{y}/\alpha - \mathbf{x}, \hat{\mathbf{x}} - \mathbf{x} \rangle \\ \|\hat{\mathbf{x}} - \mathbf{x}\|^2 &\leq \|\tilde{\mathbf{x}} - \mathbf{x}\|^2 + 2|\langle \mathbf{y}/\alpha - \mathbf{x}, \tilde{\mathbf{x}} - \mathbf{x} \rangle| + 2|\langle \mathbf{y}/\alpha - \mathbf{x}, \hat{\mathbf{x}} - \mathbf{x} \rangle|. \end{aligned}$$

Defining \mathbf{e} , \mathbf{d} and $\mathbf{e}^{(\mathbf{c})}$, $\mathbf{c} \in \mathcal{C}$, as done in the proof of Theorem 1, we have

$$\|\mathbf{e}\|^2 \leq \|\mathbf{d}\|^2 + 2|\langle \mathbf{y}/\alpha - \mathbf{x}, \mathbf{e} \rangle| + 2|\langle \mathbf{y}/\alpha - \mathbf{x}, \mathbf{d} \rangle|. \quad (54)$$

Define events

$$\mathcal{E}_1 = \left\{ \left| \sum_{i=1}^n (y_i - \alpha x_i) e_i^{(\mathbf{c})} \right| \leq t_1 : \mathbf{c} = 1, \dots, 2^R \right\}, \quad (55)$$

and

$$\mathcal{E}_2 = \left\{ \left| \sum_{i=1}^n (y_i - \alpha x_i) e_i^{(\tilde{\mathbf{x}})} \right| \leq t_2 \right\}. \quad (56)$$

Conditioned on $\mathcal{E}_1 \cap \mathcal{E}_2$, it follows that

$$\|\mathbf{e}\|^2 \leq \|\mathbf{d}\|^2 + 2(t_1 + t_2)/\alpha, \quad (57)$$

Using Lemma 2 with $y_i \sim \text{Poisson}(\alpha x_i)$ and $w_i = e_i^{(\mathbf{c})}$, it follows that

$$\mathbb{P} \left(\left| \sum_{i=1}^n (y_i - \alpha x_i) e_i^{(\mathbf{c})} \right| \geq t \right) \leq 2 \exp \left(-\frac{t^2}{2(\sigma_n^2 + w_M t/3)} \right), \quad (58)$$

where

$$\sigma_n^2 = \sum_{i=1}^n w_i^2 \alpha_i = \alpha \sum_{i=1}^n (e_i^{(\mathbf{c})})^2 x_i \leq n \alpha x_{\max}^3, \quad (59)$$

and

$$w_M = \max_{i \in \{1, \dots, n\}} |w_i| \leq x_{\max}.$$

Using the union bound and noting that $|\mathcal{C}| \leq 2^R$, we have

$$\mathbb{P}(\mathcal{E}_1^c) \leq 2^{R+1} \exp \left(-\frac{t_1^2}{2(n x_{\max}^3 \alpha + t_1 x_{\max}/3)} \right) \quad (60)$$

and

$$\mathbb{P}(\mathcal{E}_2^c) \leq 2 \exp \left(-\frac{t_2^2}{2(n x_{\max}^3 \alpha + t_2 x_{\max}/3)} \right). \quad (61)$$

Setting t_1 and t_2 such that they are both smaller than $3n x_{\max}^2 \alpha$, and noting that $x_{\max} < 1$, we have

$$\mathbb{P}(\mathcal{E}_1^c \cap \mathcal{E}_2^c) \leq 2^{R+1} \exp \left(-\frac{t_1^2}{4n\alpha} \right) + 2 \exp \left(-\frac{t_2^2}{4n\alpha} \right). \quad (62)$$

For $\eta \in (0, 1)$, set

$$t_1 = 2\sqrt{n(\ln 2)(1 + \eta)\alpha R},$$

and

$$t_2 = 2\sqrt{n\eta(\ln 2)\alpha R}.$$

Then,

$$\mathbb{P}(\mathcal{E}_1^c \cup \mathcal{E}_2^c) \leq 2^{-\eta R+2}. \quad (63)$$

Using the selected values of t_1 and t_2 in (57) yields the desired result, i.e.,

$$\begin{aligned} \|\mathbf{e}\|_2^2 &\leq \|\mathbf{d}\|_2^2 + 4\sqrt{n(\ln 2)(1 + \eta)\frac{R}{\alpha}} + 4\sqrt{n\eta(\ln 2)\frac{R}{\alpha}} \\ &\leq \|\mathbf{d}\|_2^2 + \sqrt{\frac{nR}{\alpha}} (4\sqrt{\ln 2}) (\sqrt{1 + \eta} + \sqrt{\eta} + 1). \end{aligned}$$

□




Figure 5: Neural compression network used for denoising. Conv and FC denote the convolutional and fully connected layer, respectively. GDN and ReLU are activation functions.

B Additional experiments and experimental settings

In this section, we provide the details of the networks structures and experimental settings. We also present more experiments for Poisson denoising using MLE and MSE loss functions with unknown noise level.

B.1 Network structure

For our experiments we used 3 convolutional layers in the encoder with 128 number of channels for the first two layers in the encoder (and the last two layers of decoder), see Figure 5. For color images we choose the number of channels in the last encoder (and first decoder) layer equals to 32, and for grayscale images equals to 16. The MLP-based network of the ablation study in Section 5 has 3 fully connected layers in the encoder with 1024 hidden units for the first two layers in the encoder (and the last two layers of decoder). The number of hidden units in the last encoder (and first decoder) layer equals to 16. As activation function we use GDN [32] for Conv network and ReLU for MLP.

B.2 More recent improvements on Zero-shot Noise2Noise

Zero-shot Noise2Noise (ZS-N2N) [24] learns a mapping between two downsampled noisy images extracted from the original image. This strategy of generating noisy samples was extended in two recent concurrent works. Dual-sampling Noise2Noise (DS-N2N) [29] identifies a key sub-optimality in ZS-N2N: the network learns to denoise at a lower resolution but is then directly applied to the larger, original resolution. To address this, the authors proposed to apply an additional bicubic upsampling to the downsampled images and train the network on both the low-resolution and the upsampled pairs, which boosts the performance of ZS-N2N.

In Pixel2Pixel [30], the authors identified the local sampling in ZS-N2N insufficient to break the spatial correlation of *real-world noise*. To address this issue, Pixel2Pixel first finds a “pixel bank” for each pixel based on the non-local similar patches. Then, by randomly choosing pixels from the bank, it generates multiple noisy pairs to be used as a pseudo-training dataset for zero-shot denoising.

We have compared the performance of both methods under AWGN and Poisson noise in Tables 5 and 6, respectively. Each paper’s official code was used to report the numbers in the tables. In AWGN, our compression-based denoiser ZS-NCD outperforms both methods while achieving the second-best performance in Poisson denoising. Pixel2Pixel performs well in Poisson denoising.

Table 5: Average PSNR(dB)/SSIM denoising performance on Kodak24 dataset under AWGN $\mathcal{N}(0, \sigma_z^2)$. Best results are in **bold**, second-best are underlined.

	$\sigma_z = 15$	$\sigma_z = 25$	$\sigma_z = 50$
ZS-N2N (2023)	<u>32.30 / 0.8650</u>	29.54 / 0.7798	25.82 / 0.6151
DS-N2N (2025)	32.31 / 0.8803	29.64 / 0.8044	25.42 / 0.6378
Pixel2Pixel (2025)	31.31 / 0.8707	<u>29.89 / 0.8098</u>	<u>26.55 / 0.6873</u>
ZS-NCD	33.18 / 0.9026	30.60 / 0.8144	27.89 / 0.7464

Table 6: Average PSNR(dB)/SSIM denoising performance on Kodak24 dataset under Poisson noise, Poisson(αx)/ α . Best results are in **bold**, second-best are underlined.

	$\alpha = 15$	$\alpha = 25$	$\alpha = 50$
ZS-N2N (2023)	26.80 / 0.6757	28.21 / 0.7374	30.13 / 0.8076
DS-N2N (2025)	27.29 / 0.7016	28.50 / 0.7540	30.27 / 0.8250
Pixel2Pixel (2025)	28.22 / 0.7390	29.26 / 0.7891	30.41 / 0.8372
ZS-NCD	<u>27.64 / 0.7432</u>	<u>28.77 / 0.7677</u>	30.60 / 0.8235

B.3 MSE and likelihood estimation under Poisson noise without knowing true α

We compare the MSE and MLE distortion for Poisson denoising in Table 7 using the estimated $\hat{\alpha}$ as explained in Section 4.

Table 7: Minimizing Poisson negative log-likelihood (NLL) vs. MSE with estimated $\hat{\alpha}$ for *Camera-man* image in Set11. PSNR / SSIM are reported here.

α	MSE (with estimated $\hat{\alpha}$)	NLL (with estimated $\hat{\alpha}$)
15	23.41 / 0.7554	23.13 / 0.7567
50	25.22 / 0.7961	24.88 / 0.7460

B.4 Study on factors in patch-wise compression affecting denoising

In this section, we explain the intuition behind why learning compression networks and denoising on overlapped patches is feasible. The centered pixels in the patches are better compressed as empirically observed in [58], thus they can provide better denoising performance. To study the contribution of each patch containing the single pixel to be denoised we design the experiment that, in the denoising phase, we denoise the overlapped patches, but only a single pixel at the same location from each patch is used to construct the final denoised image, instead of averaging all of them as in (9). We show the denoising performance of each pixel location in Figure 6. The PSNR at each pixel denotes the denoising performance of only using the specific pixel of each overlapped patches with stride 1. We can find that the boundary pixels give lower PSNR, which is consistent with previous research findings that the centered pixels are better compressed. Next, we analyze the effect of patch size in both learning and denoising phases. Given that scaler quantization is applied and the entropy model is learned on latent code of the patches, the compression performance on the latent code is affected by both the patch size and the number of downsampling operations in CNN-based encoder. We design the experiment that 3 downsampling operators are applied to patch size 8 and 16, where the latent code sizes are $1 \times 1 \times n_b$ and $2 \times 2 \times n_b$ respectively, where the denoising performance at each pixel location is in Figure 6 (Left) and (Middle), and if we increase the downsampling to 4 for patch size 16, which results in the latent code size to be $1 \times 1 \times n_b$, the denoising performance is in Figure 6 (Right). We find that spatial size of the latent code to be quantized matters given the scaler quantization limitation, the reconstructed output by the decoder will be restricted by the only correlated latent code as we can observe. Motivated from this, we perform the learning and denoising phases both patch-wise with proper networks structure, all pixels in each patch are used and the overlapped areas are averaged properly to reduce the variance of the compression-based estimates.

C Additional numerical results

In this section, we provide the full denoising numerical results of the denoisers on all the test images. All the experiments were run on Nvidia RTX 6000 Ada with 48 GB memory. It takes 40 minutes to denoise a grayscale image of size 256×256 , and 50 minutes for an RGB image of size 512×768 . Adam optimizer is used for training the networks over 20K steps, with initial learning rate of 5×10^{-3} decreased to 5×10^{-4} after 16K steps for the Conv-based network. The learning rate for MLP-based networks is 1×10^{-3} .




Figure 6: Denoising AWGN ($\sigma_z = 25$) of image *Parrot* by only compressing a single pixel in each overlapped patches with stride 1. The PSNR at each pixel denotes the final denoising performance by only compressing the pixel at that specific location in each patch. **Left:** patch size 8×8 , with downsampling factor equals 8 in f_{θ_1} ; **Middle:** patch size 16×16 , with downsampling factor equals 8 in f_{θ_1} ; **Right:** patch size 16×16 , with downsampling factor equals 16 in f_{θ_1} .

C.1 Set11 Dataset

For noise levels (15, 25, 50) we set $\lambda = (300, 850, 3000)$. Similar to Kodak and other experiments we set training epochs to have 20K steps of gradient back propagation. For Poisson denoising $\alpha = (15, 25, 50)$ the $\lambda = (3000, 1500, 1000)$. We report the detailed results of AWGN denoising in Table 8, and Poisson noise denoising in Table 9.

Table 8: Set11 Denoising performance comparison under AWGN $\mathcal{N}(0, \sigma_z^2 I)$.

σ	Method	256 \times 256								512 \times 512								Average
		C.man	House	Peppers	Starfish	Monarch	Airplane	Parrot		Barbara	Boats	Pirate	Couple					
15	BM3D	31.840/8974	34.940/8870	32.820/9118	31.160/9082	31.920/9409	31.090/9034	31.470/9032		33.040/9253	32.120/8604	31.940/8726	32.110/8795	32.220/8991				
	JPEG2K	27.120/7474	29.480/7621	27.960/7907	26.750/8077	26.740/8166	26.580/7664	27.300/7778		26.760/7690	27.870/7390	27.920/7471	27.440/7449	27.450/7699				
	DIP	27.940/7417	31.390/8111	29.800/8273	29.580/8605	29.930/8767	28.140/8047	28.370/7794		27.650/7538	29.480/7798	29.270/7817	28.650/7727	29.110/7990				
	DD	29.410/8099	32.830/8406	26.970/8488	29.390/8739	30.010/8957	26.440/8228	29.320/8447		24.480/7089	29.450/7883	29.780/8085	29.060/7938	28.830/8215				
	ZS-N2N	30.140/8133	32.190/8138	30.580/8264	29.520/8639	30.150/8551	29.980/8298	30.190/8290		27.700/7772	30.060/7900	30.060/7957	29.590/7913	30.010/8169				
	ZS-N2S	27.660/8272	31.080/8442	29.460/8675	28.830/8810	28.770/8961	27.340/8591	27.670/8528		28.750/8534	29.520/8139	29.410/8181	29.600/8311	28.920/8495				
	S2S	20.290/6769	32.960/8633	23.960/8387	25.500/8250	30.050/9269	28.100/8611	20.200/7132		30.350/8865	27.740/7871	29.970/8192	25.820/7754	26.810/8158				
	ZS-NCd	30.830/8554	34.450/8835	32.200/8844	31.340/8749	31.830/8966	30.070/8552	30.400/8464		31.140/8826	31.090/8014	30.820/8302	30.670/8279	31.350/8580				
	ZS-NCd (MLP)	31.180/8680	34.860/8887	32.430/9009	31.370/9053	32.040/9263	30.790/8848	31.020/8839		32.550/9123	31.820/8468	31.350/8552	31.550/8638	31.910/8851				
	BM3D	29.540/8499	32.790/8561	30.130/8705	28.580/8578	29.360/9046	28.500/8584	28.940/8561		30.630/8887	29.870/8039	29.640/8082	29.740/8206	29.790/8523				
	JPEG2K	24.490/6976	27.260/7269	24.930/7206	24.180/7167	24.060/7561	23.910/7126	24.380/7162		24.090/6825	25.610/6577	25.760/6566	25.280/6534	24.910/6997				
25	DIP	25.230/6043	28.930/7545	27.390/7579	26.390/7777	27.470/8169	25.570/6983	26.290/7409		24.750/6356	27.050/6843	27.060/6857	26.520/6847	26.600/7128				
	DD	27.240/7521	30.480/8023	25.390/7591	26.860/8051	27.690/8526	24.930/7120	27.290/7863		23.810/6455	27.490/7163	27.870/7380	27.130/7131	26.930/7530				
	ZS-N2N	27.320/7089	29.360/7276	27.460/7240	26.610/7821	27.200/7634	27.020/7463	27.160/7149		25.490/6854	27.260/6779	27.480/6931	26.630/6673	27.180/7173				
	ZS-N2S	26.240/7843	29.230/8073	27.770/8233	27.610/8463	27.350/8569	25.860/8023	26.270/7997		26.430/7759	28.230/7580	27.520/7526	27.740/7617	27.300/7971				
	S2S	16.930/5998	29.120/8275	21.880/7666	21.140/6974	25.930/8606	24.120/7350	17.090/6069		25.790/7980	23.940/7061	27.320/7403	25.290/6979	23.320/7306				
	ZS-NCd	28.780/8237	32.140/8547	29.620/8406	28.480/8134	29.020/8494	27.770/8126	28.140/8007		28.390/8192	28.850/7444	28.640/7630	28.370/7648	28.930/8079				
	ZS-NCd (MLP)	29.080/8259	32.630/8525	29.850/8574	28.730/8547	29.420/8861	28.370/8477	28.750/8431		30.010/8658	29.590/7843	29.150/7859	29.100/7958	29.520/8363				
	BM3D	26.560/7813	29.610/8029	26.850/7911	25.070/7508	25.820/8192	25.290/7713	26.020/7809		27.020/7888	26.760/7003	26.750/6962	26.370/6977	26.560/7619				
	JPEG2K	21.490/5880	24.240/6444	21.720/6077	21.390/5784	20.860/6414	21.110/6021	21.290/6035		21.650/5377	22.830/5318	23.290/5356	22.670/5025	22.050/5794				
	DIP	22.730/5846	25.670/6475	23.810/5987	22.990/6406	23.060/6293	22.640/5522	23.020/5811		22.380/5316	23.900/5371	24.430/5524	23.400/5064	23.460/5783				
	DD	23.890/6487	27.270/7282	22.950/7276	23.440/6700	23.550/7319	22.520/6652	23.870/6471		22.720/5980	24.470/6050	25.190/6340	24.300/5872	24.010/6584				
50	ZS-N2N	23.360/5174	25.170/5167	22.860/5669	22.920/6101	22.930/5988	22.870/5136	22.620/5150		23.930/5138	24.300/5318	25.300/4930	23.520/4545	23.520/4545				
	ZS-N2S	24.650/6966	26.720/7091	25.240/7297	24.050/7012	24.820/7618	24.040/7467	24.000/7078		22.810/5916	25.460/6512	25.550/6469	24.800/6197	24.740/6883				
	S2S	14.230/4809	21.140/6396	17.800/5763	15.710/4176	18.330/5955	15.700/4828	13.660/4446		17.600/4883	18.690/5264	19.550/5354	19.120/5325	17.410/5200				
	ZS-NCd	25.550/7616	28.620/7995	26.310/7604	24.590/6925	25.530/7585	24.650/7338	25.310/7228		24.060/6525	25.610/6538	25.920/6707	25.190/6519	25.580/7144				
	ZS-NCd (MLP)	25.110/7342	29.130/7881	26.300/7702	24.850/7426	25.120/7795	24.840/7497	25.550/7146		26.260/6698	26.250/6717	25.650/6560	25.890/7366					

Table 9: Set11 Denoising performance comparison under Poisson noise $\text{Poisson}(\alpha x)/\alpha$.

		256 × 256										512 × 512									
α	Method	C.man	House	Peppers	Starfish	Monarch	Airplane	Parrot		Barbara	Boats	Pirate	Couple	Average							
15	BM3D	26.640/7651	29.390/7668	27.130/7914	24.930/7519	26.320/8265	24.790/6730	26.260/7866		27.240/7860	26.820/6977	27.070/7048	26.670/7056	26.660/7505							
	JPEG2K	21.980/6032	24.350/6106	22.120/6213	21.520/5887	21.240/6493	20.870/5818	22.020/6378		21.940/5688	23.090/5346	23.750/5510	23.010/5237	22.350/5882							
	DIP	22.850/5382	26.320/6528	24.230/6138	23.230/6696	23.540/6875	22.070/4938	22.810/5723		22.590/5503	24.180/5533	24.950/5811	23.830/5362	23.690/5863							
	DD	24.450/6261	27.590/7453	23.200/7269	23.860/7164	24.660/7286	22.220/6055	24.380/6830		22.890/6081	24.640/6023	25.610/6516	24.580/5986	24.370/6629							
	ZS-N2N	24.190/5818	25.410/5346	24.650/6016	23.120/6520	23.920/6441	23.120/5565	23.830/5821		23.050/5503	24.400/5403	24.870/5684	23.940/5305	24.040/5766							
	ZS-N2S	24.940/7241	27.290/7317	25.710/7431	24.410/7417	25.370/7968	23.050/7315	22.870/6087		26.080/6696	25.940/6655	25.090/6389	25.060/7051								
	S2S	23.530/7325	22.010/7409	22.730/7390	18.200/6010	21.810/7813	16.180/5010	20.270/7304		22.100/7261	23.490/6529	24.720/6782	24.170/6843	21.750/6872							
	ZS-NCd	25.730/7660	28.870/8015	26.540/7745	24.650/6988	25.860/7791	24.210/6568	25.520/7356		24.110/6562	25.480/6510	25.930/6705	25.290/6552	25.650/7132							
	BM3D	22.690/5154	22.820/4765	22.740/5930	22.110/6947	23.440/7213	19.600/3788	23.050/5991		23.090/6508	22.770/5123	23.890/5979	23.450/5753	22.700/5741							
	JPEG2K	22.540/6267	24.970/6773	22.870/6076	22.260/6378	22.550/6641	21.590/5649	22.620/6373		22.550/5976	23.710/5685	24.200/5801	23.490/5566	23.030/6108							
25	DIP	24.210/5976	27.060/6553	25.760/6945	24.410/7312	25.210/7384	23.580/6290	24.690/6608		23.110/5903	25.240/6130	26.100/6528	24.910/5998	24.940/6512							
	DD	25.590/6695	28.470/7606	24.200/7348	25.140/7667	26.160/8022	23.240/6306	25.890/7289		23.270/6257	25.840/6517	26.760/6961	25.770/6573	25.480/7022							
	ZS-N2N	25.540/6334	27.140/6234	25.820/6522	24.330/7158	25.510/7109	24.530/6274	25.550/6617		24.090/6173	25.600/6018	26.110/6354	25.160/5958	25.400/6432							
	ZS-N2S	26.220/7776	27.810/7643	26.550/7768	24.820/7795	26.480/8254	24.770/7463	25.440/7839		23.240/6387	27.250/7143	27.130/7200	26.440/6986	26.010/7478							
	S2S	25.090/7572	24.100/7398	24.910/7733	19.110/6491	23.640/8226	17.930/6279	21.130/7692		24.010/7860	25.300/7058	26.450/7232	25.770/7360	23.400/7355							
	ZS-NCd	27.170/7635	30.090/8109	27.920/7932	26.270/7600	27.280/8093	24.930/6116	26.740/7551		26.240/7393	27.300/6993	27.320/7192	26.830/7123	27.100/7431							
	BM3D	22.940/5314	22.890/4548	23.220/5844	23.060/7150	23.870/6990	20.510/4115	23.650/6136		23.540/6545	22.820/5205	23.970/6087	23.470/5723	23.090/5787							
	JPEG2K	24.230/6635	26.870/6796	24.960/7042	24.080/7240	24.050/7568	23.400/6387	24.570/7077		23.950/6804	25.370/6414	25.730/6483	25.250/6475	24.770/6811							
	DIP	25.340/6348	28.880/7369	27.590/7559	26.180/7891	26.580/7728	24.690/6457	26.040/7062		23.880/6158	26.610/6712	27.210/7015	26.340/6747	26.300/7084							
	DD	27.240/7398	30.160/7784	25.440/7615	26.780/8127	27.820/8527	23.390/6568	27.300/7735		23.860/6518	27.000/7072	27.160/7162	26.870/7455	26.870/7455							
50	ZS-N2N	27.630/7210	29.300/7113	27.920/7424	26.350/7857	27.380/7723	26.260/707	27.680/7443		25.500/6883	27.310/6832	27.680/7075	26.840/7083	27.260/7126							
	ZS-N2S	26.820/8041	29.320/7832	27.750/8192	26.620/8243	26.830/8624	26.050/8097	26.560/8196		25.600/6606	27.920/7499	27.680/7557	25.740/7521	27.080/7855							
	S2S	26.720/8022	27.190/8016	27.830/8300	26.470/7126	26.380/8780	21.090/6607	22.610/8001		27.070/8457	27.390/7620	28.310/7574	27.640/7885	27.500/7896							
	ZS-NCd	28.240/8093	31.900/8488	29.440/8410	28.020/8064	28.380/8504	25.090/7422	27.930/8061		26.700/7920	28.160/7273	28.110/7466	27.720/7456	28.440/7919							

C.2 Set13 Dataset

All images are center-cropped at size of 192×192 . For this set of images we set $\lambda = (100, 200, 800)$ and for noise levels $\sigma_z = (15, 25, 50)$ and for Poisson denoising we have $\lambda = (900, 500, 200)$ for noise levels $\alpha = (15, 25, 50)$. We report the detailed results of AWGN denoising in Table 10, and Poisson noise denoising in Table 11.

Table 10: Set13 Denoising performance comparison under AWGN $\mathcal{N}(0, \sigma_z^2 I)$.

σ	Method	Baboon	Barbara	Bridge	Coastguard	Comic	Face	Flowers	Foreman	Man	Monarch	Peppers	PPT3	Zebra	Average
15	BM3D	28.560.7797	33.070.9151	30.390.8723	30.180.8799	28.740.9289	30.280.7665	29.620.9094	35.830.9369	29.880.8323	31.130.9361	31.770.8199	34.490.9588	31.020.9198	31.150.8808
	JPEG2K	25.120.6539	27.010.7674	26.270.7426	26.270.7426	24.860.8017	27.350.6482	25.520.7986	30.860.8397	25.970.6983	25.630.7649	28.270.7161	28.000.8183	26.330.8243	26.690.7543
	DIP	27.250.7498	30.920.8403	30.180.8692	30.790.9036	28.250.9091	29.620.7611	28.850.8317	33.810.8947	29.900.8333	31.180.9085	30.300.7782	32.220.8908	30.630.8203	30.310.8570
	DD	26.350.7029	24.270.7066	29.160.8508	28.800.8421	26.440.8932	29.590.7398	27.340.8634	34.870.9274	28.680.8073	30.020.9151	31.070.8003	33.090.9277	30.220.9052	29.220.8371
	ZS-N2N	28.630.7992	28.450.7803	32.080.9041	31.540.9141	28.700.9018	30.640.8048	29.670.8955	34.020.8817	31.630.8801	31.650.9120	31.280.8089	32.840.9009	31.750.9277	30.950.8701
	ZS-N2S	29.920.5844	31.140.5730	31.270.6468	30.780.4893	15.800.4987	22.030.5641	17.490.5154	8.430.3637	22.690.6725	12.210.6124	21.600.6438	11.580.4462	20.370.7871	18.180.5690
	S2S	22.360.5810	30.390.8769	22.740.7485	22.720.7108	17.440.7015	17.230.4383	14.100.7121	16.780.8102	15.650.4463	26.920.8838	24.200.7398	14.840.7528	14.840.5400	20.610.6879
	ZS-NCd	28.100.7831	33.850.9208	31.490.9051	32.650.9345	29.230.9358	30.510.7891	29.690.9077	35.850.9381	31.490.8867	32.210.9448	31.600.8281	35.070.9601	32.370.9488	31.930.8983
25	BM3D	26.560.6892	30.700.8824	28.060.7955	27.590.7664	25.900.8698	28.870.6978	26.930.8396	33.510.9087	27.350.7428	28.720.9056	30.110.7823	31.660.9324	28.610.8643	28.810.8213
	JPEG2K	23.980.5711	24.140.6801	24.140.6400	23.530.5801	21.360.7014	26.170.5692	22.770.6809	28.300.7960	23.510.5667	23.030.7101	26.470.6729	25.230.7707	23.580.7389	24.320.6676
	DIP	25.700.6734	27.270.7021	27.790.7999	27.860.8126	25.110.8382	28.160.6599	26.030.8164	31.170.8384	27.270.7446	29.020.8785	28.820.7313	29.200.8106	28.650.8820	27.850.7837
	DD	25.560.6589	23.560.6676	26.940.7758	26.770.7542	24.760.8397	28.370.6819	25.720.8058	32.220.8849	27.090.7366	27.680.8566	29.420.7615	29.940.8937	28.110.8641	27.400.7832
	ZS-N2N	26.850.7287	26.730.7136	29.000.8282	28.710.7250	26.670.8248	31.310.8156	28.880.7940	29.030.8683	28.880.7940	29.030.8683	28.950.7282	29.870.8085	28.540.8768	28.360.8001
	ZS-N2S	19.220.4603	15.440.3992	22.050.5839	16.640.5734	25.760.6423	17.980.6324	20.900.6792	22.300.6170	21.290.7778	22.450.6535	17.040.6502	22.410.7964	20.600.6200	20.300.6200
	S2S	18.660.4947	24.970.7969	19.430.6054	20.690.5890	16.300.6168	15.080.3846	17.570.5267	15.530.7605	14.110.3892	21.340.7849	22.400.6795	13.600.6805	13.690.4890	17.950.5998
	ZS-NCd	26.540.7128	30.650.8388	28.780.8332	29.460.8718	26.830.8847	29.140.7367	27.050.8393	32.900.8870	28.830.8027	30.180.8917	29.620.7656	31.560.8898	29.630.9023	29.330.8351
50	BM3D	24.660.5983	29.890.7854	25.310.6606	24.450.5345	22.300.7299	27.190.6140	23.520.7068	29.840.8400	24.770.6272	25.620.8492	27.340.7139	27.510.8657	25.900.7613	25.780.7134
	JPEG2K	21.990.4807	21.610.5511	21.550.4898	21.410.3641	18.500.5254	23.430.4603	19.920.5254	24.410.6873	20.740.4169	19.950.5870	23.660.5810	21.260.6490	19.810.5670	21.430.5295
	DIP	24.070.5778	23.060.5943	25.050.6773	24.100.6179	21.640.7142	26.070.5804	22.840.6871	28.340.7754	24.600.6173	25.680.7922	26.060.6398	25.600.7600	25.550.7926	24.620.6748
	DD	23.870.5745	23.060.6436	24.010.6402	23.430.4732	21.370.6956	26.830.6127	23.760.6819	29.120.8434	24.080.6011	24.080.6011	26.820.7055	25.400.8137	24.410.7659	24.560.6779
	ZS-N2N	24.410.5999	23.990.5573	25.190.6726	24.450.6513	21.870.7014	25.900.5732	22.720.6619	27.100.6419	24.910.6111	25.110.7571	25.350.5654	25.140.6243	24.620.7594	24.670.6444
	ZS-N2S	22.320.5704	16.650.5047	22.290.6475	17.750.4859	15.720.3989	25.490.6110	18.560.6121	24.160.7812	22.500.5681	19.830.6154	19.540.5440	16.260.4869	20.000.7724	20.620.5880
	S2S	14.080.3567	17.560.5046	14.950.3650	12.520.3204	13.550.3265	12.730.2916	13.310.2712	13.820.2515	12.160.2727	14.590.5444	17.160.4959	11.420.2887	12.160.2807	14.210.3938
	ZS-NCd	24.320.6035	26.840.7333	25.700.7172	25.760.7385	23.140.7768	26.900.6401	23.740.7212	28.460.7716	25.460.6710	26.290.8075	26.400.6722	27.250.7856	26.100.8113	25.870.7269

Table 11: Set13 Denoising performance comparison under Poisson noise $\text{Poisson}(\alpha x)/\alpha$.

α	Method	Baboon	Barbara	Bridge	Coastguard	Comic	Face	Flowers	Foreman	Man	Monarch	Peppers	PPT3	Zebra	Average
15	BM3D	24.460.5651	26.630.7712	25.180.6462	24.450.5186	22.070.7101	27.460.6342	23.510.7017	29.050.7747	24.980.6405	25.600.8101	27.530.7041	26.160.7288	26.190.7802	25.640.6912
	JPEG2K	22.010.4873	21.820.5515	21.770.5163	21.690.3727	18.470.5237	24.970.5289	20.420.5841	23.830.6966	21.690.4723	20.210.5689	24.170.6012	20.370.5858	21.040.6381	21.760.5494
	DIP	24.300.5907	23.050.6010	25.460.7069	24.210.5996	21.790.7282	26.990.6150	23.410.7456	28.030.7686	25.100.6444	26.340.8136	26.260.6510	26.440.7088	26.370.8180	25.510.6916
	DD	23.890.5750	23.160.6473	24.830.6513	23.800.5544	21.520.7107	27.220.6387	23.410.7239	25.250.8522	25.030.6144	24.670.7455	27.090.7156	25.430.8022	24.960.7006	25.470.6878
	ZS-N2N	24.620.6330	24.110.5768	25.820.7281	25.350.6946	21.900.7055	27.110.6779	23.670.7490	27.200.6491	26.600.7290	25.990.7740	26.010.6047	25.020.6064	26.130.8180	25.370.6878
	ZS-N2S	21.390.5390	17.460.4084	22.230.6428	21.830.5714	17.530.5711	25.140.6103	17.970.5082	24.330.7854	22.940.5959	21.150.7121	21.530.5467	16.610.6122	22.090.7725	21.230.6506
	S2S	16.580.5042	17.560.5046	18.070.6269	22.140.6206	15.010.5319	21.480.6548	18.870.6736	14.800.7562	24.780.7188	17.730.7807	21.510.6626	12.610.6130	22.090.7728	19.230.6553
	ZS-NCd	24.660.5907	27.630.7844	26.430.7443	26.440.7436	23.170.7781	27.530.6510	24.300.7680	29.580.8302	26.150.6997	26.460.7991	27.540.7196	26.680.7217	27.450.8440	26.440.7434
25	BM3D	20.610.4799	21.750.5321	22.930.6748	22.040.5620	19.940.6637	24.880.6358	22.350.7387	22.340.5167	23.110.5790	21.990.6799	22.750.5533	20.010.4929	23.480.7582	22.170.5992
	JPEG2K	22.590.5060	22.380.5742	22.810.5904	22.120.4455	19.590.6087	25.300.5533	21.510.6444	25.190.7343	22.100.4841	21.520.6226	24.780.6181	22.070.6515	22.470.6277	22.650.5952
	DIP	24.850.6243	24.020.5860	26.630.7646	25.410.6907	22.750.7640	27.420.6386	24.610.7935	29.420.8138	26.070.7119	27.020.8225	27.360.6839	26.860.7348	27.250.8467	26.130.7289
	DD	24.510.6017	23.300.6625	25.780.7389	24.960.6226	22.800.7680	28.180.6808	24.650.7872	30.340.8723	25.910.6824	26.310.8306	28.070.7242	26.920.7823	26.800.8309	26.040.7373
	ZS-N2N	25.810.6876	25.110.6387	27.170.7874	26.650.7562	23.450.7635	28.280.7277	24.950.8033	28.680.7031	28.100.7878	27.670.8276	27.470.6713	26.900.6864	27.430.8511	26.750.7455
	ZS-N2S	19.290.4409	14.360.5393	21.650.6879	20.440.4223	16.560.4988	25.410.6417	15.880.5870	25.290.8037	19.350.6591	21.090.6962	24.710.6801	22.470.7677	22.000.7806	21.190.6312
	S2S	17.400.5115	24.100.7441	18.560.6458	22.990.6967	15.480.5679	25.070.6826	19.740.7193	15.060.7385	26.120.7550	18.860.8094	23.260.6959	12.850.6302	22.890.8084	20.180.6927
	ZS-NCd	25.230.6145	29.090.8189	27.600.7969	27.660.8053	24.900.8223	28.310.6908	25.650.8231	30.990.8571	27.210.7491	27.820.8309	28.400.7464	27.980.7490	28.420.8712	27.600.8272
50	BM3D	21.780.5516	22.440.5350	23.850.7114	22.660.6103	21.360.7094	25.600.6670	23.510.7921	23.080.4721	23.860.6477	22.820.6871	23.240.5494	21.350.4511	24.070.7814	23.000.6281
	JPEG2K	23.830.5610	23.920.6565	24.260.6555	23.530.5775	21.090.6866	26.460.6041	22.880.7171	27.560.7817	23.770.5921	23.000.7123	26.680.6728	22.000.7220	24.150.6932	24.250.6696
	DIP	25.670.6702	26.620.6890	28.120.8173	27.790.8136	24.290.8131	28.630.7149	26.300.8449	30.970.8395	27.290.7409	29.100.8779	28.240.7107	28.640.7849	28.740.8816	27.720.7845
	DD	25.390.6439	25.620.6720	26.690.7441	24.400.8285	28.690.7060	26.080.8358	32.040.8838	27.450.7575	27.940.8613	29.450.7641	29.370.8636	28.330.8740	28.470.7867	27.420.7867
	ZS-N2N	27.060.7422	26.440.7151	29.330.8555	28.720.8361	25.470.8308	29.570.7822	26.980.8566	30.770.7918	30.020.8491	29.500.8730	29.030.7330	29.210.7853	29.340.8908	28.570.8112
	ZS-N2S	19.680.5011	19.800.4843	22.870.6603	19.480.4374	16.440.4783	18.210.4958	18.060.5203	25.220.8178	22.220.6118	21.950.7117	22.600.6847	19.840.7204	23.220.7586	20.750.6033
	S2S	12.600.4495	17.560.5046	18.070.6269	22.140.6206	15.010.5319	21.480.6548	18.870.6736	14.800.7562	24.780.7188	17.730.7807	21.510.6626	12.610.6130	22.090.7728	19.230.6553
	ZS-NCd	26.050.6622	30.150.8172	29.250.8571	29.830.8761	25.990.8641	26.300.7458	27.520.8790	32.510.8739	29.050.8147	27.200.8500	29.640.7627	24.530.6675	29.030.9081	29.000.8223

Table 13: Kodak24 Denoising performance performance under Poisson noise $\text{Poisson}(\alpha x)/\alpha$.

	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	Average
BM3D (c)	28.82/0.417	28.15/0.437	28.46/0.404	27.43/0.387	28.07/0.320	28.60/0.403	27.78/0.514	28.50/0.310	28.08/0.470	27.10/0.474	27.70/0.401	27.43/0.450	28.03/0.371	27.55/0.302	28.06/0.371	27.93/0.400	27.90/0.307	27.83/0.349	27.83/0.370	27.83/0.370	27.83/0.370	27.83/0.370	27.83/0.370	27.83/0.370	27.83/0.370
BM3D (f)	28.18/0.384	28.14/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	28.10/0.387	
DIP (c)	28.26/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	28.10/0.424	
DIP (f)	28.25/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	28.10/0.417	
ZS-N2N (c)	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	
ZS-N2N (f)	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	28.17/0.370	
S2S (c)	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	
S2S (f)	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	
ZS-NCD (c)	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	
ZS-NCD (f)	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	28.10/0.407	

C.4 Microscopy Mouse Nuclei Dataset

For these images with noise level $\sigma_z = (10, 20)$ we set $\lambda = (200, 600)$, we train the networks for 20K steps to obtain the results. We report the detailed denoising performance in Table 14 and 15 respectively.

Table 14: Denoising performance under AWGN $\mathcal{N}(0, \sigma^2 I)$ on fluorescence microscopy dataset: Mouse Nuclei. Images are cropped into 128×128 . Noise level $\sigma_z = 10$.

#	JPEG2K	BM3D	DIP	DD	ZS-N2N	ZS-N2S	S2S	ZS-NCD
1	32.90/0.7954	38.88/0.9631	37.31/0.8973	37.73/0.9464	36.37/0.9356	34.70/0.9410	10.88/0.1687	39.03/0.9556
2	32.32/0.8300	37.53/0.9613	35.93/0.8909	36.46/0.9560	35.26/0.9345	28.78/0.8504	13.08/0.4000	36.83/0.9546
3	32.97/0.8584	38.43/0.9690	36.17/0.8482	37.03/0.9631	35.86/0.9405	31.53/0.9307	12.76/0.3374	37.81/0.9634
4	32.57/0.8418	38.05/0.9605	35.82/0.9107	36.70/0.9478	34.86/0.9066	32.13/0.8688	14.42/0.3639	37.51/0.9303
5	34.54/0.7646	41.53/0.9596	38.09/0.8268	40.02/0.9438	39.30/0.9252	29.75/0.7976	10.22/0.1161	40.93/0.9420
6	32.02/0.8860	37.49/0.9703	35.24/0.8997	36.05/0.9628	35.38/0.9491	30.63/0.8989	14.42/0.3931	37.26/0.9588
Average	32.89/0.8294	38.65/0.9640	36.43/0.8789	37.33/0.9533	36.17/0.9319	31.26/0.8812	12.63/0.2966	<u>38.23/0.9508</u>

Table 15: Denoising performance under AWGN $\mathcal{N}(0, \sigma^2 I)$ on fluorescence microscopy dataset: Mouse Nuclei. Images are cropped into 128×128 . Noise level $\sigma_z = 20$.

#	JPEG2K	BM3D	DIP	DD	ZS-N2N	ZS-N2S	S2S	ZS-NCD
1	28.37/0.6337	35.10/0.9211	33.09/0.8485	33.70/0.8938	32.32/0.8609	32.59/0.8763	9.30/0.0240	34.98/0.8843
2	27.97/0.7255	33.80/0.9410	31.41/0.7986	32.39/0.9239	31.07/0.8421	28.42/0.8328	10.73/0.2383	33.75/0.9172
3	28.42/0.7121	34.45/0.9352	31.47/0.7642	32.64/0.9096	31.63/0.8660	31.08/0.9003	10.12/0.1807	34.20/0.9216
4	29.31/0.7557	34.30/0.9245	31.02/0.7598	32.71/0.9008	31.12/0.8168	30.60/0.8551	11.33/0.1947	33.87/0.8947
5	29.62/0.5932	38.50/0.9158	35.45/0.7763	37.18/0.9068	35.90/0.8585	32.89/0.8640	8.23/0.0650	37.70/0.9137
6	27.71/0.7713	33.61/0.9399	31.48/0.7860	32.41/0.9206	31.43/0.8750	26.87/0.8312	10.83/0.2328	33.72/0.9245
Average	28.57/0.6986	34.96/0.9296	32.32/0.7889	33.50/0.9092	32.25/0.8532	30.41/0.8600	10.09/0.1559	<u>34.71/0.9093</u>

C.5 Real Camera Noise Dataset PolyU

For these images with unknown noise model/level $\lambda = 25$. Also for BM3D the best performance was achieved with setting $\sigma_{\text{BM3D}} = 15$. We report the detailed denoising performance in Table 16.

Table 16: Real camera denoising performance on camera image dataset: PolyU. The dataset includes photos taken from 3 brands of cameras. Randomly selected 6 images are cropped into 512×512 .

Models	C.plug11	C.bike10	N.flower1	N.plant10	S.plant13	S.door10	Average
JPEG2K	36.26 / 0.9615	34.23 / 0.9371	33.55 / 0.9194	36.74 / 0.9157	30.39 / 0.9001	34.84 / 0.9012	34.33 / 0.9225
BM3D	<u>37.15</u> / 0.9758	34.85 / 0.9615	35.81 / 0.9504	<u>38.40 / 0.9410</u>	31.65 / <u>0.9465</u>	36.43 / 0.9285	35.71 / 0.9506
DIP	37.62 / 0.9724	<u>34.85</u> / 0.9534	34.93 / 0.9396	<u>37.64</u> / 0.9256	31.50 / 0.9396	36.02 / 0.9145	35.43 / 0.9408
DD	36.79 / 0.9722	34.73 / 0.9566	34.85 / 0.9366	37.84 / 0.9327	30.91 / 0.9305	33.88 / 0.9084	34.83 / 0.9395
ZS-N2N	36.30 / 0.9621	33.18 / 0.8853	33.28 / 0.8974	36.21 / 0.8862	30.57 / 0.9052	34.89 / 0.8804	34.07 / 0.9028
ZS-N2S	22.76 / 0.9119	20.36 / 0.8133	25.20 / 0.8670	33.63 / 0.8920	21.33 / 0.8256	18.39 / 0.6966	23.61 / 0.8344
S2S	37.75 / 0.9765	33.56 / 0.9545	<u>35.78</u> / 0.9537	38.30 / 0.9398	31.93 / 0.9483	36.65 / 0.9433	35.66 / 0.9527
ZS-NCD	36.99 / <u>0.9763</u>	34.79 / <u>0.9586</u>	35.43 / 0.9489	38.65 / 0.9449	<u>31.79</u> / 0.9464	37.42 / 0.9451	35.84 / 0.9534

D Visual Comparisons

In this section, we provide more visualization comparison of the zero-shot denoisers. The reconstruction PSNR and SSIM are above the images.

D.1 Kodak24

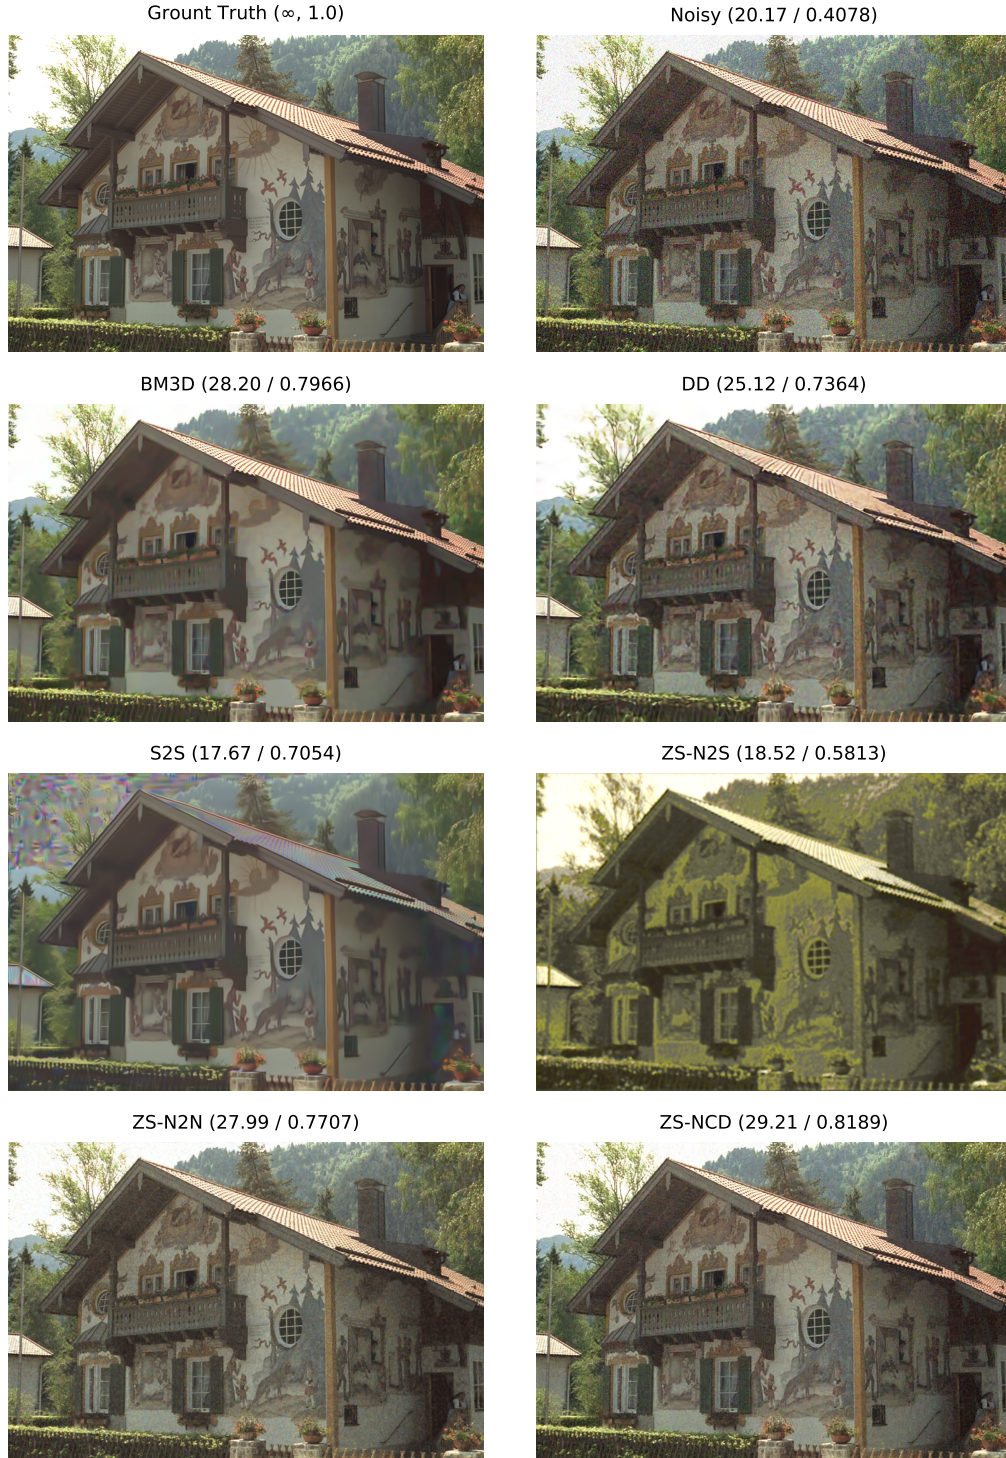


Figure 7: Kodim24 under additive white Gaussian noise ($\sigma_z = 25$).

Ground Truth (∞ , 1.0)



Noisy (17.80 / 0.3365)



BM3D (23.97 / 0.6702)



DD (24.57 / 0.7007)



S2S (18.05 / 0.7221)



ZS-N2S (14.82 / 0.5742)



ZS-N2N (26.39 / 0.7239)



ZS-NCD (27.82 / 0.8012)



Figure 8: *Kodim24* under Poisson noise ($\alpha = 25$).

D.2 Mouse Nuclei




Figure 9: Mouse nuclei reconstruction comparison under additive white Gaussian noise ($\sigma_z = 20$).

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: [Yes]

Justification:

- We claim that our proposed method is zero-shot, the details of how to crop patches from the single given noisy image can be found in Figure 1 and Section 4.
- We provide the theory connects compression and denoising, which can be found in Section 3, we characterize the performance for both AWGN and Poisson noise.
- We claim that the proposed ZS-NCD is the state-of-the-art zero-shot denoiser over both natural and non-natural images, this can be found empirically in Table 2 and 3.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We have a section 6 which includes the discussion on the limitation of our work.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren’t acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification:

- We assume that the compression code for \mathcal{Q} is defined by an encoder-decoder pair (f, g) characterized by R and δ , these appear in the theoretical results we provide.
- The proof of all the proposed theorems and corollary are provided in Appendix [A](#).

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification:

- The neural networks structures we used are provided in Appendix [B](#). The learning parameters including learning rate and optimizer are also included. The hyperparameter λ for each noise level are introduced in Appendix [C](#).
- The code of our proposed method is included in the supplementary.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.

- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification:

- We provide the implementation code of our method in ZS-NCD. We provide the instruction and comments in how to run the code. We also provide the requirements of the environments to run the code.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide the networks structures, hyperparameter λ , learning rate and optimizer in training the neural networks.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No] .

Justification: Our experiments follow standard practice in image denoising, where evaluation is performed on fixed noisy images and results are reported using metrics such as PSNR.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide the information of computing resources and the running time in Appendix C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: Our research on zero-shot neural compression based denoising conducted in the paper conform with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.

- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: The paper proposes a novel way of denoising. Zero-shot denoiser has great potential to be used for imaging systems without ground truth and large amount of noisy images, such as microcopy and medical imaging. Providing the good denoised version of the acquired noisy images is beneficial to the downstream tasks.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We have cited properly about the datasets we used in the paper.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: Our paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in our research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.