An Unpaired Learning-based Method for Image Despeckling

CISA 2025

Ali Zafari Shirin Jalali

Department of Electrical and Computer Engineering
Rutgers University

Outline

Outline

Speckle noise and despeckling

QMAP: Theoretically-founded learning-based denoiser

Bayesian despeckling via QMAP

BD-QMAP for Images: Unpaired Learning-based Despeckler

Final Remarks

Speckle in coherent imaging systems

- Synthetic Aperture Radar
- Digital Holography
- Optical Coherence Tomography

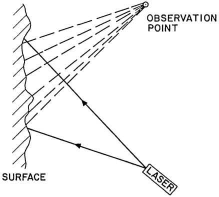

Image adapted from [Goodman, 1976].

Speckle in coherent imaging systems

- Synthetic Aperture Radar
- Digital Holography
- Optical Coherence Tomography

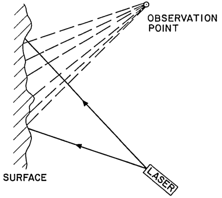

Image adapted from [Goodman, 1976].

- Speckle is modeled as multiplicative noise

$$Y = XW,$$

$$W \sim \mathcal{N}(0, \sigma^2).$$

Clean

Speckled

Despeckling

- **Goal**: Recover $X \in \mathbb{R}^n$ from $Y \in \mathbb{R}^n$:

$$Y = X \odot W,$$

$W \sim \mathcal{N}(0, \sigma^2 I)$ and $Y_i = X_i W_i$.

Despeckling

- **Goal**: Recover $X \in \mathbb{R}^n$ from $Y \in \mathbb{R}^n$:

$$Y = X \odot W,$$

$W \sim \mathcal{N}(0, \sigma^2 I)$ and $Y_i = X_i W_i$.

- *Maximum likelihood estimation*:

$$\textbf{MLE}: \quad \widehat{X} = \underset{u \in \mathbb{R}^n}{\arg\min} \; -\log \mathbb{P}(Y \mid X = u) = \frac{|Y|}{\sigma}$$

Despeckling

- **Goal**: Recover $X \in \mathbb{R}^n$ from $Y \in \mathbb{R}^n$:

$$Y = X \odot W,$$

$W \sim \mathcal{N}(0, \sigma^2 I)$ and $Y_i = X_i W_i$.

- *Maximum likelihood estimation*:

$$\textbf{MLE}: \quad \widehat{X} = \operatorname*{arg\,min}_{u \in \mathbb{R}^n} \; -\log \mathbb{P}(Y \mid X = u) = \frac{|Y|}{\sigma}$$

- Some knowledge/assumption over X is required:
  - Classical algorithms (*e.g.*, LMMSE in a transform domain, sparsity *etc.*)
  - Learning-based techniques (*e.g.*, CNN trained on a dataset)

From additive denoising to despeckling

	AWGN $Y = X + Z$	Speckle $Y = X \odot W$
Classic Methods	LMMSE Filtering [1] NLM [2] BM3D [3]	Lee/Kuan [9] PPB [10] SAR-BM3D [11]
Paired Learning-based	DnCNN [4] Restormer [5] DDRM (diffusion-based) [6] Noise2Noise [7]	ID-CNN [12] Transformer-based Despeckler [13] SAR-DDPM [14] SAR2SAR [15]
Unpaired Learning-based	QMAP [8]	BD-QMAP [16]

[1] Wallis (1976) [2] Buades (2005) [3] Dabov (2007) [4] Zhang (2016)
[5] Zamir (2022) [6] Kawar (2022) [7] Lehtinen (2018) [8] Zhou (2023)

[9] Kuan (1985) [10] Deledalle (2009) [11] Parrilli (2011) [12] Wang (2017)
[13] Perera (2022) [14] Perera (2023) [15] Dalsasso (2021) [16] Zafari (2025)

From additive denoising to despeckling

	AWGN $Y = X + Z$	**Speckle** $Y = X \odot W$
Classic Methods	LMMSE Filtering [1] NLM [2] BM3D [3]	Lee/Kuan [9] PPB [10] SAR-BM3D [11]
Paired Learning-based	DnCNN [4] Restormer [5] DDRM (diffusion-based) [6] Noise2Noise [7]	ID-CNN [12] Transformer-based Despeckler [13] SAR-DDPM [14] SAR2SAR [15]
Unpaired Learning-based	QMAP [8]	BD-QMAP [16]

✗ classic methods : strong assumptions on source model

From additive denoising to despeckling

	AWGN $Y = X + Z$	Speckle $Y = X \odot W$
Classic Methods	LMMSE Filtering [1] NLM [2] BM3D [3]	Lee/Kuan [9] PPB [10] SAR-BM3D [11]
Paired Learning-based	DnCNN [4] Restormer [5] DDRM (diffusion-based) [6] Noise2Noise [7]	ID-CNN [12] Transformer-based Despeckler [13] SAR-DDPM [14] SAR2SAR [15]
Unpaired Learning-based	QMAP [8]	BD-QMAP [16]

✗ classic methods : strong assumptions on source model

✗ paired learning-based : requires access to a labeled dataset

Outline

Towards a theoretically-founded denoiser

Minimum mean squared error denoiser:

$$\textbf{MMSE}: \quad \widehat{X} = \mathbb{E}[X \mid Y]$$

Towards a theoretically-founded denoiser

Minimum mean squared error denoiser:

$$\textbf{MMSE}: \quad \widehat{X} = \mathbb{E}[X|Y]$$

✗ Access to $\mathbb{P}(X)$ is often non-trivial.

Towards a theoretically-founded denoiser

Minimum mean squared error denoiser:

$$\textbf{MMSE}: \quad \widehat{X} = \mathbb{E}[X|Y]$$

- ✗ Access to $\mathbb{P}(X)$ is often non-trivial.
- ✗ Given $\mathbb{P}(X)$, computing $\mathbb{E}[X|Y]$ is intractable.

Towards a theoretically-founded denoiser

Minimum mean squared error denoiser:

$$\textbf{MMSE}: \quad \widehat{X} = \mathbb{E}[X|Y]$$

- ✗ Access to $\mathbb{P}(X)$ is often non-trivial.
- ✗ Given $\mathbb{P}(X)$, computing $\mathbb{E}[X|Y]$ is intractable.
- ✗ Theoretical analysis is hard, especially for high-dimensional X.

Towards a theoretically-founded denoiser

Minimum mean squared error denoiser:

$$\textbf{MMSE}: \quad \widehat{X} = \mathbb{E}[X|Y]$$

- ✗ Access to $\mathbb{P}(X)$ is often non-trivial.
- ✗ Given $\mathbb{P}(X)$, computing $\mathbb{E}[X|Y]$ is intractable.
- ✗ Theoretical analysis is hard, especially for high-dimensional X.

Same challenges for *maximum a'posteriori* denoiser:

$$\textbf{MAP}: \quad \widehat{X} = \underset{u \in \mathbb{R}^n}{\arg\min} \; -\log \mathbb{P}(Y \mid X = u) - \log \mathbb{P}(X = u)$$

QMAP: Theoretically-founded additive denoiser

MAP: $\widehat{X} = \arg\min_{u \in \mathbb{R}^n} \underbrace{-\log \mathbb{P}(Y \mid X = u)}_{\text{measurement-alignment}} + \underbrace{-\log \mathbb{P}(X = u)}_{\text{structure-consistency}}$

QMAP: Theoretically-founded additive denoiser

$$\textbf{QMAP}: \ \widehat{X} = \arg\min_{u \in \mathbb{R}^n} \ \underbrace{-\log \mathbb{P}(Y \mid X = u)}_{\text{measurement-alignment}} \ + \ \lambda \ \underbrace{c_{\mathbf{w}}(u)}_{\text{structure-consistency}}$$

QMAP: Theoretically-founded additive denoiser

$$\textbf{QMAP}: \ \widehat{X} = \arg\min_{u \in \mathbb{R}^n} \underbrace{-\log \mathbb{P}(Y \mid X = u)}_{\text{measurement-alignment}} + \lambda \underbrace{c_{\mathbf{w}}(u)}_{\text{structure-consistency}}$$

Measure of *quantized structural consistency* for $[u_i \cdots u_{k+i-1}] \in \mathbb{R}^k$

$$w_i := -\log \mathbb{P}(\lfloor X_i \cdots X_{k+i-1} \rceil_b = [u_i \cdots u_{k+i-1}]),$$

where $\lfloor \cdot \rceil_b$ quantizes to b bits.

QMAP: Theoretically-founded additive denoiser

$$\textbf{QMAP}: \ \widehat{X} = \arg\min_{u \in \mathbb{R}^n} \underbrace{-\log \mathbb{P}(Y \mid X = u)}_{\text{measurement-alignment}} + \lambda \underbrace{c_{\mathbf{w}}(u)}_{\text{structure-consistency}}$$

Measure of *quantized structural consistency* for $[u_i \cdots u_{k+i-1}] \in \mathbb{R}^k$

$$w_i := -\log \mathbb{P}(\lfloor X_i \cdots X_{k+i-1} \rceil_b = [u_i \cdots u_{k+i-1}]),$$

$$c_{\mathbf{w}}(u) := \sum_{i=1} w_i$$

QMAP: Theoretically-founded additive denoiser

QMAP: $\widehat{X} = \arg\min_{u \in \mathbb{R}^n} \underbrace{-\log \mathbb{P}(Y \mid X = u)}_{\text{measurement-alignment}} + \lambda \underbrace{c_{\mathbf{w}}(u)}_{\text{structure-consistency}}$

Measure of *quantized structural consistency* for $[u_i \cdots u_{k+i-1}] \in \mathbb{R}^k$

$$w_i := -\log \mathbb{P}(\lfloor X_i \cdots X_{k+i-1} \rceil_b = [u_i \cdots u_{k+i-1}]),$$

$$c_{\mathbf{w}}(u) := \sum_{i=1} w_i$$

Key observations [Zhou et al., 2023]:

- Asymptotically optimal in high SNR regime.

QMAP: Theoretically-founded additive denoiser

$$\textbf{QMAP}: \ \widehat{X} = \arg\min_{u \in \mathbb{R}^n} \underbrace{-\log \mathbb{P}(Y \mid X = u)}_{\text{measurement-alignment}} + \lambda \underbrace{c_{\mathbf{w}}(u)}_{\text{structure-consistency}}$$

Measure of *quantized structural consistency* for $[u_i \cdots u_{k+i-1}] \in \mathbb{R}^k$

$$w_i := -\log \mathbb{P}(\lfloor X_i \cdots X_{k+i-1} \rceil_b = [u_i \cdots u_{k+i-1}]),$$

$$c_{\mathbf{w}}(u) := \sum_{i=1} w_i$$

Key observations [Zhou et al., 2023]:

- Asymptotically optimal in high SNR regime.
- A small subset of weights w_i capture most source information.
- The set of weights w_i can be found from b-bit quantized k-th order empirical distribution learned from a set of realizations.

QMAP for piecewise constant source

Consider Markov source ($p \ll 1$):

$$X_{i+1} = \begin{cases} X_i & \text{with probability } 1-p, \\ \text{Uniform}(0,1) & \text{with probability } p. \end{cases}$$

QMAP for piecewise constant source

Consider Markov source ($p \ll 1$):

$$X_{i+1} = \begin{cases} X_i & \text{with probability } 1 - p, \\ \text{Uniform}(0, 1) & \text{with probability } p. \end{cases}$$

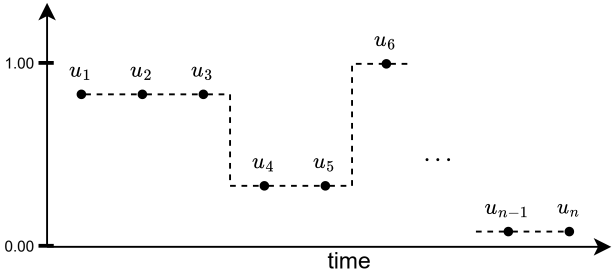

QMAP for piecewise constant source

Consider Markov source ($p \ll 1$):

$$X_{i+1} = \begin{cases} X_i & \text{with probability } 1-p, \\ \text{Uniform}(0,1) & \text{with probability } p. \end{cases}$$

$b = 2$

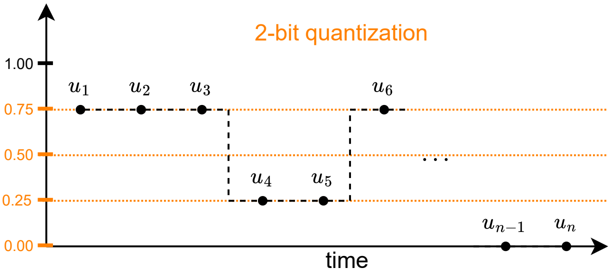

QMAP for piecewise constant source

Consider Markov source ($p \ll 1$):

$$X_{i+1} = \begin{cases} X_i & \text{with probability } 1-p, \\ \text{Uniform}(0,1) & \text{with probability } p. \end{cases}$$

$b = 2$, $k = 2$.

$$w_i = -\log \mathbb{P}(\lfloor X_i, X_{i+1} \rceil_b = [u_i, u_{i+1}]) \propto \begin{cases} -\log(1-p) & u_i = u_{i+1} \\ -\log(p) & u_i \neq u_{i+1} \end{cases}$$

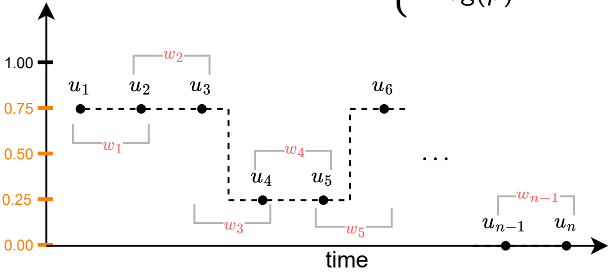

QMAP for piecewise constant source

Consider Markov source ($p \ll 1$):

$$X_{i+1} = \begin{cases} X_i & \text{with probability } 1-p, \\ \text{Uniform}(0,1) & \text{with probability } p. \end{cases}$$

$b = 2$, $k = 2$.

$$c_{\mathbf{w}}(u) = \sum_{i=1}^{n-1} w_i$$

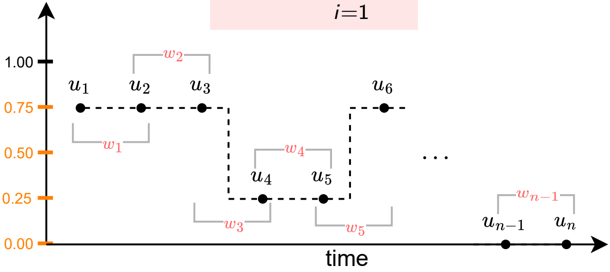

Outline

Despeckling via QMAP: BD-QMAP [Zafari and Jalali, 2025]

$Y = X \odot W$:

$$\widehat{X} = \underset{u \in \mathbb{R}^n}{\arg\min} \underbrace{-\log \mathbb{P}(Y \mid X = u)}_{\text{measurement-alignment}} + \lambda \underbrace{c_{\mathbf{w}}(u)}_{\text{structure-consistency}}$$

Despeckling via QMAP: BD-QMAP [Zafari and Jalali, 2025]

$Y = X \odot W$:

$$\widehat{X} = \arg\min_{u \in \mathbb{R}^n} \underbrace{\sum_{i=1}^{n} \log u_i^2 + \frac{Y_i^2}{u_i^2}}_{\text{measurement-alignment}} + \lambda \underbrace{c_{\mathbf{w}}(u)}_{\text{structure-consistency}}$$

Despeckling via QMAP: BD-QMAP [Zafari and Jalali, 2025]

$Y = X \odot W$:

$$\widehat{X} = \arg\min_{u \in \mathbb{R}^n} \underbrace{\sum_{i=1}^{n} \log u_i^2 + \frac{Y_i^2}{u_i^2}}_{\text{measurement-alignment}} + \lambda \underbrace{c_{\mathbf{w}}(u)}_{\text{structure-consistency}}$$

- Same set of weights can be used for despeckling/denoising.

Despeckling via QMAP: BD-QMAP [Zafari and Jalali, 2025]

$Y = X \odot W$:

$$\widehat{X} = \arg\min_{u \in \mathbb{R}^n} \underbrace{\sum_{i=1}^{n} \log u_i^2 + \frac{Y_i^2}{u_i^2}}_{\text{measurement-alignment}} + \lambda \underbrace{c_{\mathbf{w}}(u)}_{\text{structure-consistency}}$$

- Same set of weights can be used for despeckling/denoising.
- For piecewise constant source empirical distribution of quantized X can be estimated from a dataset.

Despeckling via QMAP: BD-QMAP [Zafari and Jalali, 2025]

$Y = X \odot W$:

$$\widehat{X} = \arg\min_{u \in \mathbb{R}^n} \underbrace{\sum_{i=1}^{n} \log u_i^2 + \frac{Y_i^2}{u_i^2}}_{\text{measurement-alignment}} + \lambda \underbrace{c_{\mathbf{w}}(u)}_{\text{structure-consistency}}$$

- Same set of weights can be used for despeckling/denoising.
- For piecewise constant source empirical distribution of quantized X can be estimated from a dataset.
- While MSE can still be used as the reconstruction loss, likelihood is a more natural choice.

BD-QMAP on piecewise constant source

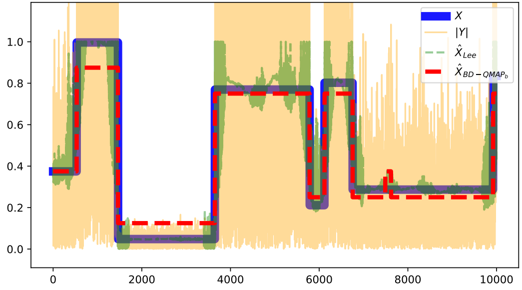

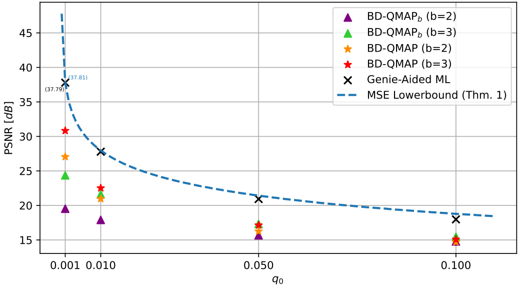

For more theoretical analysis of BD-QMAP look at

Zafari, A., & Jalali, S. (2025) *Bayesian Despeckling of Structured Sources.* arXiv:2501.11860.

Outline

BD-QMAP extension for image data: Challenges

BD-QMAP extension for image data: Challenges

Challenge 1 Finding weights $c_{\mathbf{w}}(u)$

- Quantizing and counting less effective as $n \nearrow$

BD-QMAP extension for image data: Challenges

Challenge 1 Finding weights $c_{\mathbf{w}}(u)$

- Quantizing and counting less effective as $n \nearrow$
- Recall: only a small subset of weights are important.

BD-QMAP extension for image data: Challenges

Challenge 1 Finding weights $c_{\mathbf{w}}(u)$

- Quantizing and counting less effective as $n \nearrow$
- Recall: only a small subset of weights are important.

Solution Quantized representation in a transform domain.

BD-QMAP extension for image data: Challenges

Challenge 1 Finding weights $c_{\mathbf{w}}(u)$

- Quantizing and counting less effective as $n \nearrow$
- Recall: only a small subset of weights are important.

Solution Quantized representation in a transform domain.

Challenge 2 Joint optimization of all pixels in $u \in \mathbb{R}^n$

$$\widehat{X} = \underset{u \in \mathbb{R}^n}{\arg\min} \quad \sum_{i=1}^{n} \log u_i^2 + \frac{Y_i^2}{u_i^2} \quad + \quad \lambda \quad c_{\mathbf{w}}(u)$$

BD-QMAP extension for image data: Challenges

Challenge 1 Finding weights $c_{\mathbf{w}}(u)$

- Quantizing and counting less effective as $n \nearrow$
- Recall: only a small subset of weights are important.

Solution Quantized representation in a transform domain.

Challenge 2 Joint optimization of all pixels in $u \in \mathbb{R}^n$

$$\widehat{X} = \arg\min_{u \in \mathbb{R}^n} \quad \sum_{i=1}^{n} \log u_i^2 + \frac{Y_i^2}{u_i^2} \quad + \quad \lambda \quad c_{\mathbf{w}}(u)$$

- *e.g.*, $n = 256 \times 256 = 2^{16}$ $\rightarrow$ Practically infeasible ($c_{\mathbf{w}}(u)$ non-convex)

BD-QMAP extension for image data: Challenges

Challenge 1 Finding weights $c_{\mathbf{w}}(u)$

- Quantizing and counting less effective as $n \nearrow$
- Recall: only a small subset of weights are important.

Solution Quantized representation in a transform domain.

Challenge 2 Joint optimization of all pixels in $u \in \mathbb{R}^n$

$$\widehat{X} = \arg\min_{u \in \mathbb{R}^n} \quad \sum_{i=1}^{n} \log u_i^2 + \frac{Y_i^2}{u_i^2} \quad + \quad \lambda \quad c_{\mathbf{w}}(u)$$

- *e.g.*, $n = 256 \times 256 = 2^{16}$ $\rightarrow$ Practically infeasible ($c_{\mathbf{w}}(u)$ non-convex)

Solution (sub-optimal) patch-wise optimization

BD-QMAP extension for image data: Challenges

Challenge 1 Finding weights $c_{\mathbf{w}}(u)$

- Quantizing and counting less effective as $n \nearrow$
- Recall: only a small subset of weights are important.

Solution Quantized representation in a transform domain.

Challenge 2 Joint optimization of all pixels in $u \in \mathbb{R}^n$

$$\widehat{X} = \arg\min_{u \in \mathbb{R}^n} \sum_{i=1}^{n} \log u_i^2 + \frac{Y_i^2}{u_i^2} + \lambda \, c_{\mathbf{w}}(u)$$

- *e.g.*, $n = 256 \times 256 = 2^{16}$ $\rightarrow$ Practically infeasible ($c_{\mathbf{w}}(u)$ non-convex)

Solution (sub-optimal) patch-wise optimization

Next 2 slides detail both solutions.

BD-QMAP for image data: learning weights

Challenge 1

- Prohibitively large number of patches $\propto 2^{b \times k^2}$
- Recall: small subset of $\mathbf{w}$ capture most of the source information.

BD-QMAP for image data: learning weights

Challenge 1

- Prohibitively large number of patches $\propto 2^{b \times k^2}$
- Recall: small subset of $\mathbf{w}$ capture most of the source information.

Autoencoder with binary bottleneck trained to minimize MSE.

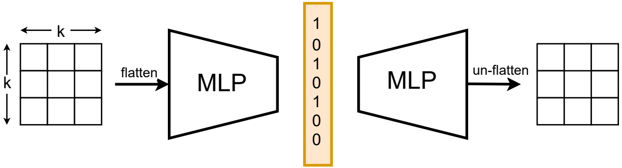

BD-QMAP extension for image data: patch-wise

Challenge 2 Despeckle $k \times k$ patches ($k \ll n$)

$$\widehat{X} = \arg\min_{u \in \mathbb{R}^{k \times k}} \underbrace{\sum_{i=1}^{k^2} \log u_i^2 + \frac{Y_i^2}{u_i^2}}_{\text{measurement-alignment}} + \lambda \underbrace{w(u)}_{\text{structure-consistency}}$$

Challenge 2 Despeckle $k \times k$ patches ($k \ll n$)

$$\widehat{X} = \arg\min_{u \in \mathbb{R}^{k \times k}} \underbrace{\sum_{i=1}^{k^2} \log u_i^2 + \frac{Y_i^2}{u_i^2}}_{\text{measurement-alignment}} + \lambda \underbrace{w(u)}_{\text{structure-consistency}}$$

To despeckle a pixel:

1. Extract all patches including the pixel
2. Despeckle each patch (parallel in a batch)
3. Average proportionally

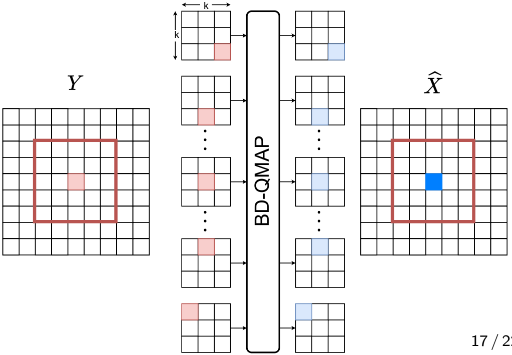

BD-QMAP: Synthetic speckle

Despeckling performance on *Set11* images.

Despeckling Method	PSNR	SSIM
speckled image	09.39	0.14
box car	17.11	0.42
Kuan (enhanced)	20.20	0.41
SAR-BM3D	22.75	0.60
ID-CNN	23.55	0.60
BD-QMAP (7×7)	21.46	0.54

BD-QMAP: Sample visual comparison

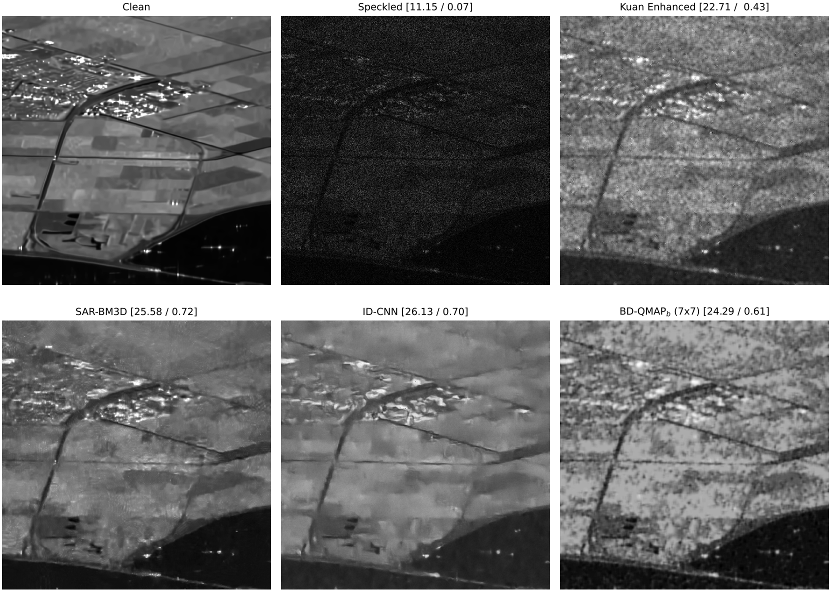

Outline

Key takeaways & future directions

- Broad applicability of the despeckler for different noise settings.

Key takeaways & future directions

- Broad applicability of the despeckler for different noise settings.
- Generalizability across diverse sources, including 1D and 2D data.

Key takeaways & future directions

- Broad applicability of the despeckler for different noise settings.
- Generalizability across diverse sources, including 1D and 2D data.
- Patch-wise despeckling offers room for further investigation:

	Patch-wise		Joint
Piecewise Constant	14.14 dB	$\xrightarrow{+\,8\text{ dB}}$	22.25 dB
Set11 Images (average)	21.46 dB	$\xrightarrow{+\,?\text{ dB}}$	? dB

Key takeaways & future directions

- Broad applicability of the despeckler for different noise settings.
- Generalizability across diverse sources, including 1D and 2D data.
- Patch-wise despeckling offers room for further investigation:

	Patch-wise		Joint
Piecewise Constant	14.14 dB	$\xrightarrow{+\,8\text{ dB}}$	22.25 dB
Set11 Images (average)	21.46 dB	$\xrightarrow{+\,?\text{ dB}}$	? dB

- Enhanced weight-learning strategies for image data have yet to be explored.

References

Goodman, Joseph W [1976]. "Some fundamental properties of speckle". In: *Journal of the Optical Society of America.*

Zafari, Ali and Jalali, Shirin [2025]. "Bayesian despeckling of structured sources". In: *arXiv preprint arXiv:2501.11860.*

Zhou, Wenda and Wabnig, Joachim and Jalali, Shirin [2023]. "Bayesian denoising of structured sources and its implications on learning-based denoising". In: *Information and Inference: A Journal of the IMA.*