Dark Noise Diffusion: Synthèse de bruit par modèle de diffusion pour la restauration d'images en basse luminosité

Liying Lu, Raphaël Achddou, Sabine Süsstrunk

IVRL, EPFL

Low-Light Denoising


Noisy image

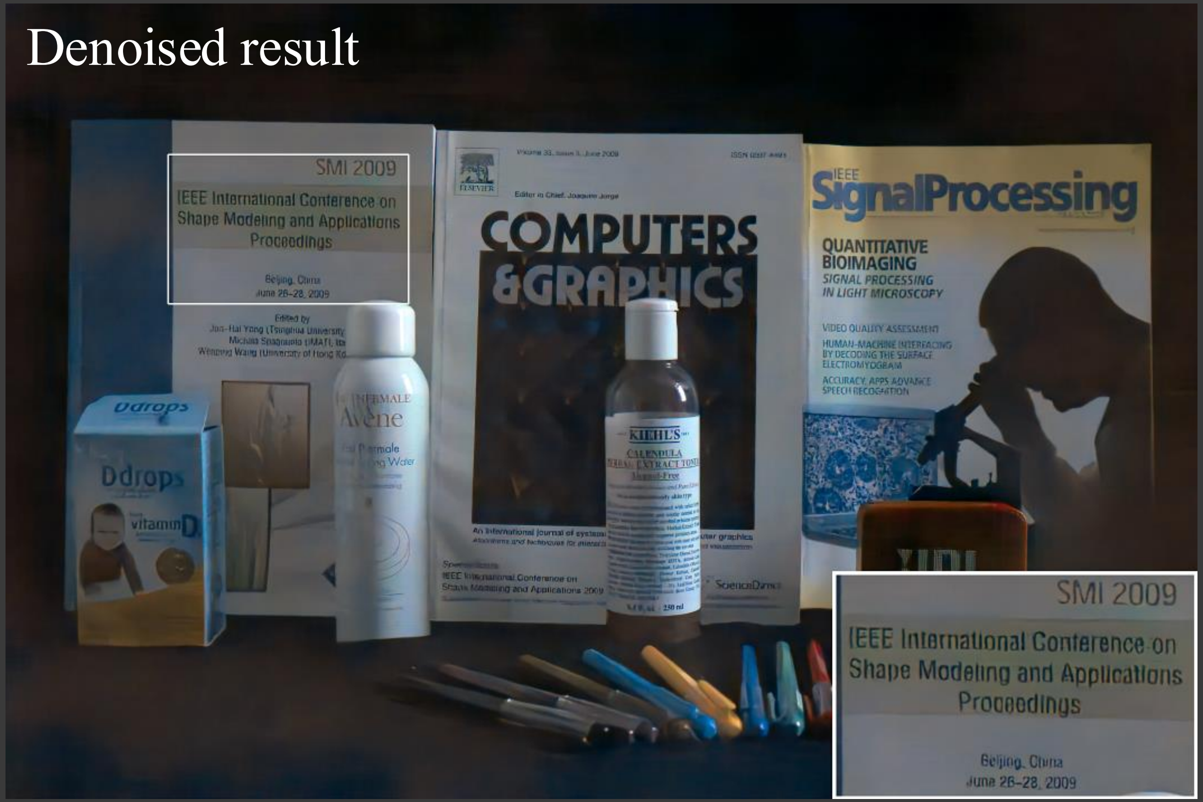
Denoised result

Noisy image from: Learning to see in the dark, Chen *et al*, CVPR2018

Low-Light Denoising

Noisy image

Denoised result

Camera noise is far from an iid Gaussian distribution!

Noisy image from: Learning to see in the dark, Chen *et al*, CVPR2018

Noise Source

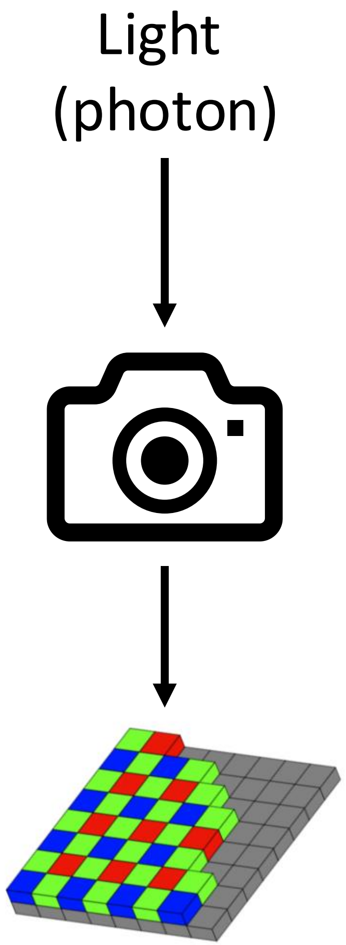

Photon (shot) noise

Dark current noise

Read noise

Banding pattern noise

Quantization noise

…

Real Paired Data

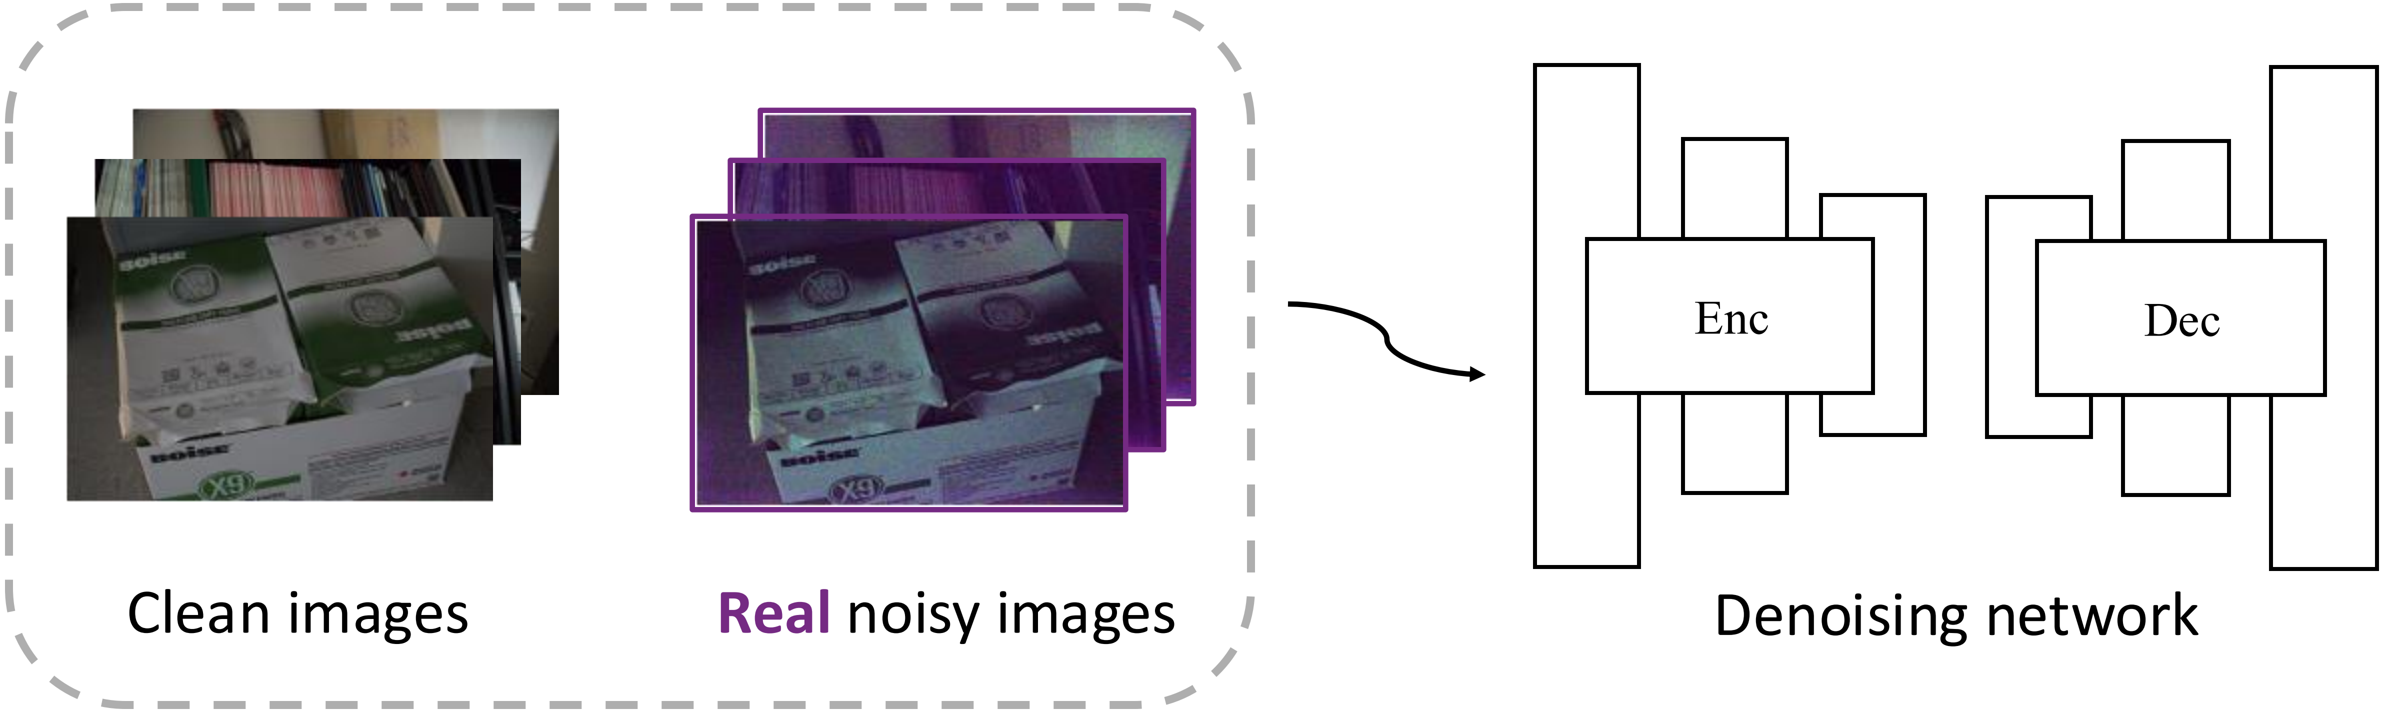

How Do We Collect Those Paired Data?

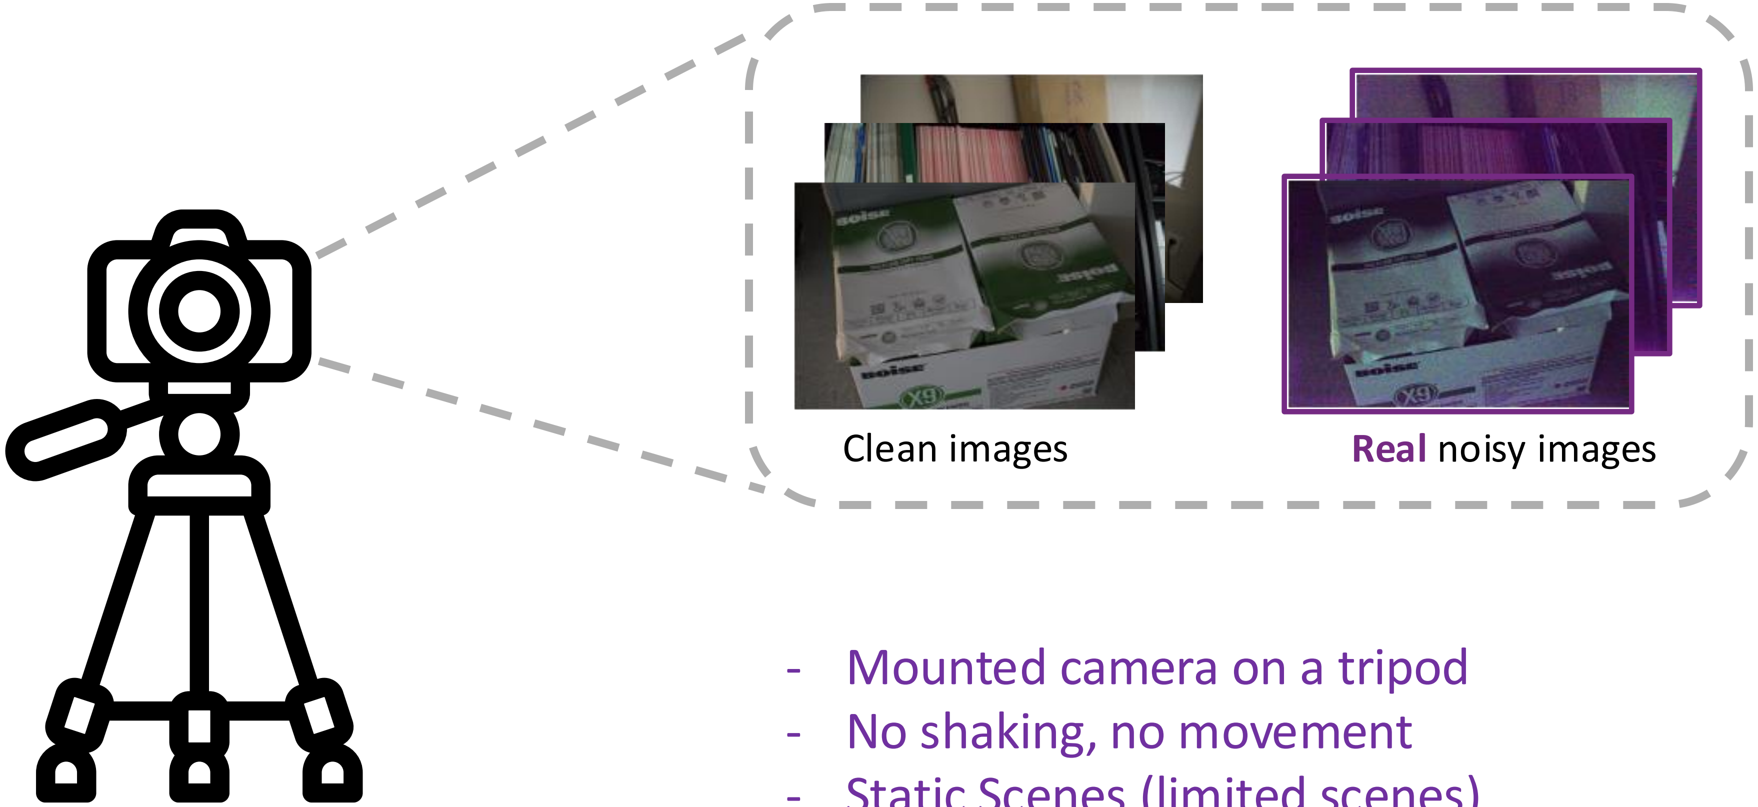

- Mounted camera on a tripod
- No shaking, no movement
- Static Scenes (limited scenes)
- Alignment postprocessing

…

Learning to see in the dark, Chen *et al*, CVPR2018
A high-quality denoising dataset for smartphone cameras, Abdelrahman *et al*, CVPR 2018.

Synthetic Paired Data

- Generate as much data as we want!
- Not limited to static scenes
- Less labor-intensive

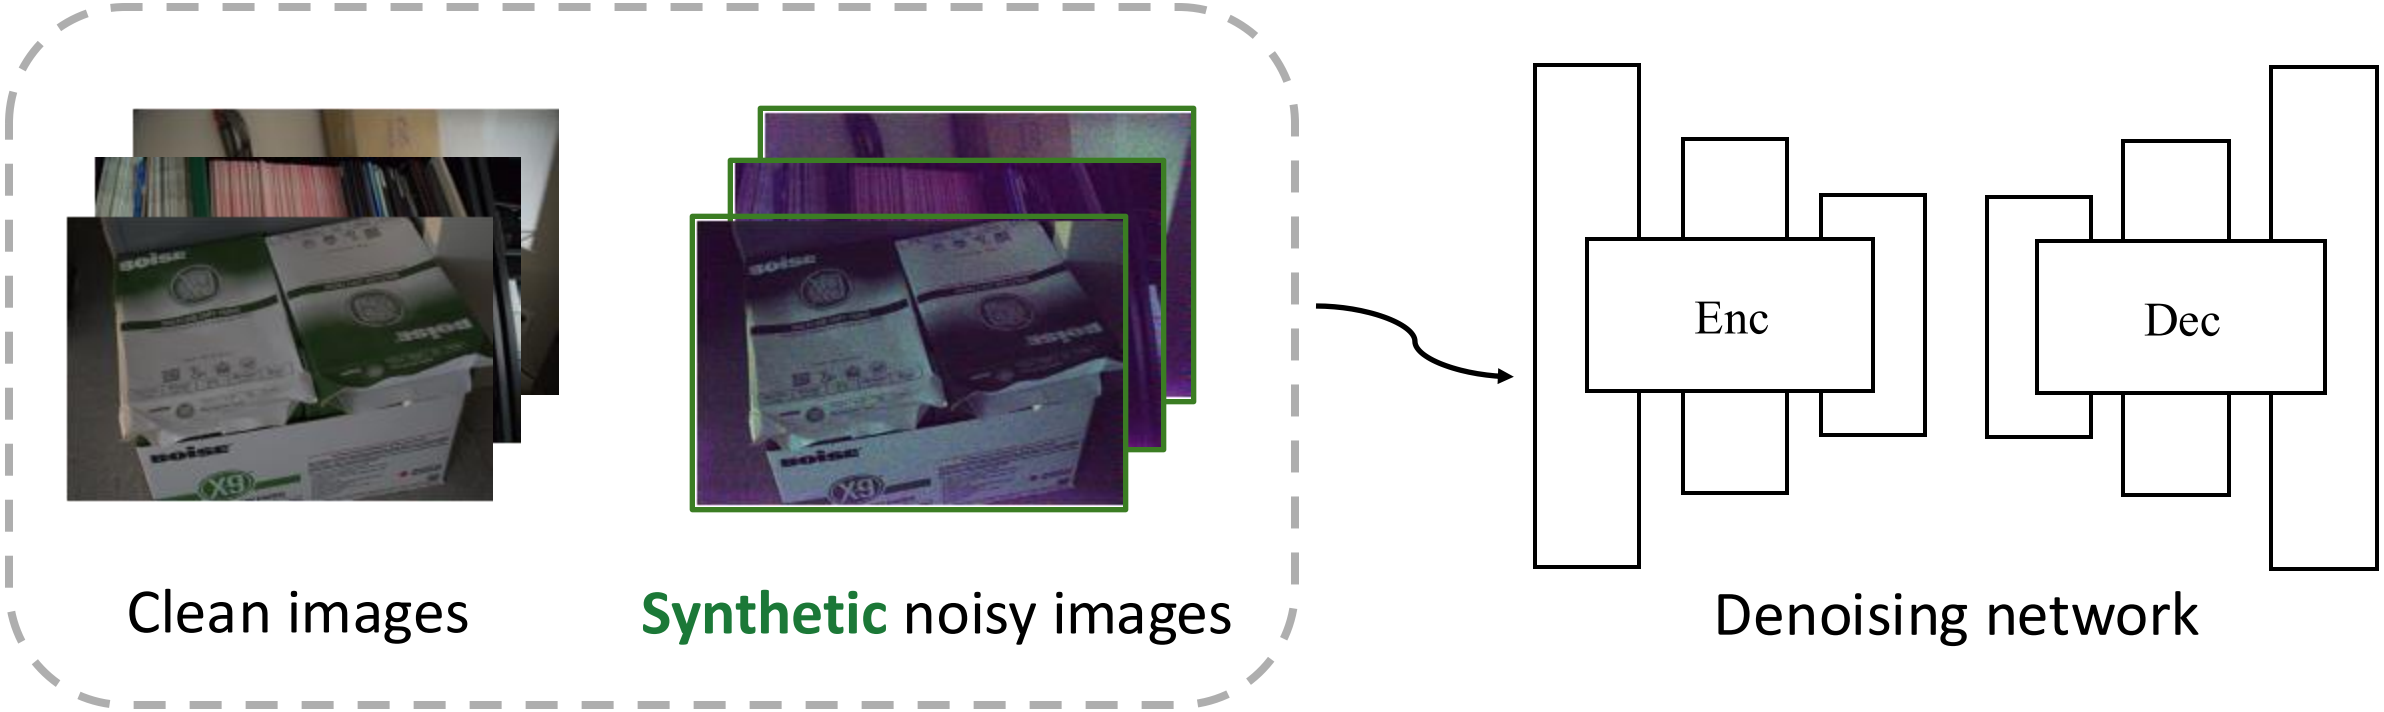

How Do We Synthesize Clean-Noisy Pairs?

Clean images

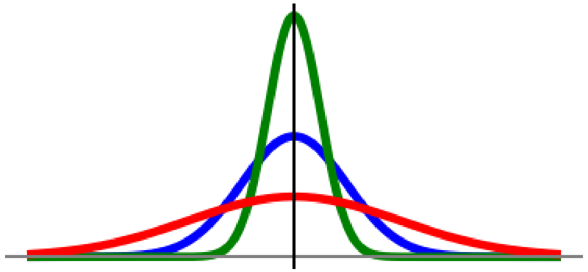

Physics-based (Poisson, Gaussian …)

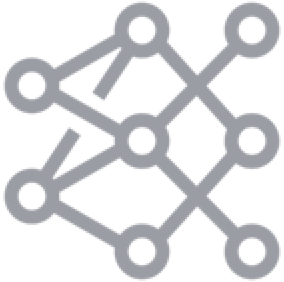

Networks (GAN, Normalizing flow…)

…

Synthetic noisy images

Practical Poissonian-Gaussian noise modeling and fitting for single-image raw-data, Alessandro *et al*, TIP2008
A Physics-based Noise Formation Model for Extreme Low-light Raw Denoising, Wei *et al*, CVPR2020
Noise Flow: Noise Modeling with Conditional Normalizing Flows, Abdelhamed *et al*, ICCV2019
Towards General Low-Light Raw Noise Synthesis and Modeling, Zhang *et al*, ICCV2023
Learnability Enhancement for Low-light Raw Denoising: A Data Perspective, Feng *et al*, TPAMI 2024

How Do We Synthesize Clean-Noisy Pairs?

Clean images

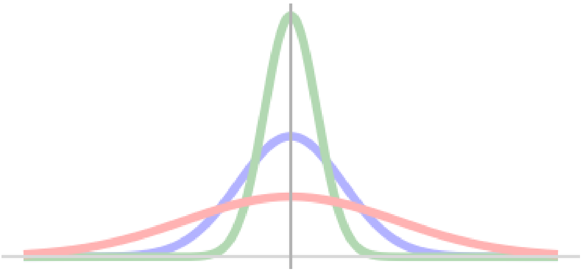

Physics-based (Poisson, Gaussian …)

Networks (GAN, Normalizing flow…)

…

Synthetic noisy images

Practical Poissonian-Gaussian noise modeling and fitting for single-image raw-data, Alessandro *et al*, TIP2008
A Physics-based Noise Formation Model for Extreme Low-light Raw Denoising, Wei *et al*, CVPR2020
Noise Flow: Noise Modeling with Conditional Normalizing Flows, Abdelhamed *et al*, ICCV2019
Towards General Low-Light Raw Noise Synthesis and Modeling, Zhang *et al*, ICCV2023
Learnability Enhancement for Low-light Raw Denoising: A Data Perspective, Feng *et al*, TPAMI 2024

How Do We Synthesize Clean-Noisy Pairs?

Clean images

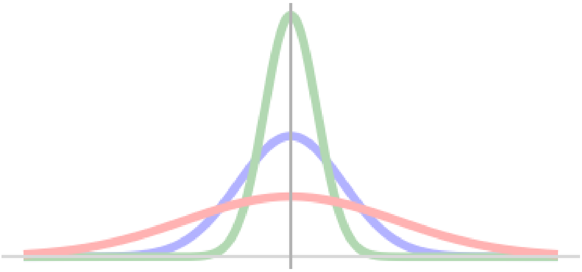

Physics-based (Poisson, Gaussian …)

Diffusion models?

Synthetic noisy images

Practical Poissonian-Gaussian noise modeling and fitting for single-image raw-data, Alessandro *et al*, TIP2008
A Physics-based Noise Formation Model for Extreme Low-light Raw Denoising, Wei *et al*, CVPR2020
Noise Flow: Noise Modeling with Conditional Normalizing Flows, Abdelhamed *et al*, ICCV2019
Towards General Low-Light Raw Noise Synthesis and Modeling, Zhang *et al*, ICCV2023
Learnability Enhancement for Low-light Raw Denoising: A Data Perspective, Feng *et al*, TPAMI 2024

Framework

Diffusion Models basics

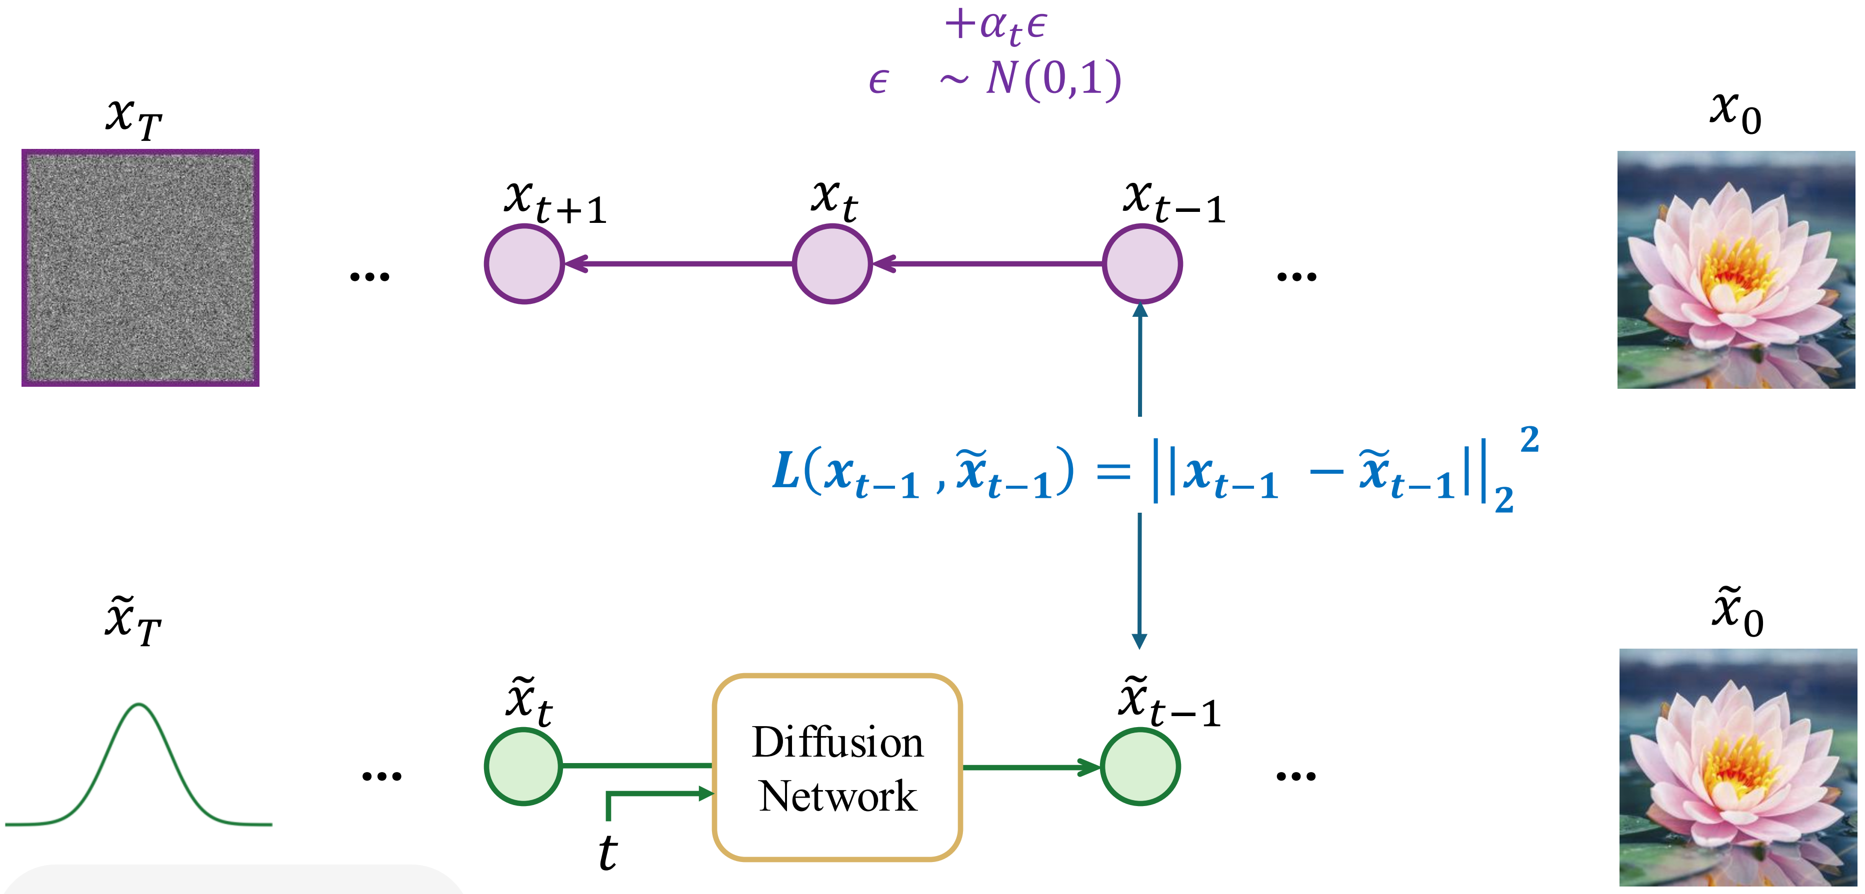

t: diffusion timestep

← Forward Process

→ Reverse Process

Diffusion Framework

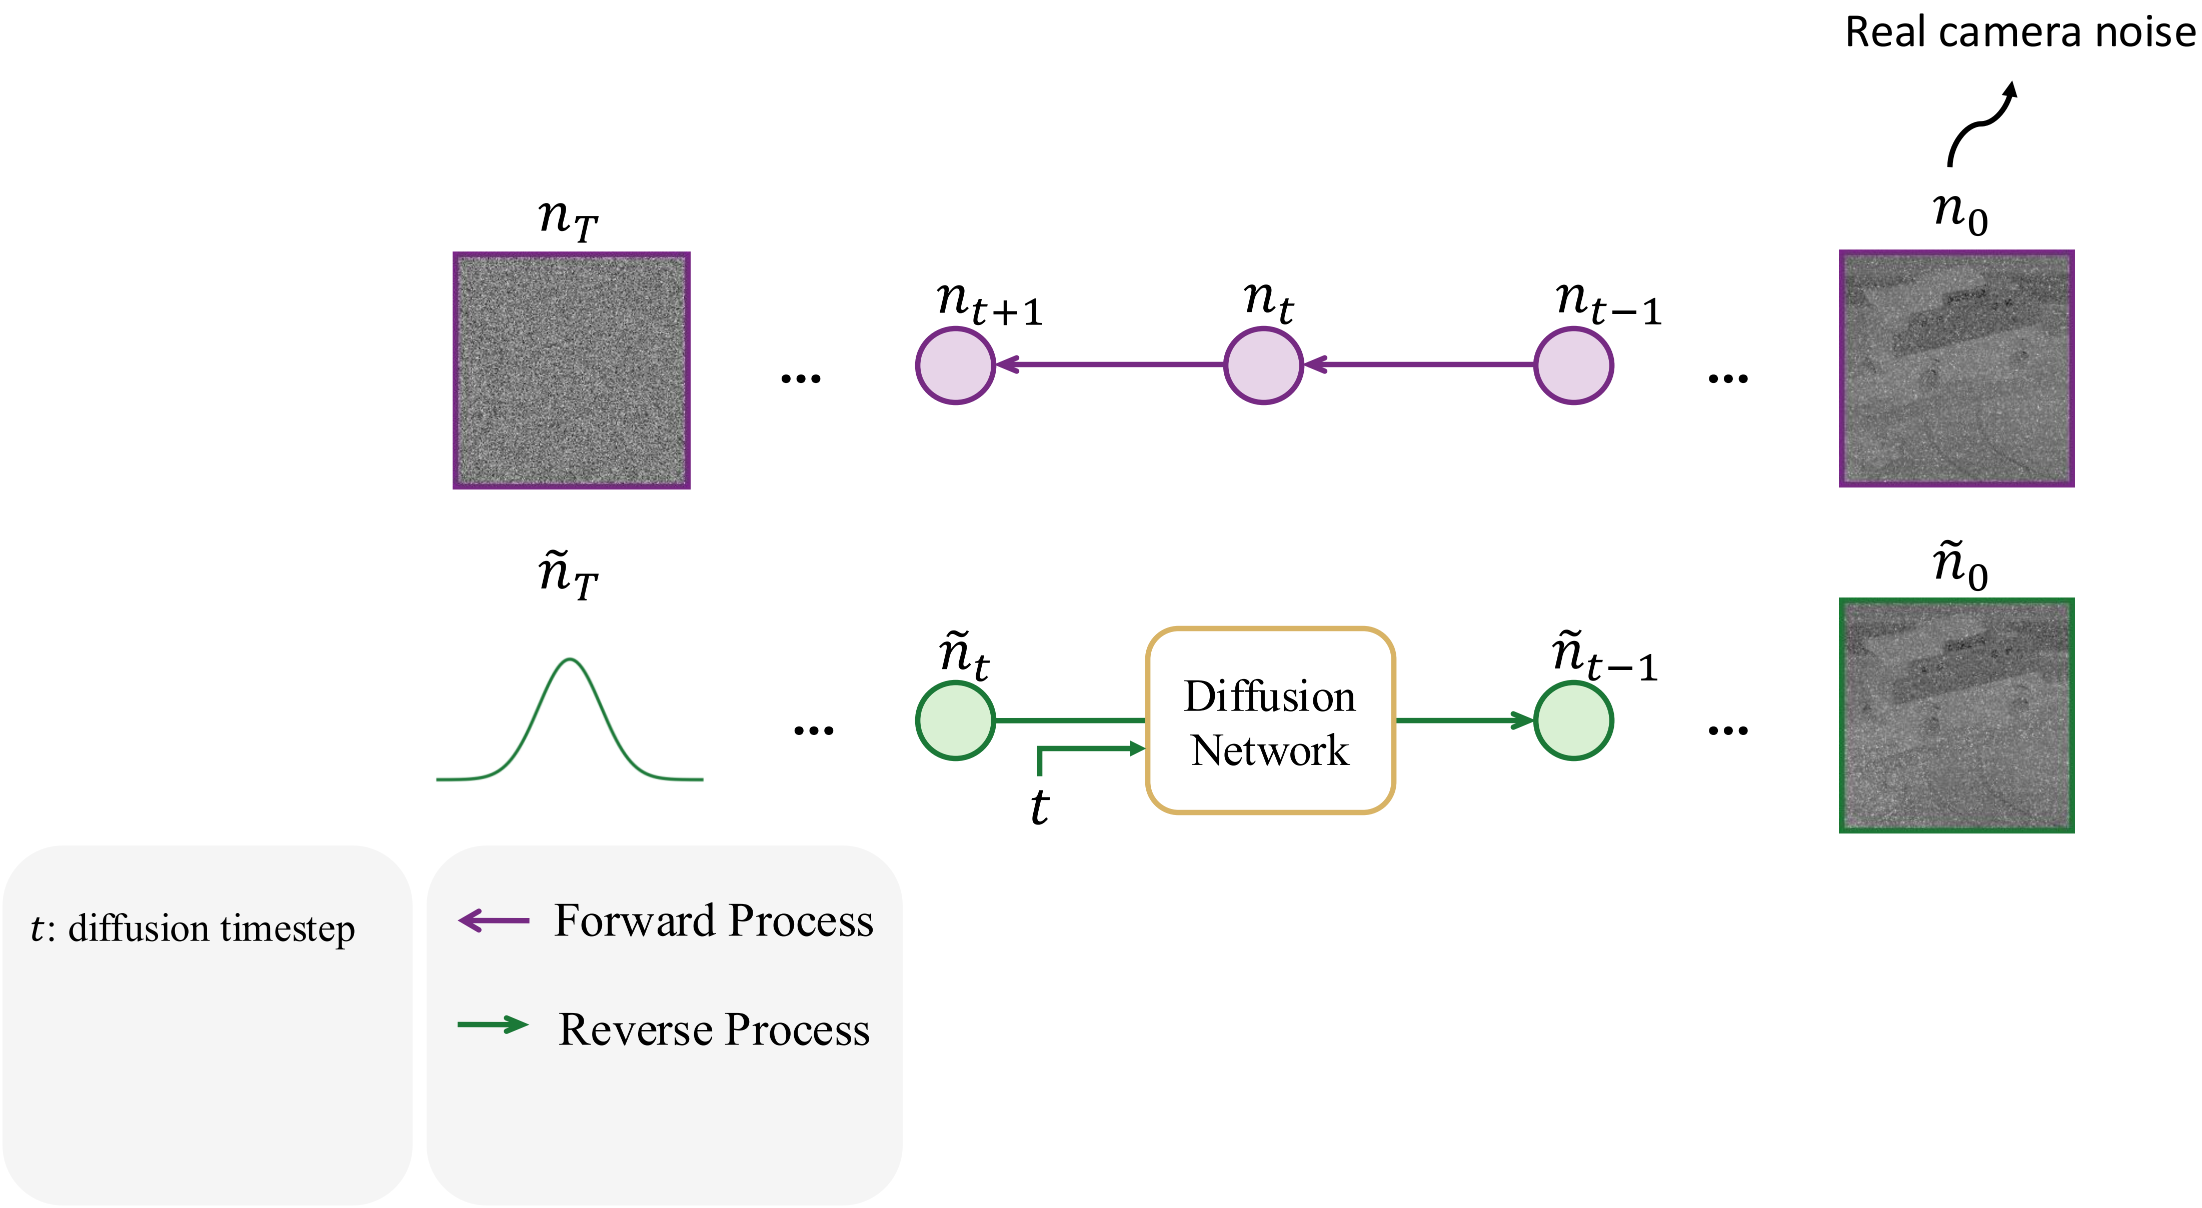

Diffusion Framework

Signal dependence | Camera setting aware | Spatially-correlated noise modeling

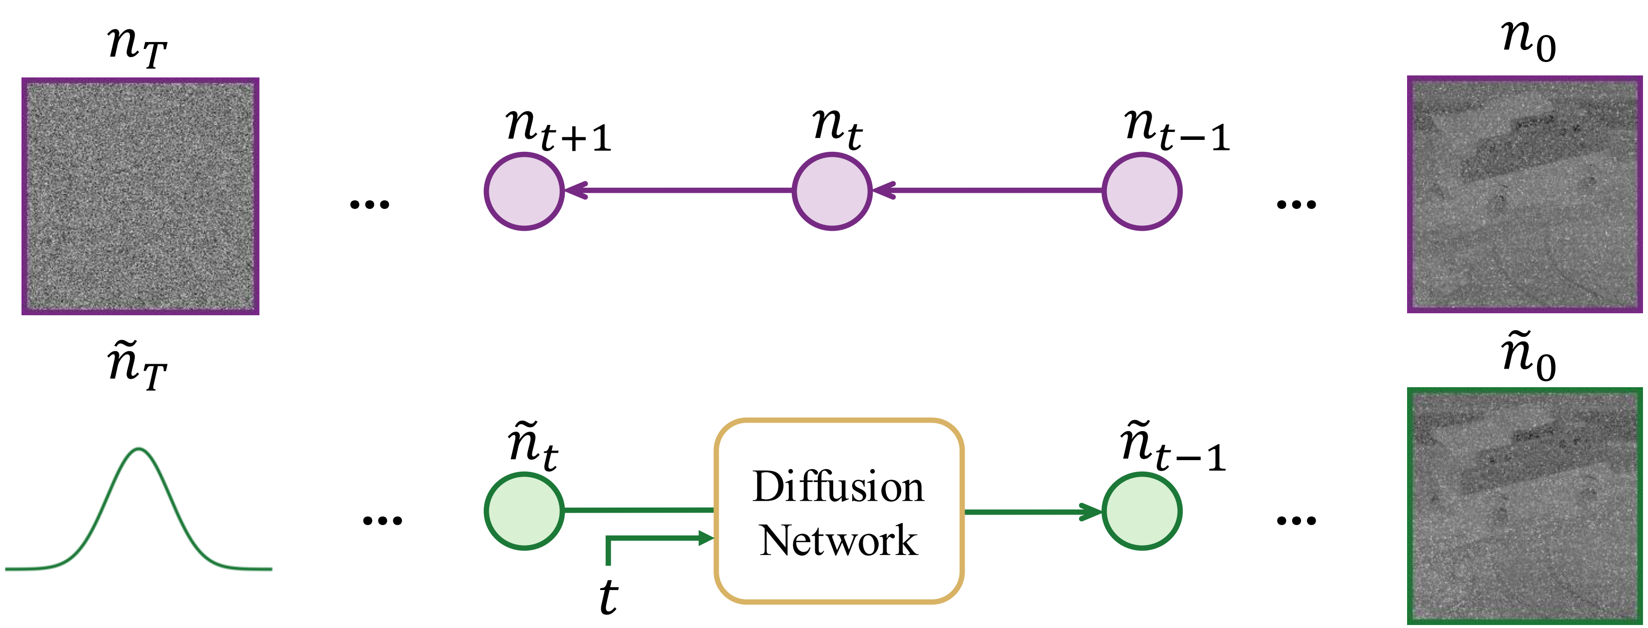

t: diffusion timestep

Forward Process

Reverse Process

Diffusion Framework

Signal dependence

Camera setting aware

Spatially-correlated noise modeling

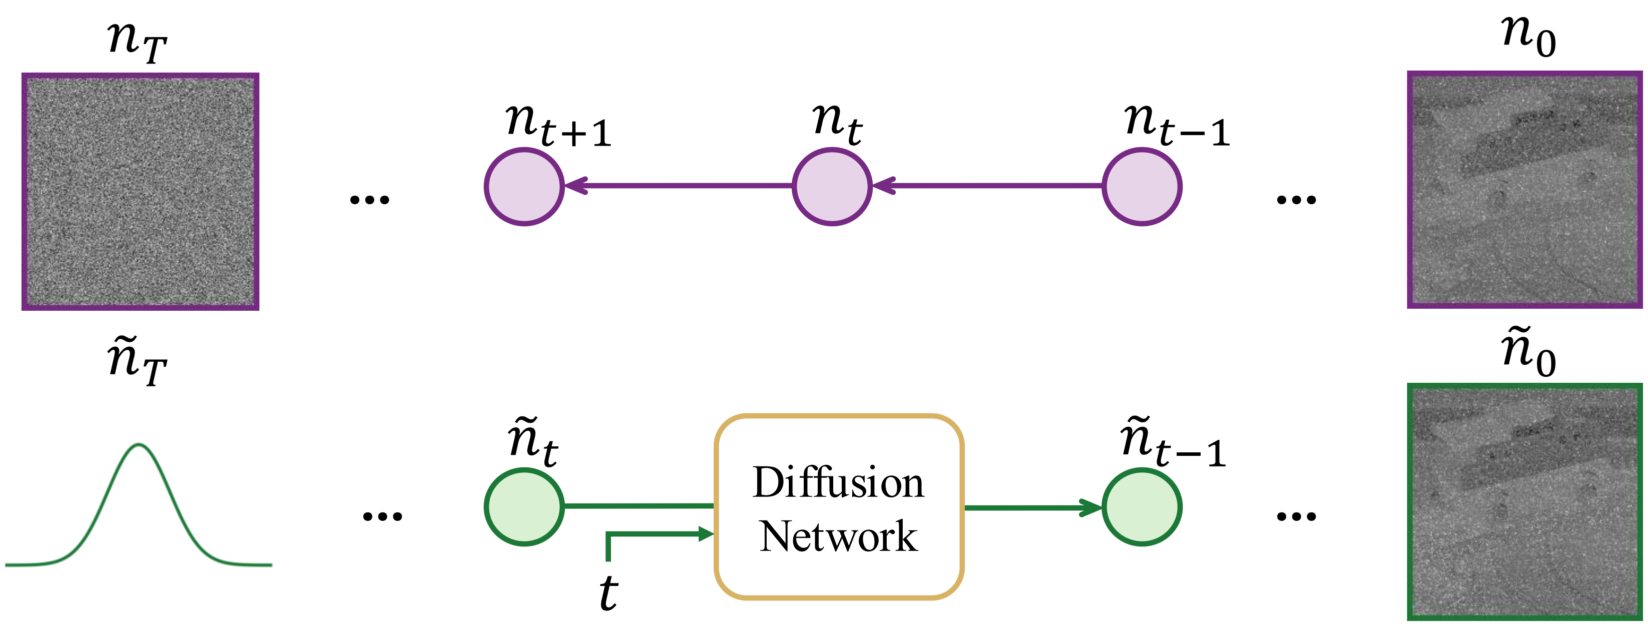

t: diffusion timestep

Forward Process

Reverse Process

Diffusion Framework

Signal dependence ✓

Camera setting aware

Spatially-correlated noise modeling

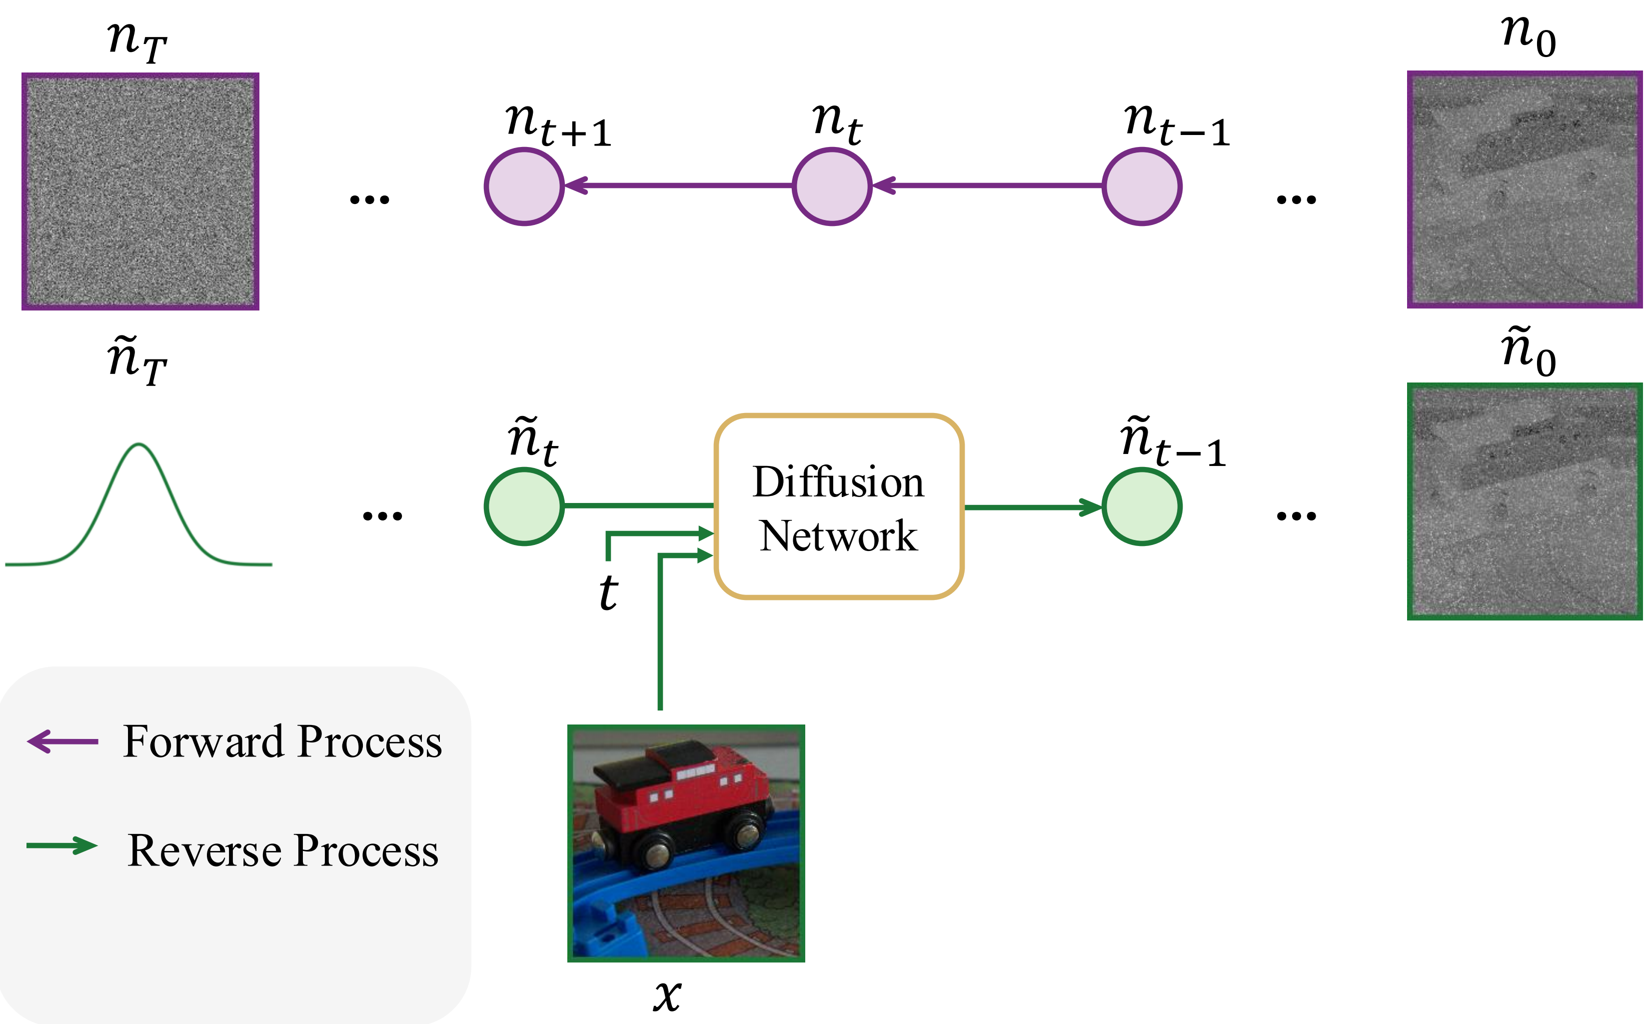

t: diffusion timestep
x: clean image

Forward Process

Reverse Process

Diffusion Framework

Signal dependence ✓ | Camera setting aware | Spatially-correlated noise modeling

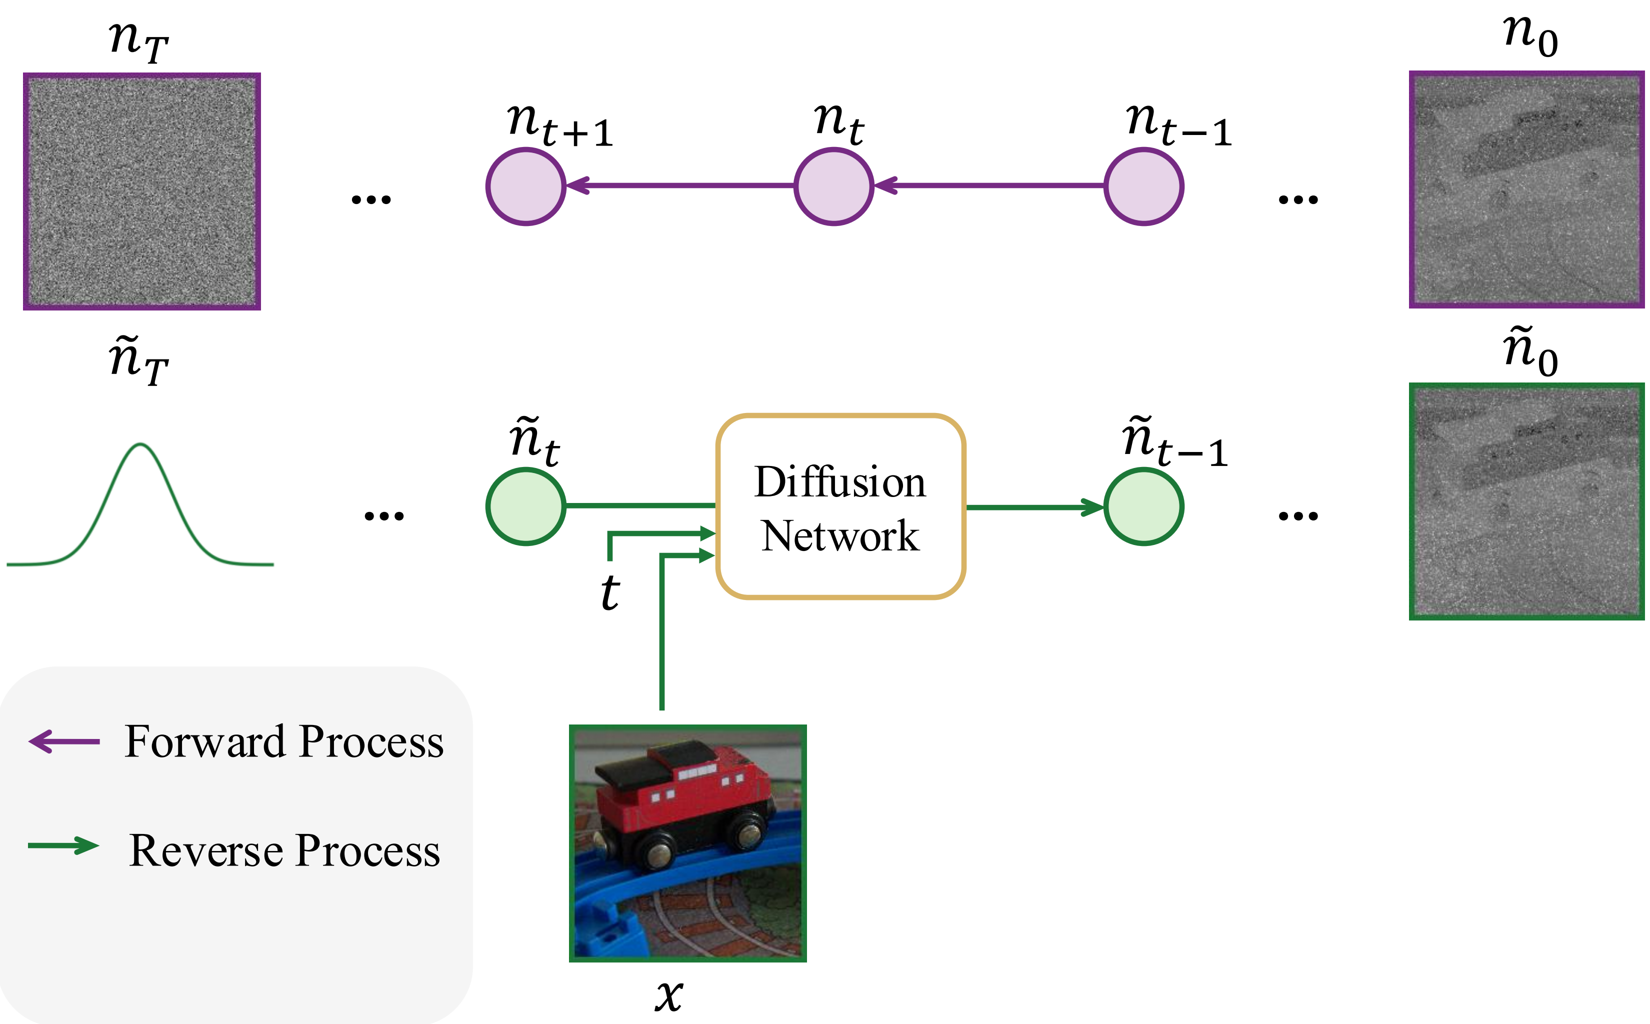

t: diffusion timestep
x: clean image

Diffusion Framework

Signal dependence ✓ Camera setting aware ✓ Spatially-correlated noise modeling

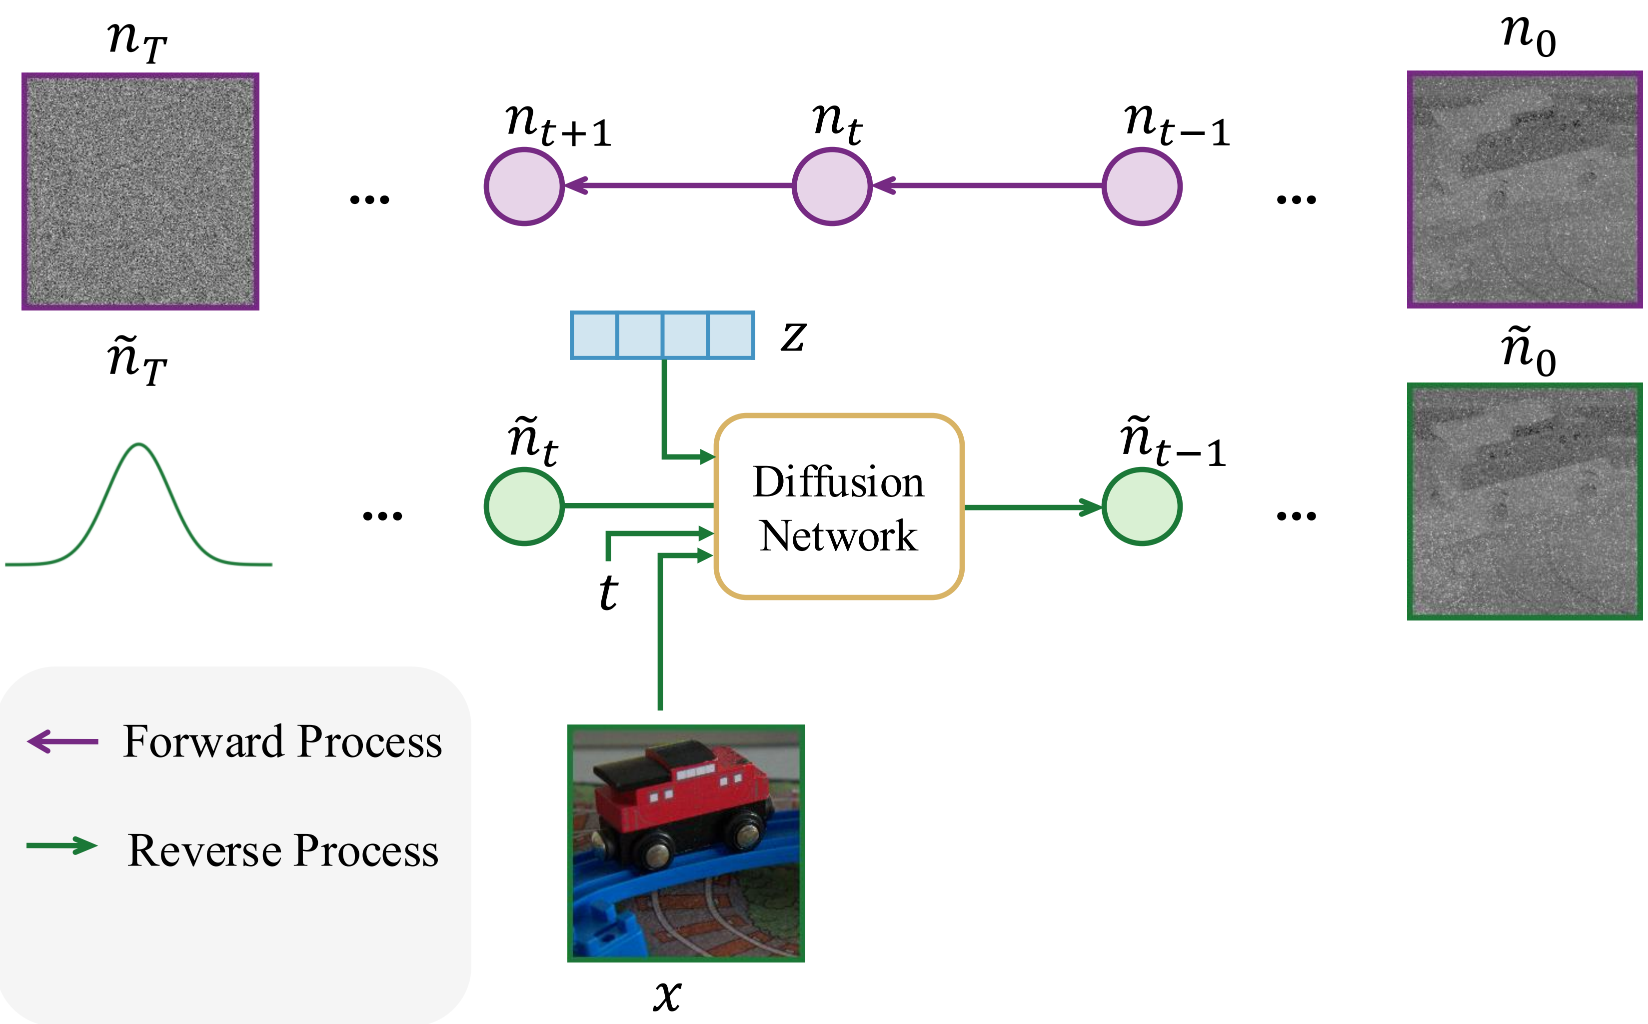

t: diffusion timestep
x: clean image
z: camera setting

← Forward Process
→ Reverse Process

Diffusion Framework

Signal dependence ✓ Camera setting aware ✓ Spatially-correlated noise modeling

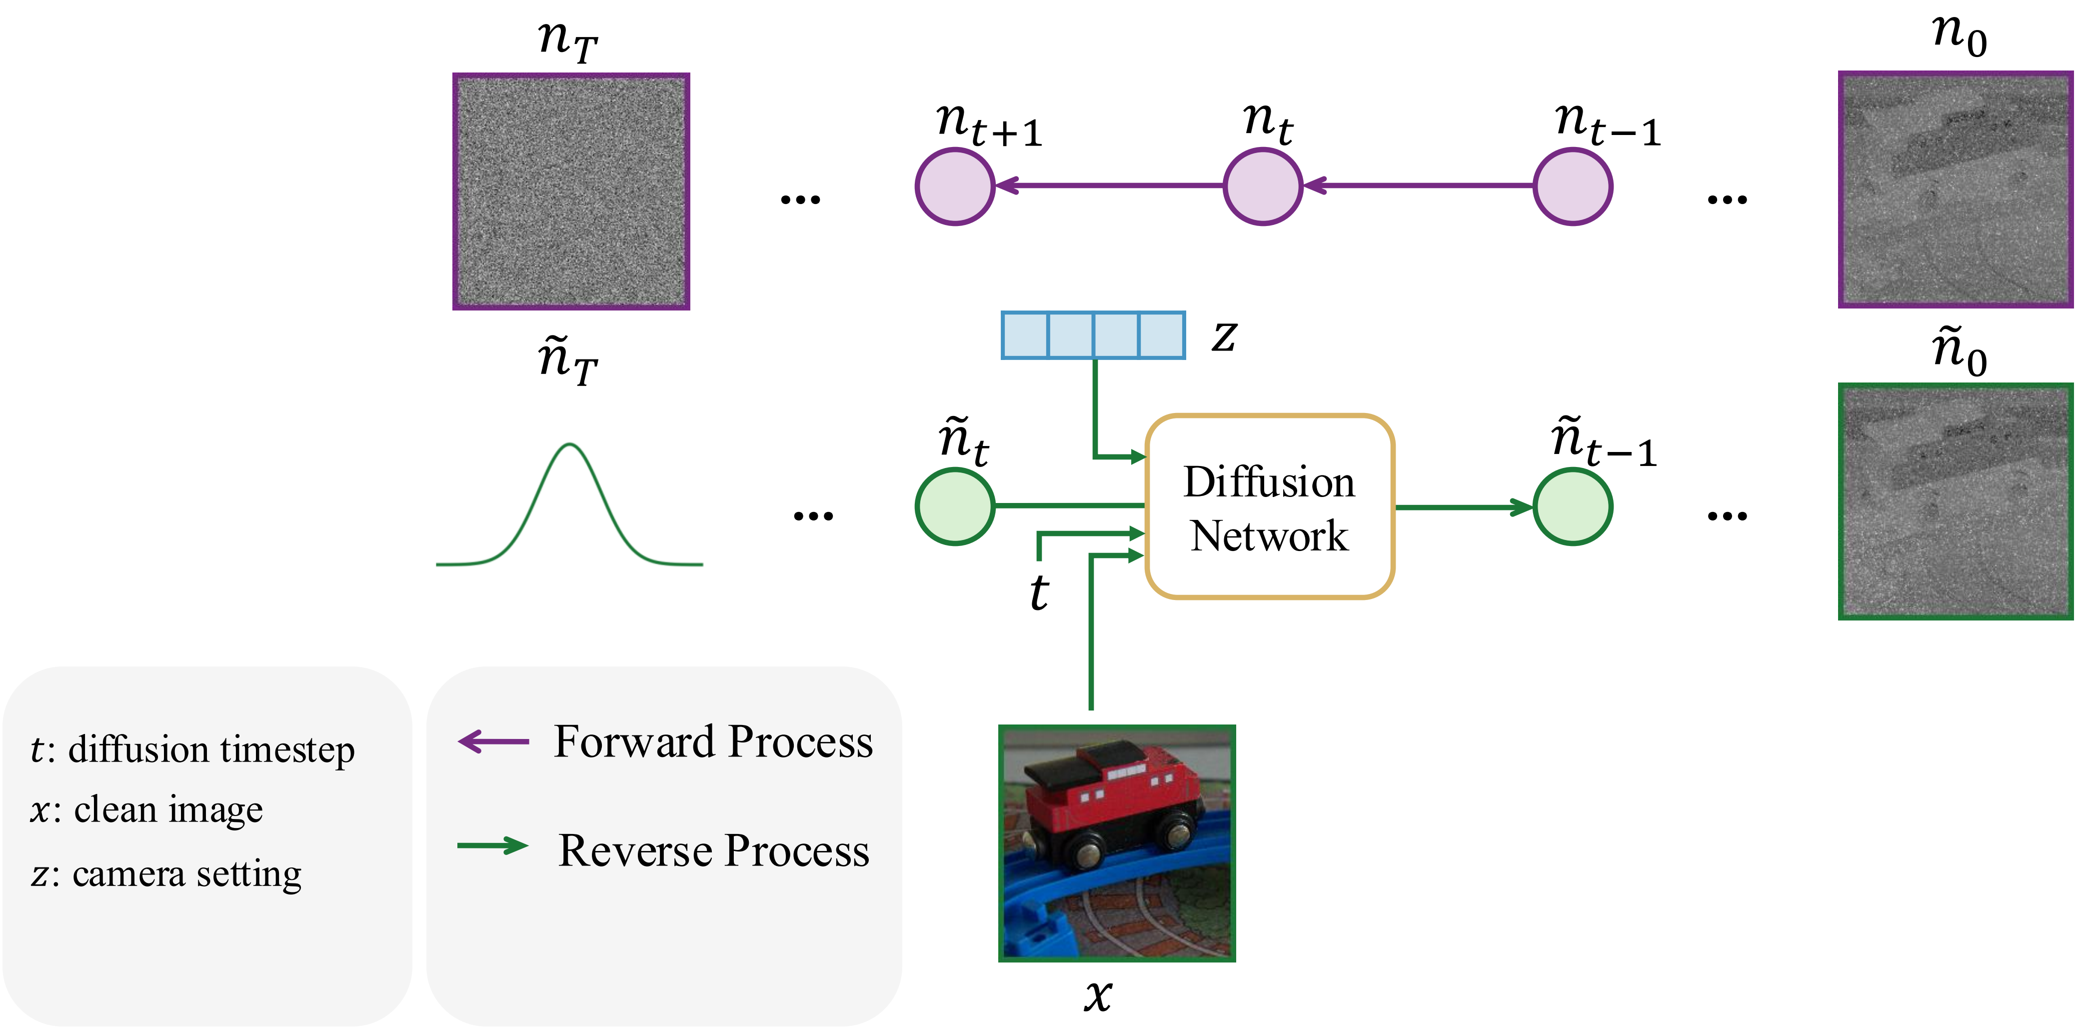

Diffusion Framework

Spatially-correlated noise

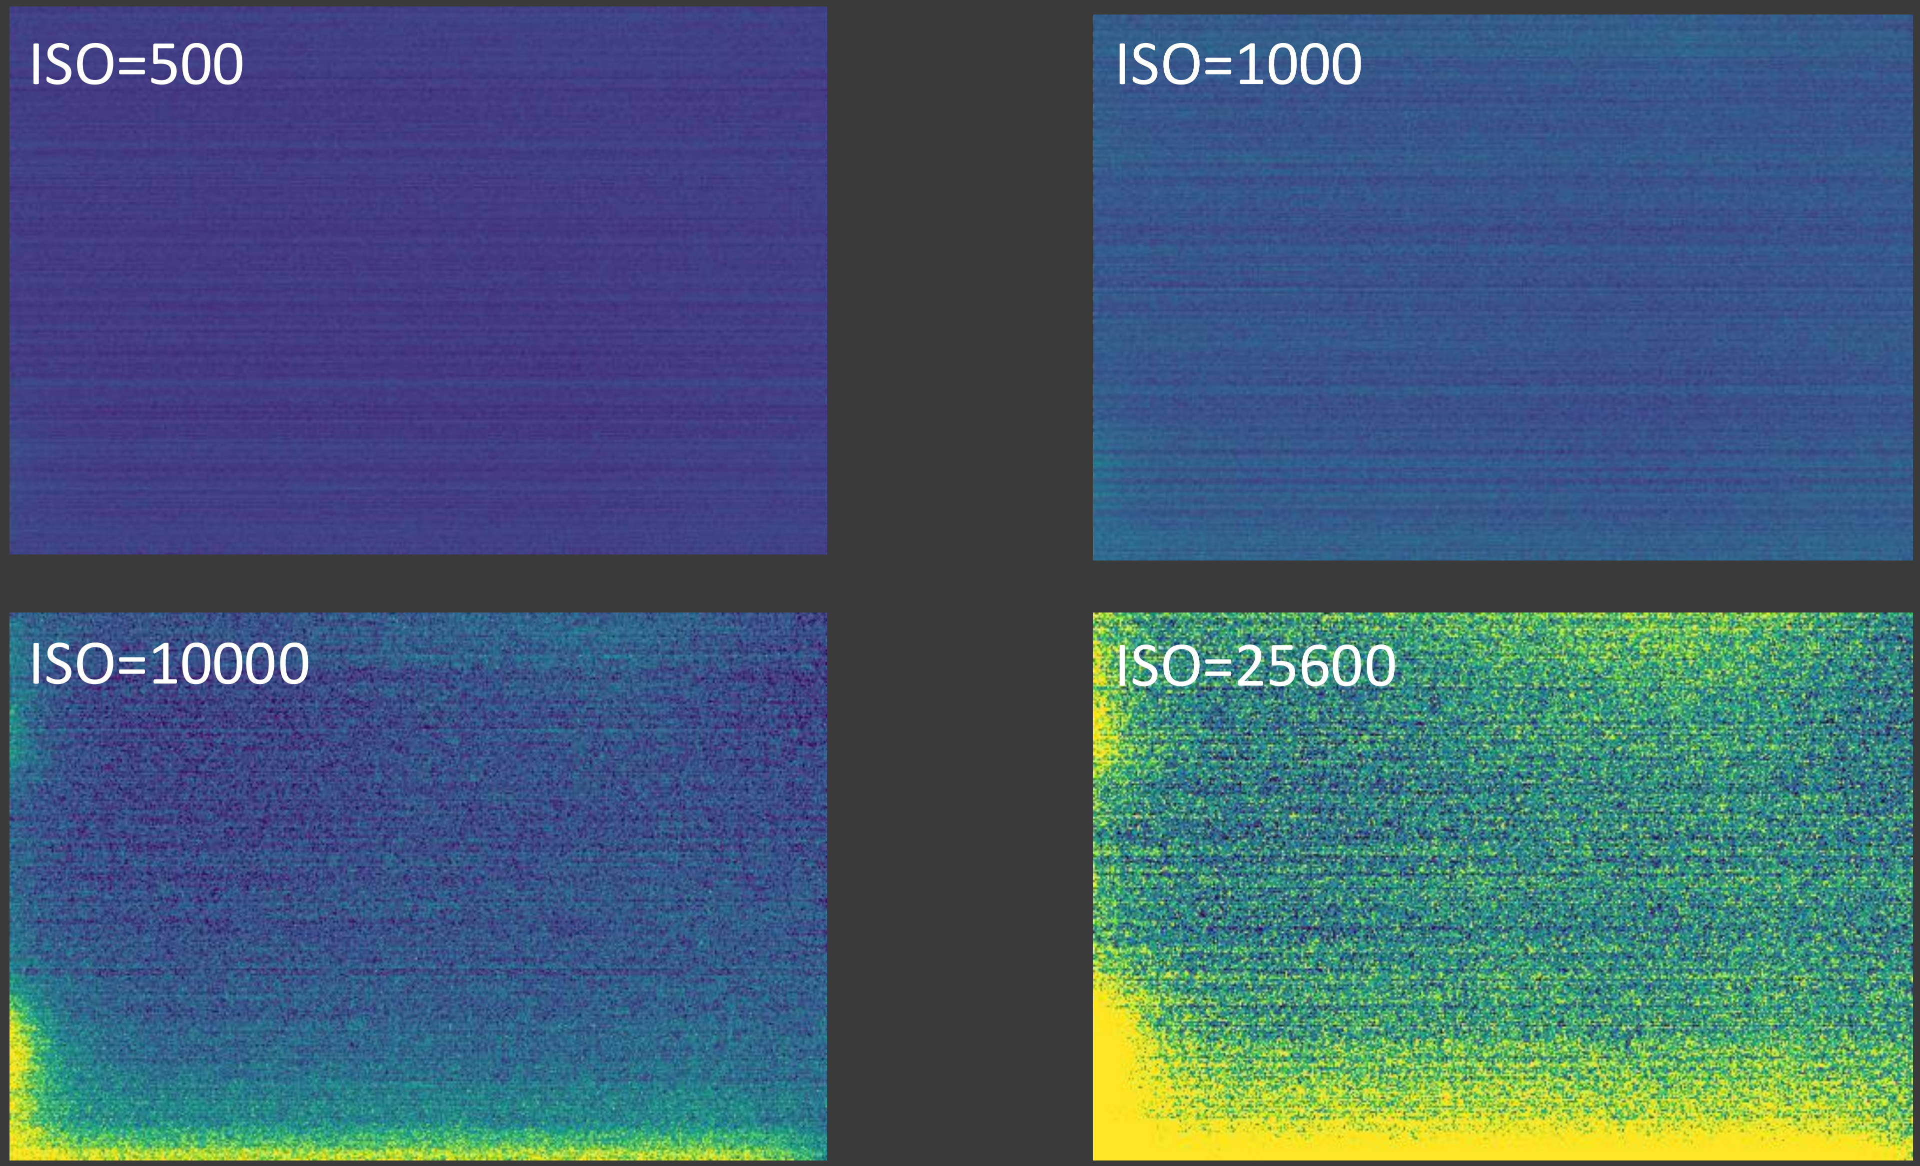

Diffusion Framework

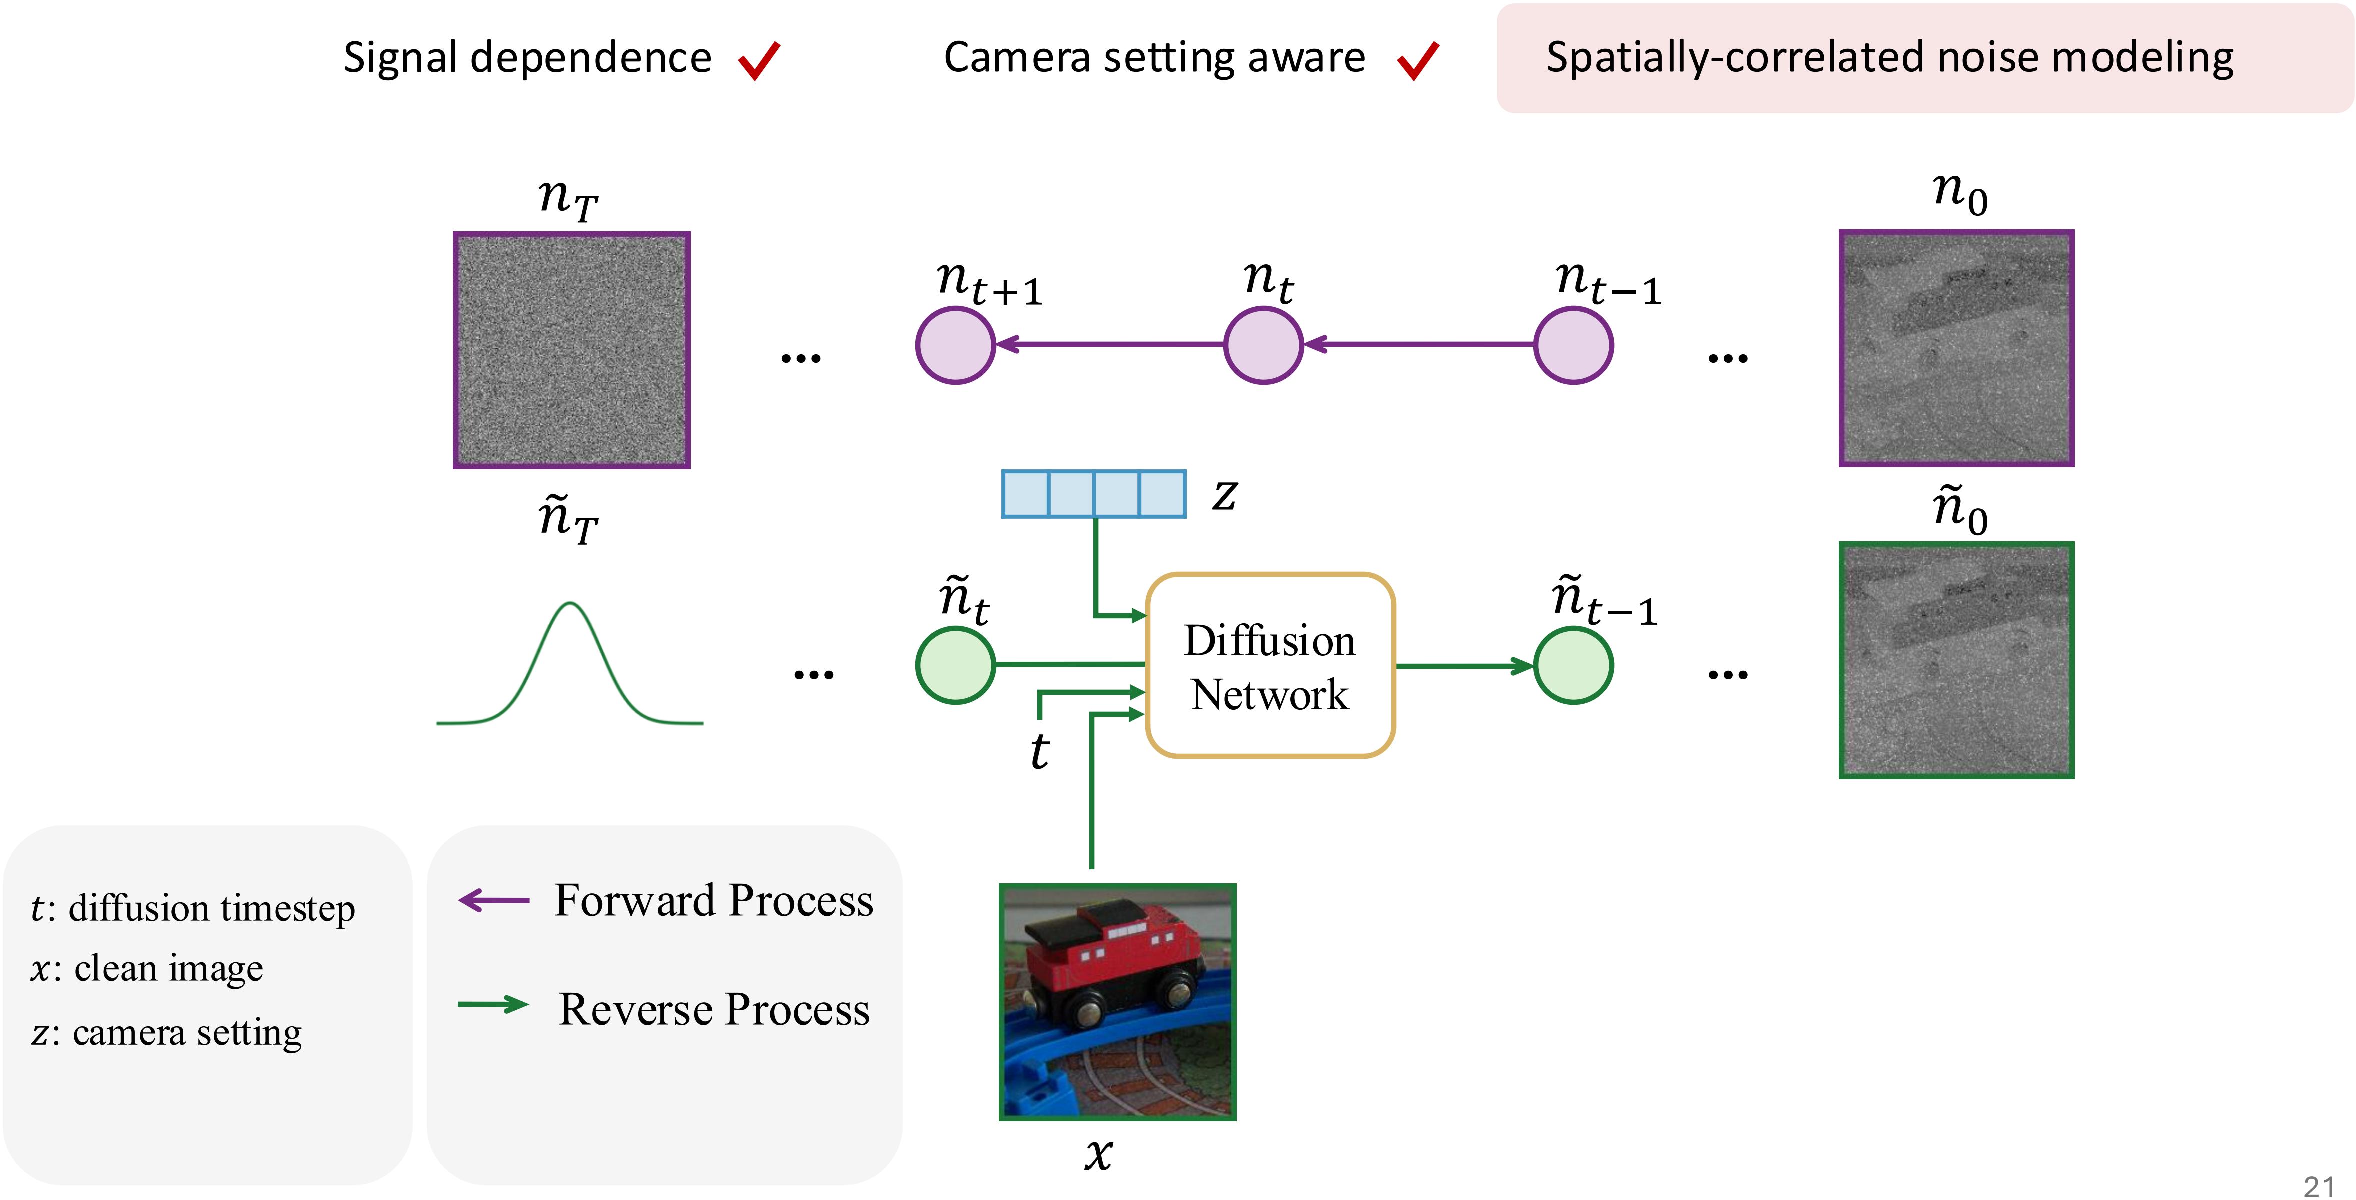

Diffusion Framework

Signal dependence ✓ Camera setting aware ✓ Spatially-correlated noise modeling ✓

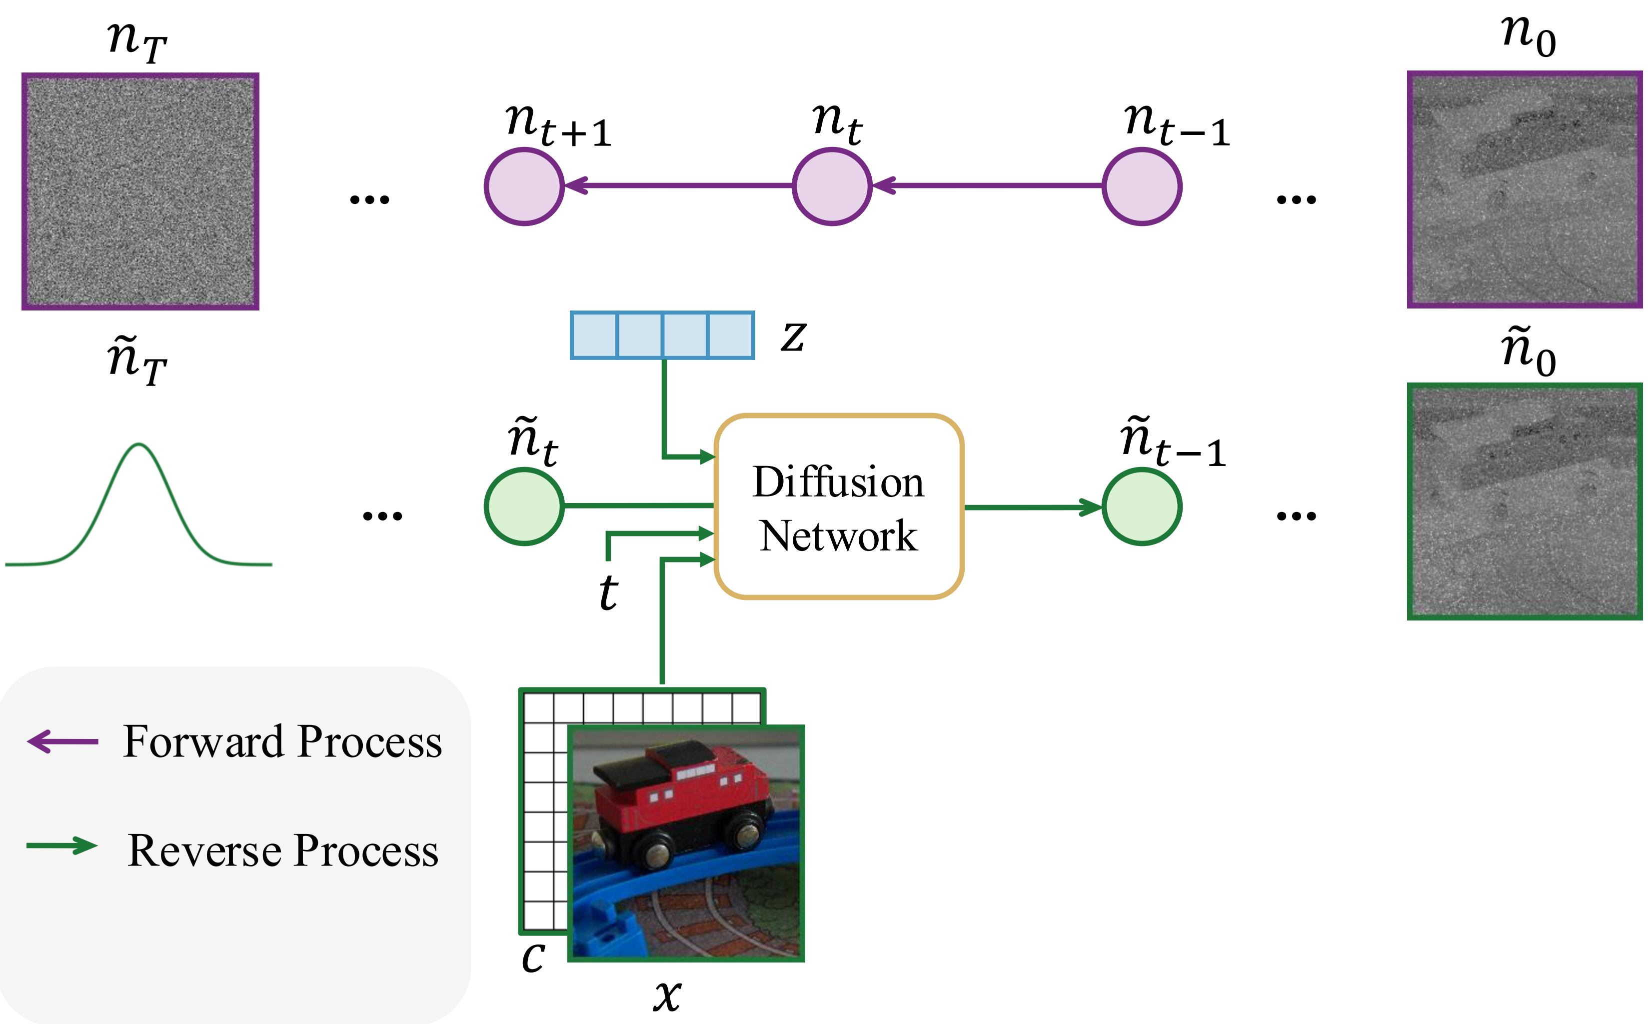

t: diffusion timestep
x: clean image
z: camera setting
c: positional encoding

← Forward Process
→ Reverse Process

Diffusion Framework

Signal dependence ✓ Camera setting aware ✓ Spatially-correlated noise modeling ✓

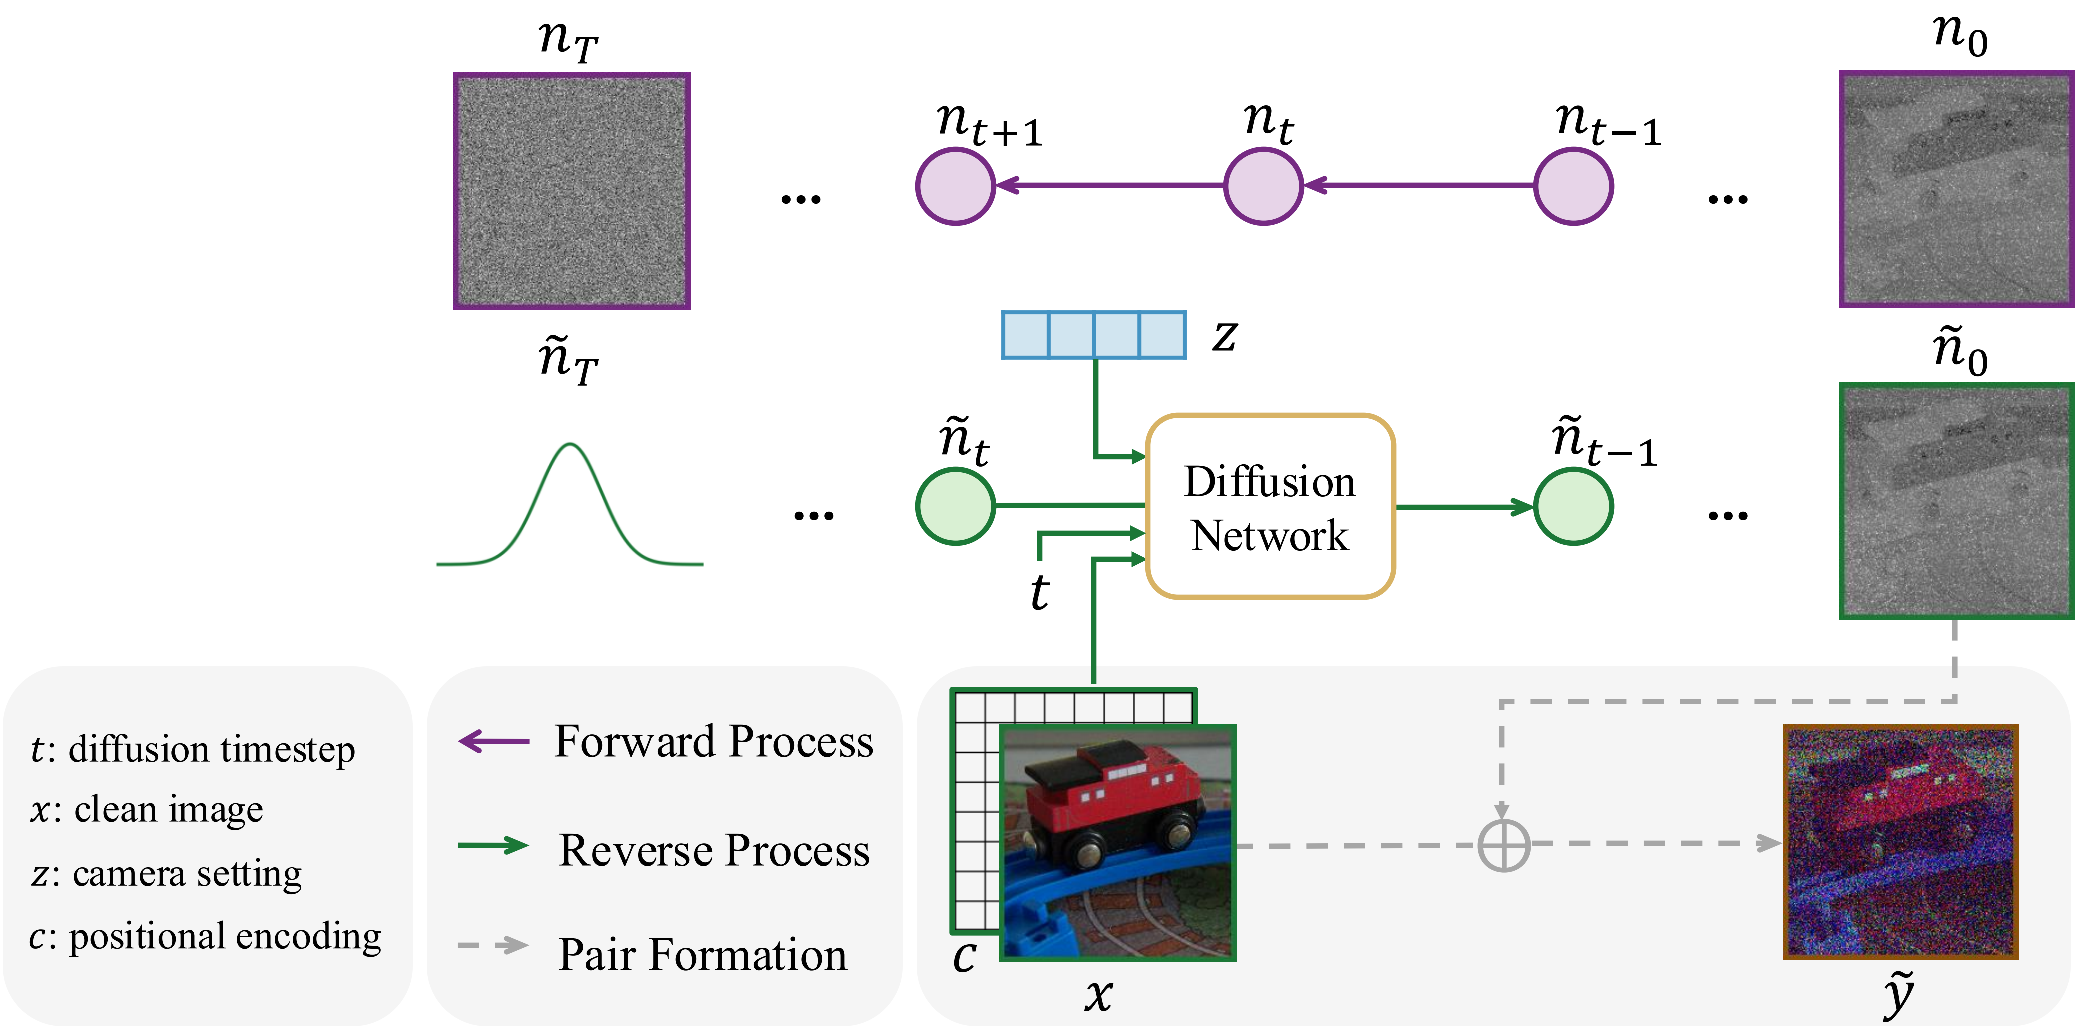

Two-Branch Network Architecture

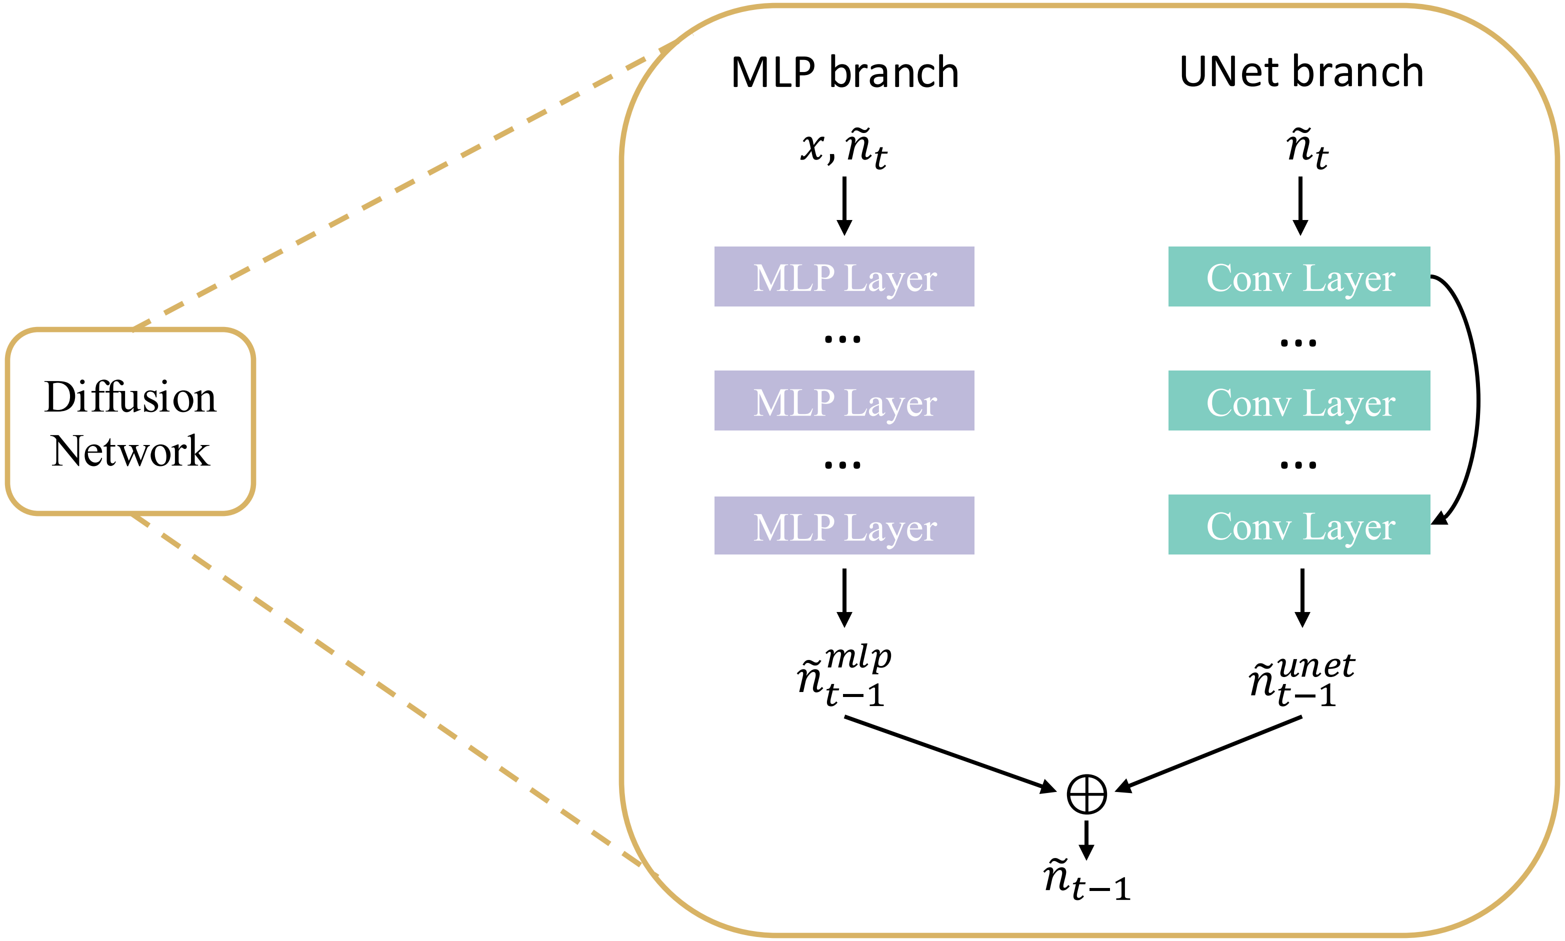

Diffusion Model Tends to Smoothen the Noise Distribution

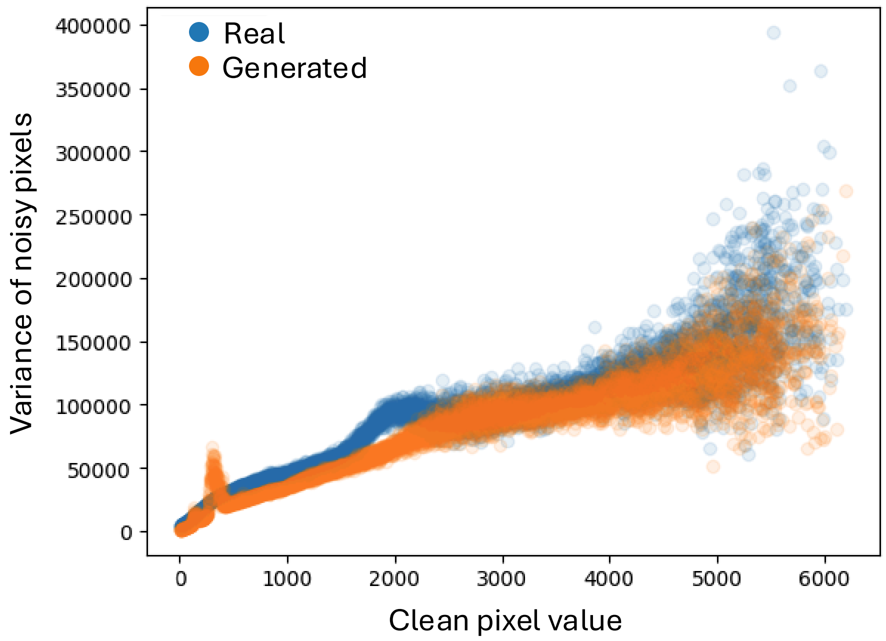

The diffusion model tends to smoothen the generated low-light noise distribution.
The variance of the generated noise data is smaller than that of the real noise data

Different Diffusion Noise Schedules on 1D Poisson Distribution

A toy example

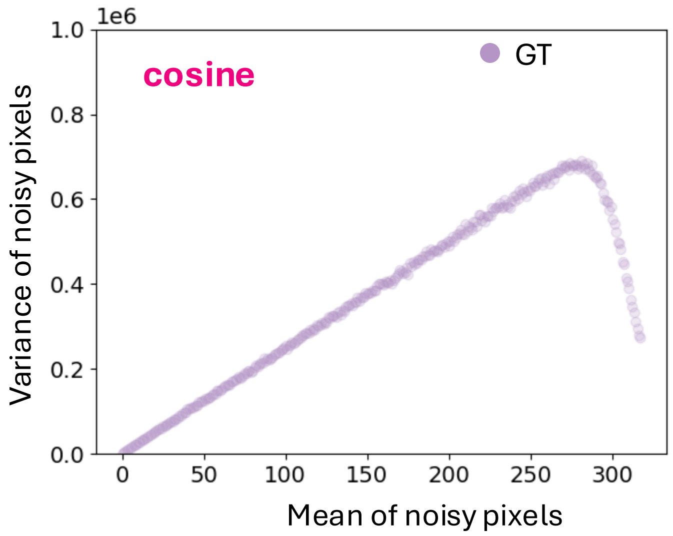

Different Diffusion Noise Schedules on 1D Poisson Distribution

A toy example

$$\begin{aligned}
\mathbf{x}_t &= \sqrt{\alpha_t}\mathbf{x}_{t-1} + \sqrt{1-\alpha_t}\boldsymbol{\epsilon}_{t-1} \\
&= \sqrt{\alpha_t\alpha_{t-1}}\mathbf{x}_{t-2} + \sqrt{1-\alpha_t\alpha_{t-1}}\bar{\boldsymbol{\epsilon}}_{t-2} \\
&= \dots \\
&= \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \boxed{\sqrt{1-\bar{\alpha}_t}}\boldsymbol{\epsilon}
\end{aligned}$$

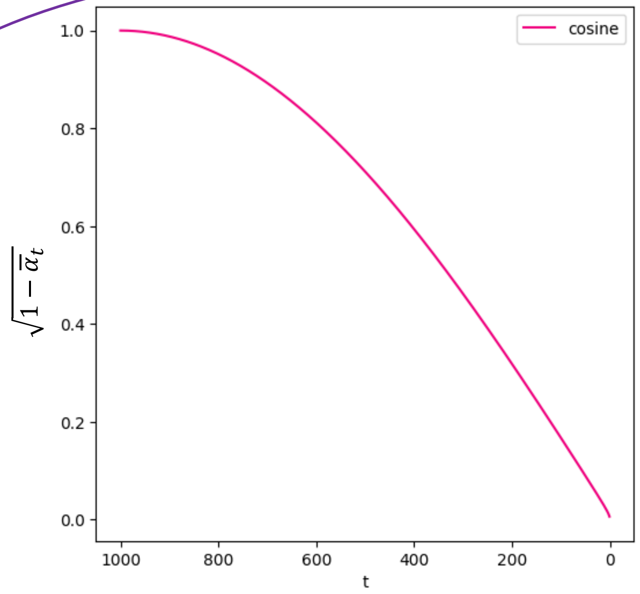

Reverse Process

x_{999} x_t x_0

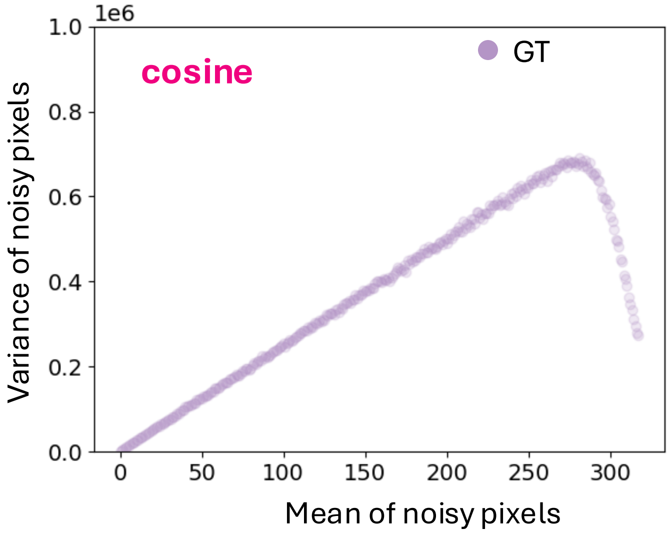

Different Diffusion Noise Schedules on 1D Poisson Distribution

A toy example

$$\begin{aligned}\mathbf{x}_t &= \sqrt{\alpha_t}\mathbf{x}_{t-1} + \sqrt{1-\alpha_t}\boldsymbol{\epsilon}_{t-1} \\ &= \sqrt{\alpha_t\alpha_{t-1}}\mathbf{x}_{t-2} + \sqrt{1-\alpha_t\alpha_{t-1}}\bar{\boldsymbol{\epsilon}}_{t-2} \\ &= \dots \\ &= \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \boxed{\sqrt{1-\bar{\alpha}_t}}\boldsymbol{\epsilon}\end{aligned}$$

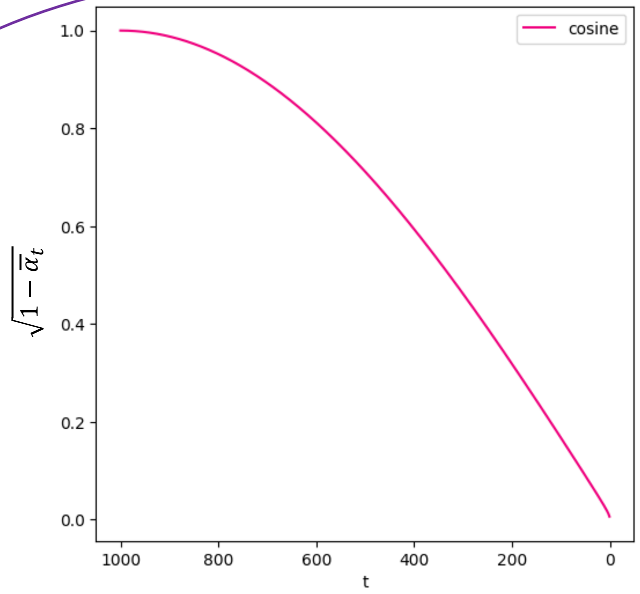

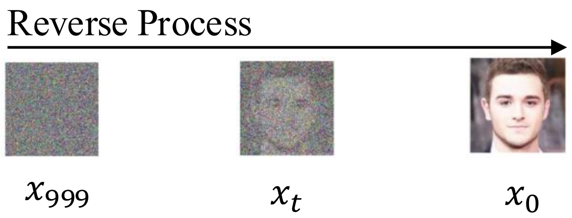

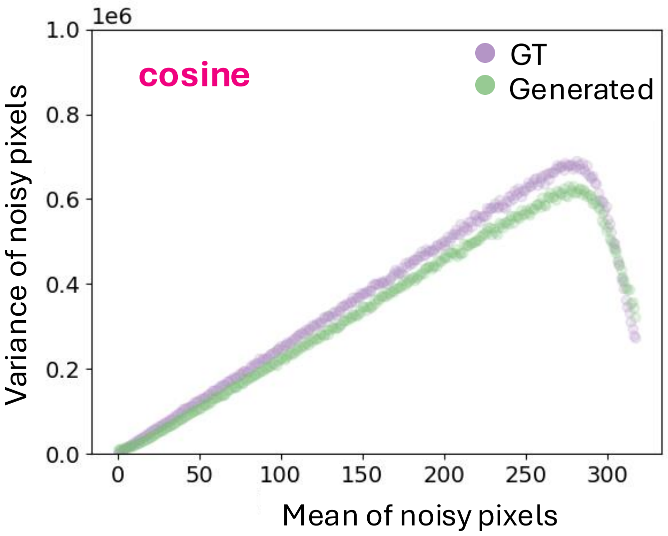

Different Diffusion Noise Schedules on 1D Poisson Distribution

A toy example

$$
\begin{aligned}
\mathbf{x}_t &= \sqrt{\alpha_t}\mathbf{x}_{t-1} + \sqrt{1-\alpha_t}\boldsymbol{\epsilon}_{t-1} \\
&= \sqrt{\alpha_t \alpha_{t-1}}\mathbf{x}_{t-2} + \sqrt{1-\alpha_t \alpha_{t-1}}\bar{\boldsymbol{\epsilon}}_{t-2} \\
&= \dots \\
&= \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \boxed{\sqrt{1-\bar{\alpha}_t}}\boldsymbol{\epsilon}
\end{aligned}
$$

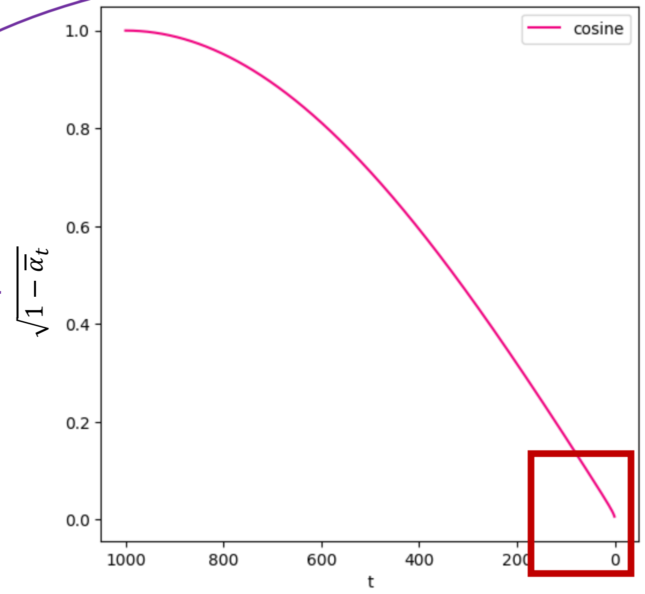

Reverse Process

x_{999} x_t x_0

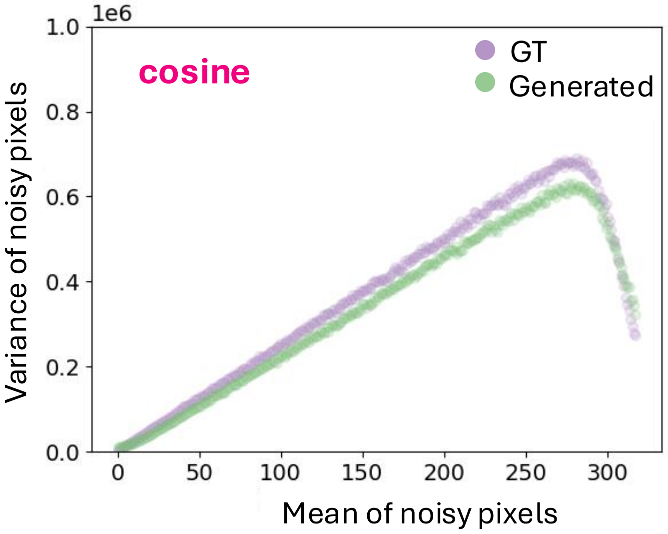

$$f_{MMSE} = \mathbb{E}[x_0 | x_t]$$

Different Diffusion Noise Schedules on 1D Poisson Distribution

A toy example

$$
\begin{aligned}
\mathbf{x}_t &= \sqrt{\alpha_t}\mathbf{x}_{t-1} + \sqrt{1-\alpha_t}\boldsymbol{\epsilon}_{t-1} \\
&= \sqrt{\alpha_t\alpha_{t-1}}\mathbf{x}_{t-2} + \sqrt{1-\alpha_t\alpha_{t-1}}\bar{\boldsymbol{\epsilon}}_{t-2} \\
&= \dots \\
&= \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \boxed{\sqrt{1-\bar{\alpha}_t}}\boldsymbol{\epsilon}
\end{aligned}
$$

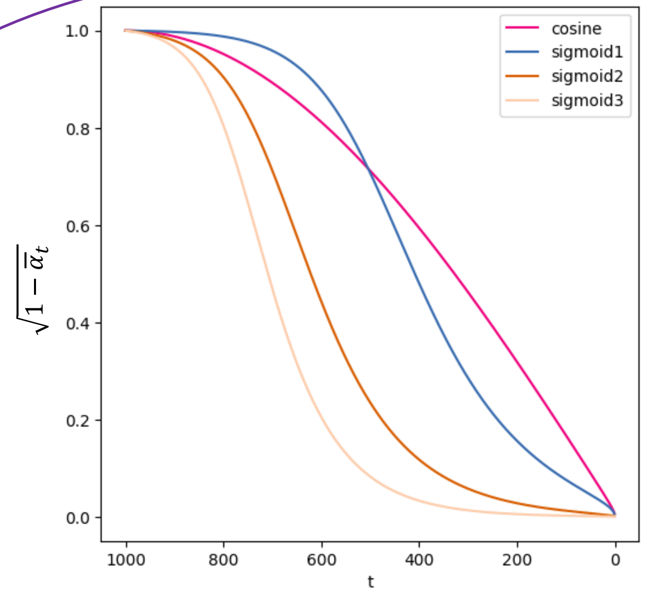

Reverse Process

x_{999} x_t x_0

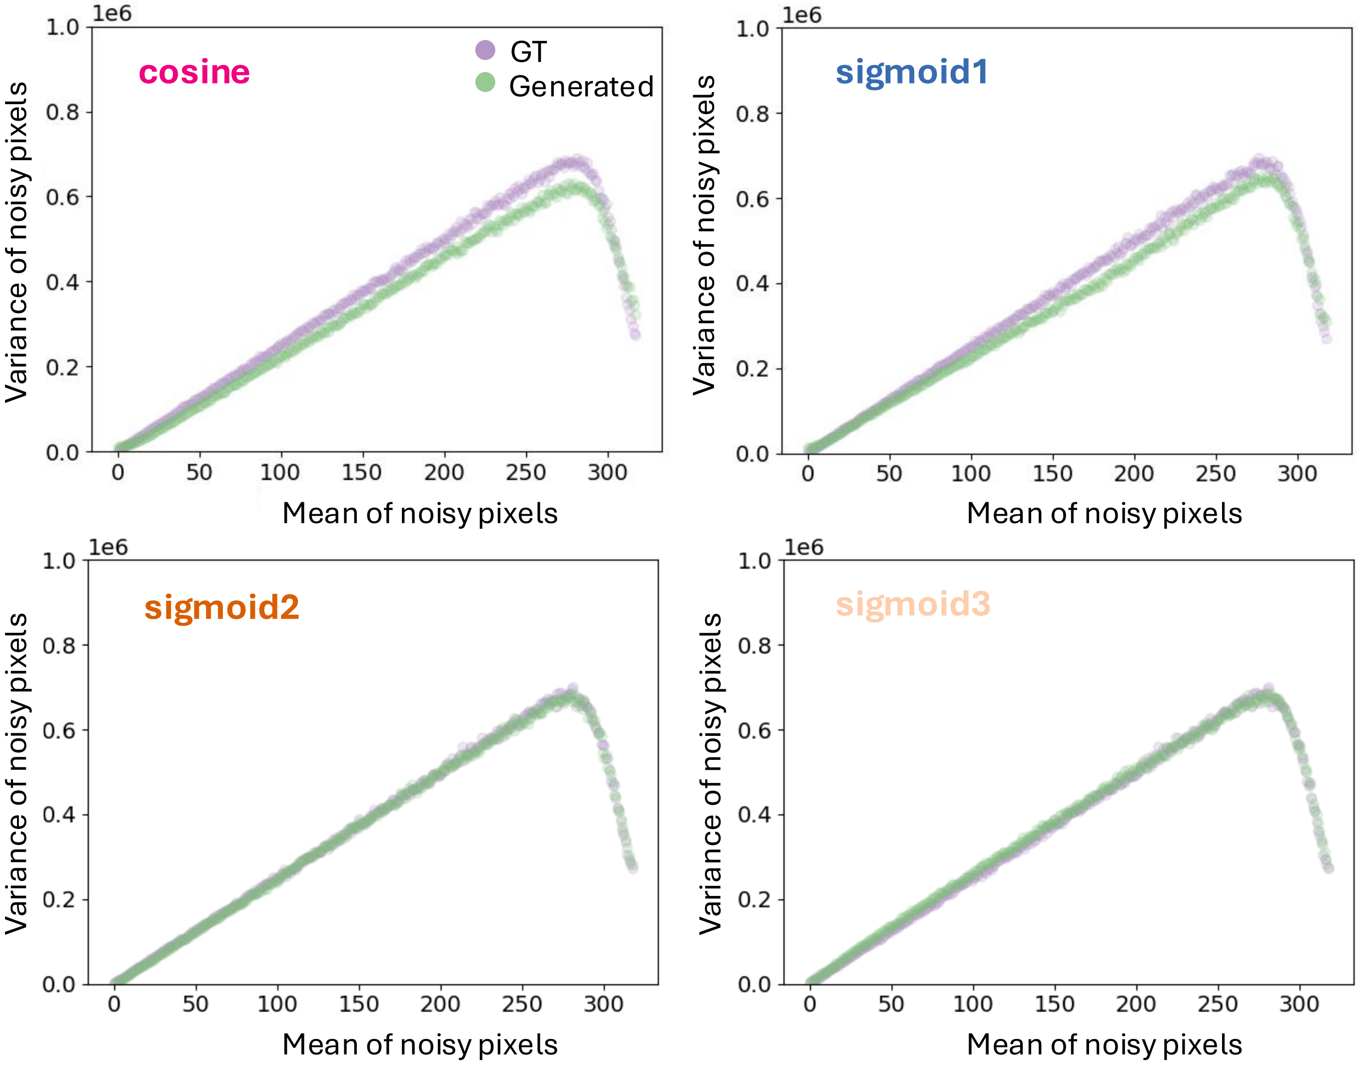

Different Diffusion Noise Schedules on 1D Poisson Distribution

A toy example

$$
\begin{aligned}
\mathbf{x}_t &= \sqrt{\alpha_t}\mathbf{x}_{t-1} + \sqrt{1-\alpha_t}\boldsymbol{\epsilon}_{t-1} \\
&= \sqrt{\alpha_t\alpha_{t-1}}\mathbf{x}_{t-2} + \sqrt{1-\alpha_t\alpha_{t-1}}\bar{\boldsymbol{\epsilon}}_{t-2} \\
&= \dots \\
&= \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \boxed{\sqrt{1-\bar{\alpha}_t}}\boldsymbol{\epsilon}
\end{aligned}
$$

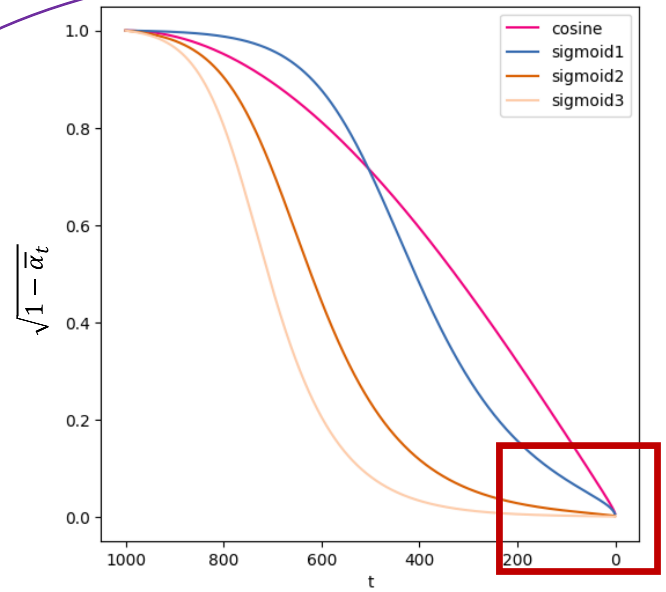

Reverse Process

x_{999} x_t x_0

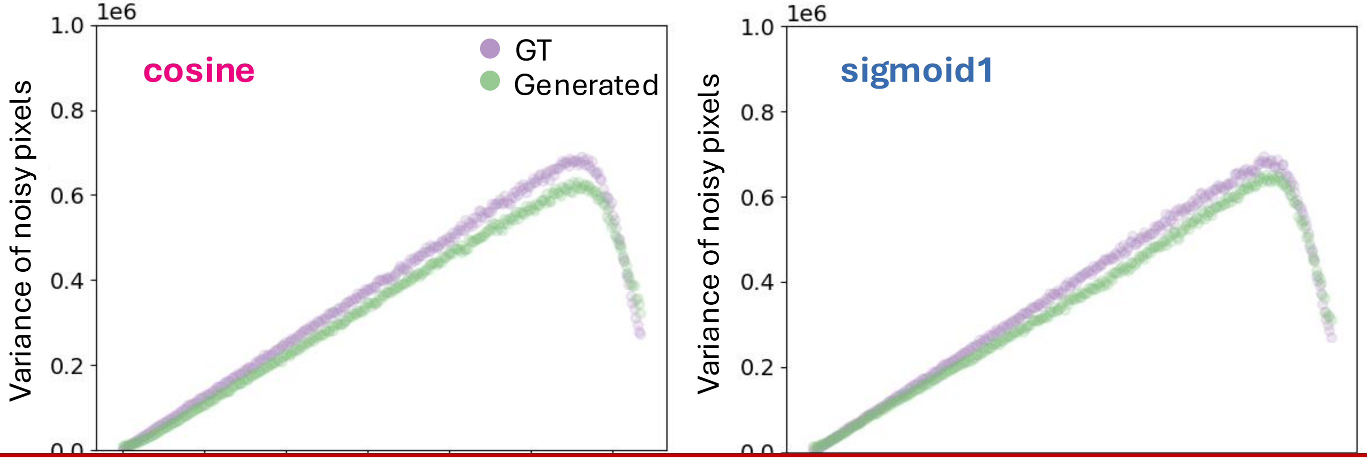

Noise schedules with lower slopes yield better results

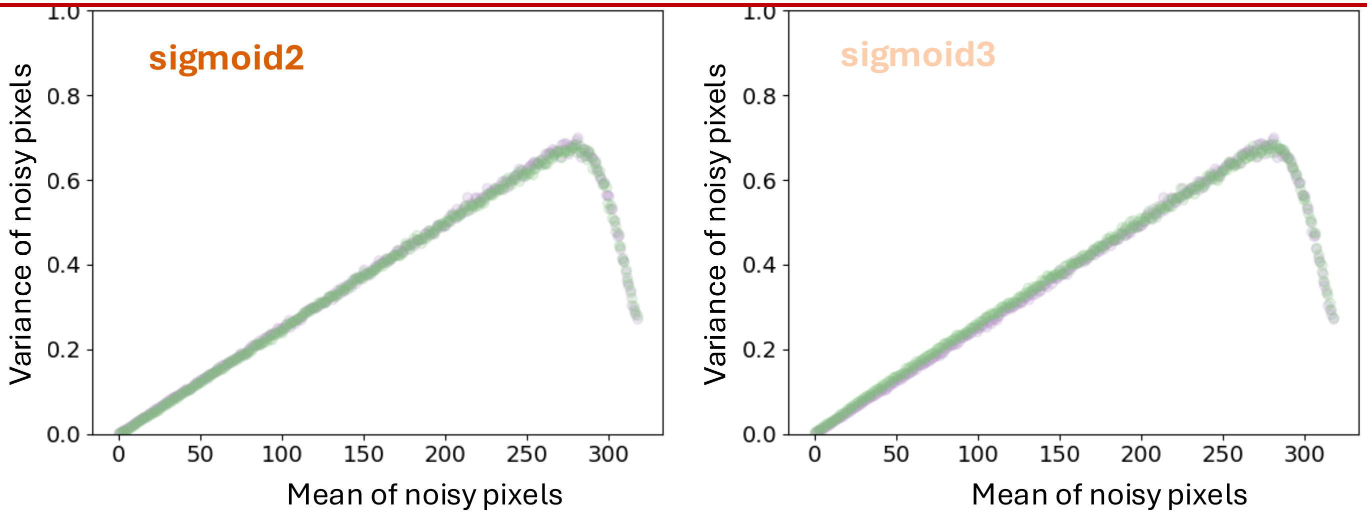

Generated Results

Generated-Noise Visualization ISO 6400

Clean image

Real noisy image

Generated noisy image

Real noise image

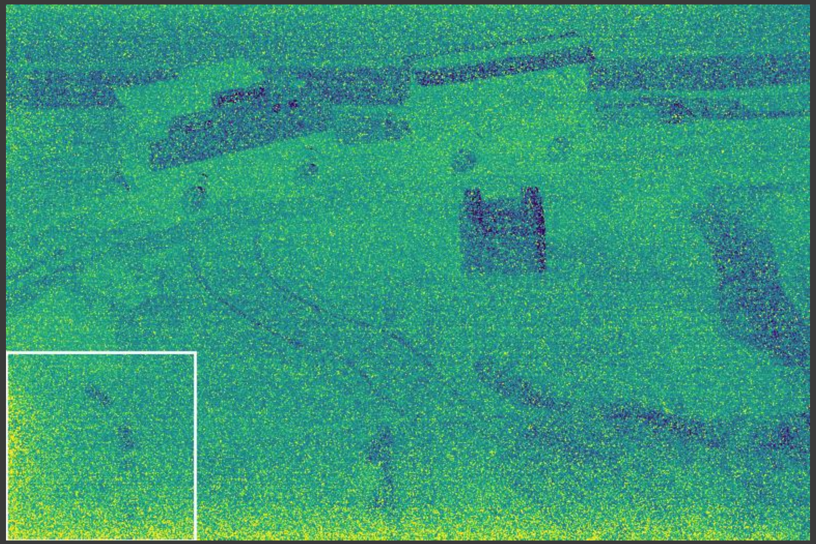

Generated noise image

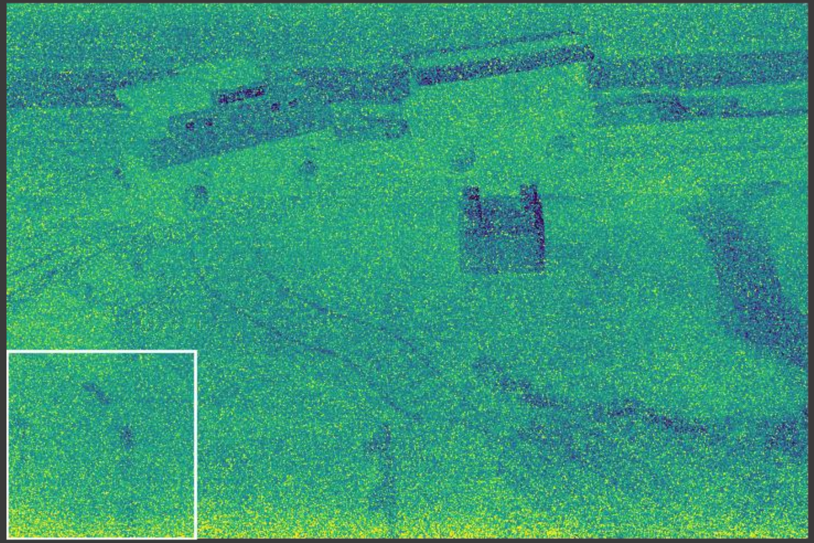

Generated-Noise Visualization ISO 6400 (Zoomed In)

Clean image

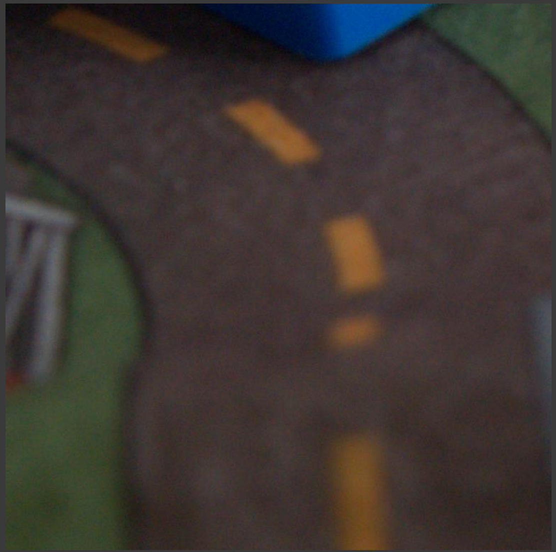

Real noisy image

Generated noisy image

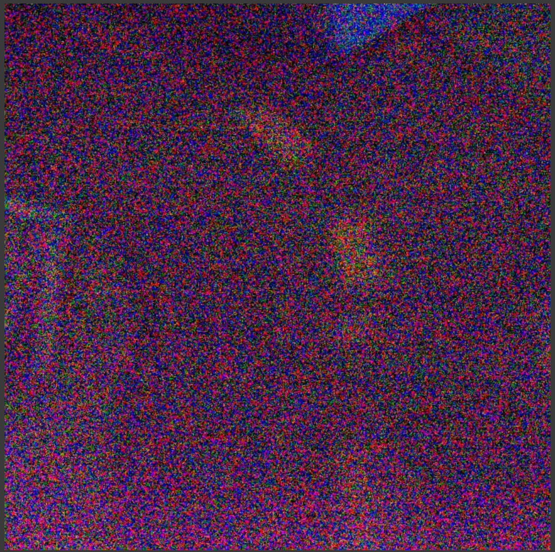

Real noise image

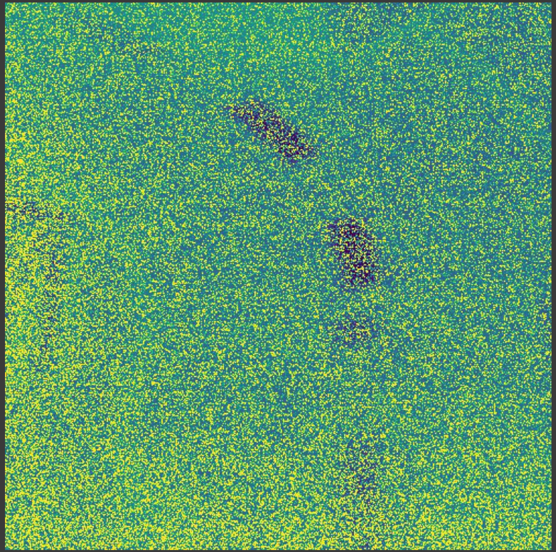

Generated noise image

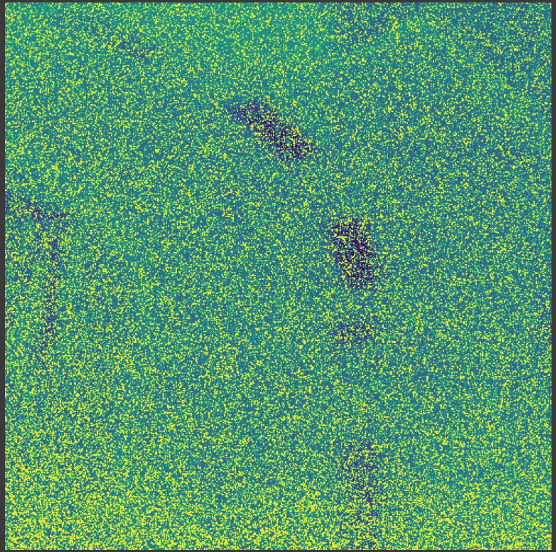

Generated-Noise Visualization ISO 200

Clean image

Real noisy image

Generated noisy image

Real noise image

Generated noise image

Generated-Noise Visualization ISO 200 (Zoomed In)

Clean image

Real noisy image

Generated noisy image

Real noise image

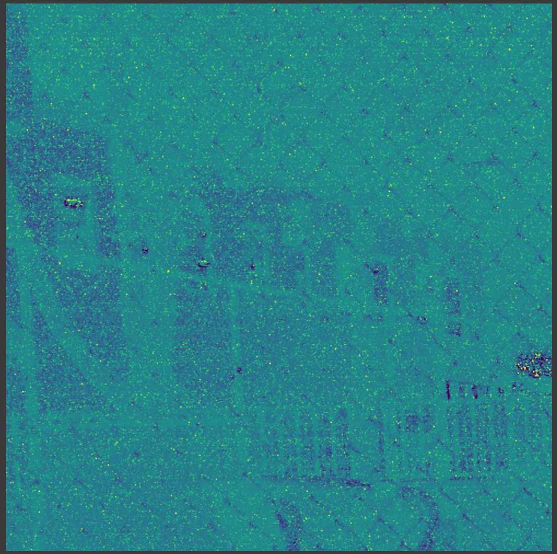

Generated noise image

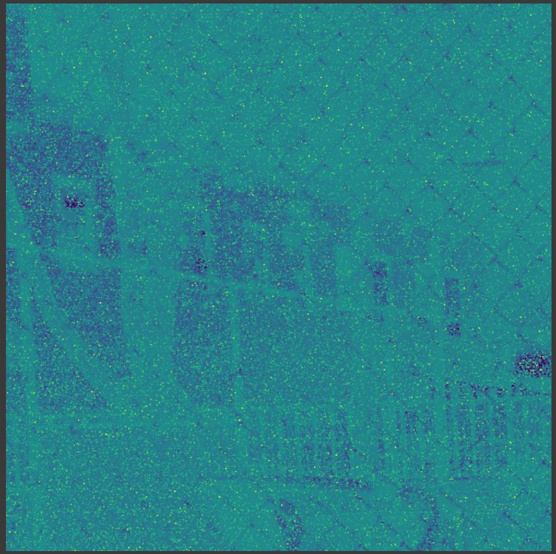

Denoising Training with Synthetic data

Denoising Results

Denoising network trained with synthetic data

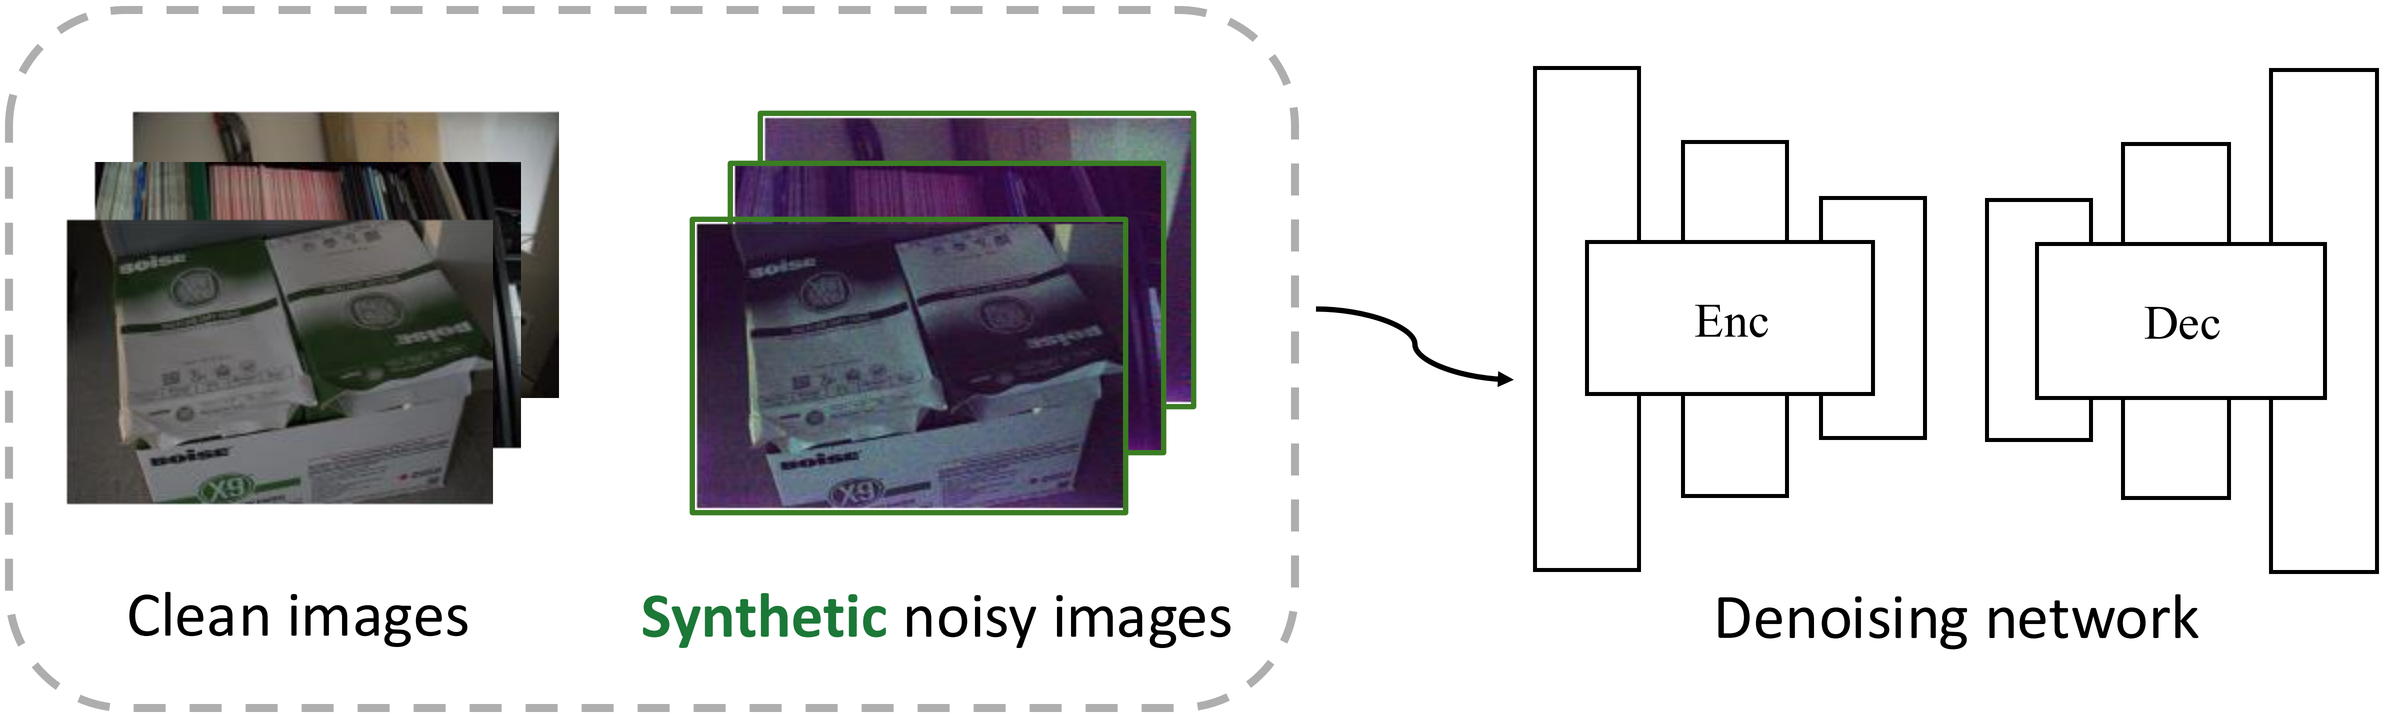

Denoising Results   Denoising network trained with synthetic data

PSNR / SSIM:

Dataset	Ratio	*Real Paired*	Poisson-Gaussian	ELD	Starlight	NoiseFlow	Ours
SID	× 100	*42.95 / 0.958*	41.05 / 0.936	<u>41.95</u> / <u>0.953</u>	40.47 / 0.926	40.20 / 0.925	**43.30** / **0.958**
	× 250	*40.27 / 0.943*	36.63 / 0.885	<u>39.44</u> / <u>0.931</u>	36.25 / 0.858	36.15 / 0.870	**40.53** / **0.944**
	× 300	*37.32 / 0.928*	33.34 / 0.811	<u>36.36</u> / <u>0.911</u>	32.99 / 0.780	33.27 / 0.803	**37.68** / **0.928**
ELD	× 100	*45.52 / 0.977*	44.28 / 0.936	<u>45.45</u> / **0.975**	43.80 / 0.936	43.31 / 0.941	**45.79** / <u>0.972</u>
	× 200	*41.70 / 0.912*	41.16 / 0.885	**43.43** / **0.954**	40.86 / 0.884	40.26 / 0.885	<u>42.25</u> / <u>0.924</u>

ELD: A Physics-based Noise Formation Model for Extreme Low-light Raw Denoising, Wei *et al*, CVPR2020
Starlight: Dancing under the stars: video denoising in starlight, Kristina, et al, CVPR2022.
Noise Flow: Noise Modeling with Conditional Normalizing Flows, Abdelhamed *et al*, ICCV2019
LRD: Towards General Low-Light Raw Noise Synthesis and Modeling, Zhang *et al*, ICCV2023
PMN: Learnability Enhancement for Low-light Raw Denoising: A Data Perspective, Feng *et al*, TPAMI 2024

Denoising Results

Denoising network trained with synthetic data

PSNR / SSIM (with calibration data):

Dataset	Ratio	PMN	LRD	Ours*
SID	× 100	43.47 / **0.961**	43.16 / 0.958	**43.92** / **0.961**
	× 250	41.04 / **0.947**	40.69 / 0.941	**41.28** / 0.946
	× 300	37.87 / **0.934**	37.48 / 0.919	**37.94** / 0.930
ELD	× 100	**46.99** / **0.984**	46.16 / 0.983	46.95 / 0.978
	× 200	44.85 / 0.969	43.91 / 0.968	**45.11** / **0.971**

ELD: A Physics-based Noise Formation Model for Extreme Low-light Raw Denoising, Wei *et al*, CVPR2020
Starlight: Dancing under the stars: video denoising in starlight, Kristina, et al, CVPR2022.
Noise Flow: Noise Modeling with Conditional Normalizing Flows, Abdelhamed *et al*, ICCV2019
LRD: Towards General Low-Light Raw Noise Synthesis and Modeling, Zhang *et al*, ICCV2023
PMN: Learnability Enhancement for Low-light Raw Denoising: A Data Perspective, Feng *et al*, TPAMI 2024

Denoising Results

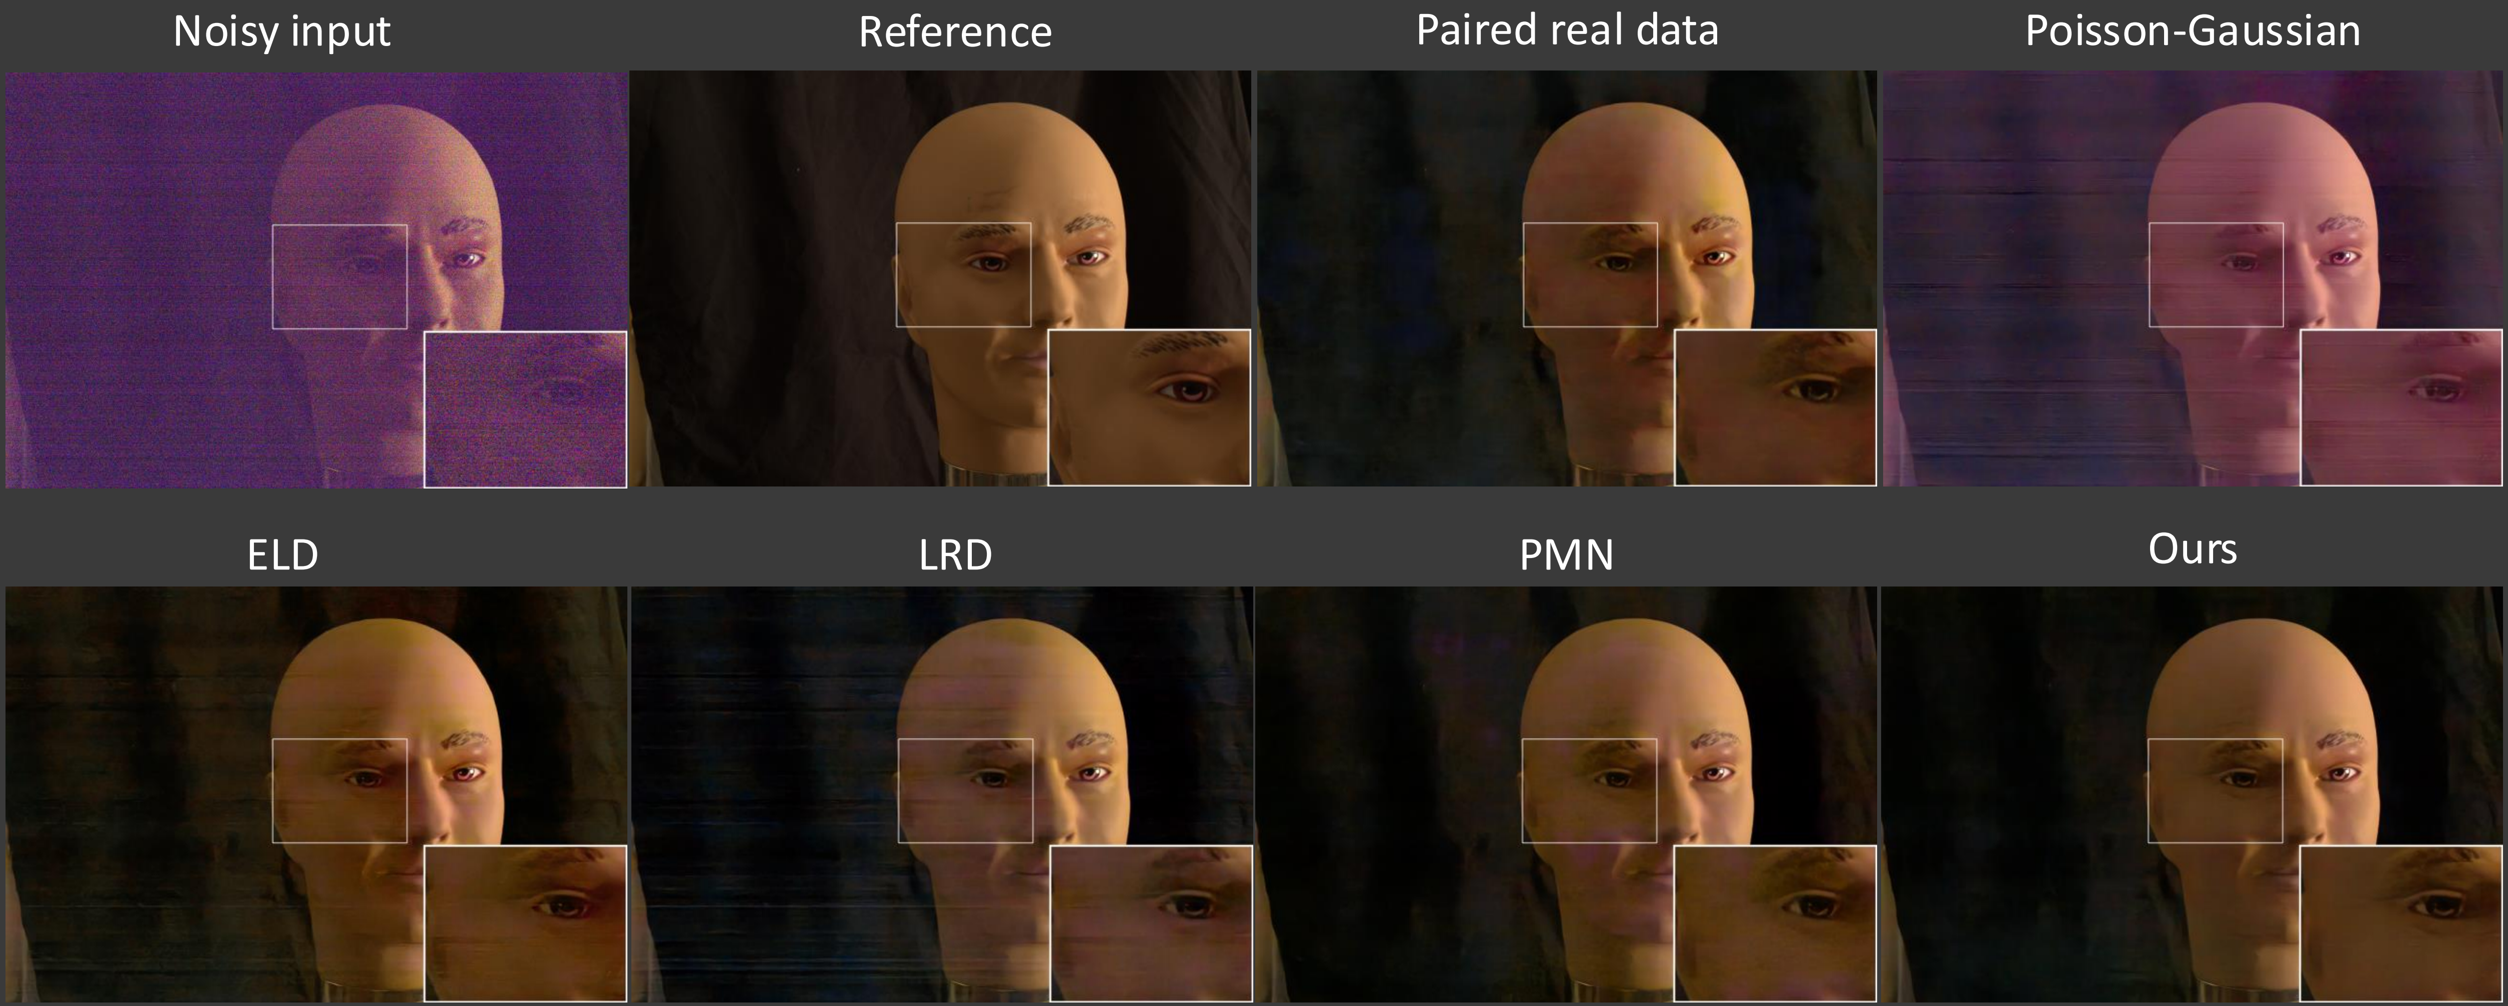

ELD: A Physics-based Noise Formation Model for Extreme Low-light Raw Denoising, Wei *et al*, CVPR2020
LRD: Towards General Low-Light Raw Noise Synthesis and Modeling, Zhang *et al*, ICCV2023
PMN: Learnability Enhancement for Low-light Raw Denoising: A Data Perspective, Feng *et al*, TPAMI 2024

Denoising Results (Zoomed In)

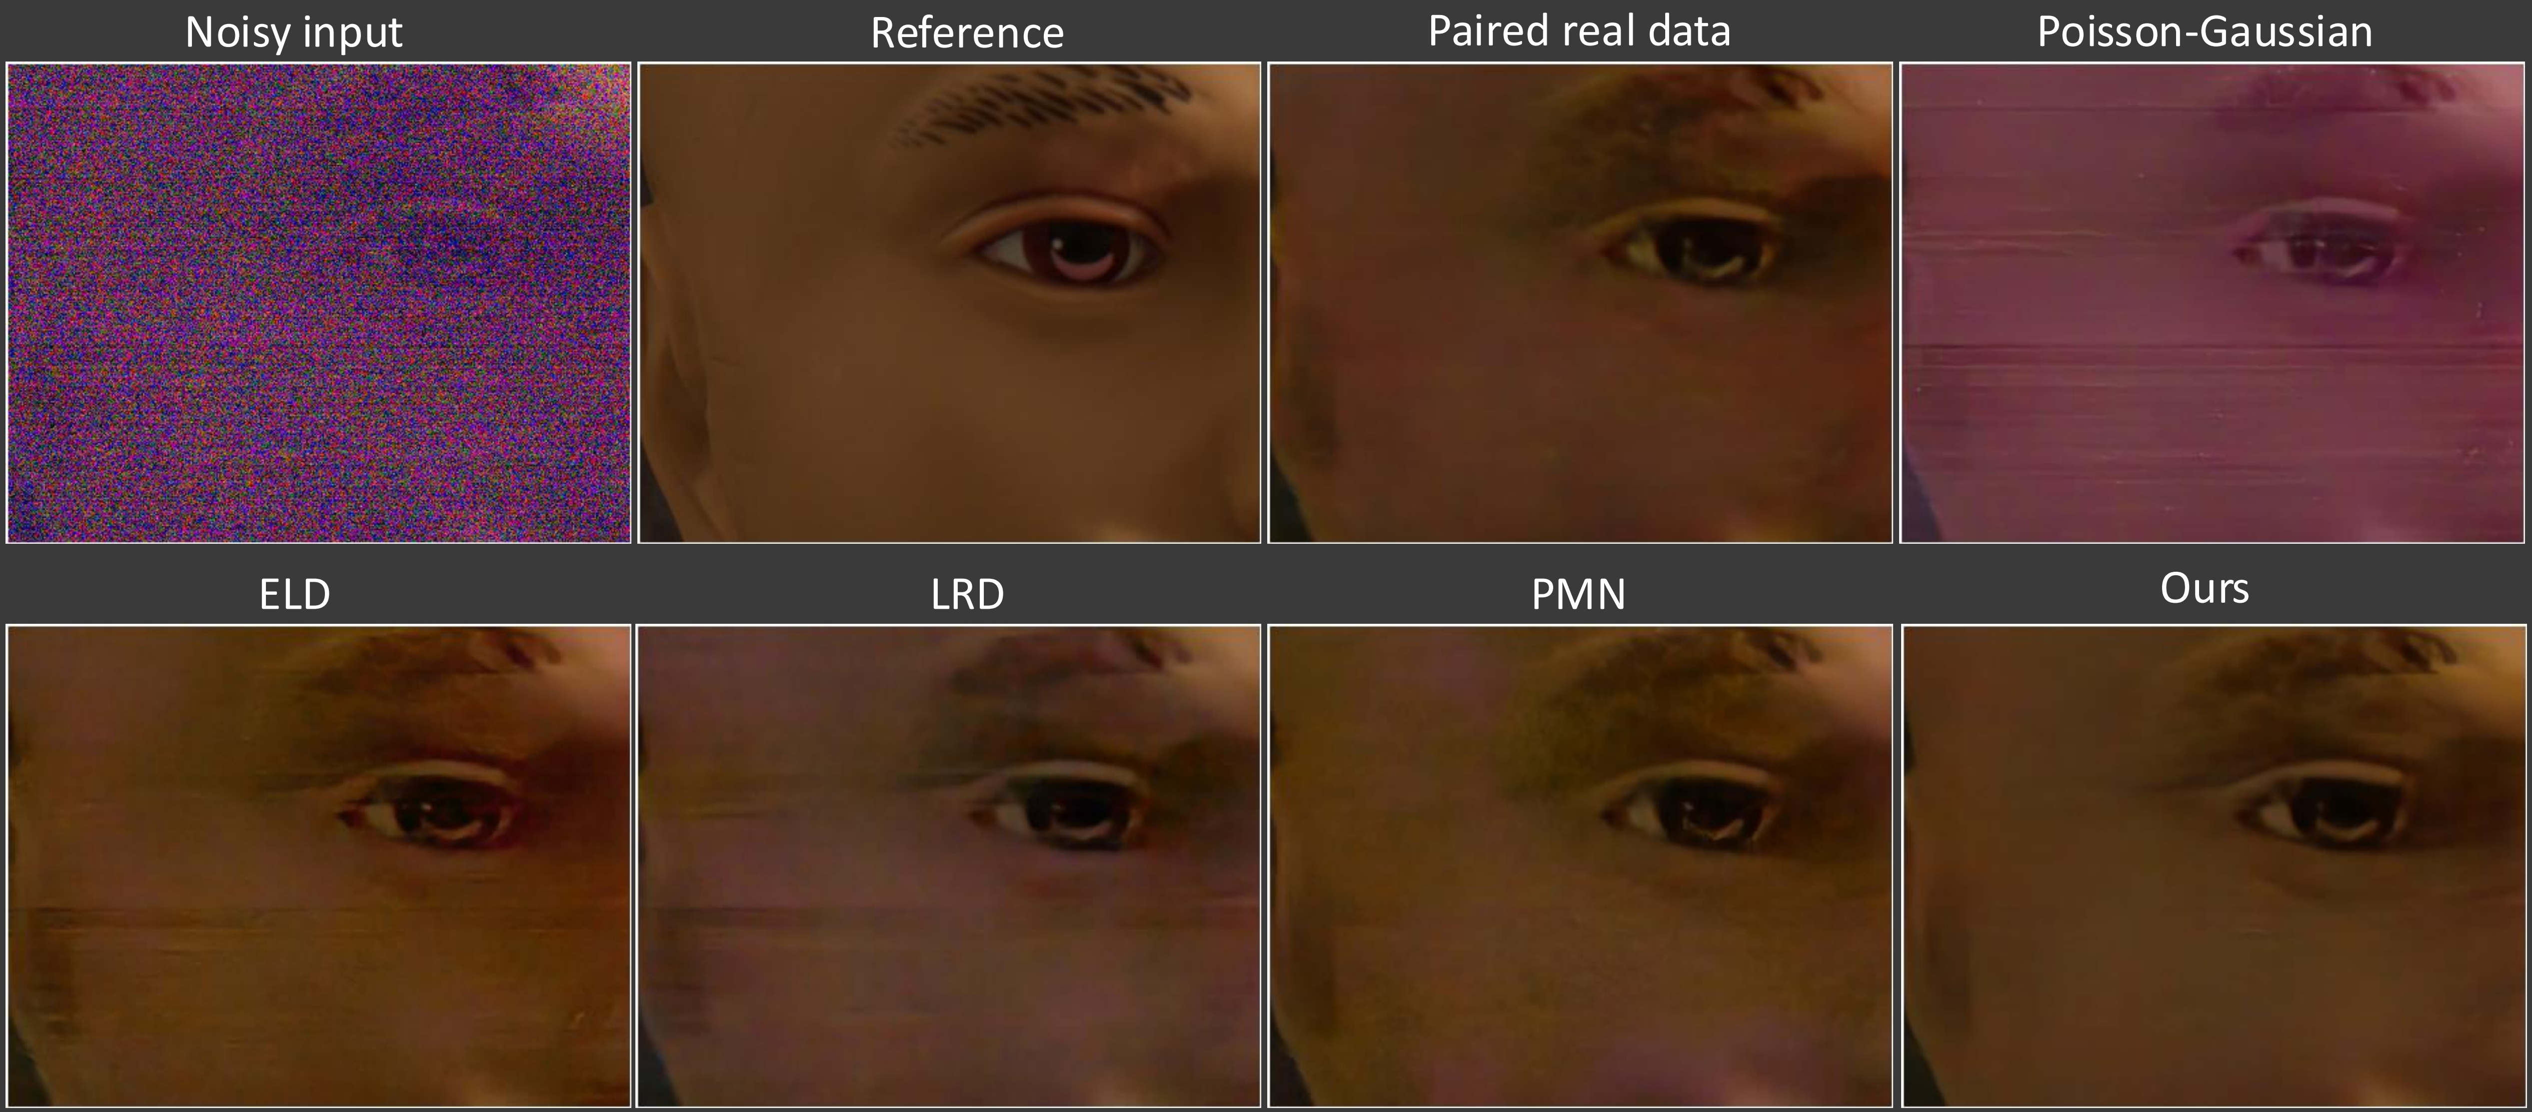

ELD: A Physics-based Noise Formation Model for Extreme Low-light Raw Denoising, Wei *et al*, CVPR2020
LRD: Towards General Low-Light Raw Noise Synthesis and Modeling, Zhang *et al*, ICCV2023
PMN: Learnability Enhancement for Low-light Raw Denoising: A Data Perspective, Feng *et al*, TPAMI 2024

Denoising Results

Noisy input

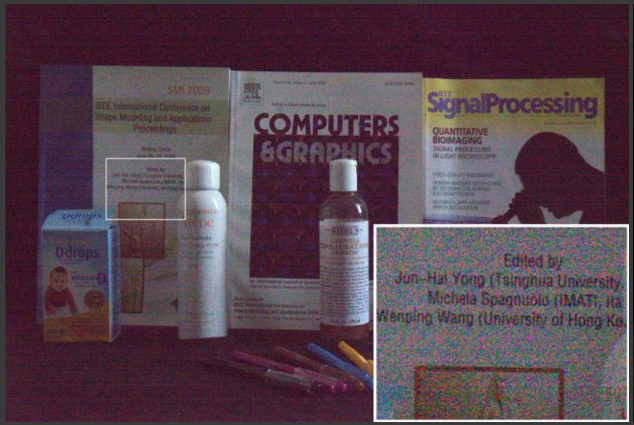

Reference

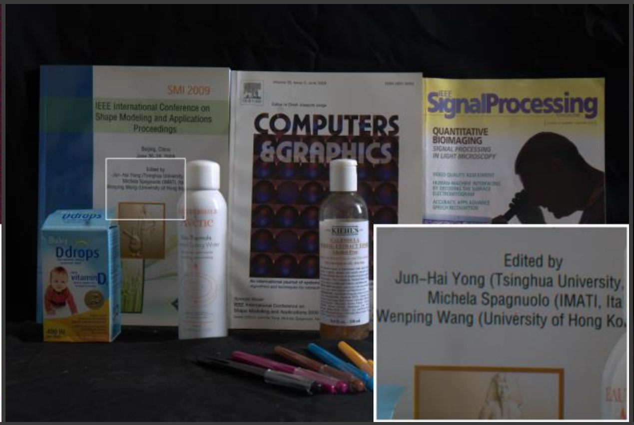

Paired real data

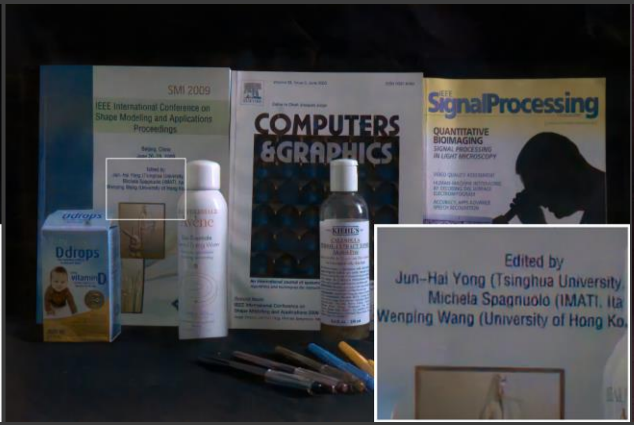

Poisson-Gaussian

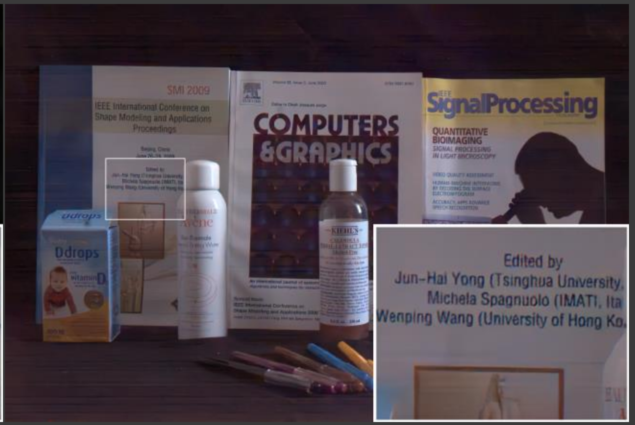

ELD

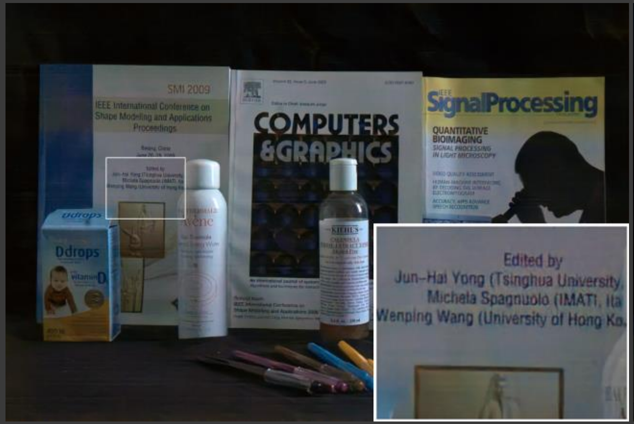

LRD

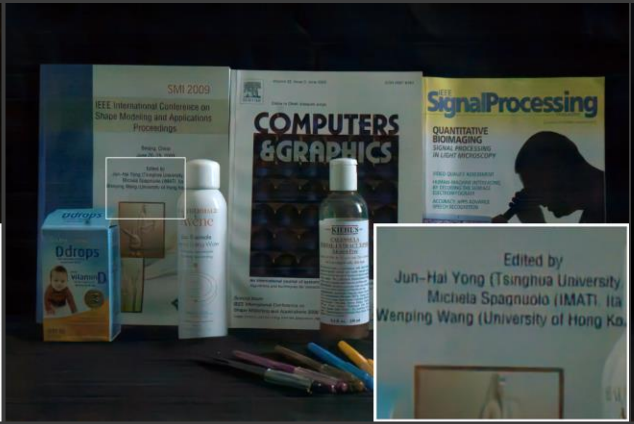

PMN

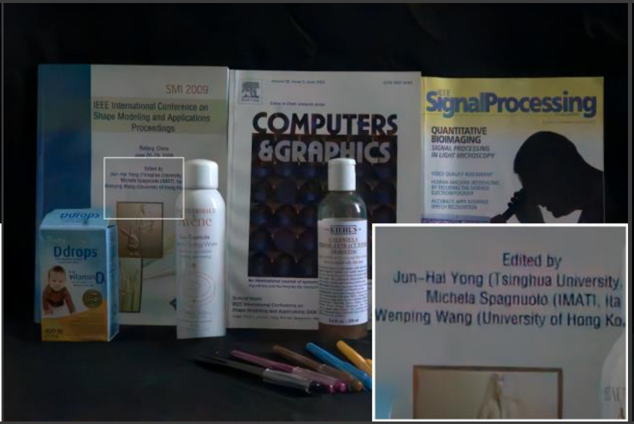

Ours

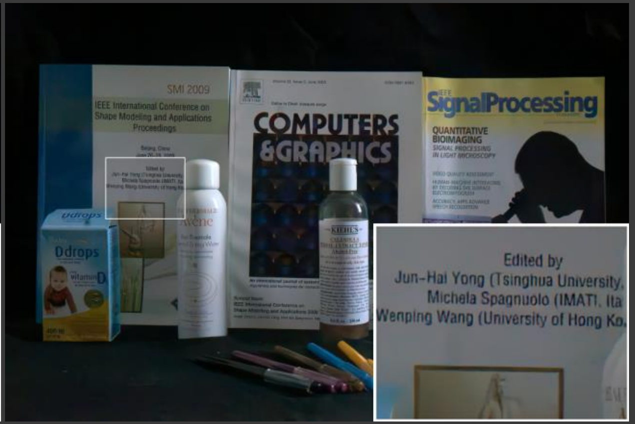

ELD: A Physics-based Noise Formation Model for Extreme Low-light Raw Denoising, Wei *et al*, CVPR2020
LRD: Towards General Low-Light Raw Noise Synthesis and Modeling, Zhang *et al*, ICCV2023
PMN: Learnability Enhancement for Low-light Raw Denoising: A Data Perspective, Feng *et al*, TPAMI 2024

Denoising Results (Zoomed In)

Noisy input

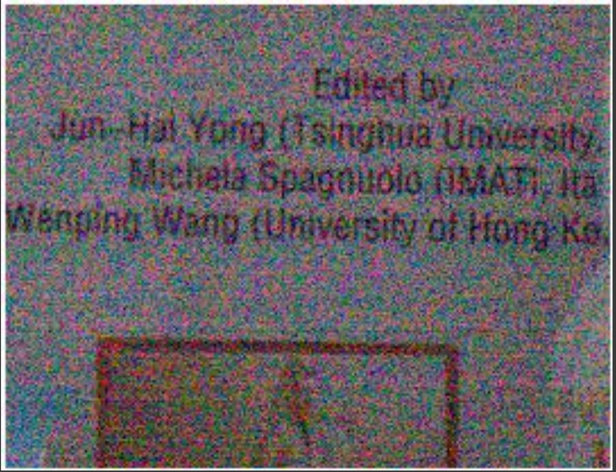

Reference

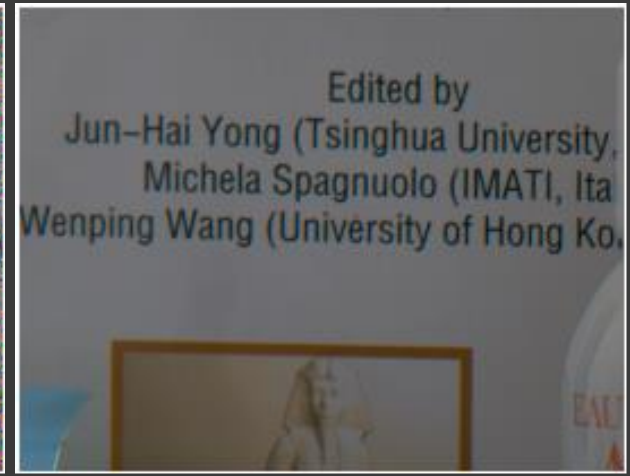

Paired real data

Poisson-Gaussian

ELD

LRD

PMN

Ours

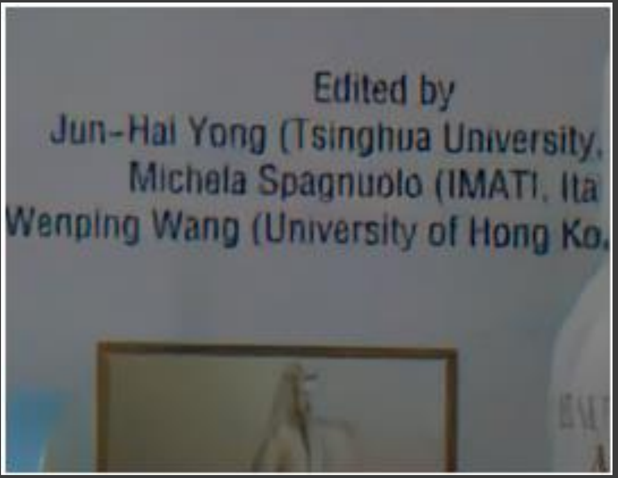

ELD: A Physics-based Noise Formation Model for Extreme Low-light Raw Denoising, Wei *et al*, CVPR2020
LRD: Towards General Low-Light Raw Noise Synthesis and Modeling, Zhang *et al*, ICCV2023
PMN: Learnability Enhancement for Low-light Raw Denoising: A Data Perspective, Feng *et al*, TPAMI 2024

Denoising Results

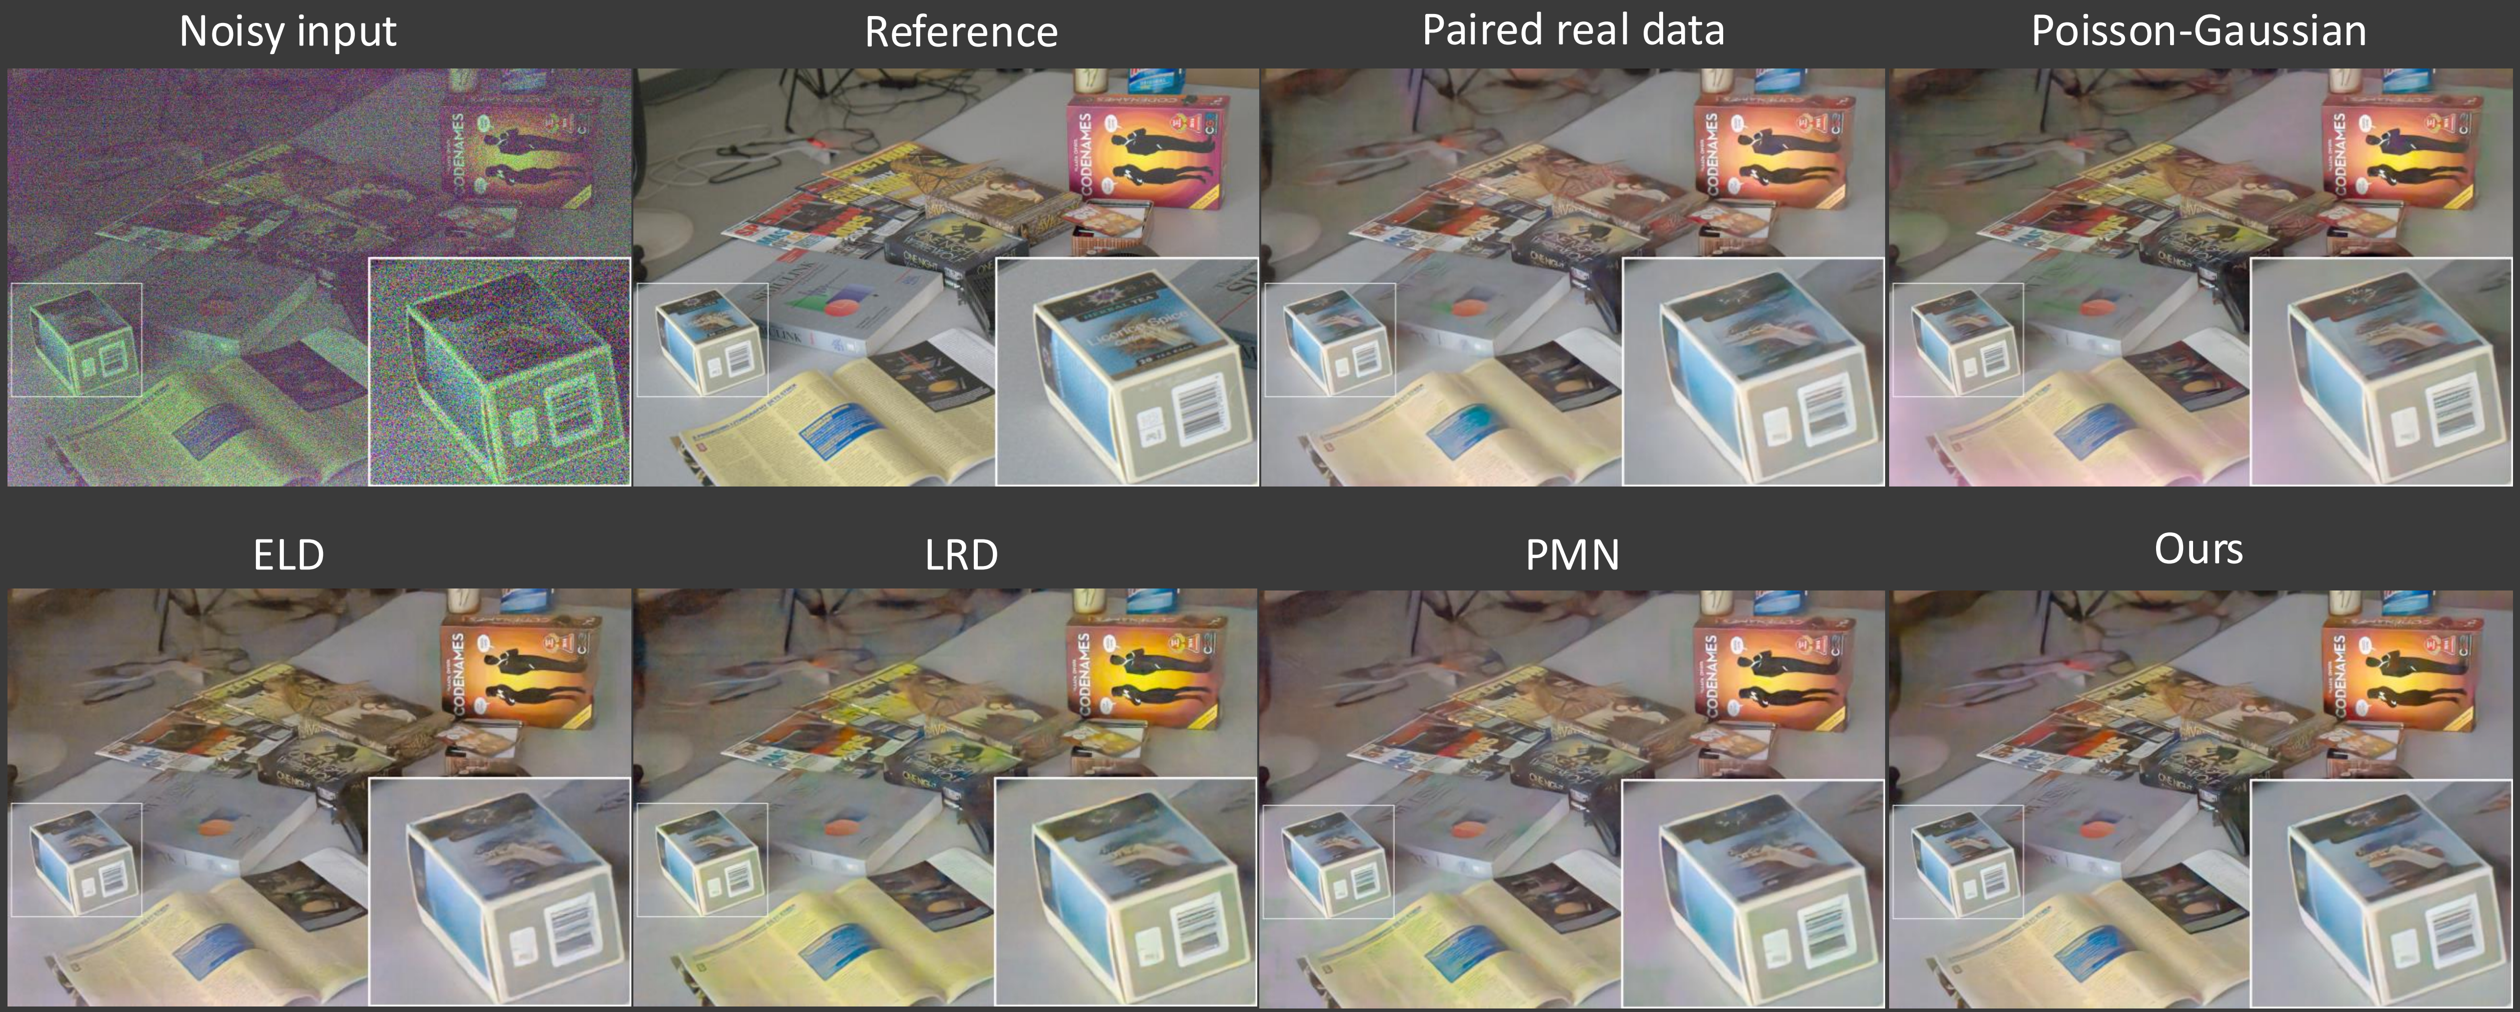

ELD: A Physics-based Noise Formation Model for Extreme Low-light Raw Denoising, Wei *et al*, CVPR2020
LRD: Towards General Low-Light Raw Noise Synthesis and Modeling, Zhang *et al*, ICCV2023
PMN: Learnability Enhancement for Low-light Raw Denoising: A Data Perspective, Feng *et al*, TPAMI 2024

Denoising Results (Zoomed In)

Noisy input

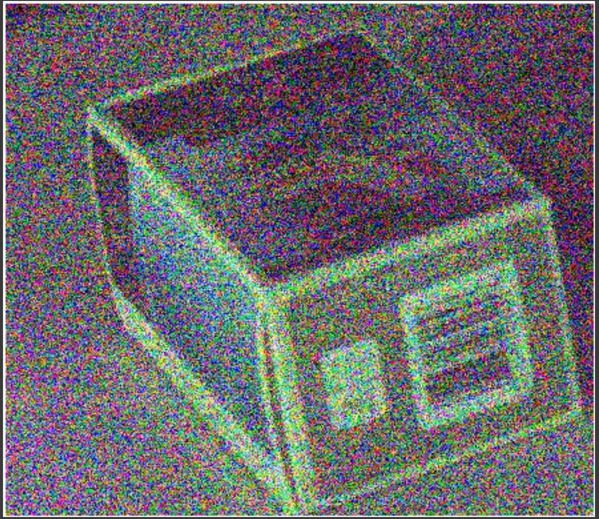

Reference

Paired real data

Poisson-Gaussian

ELD

LRD

PMN

Ours

ELD: A Physics-based Noise Formation Model for Extreme Low-light Raw Denoising, Wei *et al*, CVPR2020
LRD: Towards General Low-Light Raw Noise Synthesis and Modeling, Zhang *et al*, ICCV2023
PMN: Learnability Enhancement for Low-light Raw Denoising: A Data Perspective, Feng *et al*, TPAMI 2024

Denoising Results

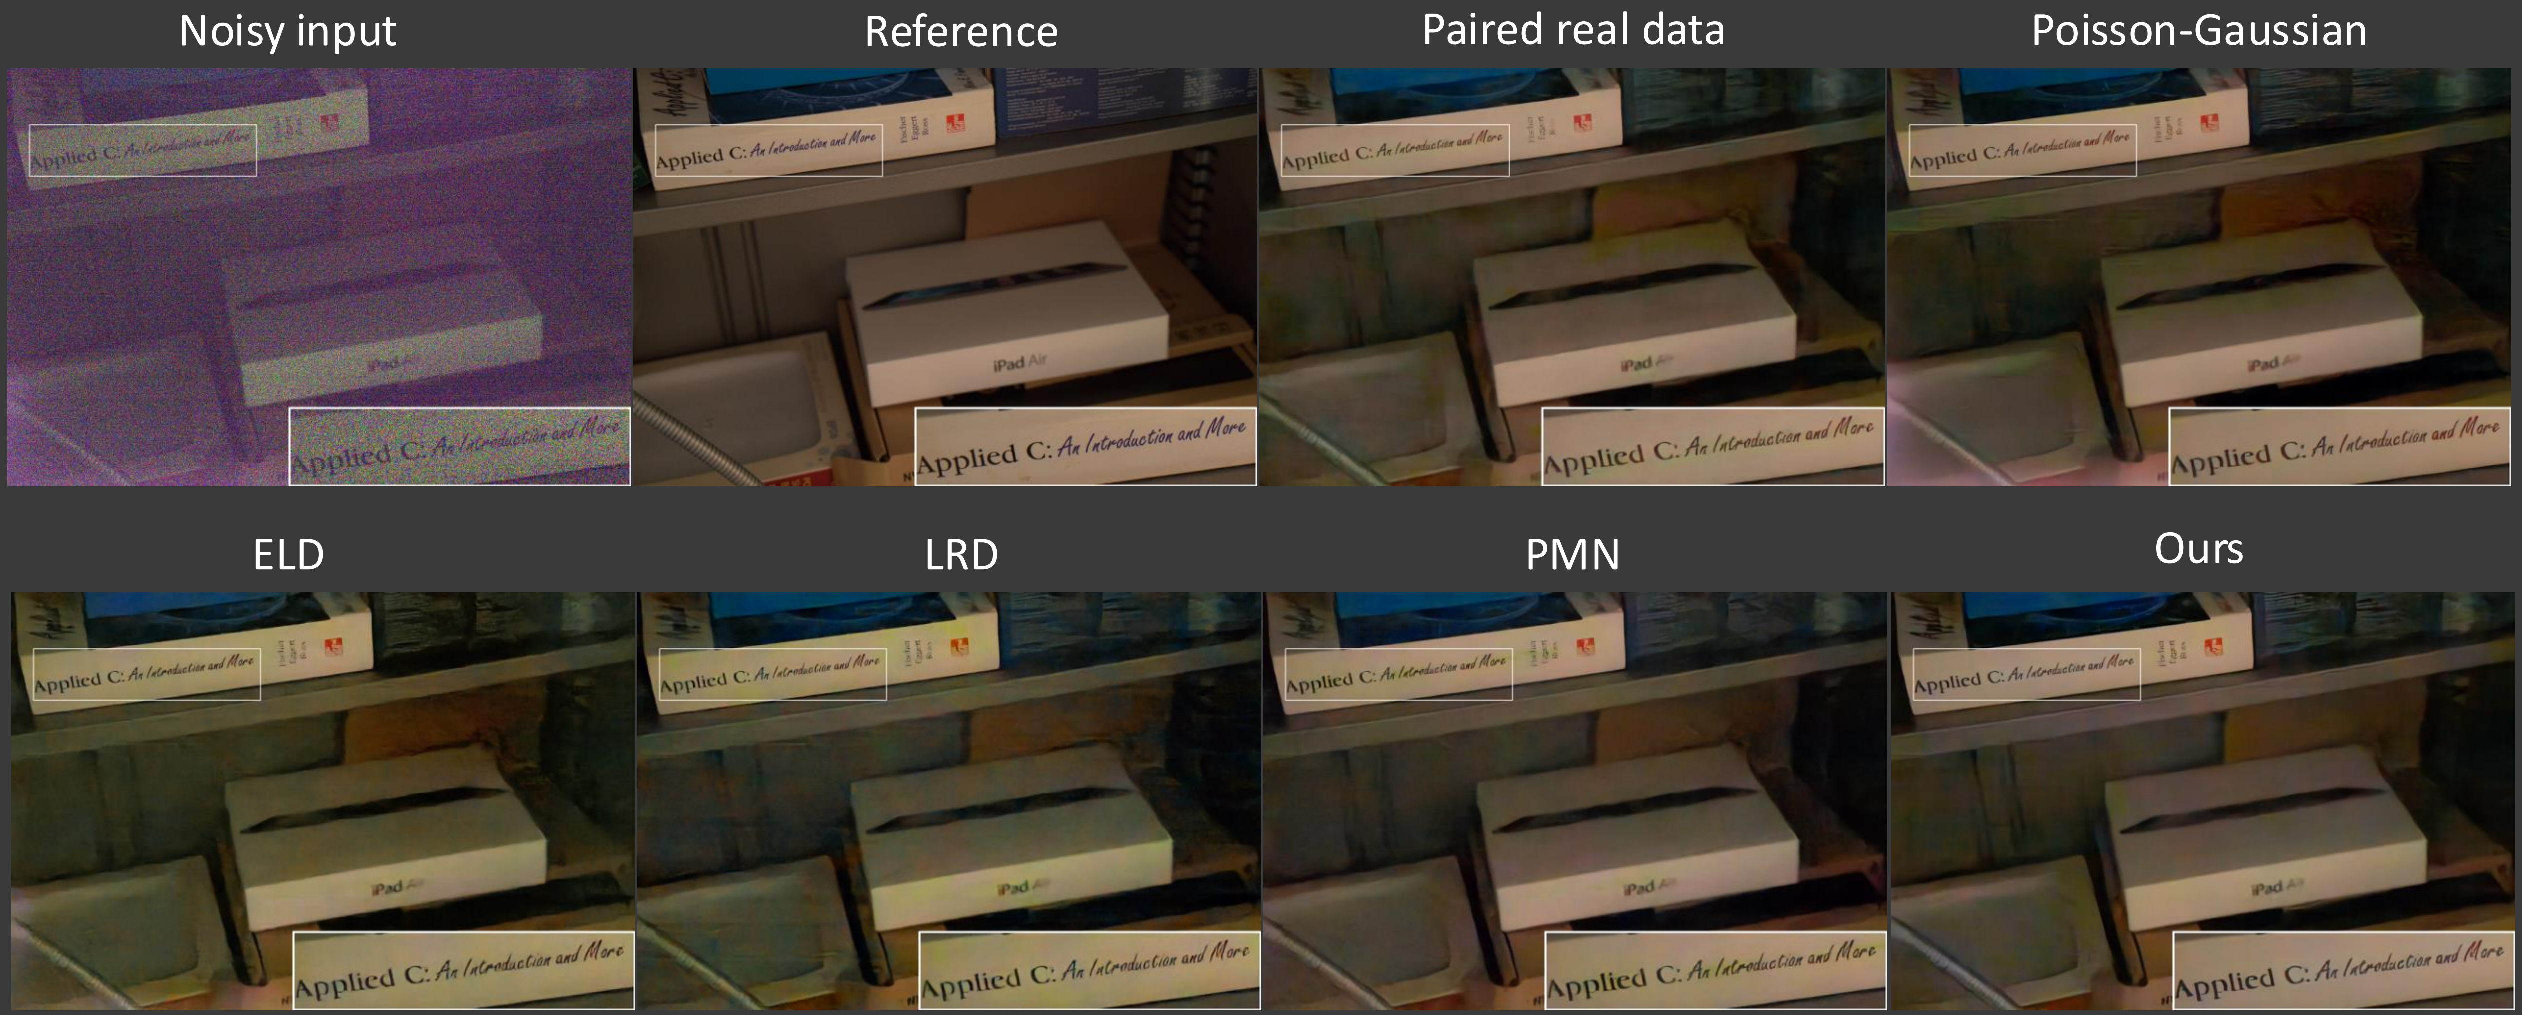

ELD: A Physics-based Noise Formation Model for Extreme Low-light Raw Denoising, Wei *et al*, CVPR2020
LRD: Towards General Low-Light Raw Noise Synthesis and Modeling, Zhang *et al*, ICCV2023
PMN: Learnability Enhancement for Low-light Raw Denoising: A Data Perspective, Feng *et al*, TPAMI 2024

Denoising Results (Zoomed In)

Noisy input

ELD

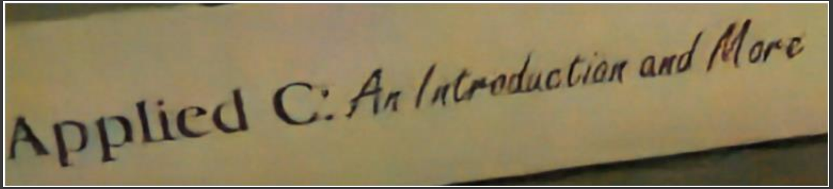

Reference

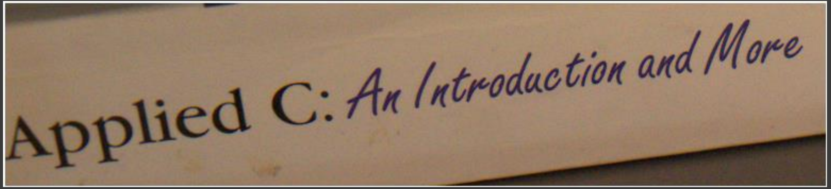

LRD

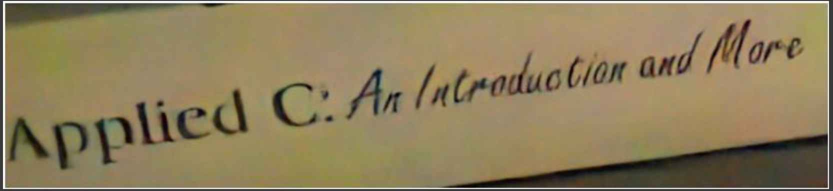

Paired real data

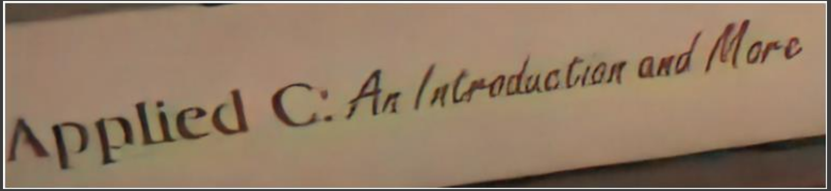

PMN

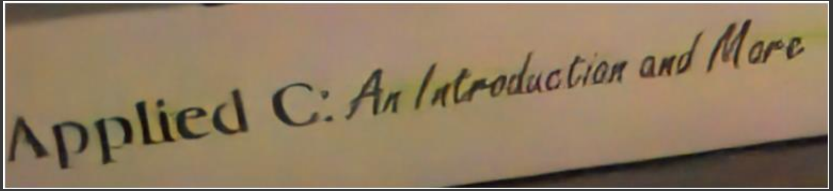

Poisson-Gaussian

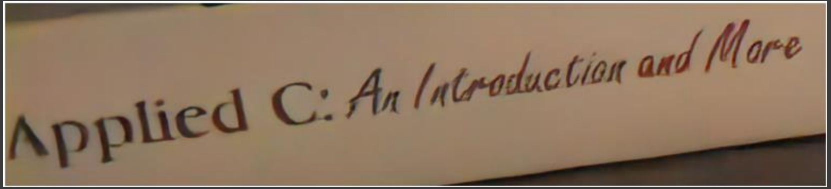

Ours

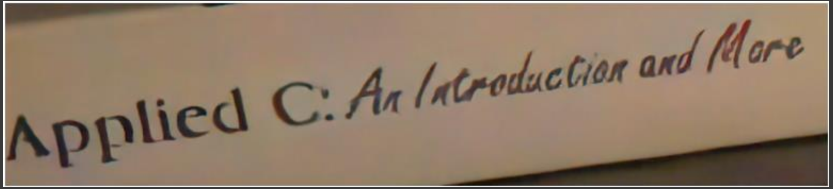

ELD: A Physics-based Noise Formation Model for Extreme Low-light Raw Denoising, Wei *et al*, CVPR2020
LRD: Towards General Low-Light Raw Noise Synthesis and Modeling, Zhang *et al*, ICCV2023
PMN: Learnability Enhancement for Low-light Raw Denoising: A Data Perspective, Feng *et al*, TPAMI 2024

Conclusions

- *A novel diffusion framework for low-light noise generation*
  the denoising network trained with our synthetic data achieves the **best results**

- *Two-branch network*
  effectively model **different noise components**

- *Positional encoding*
  helpful for modeling **spatially-correlated noise**

- *Proper diffusion noise schedule*
  essential for preserving **noise variance**

- *Limitation*
  still require **real data pairs** for training the diffusion network

GRETSI Paper Link

Extended Paper Link

Code Link

Thank you!

Liying Lu*

Raphaël Achddou

Sabine Süsstrunk

Ablations Diffusion network: with positional encoding

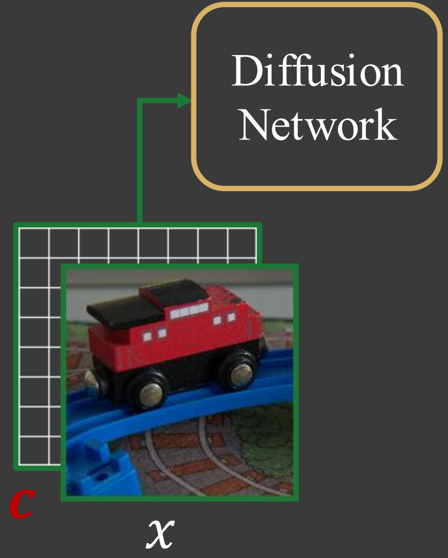

c: positional encoding
x: clean image

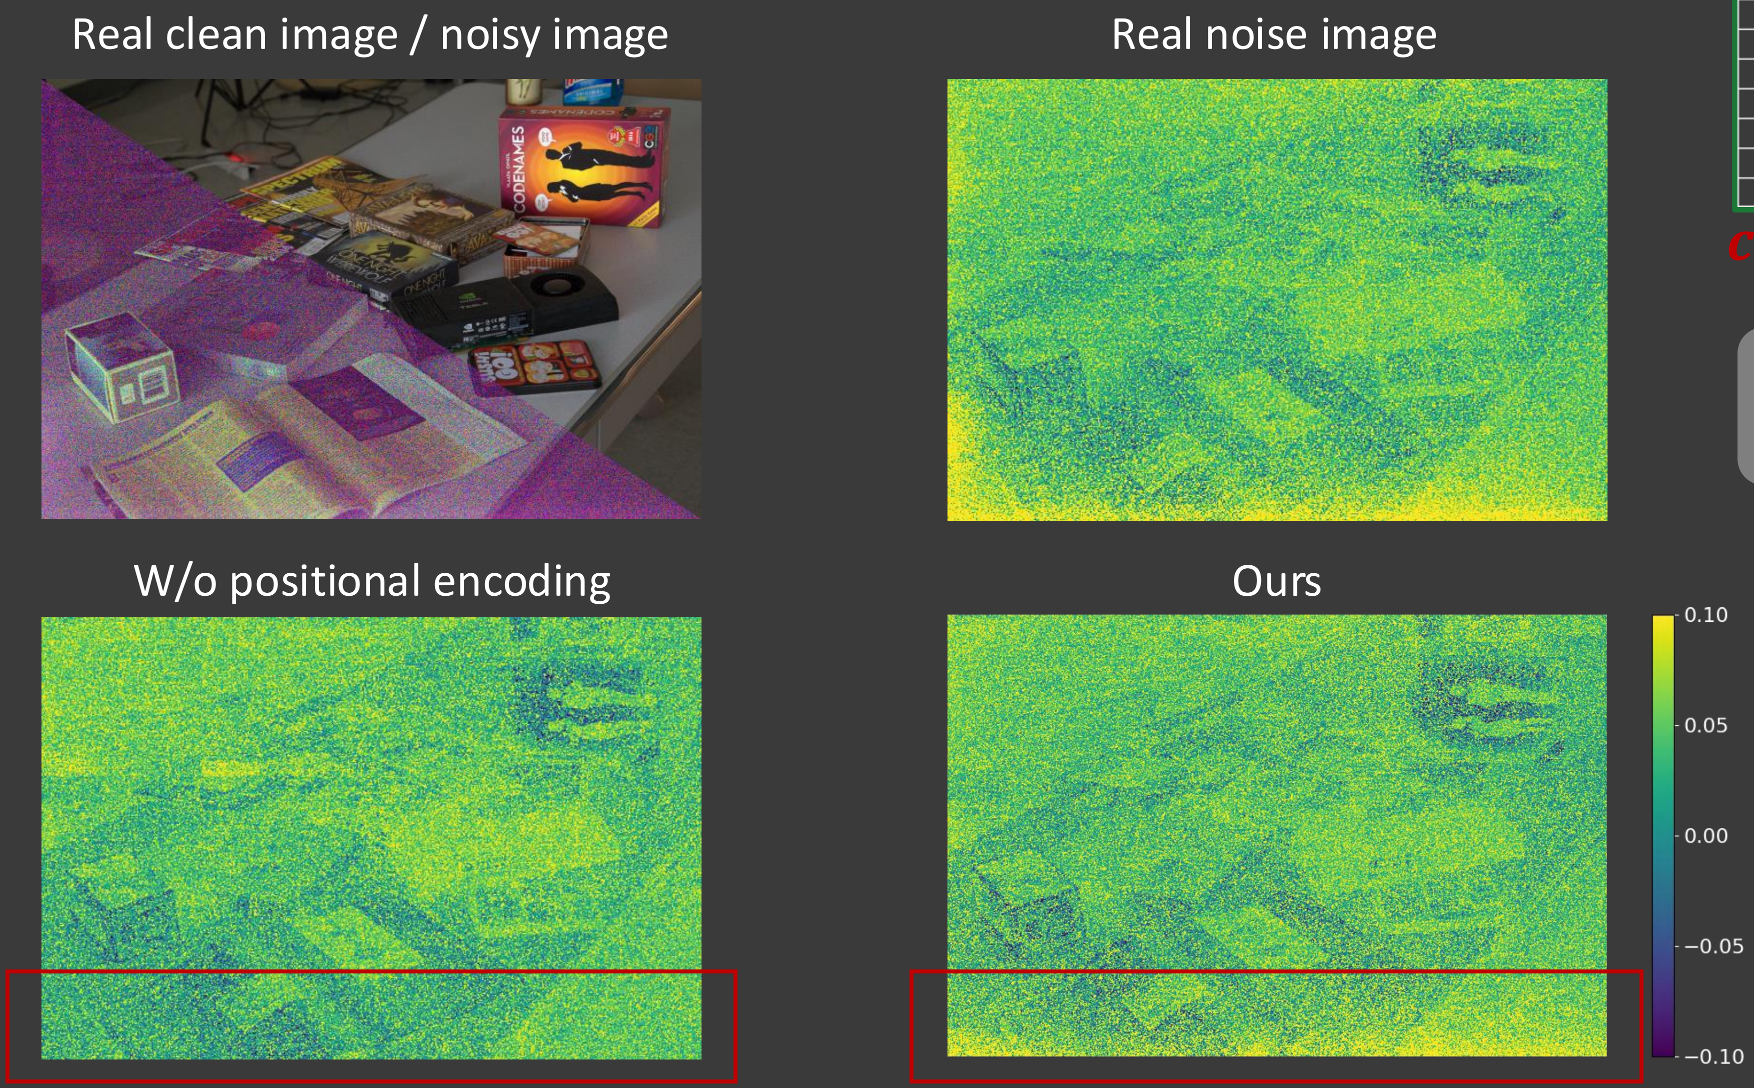

Ablations   Diffusion network: two-branch architecture

Real noise image

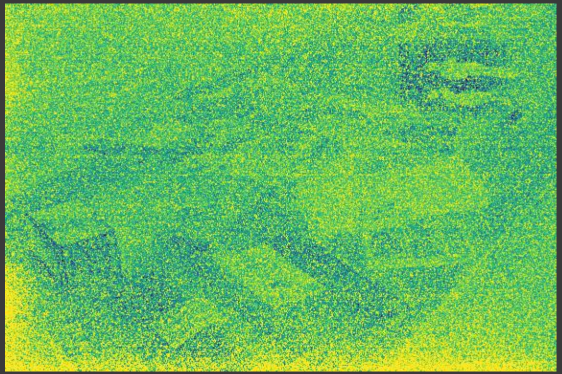

UNet-only

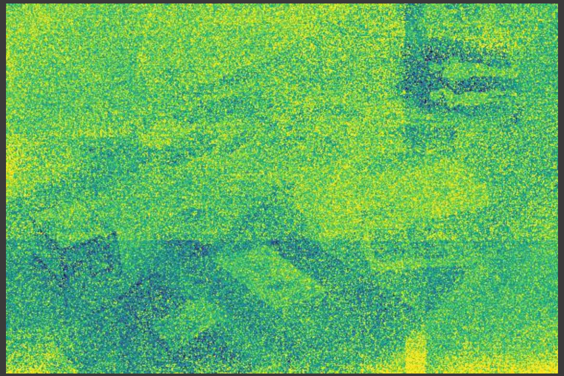

UNet+MLP

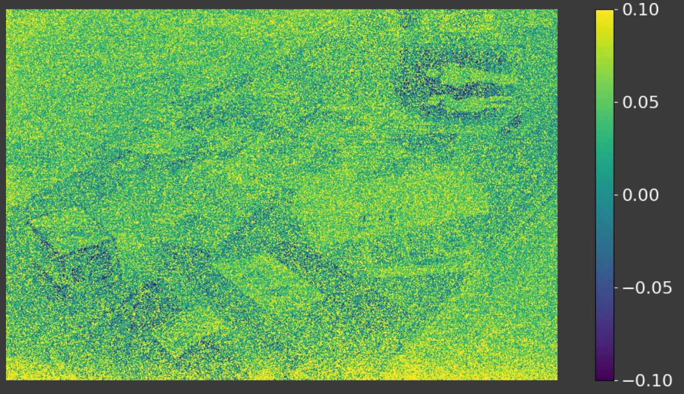

Real clean / noisy image