Zero-shot Denoising via Neural Compression Theoretical and algorithmic framework

Ali Zafari* Xi Chen* Shirin Jalali

Department of Electrical and Computer Engineering
Rutgers University

NeurIPS 2025 (Spotlight)

RUTGERS
NEW BRUNSWICK

NEURAL INFORMATION
PROCESSING SYSTEMS

Denoising

- Decades-old fundamental problem in signal processing and machine learning:
  - Recover $\boldsymbol{x} = (x_1, \ldots, x_n)$ from noisy observations $\boldsymbol{y} = (y_1, \ldots, y_n)$, where $\boldsymbol{y} \sim \prod_{i=1}^{n} p(y_i \mid x_i)$

Denoising

- Decades-old fundamental problem in signal processing and machine learning:
  - Recover $\boldsymbol{x} = (x_1, \ldots, x_n)$ from noisy observations $\boldsymbol{y} = (y_1, \ldots, y_n)$, where $\boldsymbol{y} \sim \prod_{i=1}^{n} p(y_i \mid x_i)$
- Deep learning has led to a renewed interest in denoising
  - Learning-based denoising achieves state-of-the-art performance across a wide range of denoising tasks

Zero-shot denoising

- Learning-based denoisers require large datasets of paired (noisy, clean) images for training.
- In many imaging domains (e.g., astronomy, biomedical):
  - data acquisition is expensive;
  - no direct access to ground truth.

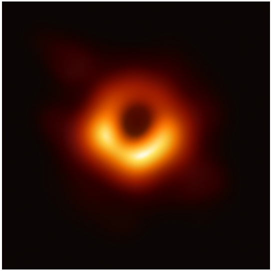

Black Hole Imaging (EHT)

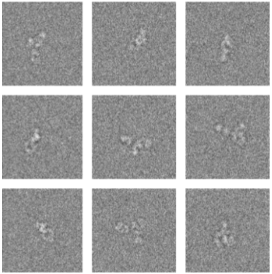

Cryo-EM

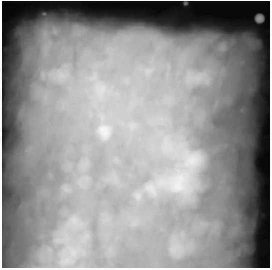

X-ray Nanotomography

Literature

Zero-shot denoising $\equiv$ Learning based denoising without access to any separate training dataset

- Zero-shot denoisers
  - Neural net as regularizer
    Deep Image Prior [Dmitry et al., 2020], Deep Decoder [Heckel et al., 2019]
  - Masked pixel training
    Zero-shot Noise2Self [Batson et al., 2019], Self2Self [Quan et al., 2020]
  - Pseudo-generated noisy pairs
    Zero-shot Noise2Noise [Mansour et al., 2023], Pixel2Pixel [Ma et al., 2025]

Literature

Zero-shot denoising $\equiv$ Learning based denoising without access to any separate training dataset

- Zero-shot denoisers
  - Neural net as regularizer
    Deep Image Prior [Dmitry et al., 2020], Deep Decoder [Heckel et al., 2019]
  - Masked pixel training
    Zero-shot Noise2Self [Batson et al., 2019], Self2Self [Quan et al., 2020]
  - Pseudo-generated noisy pairs
    Zero-shot Noise2Noise [Mansour et al., 2023], Pixel2Pixel [Ma et al., 2025]

- Drawbacks
  - Sub-optimal performance
  - Overfitting issue in neural net regularization methods
  - Discarding most useful info in mask-based approaches
  - Resolution mismatch in pseudo-generated pairs method

Compression-based denoising

Lossy compression of noisy signal at a distortion adjusted to noise power $\equiv$ denoising

- Theoretical foundations of compression-based denoising:
  - Minimum Kolmogorov Complexity Estimator [Donoho, 2002]
  - Universal denoising of discrete-valued sources [Weissman et al., 2005]

Compression-based denoising

Lossy compression of noisy signal at a distortion adjusted to noise power	$\equiv$	denoising

- Theoretical foundations of compression-based denoising:
  - Minimum Kolmogorov Complexity Estimator [Donoho, 2002]
  - Universal denoising of discrete-valued sources [Weissman et al., 2005]
- Despite theoretical foundations, it has not led to practical denoising methods.

Compression-based denoising

Lossy compression of noisy signal at a distortion adjusted to noise power $\equiv$ denoising

- Theoretical foundations of compression-based denoising:
  - Minimum Kolmogorov Complexity Estimator [Donoho, 2002]
  - Universal denoising of discrete-valued sources [Weissman et al., 2005]
- Despite theoretical foundations, it has not led to practical denoising methods.
- Emergence of neural compression encourages us to revisit the compression-based denoising.

Proposed method: ZS-NCD

Zero Shot – Neural Compression Denoiser (ZS-NCD)

- *Training phase.* Train a neural compression on patches (e.g., 8×8) extracted from noisy image y

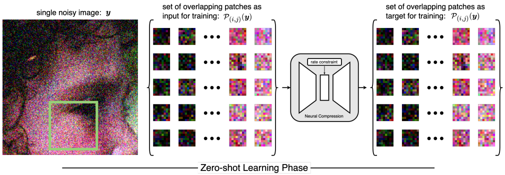

- *Denoising phase.* Independently compress patches and reconstruct by averaging overlaps

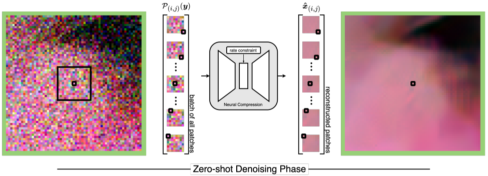

Our contributions

Theoretical contributions:

- Finite-sample error bounds for compression-based denoising
  - Additive white Gaussian Noise (AWGN)
  - Signal-dependent Poisson noise
  - Matching the minimax rate for k-sparse vectors under AWGN

Our contributions

Theoretical contributions:

- Finite-sample error bounds for compression-based denoising
  - Additive white Gaussian Noise (AWGN)
  - Signal-dependent Poisson noise
  - Matching the minimax rate for k-sparse vectors under AWGN

Algorithmic contributions:

- Introduced ZS-NCD that achieves state-of-the-art performance in both AWGN and Poisson denoising with synthesized and real-world noise.

ZS-NCD vs. Other zero-shot denoisers

ZS-NCD:

- Enjoys extended training iterations (robust to overfitting)
- Achieves state-of-the-art performance
- Maintains superior performance under distribution shifts

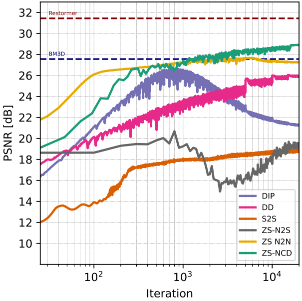

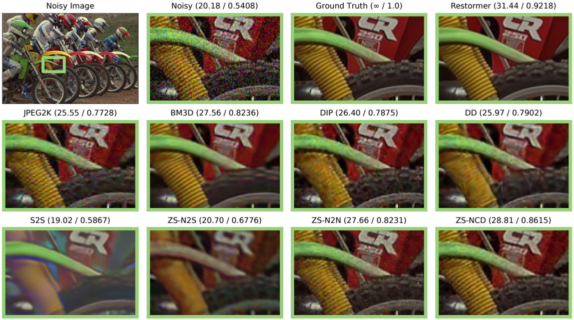

Extensive experiments on (non-)natural images

Noise Parameter		AWGN, $\mathcal{N}(0,\sigma^2)$			Poisson, $\mathrm{Poisson}(\alpha\boldsymbol{x})/\alpha$		
σ or α	Method	Set11	Set13	Kodak24	Set11	Set13	Kodak24
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**	28.81 / 0.8213	29.98 / 0.8092	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 / 0.7674
	ZS-N2S	27.30 / 0.7971	20.39 / 0.6200	20.89 / 0.6156	26.01 / 0.7478	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	26.75 / 0.7455	28.21 / 0.7374
	ZS-NCD	28.93 / 0.8079	**29.33 / 0.8351**	**30.60 / 0.8144**	**27.10 / 0.7431**	**27.60 / 0.7827**	**28.77 / 0.7677**

more noise levels can be found in the paper.

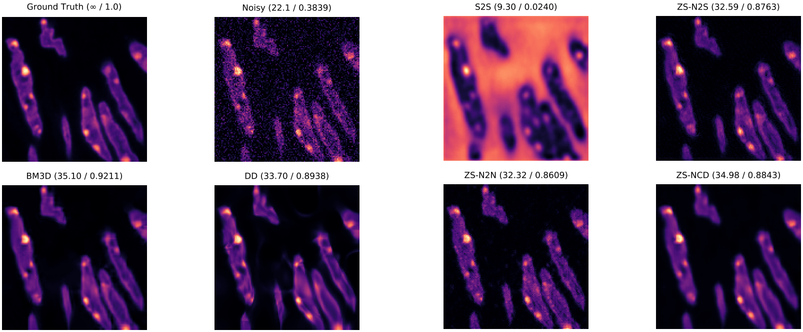

Takeaways

- Zero-shot denoising via neural compression provides an effective denoising solution, especially in applications with no paired training data.
- Opens a new research direction: Applying neural compression to a wide range of inverse problems.

github.com/Computational-Imaging-RU/ZS-NCDenoiser