Unsupervised Denoising Requires Unsupervised Metrics

Carlos Fernandez-Granda

www.cims.nyu.edu/~cfgranda

10/18/2022

Acknowledgements

This work was supported by NSF grants OAC-1940097, OAC-2103936, and NRT-1922658

Joint work with

Matan Leibovich, Adria Marcos-Morales, Sreyas Mohan (NYU)

Peter Crozier, Piyush Haluai, Mai Tan, Joshua Vincent (ASU)

Motivation: Studying catalysis

90% of all manufactured goods involve catalytic processes somewhere in their production chain

Considerable impact in energy, healthcare (pharmaceuticals), new material (polymers), transport, and the environment (water, air-quality, renewable and bio-produced materials)

Motivation: Studying catalysis

90% of all manufactured goods involve catalytic processes somewhere in their production chain

Considerable impact in energy, healthcare (pharmaceuticals), new material (polymers), transport, and the environment (water, air-quality, renewable and bio-produced materials)

To understand catalysis we need to see what is going on

Motivation: Studying catalysis

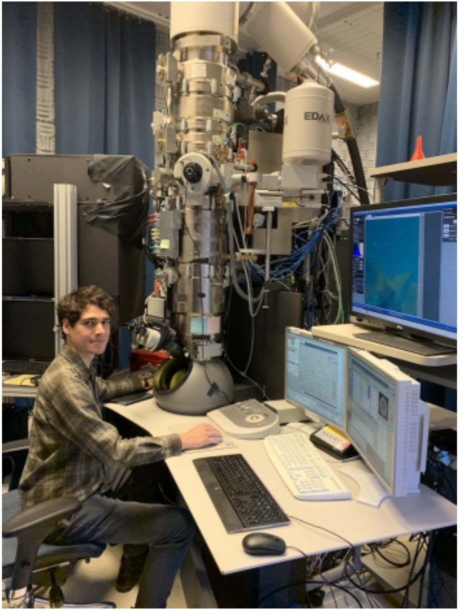

Electron microscope image

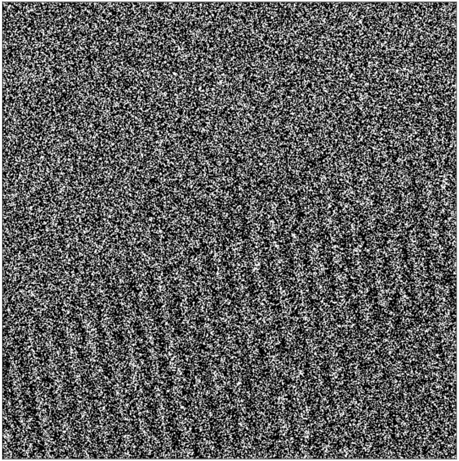

Electron microscope image

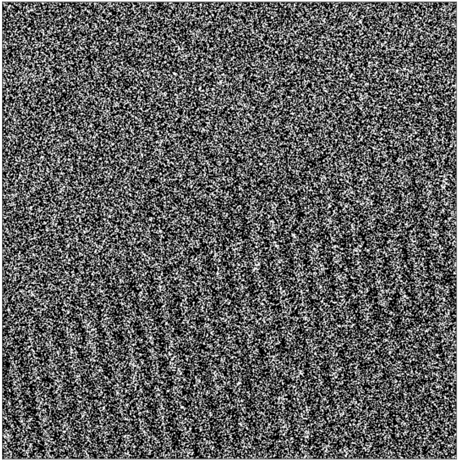

We need to denoise

Electron microscope image

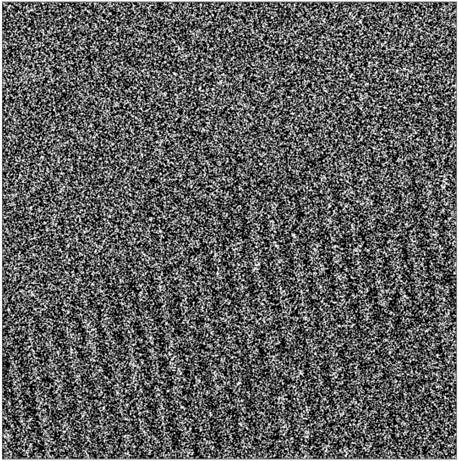

We need to denoise and know how well we are denoising!

Plan

Denoising via deep learning

Unsupervised denoising

Unsupervised metrics

Denoising via deep learning

Unsupervised denoising

Unsupervised metrics

The denoising problem

Estimate this

The denoising problem

Estimate this

From this

Convolutional estimation

Challenge: There are many pixels! (at least 10^4, often 10^6)

Convolutional estimation

Challenge: There are many pixels! (at least 10^4, often 10^6)

Solution: Exploit translation-invariant statistics

Covariance for dataset of natural images:

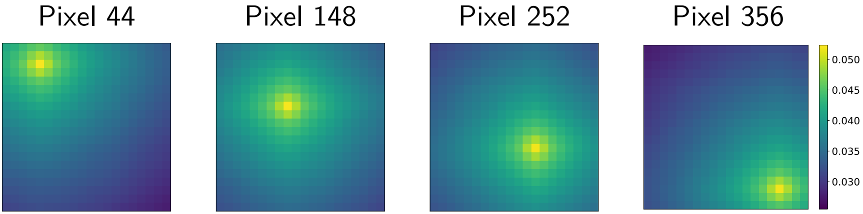

Convolutional filter

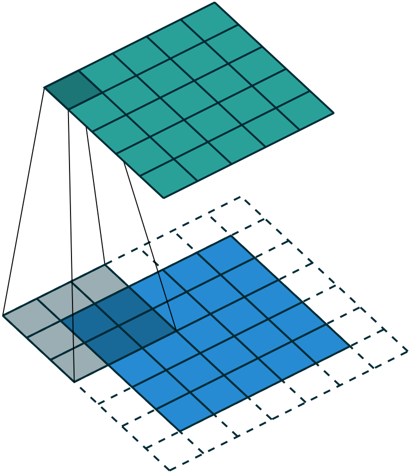

Image from *A guide to convolution arithmetic for deep learning*, Dumoulin & Visin, 2016.

Convolutional filter

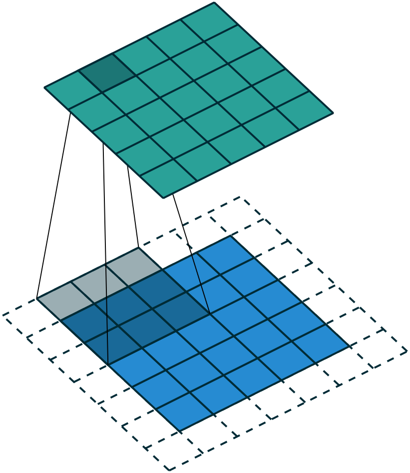

Image from *A guide to convolution arithmetic for deep learning*, Dumoulin & Visin, 2016.

Convolutional filter

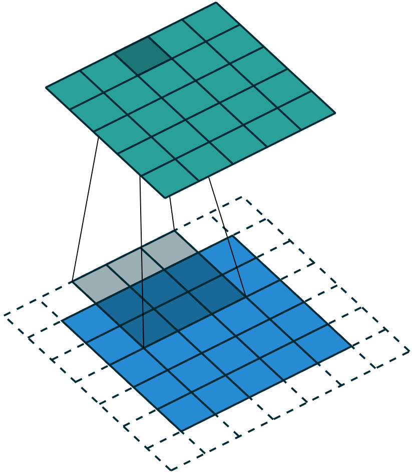

Image from *A guide to convolution arithmetic for deep learning*, Dumoulin & Visin, 2016.

Convolutional filter

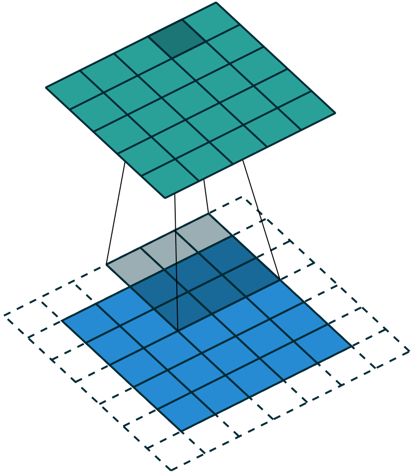

Image from *A guide to convolution arithmetic for deep learning*, Dumoulin & Visin, 2016.

Cost function

Supervised mean squared error

$\Bigg\|$  $-$  $\Bigg\|^2$

Denoised

Clean

Linear estimate (low noise level)

Example noisy image

Learned weights

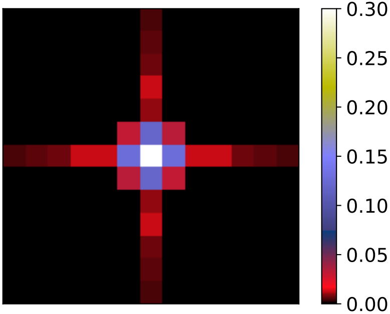

Linear estimate (medium noise level)

Example noisy image

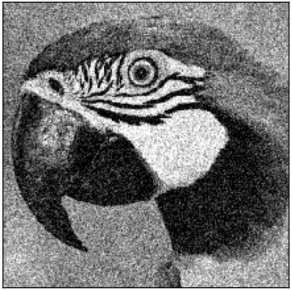

Learned weights

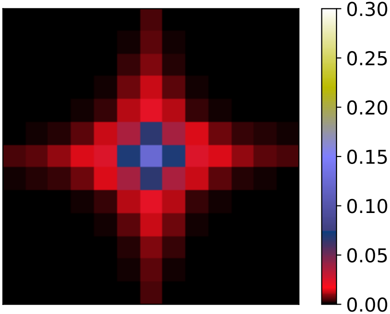

Linear estimate (high noise level)

Example noisy image

Learned weights

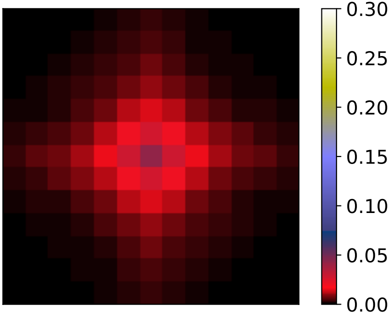

Limitations of linearity

Problem: Same estimate for each pixel

Limitations of linearity

Problem: Same estimate for each pixel

Blurs edges and other features

Limitations of linearity

Problem: Same estimate for each pixel

Blurs edges and other features

Pre-deep-learning solutions:

Adapt filter locally (e.g. bilateral filter [Tomasi and Manduchi 1998, Milanfar 2013])

Design/learn sparsifying transforms (wavelets, dictionary learning)

Results on electron microscopy

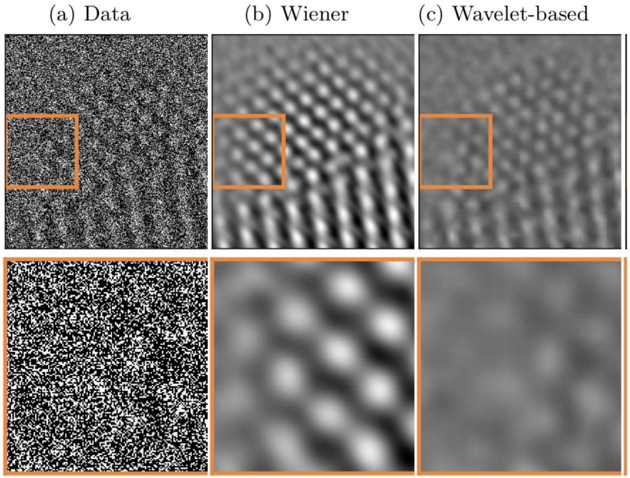

Deep-learning solution

Learn nonlinear convolutional model

Deep learning for image denoising

Denoising Convolutional Neural Network (DnCNN)[1]

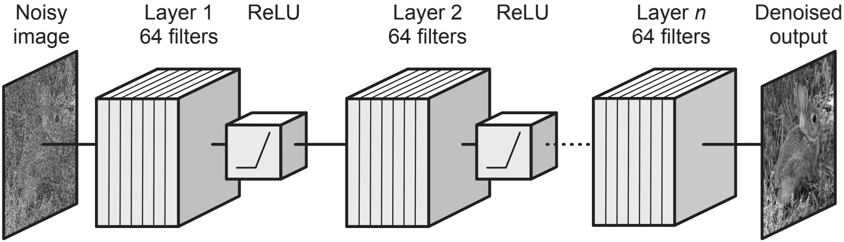

[1] *Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising*. K. Zhang, W. Zuo, Y. Chen, D. Meng, L. Zhang. IEEE Transactions in Image Processing (2017)

Deep learning for image denoising

- Gather dataset of natural images

Deep learning for image denoising

- Gather dataset of natural images
- Add synthetic Gaussian noise to generate noisy images

$+$

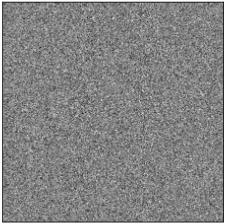

$=$

Deep learning for image denoising

- Gather dataset of natural images
- Add synthetic Gaussian noise to generate noisy images
- Train CNN to estimate clean image minimizing mean squared error

Works very well (state of the art)

Test image

CNN

Why?

Gradient of output pixels with respect to input image reveals learned function[2]

[2] *Robust and interpretable blind image denoising via bias-free convolutional neural networks*. S. Mohan, Z. Kadkhodaie, E. Simoncelli, C. Fernandez-Granda. ICLR 2020

Low noise

Noisy image

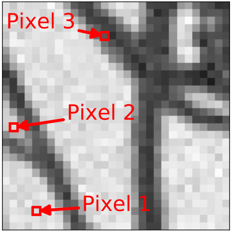

Denoised

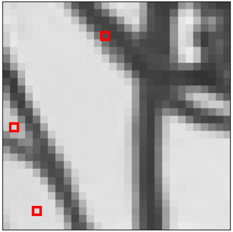

Low noise

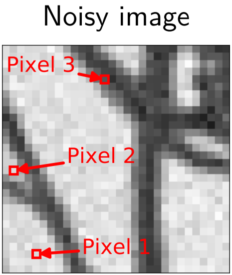

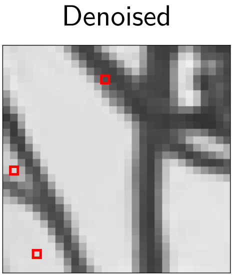

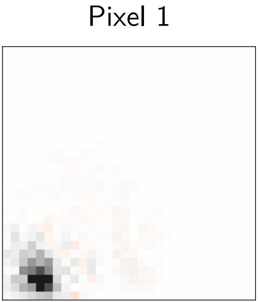

Low noise

Noisy image

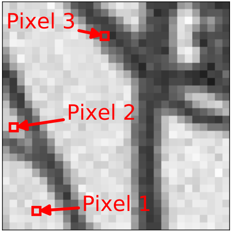

Denoised

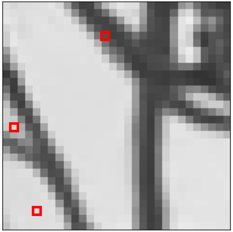

Pixel 1

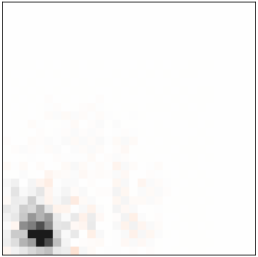

Pixel 2

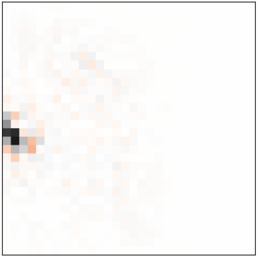

Low noise

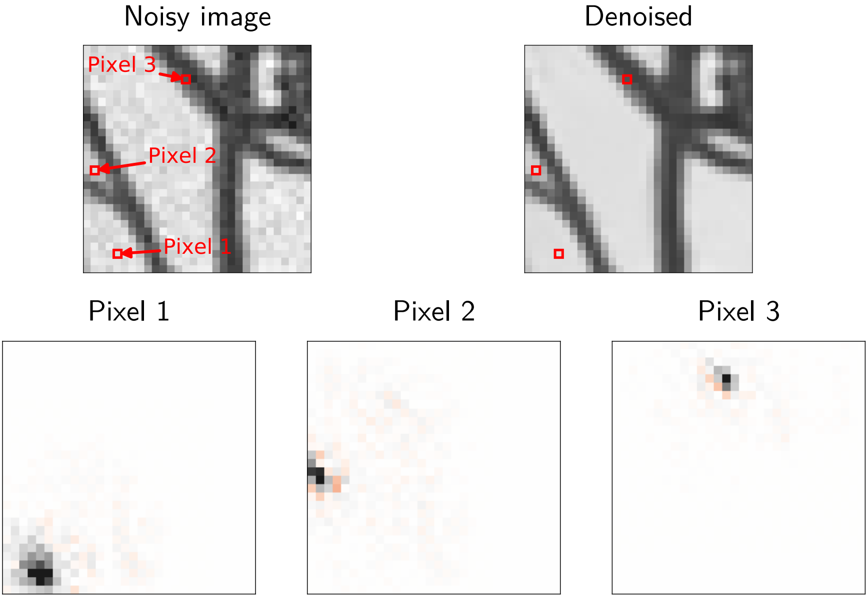

Noisy image
Denoised
Pixel 3
Pixel 2
Pixel 1
Pixel 1
Pixel 2
Pixel 3

Medium noise

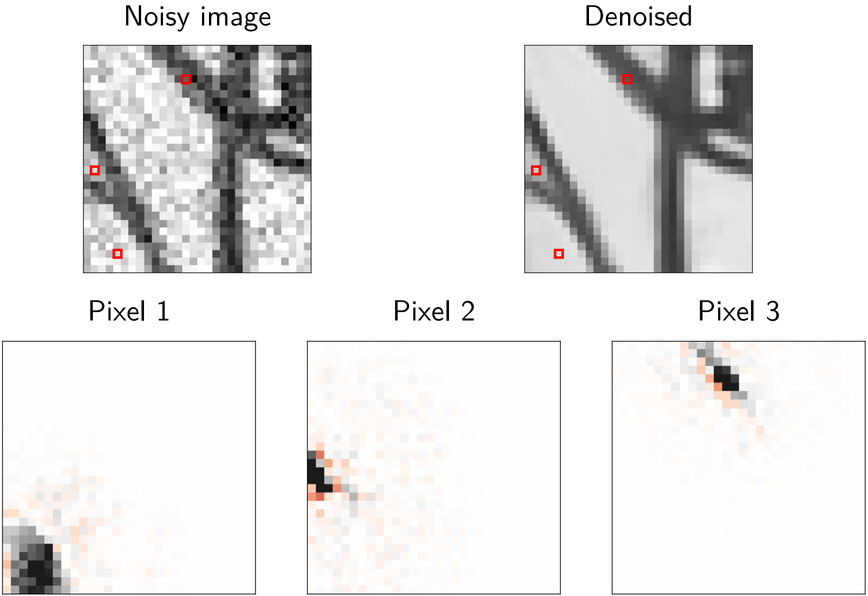

High noise

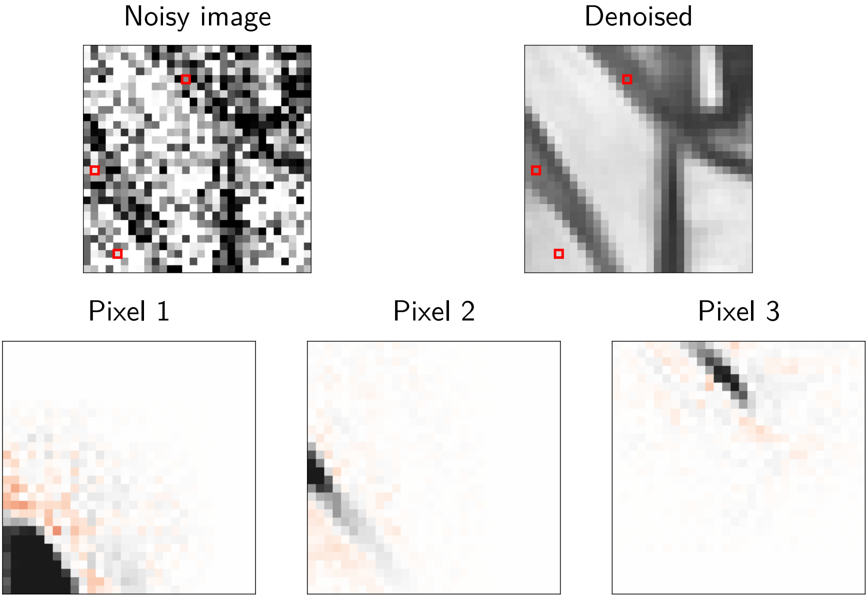

Application to electron microscopy[3]

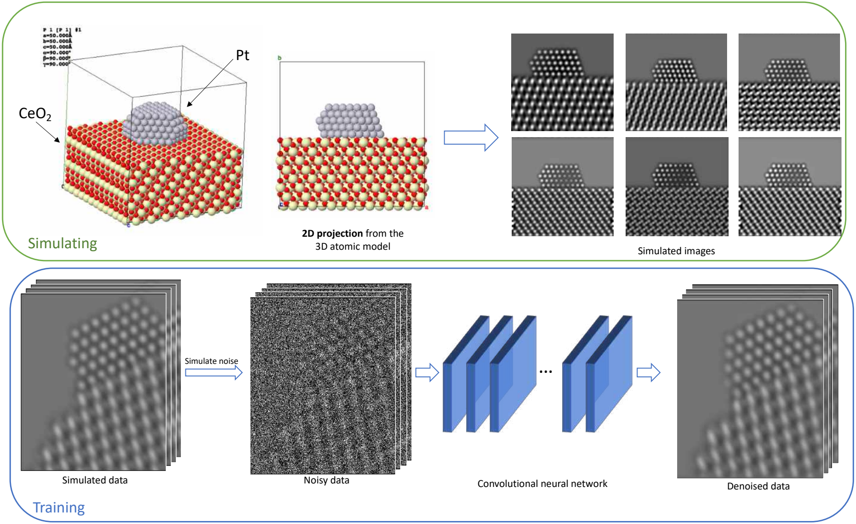

[3] *Deep denoising for scientific discovery: A case study in electron microscopy*. S. Mohan, R. Manzorro, J. L. Vincent, B. Tang, D. Y. Sheth, D. S. Matteson, E. P. Simoncelli, P. A. Crozier, C. Fernandez-Granda. IEEE Transactions on Computational Imaging 2022

Results

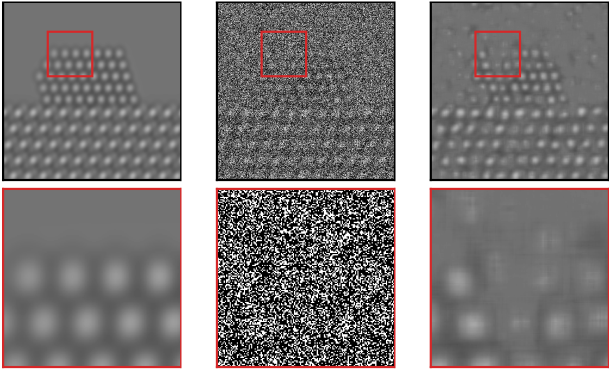

Simulated clean image

Noisy image

DnCNN

Results

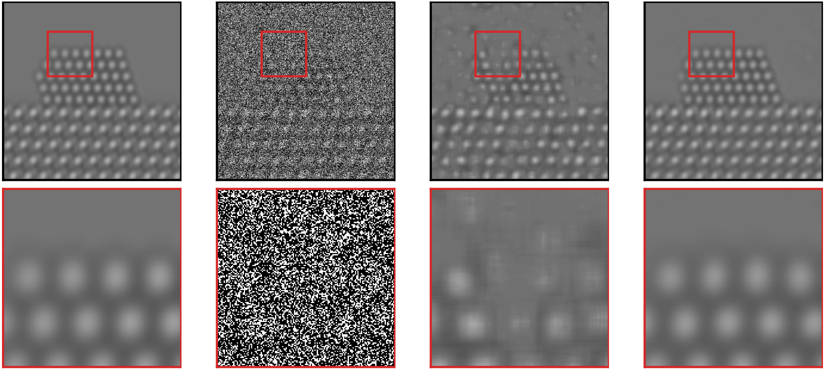

Simulated clean image | Noisy image | DnCNN | UNet (large receptive field)

Gradient

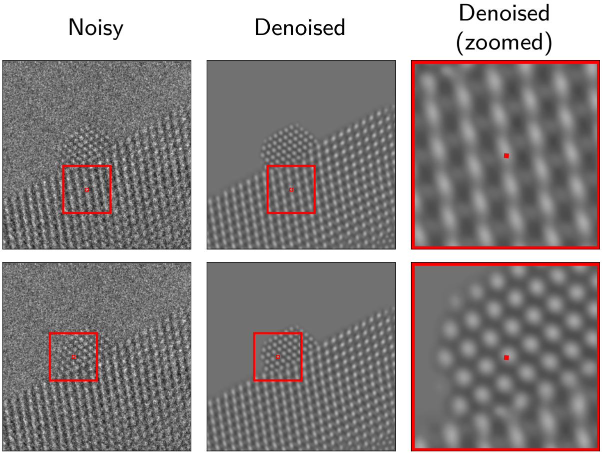

Gradient

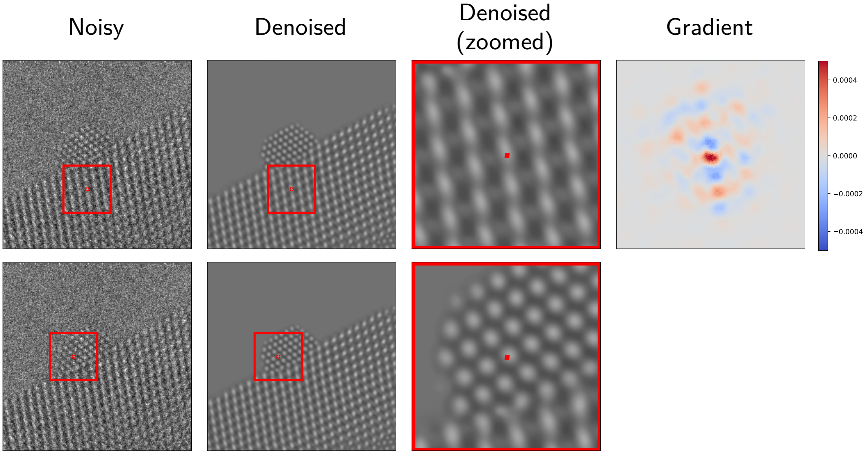
Noisy
Denoised
Denoised (zoomed)
Gradient
0.0004
0.0002
0.0000
−0.0002
−0.0004

Gradient

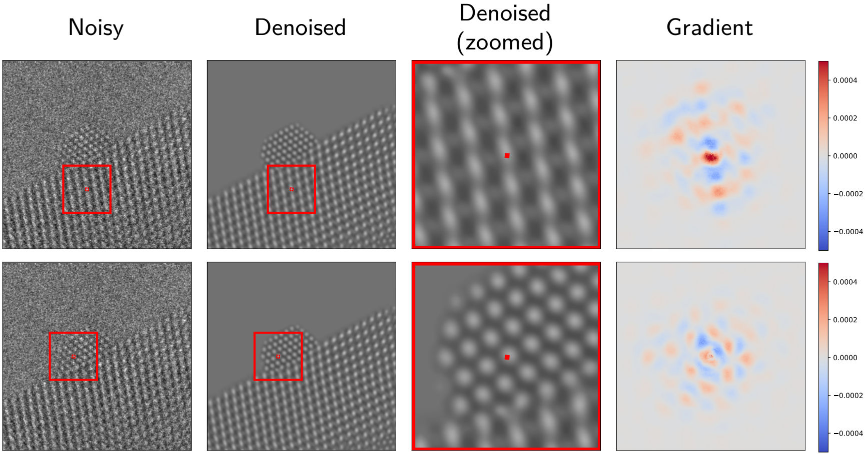

Noisy
Denoised
Denoised (zoomed)
Gradient
0.0004
0.0002
0.0000
−0.0002
−0.0004

Denoising via deep learning

Unsupervised denoising

Unsupervised metrics

Challenge

We often cannot simulate ground truth (because we don't know it!)

Supervised MSE

Supervised mean squared error

Denoised Clean

Supervised MSE

Supervised mean squared error

$\Big\|$  $-$  $\Big\|^2$

Denoised — Clean

Problem: We don't have clean images...

Noise2Noise[4]

Solution: Just use noisy images!

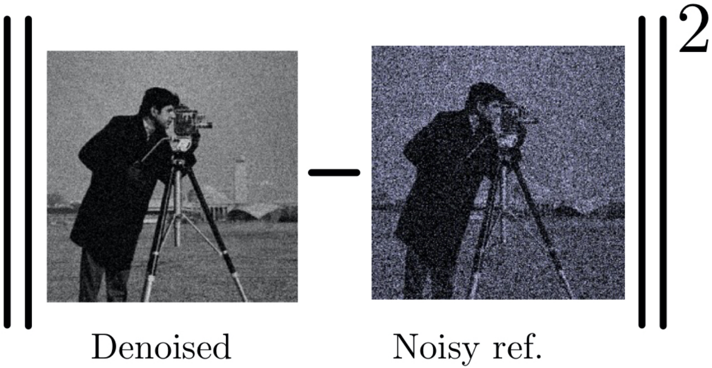

[4] *Noise2noise: Learning image restoration without clean data*. Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., Aila, T. ICML 2018

Noise2Noise[4]

Solution: Just use noisy images!

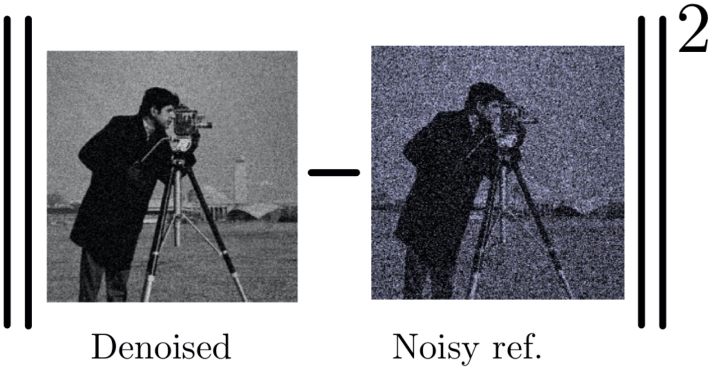

Requires multiple copies of clean image with independent noise

[4] *Noise2noise: Learning image restoration without clean data*. Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., Aila, T. ICML 2018

Neighbor2Neighbor[5]

Obtains copies from single image via spatial subsampling

[5] *Neighbor2Neighbor: Self-Supervised Denoising from Single Noisy Images*. T. Huang, S. Li, X. Jia, H. Lu, J. Liu CVPR 2021

Blind-spot denoising[6]

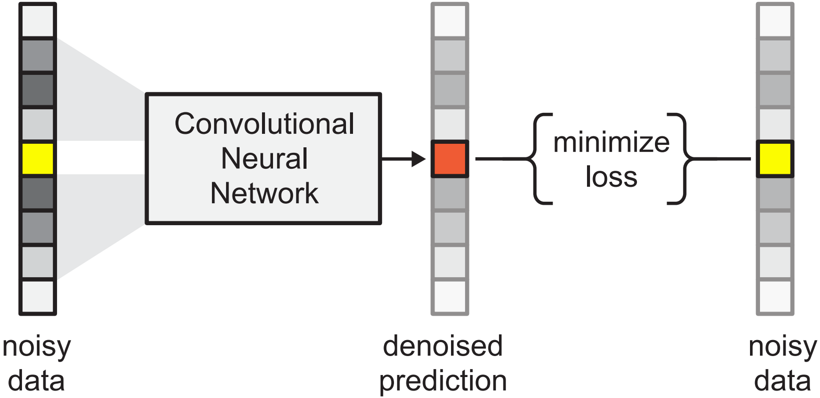

[6] *Noise2self: Blind denoising by self-supervision*. J. Batson, L. Royer. ICML 2019
Noise2void- Learning denoising from single noisy images A. Krull, T. Buchholz, F. Jug. CVPR 2019
High-quality self-supervised deep image denoising S. Laine, T. Karras, J. Lehtinen, T. Aila. Neurips 2019

Electron microscopy

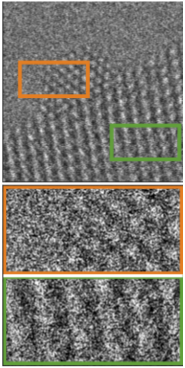

Noisy image

Electron microscopy

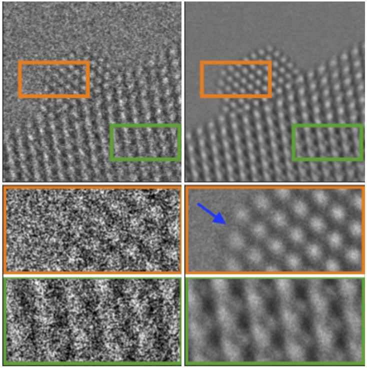

Noisy image Reference

Electron microscopy

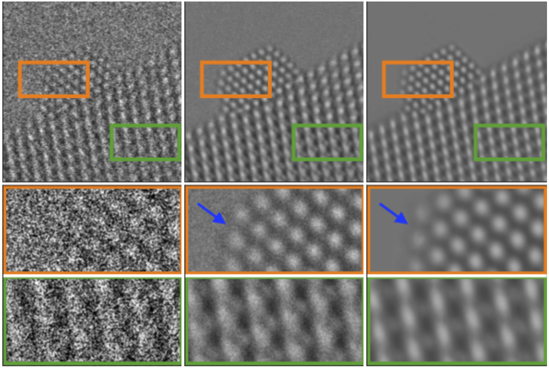

Noisy image Reference Supervised

Electron microscopy

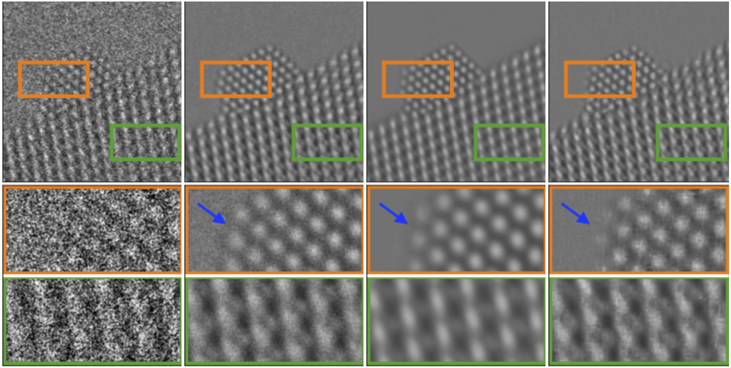

Noisy image | Reference | Supervised | Unsupervised[7]

[7] *Unsupervised Deep Video Denoising* D. Sheth, S. Mohan, J. Vincent, R. Manzorro, P. Crozier, M. Khapra, E. Simoncelli, C. Fernandez-Granda. ICCV 2021

Electron microscopy

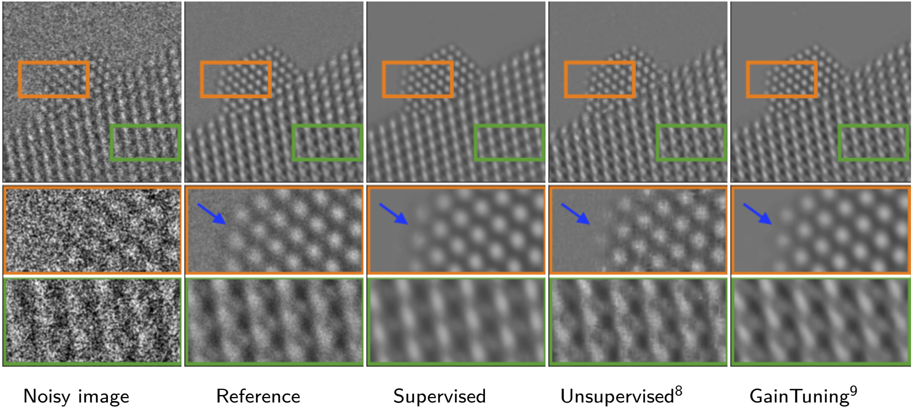

[8] *Unsupervised Deep Video Denoising* D. Sheth, S. Mohan, J. Vincent, R. Manzorro, P. Crozier, M. Khapra, E. Simoncelli, C. Fernandez-Granda. ICCV 2021

[9] *Adaptive Denoising via GainTuning* S. Mohan, J. Vincent, R. Manzorro, P. Crozier, C. Fernandez-Granda, E. Simoncelli. NeurIPS 2021

Denoising via deep learning

Unsupervised denoising

Unsupervised metrics

Unsupervised metrics

In existing work, unsupervised methods are evaluated:

Unsupervised metrics

In existing work, unsupervised methods are evaluated:

- On simulated data with known clean images

Unsupervised metrics

In existing work, unsupervised methods are evaluated:

- On simulated data with known clean images
- By visual inspection

Unsupervised metrics

In existing work, unsupervised methods are evaluated:

- On simulated data with known clean images
- By visual inspection
- By comparing to *clean* images estimated via averaging

Unsupervised metrics

In existing work, unsupervised methods are evaluated:

- On simulated data with known clean images
- By visual inspection
- By comparing to *clean* images estimated via averaging

Goal: Metric for quantitative evaluation without clean images

Idea

Compare to a noisy reference as in the Noise2Noise cost function

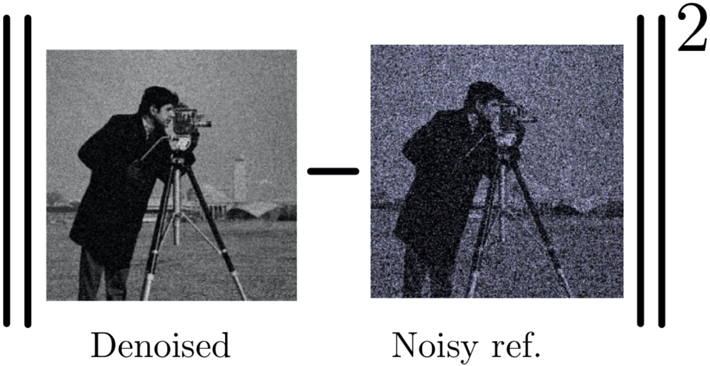

Additive Gaussian noise with variance σ^2

Clean image: x Data: $y = x + z$ Denoised estimate: $f(y)$

Additive Gaussian noise with variance σ^2

Clean image: x Data: $y = x + z$ Denoised estimate: $f(y)$

$$\mathrm{MSE} := \frac{1}{n} \sum_{i=1}^{n} (x_i - f(y)_i)^2$$

Additive Gaussian noise with variance σ^2

Clean image: x Data: $y = x + z$ Denoised estimate: $f(y)$

$$\mathrm{MSE} := \frac{1}{n}\sum_{i=1}^{n}(x_i - f(y)_i)^2$$

Noisy reference: $a = x + w$

Additive Gaussian noise with variance σ^2

Clean image: x Data: $y = x + z$ Denoised estimate: $f(y)$

$$\mathrm{MSE} := \frac{1}{n} \sum_{i=1}^{n} \left(x_i - f(y)_i\right)^2$$

Noisy reference: $a = x + w$

$$\frac{1}{n} \sum_{i=1}^{n} \left(a_i - f(y)_i\right)^2$$

Additive Gaussian noise with variance σ^2

Clean image: x Data: $y = x + z$ Denoised estimate: $f(y)$

$$\mathrm{MSE} := \frac{1}{n}\sum_{i=1}^{n}(x_i - f(y)_i)^2$$

Noisy reference: $a = x + w$

$$\frac{1}{n}\sum_{i=1}^{n}(a_i - f(y)_i)^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i + w_i - f(y)_i)^2$$

Additive Gaussian noise with variance σ^2

Clean image: x Data: $y = x + z$ Denoised estimate: $f(y)$

$$\mathrm{MSE} := \frac{1}{n} \sum_{i=1}^{n} \left(x_i - f(y)_i \right)^2$$

Noisy reference: $a = x + w$

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^{n} \left(a_i - f(y)_i \right)^2 &= \frac{1}{n} \sum_{i=1}^{n} \left(x_i + w_i - f(y)_i \right)^2 \\ &\approx \frac{1}{n} \sum_{i=1}^{n} \left(x_i - f(y)_i \right)^2 + \frac{1}{n} \sum_{i=1}^{n} w_i^2 \end{aligned}$$

Additive Gaussian noise with variance σ^2

Clean image: x Data: $y = x + z$ Denoised estimate: $f(y)$

$$\mathrm{MSE} := \frac{1}{n}\sum_{i=1}^{n}(x_i - f(y)_i)^2$$

Noisy reference: $a = x + w$

$$\begin{aligned}
\frac{1}{n}\sum_{i=1}^{n}(a_i - f(y)_i)^2 &= \frac{1}{n}\sum_{i=1}^{n}(x_i + w_i - f(y)_i)^2 \\
&\approx \frac{1}{n}\sum_{i=1}^{n}(x_i - f(y)_i)^2 + \frac{1}{n}\sum_{i=1}^{n} w_i^2 \\
&\approx \mathrm{MSE} + \sigma^2
\end{aligned}$$

Correction term

$$\frac{1}{n}\sum_{i=1}^{n}(a_i - f(y)_i)^2 \approx \mathrm{MSE} + \sigma^2$$

Additional noisy references: $b := x + v$ $\quad c := x + u$

Correction term

$$\frac{1}{n}\sum_{i=1}^{n}(a_i - f(y)_i)^2 \approx \mathrm{MSE} + \sigma^2$$

Additional noisy references: $\quad b := x + v \qquad c := x + u$

$$\frac{1}{n}\sum_{i=1}^{n}\frac{(b_i - c_i)^2}{2} = \frac{1}{n}\sum_{i=1}^{n}\frac{(v_i - u_i)^2}{2}$$

Correction term

$$\frac{1}{n}\sum_{i=1}^{n}(a_i - f(y)_i)^2 \approx \mathrm{MSE} + \sigma^2$$

Additional noisy references: $\quad b := x + v \qquad c := x + u$

$$\begin{aligned}\frac{1}{n}\sum_{i=1}^{n}\frac{(b_i - c_i)^2}{2} &= \frac{1}{n}\sum_{i=1}^{n}\frac{(v_i - u_i)^2}{2} \\ &\approx \frac{1}{2n}\sum_{i=1}^{n} v_i^2 + \frac{1}{2n}\sum_{i=1}^{n} u_i^2\end{aligned}$$

Correction term

$$\frac{1}{n}\sum_{i=1}^{n}(a_i - f(y)_i)^2 \approx \mathrm{MSE} + \sigma^2$$

Additional noisy references: $\quad b := x + v \qquad c := x + u$

$$\begin{aligned}
\frac{1}{n}\sum_{i=1}^{n}\frac{(b_i - c_i)^2}{2} &= \frac{1}{n}\sum_{i=1}^{n}\frac{(v_i - u_i)^2}{2} \\
&\approx \frac{1}{2n}\sum_{i=1}^{n} v_i^2 + \frac{1}{2n}\sum_{i=1}^{n} u_i^2 \\
&\approx \sigma^2
\end{aligned}$$

Correction term

$$\frac{1}{n}\sum_{i=1}^{n}(a_i - f(y)_i)^2 \approx \mathrm{MSE} + \sigma^2$$

Additional noisy references: $\quad b := x + v \qquad c := x + u$

$$\begin{aligned}\frac{1}{n}\sum_{i=1}^{n}\frac{(b_i - c_i)^2}{2} &= \frac{1}{n}\sum_{i=1}^{n}\frac{(v_i - u_i)^2}{2}\\ &\approx \frac{1}{2n}\sum_{i=1}^{n}v_i^2 + \frac{1}{2n}\sum_{i=1}^{n}u_i^2\\ &\approx \sigma^2\end{aligned}$$

$$\text{uMSE} := \frac{1}{n}\sum_{i=1}^{n}(a_i - f(y)_i)^2 - \frac{(b_i - c_i)^2}{2}$$

uMSE and uPSNR

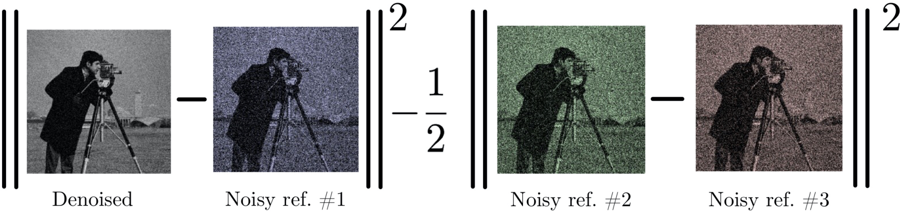

$$\mathrm{uMSE} := \frac{1}{n}\sum_{i=1}^{n}\left(a_i - f(y)_i\right)^2 - \frac{\left(b_i - c_i\right)^2}{2}$$

uMSE and uPSNR

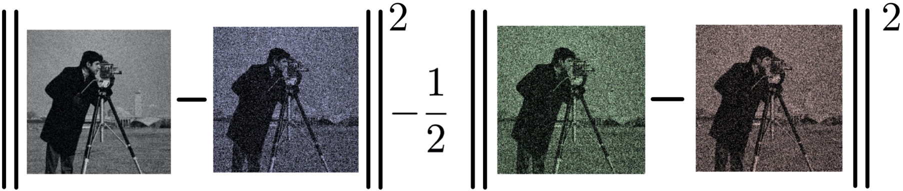

Denoised Noisy ref. #1 Noisy ref. #2 Noisy ref. #3

$$\text{uMSE} := \frac{1}{n}\sum_{i=1}^{n} (a_i - f(y)_i)^2 - \frac{(b_i - c_i)^2}{2}$$

$$\text{uPSNR} := 10 \log \left(\frac{255^2}{\text{uMSE}} \right)$$

Statistical properties

If noisy references correspond to the same clean image and noise is pixel-wise independent

Statistical properties

If noisy references correspond to the same clean image and noise is pixel-wise independent

- The uMSE and uPSNR are unbiased

Statistical properties

If noisy references correspond to the same clean image and noise is pixel-wise independent

- The uMSE and uPSNR are unbiased
- The uMSE and uPSNR are consistent

Statistical properties

If noisy references correspond to the same clean image and noise is pixel-wise independent

- The uMSE and uPSNR are unbiased
- The uMSE and uPSNR are consistent
- The uMSE is asymptotically Gaussian

Simulations

20 pixels

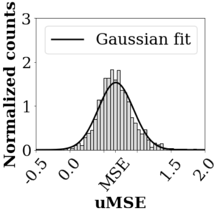

Simulations

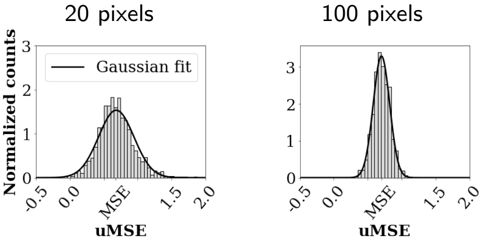

Simulations

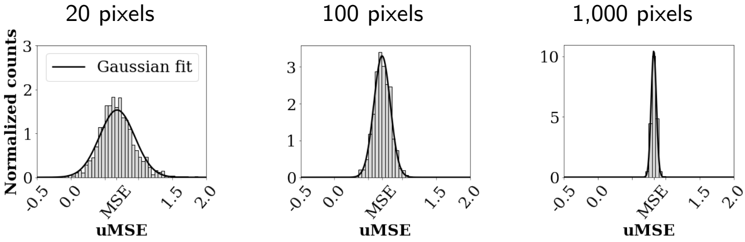

Simulations

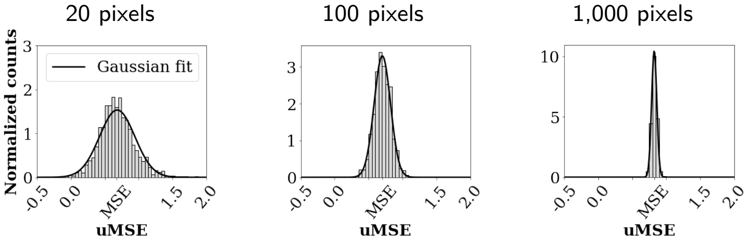

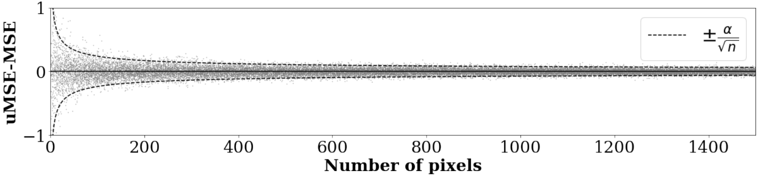

Confidence intervals

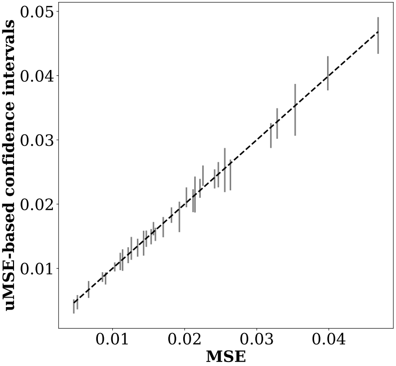

Comparison to averaging approach

Existing works compute MSE using average of noisy images as *clean image*

Comparison to averaging approach

Existing works compute MSE using average of noisy images as *clean image*

Requires many noisy images to converge to the true MSE

Comparison to averaging approach

Existing works compute MSE using average of noisy images as *clean image*

Requires many noisy images to converge to the true MSE

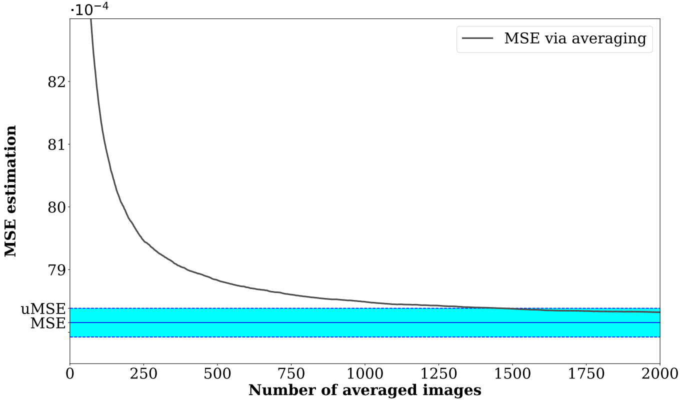

How do we compute the noisy references?

How do we compute the noisy references?

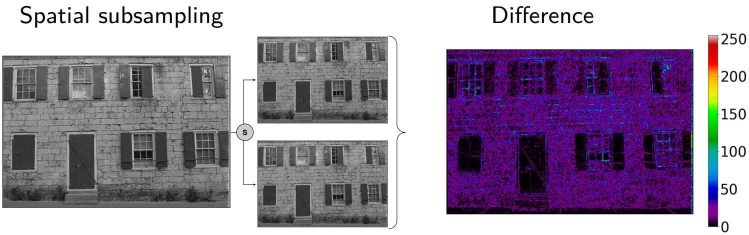

How do we compute the noisy references?

Spatial subsampling

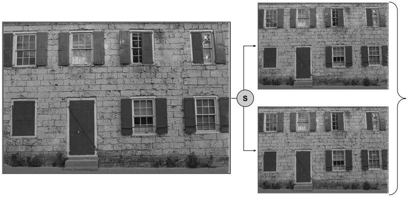

Difference

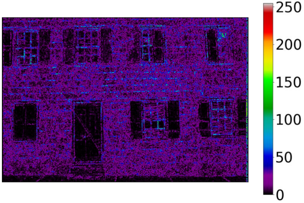

Consecutive frames

Difference

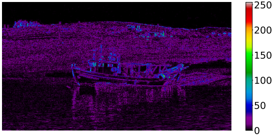

Bias

Natural images (Gaussian noise)

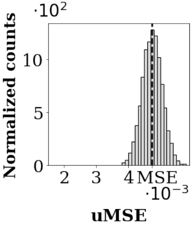

Bias

Natural images (Gaussian noise)

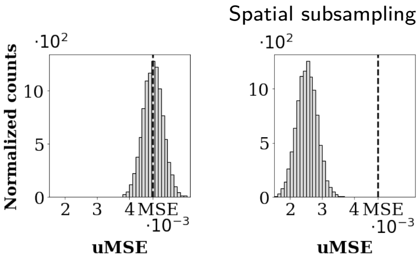

Bias

Natural images (Gaussian noise) **Electron microscopy (Poisson noise)**

Spatial subsampling

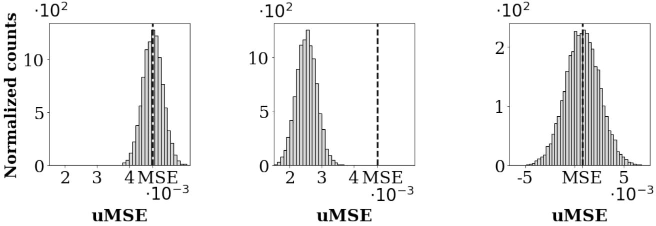

Bias

Natural images (Gaussian noise) **Electron microscopy (Poisson noise)**

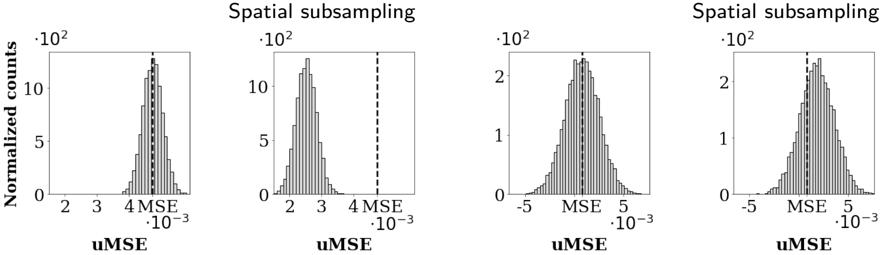

Natural images (Gaussian noise)

	$\sigma = 25$	$\sigma = 50$	$\sigma = 75$
Method	PSNR	PSNR	PSNR
Bilateral	24.20	21.84	19.14
DenseNet	26.54	23.98	22.75
DnCNN	26.19	23.95	22.72
UNet	26.29	23.92	22.68

Natural images (Gaussian noise)

	$\sigma = 25$		$\sigma = 50$		$\sigma = 75$	
Method	PSNR	uPSNR	PSNR	uPSNR	PSNR	uPSNR
Bilateral	24.20	24.18	21.84	21.86	19.14	19.17
DenseNet	26.54	26.51	23.98	24.06	22.75	23.00
DnCNN	26.19	26.21	23.95	24.02	22.72	22.75
UNet	26.29	26.28	23.92	24.01	22.68	22.70

Natural images (Gaussian noise)

	$\sigma = 25$			$\sigma = 50$			$\sigma = 75$		
Method	PSNR	uPSNR	uPSNR_S	PSNR	uPSNR	uPSNR_S	PSNR	uPSNR	uPSNR_S
Bilateral	24.20	24.18	26.20	21.84	21.86	22.90	19.14	19.17	19.58
DenseNet	26.54	26.51	27.61	23.98	24.06	26.28	22.75	23.00	24.69
DnCNN	26.19	26.21	28.14	23.95	24.02	26.08	22.72	22.75	24.59
UNet	26.29	26.28	27.98	23.92	24.01	26.25	22.68	22.70	24.84

Electron microscopy (Poisson noise)

Bilateral	Supervised	Unsupervised
PSNR	PSNR	PSNR
20.18	25.74	24.86

Electron microscopy (Poisson noise)

Bilateral		Supervised		Unsupervised	
PSNR	uPSNR	PSNR	uPSNR	PSNR	uPSNR
20.18	20.20	25.74	25.68	24.86	24.87

Electron microscopy (Poisson noise)

Bilateral			Supervised			Unsupervised		
PSNR	uPSNR	uPSNR_S	PSNR	uPSNR	uPSNR_S	PSNR	uPSNR	uPSNR_S
20.18	20.20	20.21	25.74	25.68	25.86	24.86	24.87	24.74

Real electron-microscope data

Inter-pixel correlation is non-negligeable

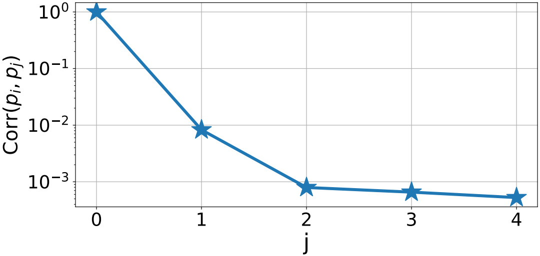

Real electron-microscope data

Inter-pixel correlation is non-negligeable

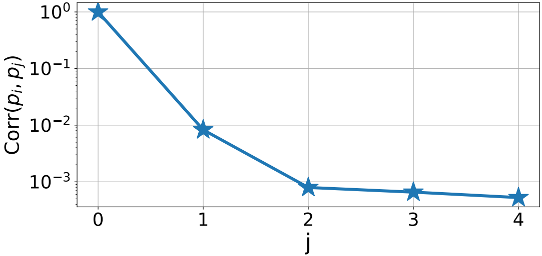

Solution: Spatial subsampling

Real electron-microscope data

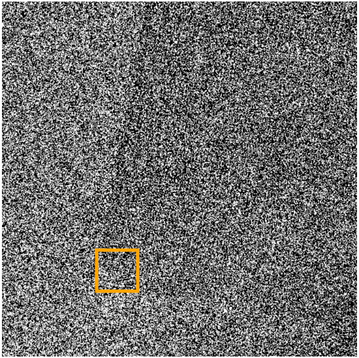

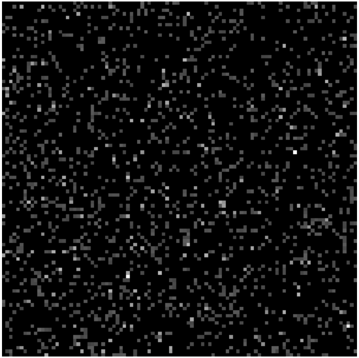

Gaussian smoothing, uPSNR: 20.4 dB

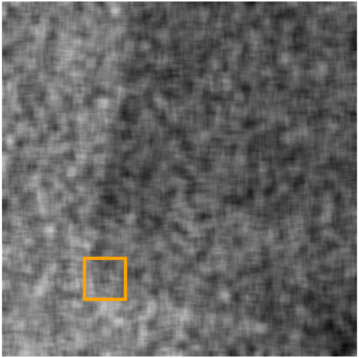

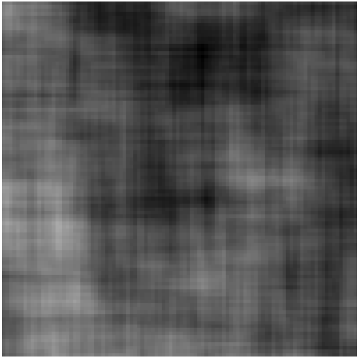

Neural network (Neighbor2neighbor), uPSNR: 26.9 dB

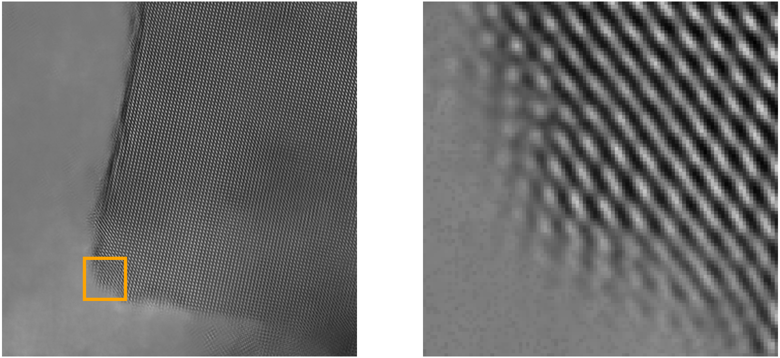

Conclusion

uMSE/uPSNR are a consistent estimator of MSE/PSNR

Conclusion

uMSE/uPSNR are a consistent estimator of MSE/PSNR

Open questions:

Conclusion

uMSE/uPSNR are a consistent estimator of MSE/PSNR

Open questions:

- How to address spatial-subsampling bias?

Conclusion

uMSE/uPSNR are a consistent estimator of MSE/PSNR

Open questions:

- How to address spatial-subsampling bias?
- How to deal with correlated noise?

Conclusion

uMSE/uPSNR are a consistent estimator of MSE/PSNR

Open questions:

- How to address spatial-subsampling bias?
- How to deal with correlated noise?
- Extension to inverse problems beyond denoising

For more information

Robust and interpretable blind image denoising via bias-free convolutional neural networks
Mohan & Kadkhodaie et. al. ICLR 2020

Unsupervised deep video denoising
Sheth & Mohan et. al. ICCV 2021

Adaptive denoising via GainTuning
Mohan et. al. NeurIPS 2021

Deep denoising for scientific discovery: A case study in electron microscopy
Mohan et. al. IEEE Transactions on Computational Imaging 2022

Developing and Evaluating Deep Neural Network-based denoising for Nanoparticle TEM Images with Ultra-low Signal-to-Noise
Vincent et. al. Microscopy & Microanalysis 2021

Evaluating Unsupervised Denoising Requires Unsupervised Metrics
Marcos-Morales et. al. Preprint 2022