@anirbanray

@anirbanray_

ResMatching: Noise-Resilient Computational Super-Resolution via Guided Conditional Flow Matching

Anirban Ray, Vera Galinova, and Florian Jug

Human Technopole, Italy

Technische Universität Dresden, Germany

Triangle of Frustration (in Light Microscopy)

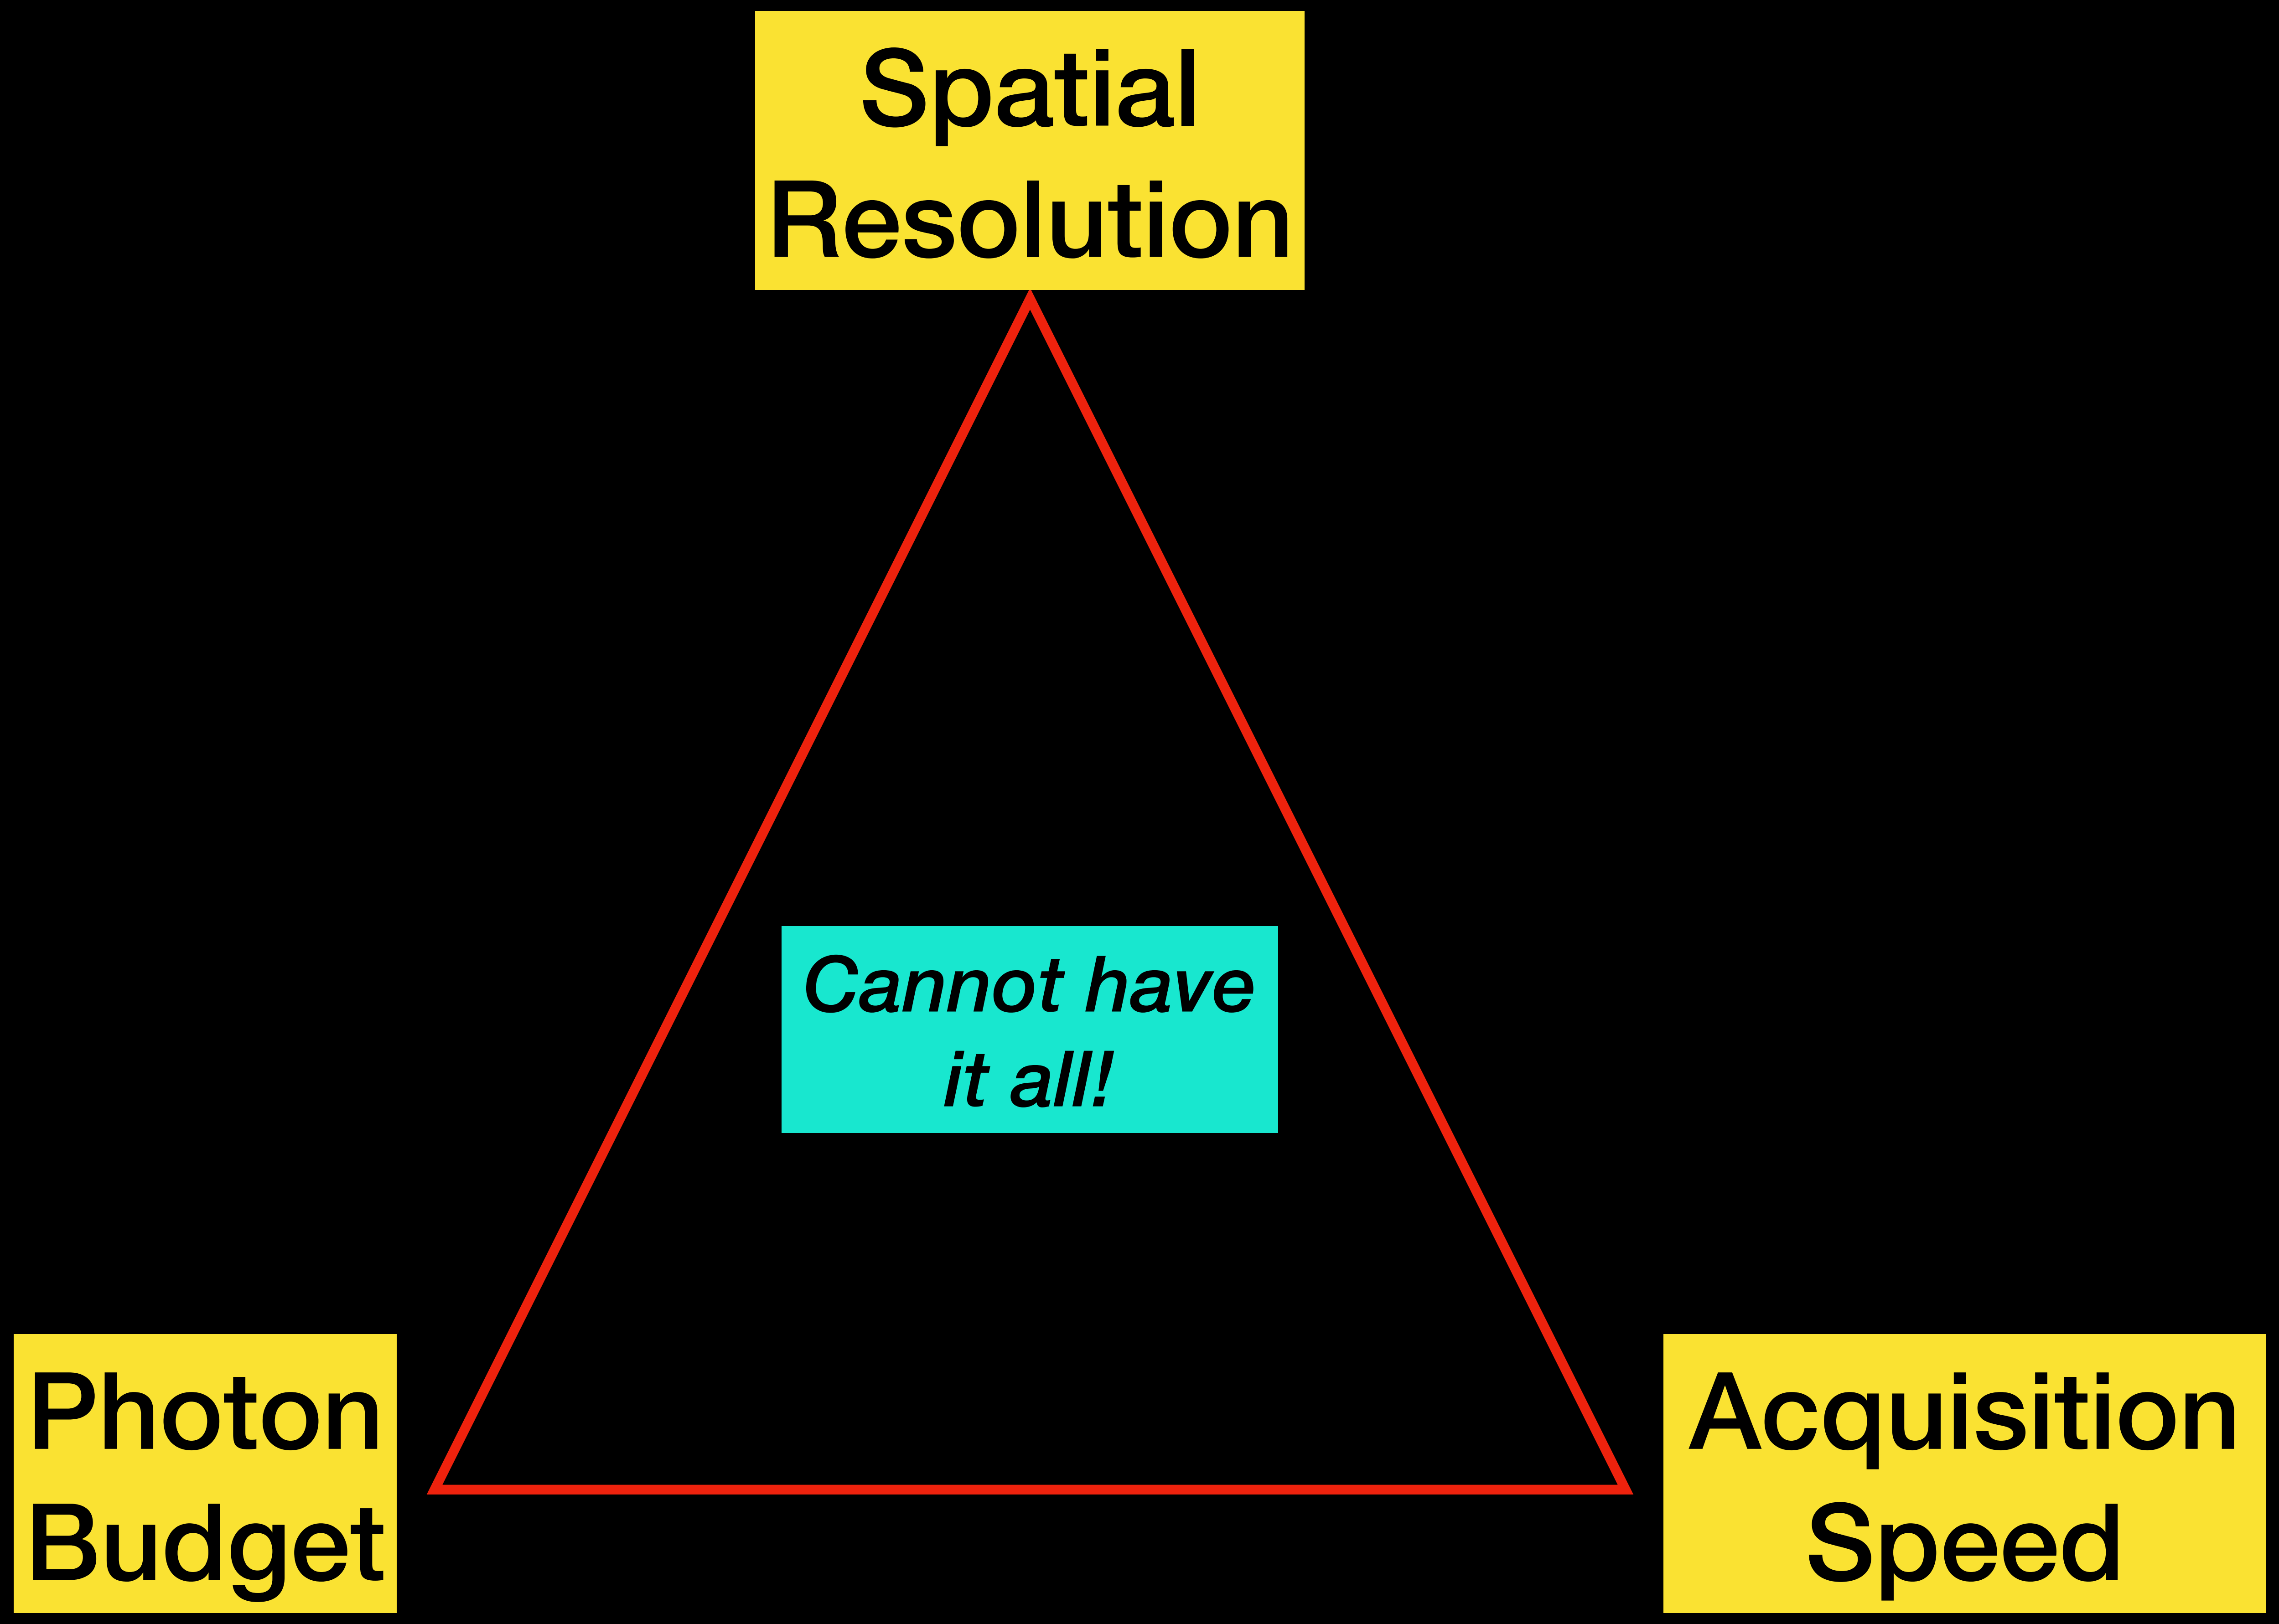

Triangle of Frustration (in Light Microscopy)

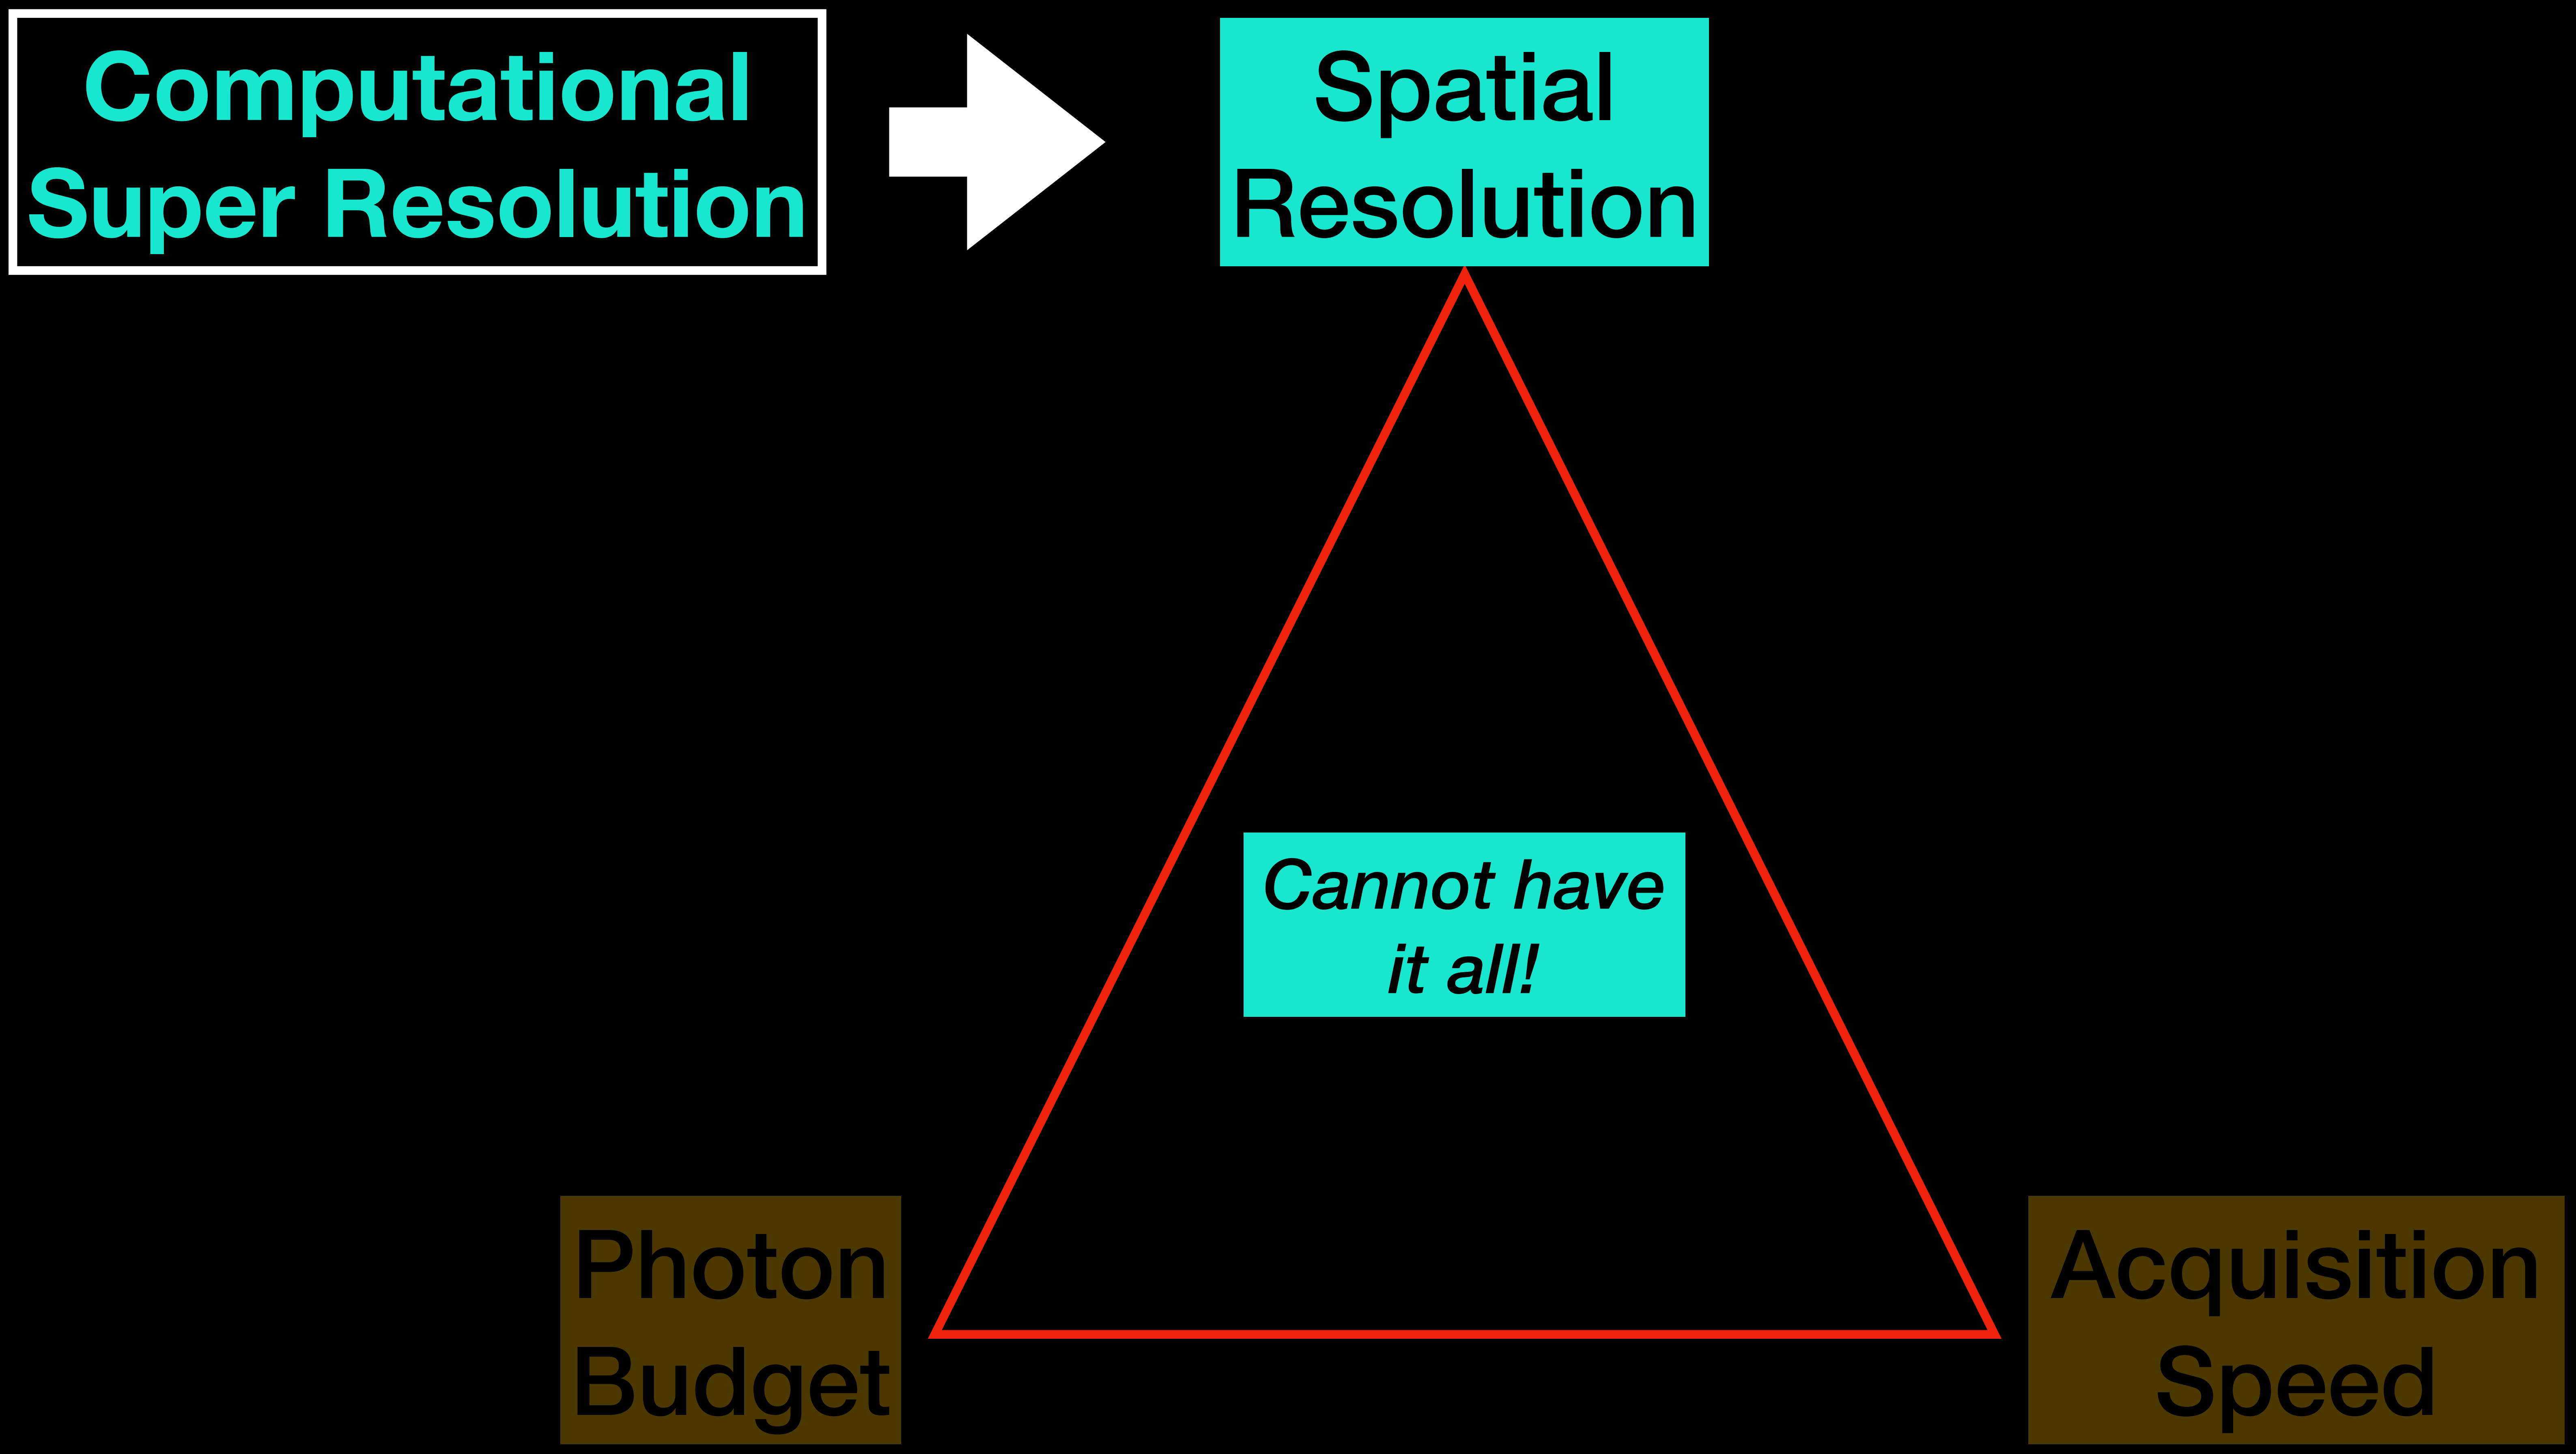

Computational Super Resolution (CSR)

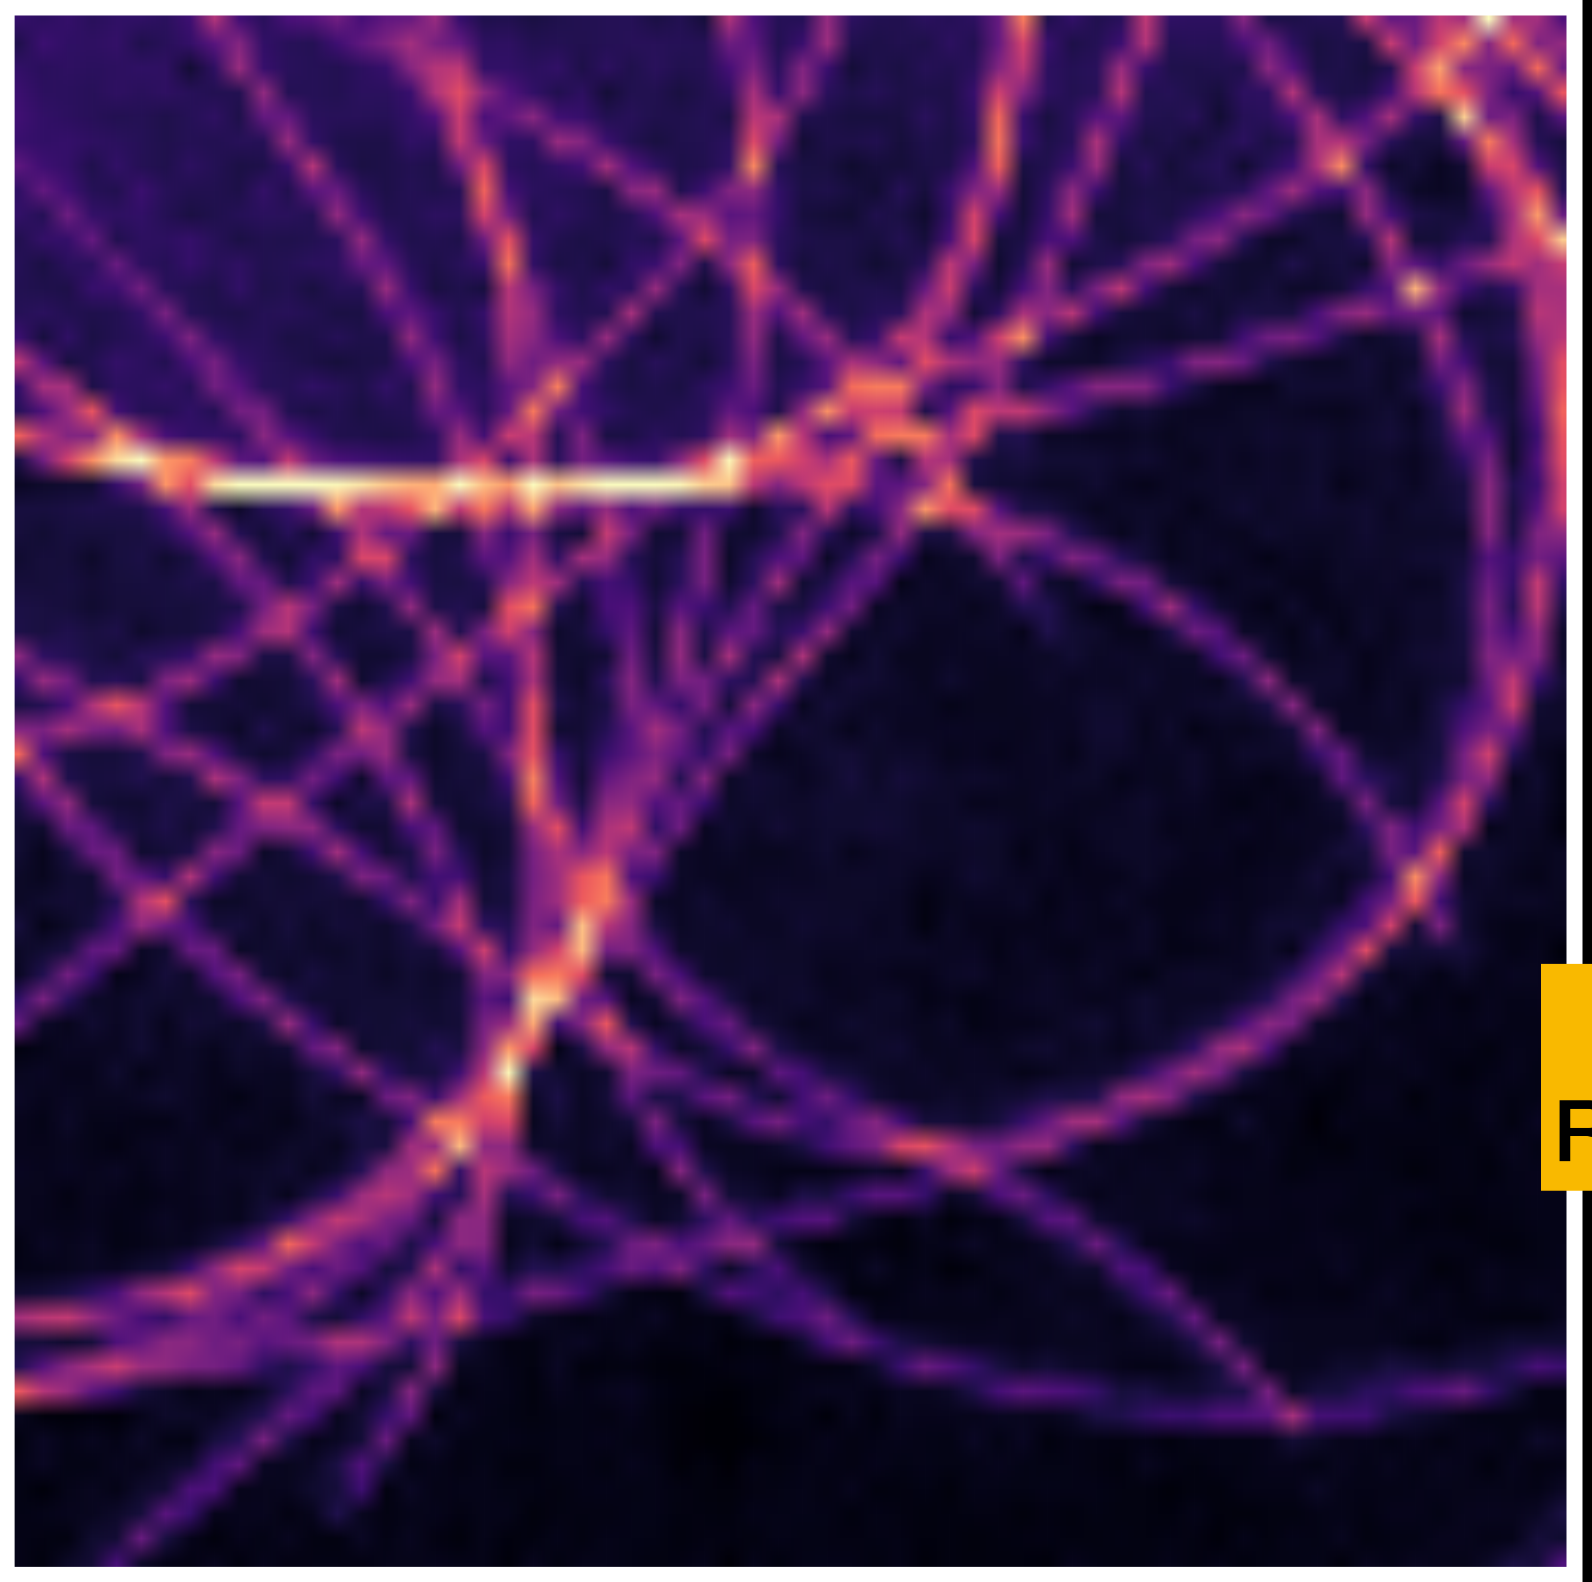

Gain Resolution

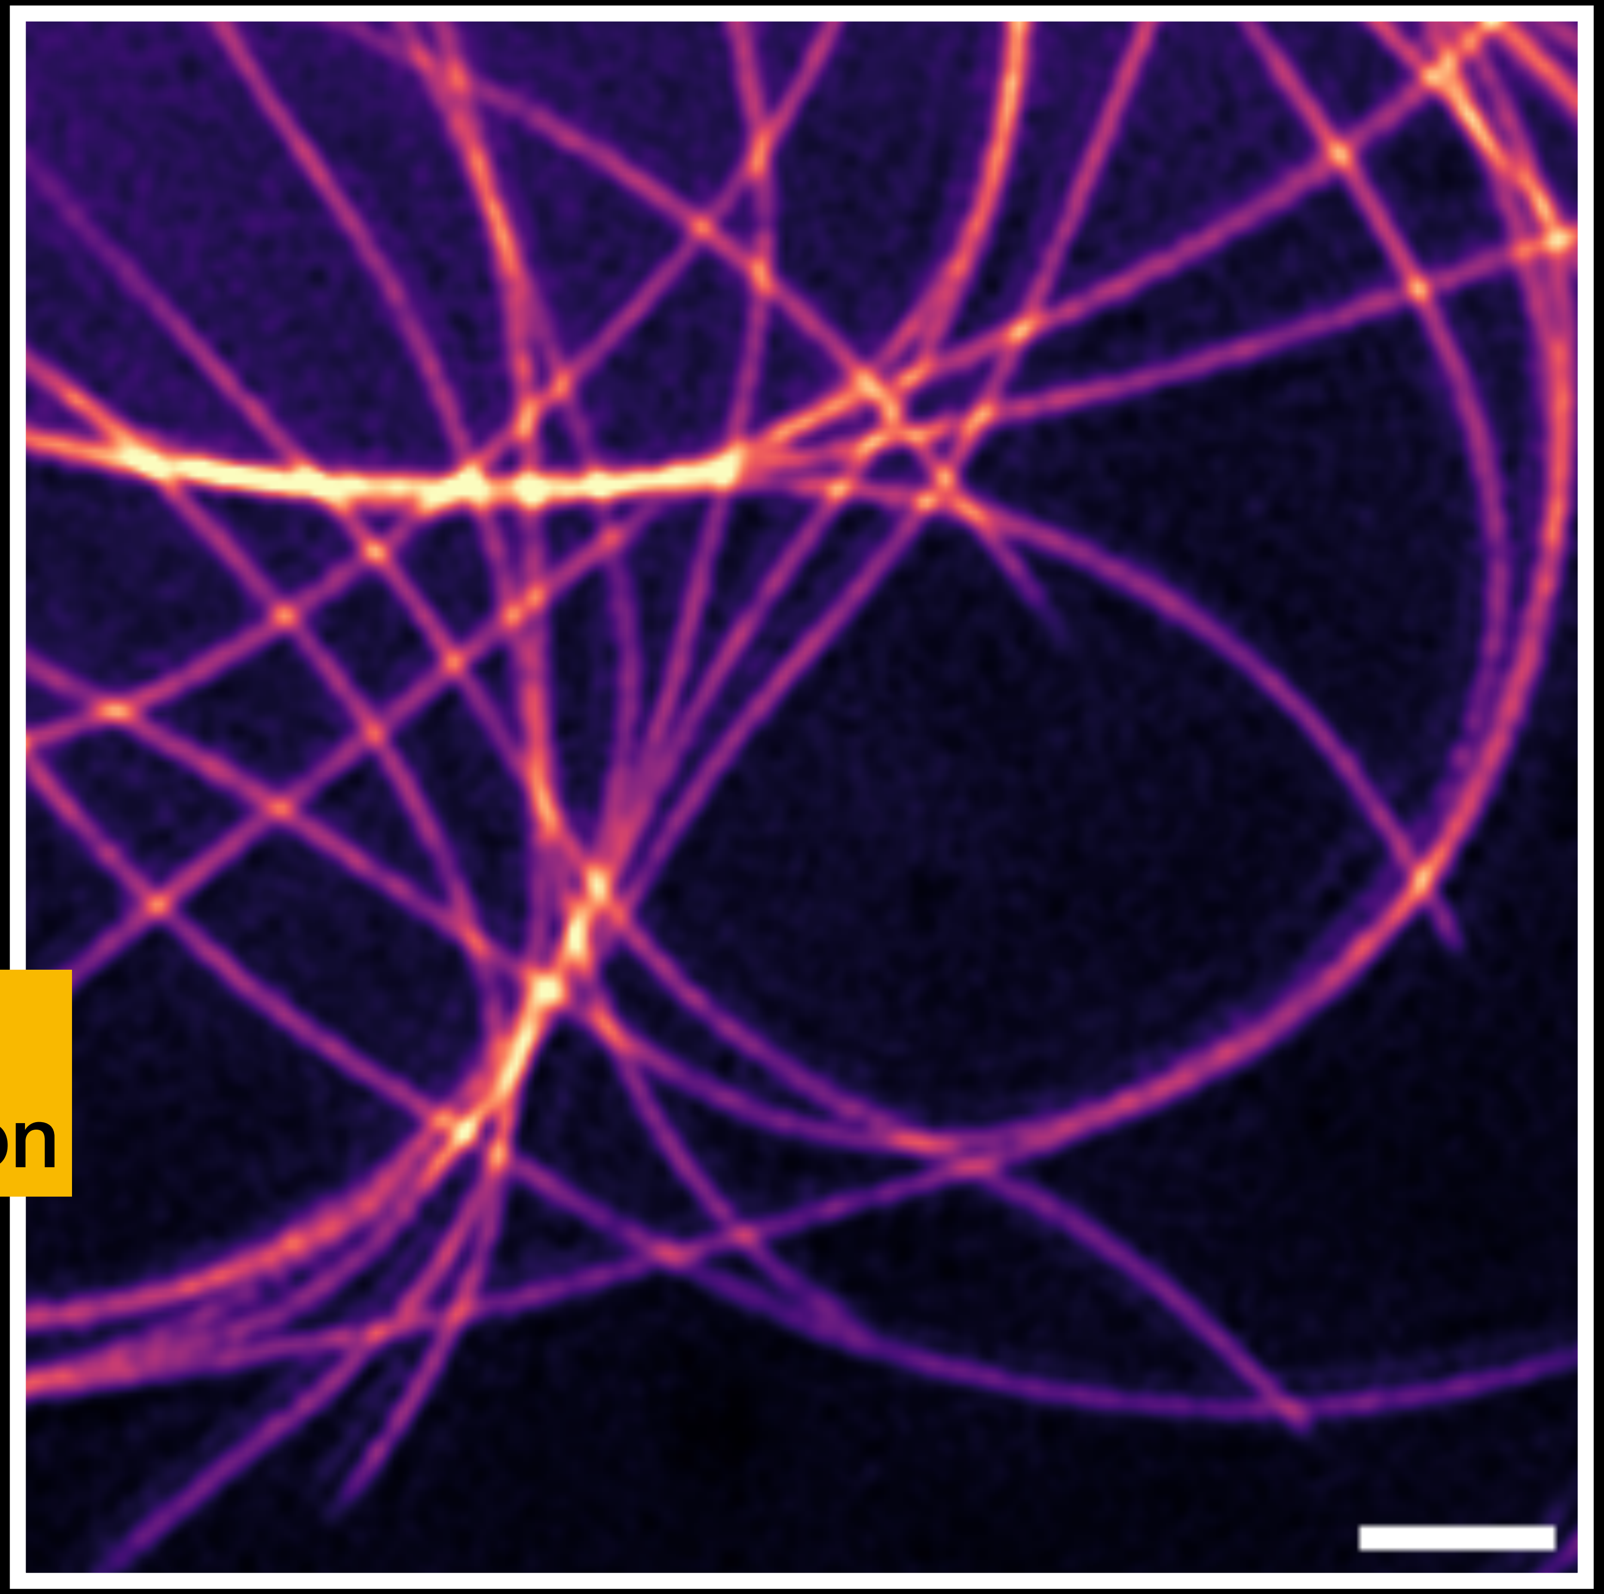

Low Resolution Acquired Image

High Resolution Image

- CSR aims to gain resolution by recovering missing high-frequency information.
- CSR is an ill-posed inverse problem.
- Deep learning enables capturing strong data-driven priors.

Data: BioSR Microtubule, Scale: $1\mu m$

CSR under heavy noise?

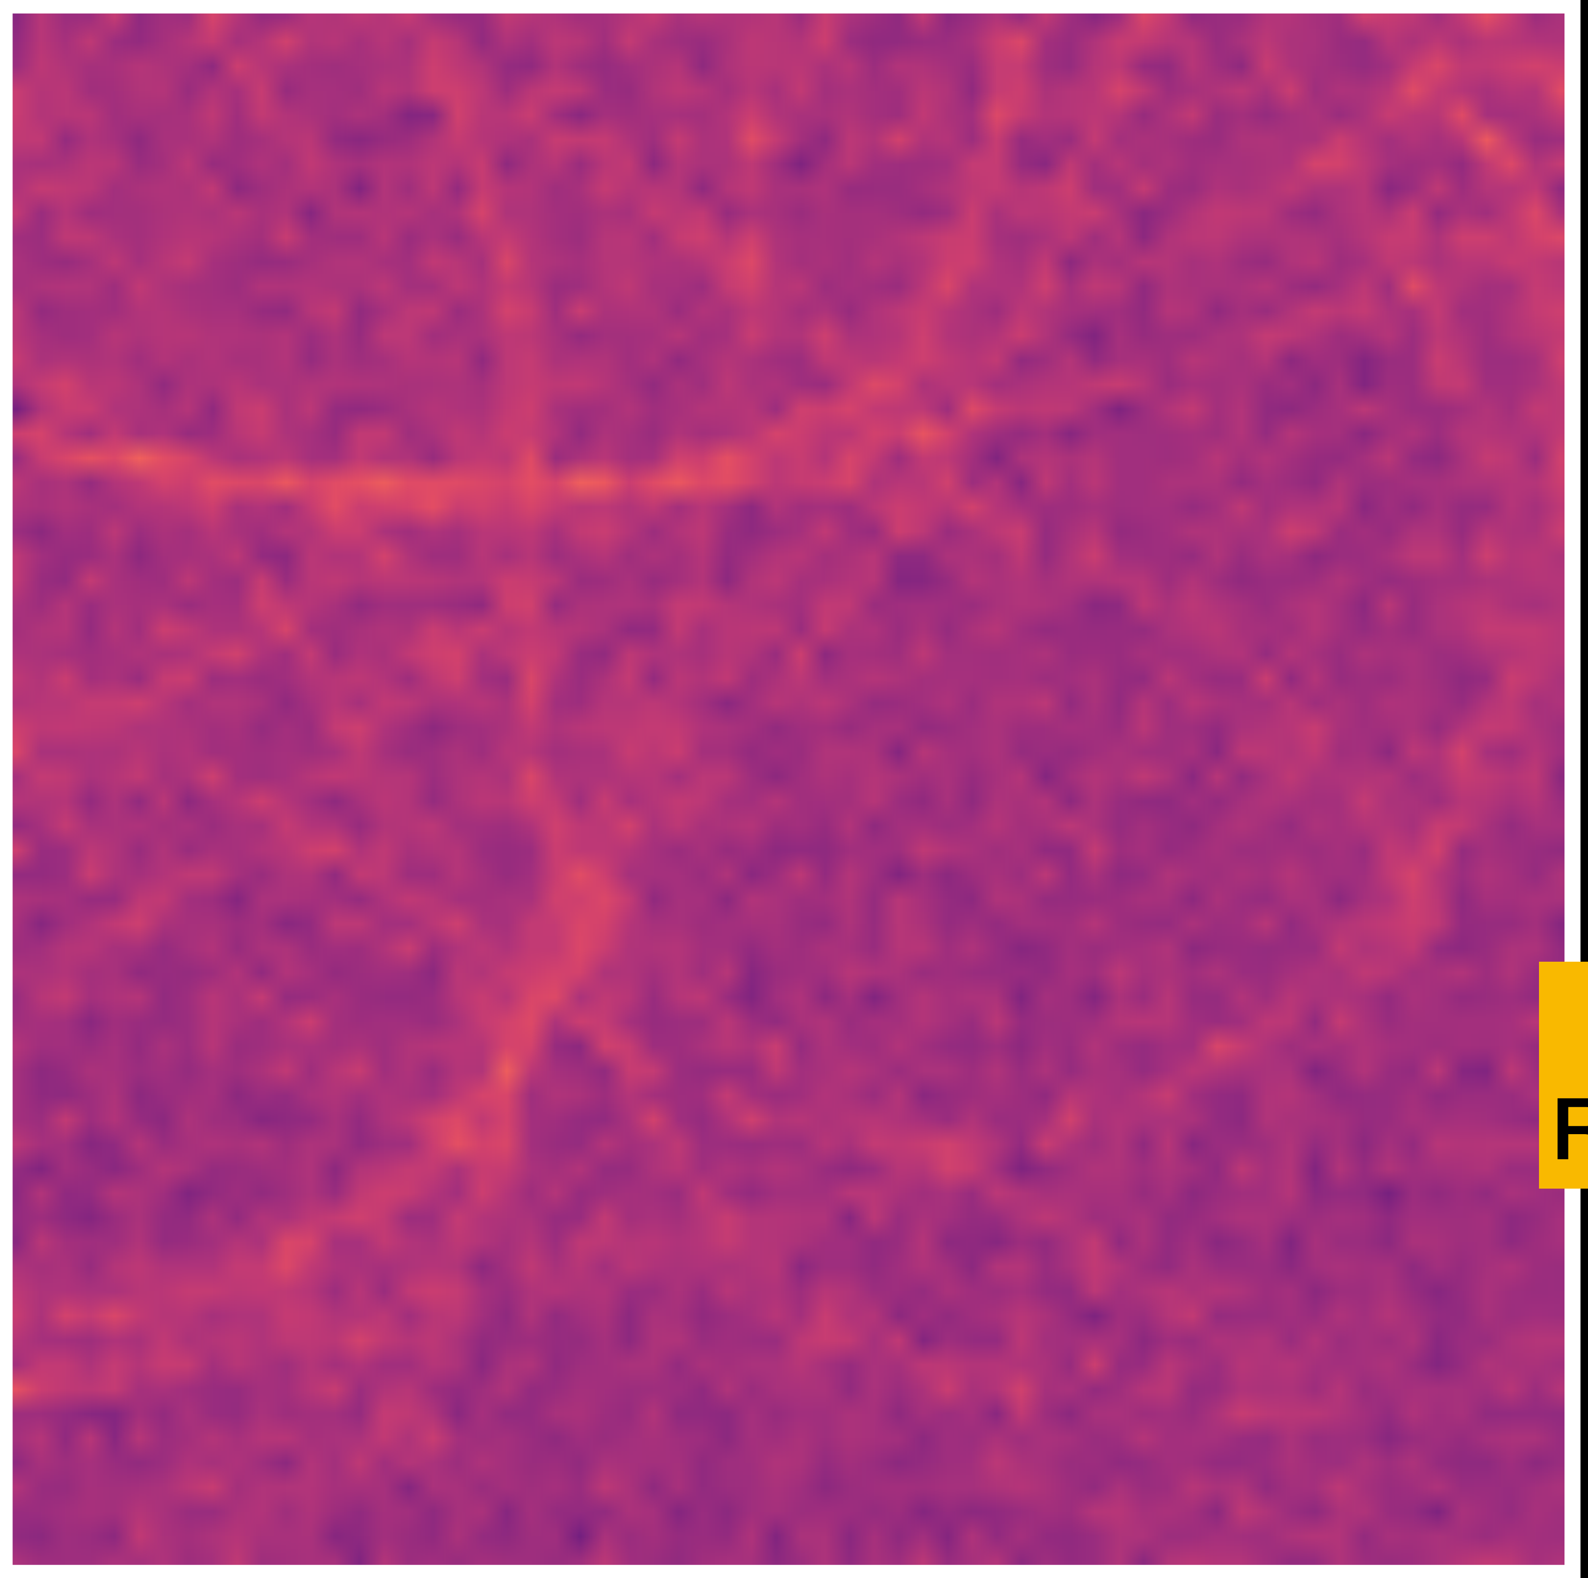

Noisy Low Resolution Acquired Image

Gain Resolution

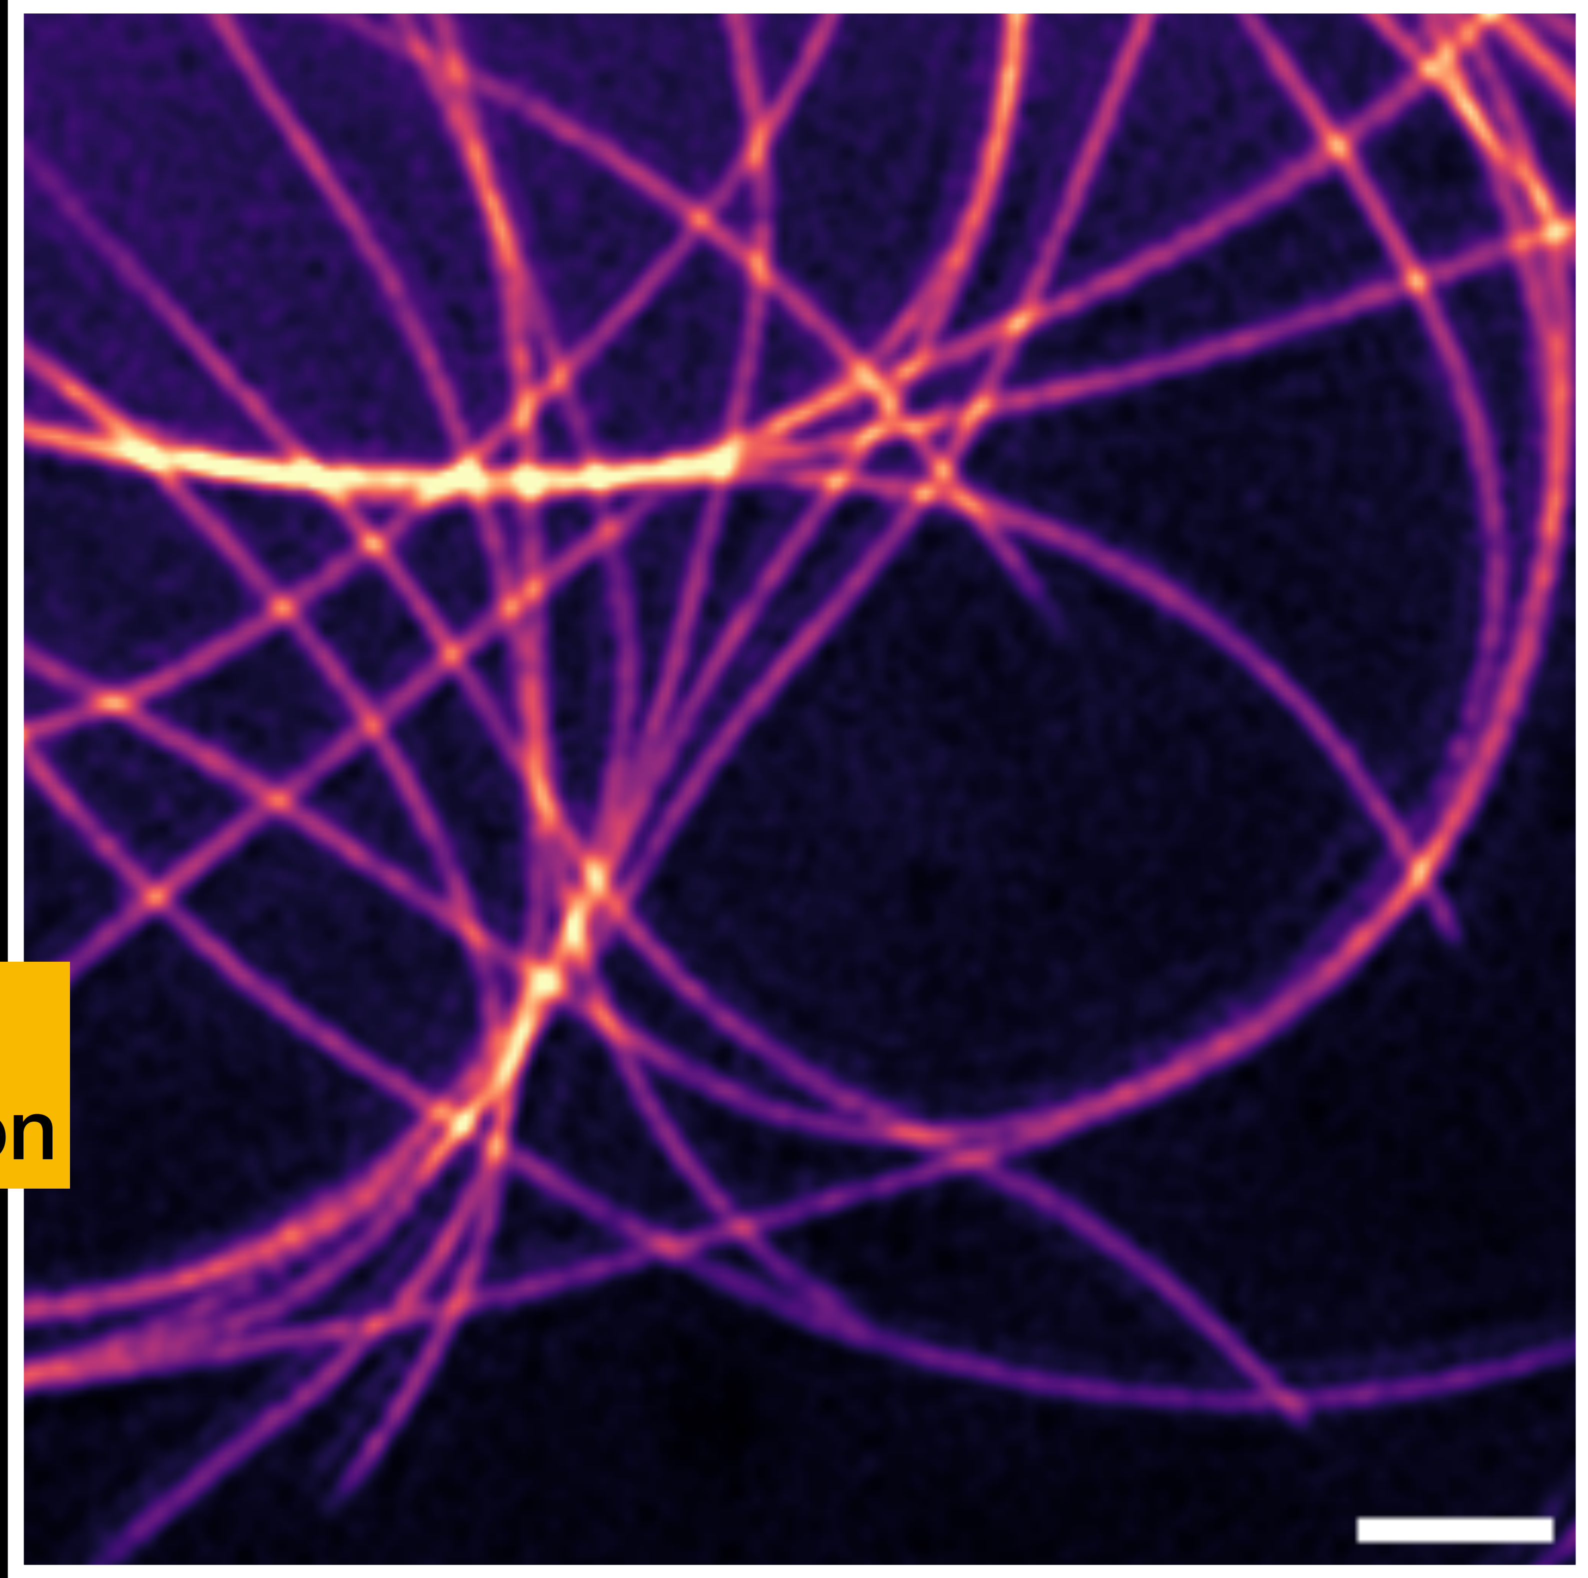

High Resolution Image

- Enable **imaging at low light** (less phototoxicity & photobleaching).
- Improve **temporal resolution** in live-cell imaging.
- …

Stress test Generative Image Restoration

Data: BioSR Microtubule, Scale: $1 \mu m$

Traditional Deep Learning (CSR) Method

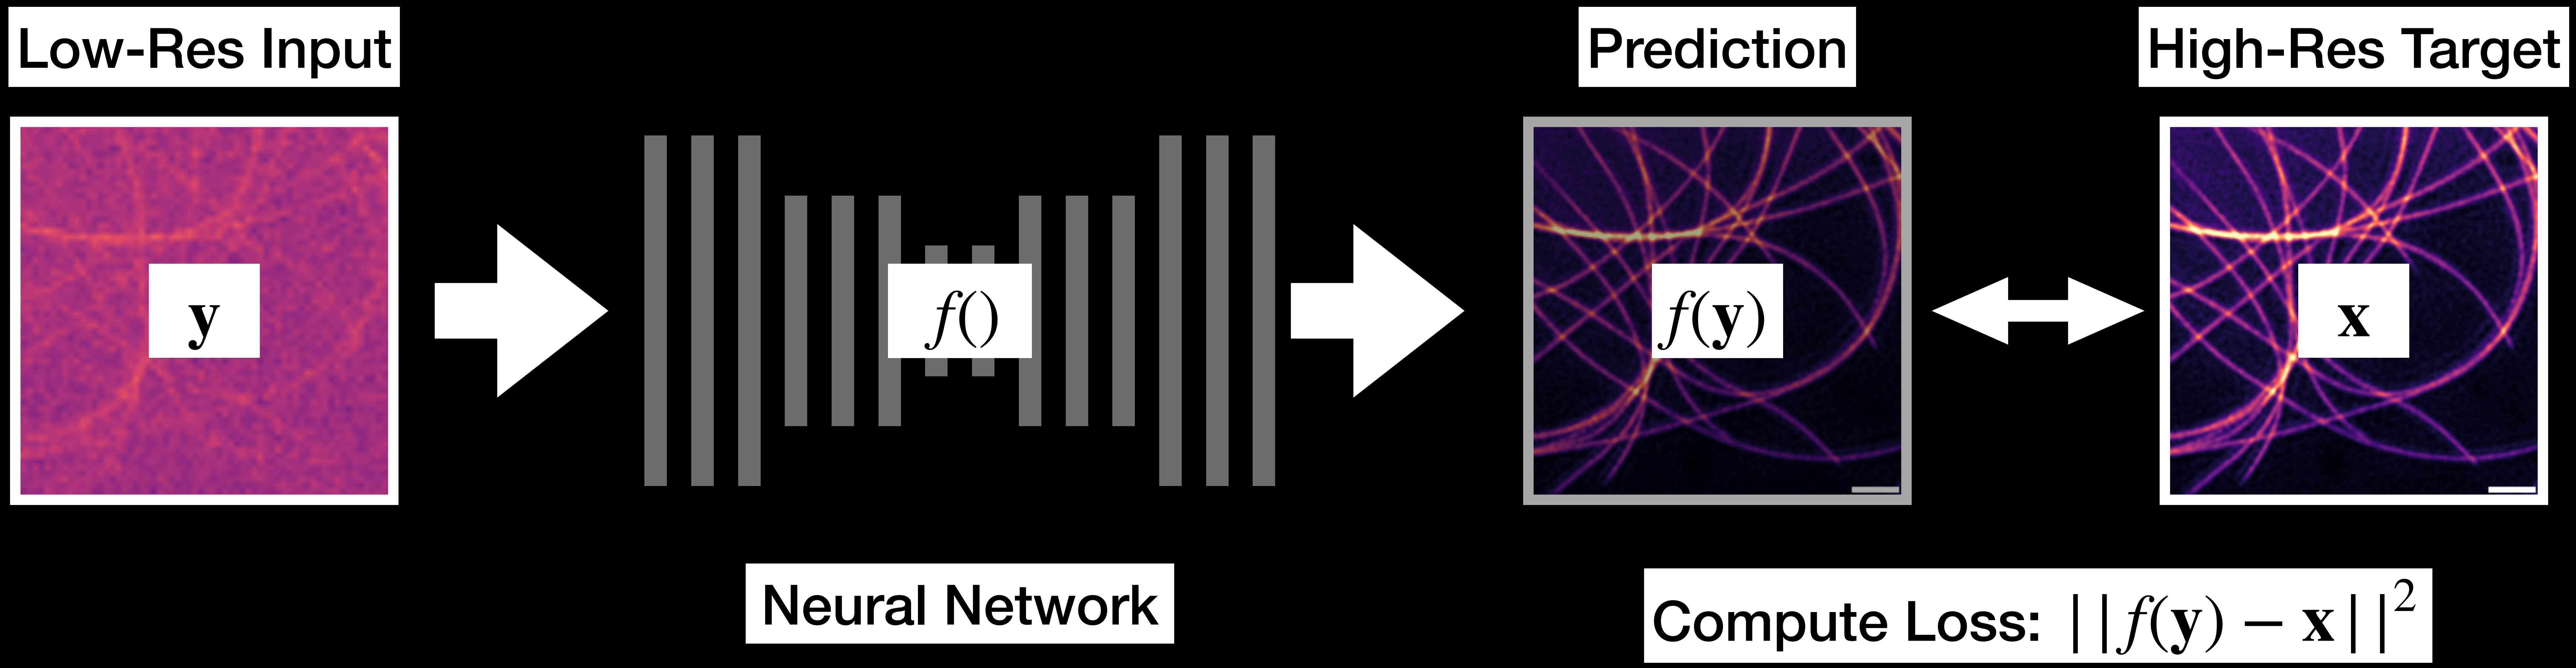

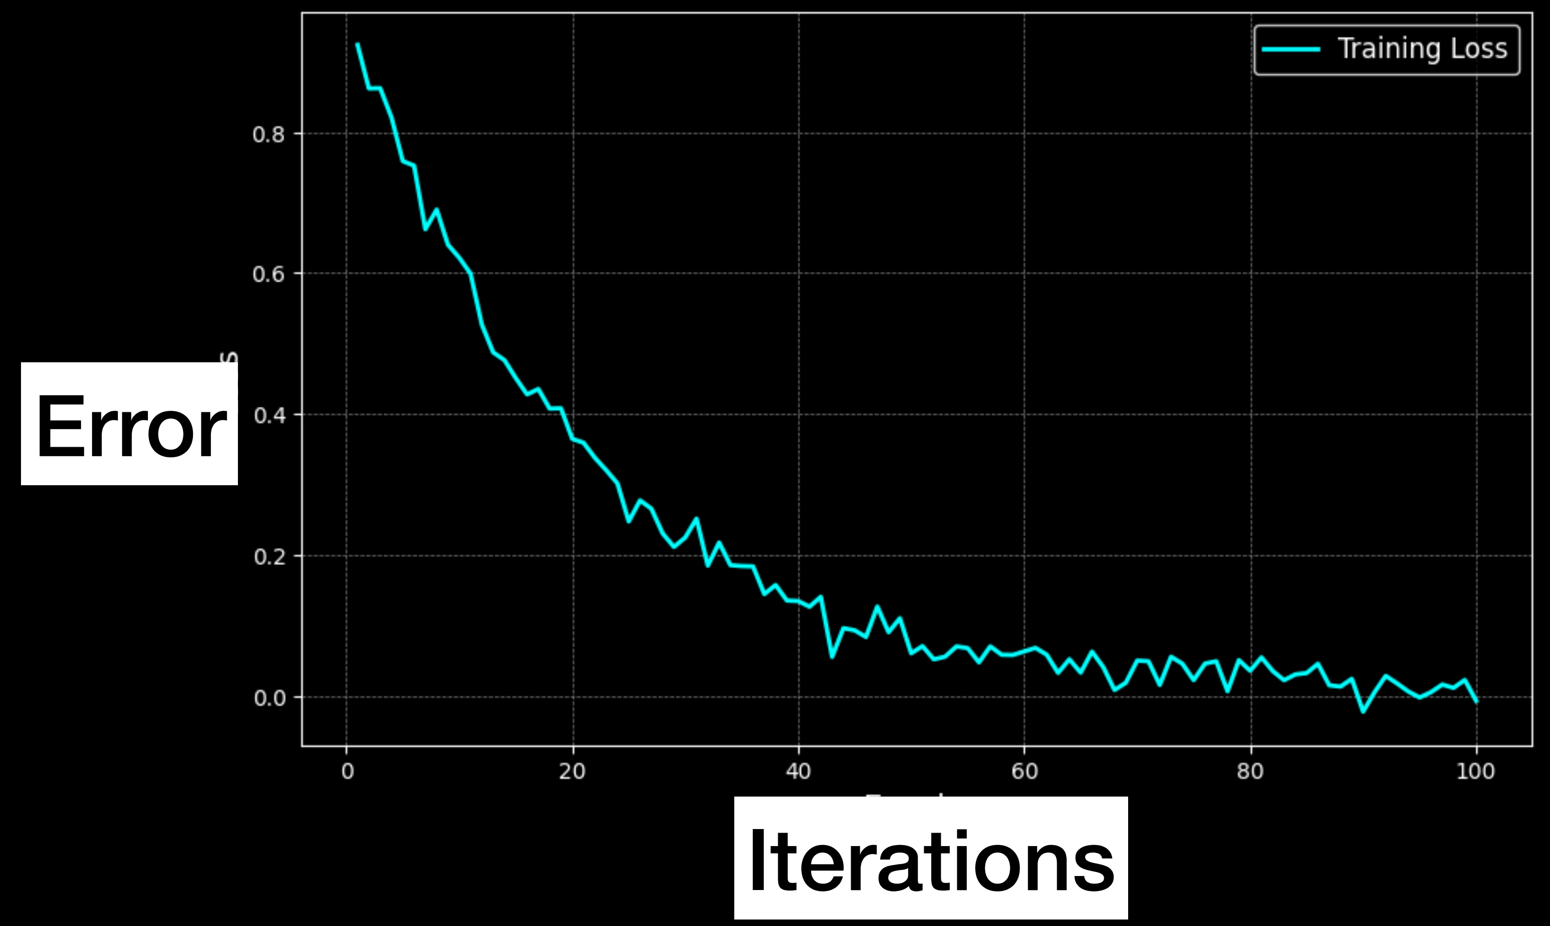

Traditional Deep Learning (CSR) Method

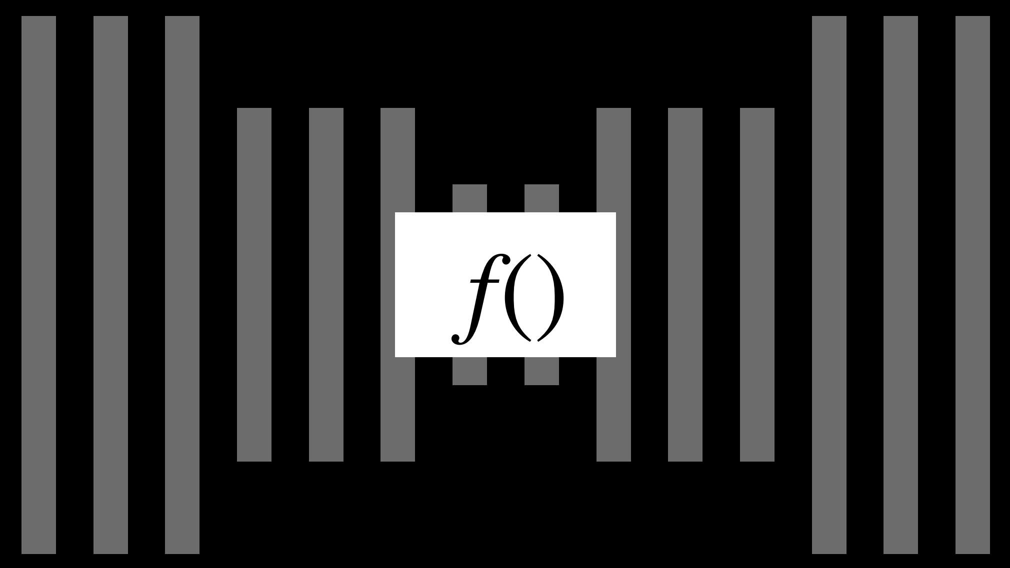

Neural Network

- CNN / attention
- Physics informed
- Bayesian
- GAN
- Transformer based
- Frequency based
- …

Various flavors

Traditional Deep Learning (CSR) Method

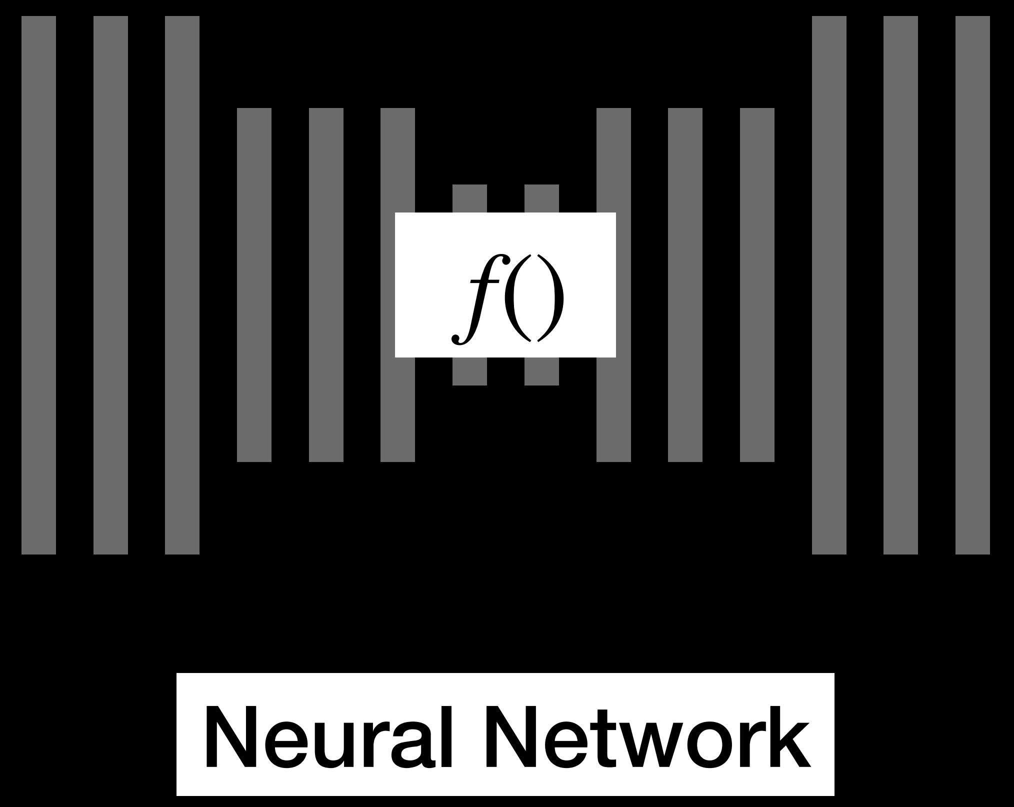

- CNN / attention
- Physics informed
- Bayesian
- GAN
- Transformer based
- Frequency based
- …

Various flavors

One-Shot Mapping!

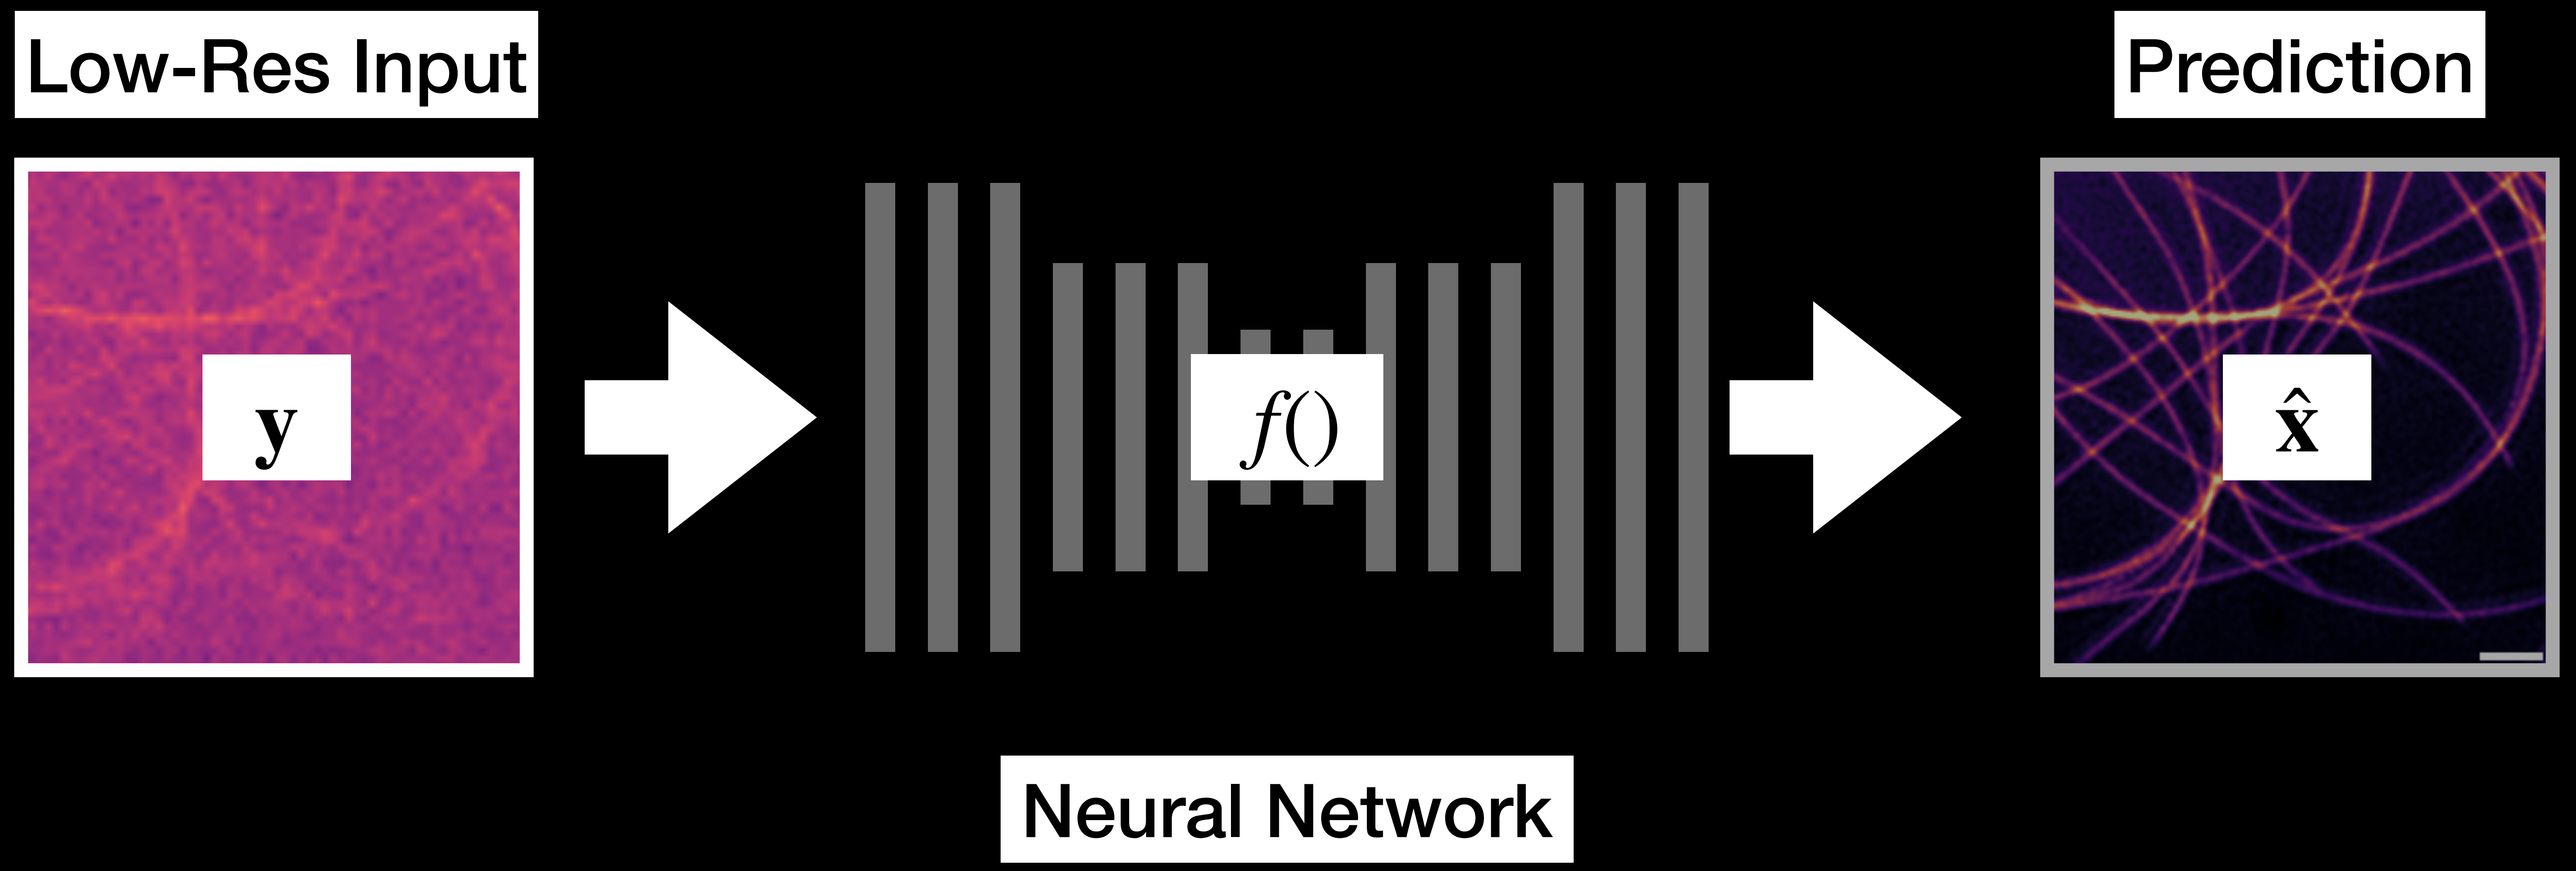

Prediction is usually done in a single step

Perception Distortion Tradeoff (PDT)

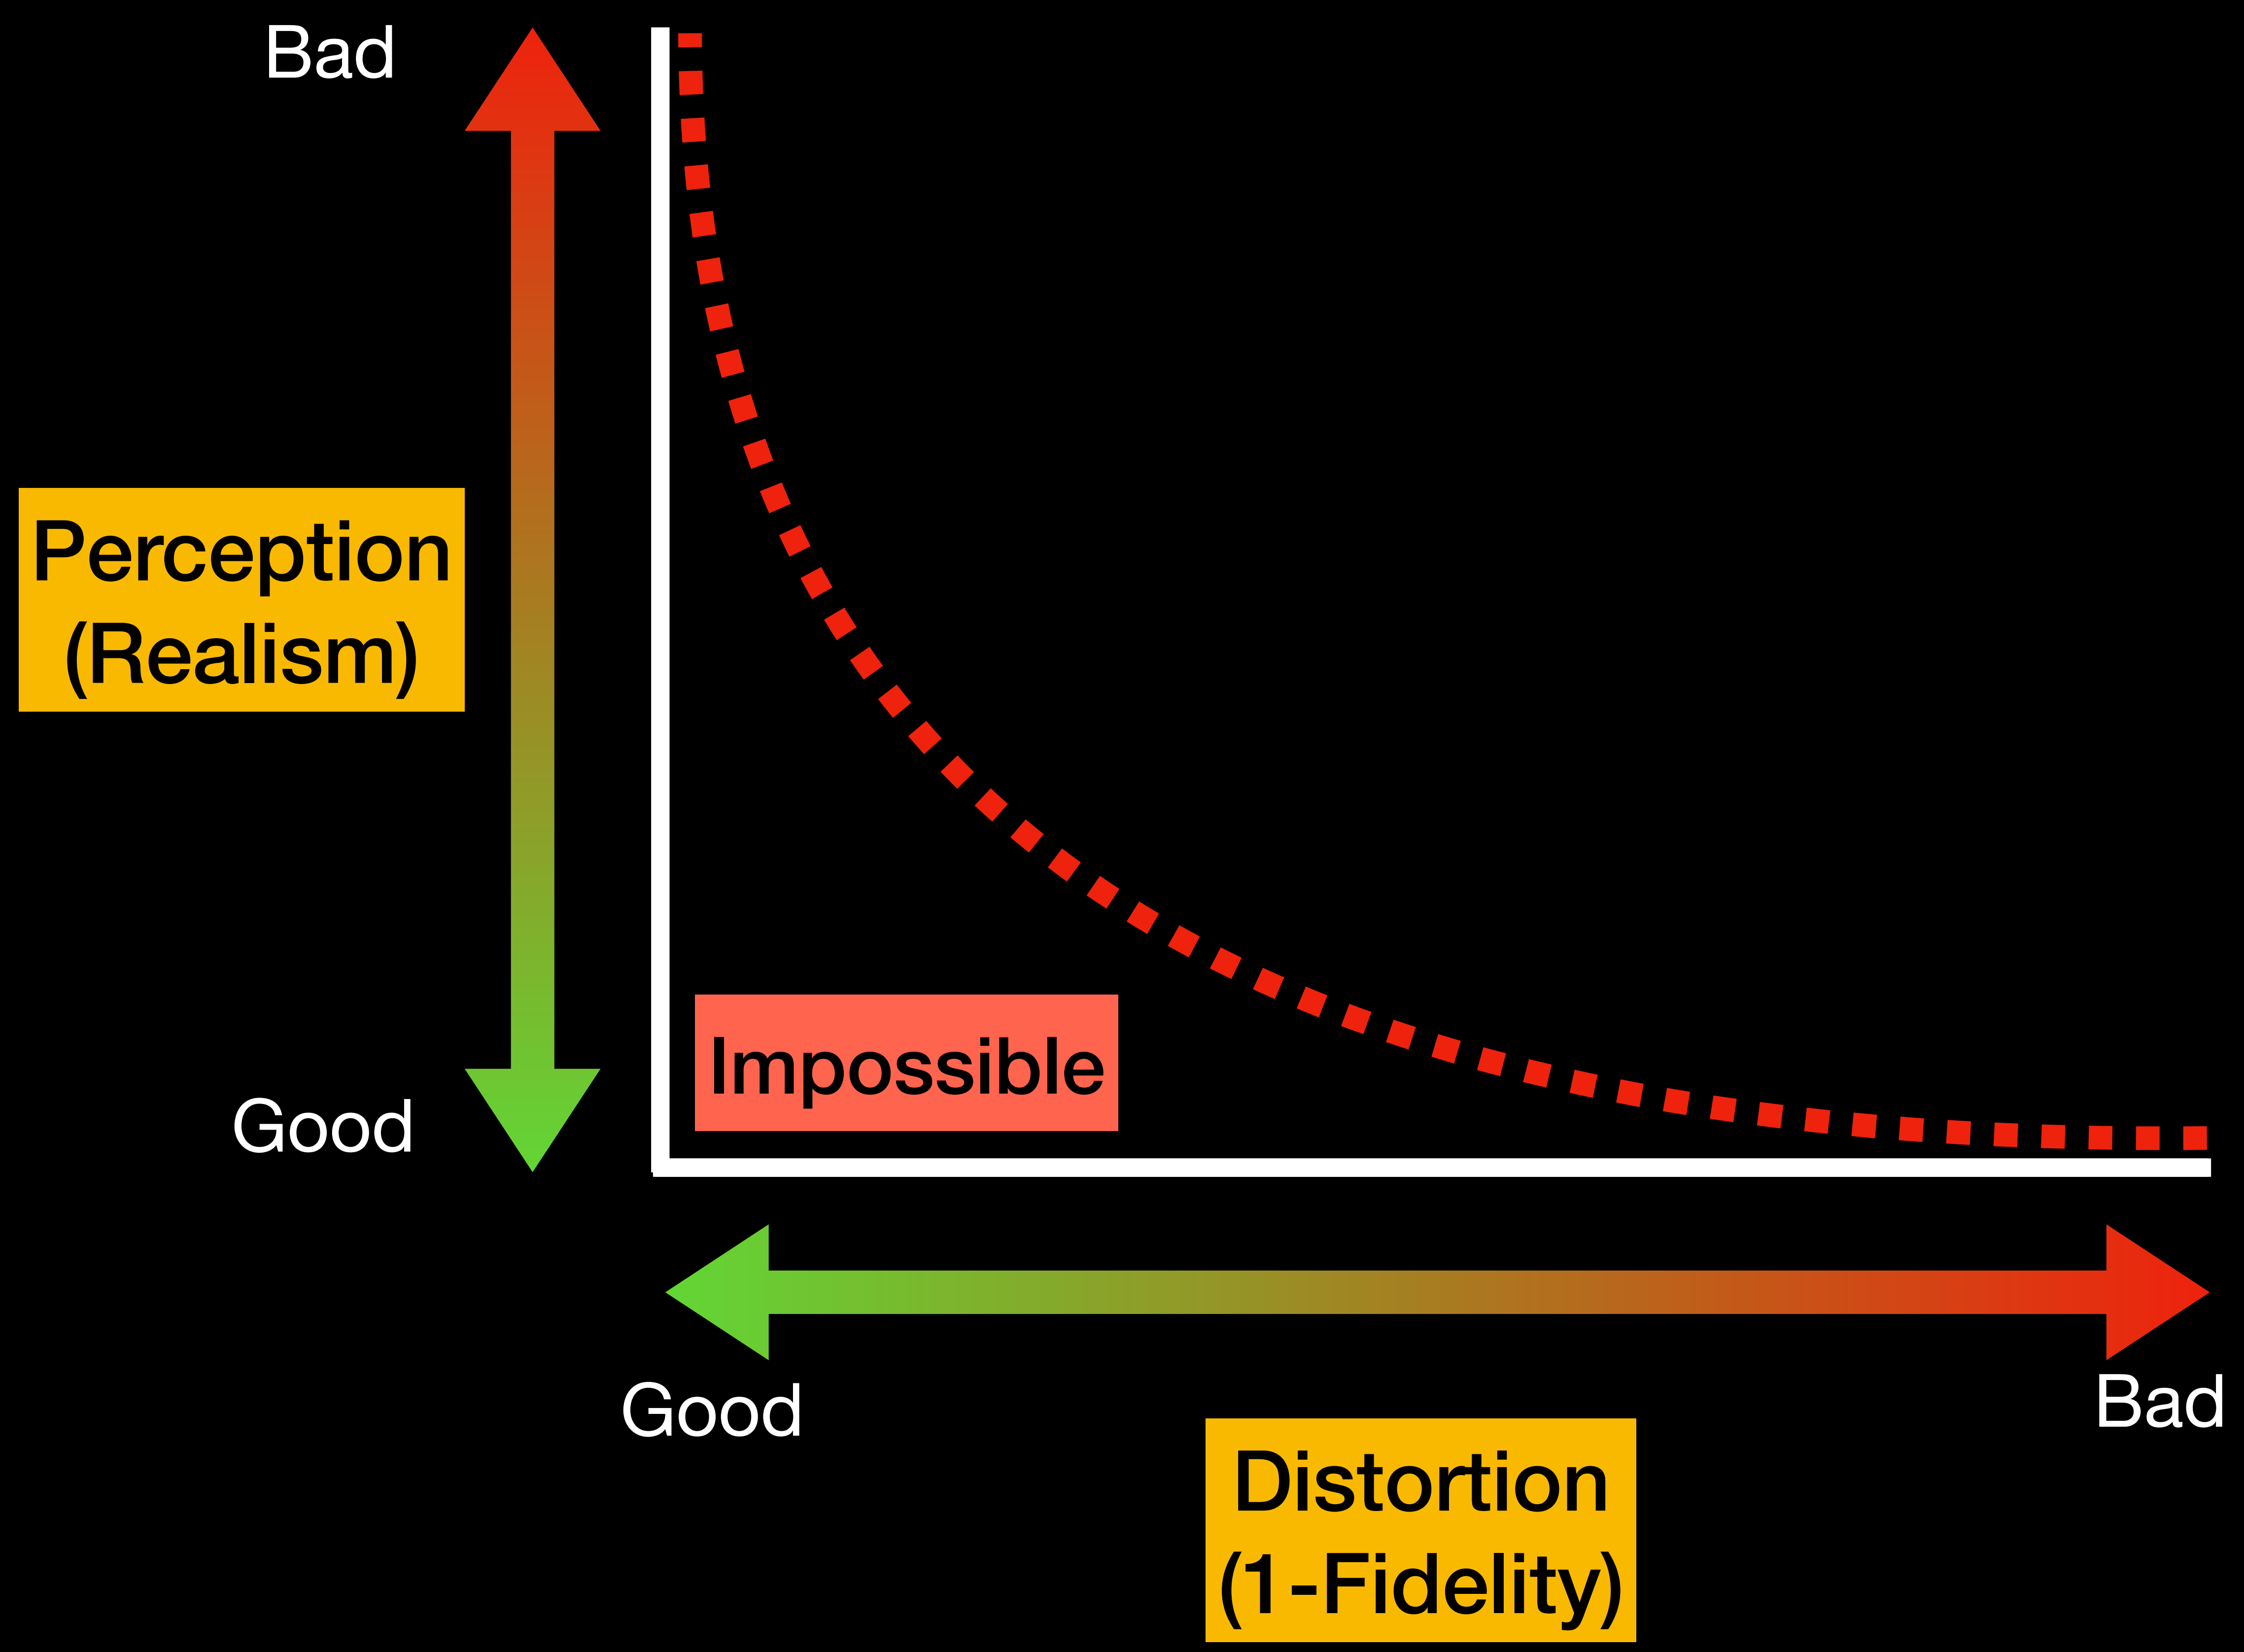

*The Perception-Distortion Tradeoff, Blau et al. - CVPR 2018

Perception Distortion Tradeoff (PDT)

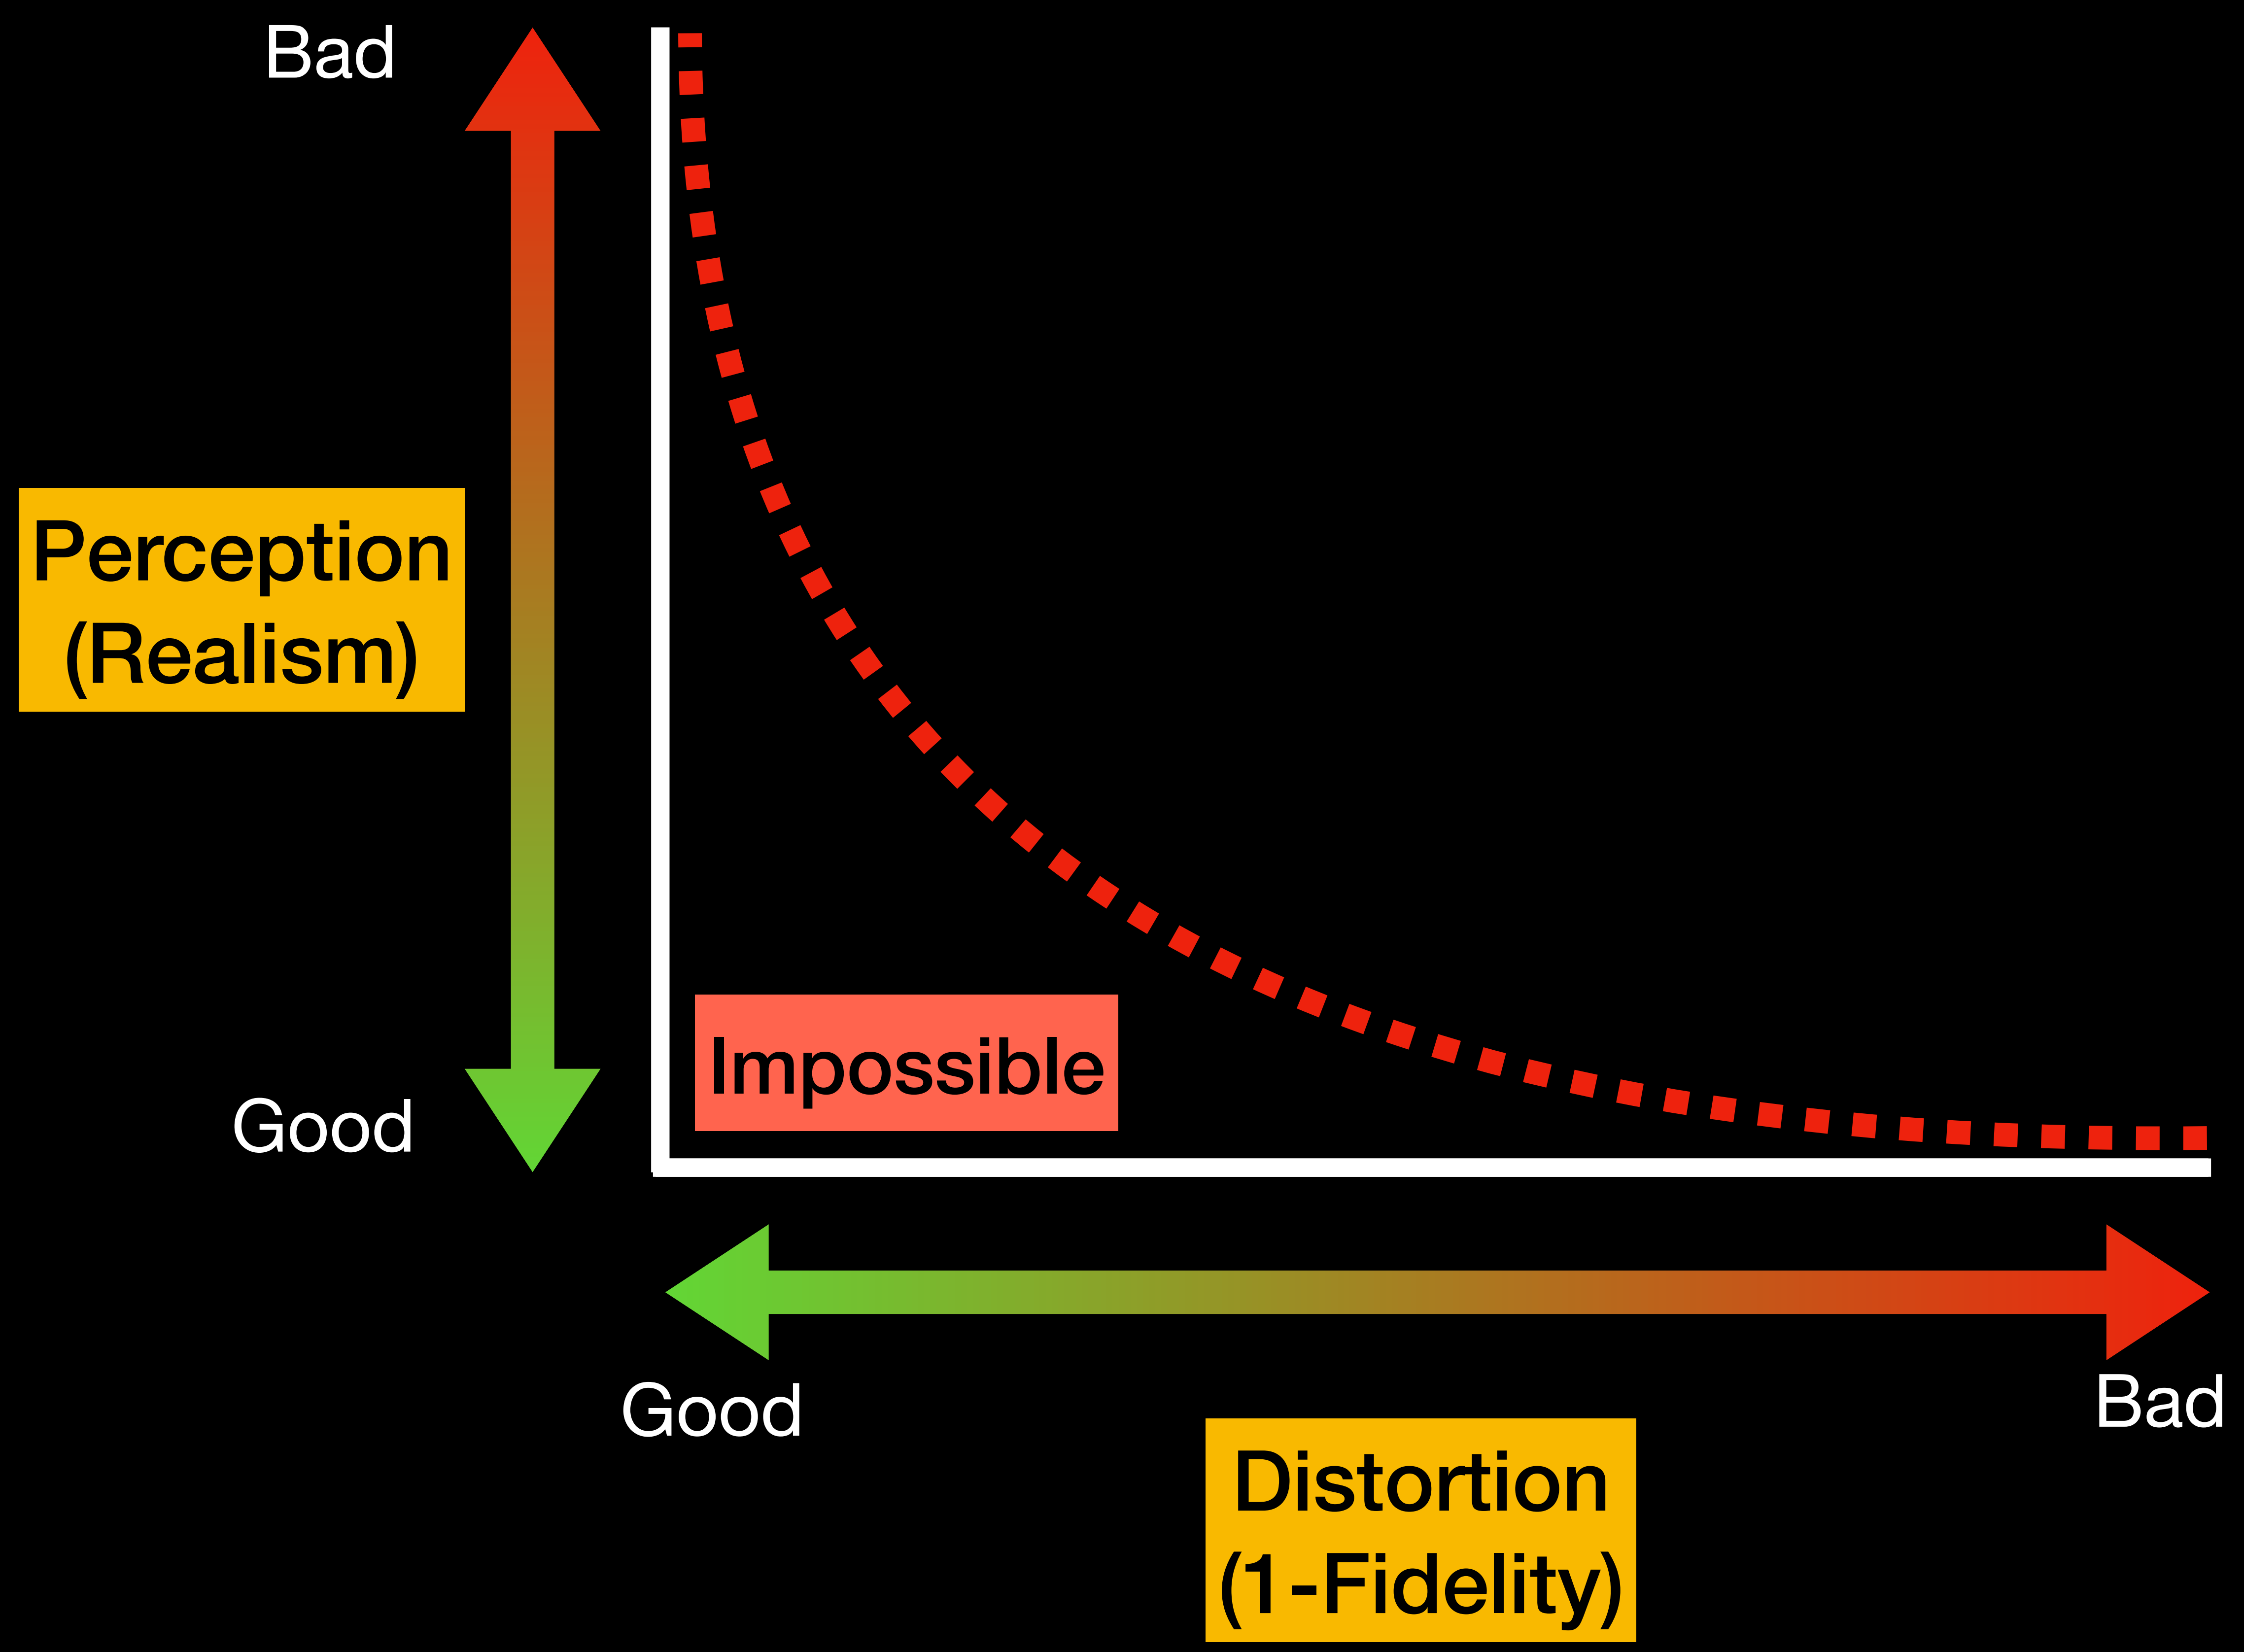

An model cannot be optimized for the *best perception* and *best distortion* at the same time!

One-shot models can be only optimized for one of these axes!

*The Perception-Distortion Tradeoff, Blau et al. - CVPR 2018

Perception Distortion Tradeoff (PDT)

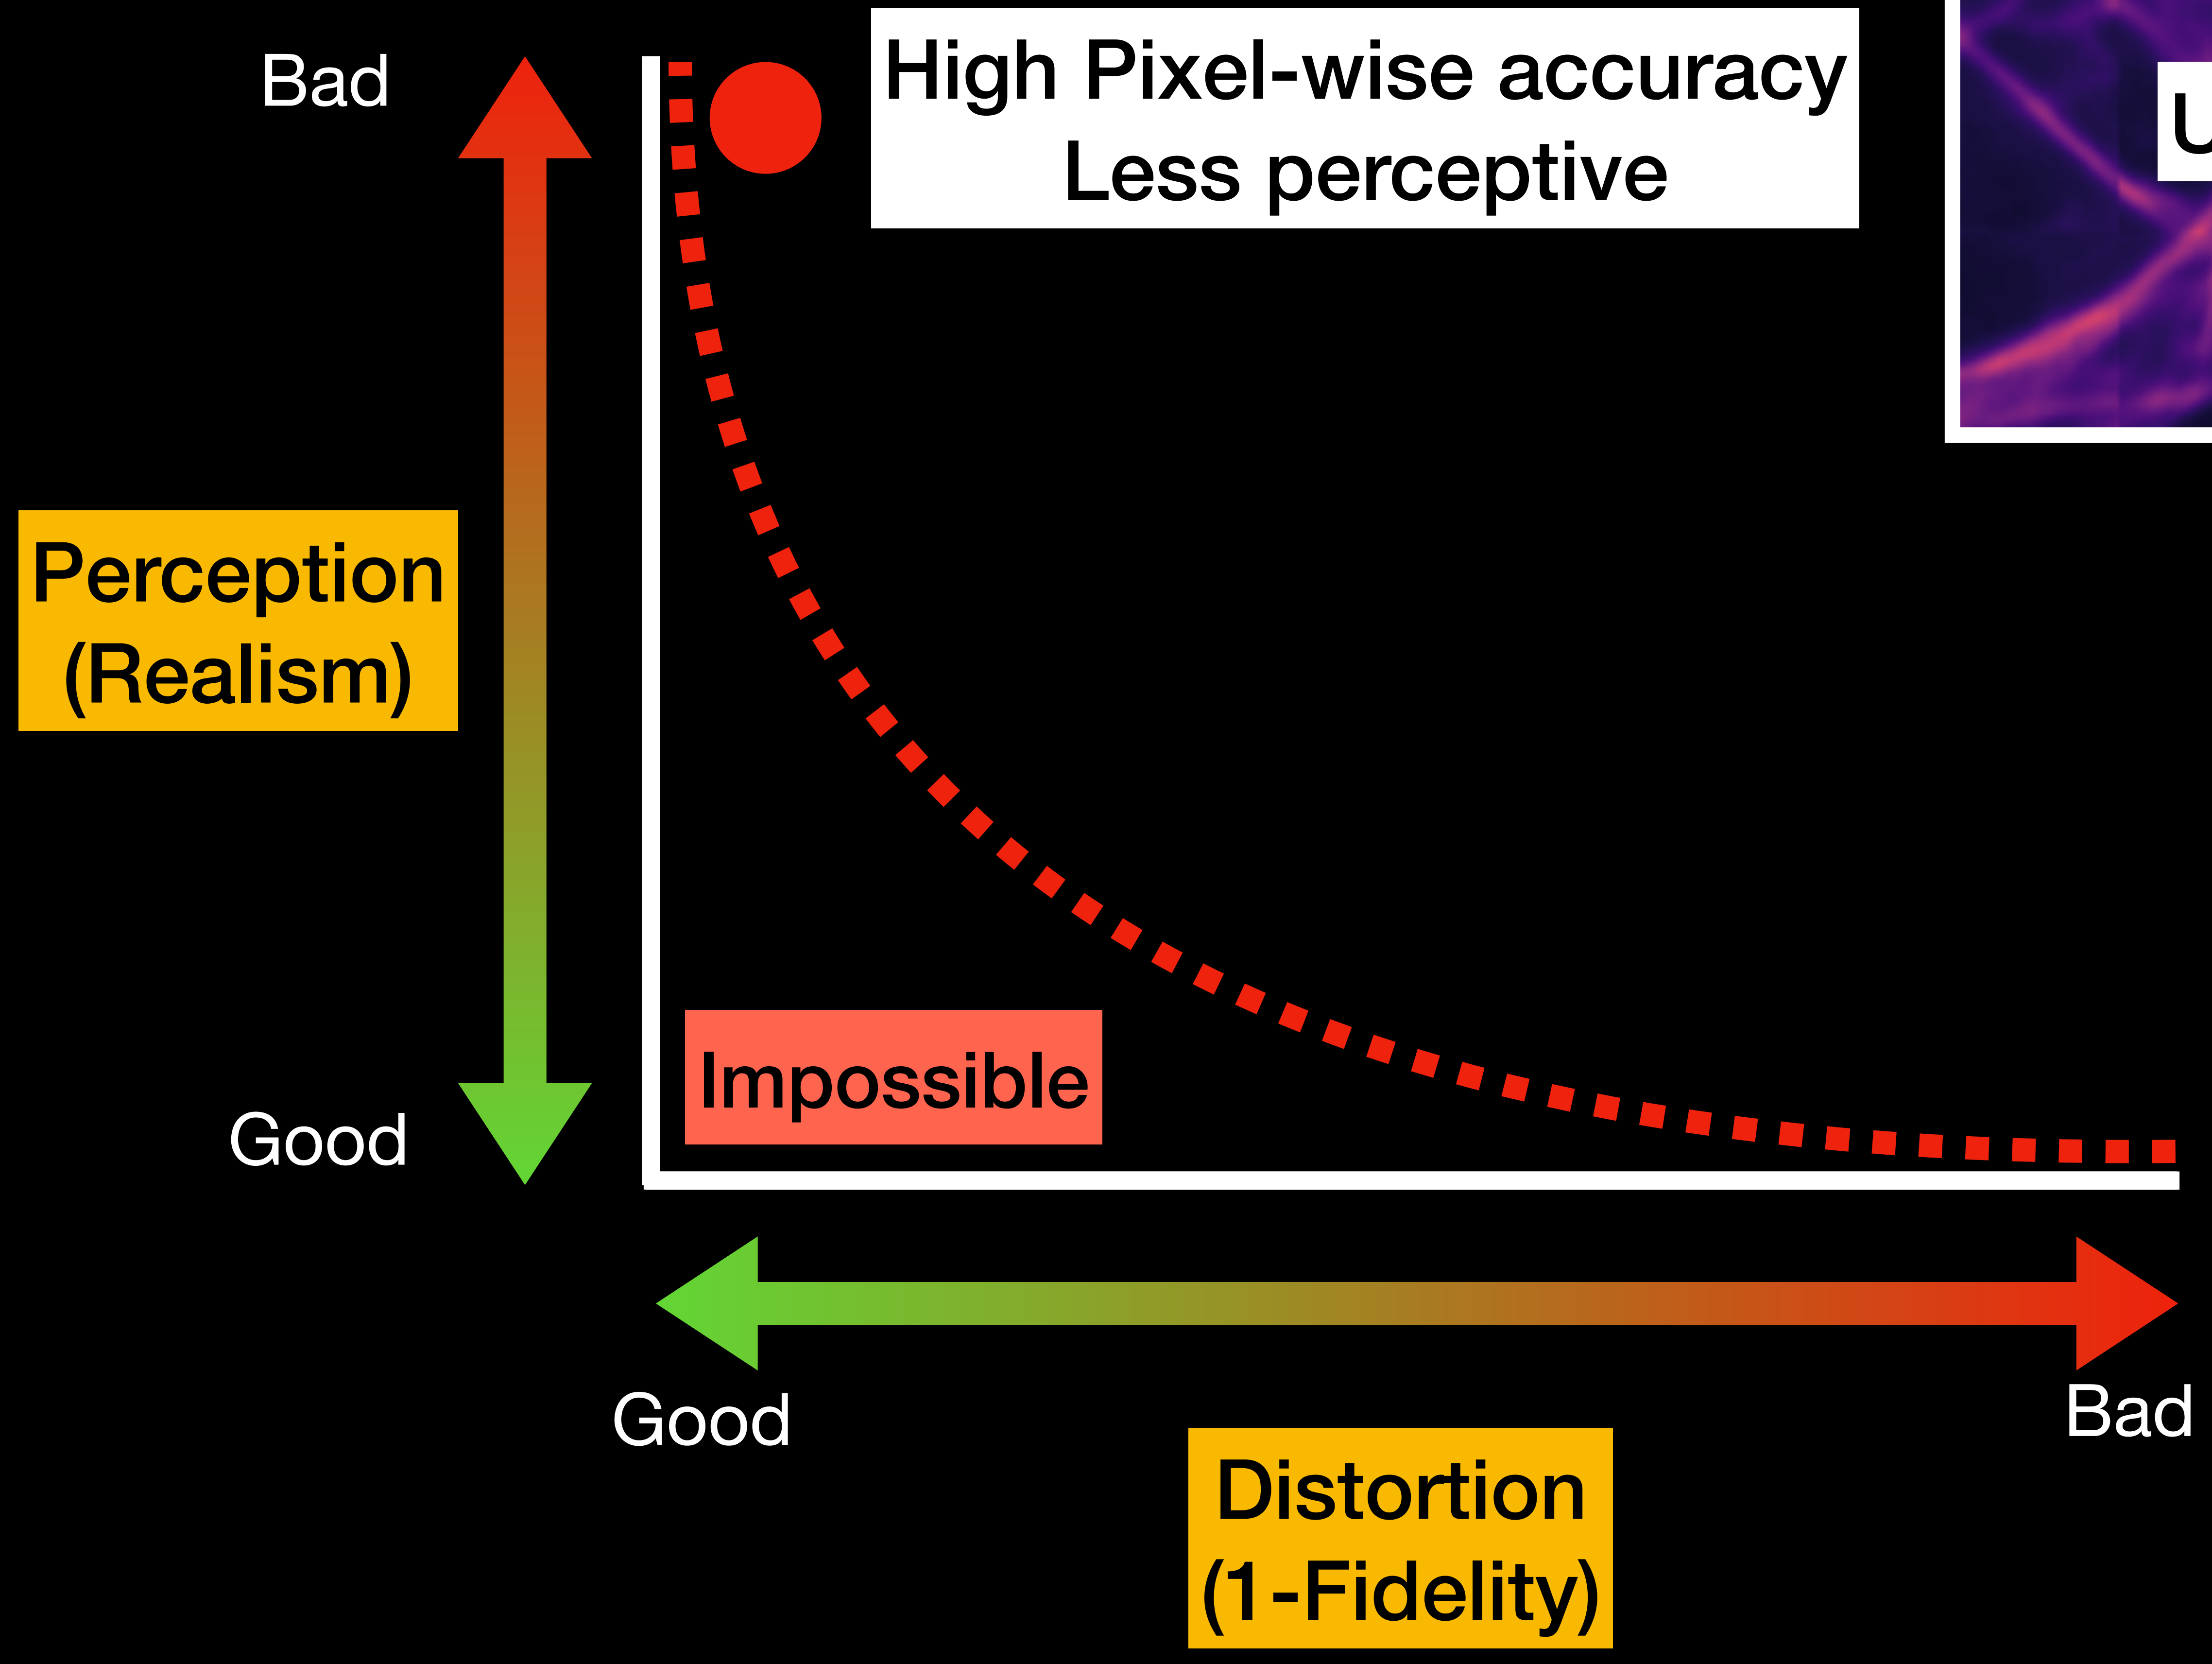

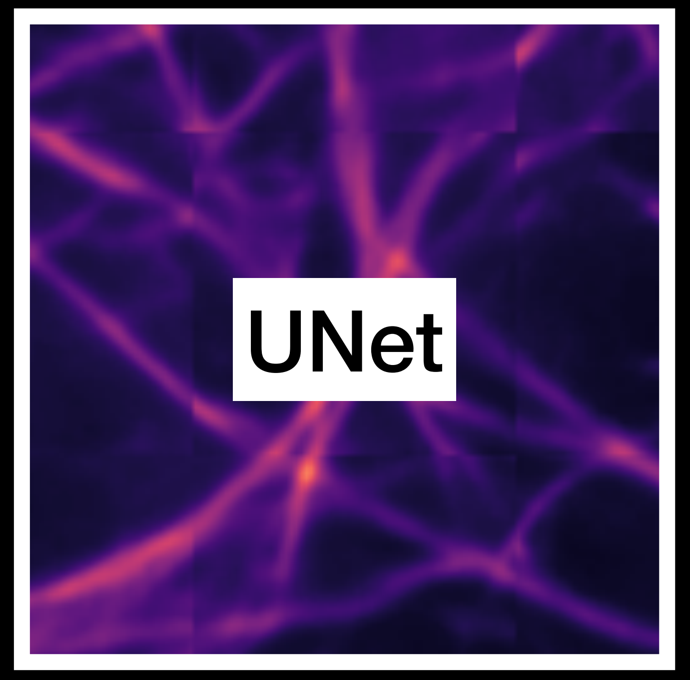

An model cannot be optimized for the *best perception* and *best distortion* at the same time!

One-shot models can be only optimized for one of these axes!

*The Perception-Distortion Tradeoff, Blau et al. - CVPR 2018

Perception Distortion Tradeoff (PDT)

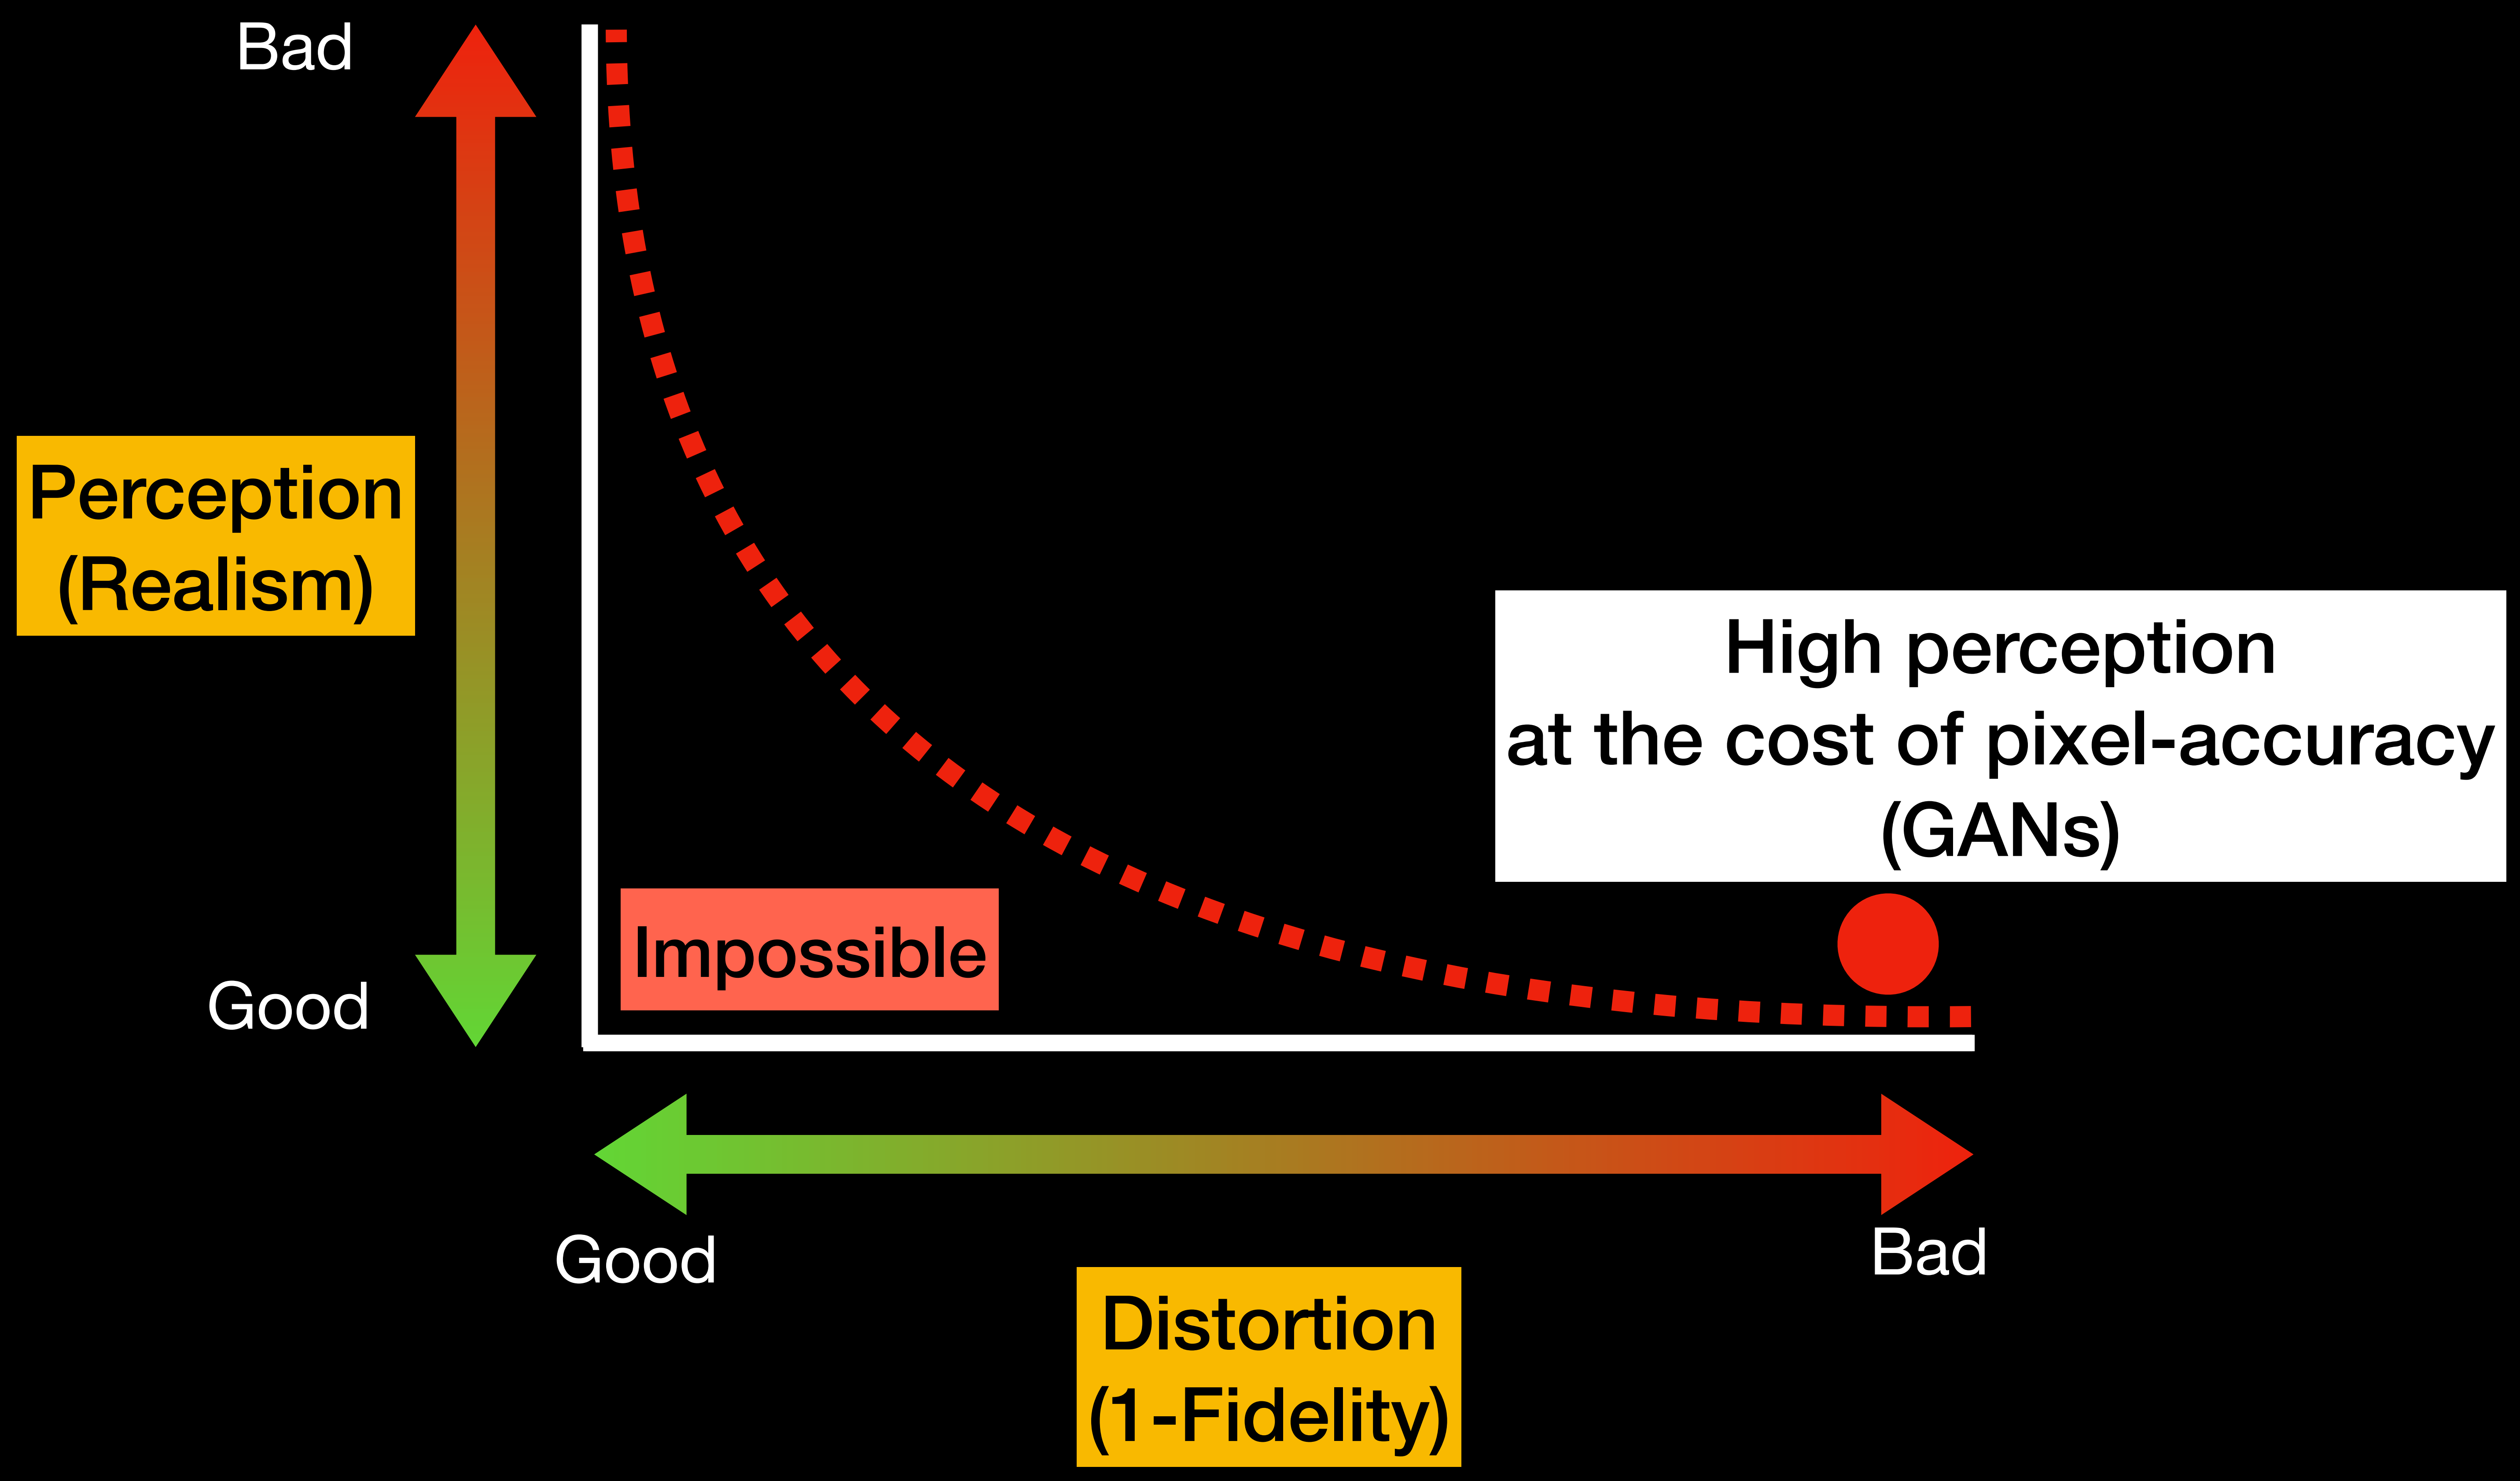

An model cannot be optimized for the *best perception* and *best distortion* at the same time!

One-shot models can be only optimized for one of these axes!

*The Perception-Distortion Tradeoff, Blau et al. - CVPR 2018

Perception Distortion Tradeoff (PDT)

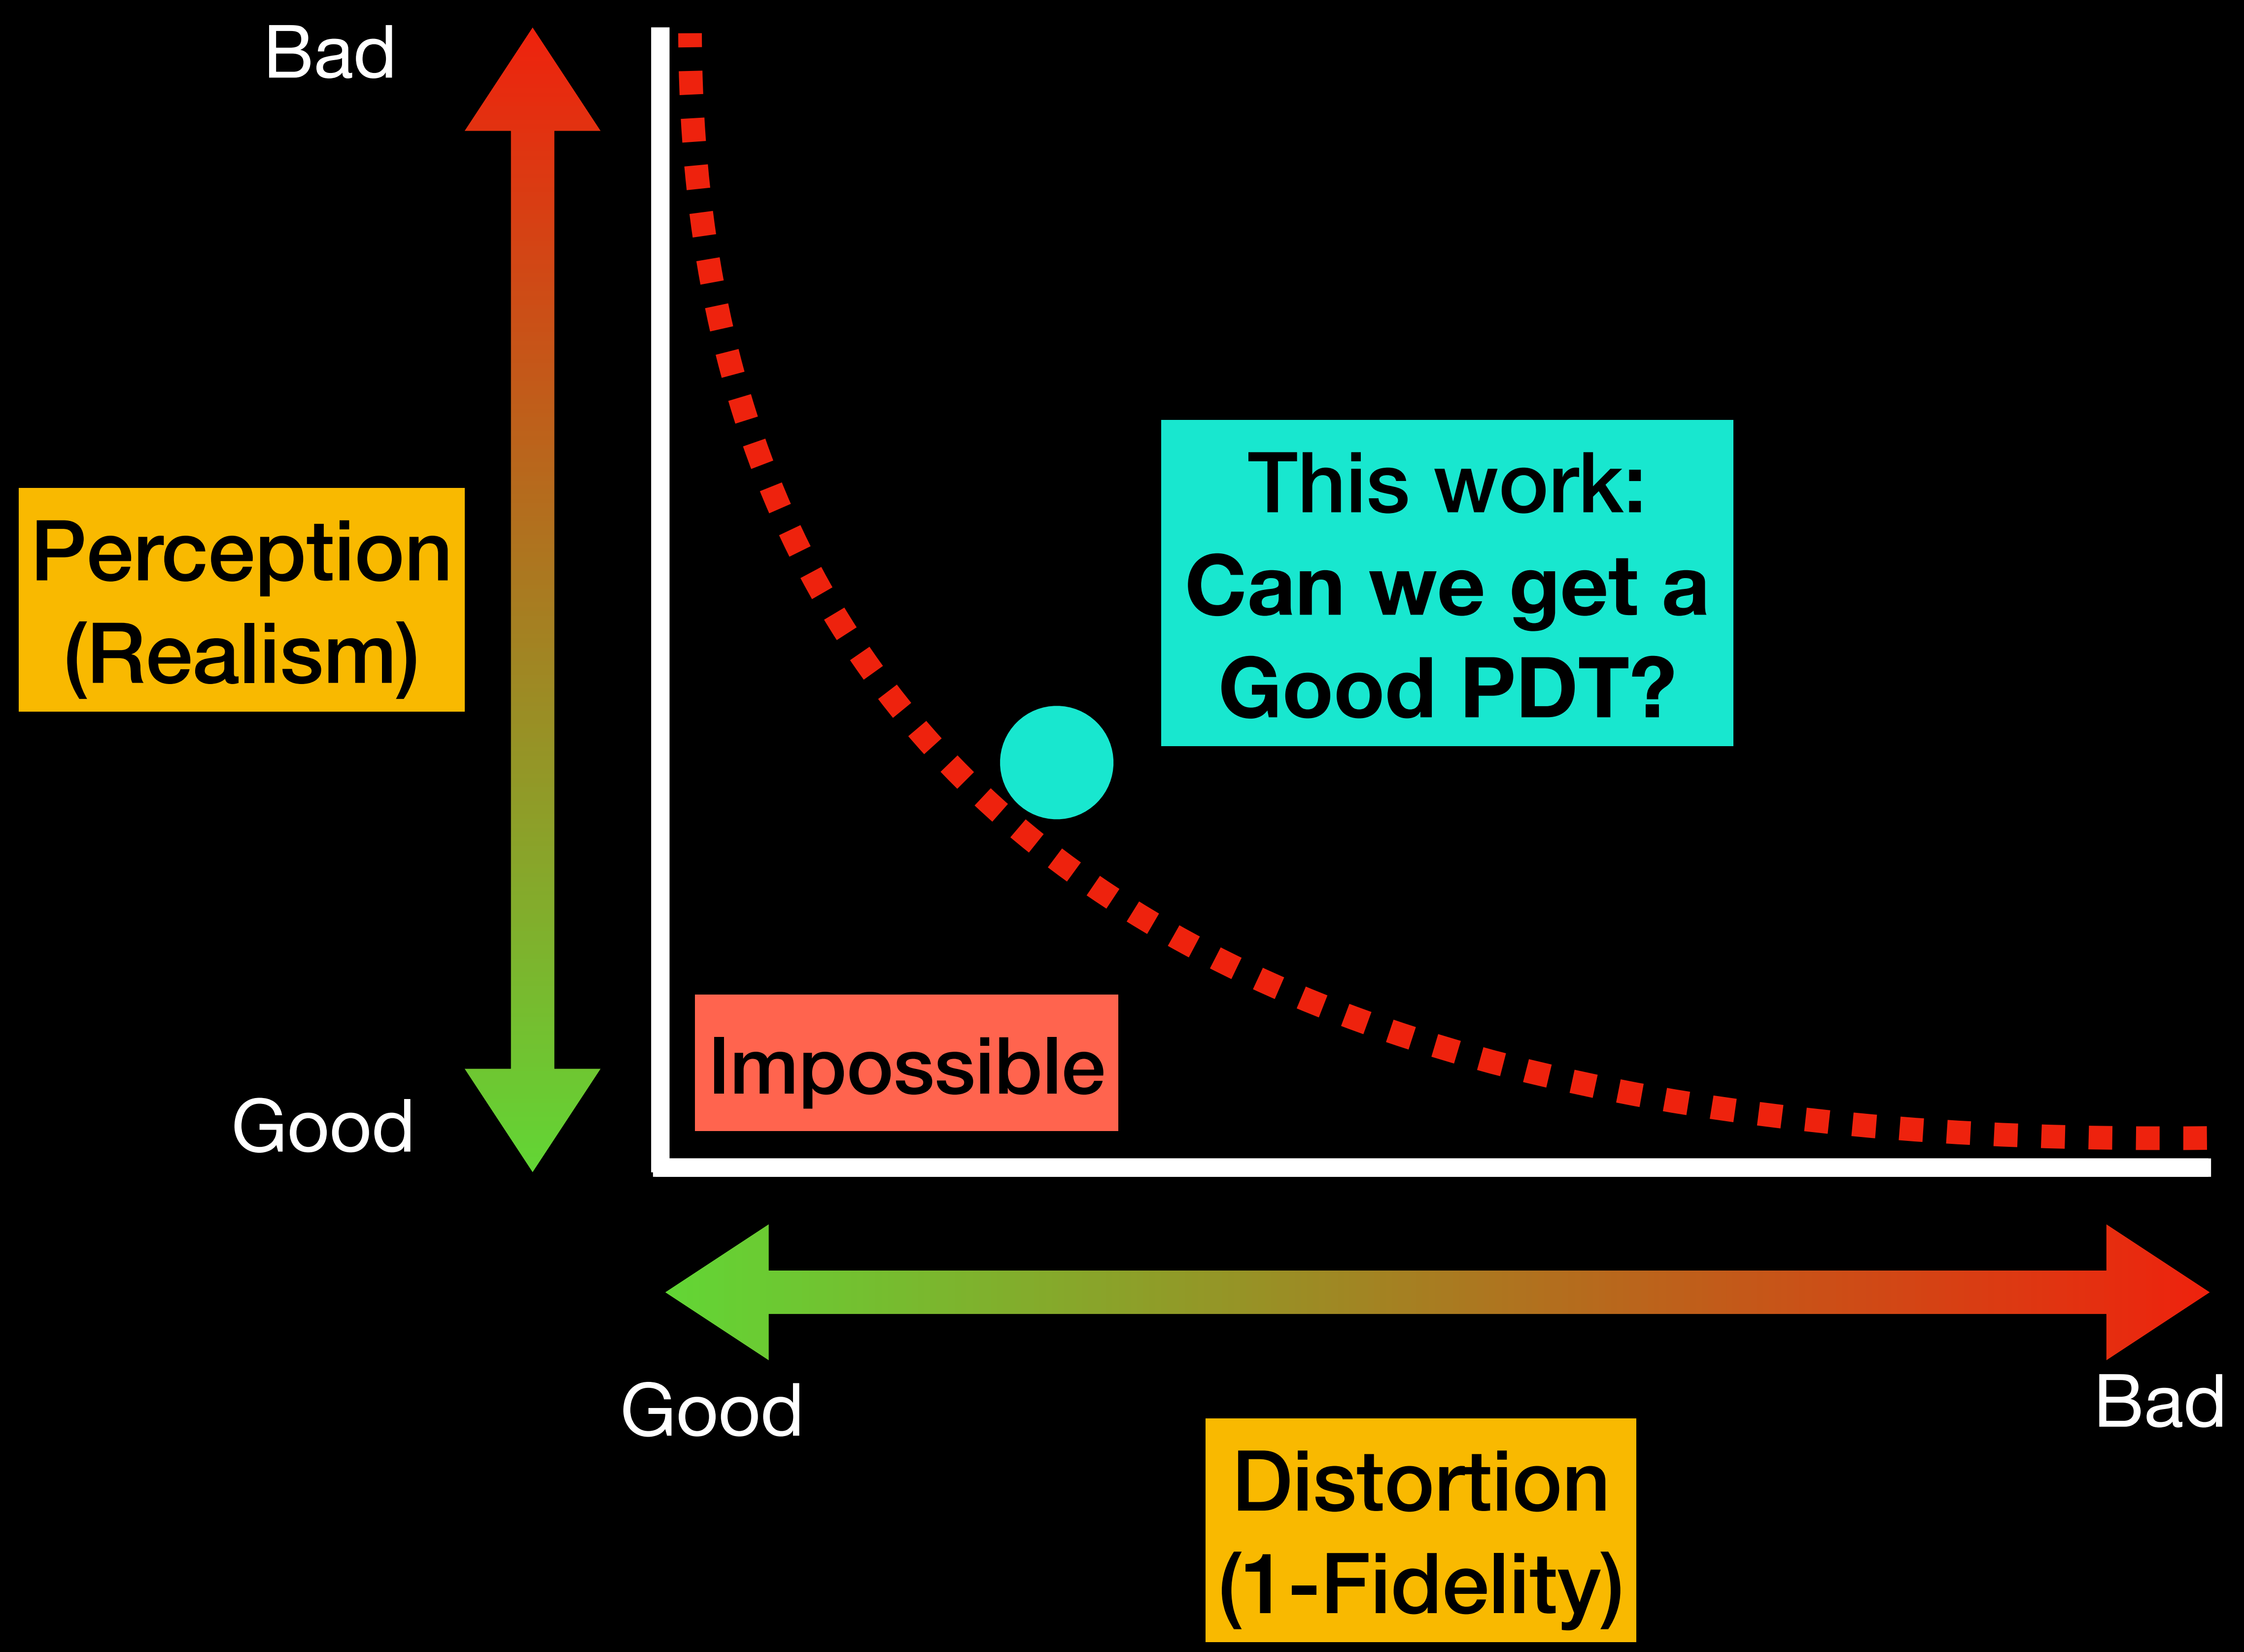

*The Perception-Distortion Tradeoff, Blau et al. - CVPR 2018

ResMatching

A guided conditional flow matching approach

ResMatching

A guided conditional flow matching approach

- **Based on HazeMatching*:** Best PDT in Microscopy Image Dehazing
- **Iterative:** plausible samples, instead of blurry MMSE only
- **Posterior Sampling:** Access to MMSE, model calibration
- **ODE-based Generative Model:** Fast! (Fewer NFEs compared to Diffusion)

HazeMatching, Ray et al. - CVPR 2026 (Findings)

ResMatching - Training

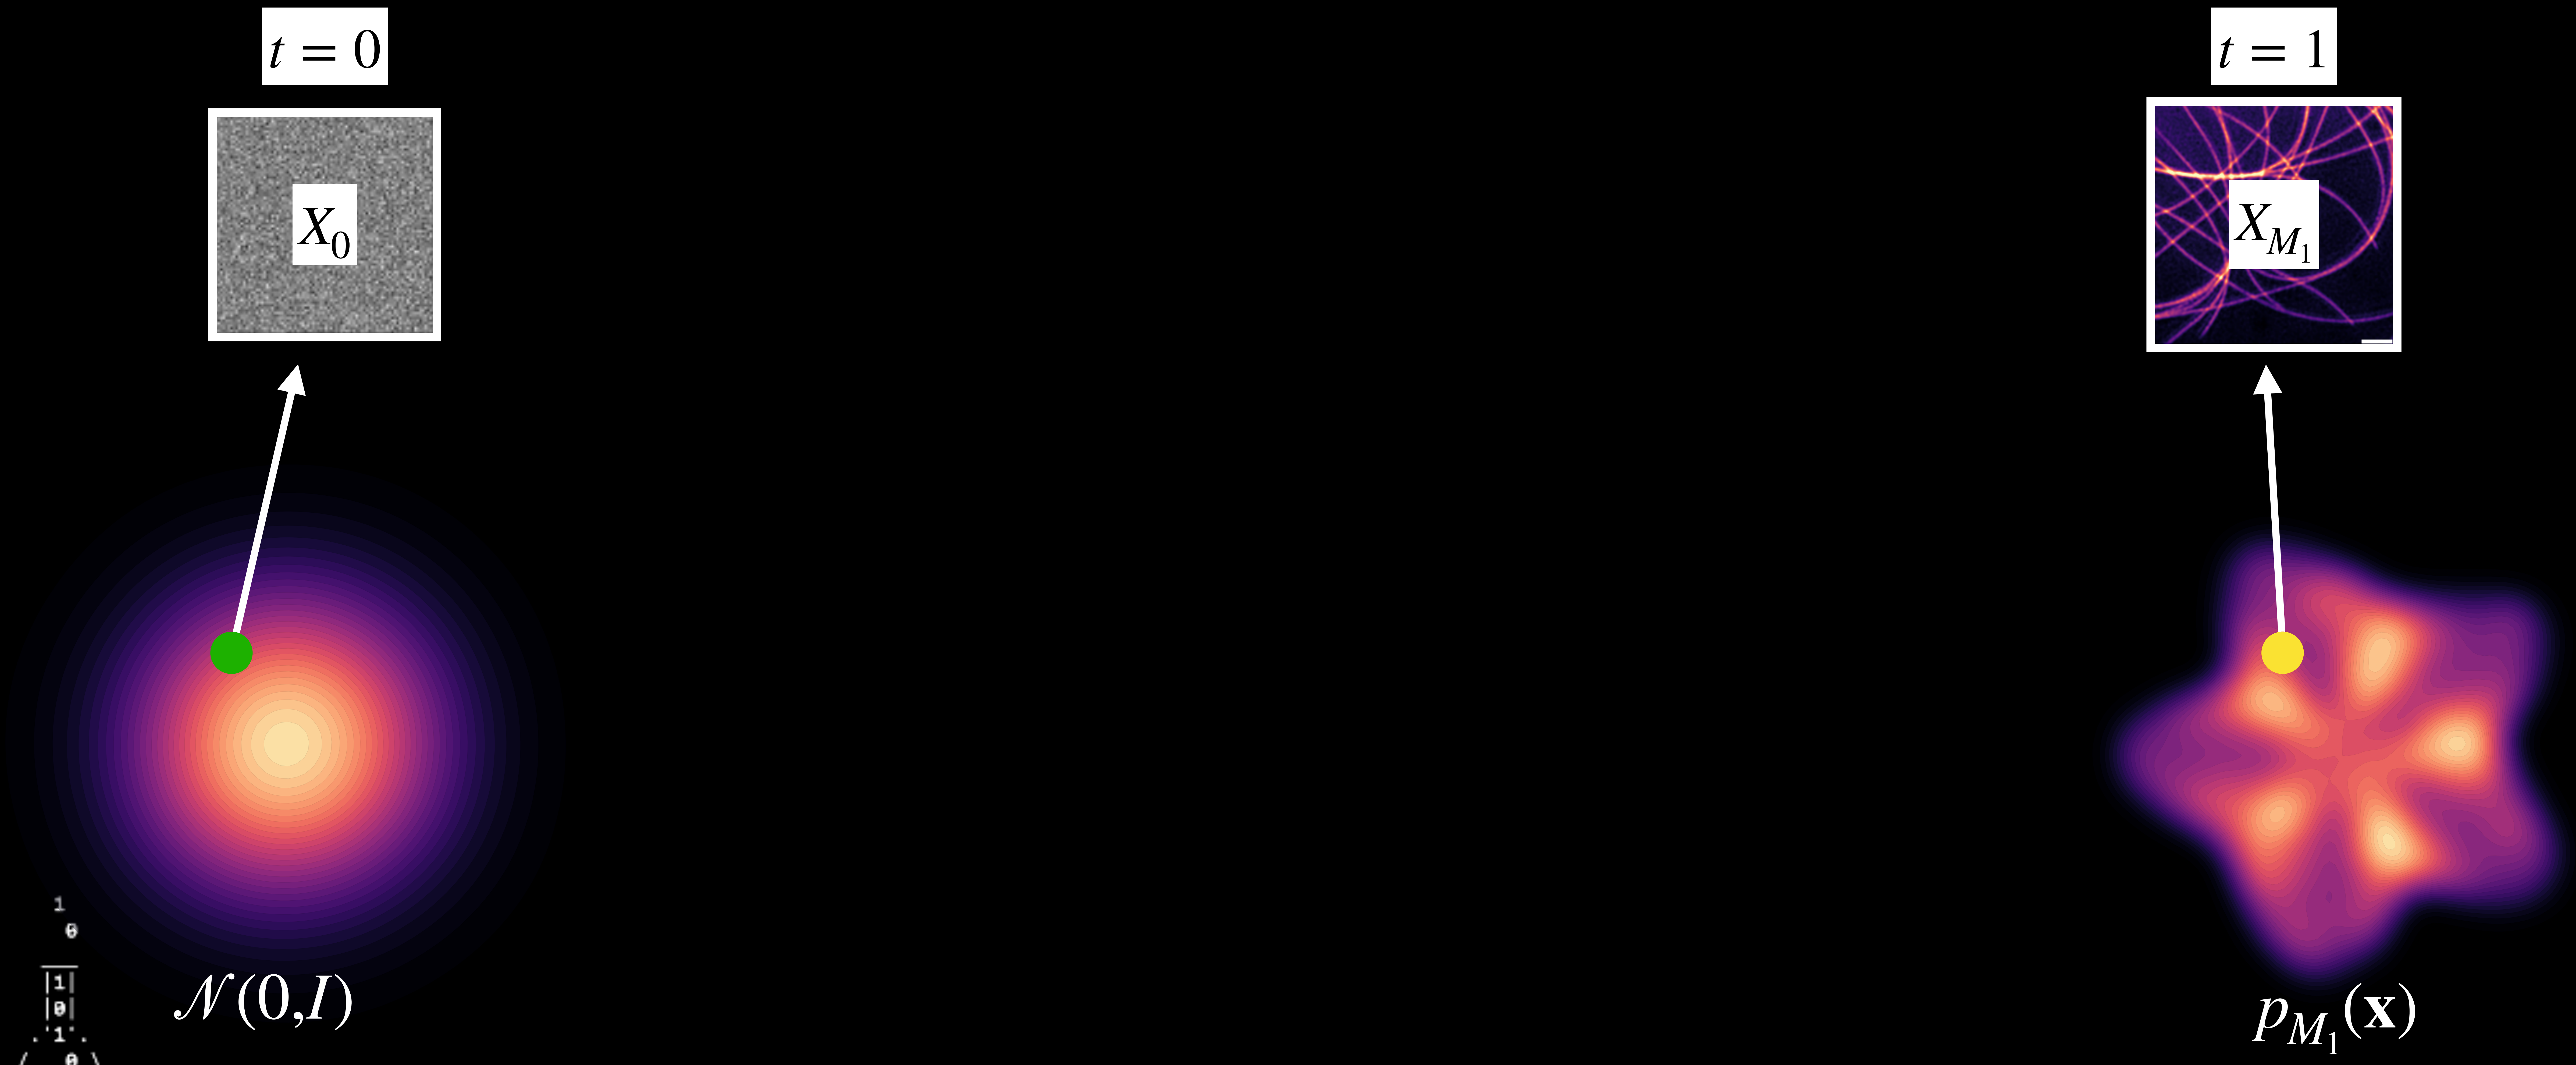

ResMatching - Training

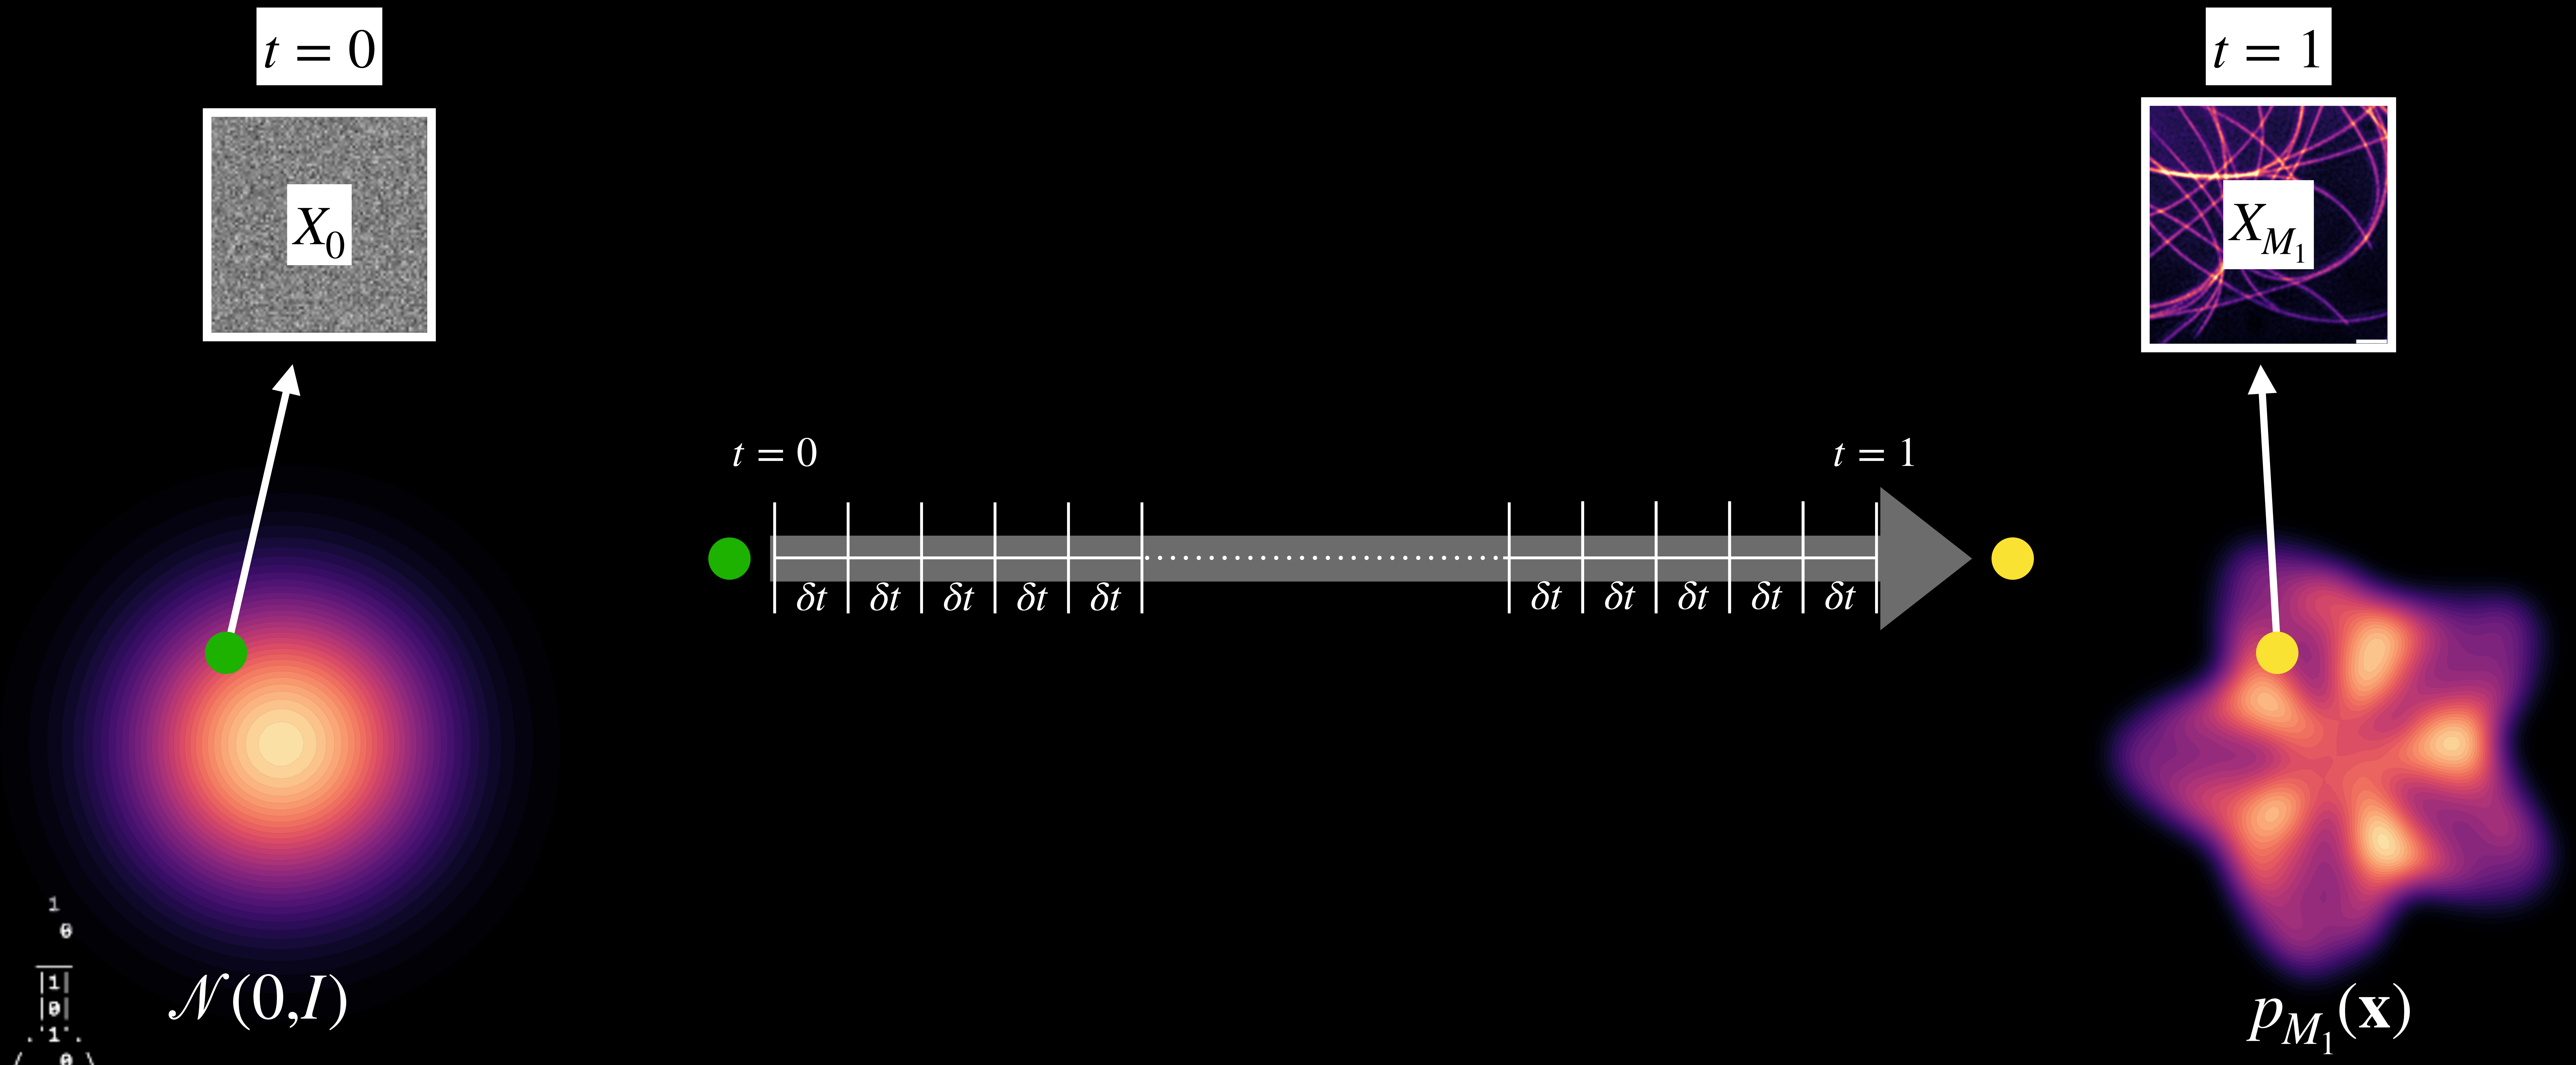

ResMatching - Training

ResMatching - Training

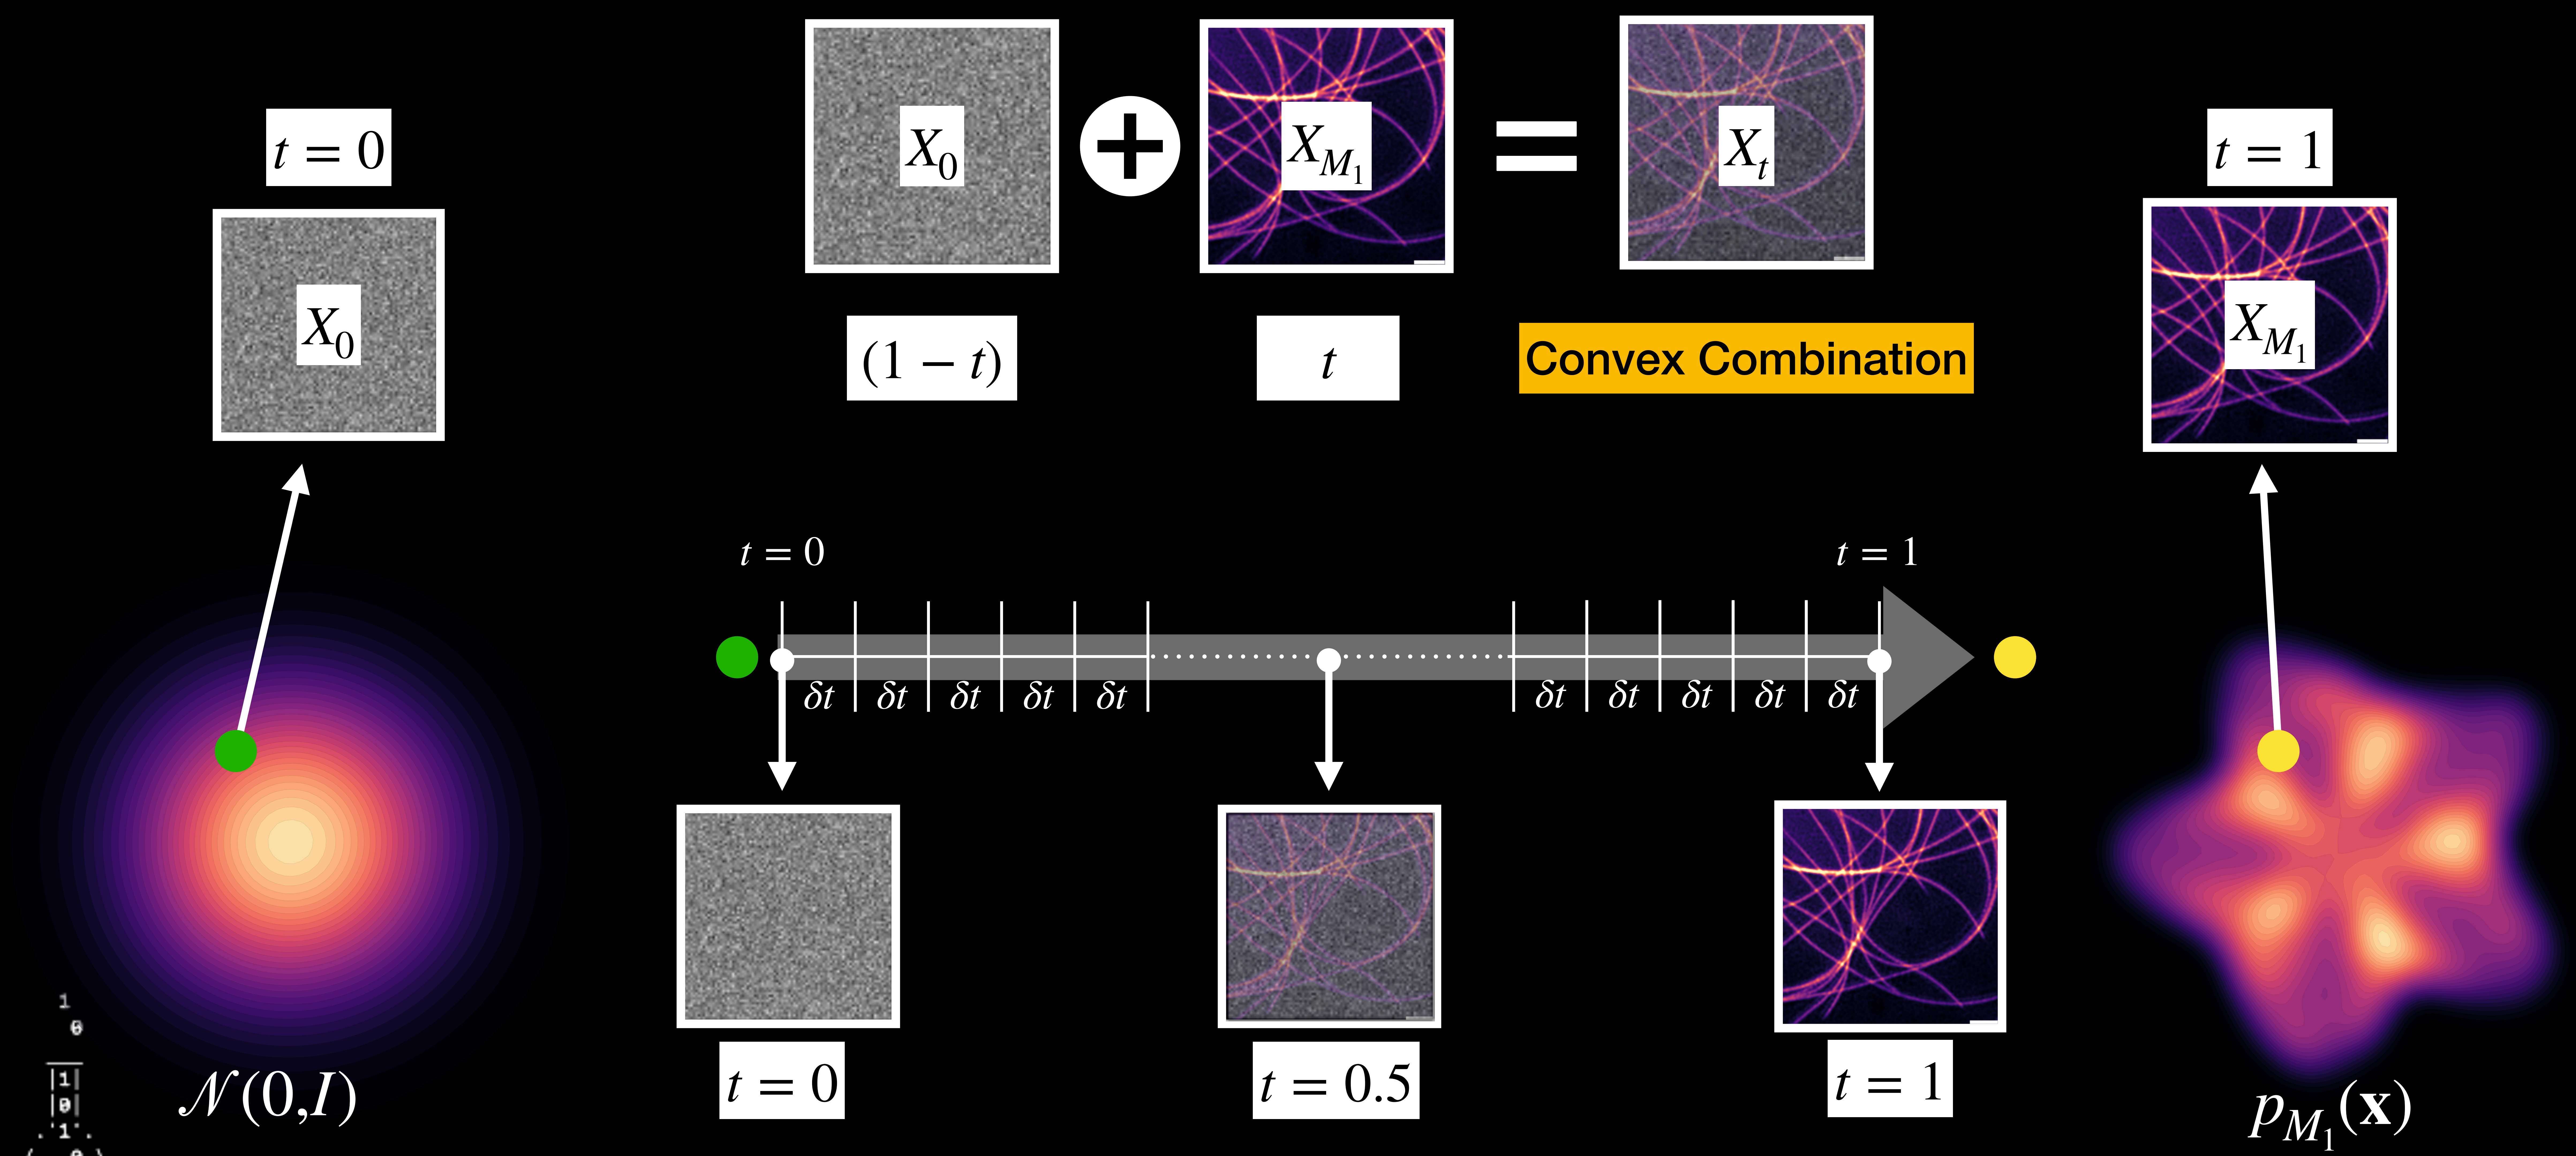

ResMatching - Training

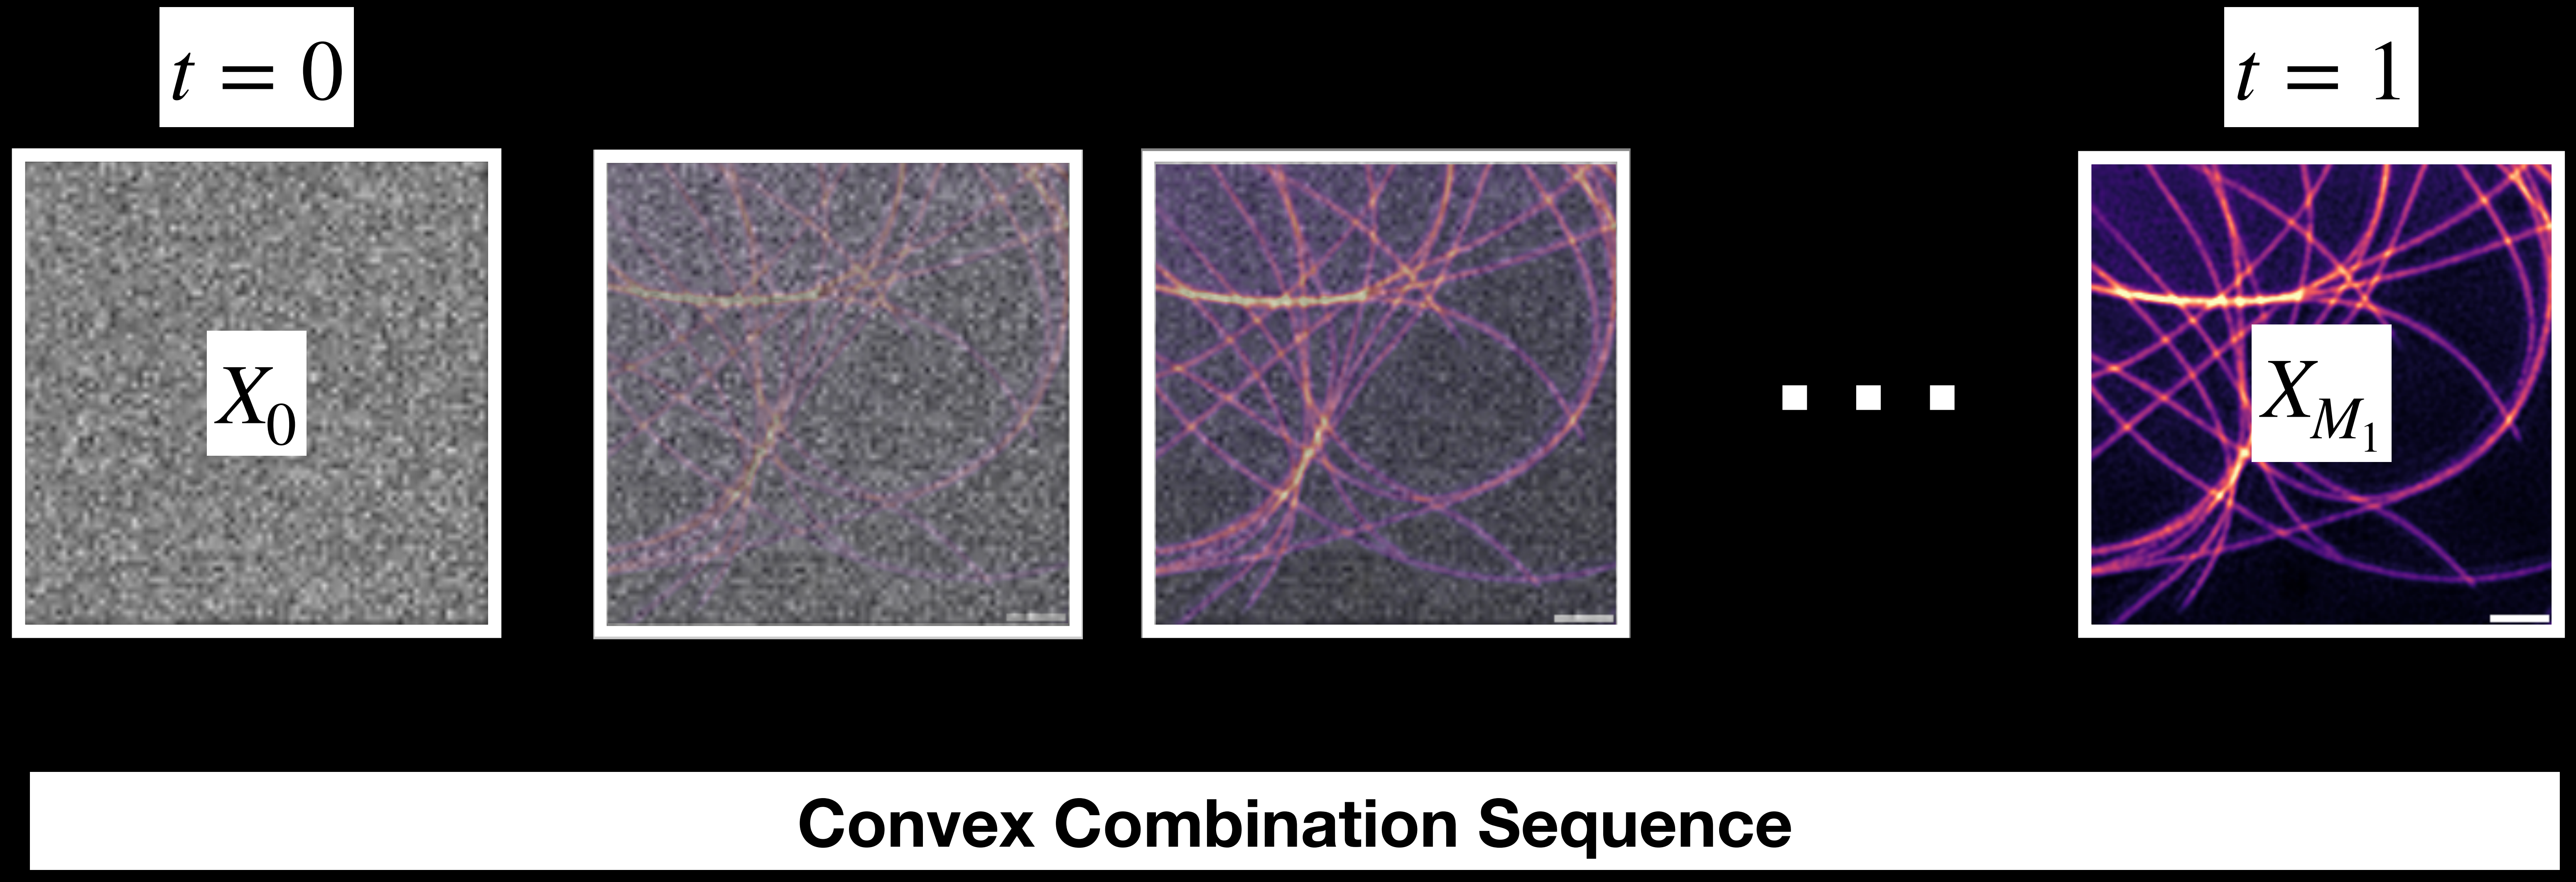

ResMatching - Training

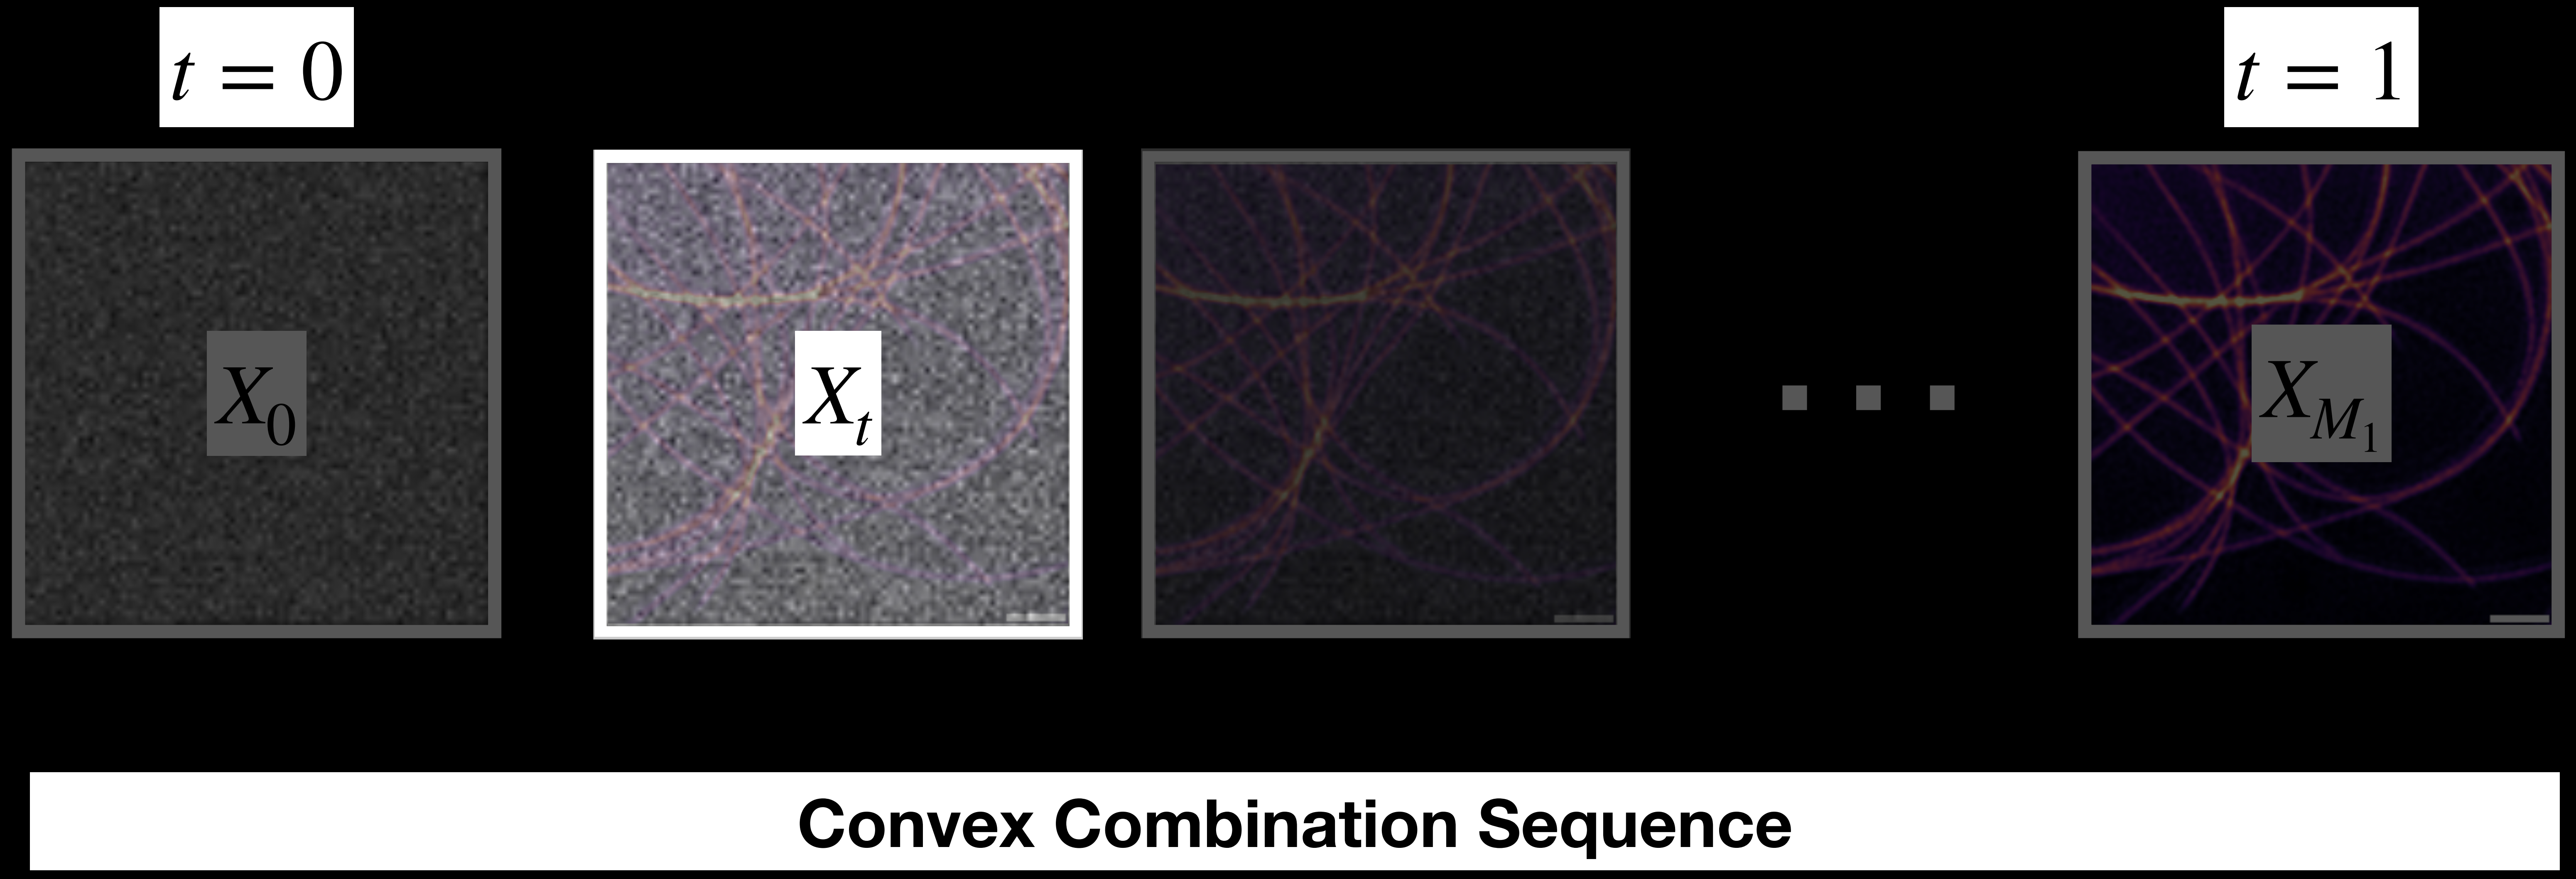

ResMatching - Training

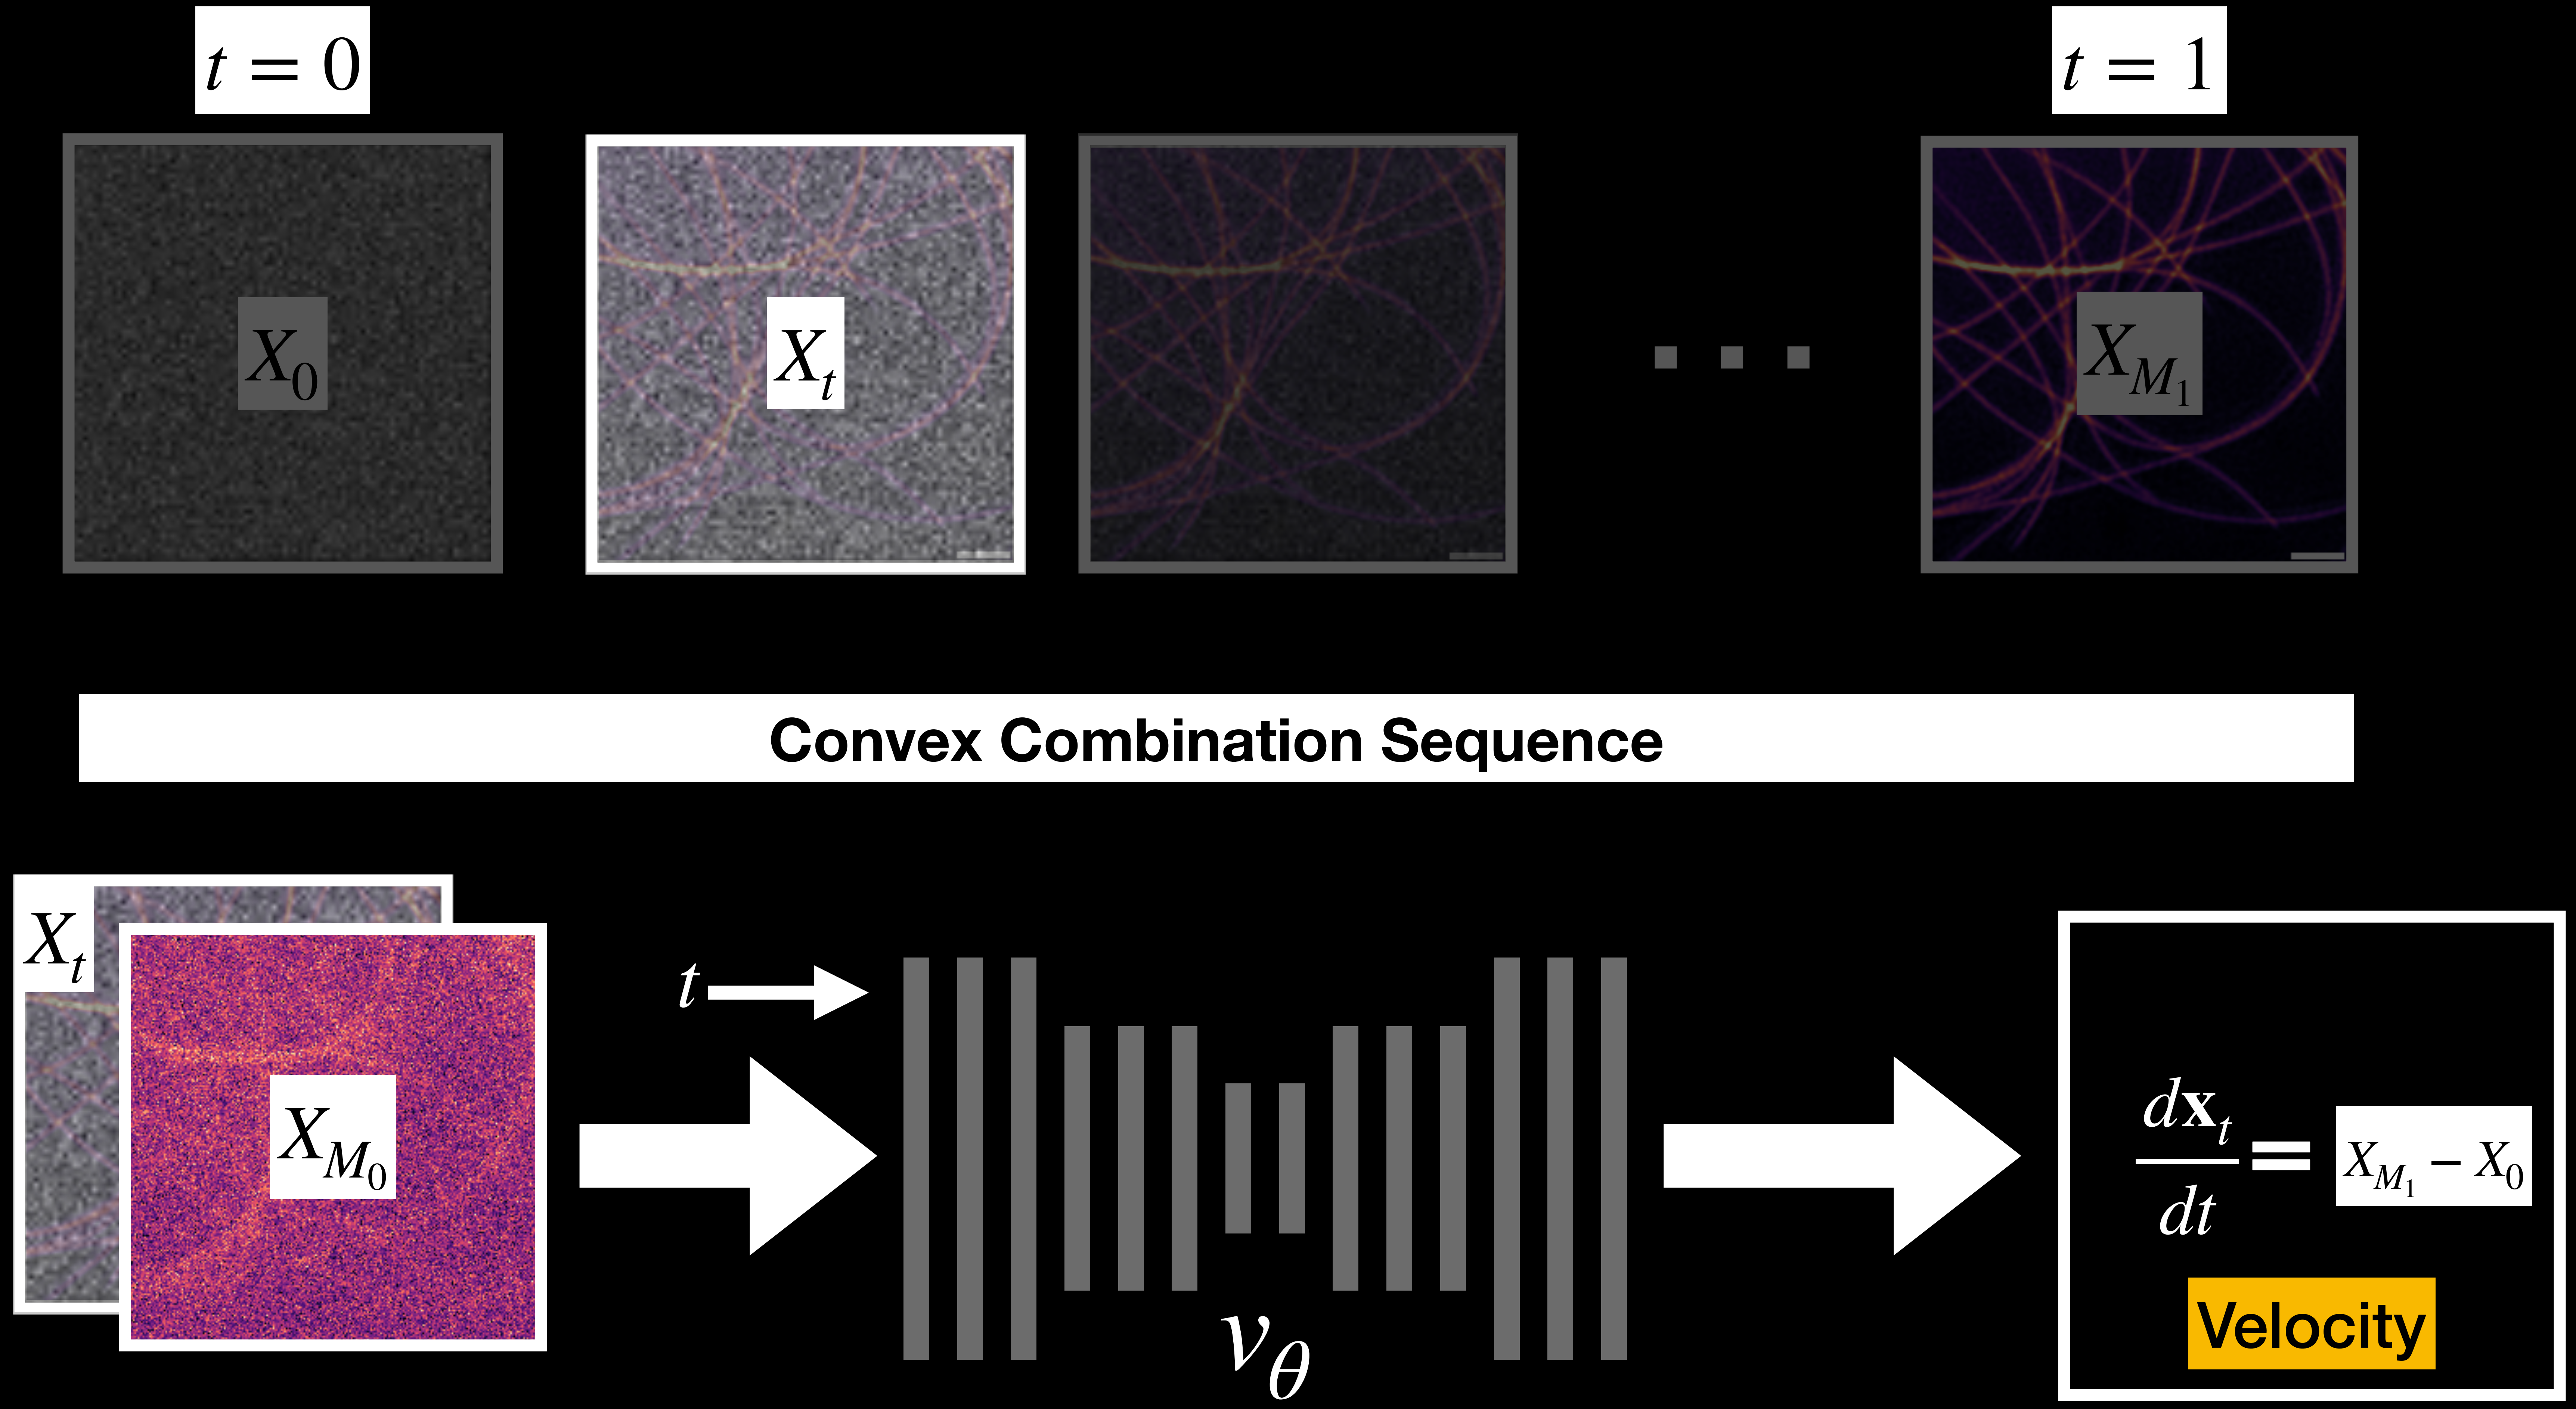

Condition on Degraded Observations X_{M_0}

Flow Model

$$\min_\theta \mathbb{E}\,||\,v_\theta(t, \mathbf{x}_t, \mathbf{x}_{M_0}) - (\mathbf{x}_{M_1} - \mathbf{x}_0)\,||^2$$

ResMatching - *Iterative* Inference

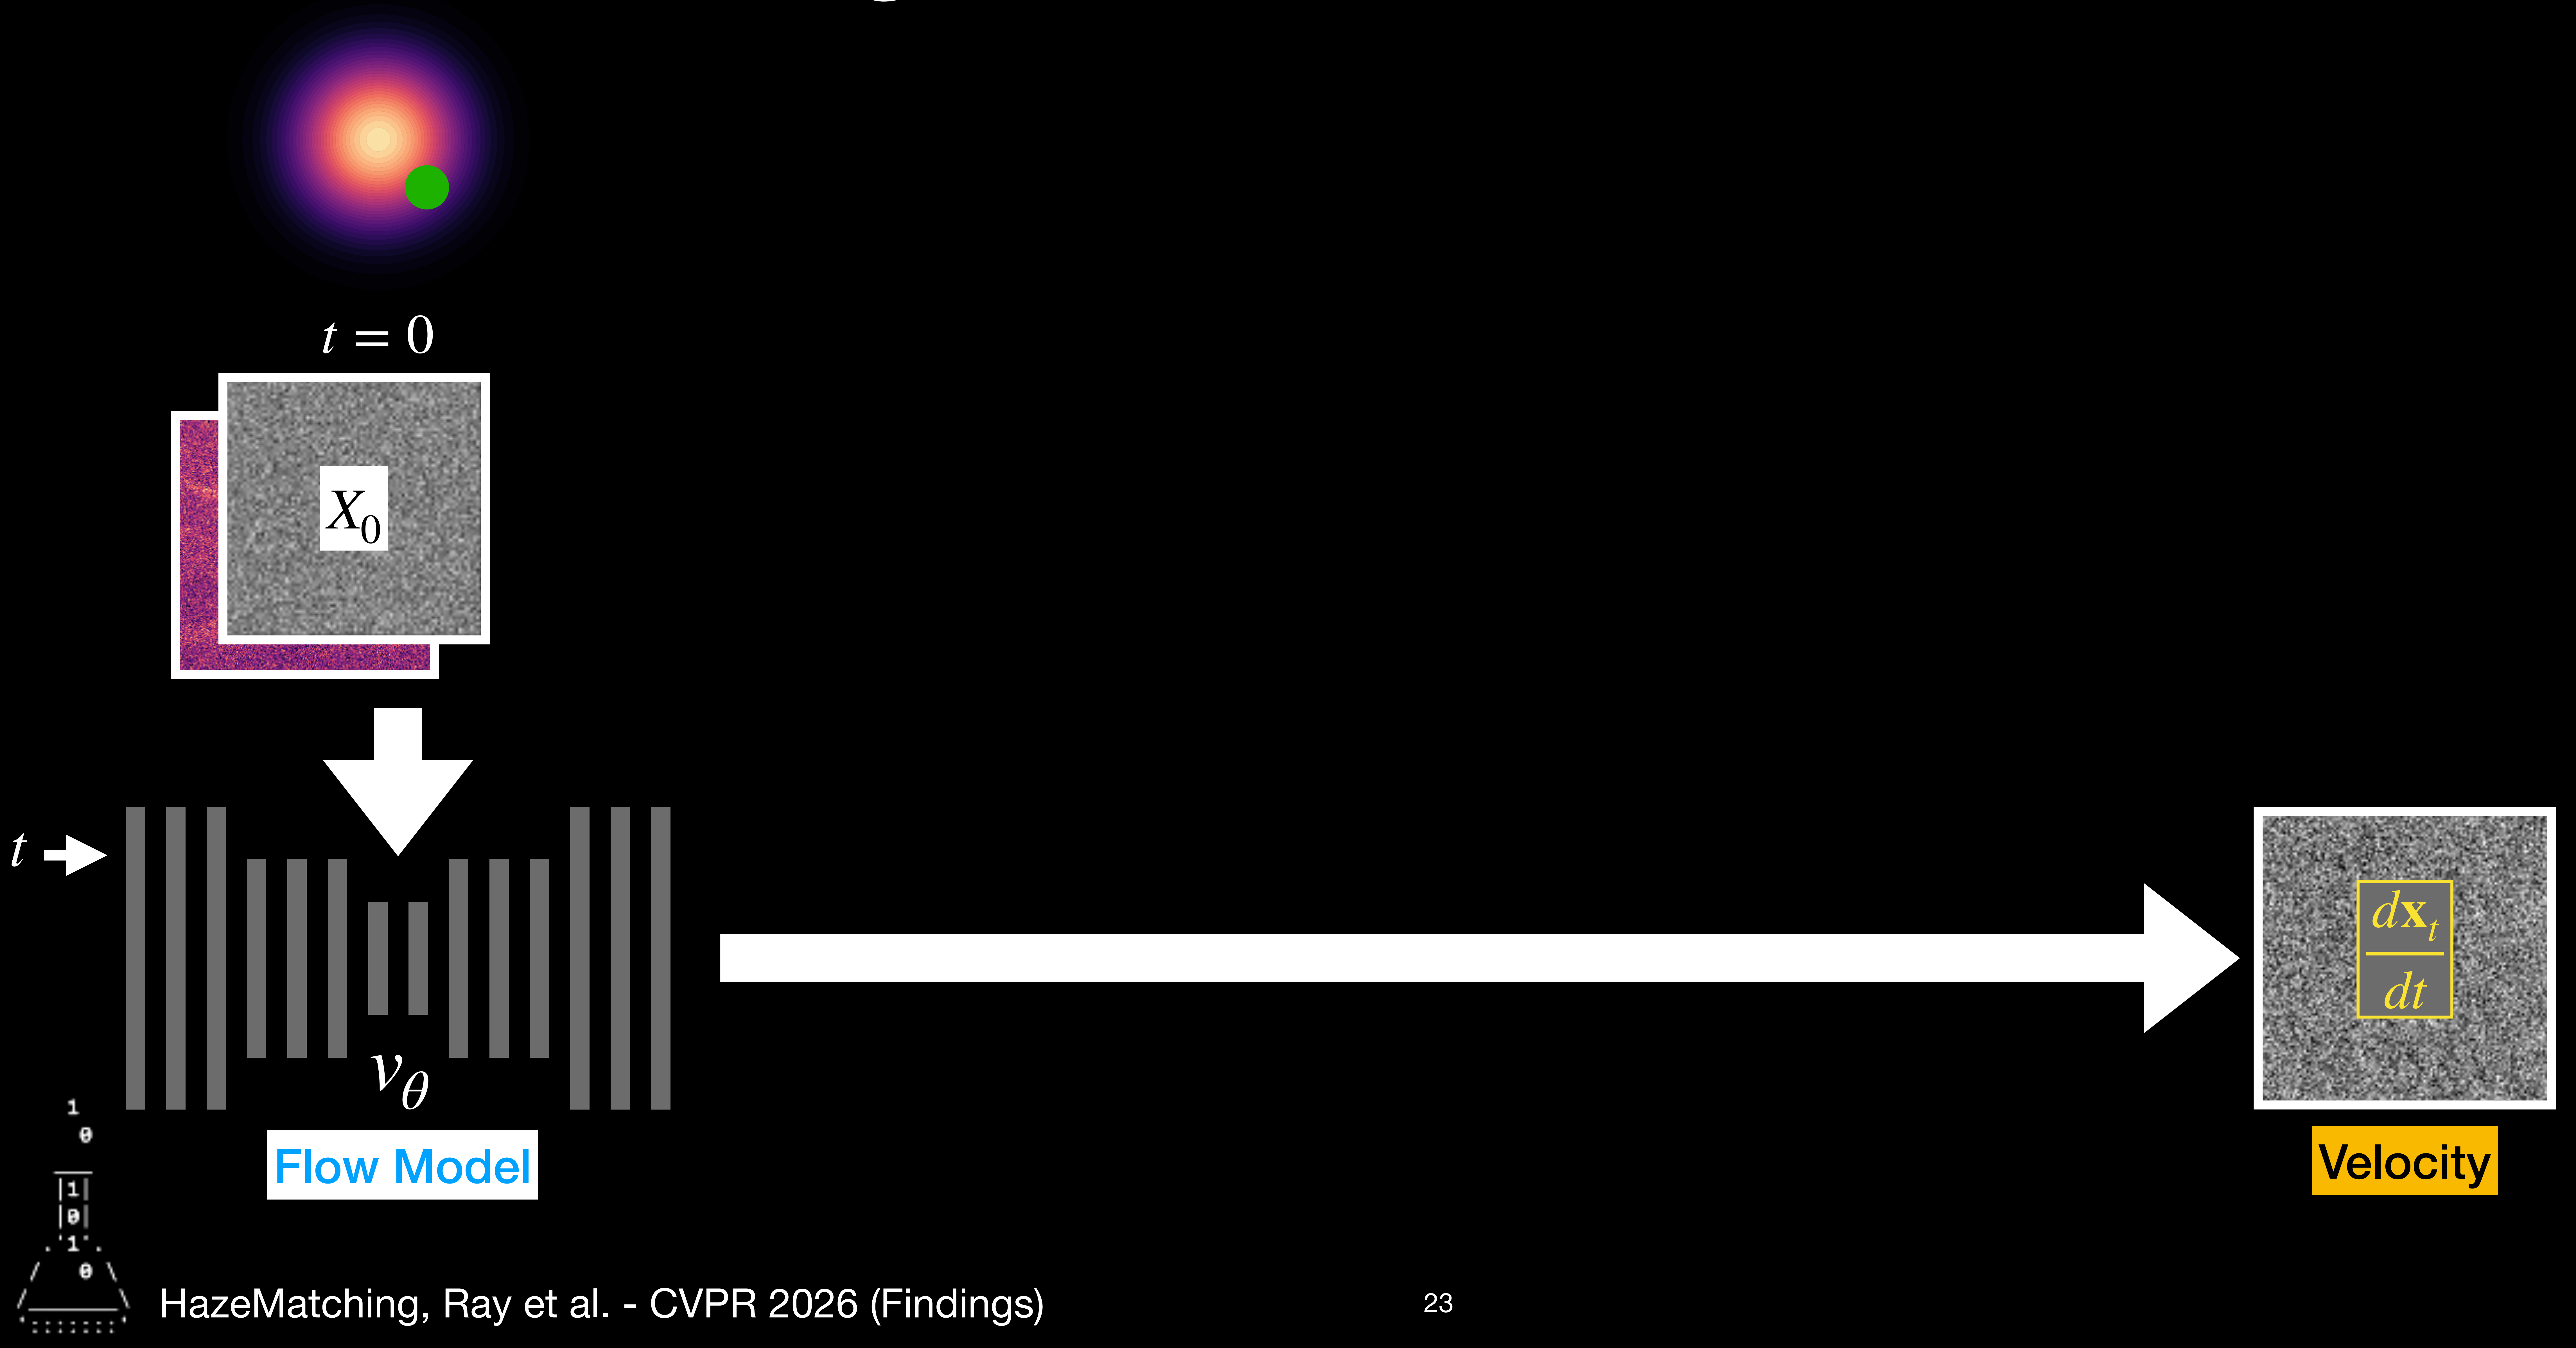

HazeMatching, Ray et al. - CVPR 2026 (Findings)

ResMatching - *Iterative* Inference

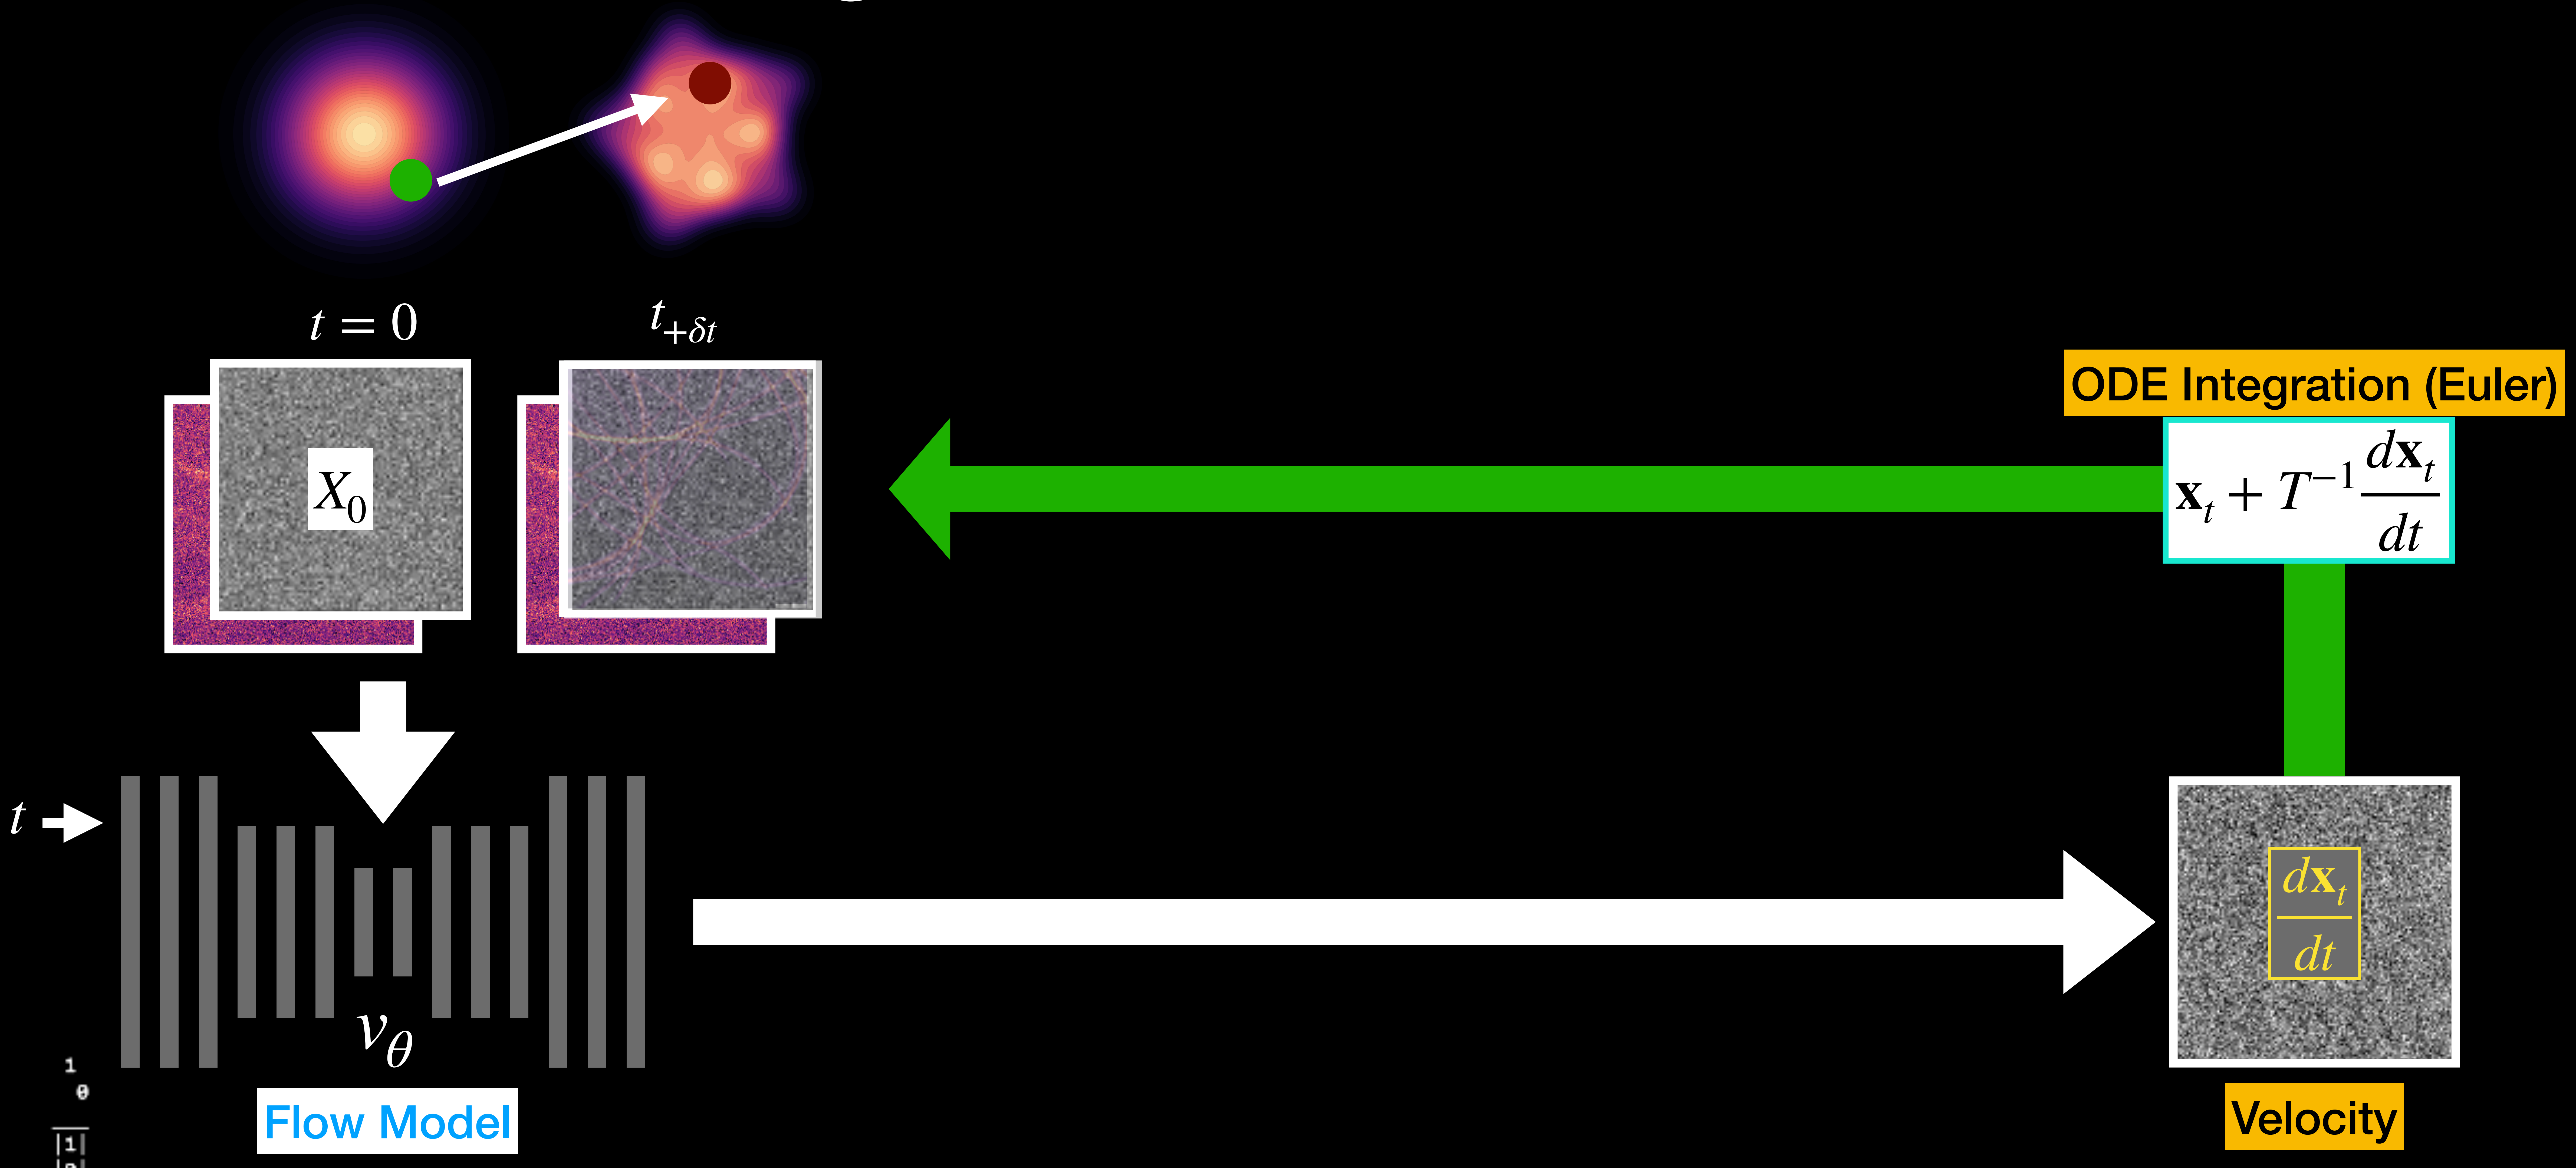

ResMatching - *Iterative* Inference

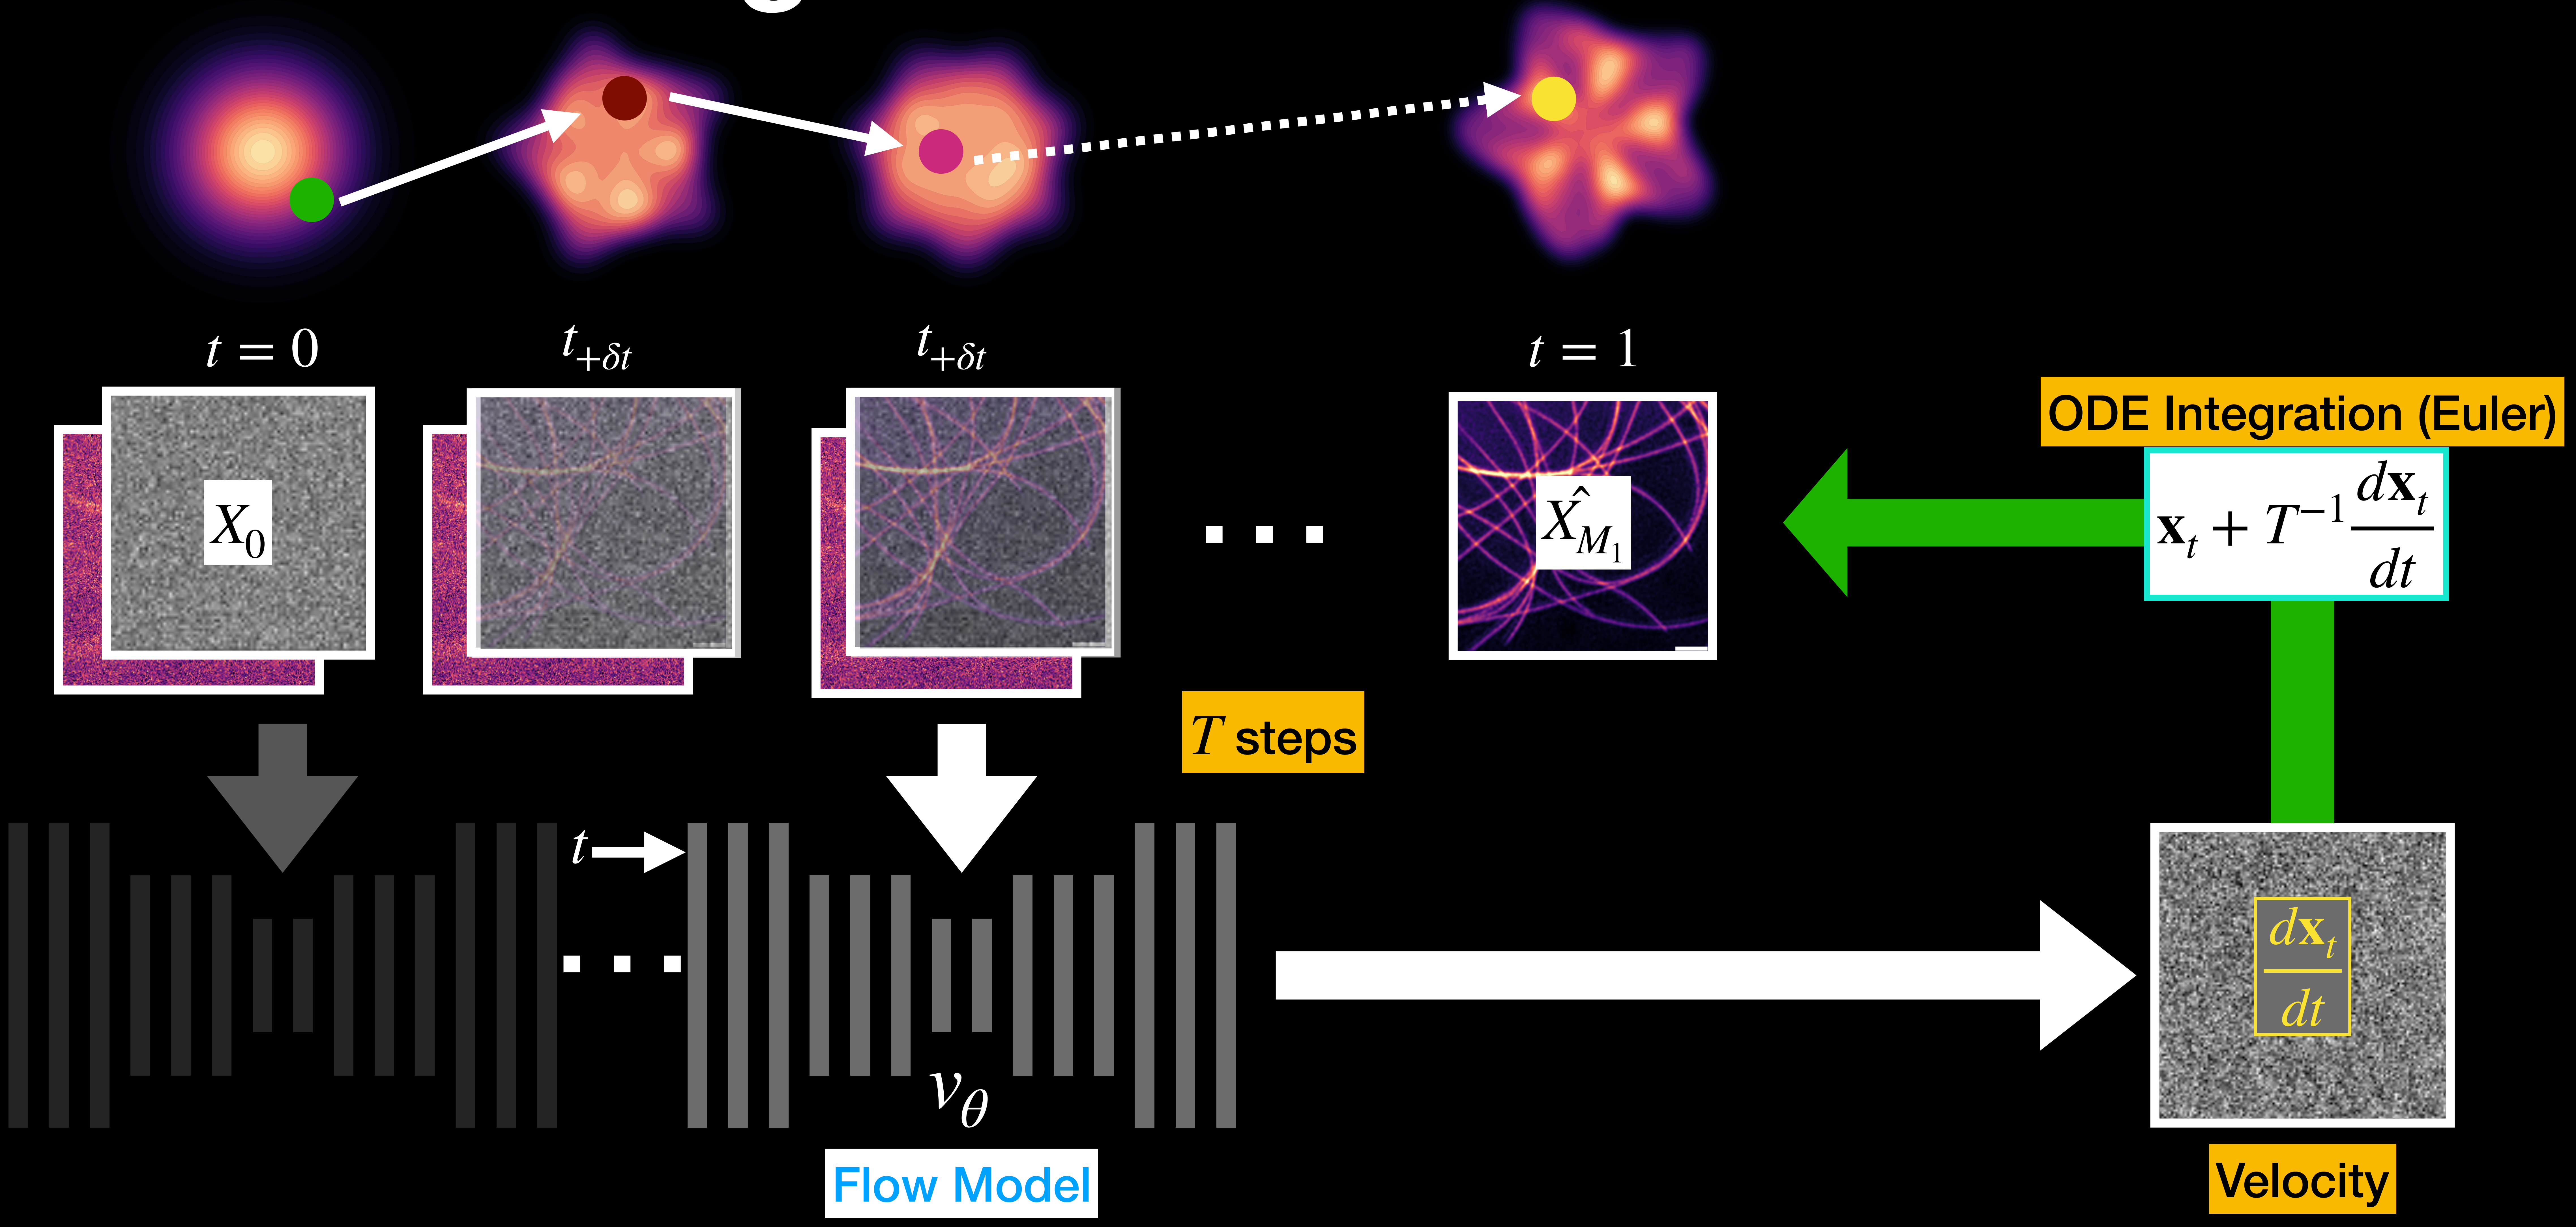

ResMatching - *Iterative* Inference

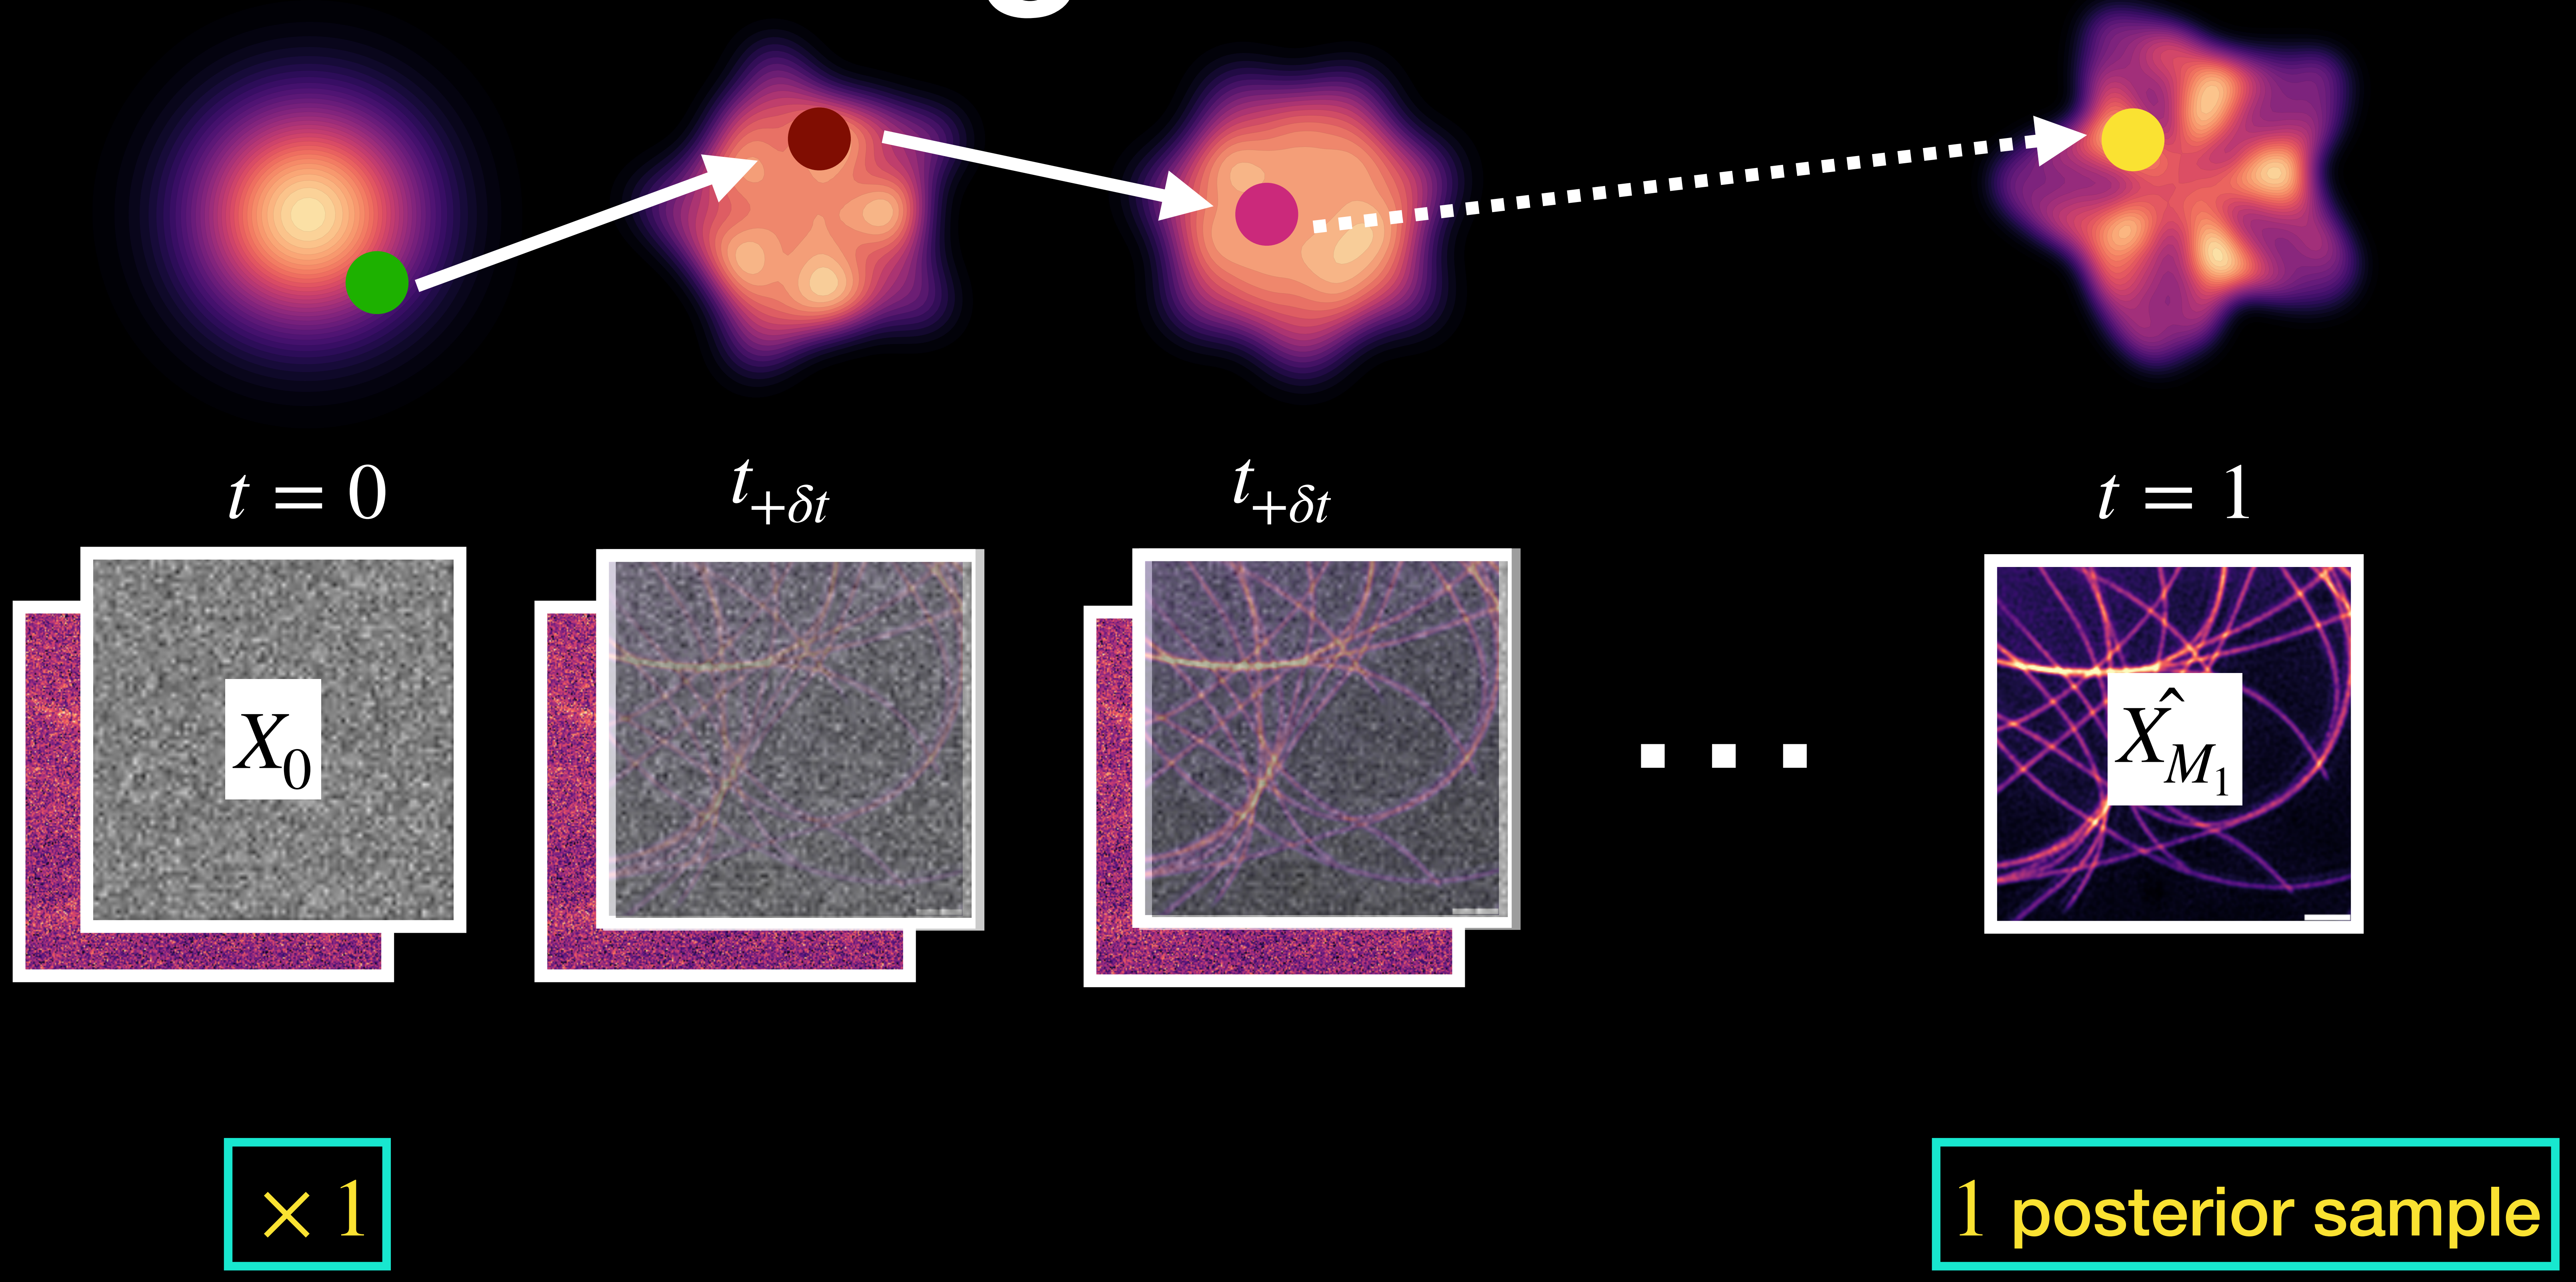

ResMatching - Posterior Sampling

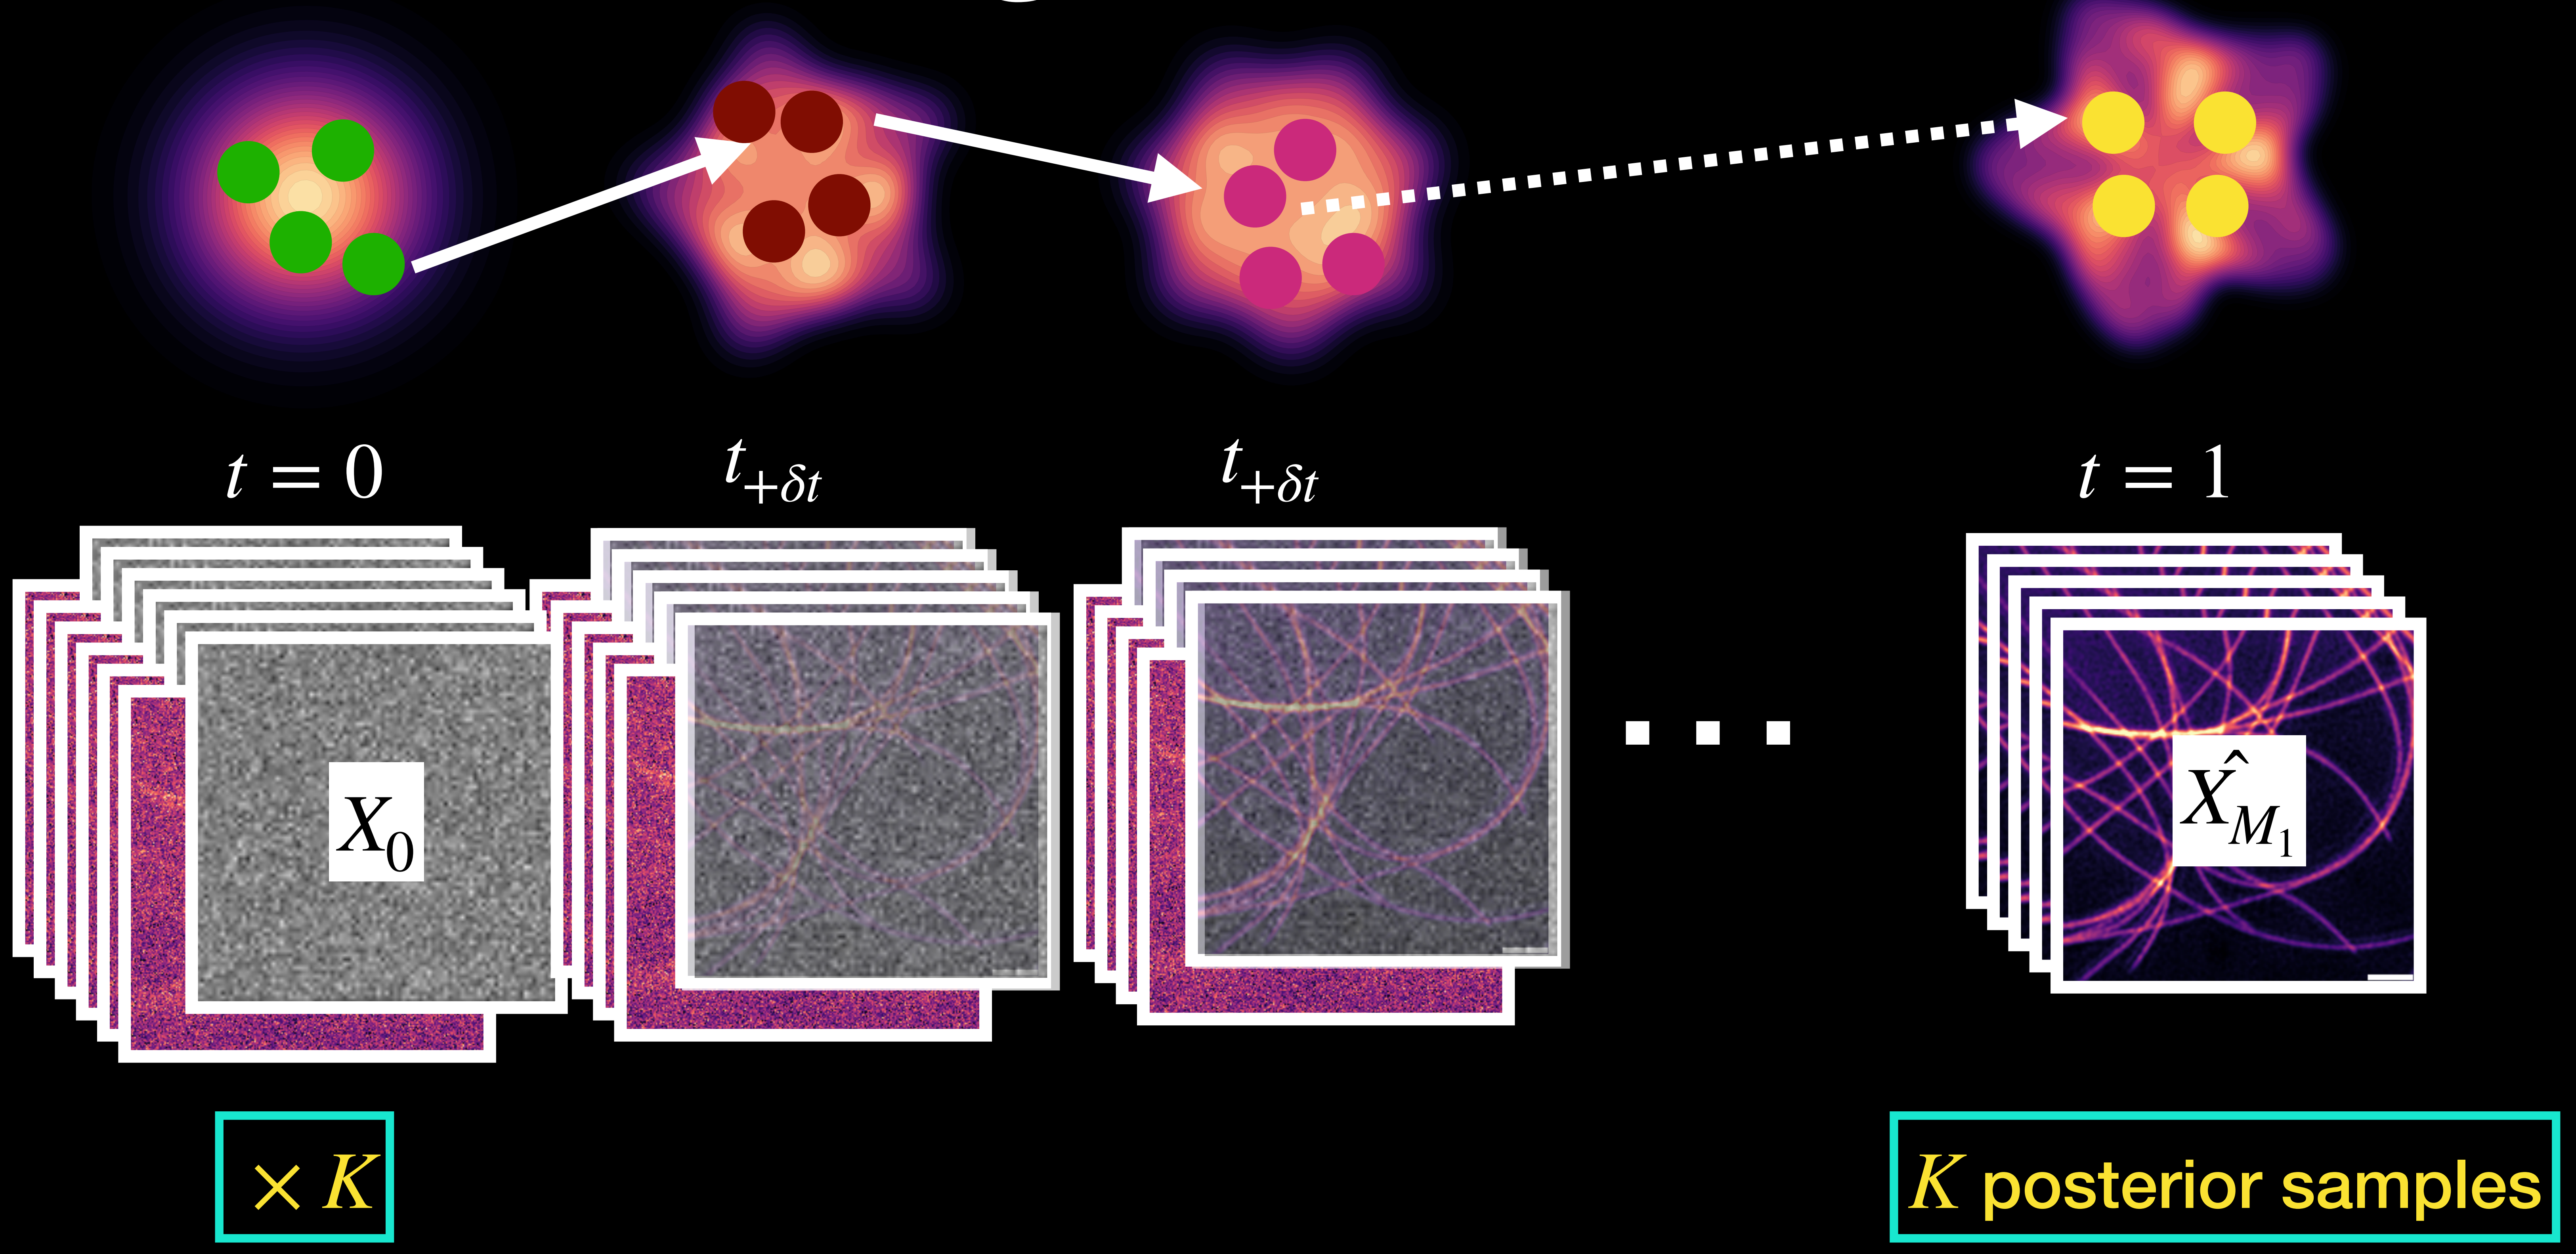

Multiple inference cycles generates posterior samples

$\frac{1}{K}\sum K$

MMSE

$Var(K)$

Posterior variance

BioSR Data (heavy Noise)

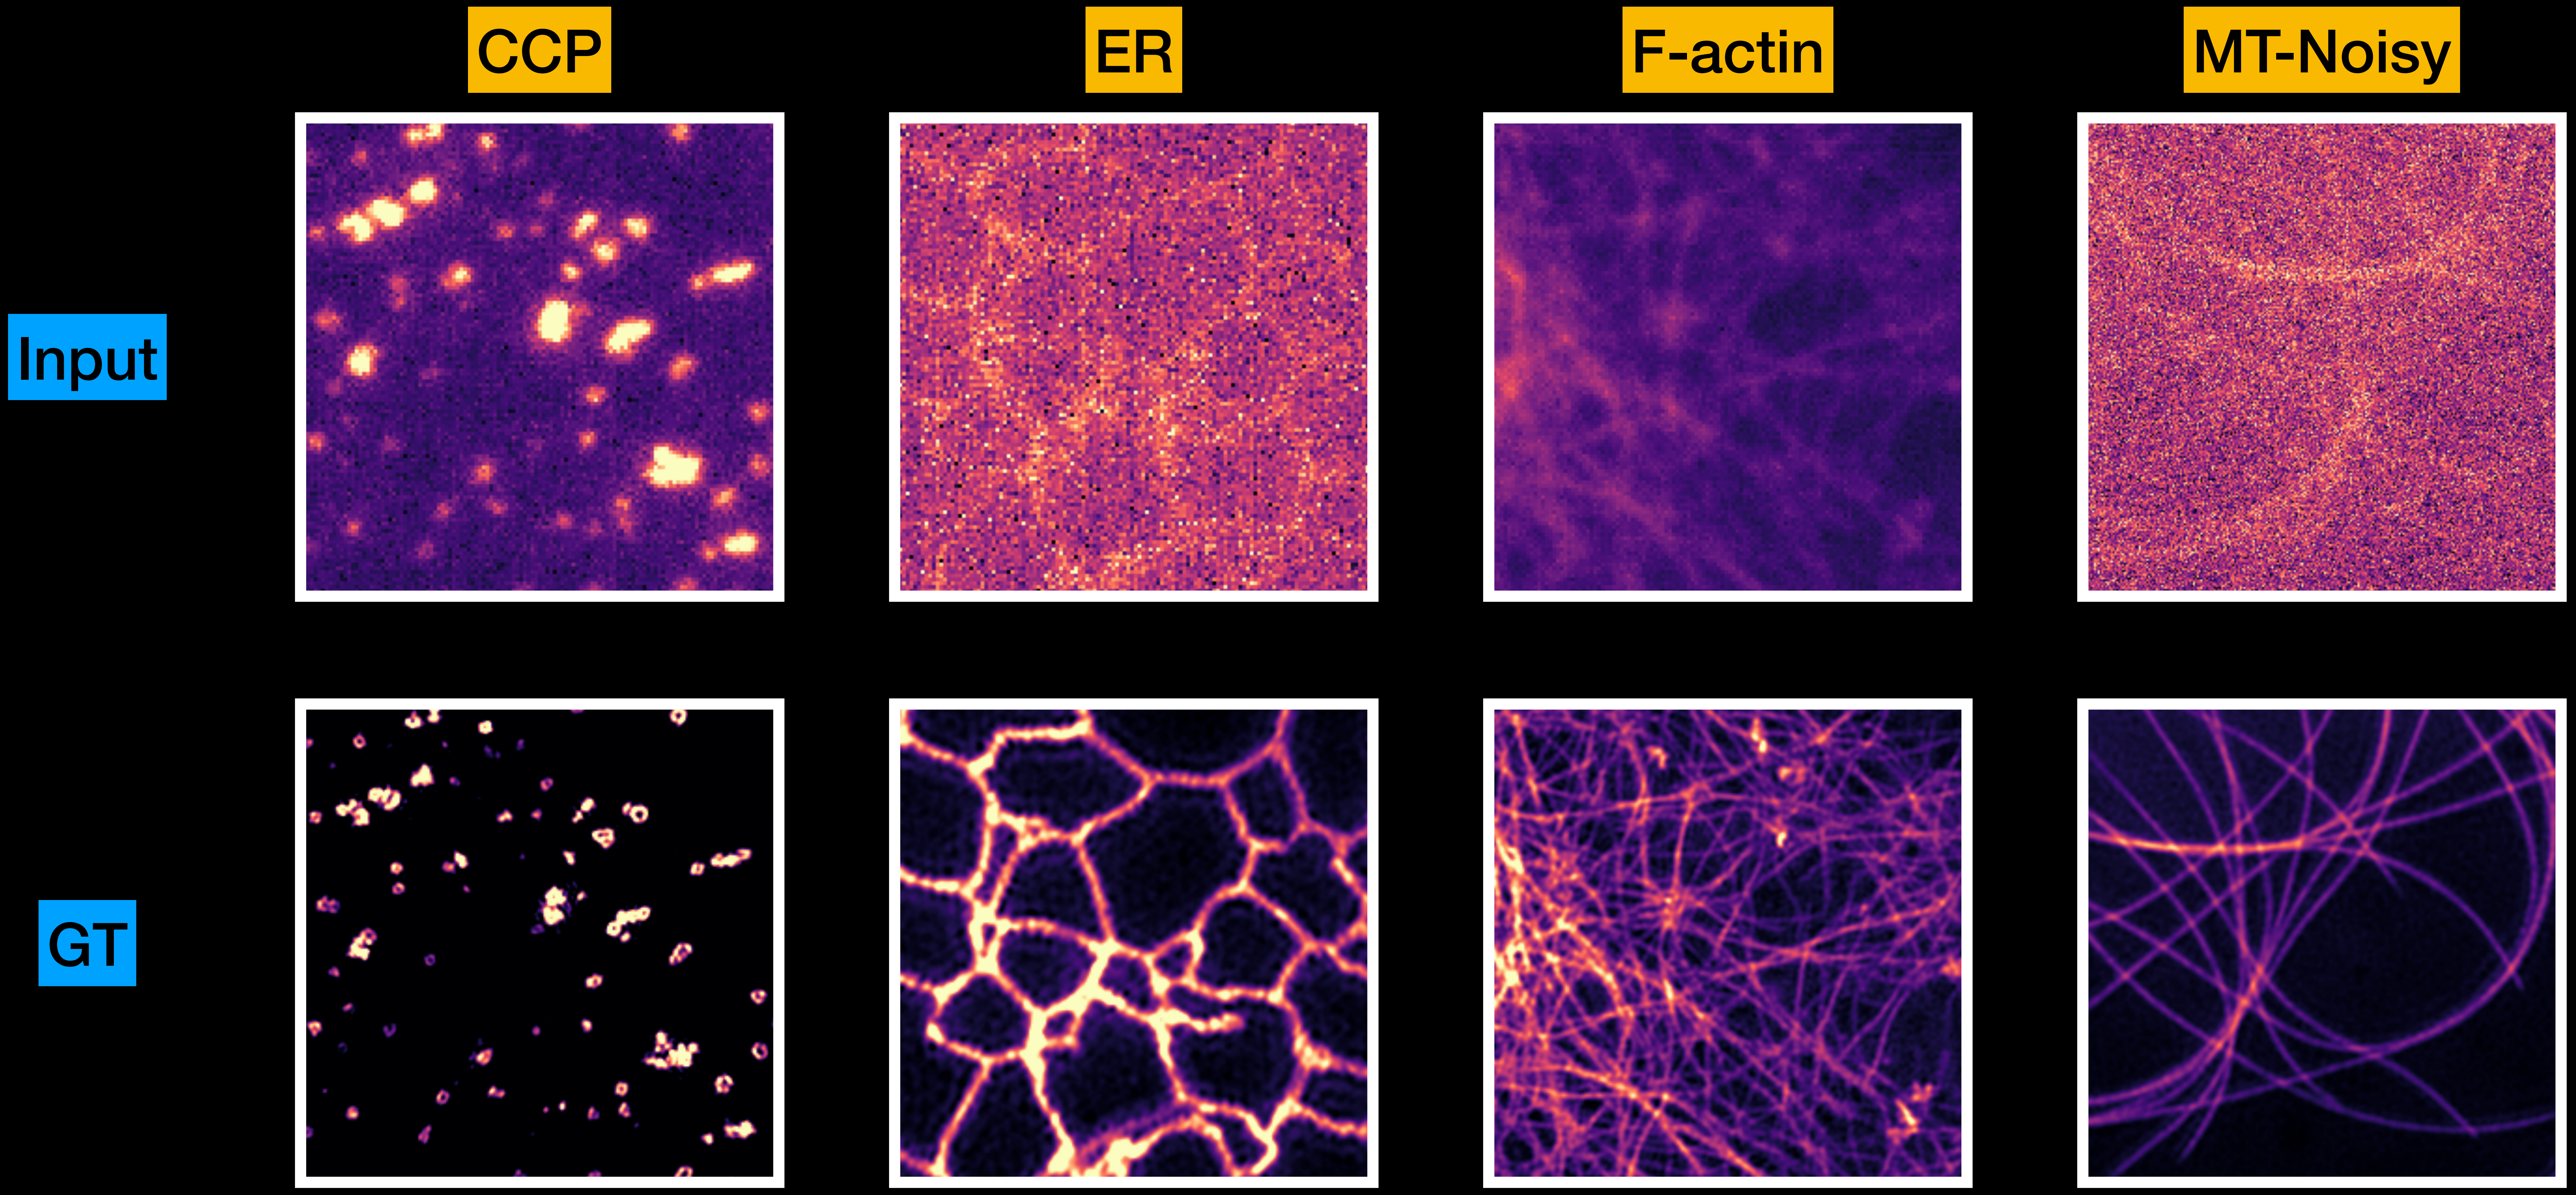

$2\times$ Upsampling under extreme uncertainty (noise)

Baselines

	Single Prediction	Posterior Sampling
One-shot	✅ UNet, RCAN, ESRGAN, InDI-1	✅ LVAE
Iterative	✅ InDI-20	✅ SIFM, ResMatching

Metrics

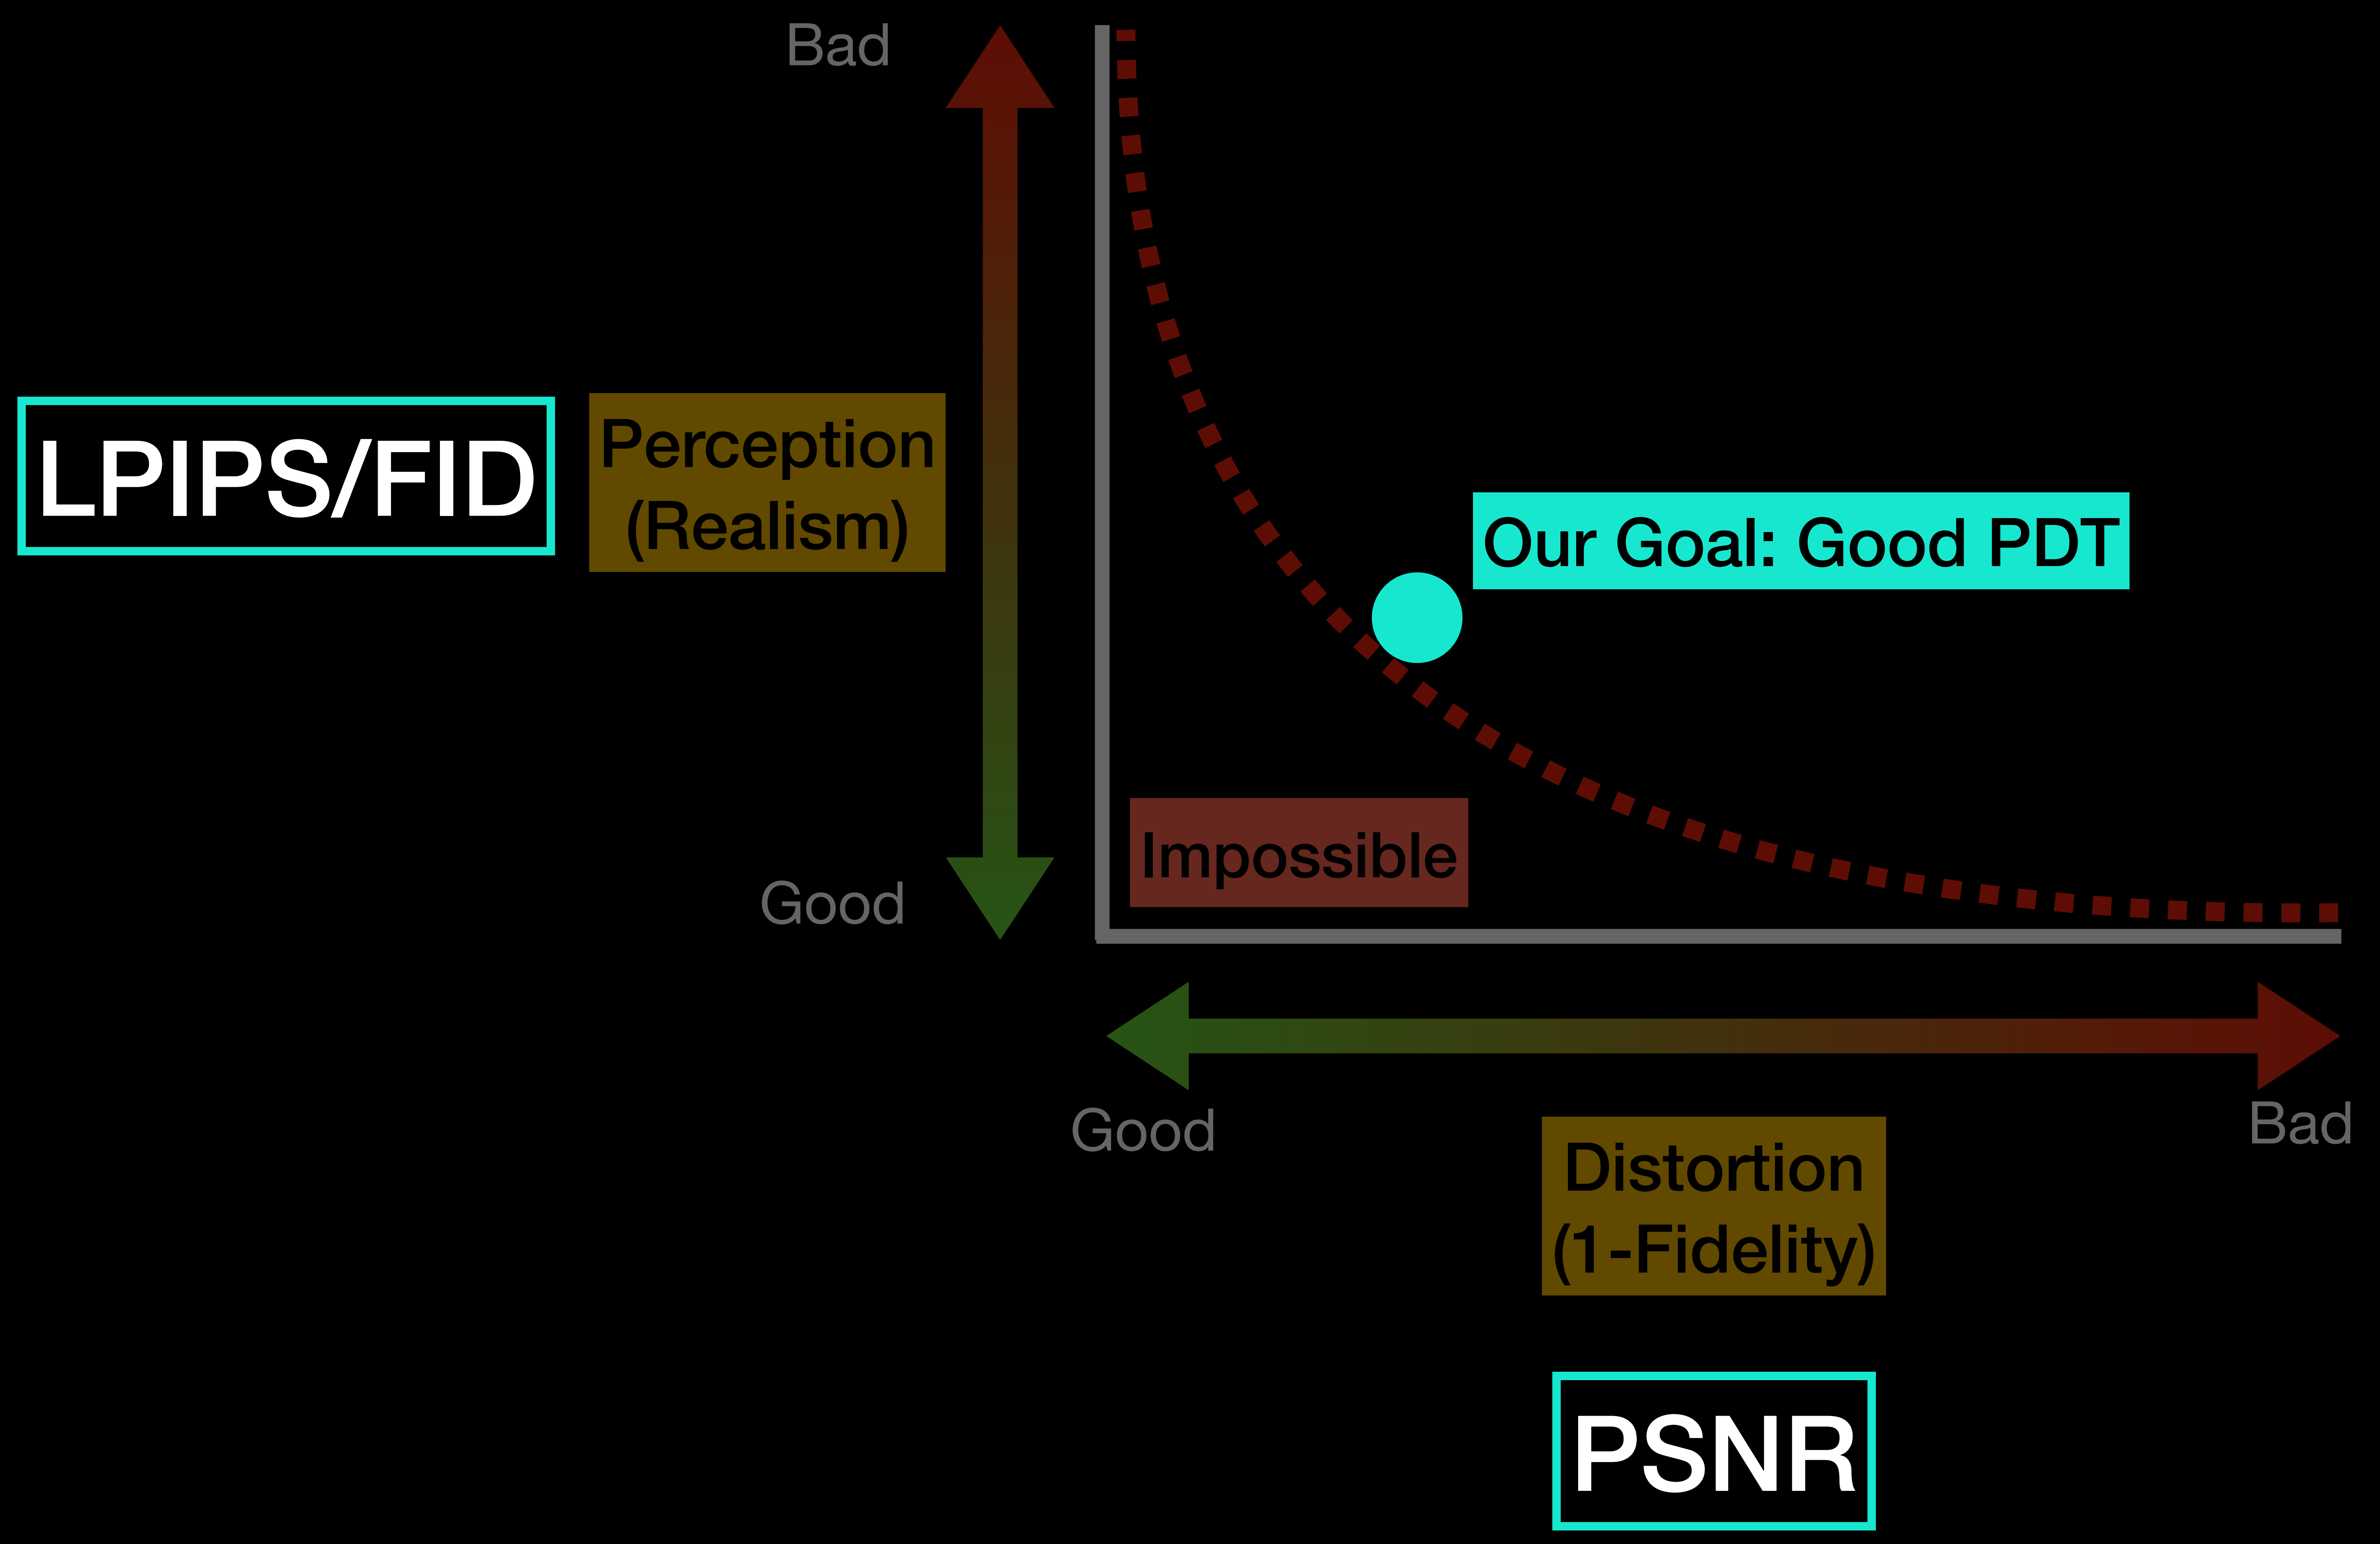
LPIPS/FID
Perception
(Realism)
Bad
Good
Our Goal: Good PDT
Impossible
Good
Bad
Distortion
(1-Fidelity)
PSNR

Metrics

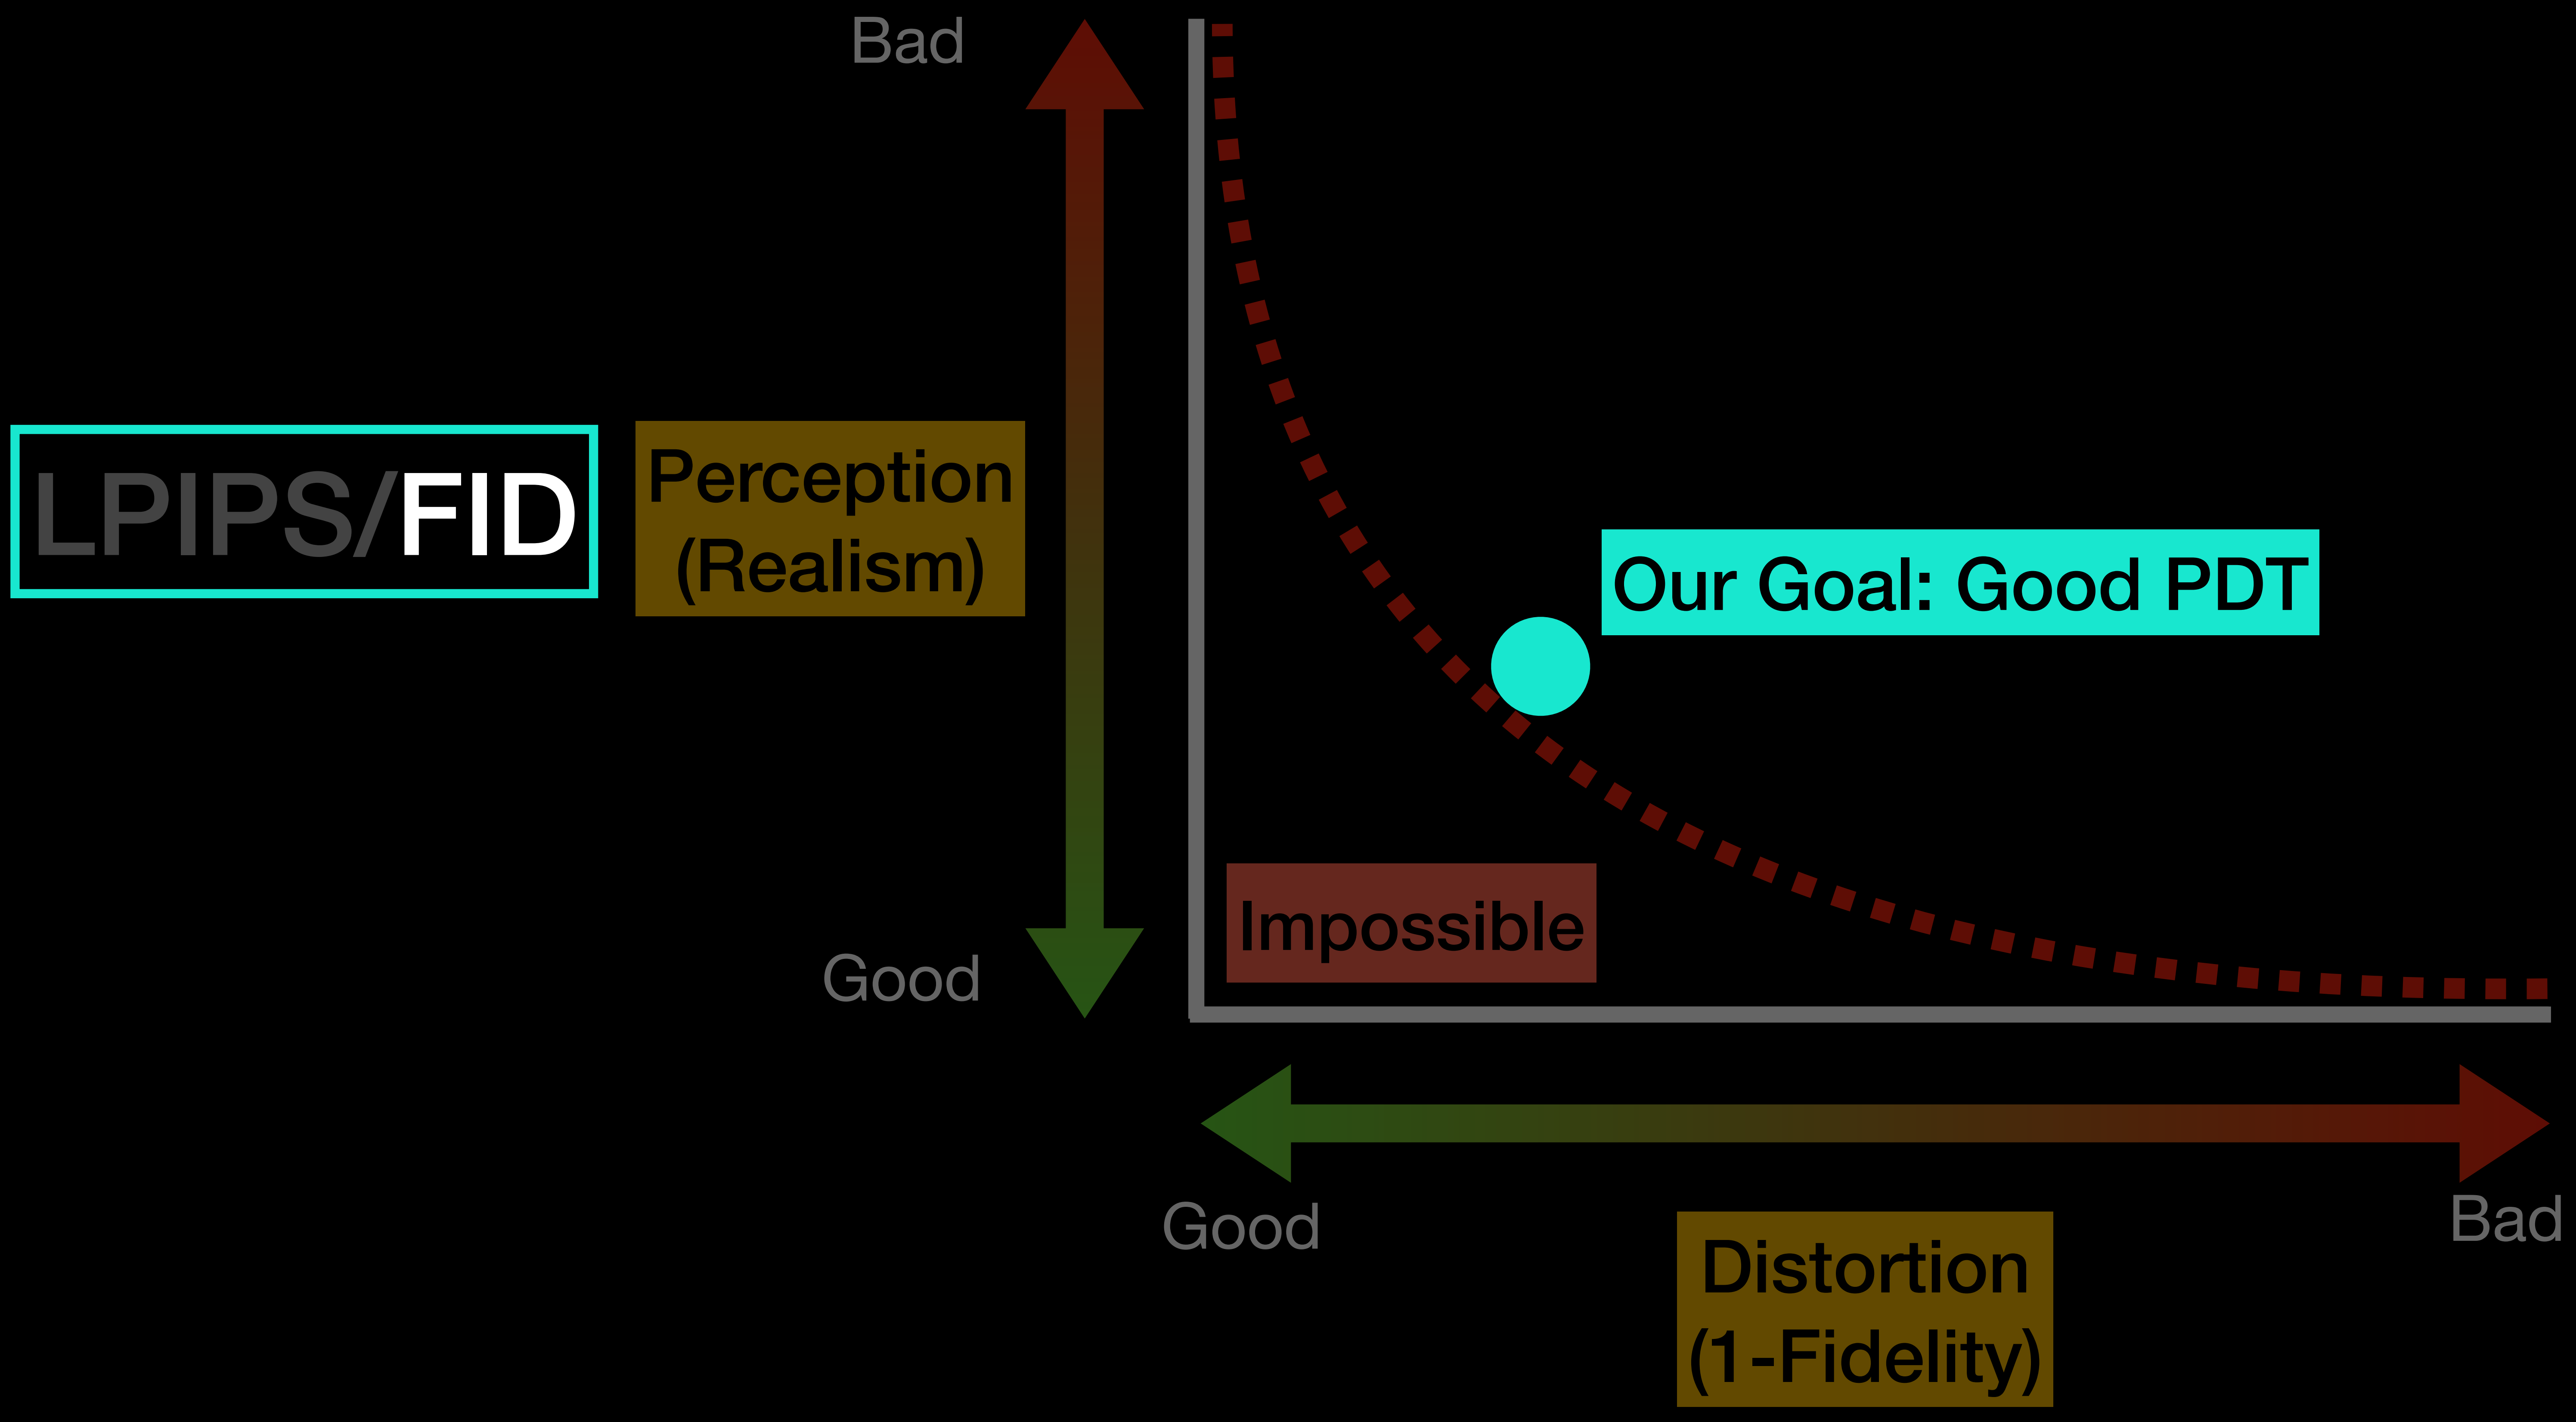

PSNR

Quantitative: FID vs PSNR

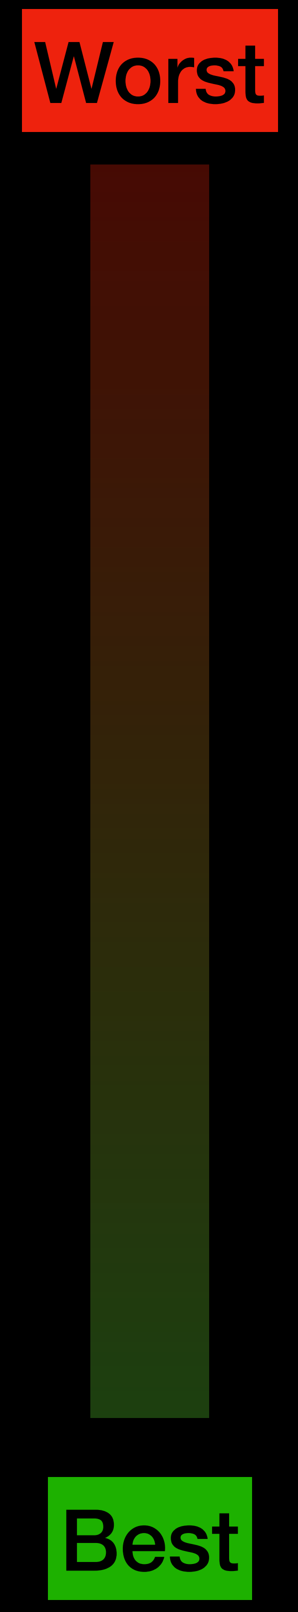

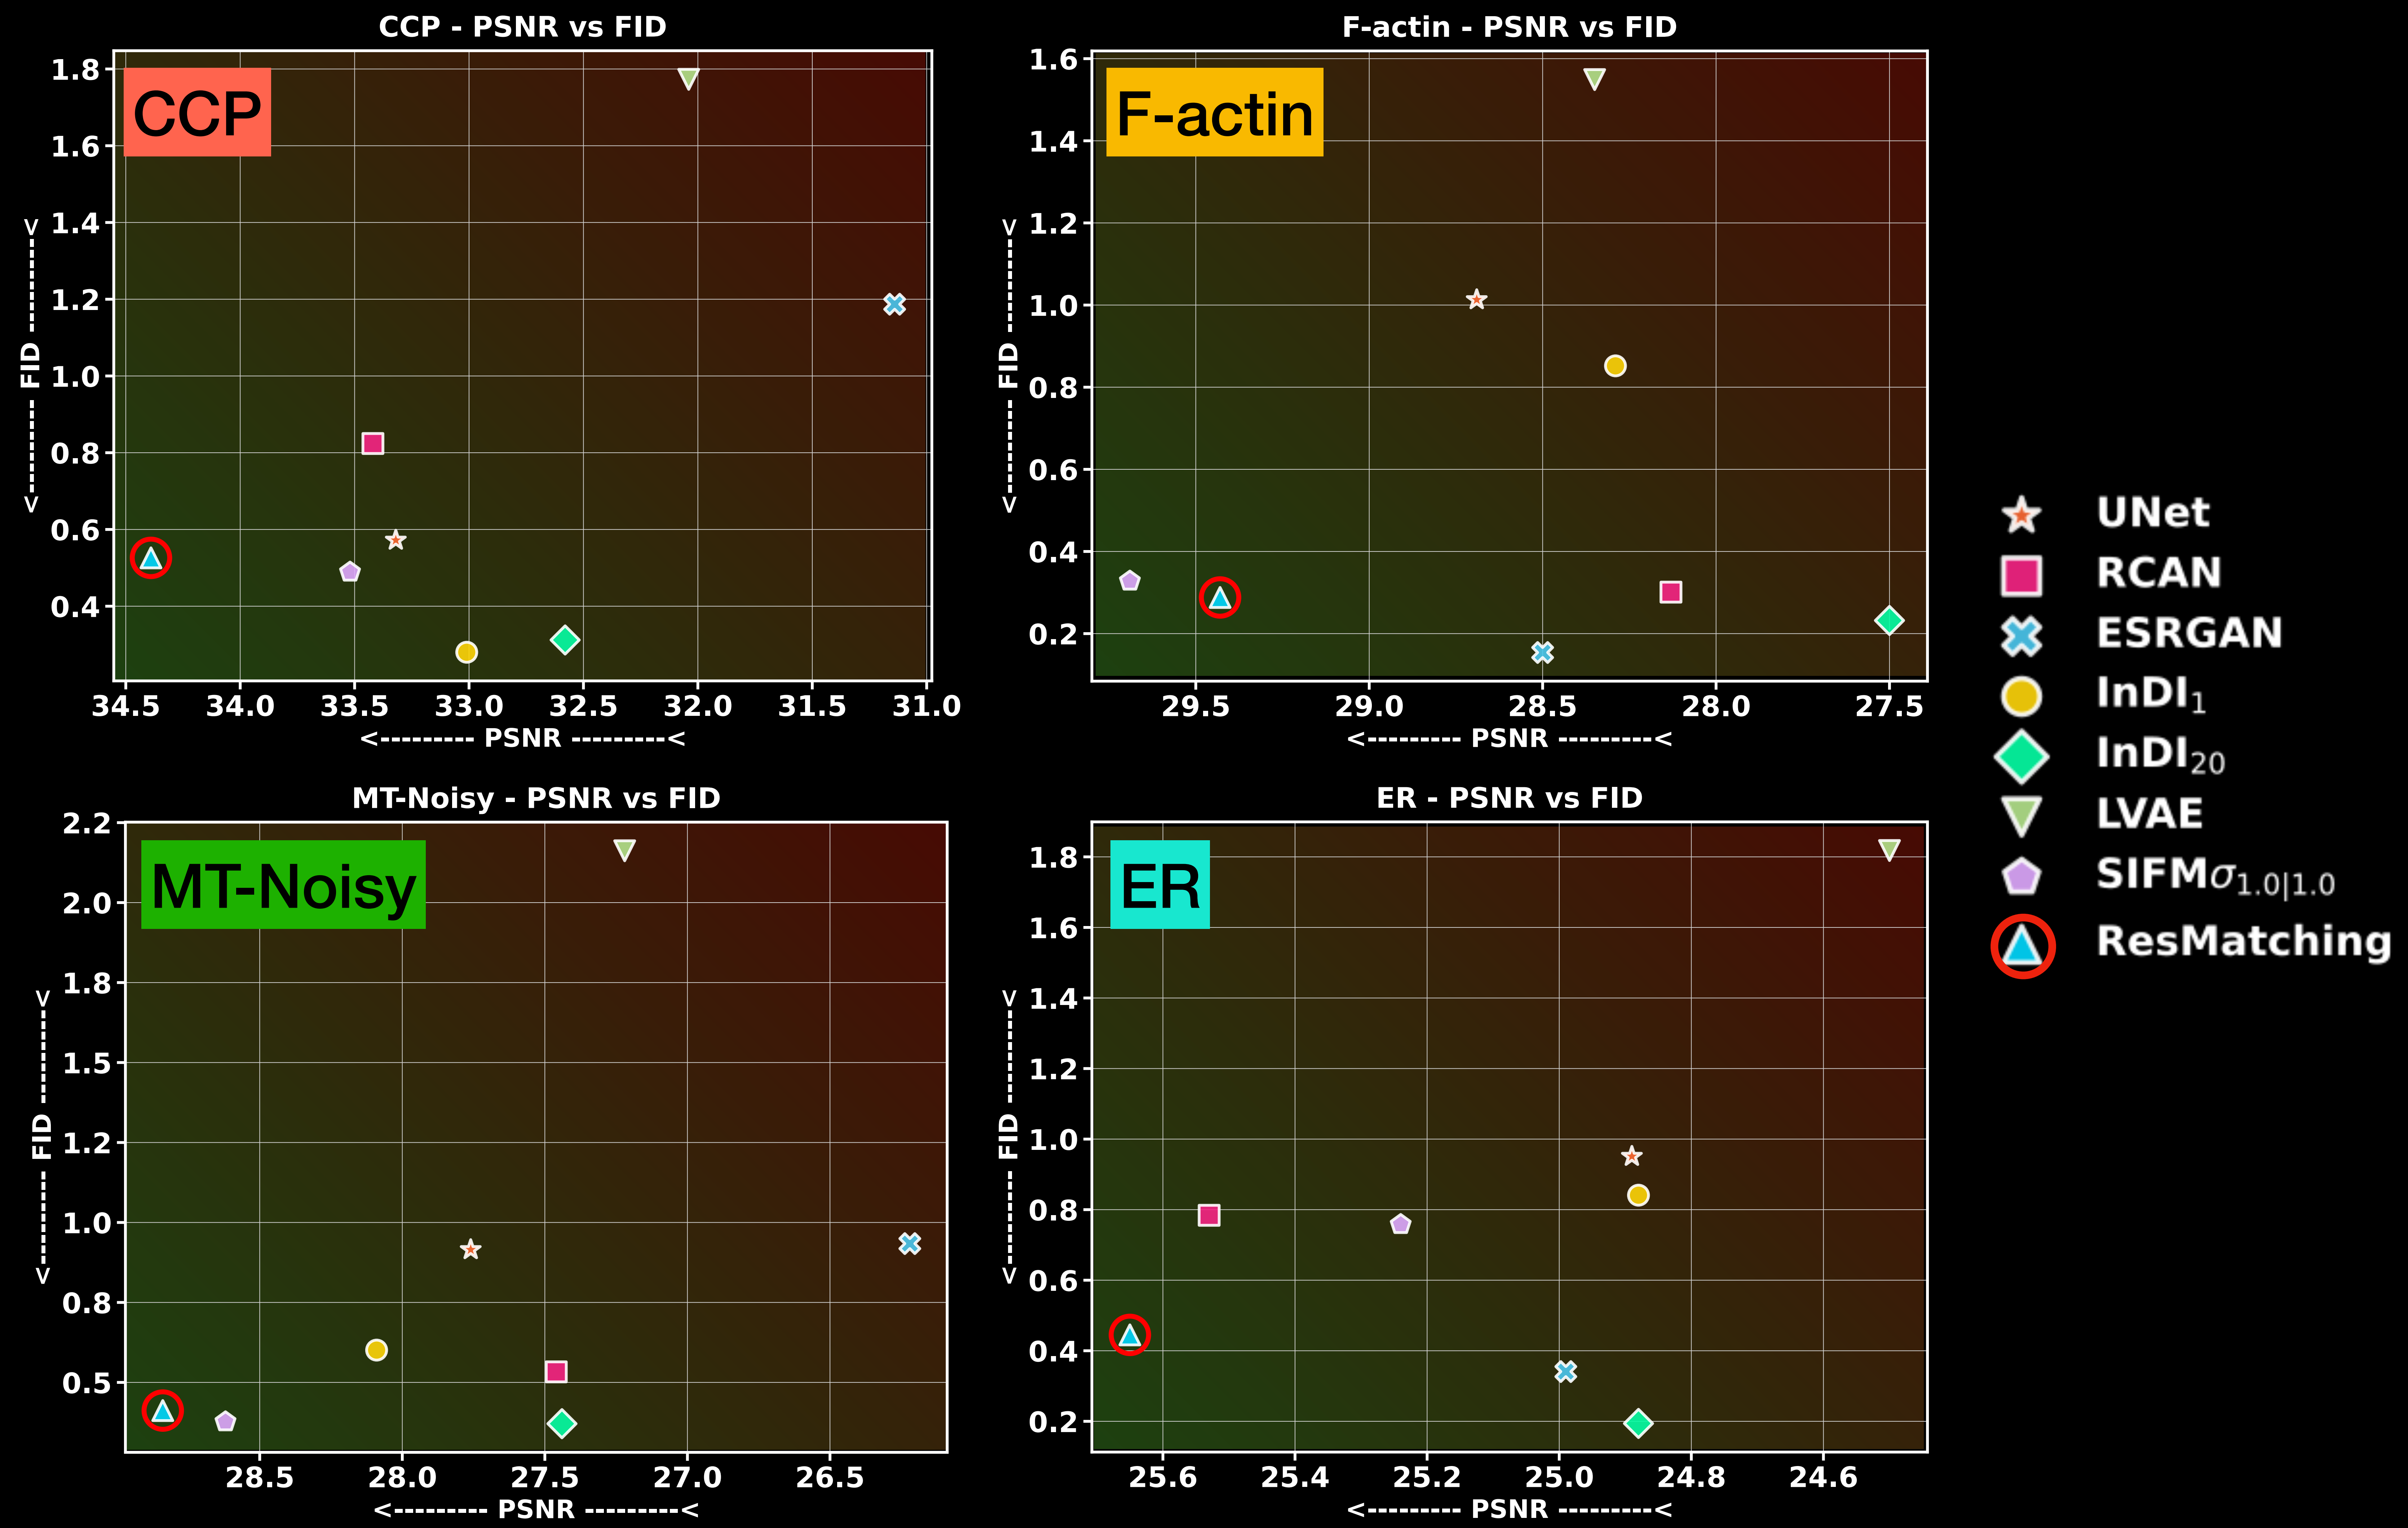

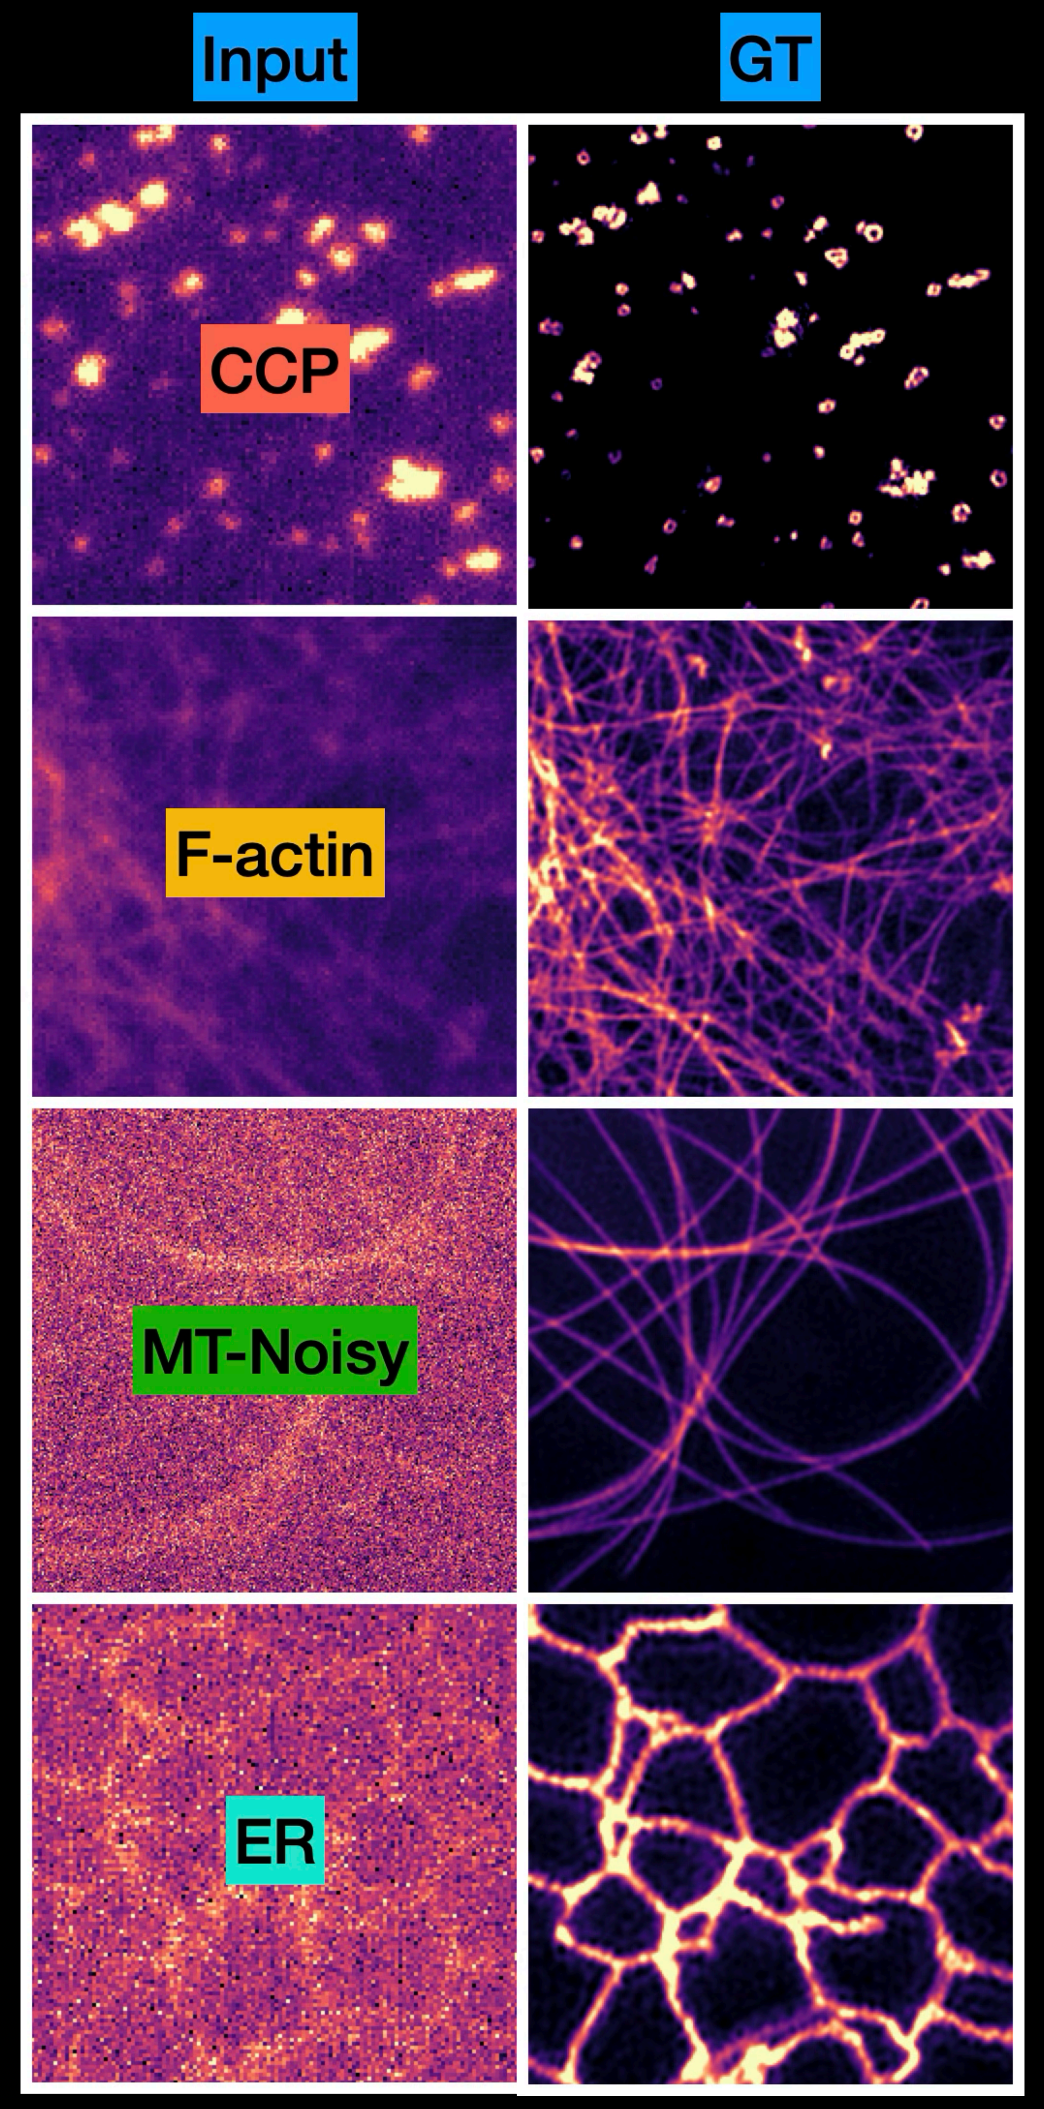

ResMatching Qualitative Results: MT-Noisy

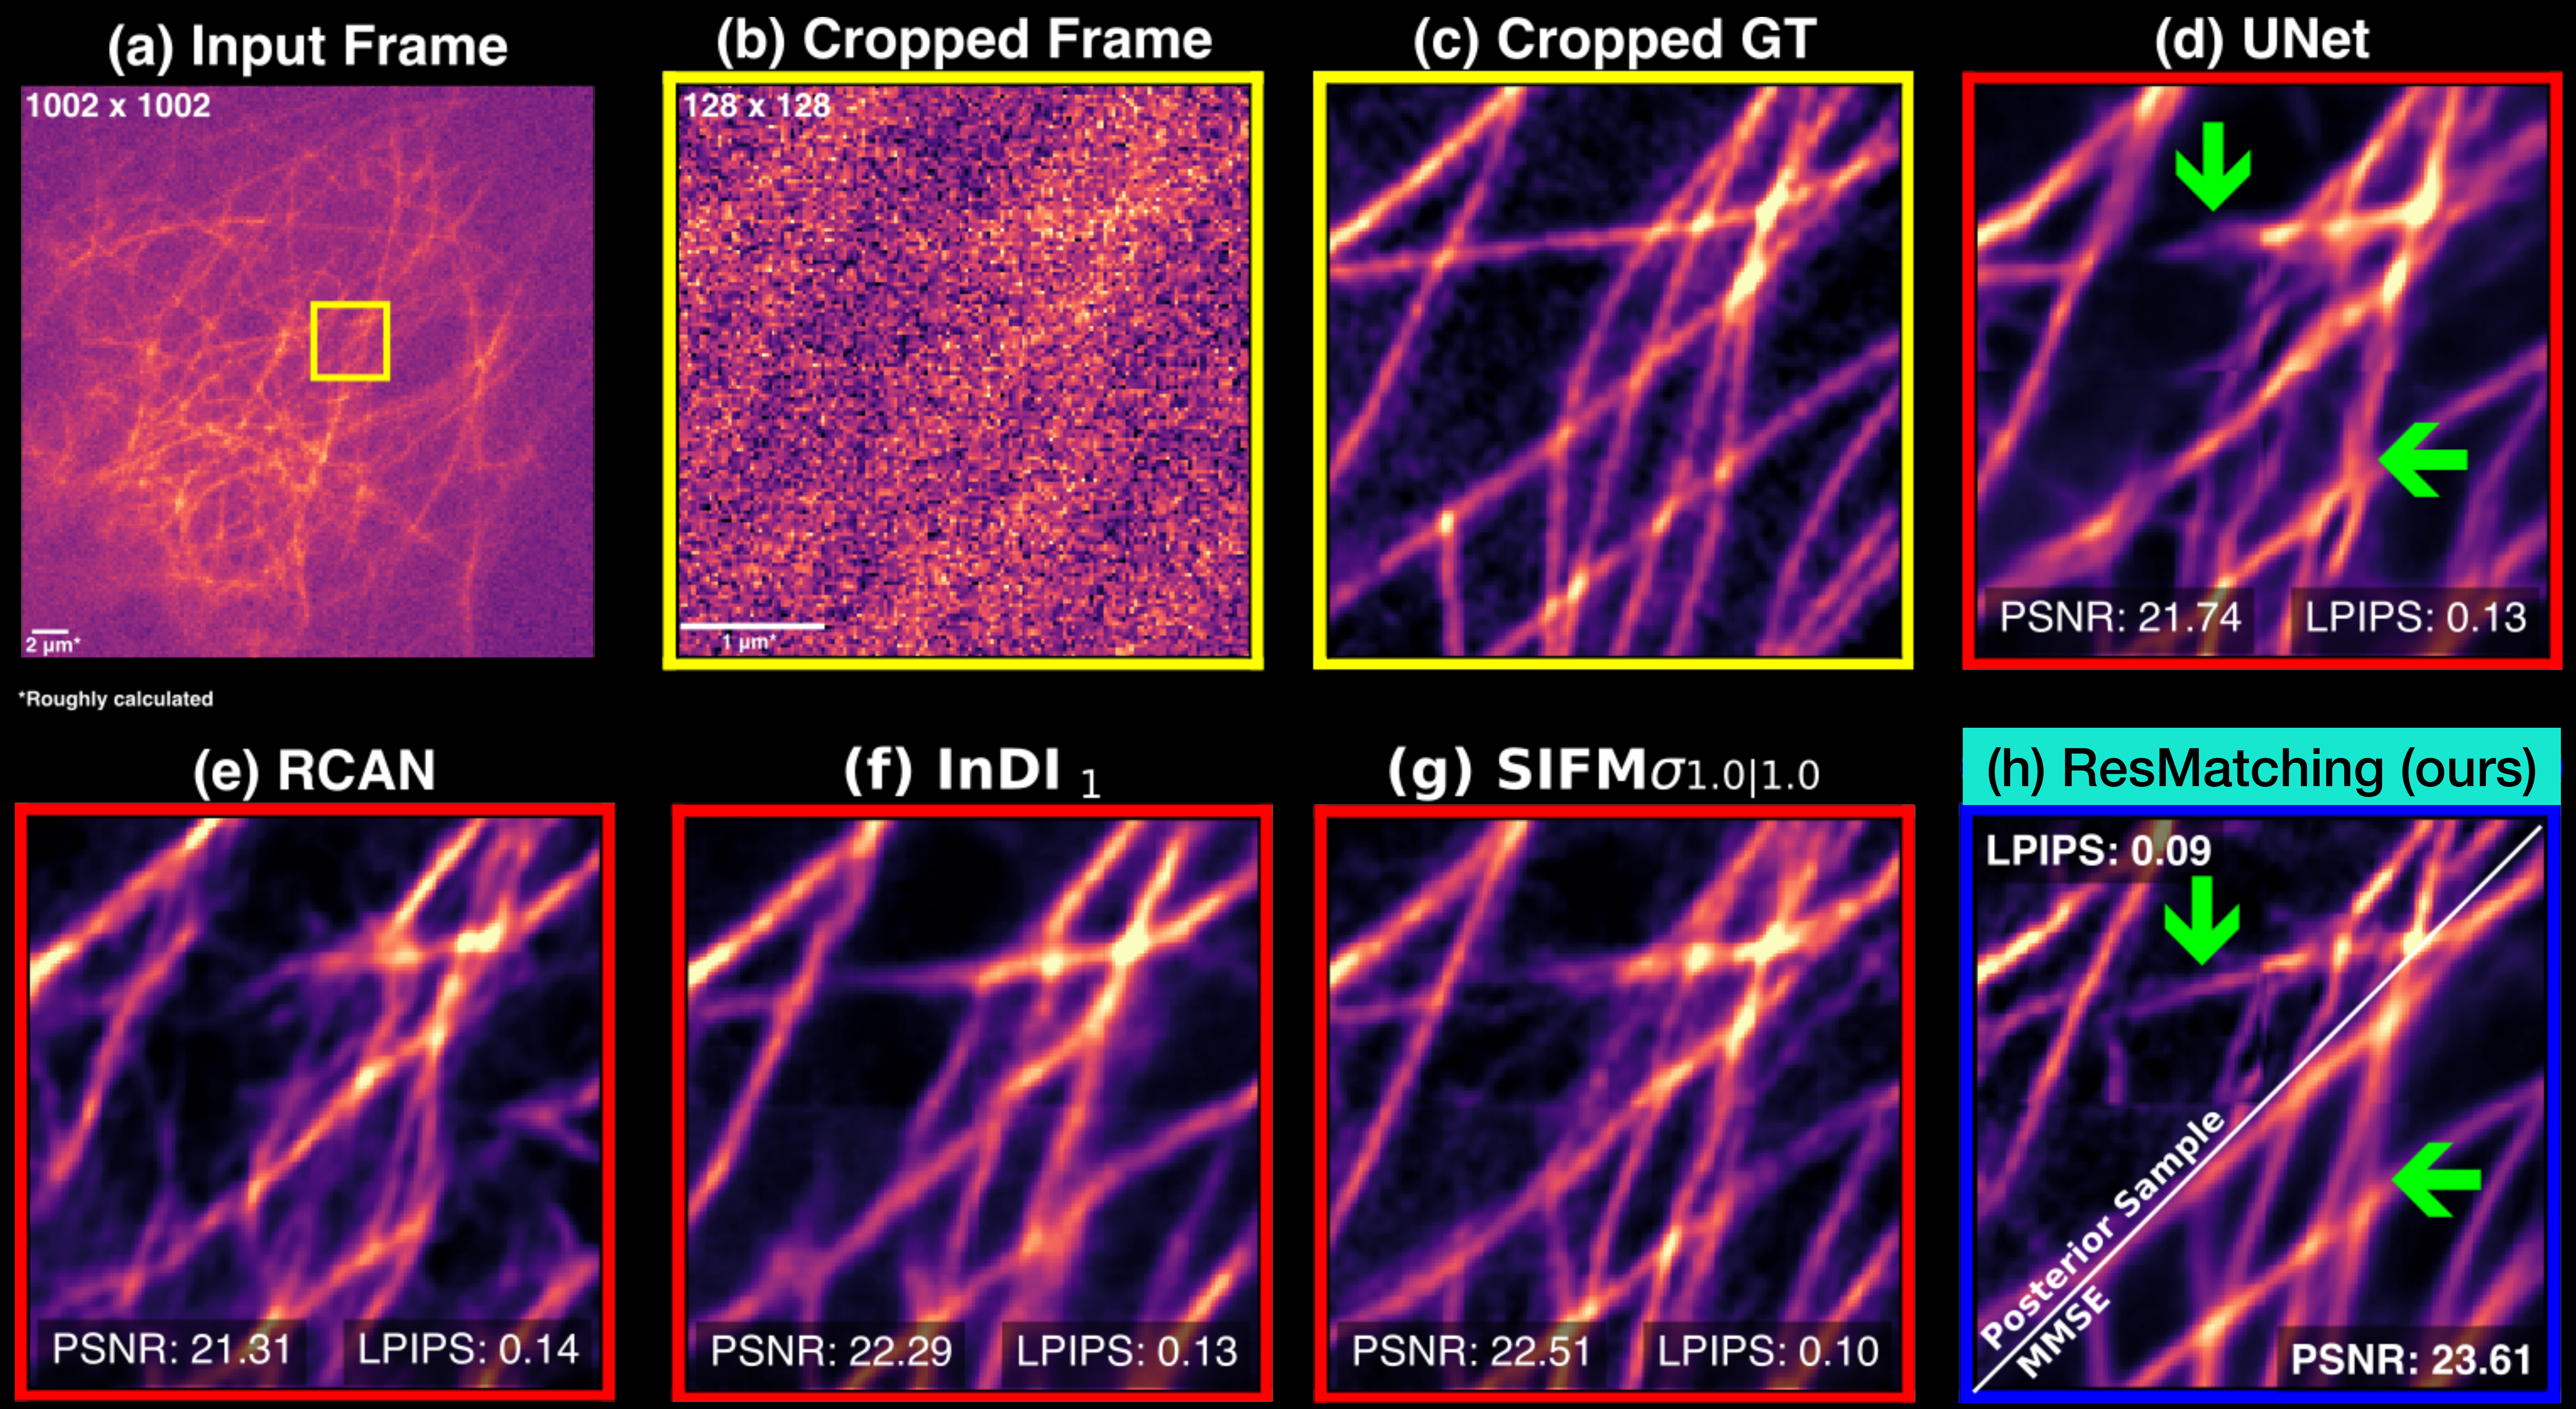

How well does it handle uncertainty?

MT vs MT-Noisy: How well does it handle uncertainty?

BioSR MT LR Data

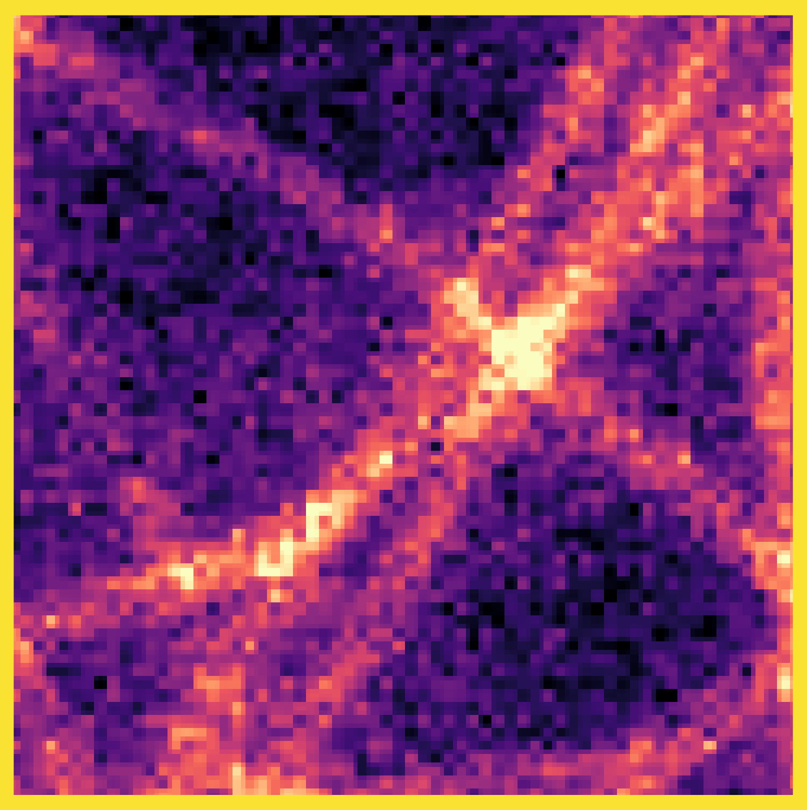

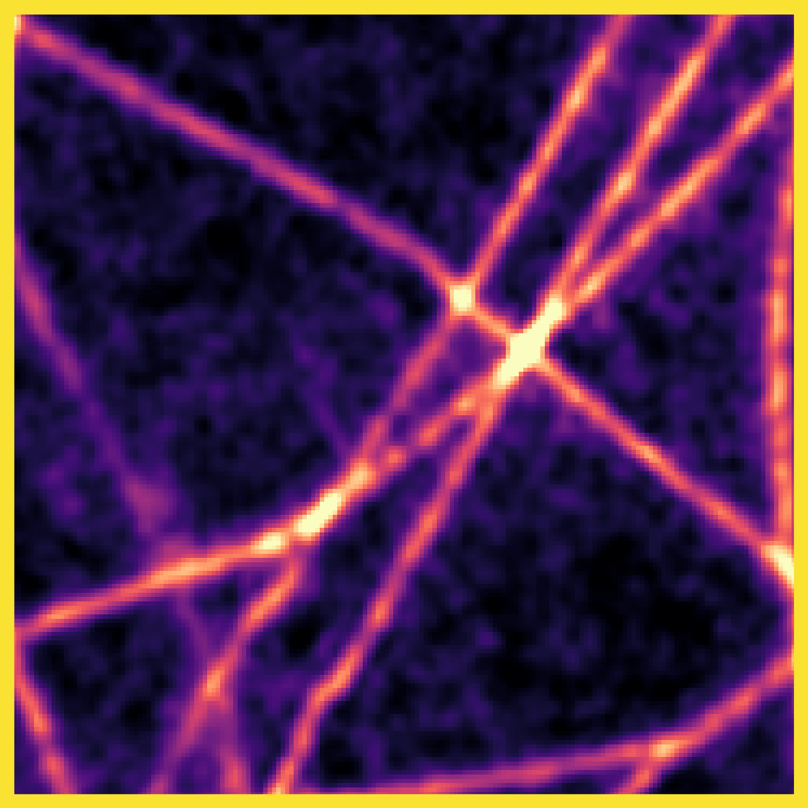

MT Data GT

MT vs MT-Noisy: How well does it handle uncertainty?

BioSR MT LR Data

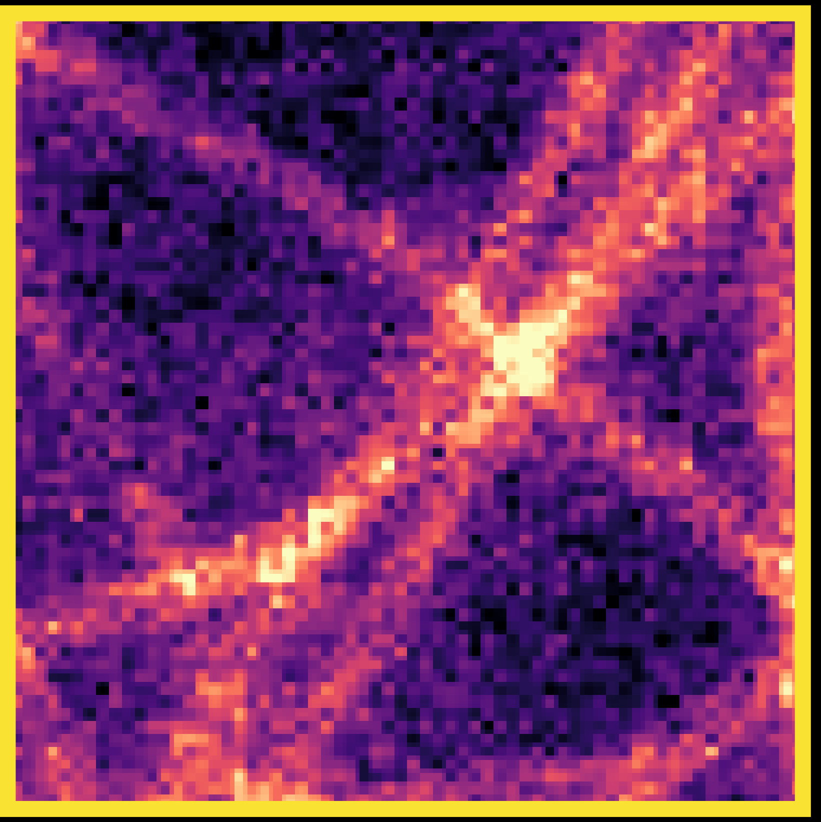

ResMatching Pred

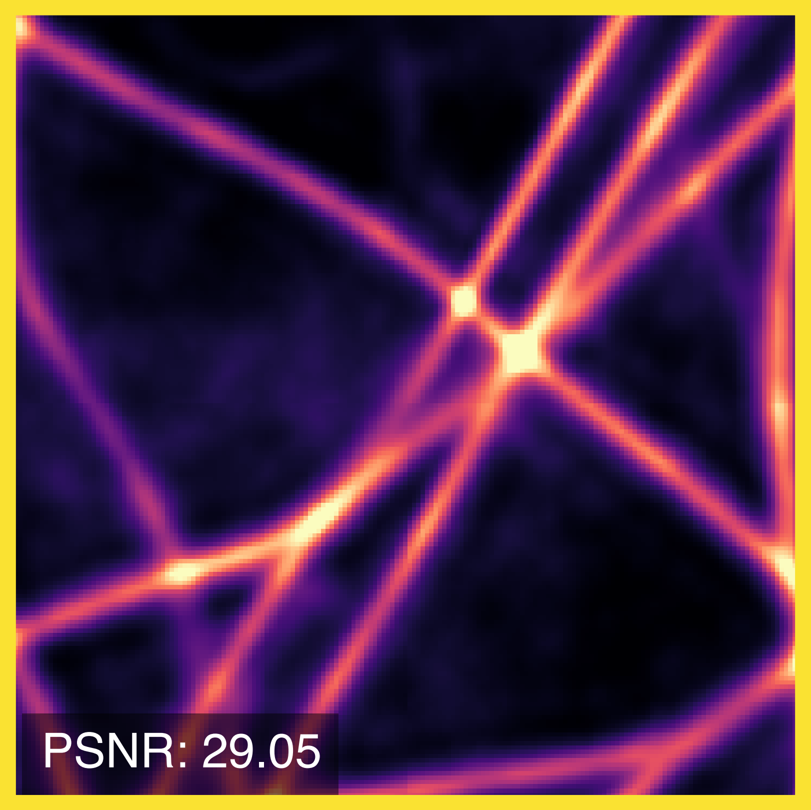

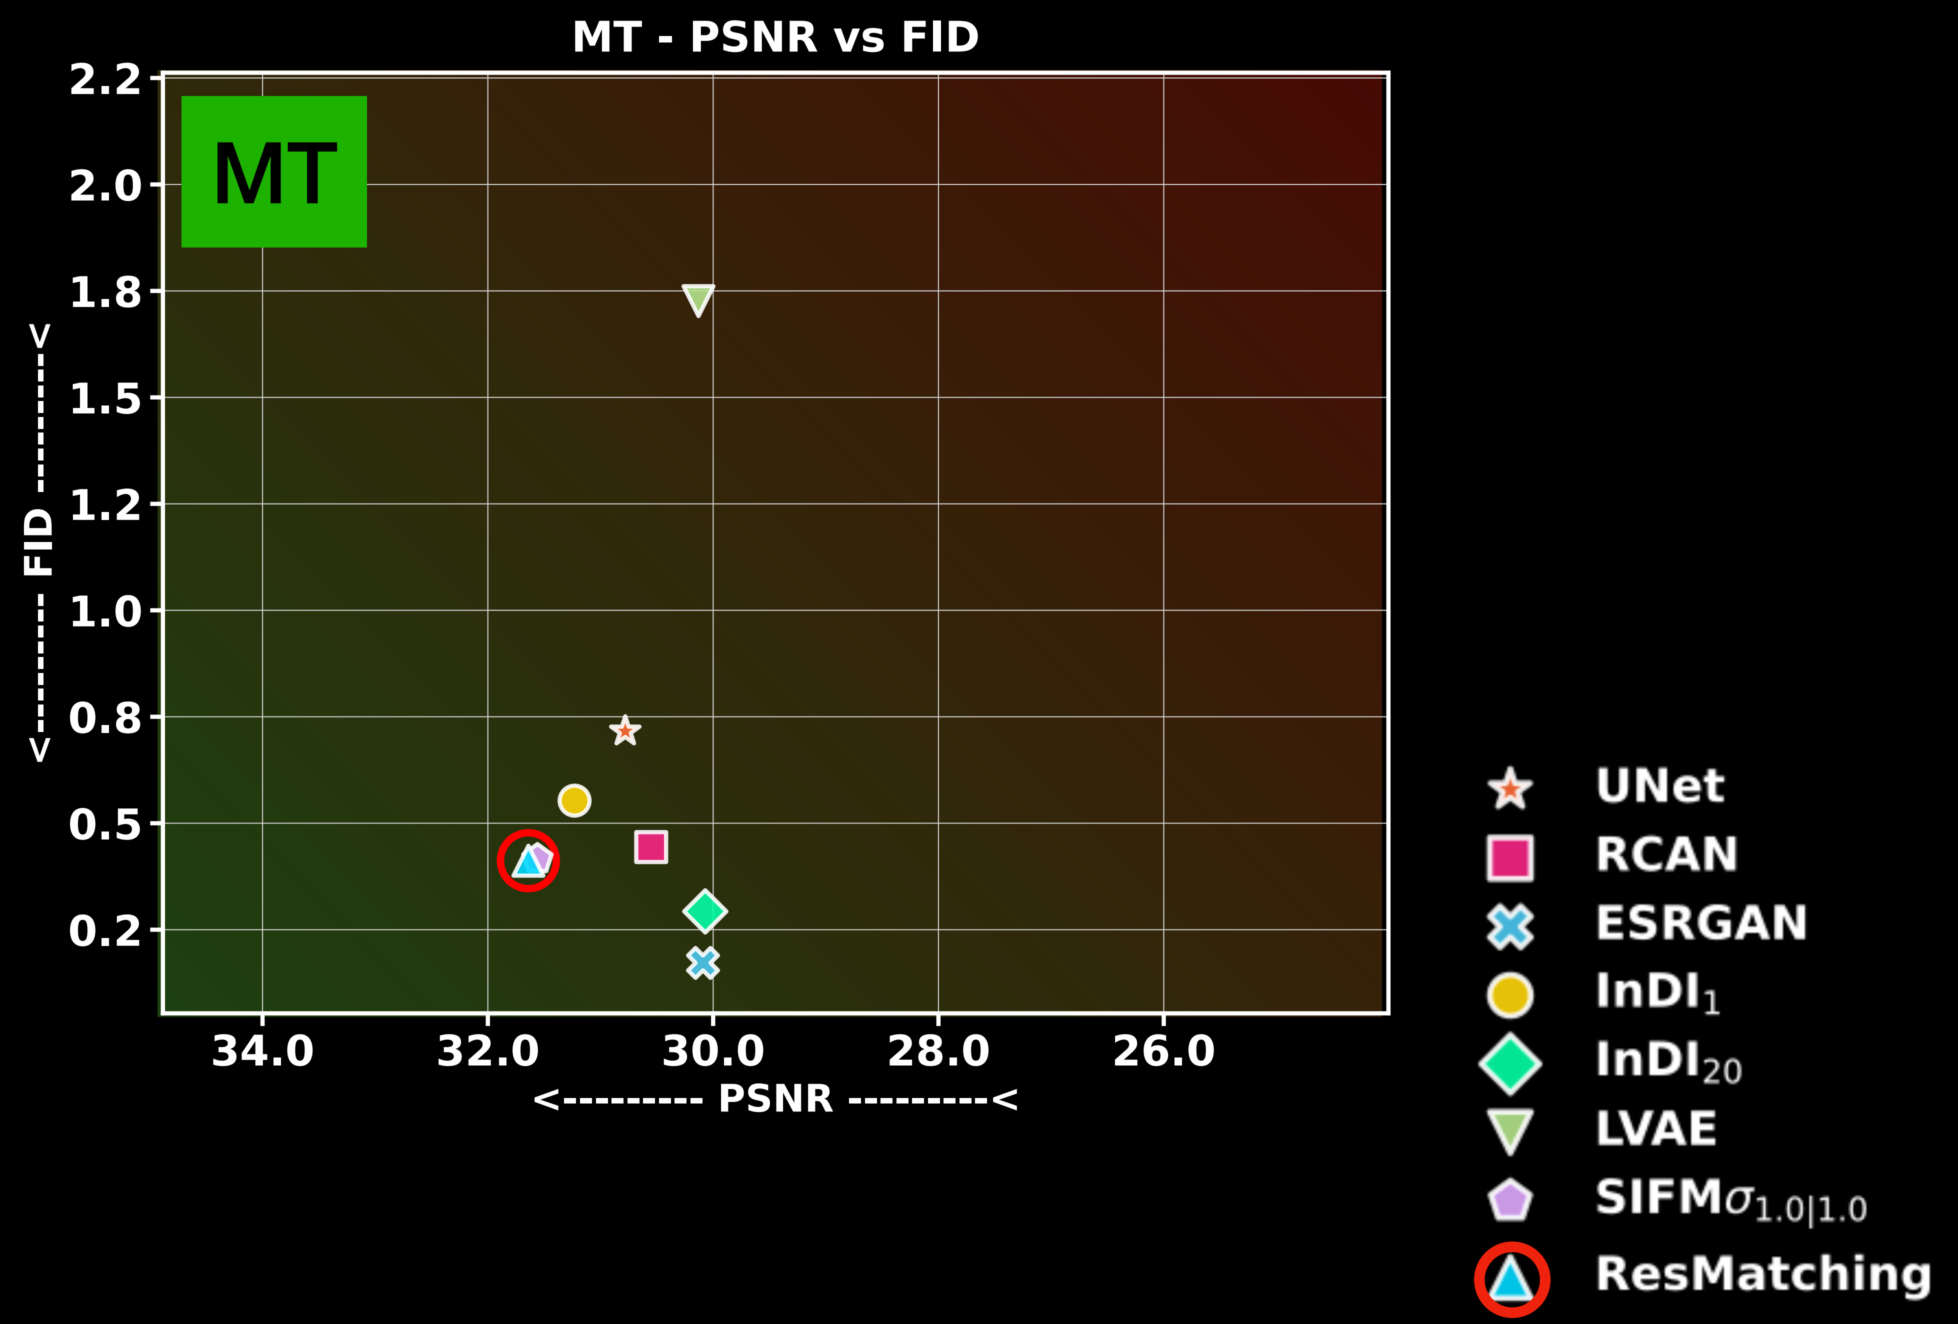

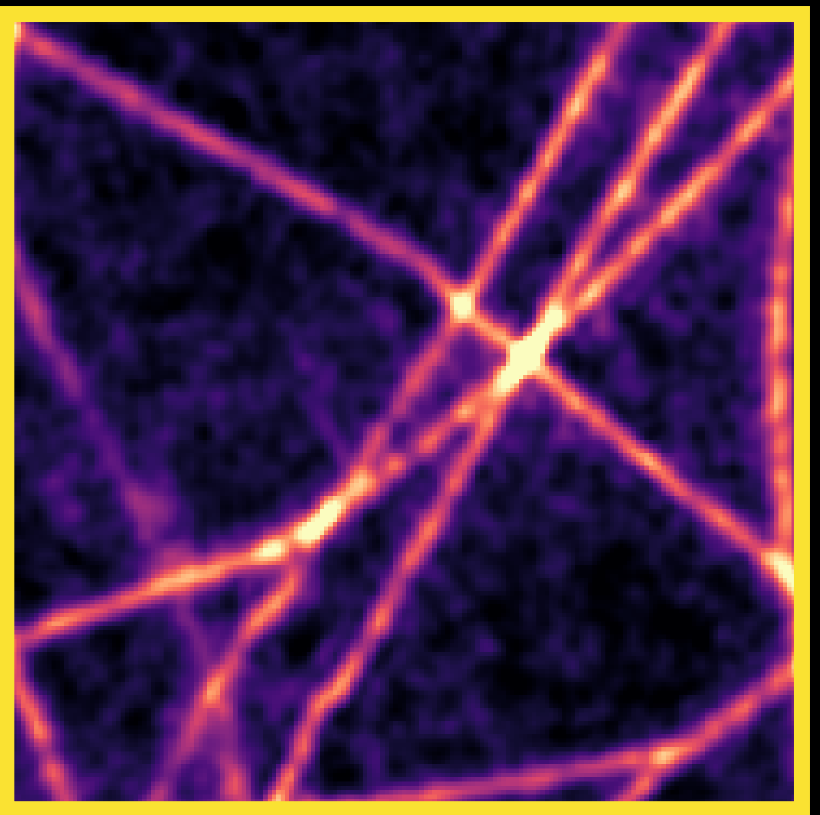

MT Data GT

MT vs MT-Noisy: How well does it handle uncertainty?

BioSR MT LR Data

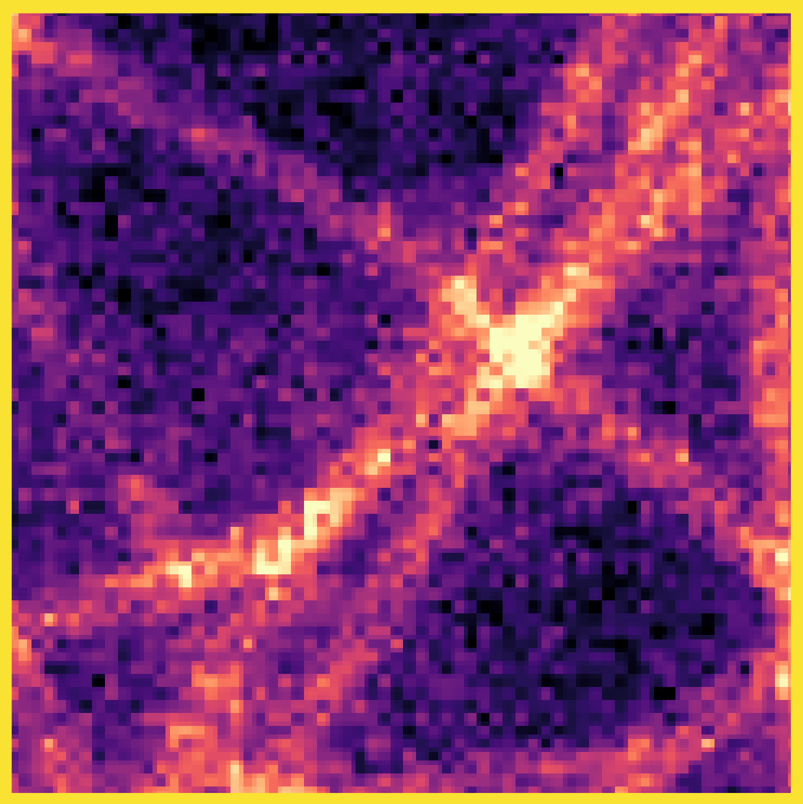

ResMatching Pred

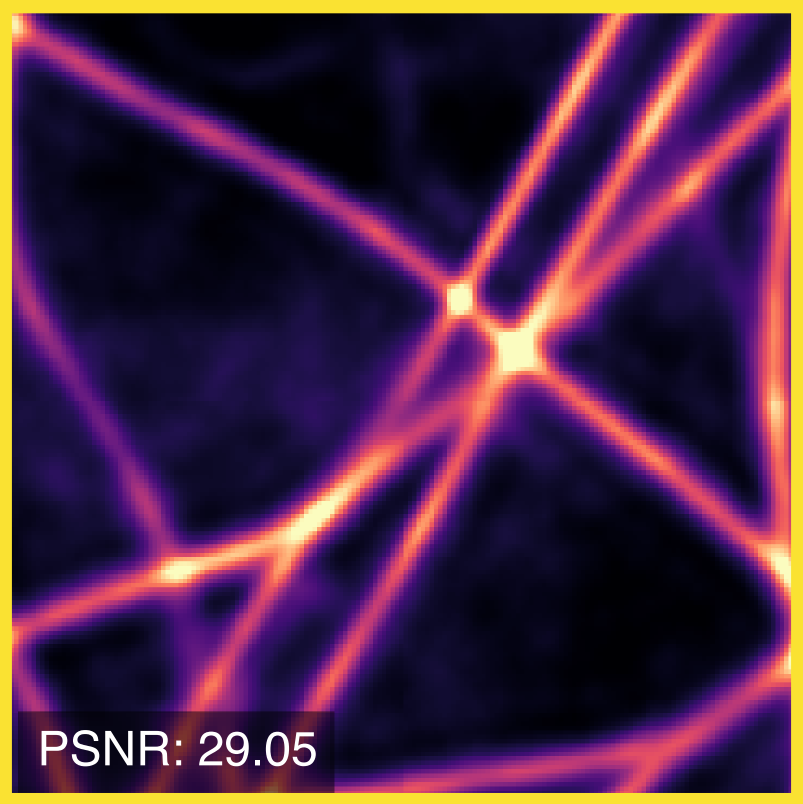

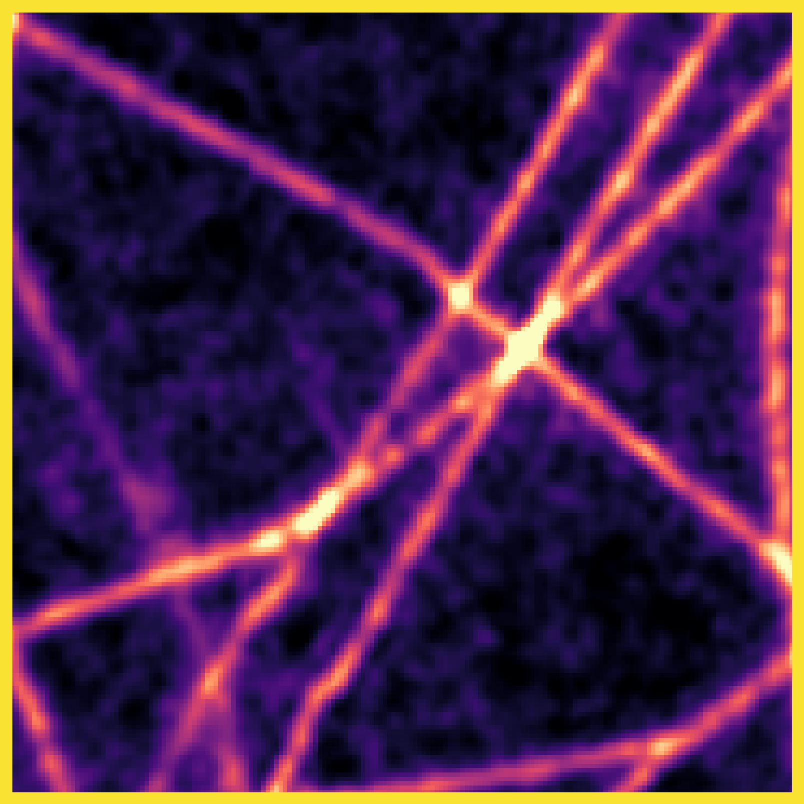

MT Data GT

MT-Noisy LR Data

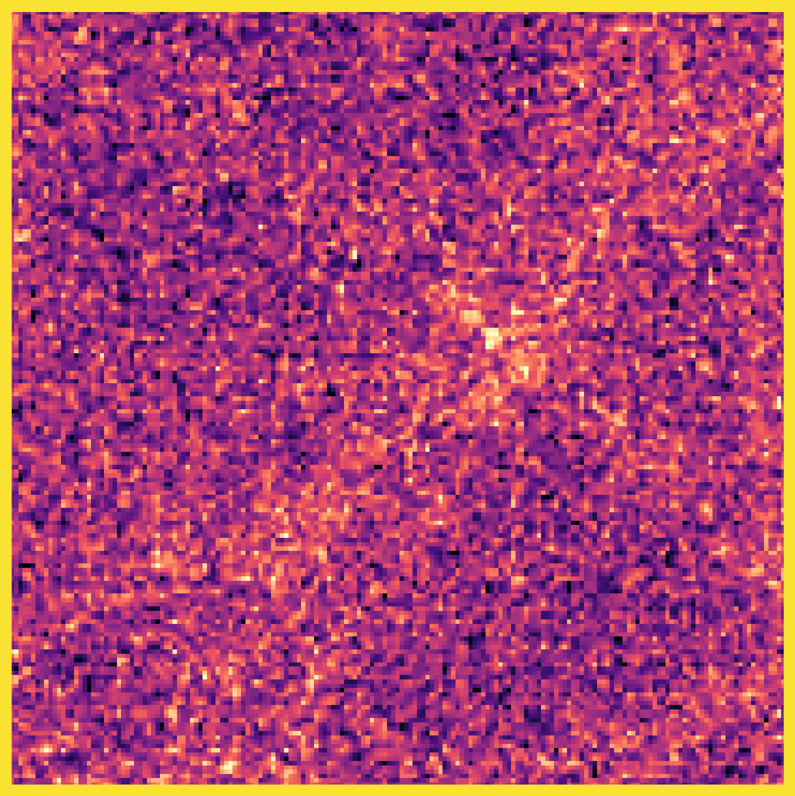

Additional Synthetic Noise

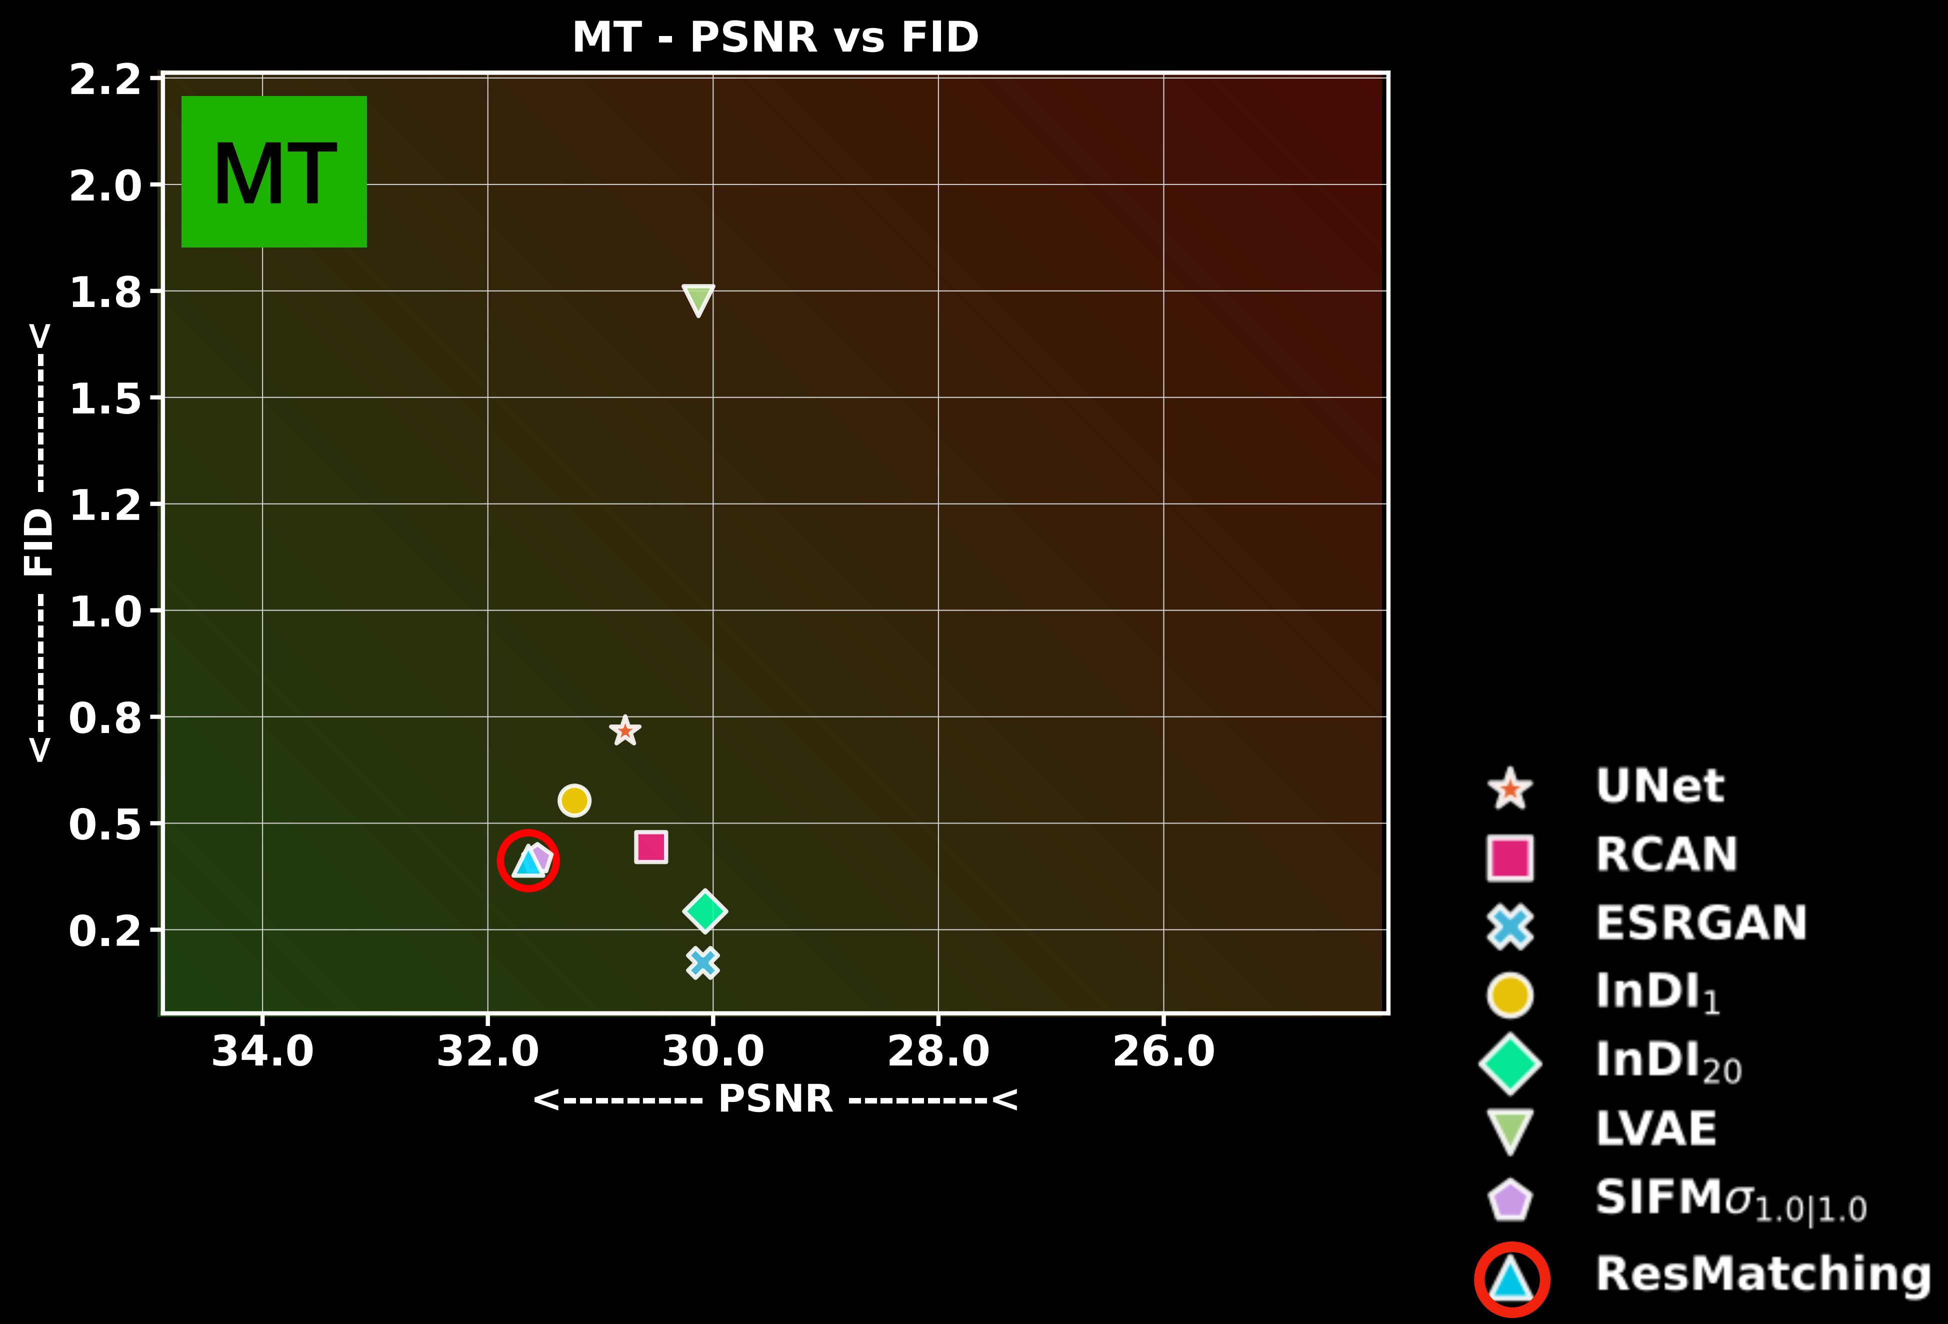

MT vs MT-Noisy: How well does it handle uncertainty?

BioSR MT LR Data

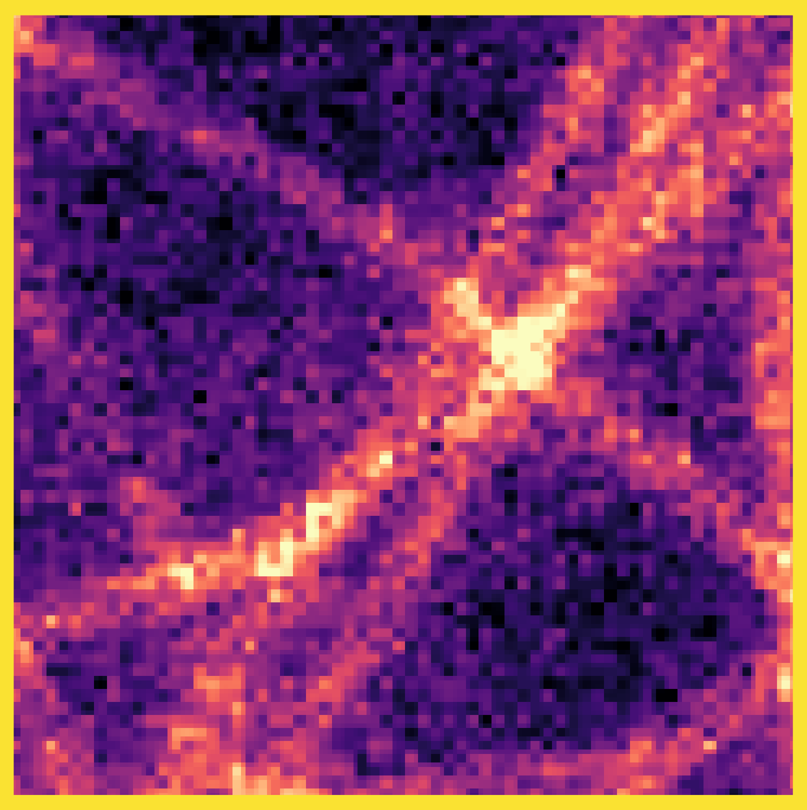

ResMatching Pred

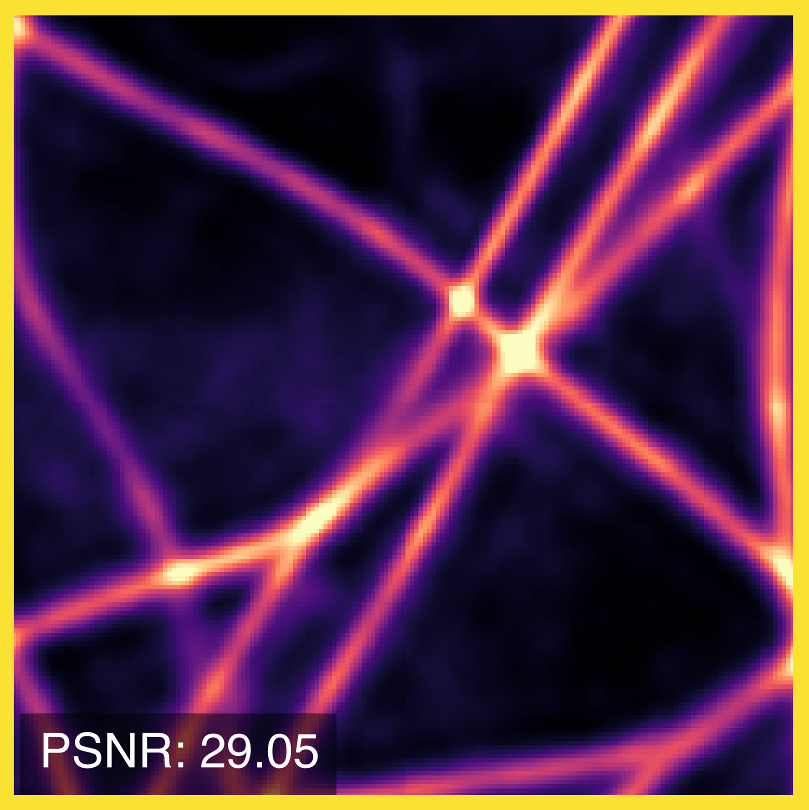

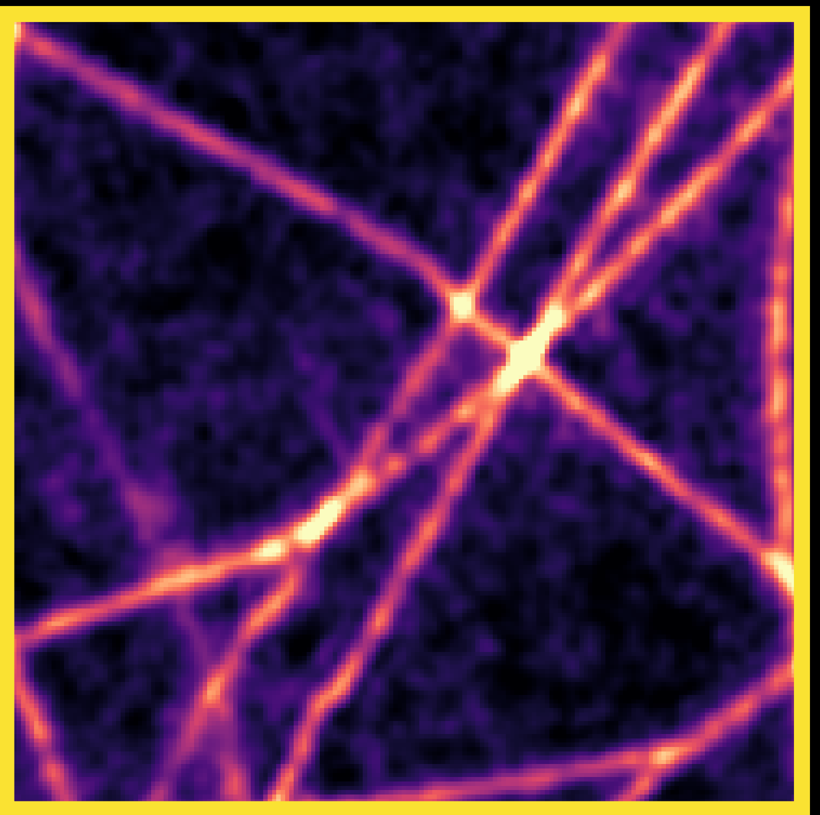

MT Data GT

MT-Noisy LR Data

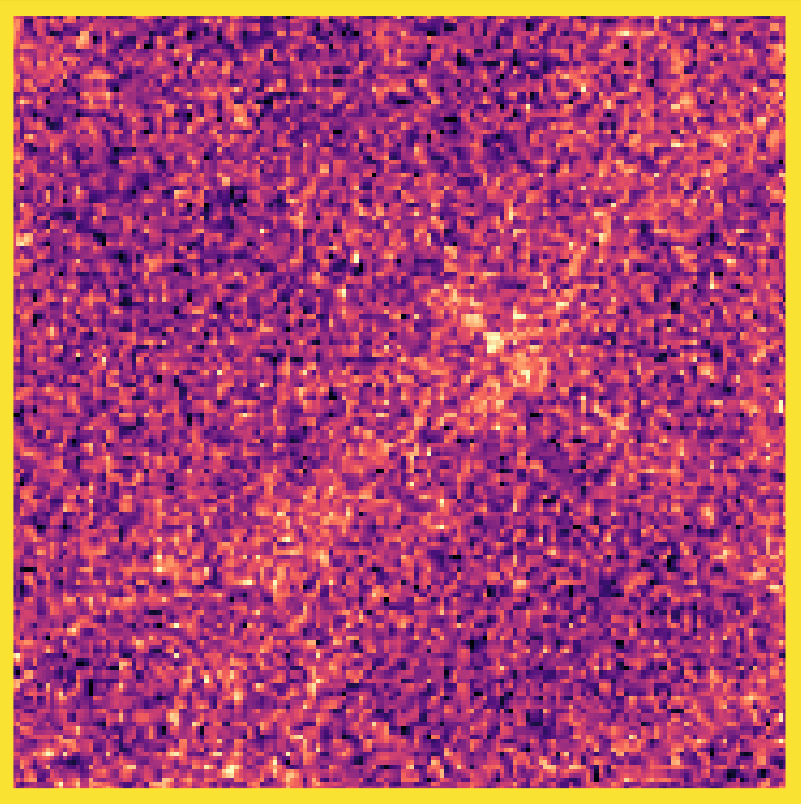

Additional Synthetic Noise

ResMatching Pred

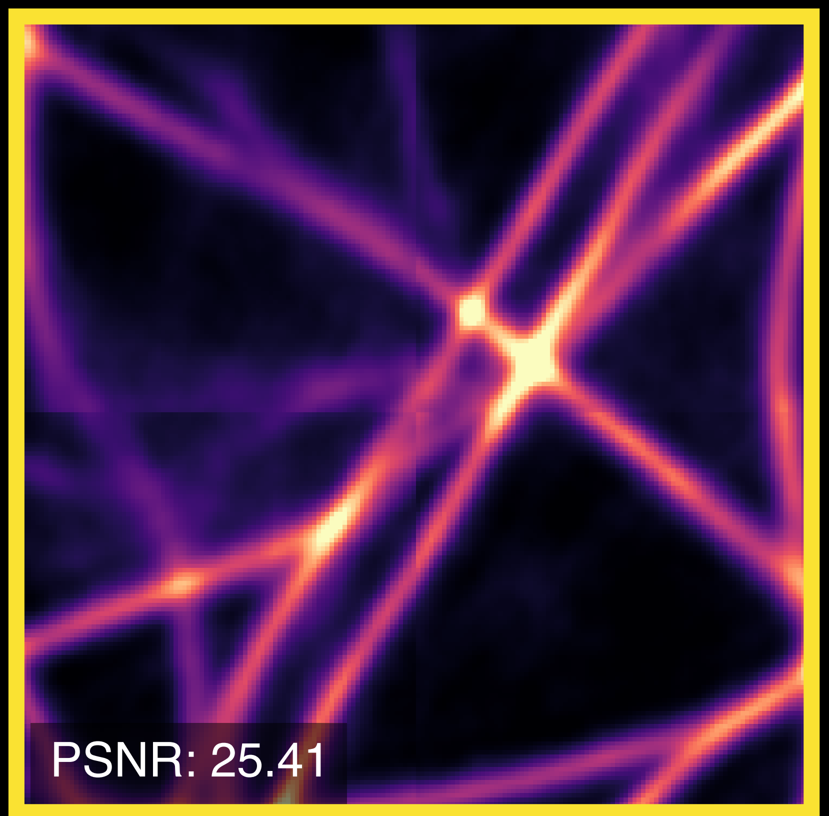

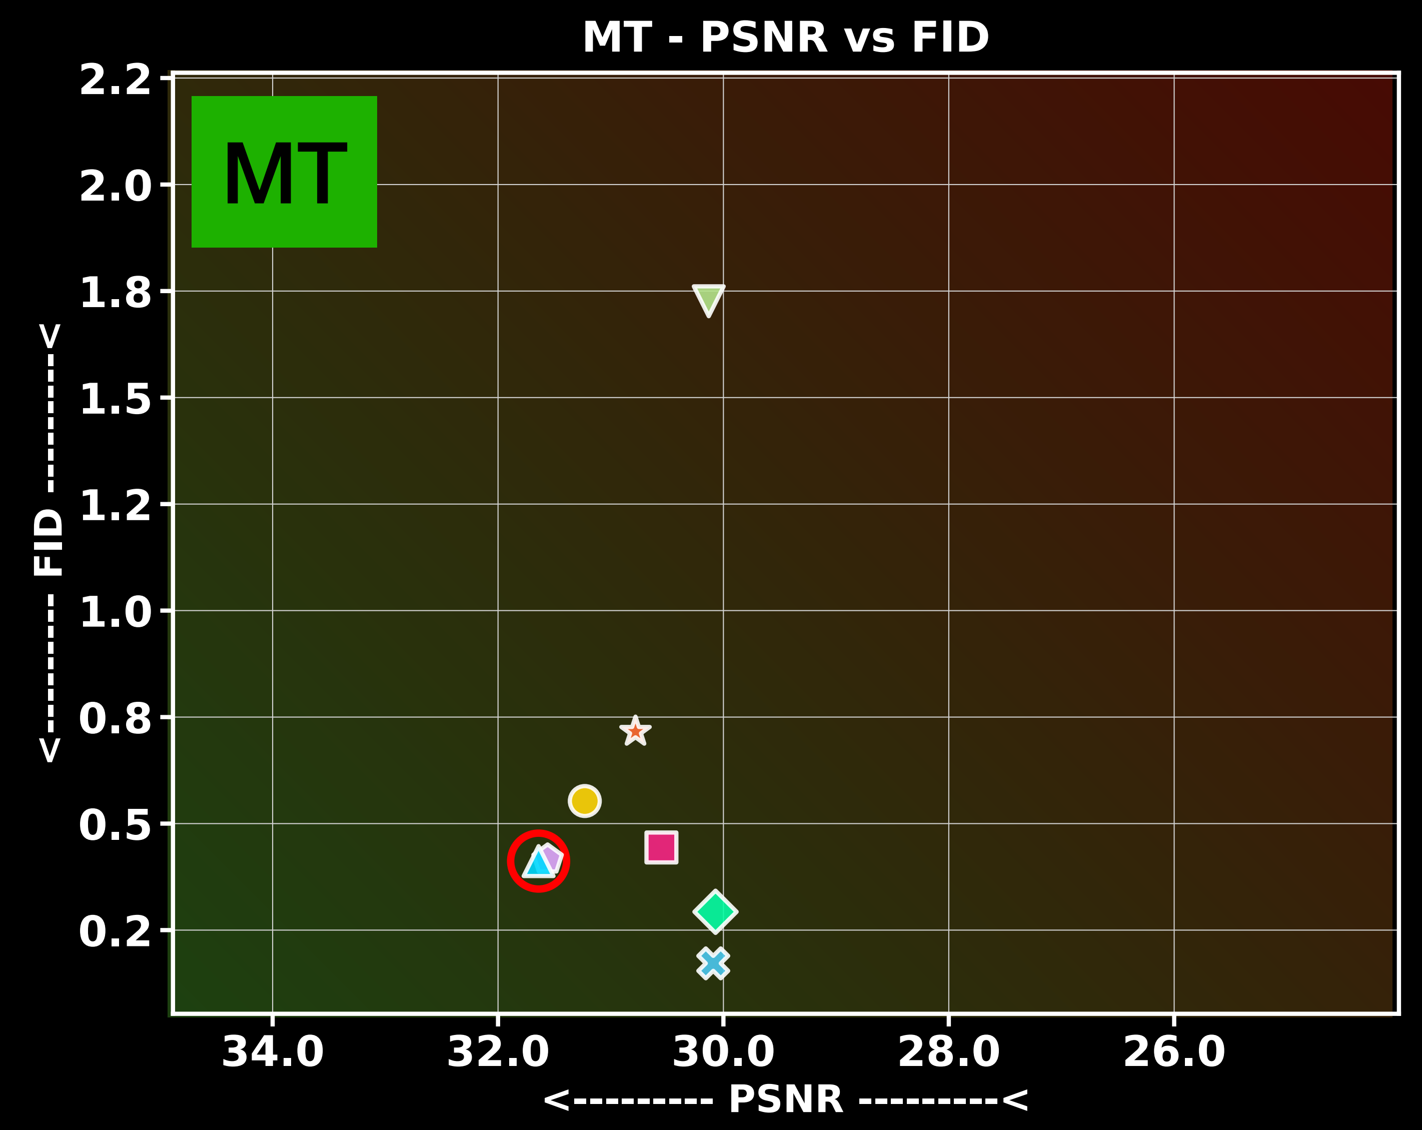

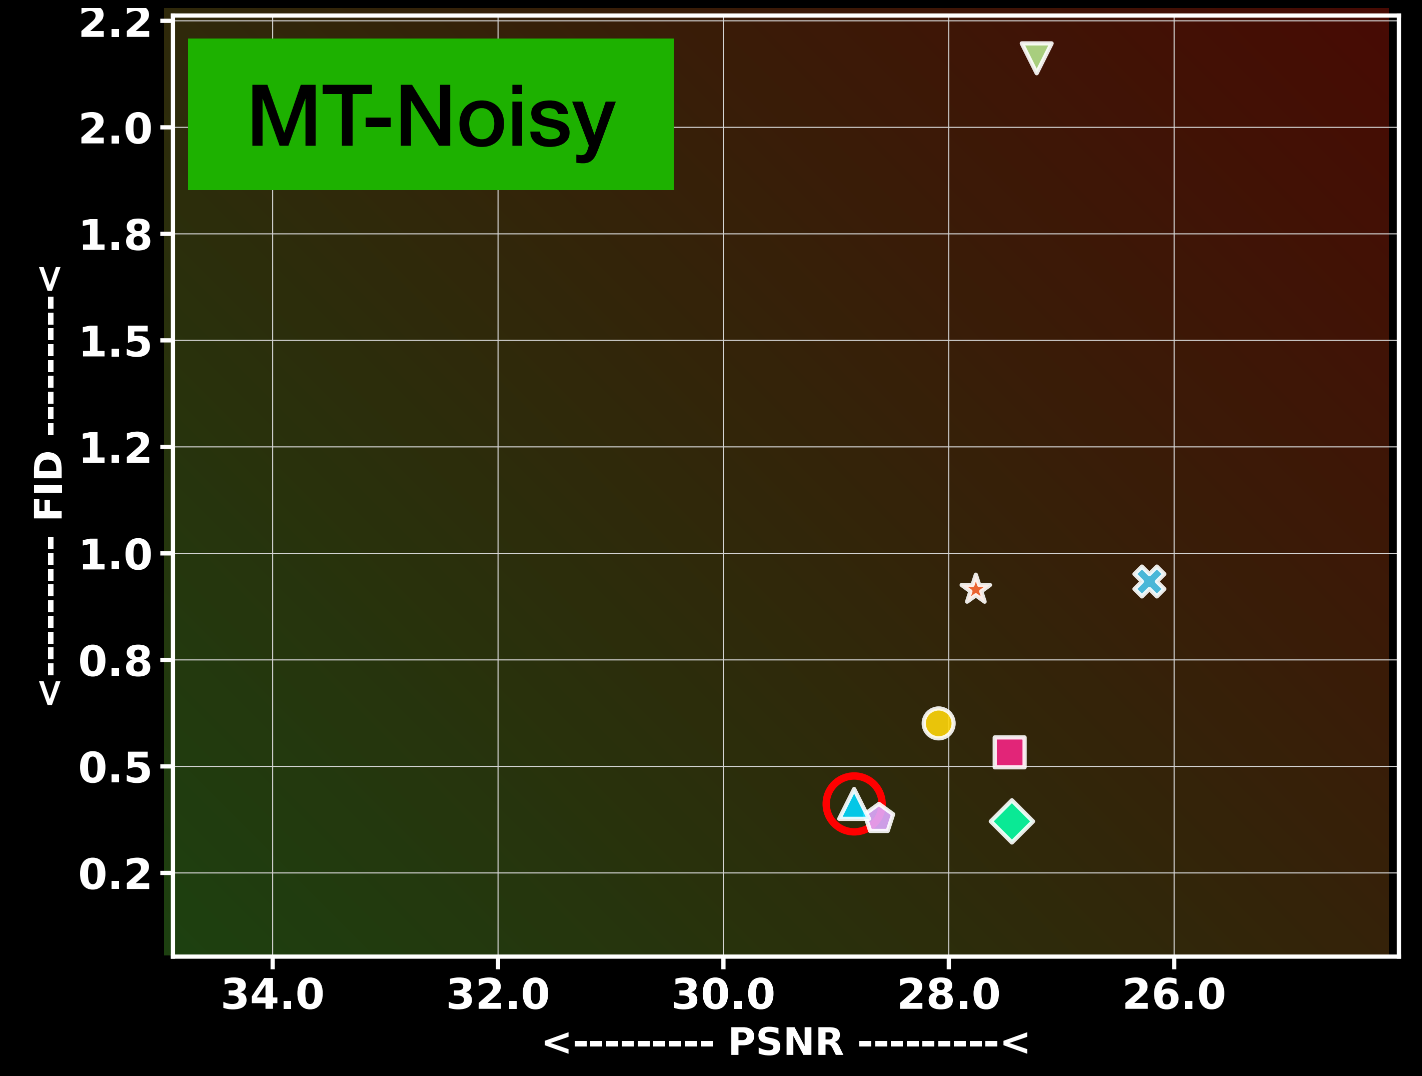

Model Calibration

ResMatching - Model Calibration

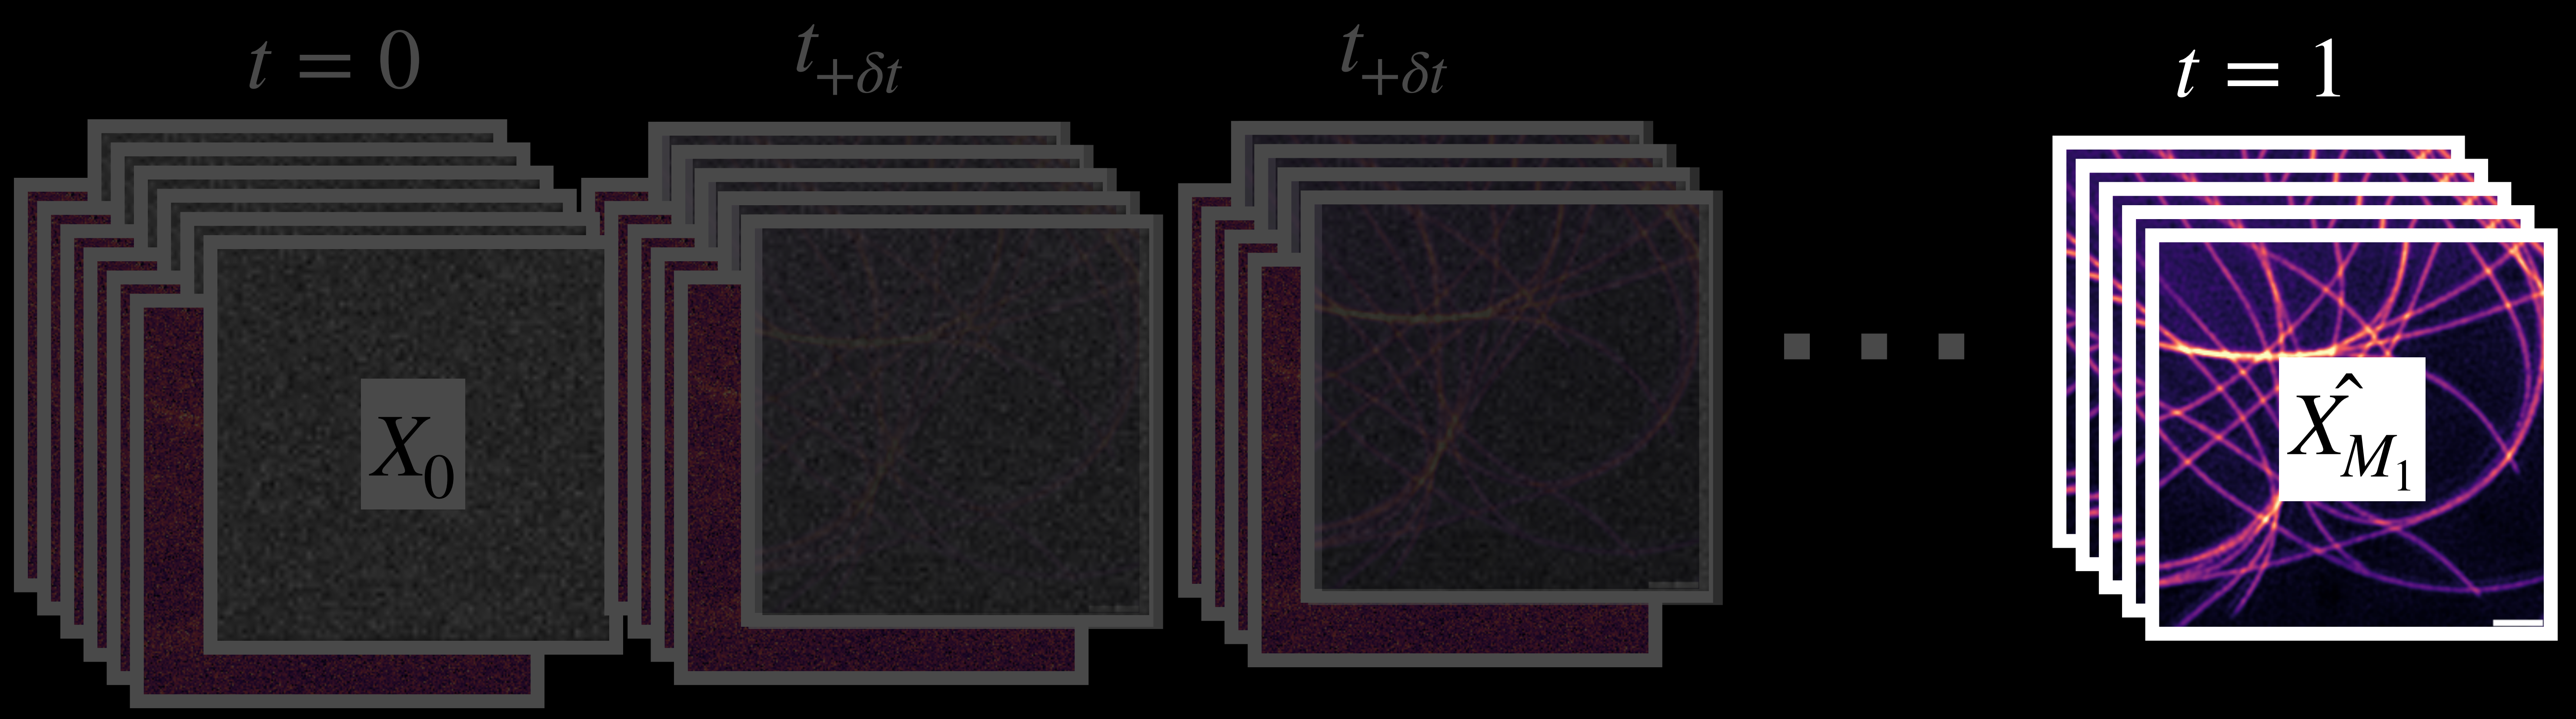

$\times K$

K posterior samples

$\frac{1}{K}\sum K$ — MMSE

$Var(K)$ — **Posterior variance**

ResMatching - Model Calibration

$t = 1$

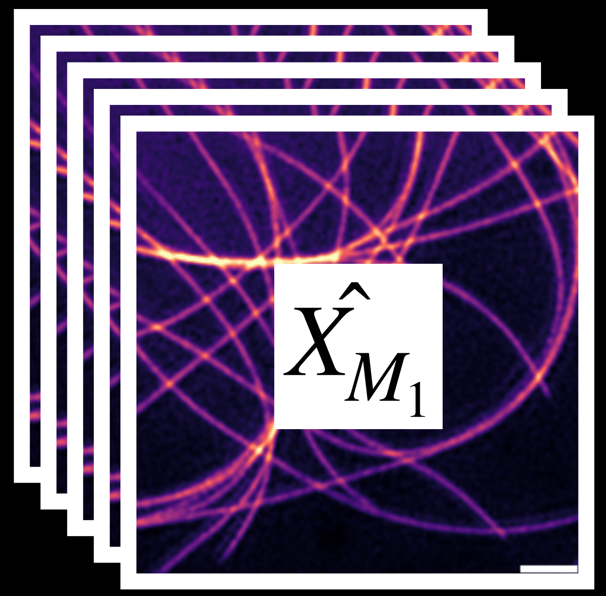

Is the ***Posterior Variance*** indicative of the Model's ***Prediction Error?***

K posterior samples

$\frac{1}{K}\sum K$

MMSE

$Var(K)$

Posterior variance

ResMatching - Model Calibration

ResMatching - Model Calibration

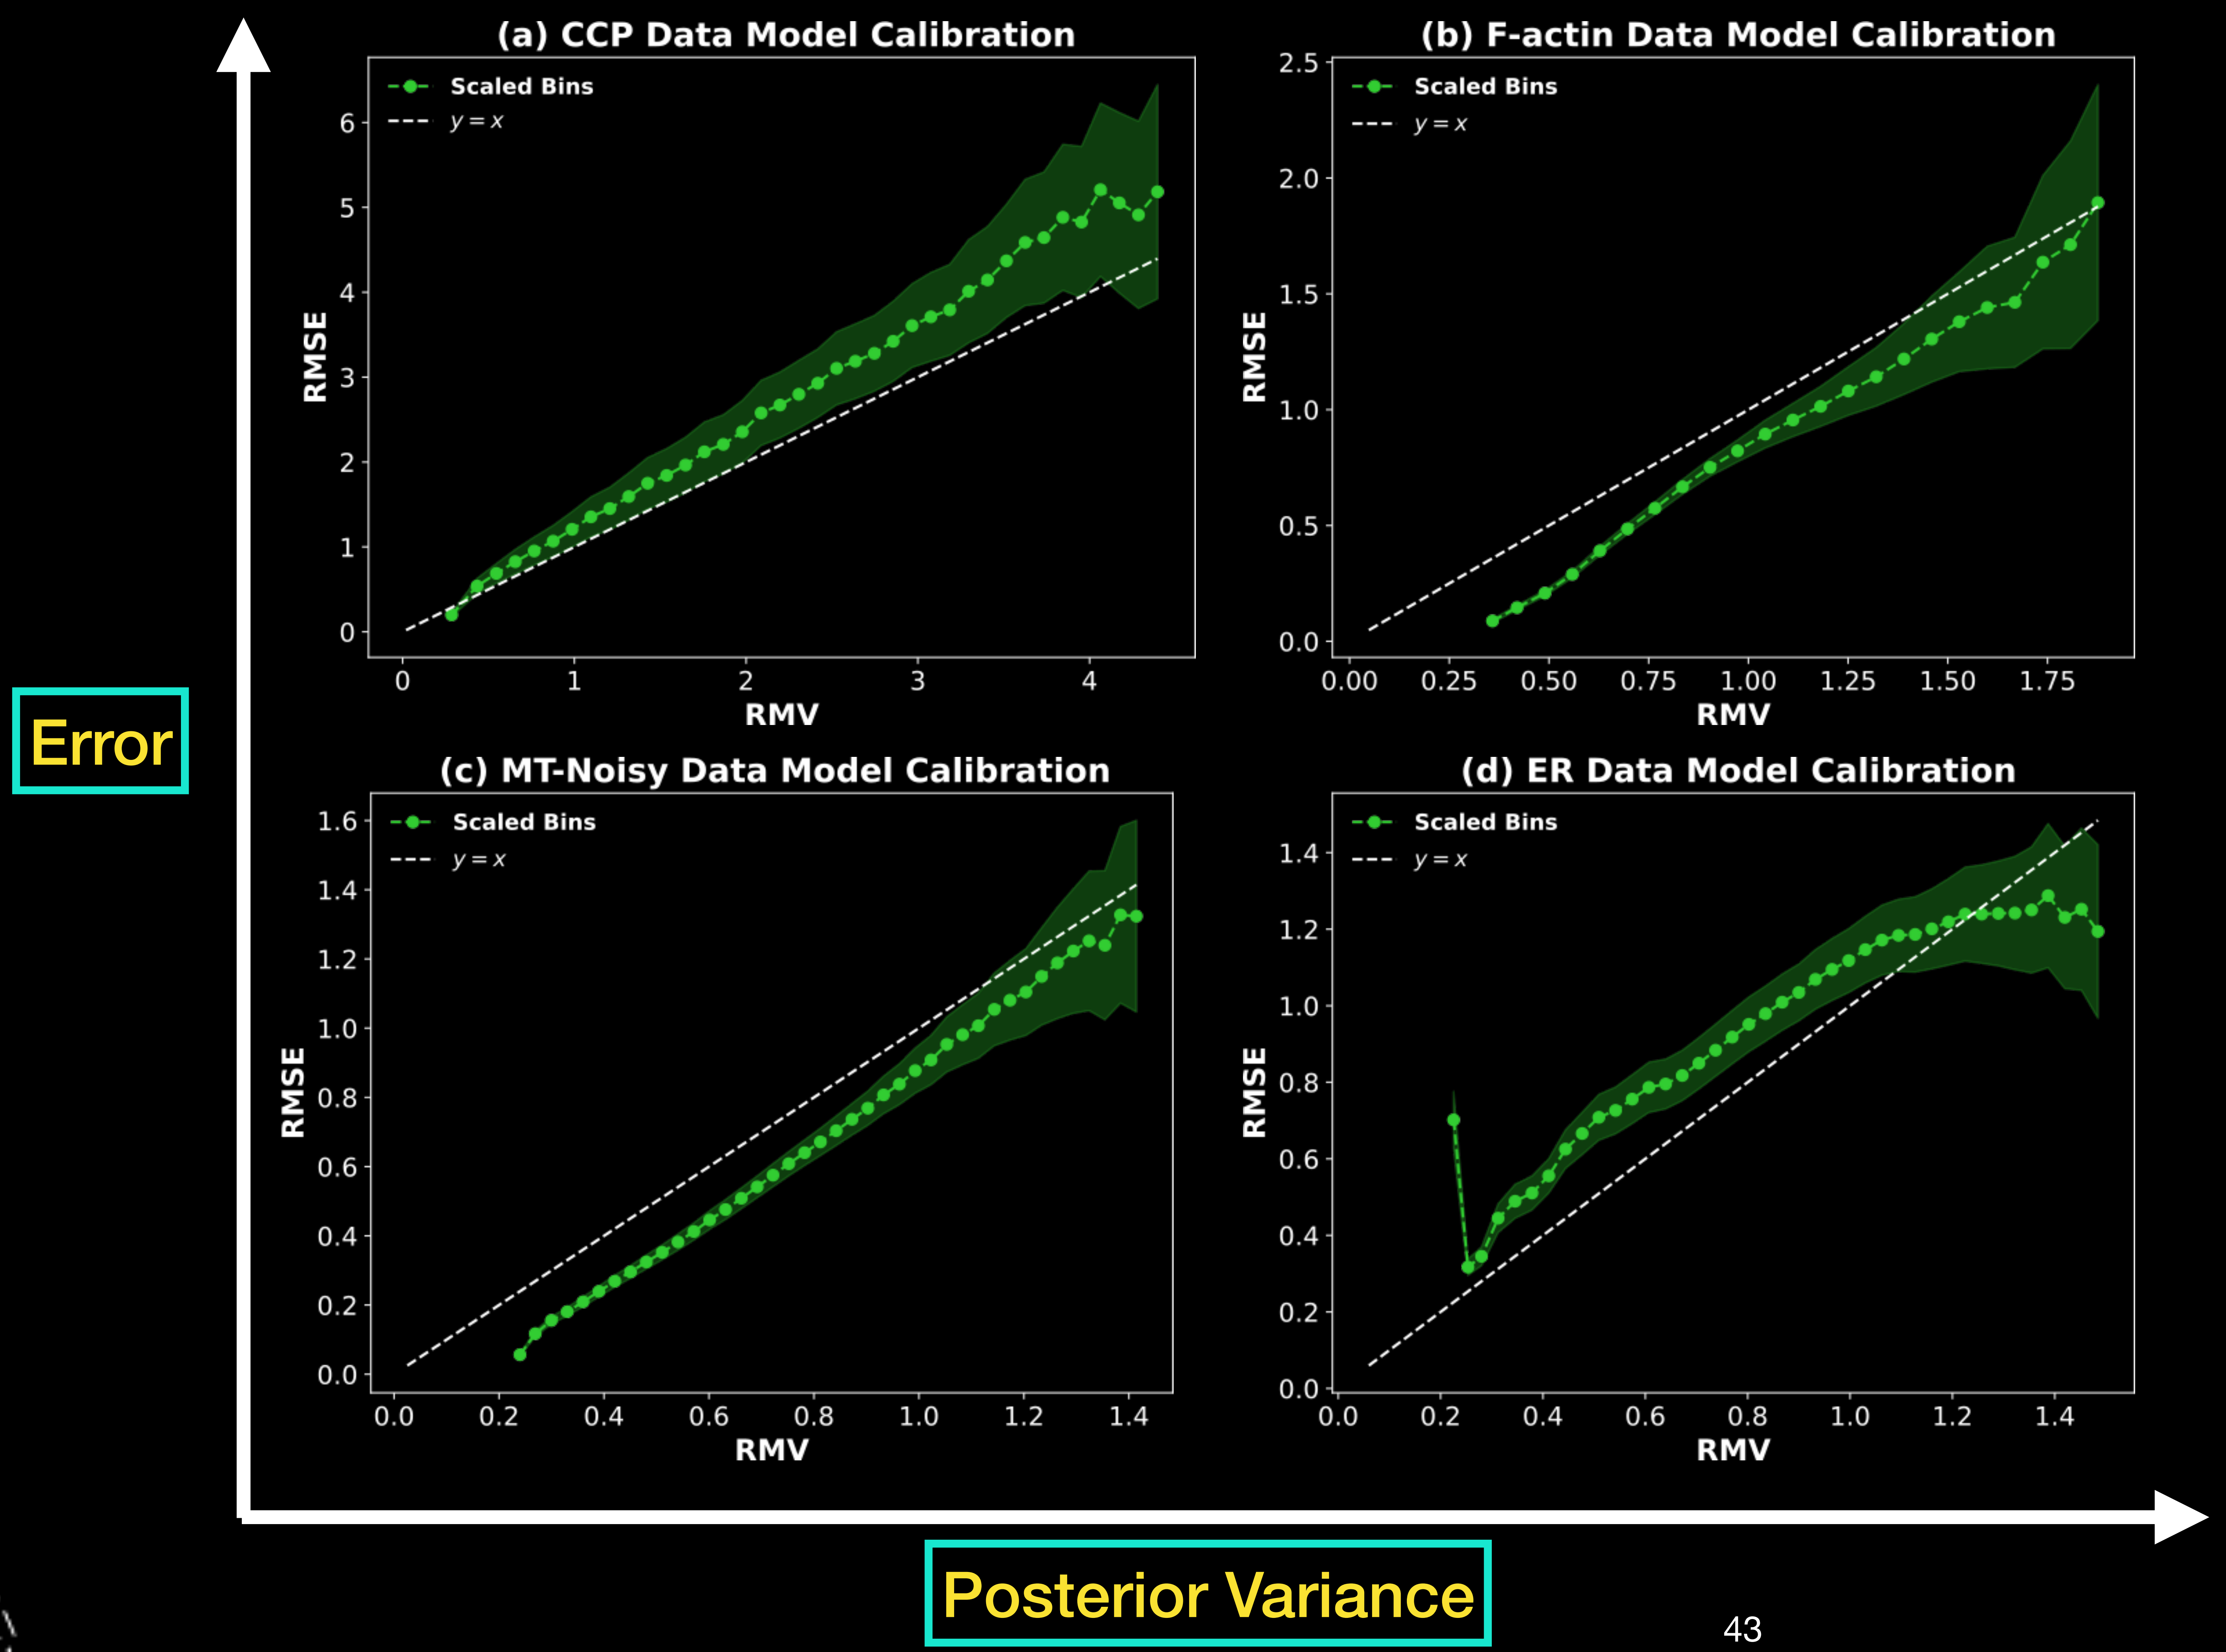

The models are well calibrated: Errors scales ~linearly with posterior variance

trained posterior's uncertainty is in-line with the uncertainty in the data!

Posterior variance can potentially be a surrogate for model's error when no Ground Truth is available

Summary

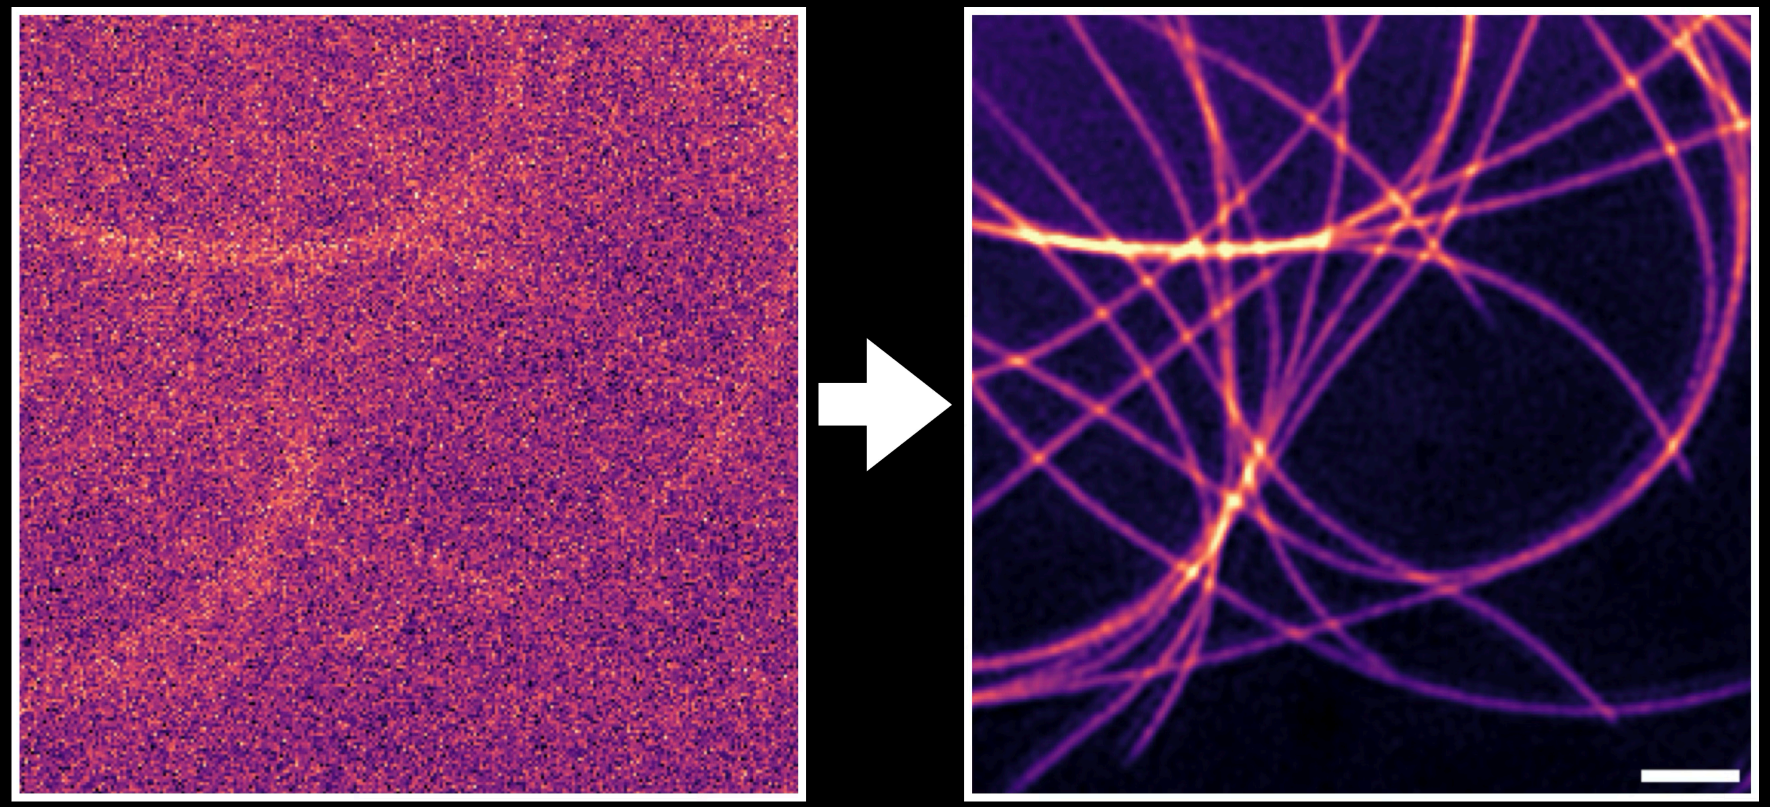

Summary

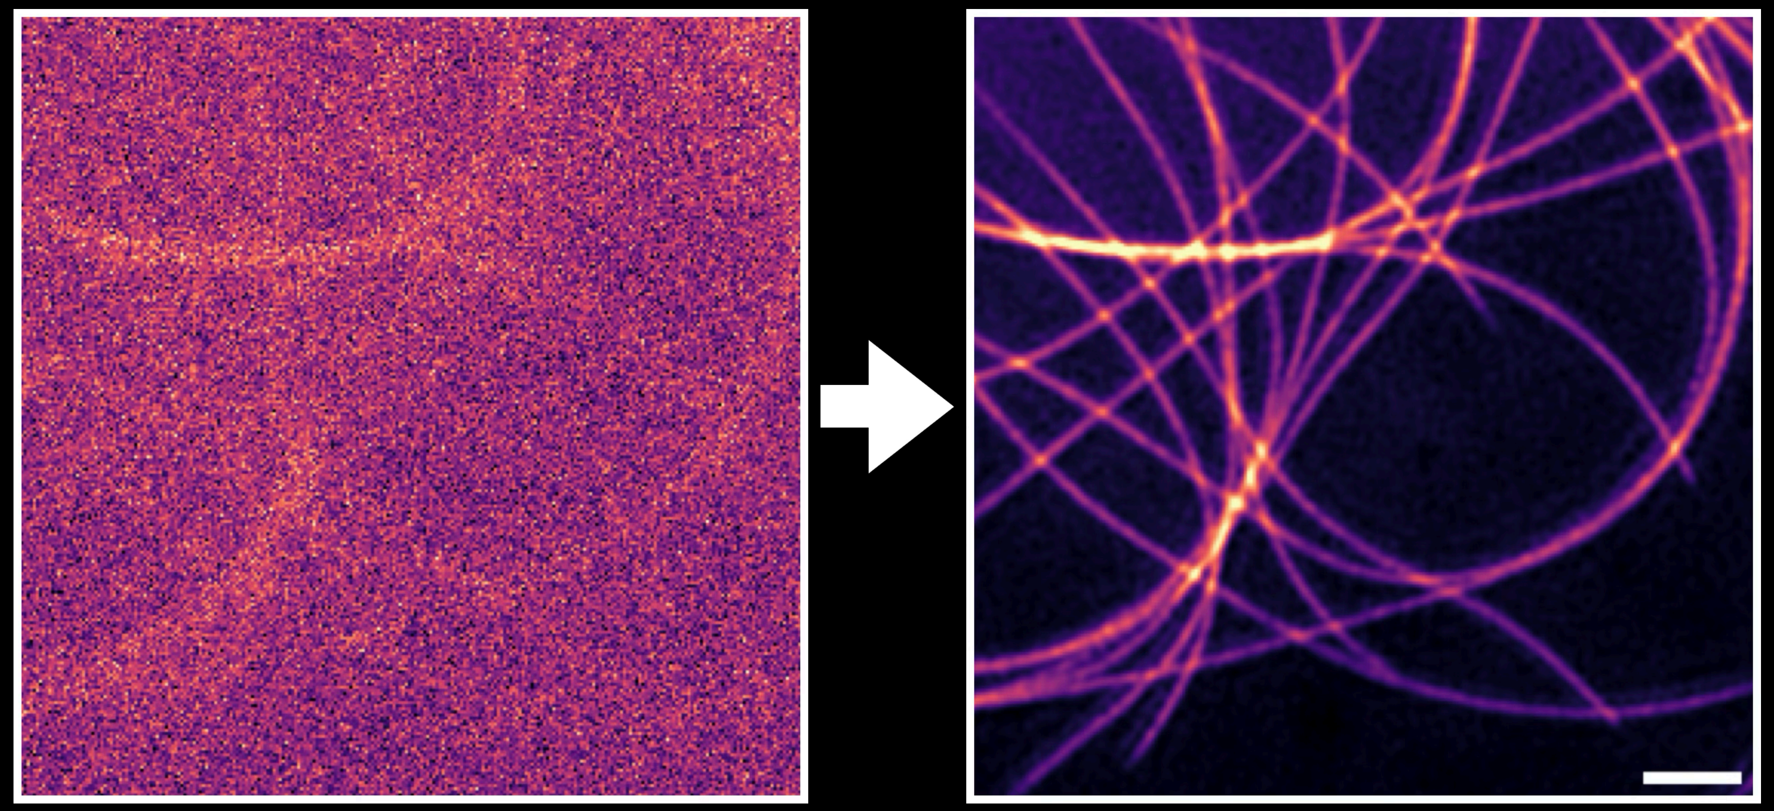

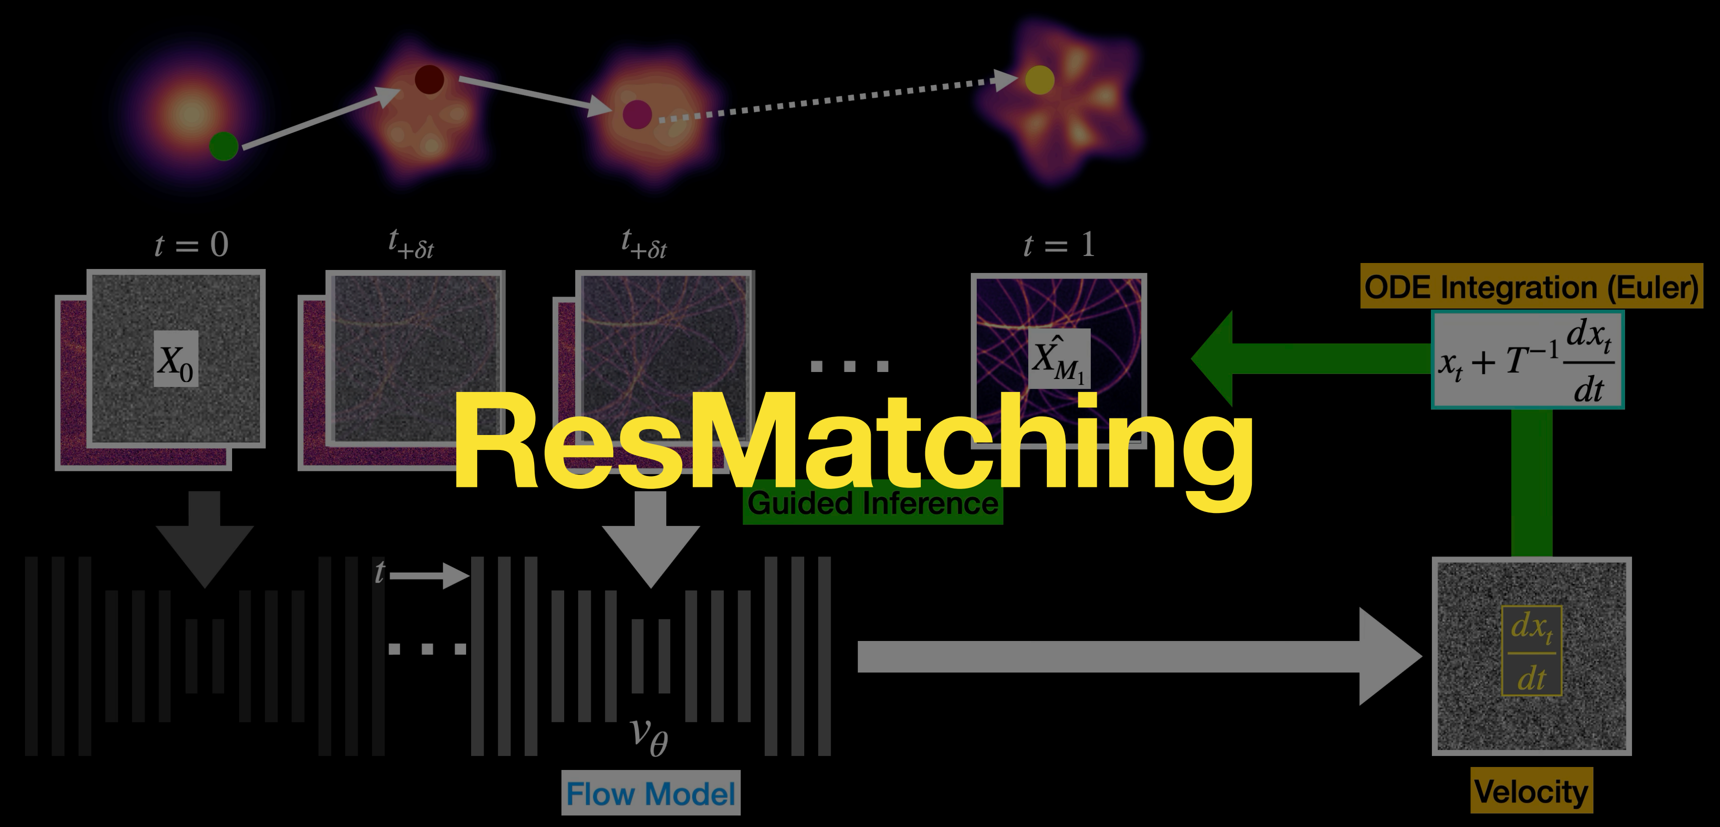
ResMatching
$t=0$
$t_{+\delta t}$
$t_{+\delta t}$
$t=1$
X_0
$\hat{X}_{M_1}$
ODE Integration (Euler)
$x_t + T^{-1}\frac{dx_t}{dt}$
Guided Inference
v_θ
Flow Model
$\frac{dx_t}{dt}$
Velocity

Summary

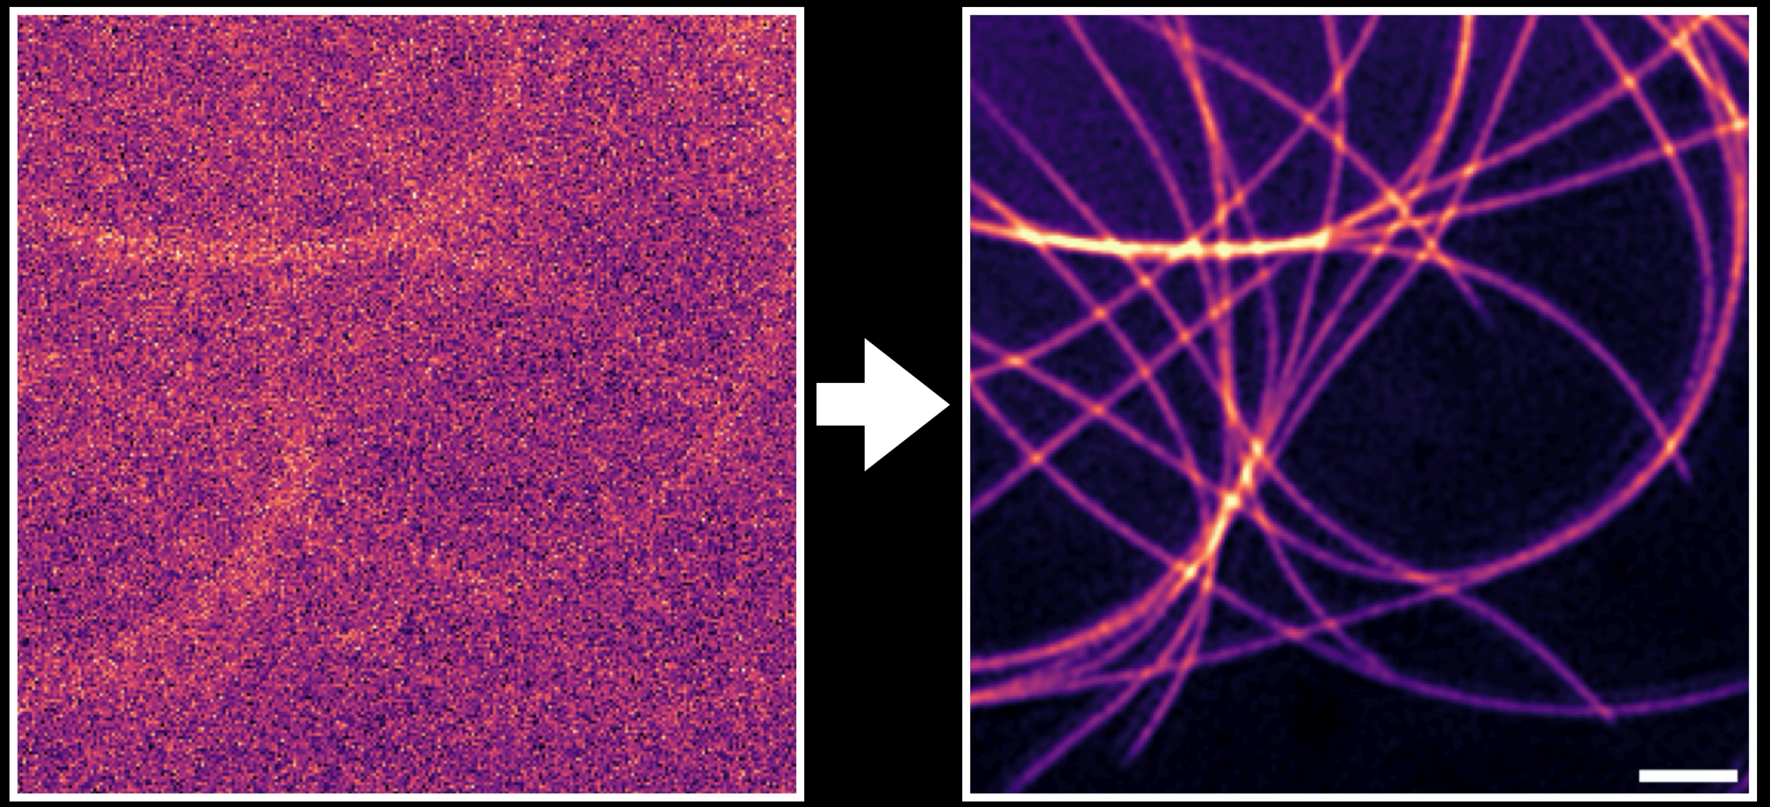

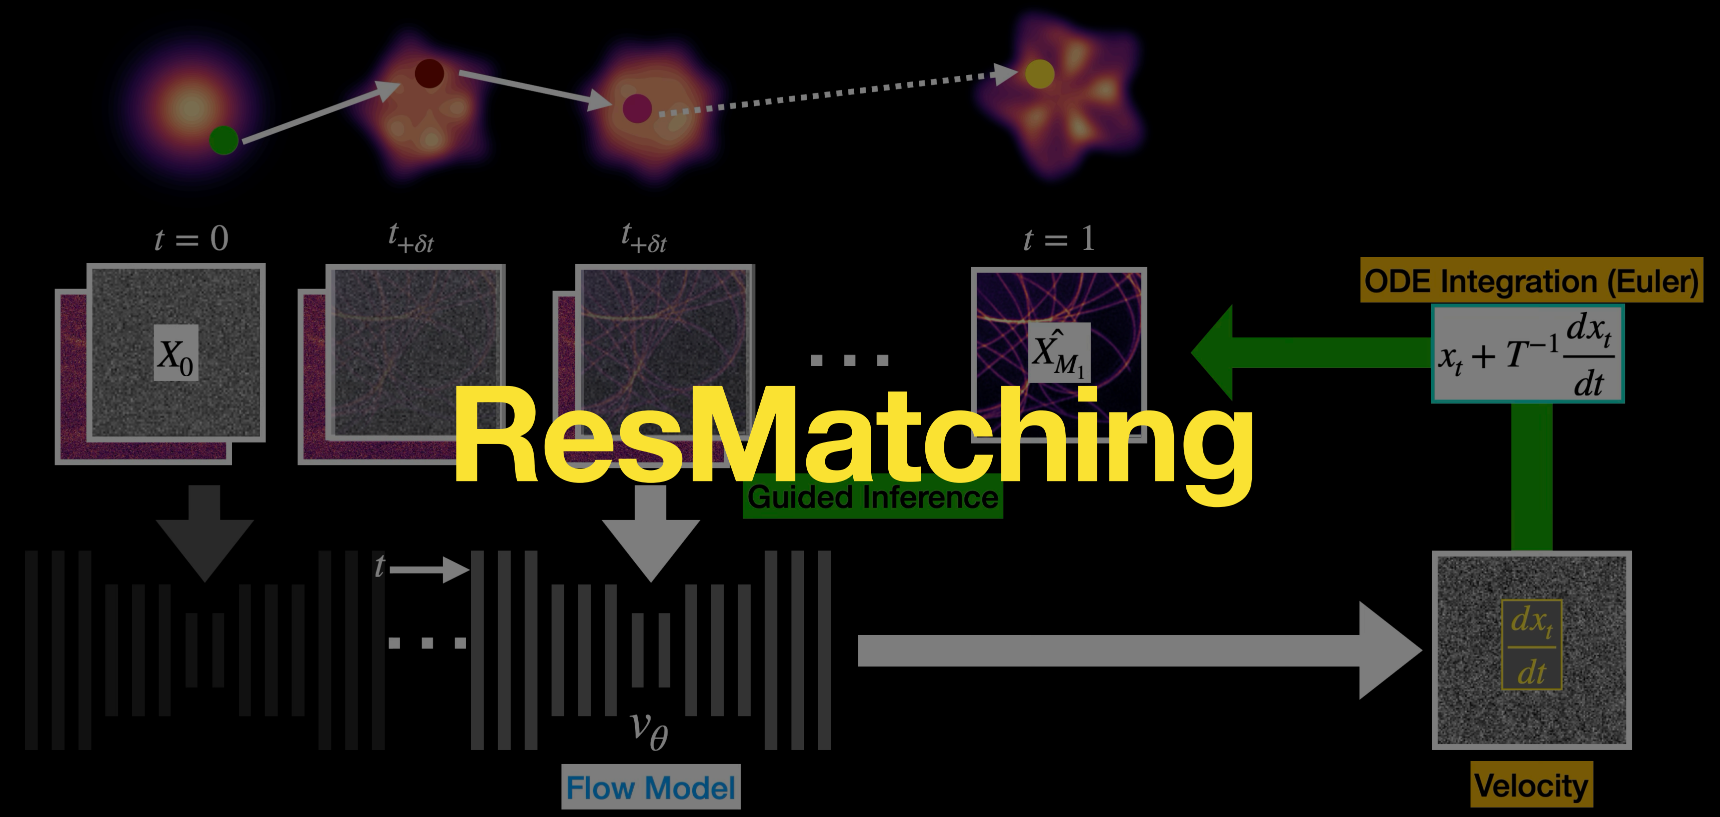
$t = 0$
$t_{+\delta t}$
$t_{+\delta t}$
$t = 1$
X_0
$\hat{X}_{M_1}$
ODE Integration (Euler)
$x_t + T^{-1}\frac{dx_t}{dt}$
ResMatching
Guided Inference
v_θ
Flow Model
$\frac{dx_t}{dt}$
Velocity

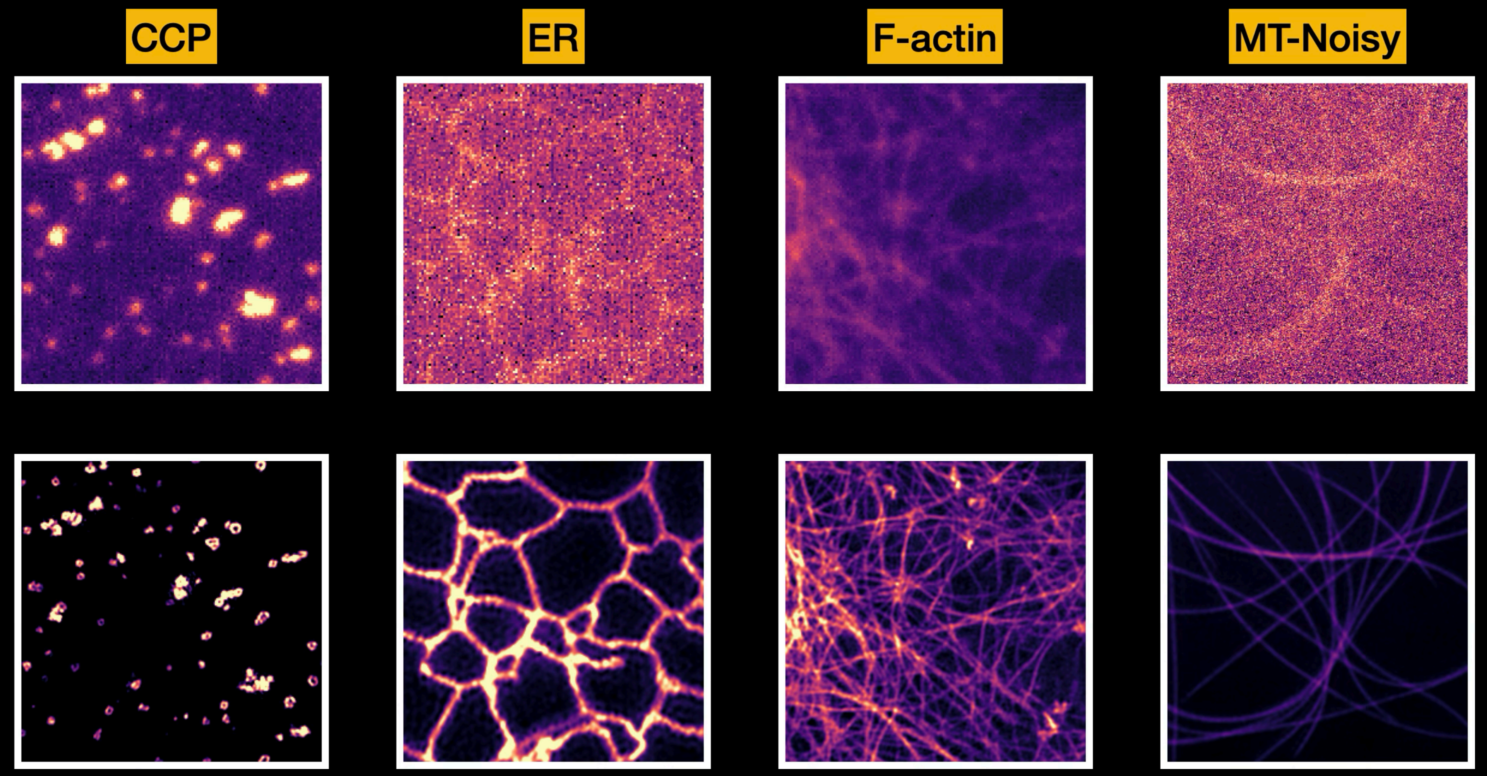
CCP
ER
F-actin
MT-Noisy

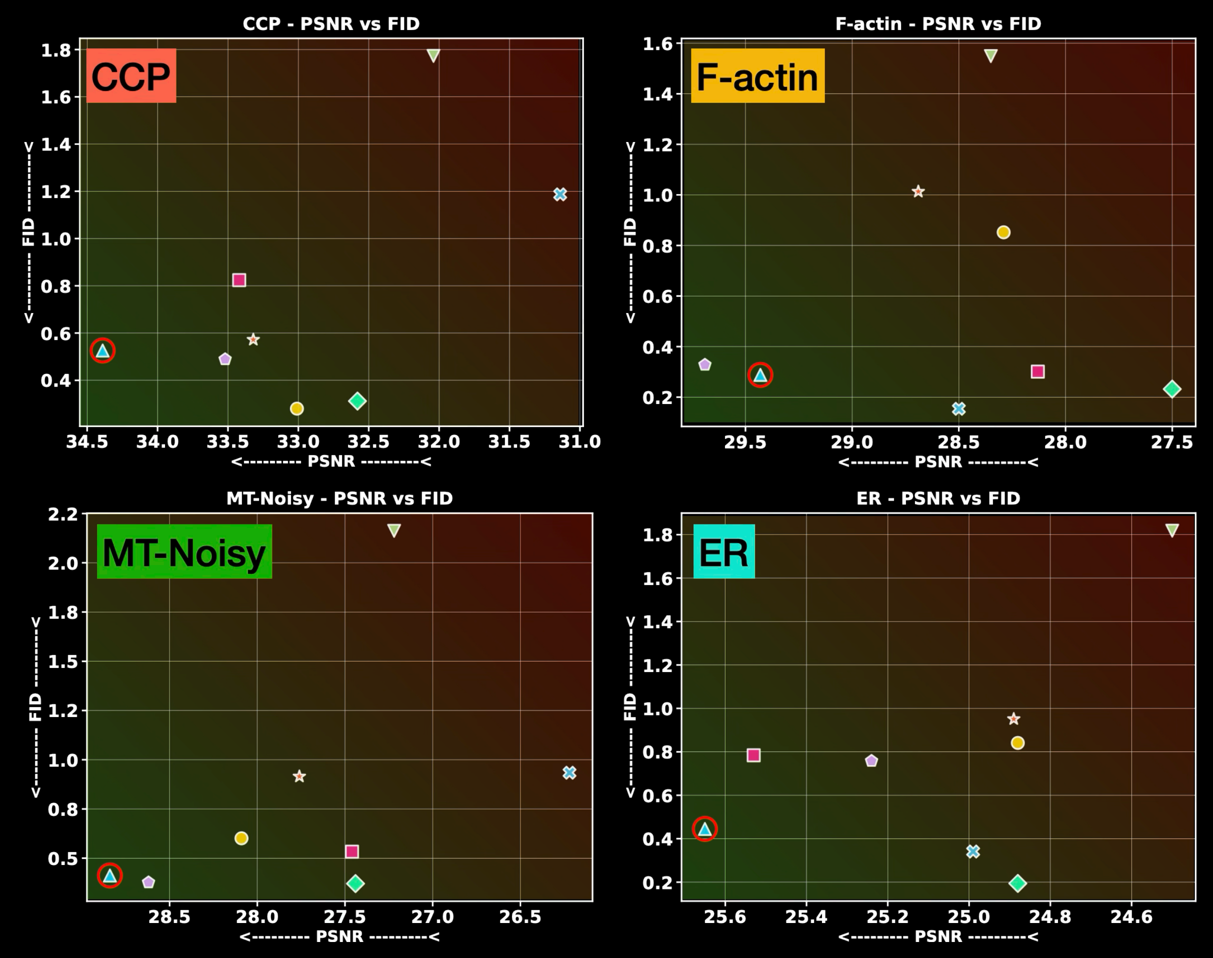
CCP - PSNR vs FID
F-actin - PSNR vs FID
MT-Noisy - PSNR vs FID
ER - PSNR vs FID
CCP
F-actin
MT-Noisy
ER
FID
PSNR

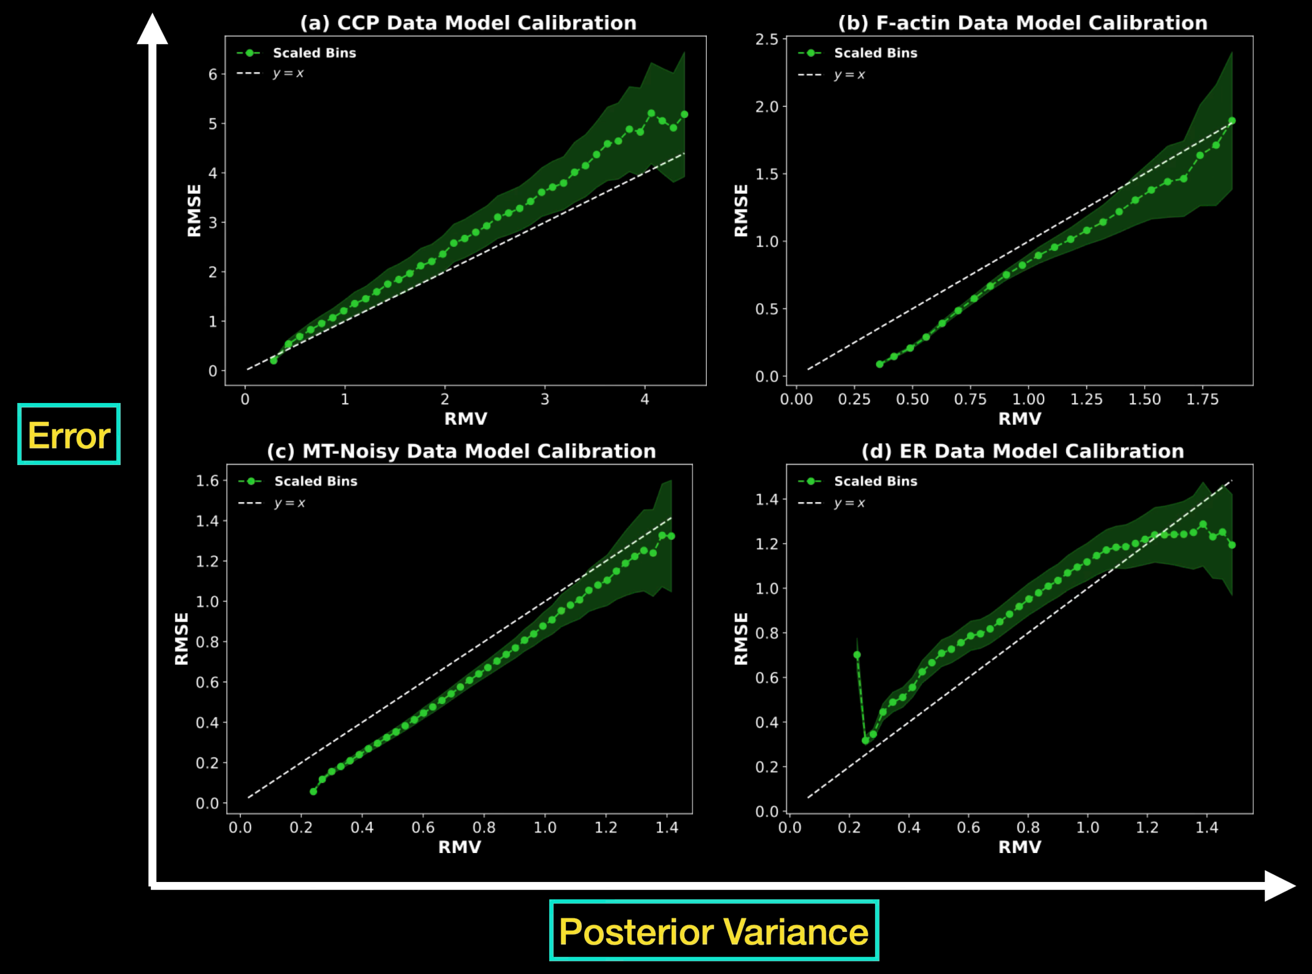
Error
Posterior Variance
(a) CCP Data Model Calibration
(b) F-actin Data Model Calibration
(c) MT-Noisy Data Model Calibration
(d) ER Data Model Calibration
Scaled Bins
y = x
RMSE
RMV

A Cautionary note

Image Dehazing*

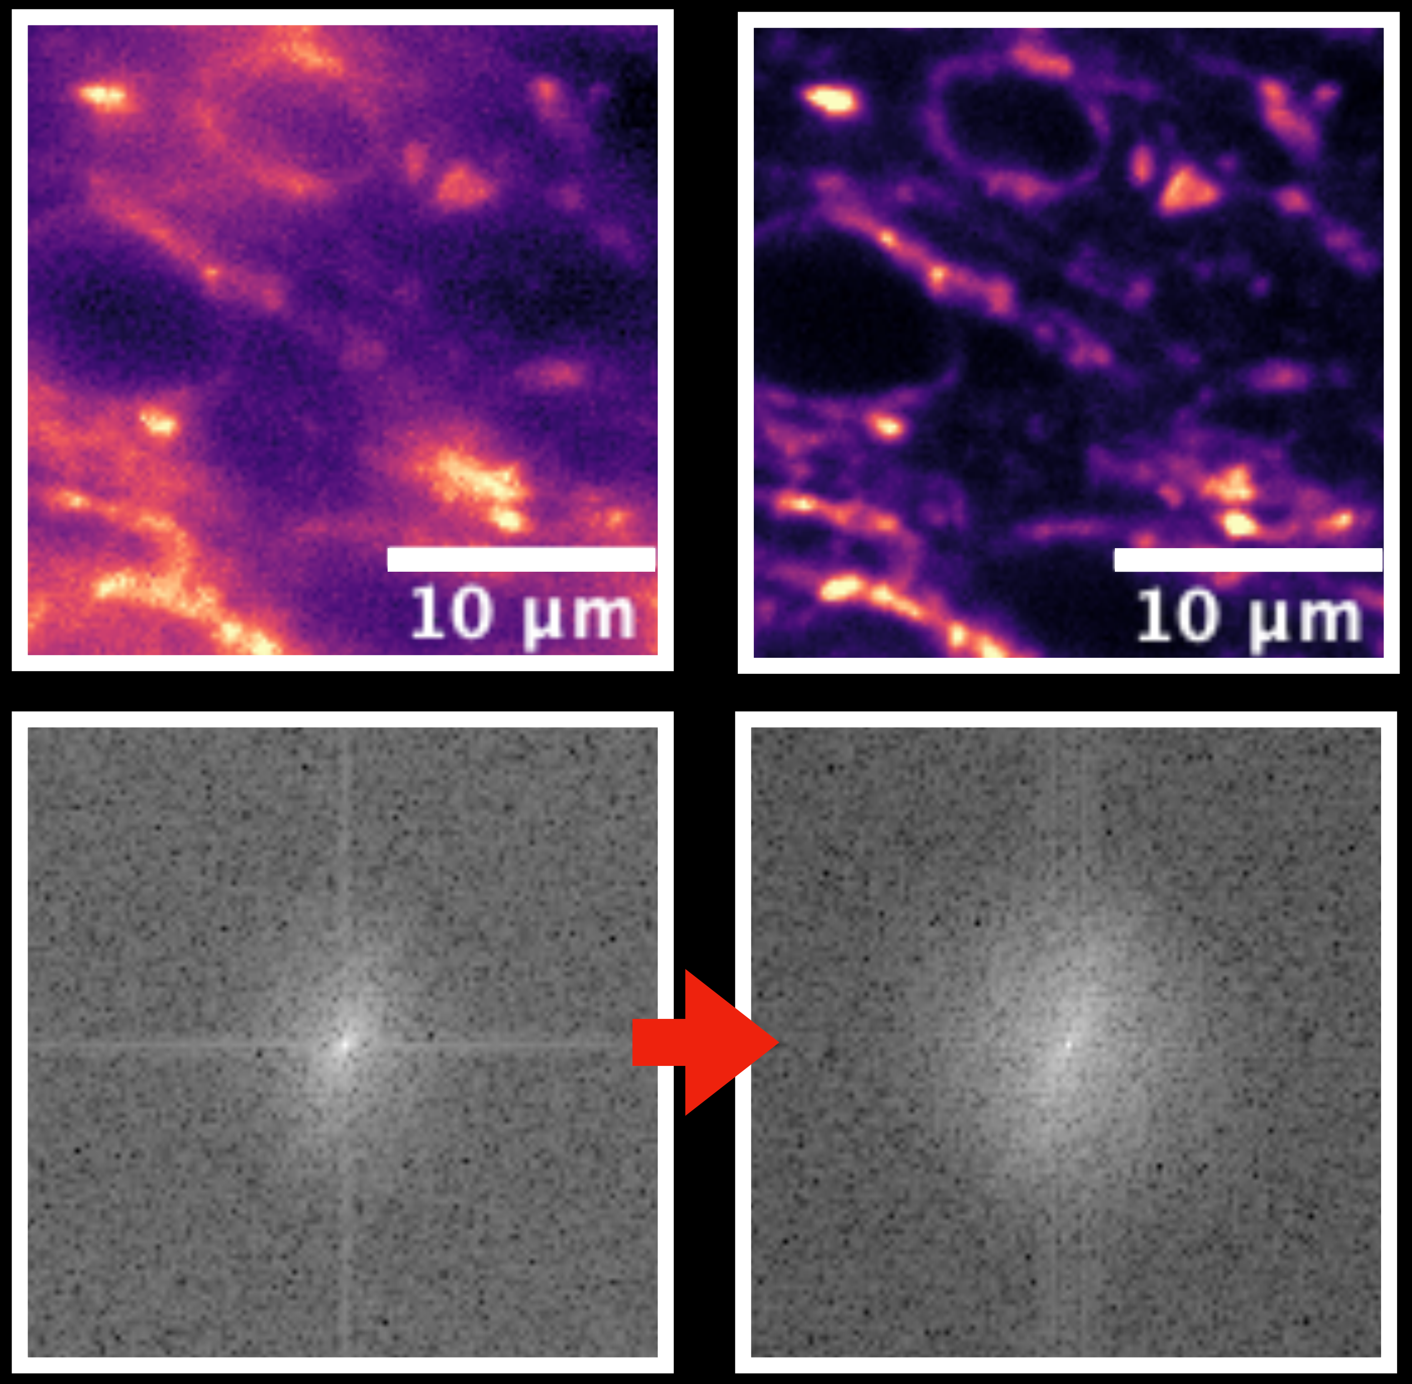

Widefield Image | Confocal Image

*HazeMatching, Ray et al. - CVPR 2026 (Findings)

A Cautionary note

Image Dehazing*

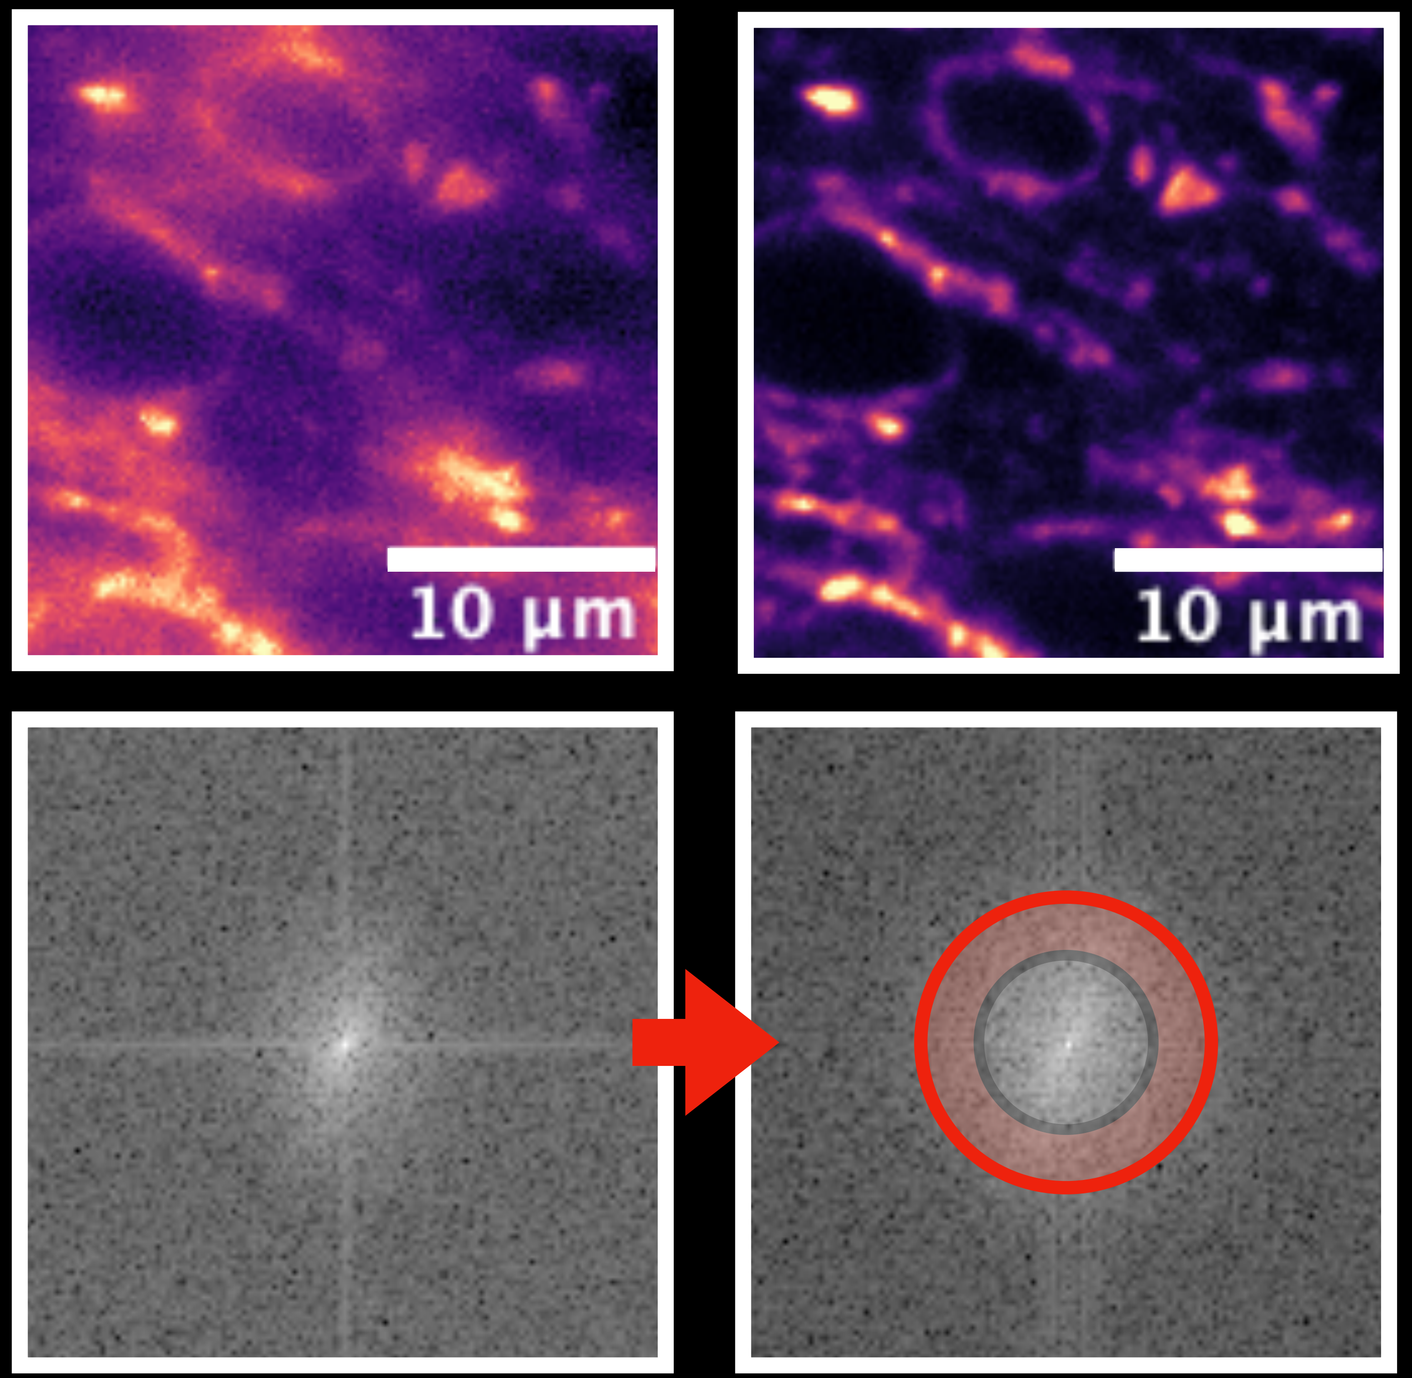

Widefield Image

Confocal Image

Manipulation of observable frequency ranges

Interpolation

*HazeMatching, Ray et al. - CVPR 2026 (Findings)

A Cautionary note

Image Dehazing* | **CSR**

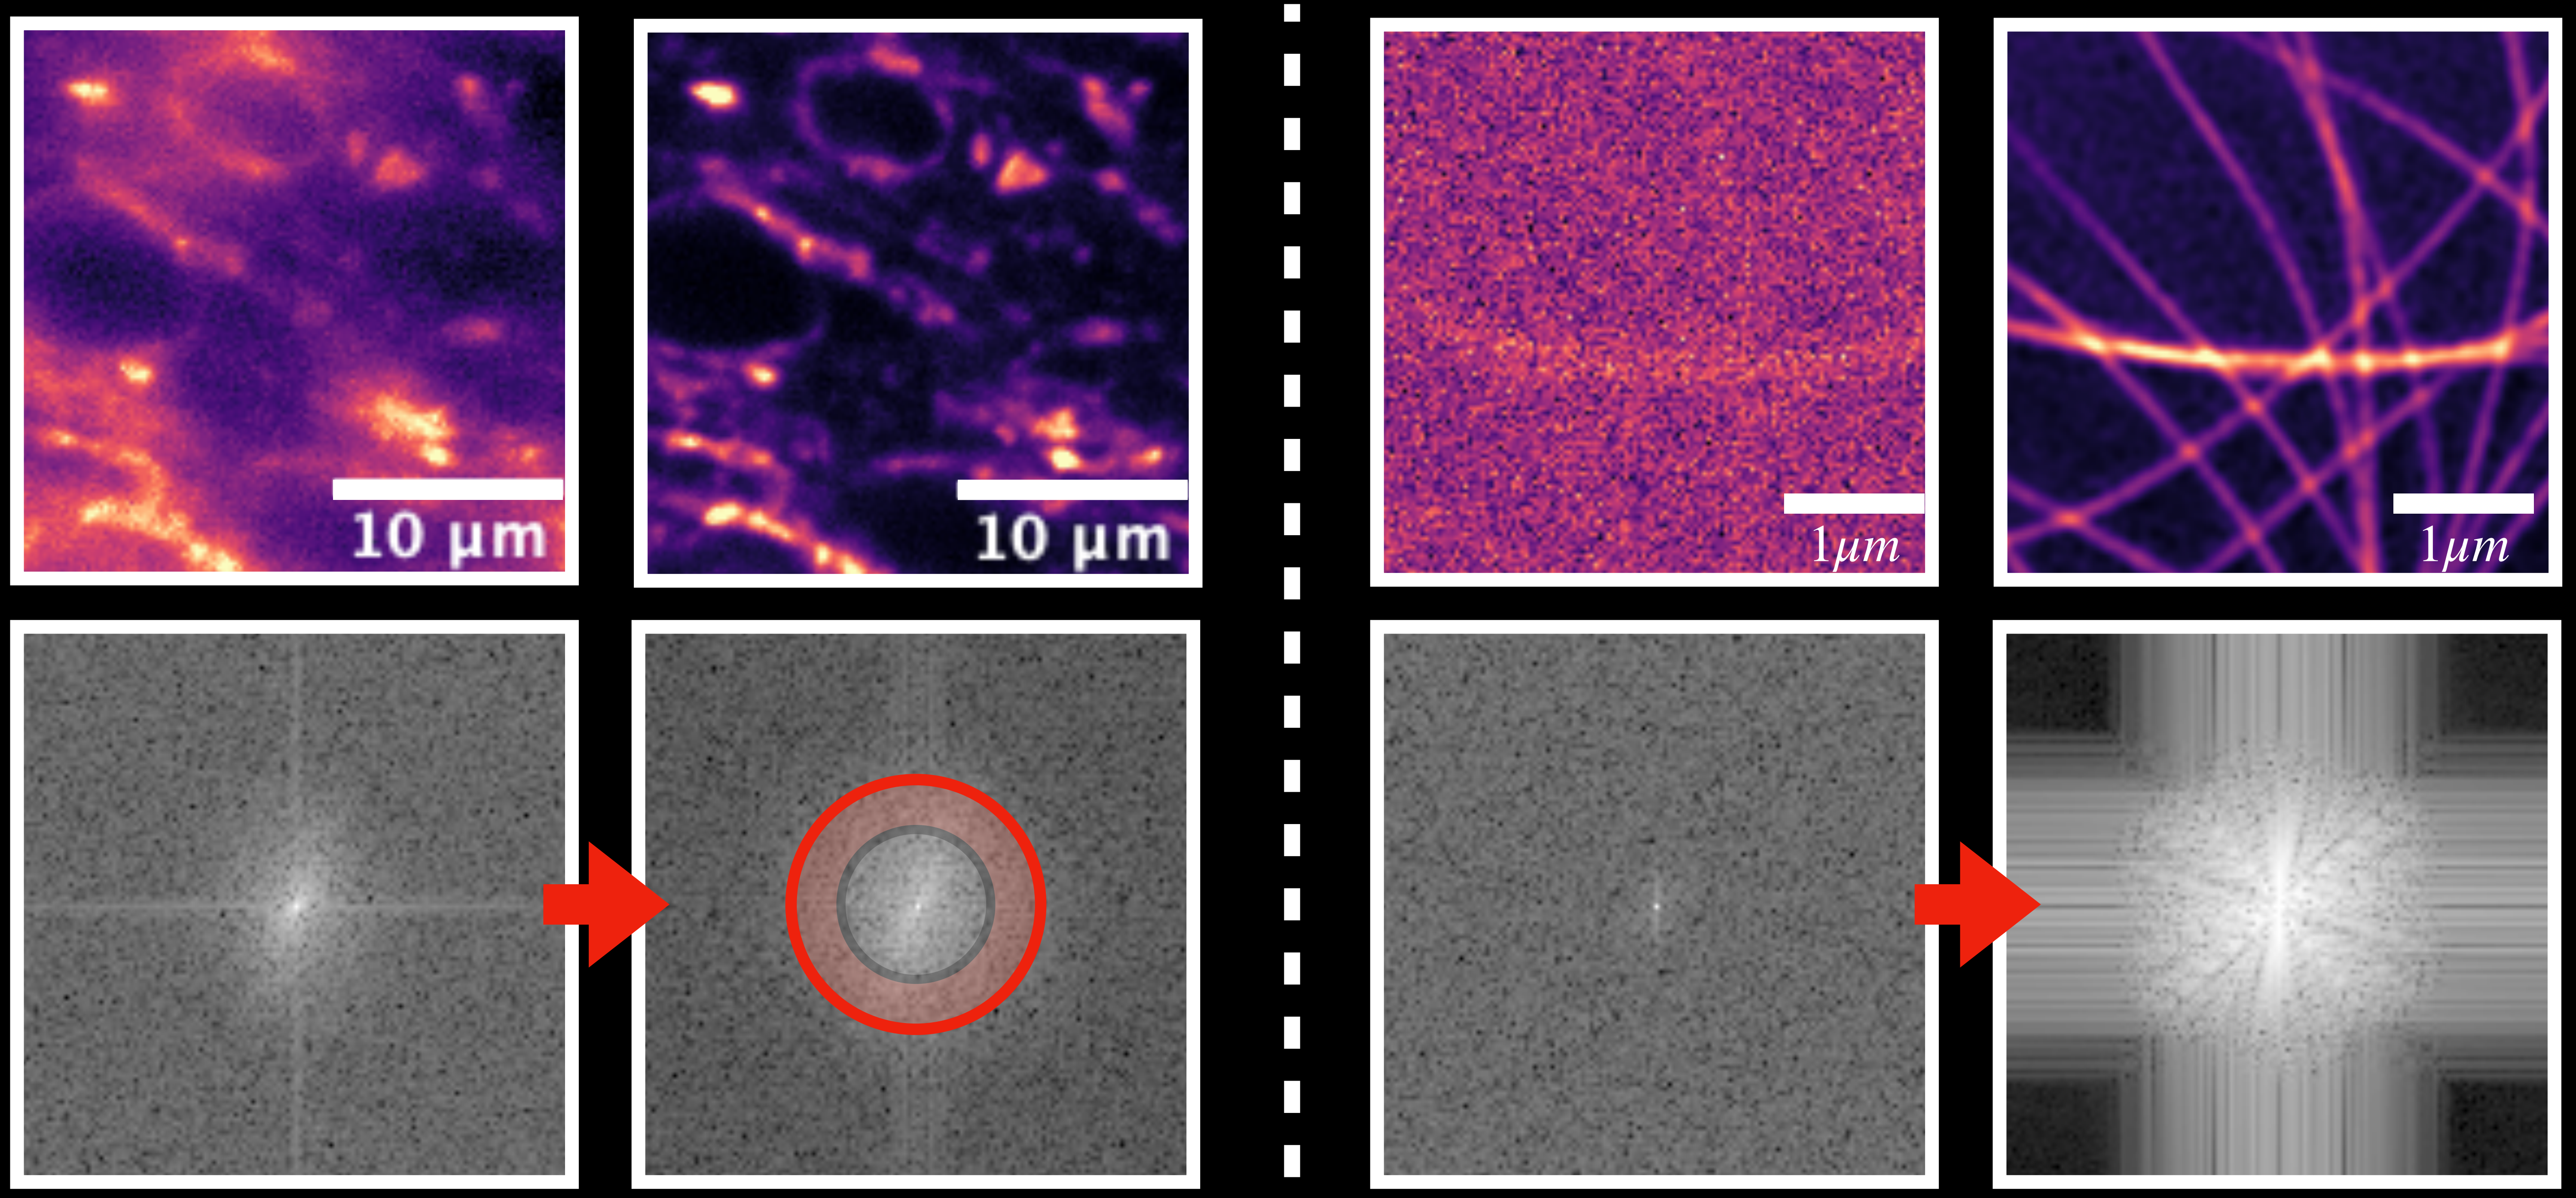

Widefield Image | Confocal Image | Noisy LR Image | SR Image

Manipulation of observable frequency ranges

Interpolation

*HazeMatching, Ray et al. - CVPR 2026 (Findings)

A Cautionary note

Image Dehazing* | **CSR**

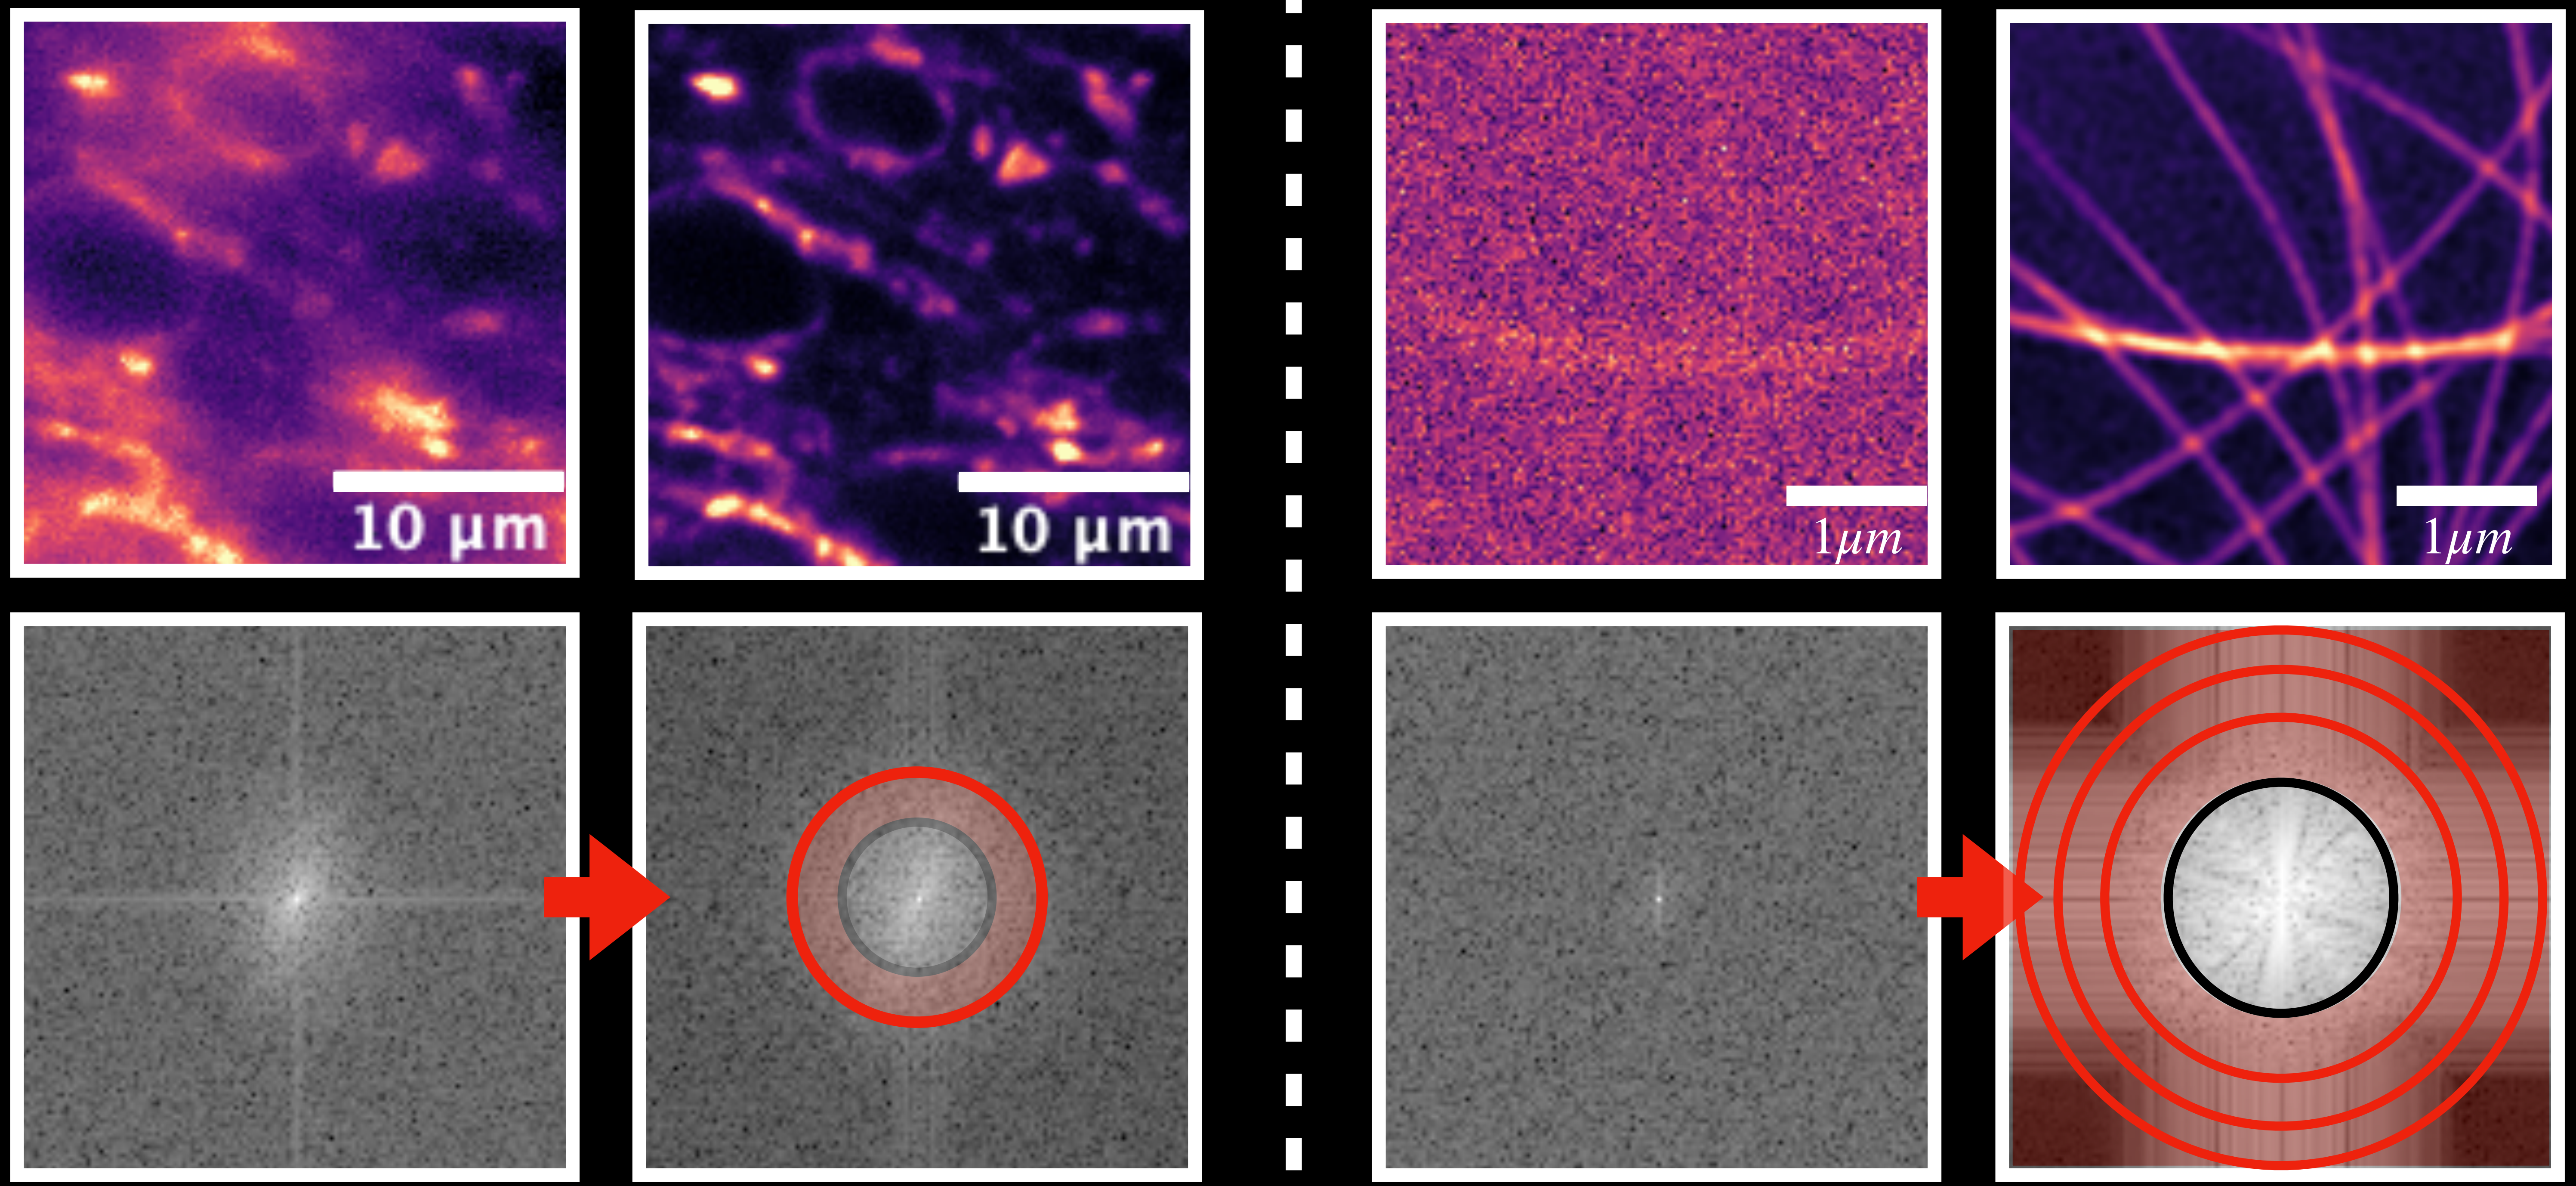

Widefield Image | Confocal Image | Noisy LR Image | SR Image

Manipulation of observable frequency ranges

Interpolation

Extrapolation to non-observed frequency ranges

Extrapolation

*HazeMatching, Ray et al. - CVPR 2026 (Findings)

A Cautionary note

Image Dehazing* | **CSR**

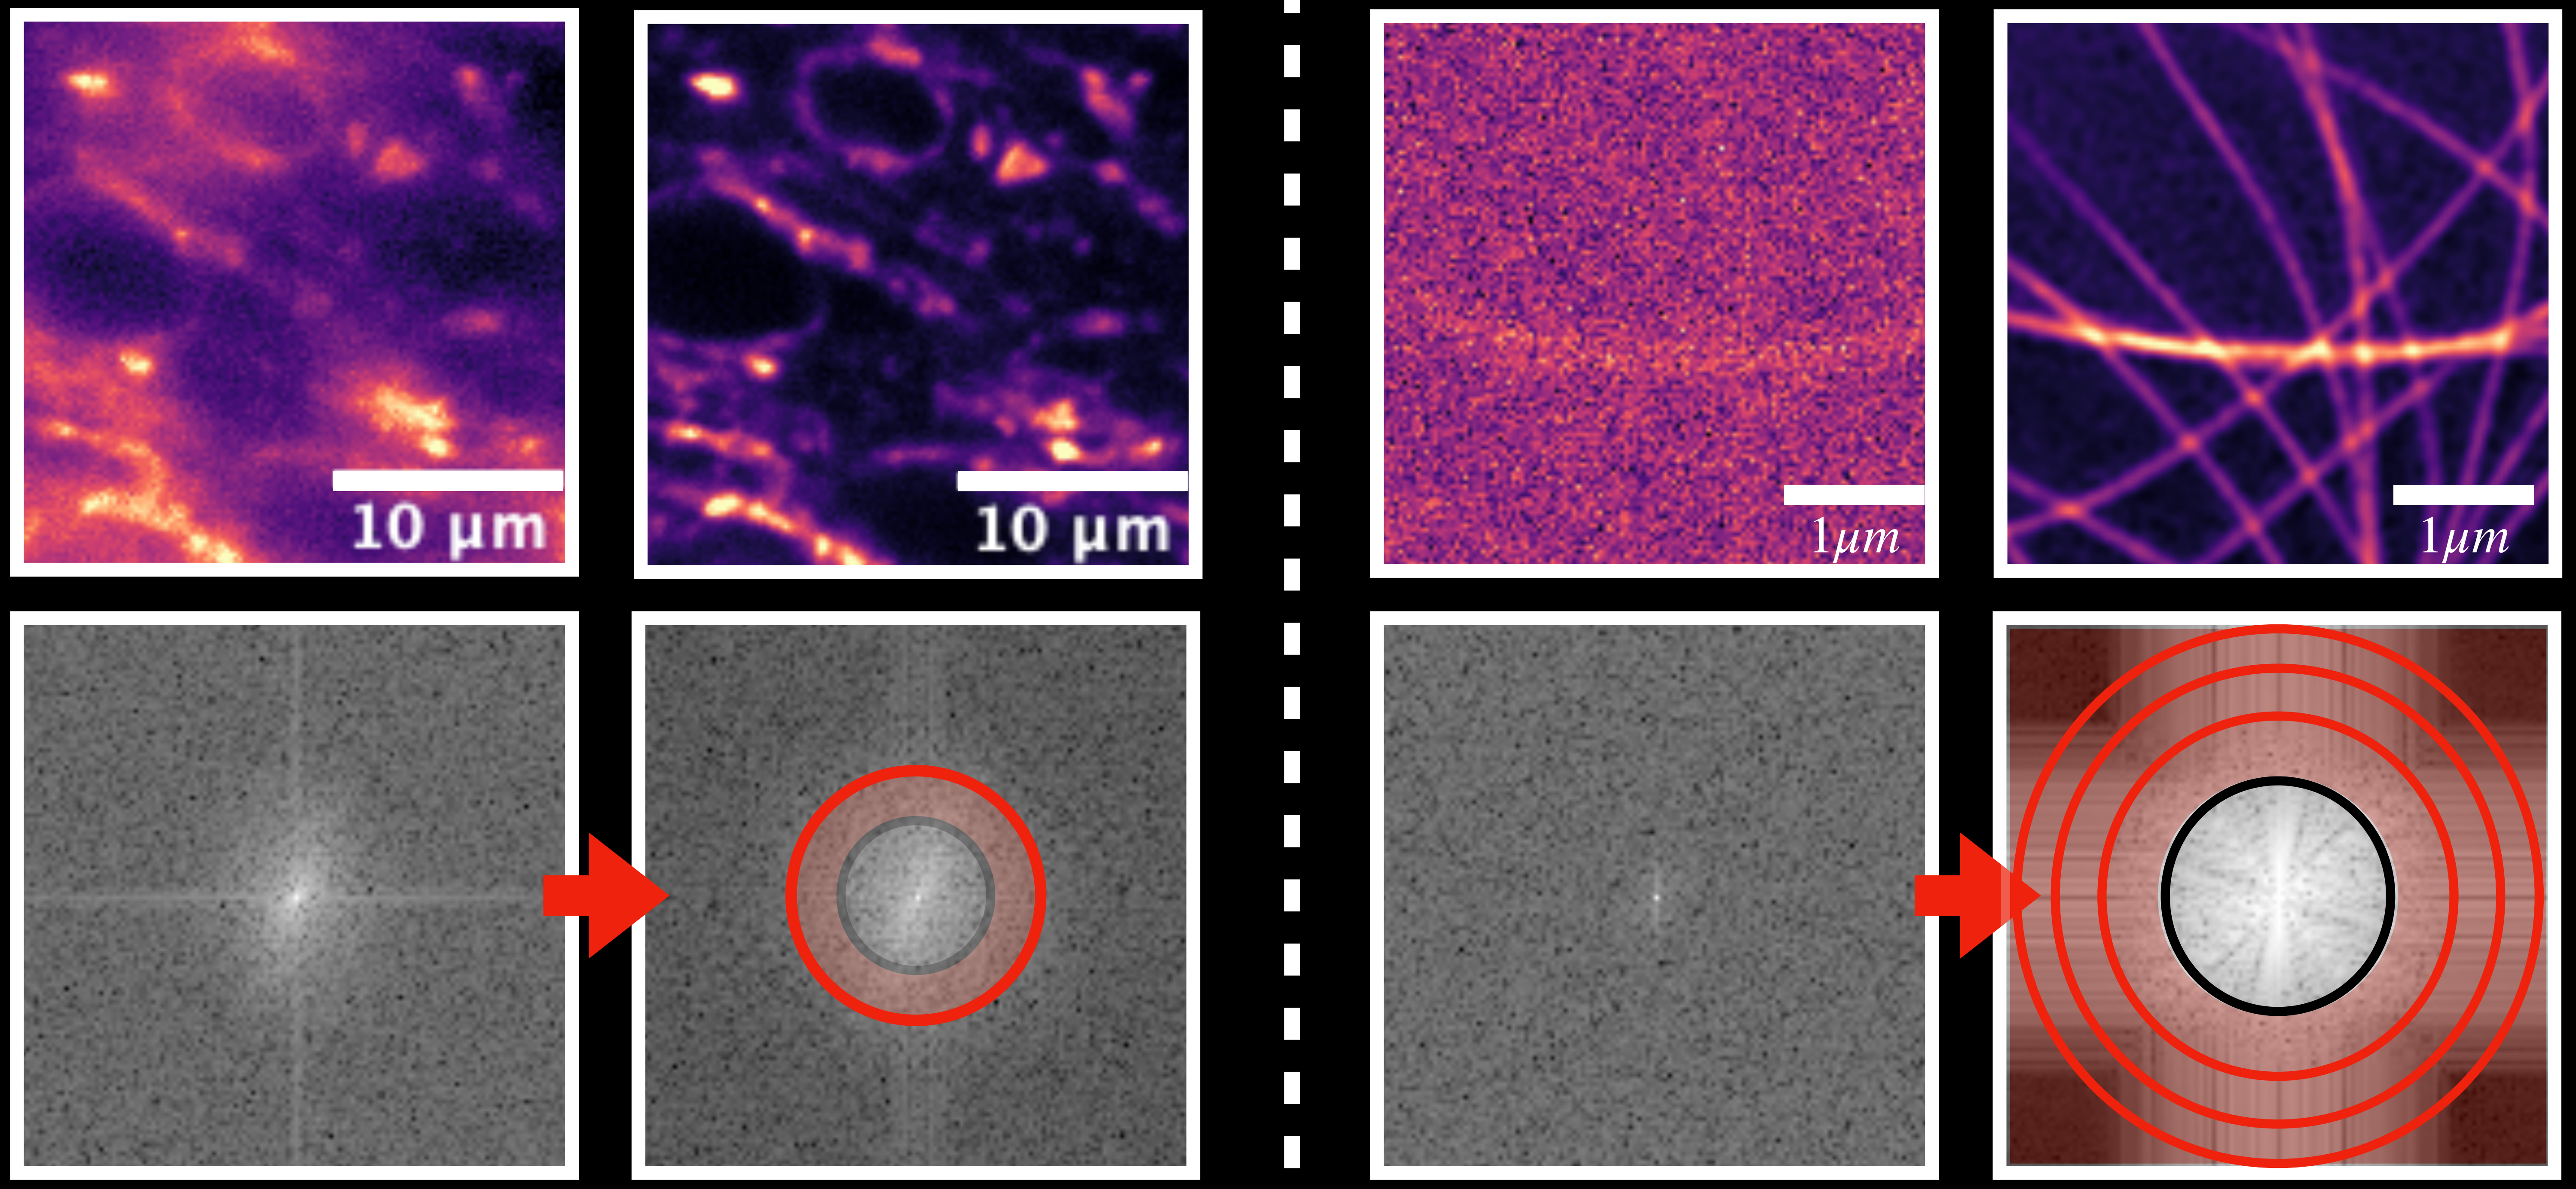

Widefield Image | Confocal Image | Noisy LR Image | SR Image

Manipulation of observable frequency ranges | Extrapolation to non-observed frequency ranges

Interpolation | Extrapolation

- CSR is an inverse problem that predicts **extrapolated** high frequencies.
- It is critical for CSR to have the **best possible prior**.
- **Calibration** plots can tell if the learned prior represents the actual uncertainty well, but this is an **active area of research**.
- Hence, we believe that **CSR should be used with caution** only when gains can justify the uncertainty levels in predictions.

*HazeMatching, Ray et al. - CVPR 2026 (Findings)

Conclusion

- Flow Matching transforms **low-resolution images to high-resolution images under heavy noise.**

- We propose practical way of using **GenAI for scientific image restoration** (with a cautionary note).

Conclusion

- Flow Matching transforms **low-resolution images to high-resolution images under heavy noise.**
- We propose practical way of using **GenAI for scientific image restoration** (with a cautionary note).

Paper, more results and interactive tool

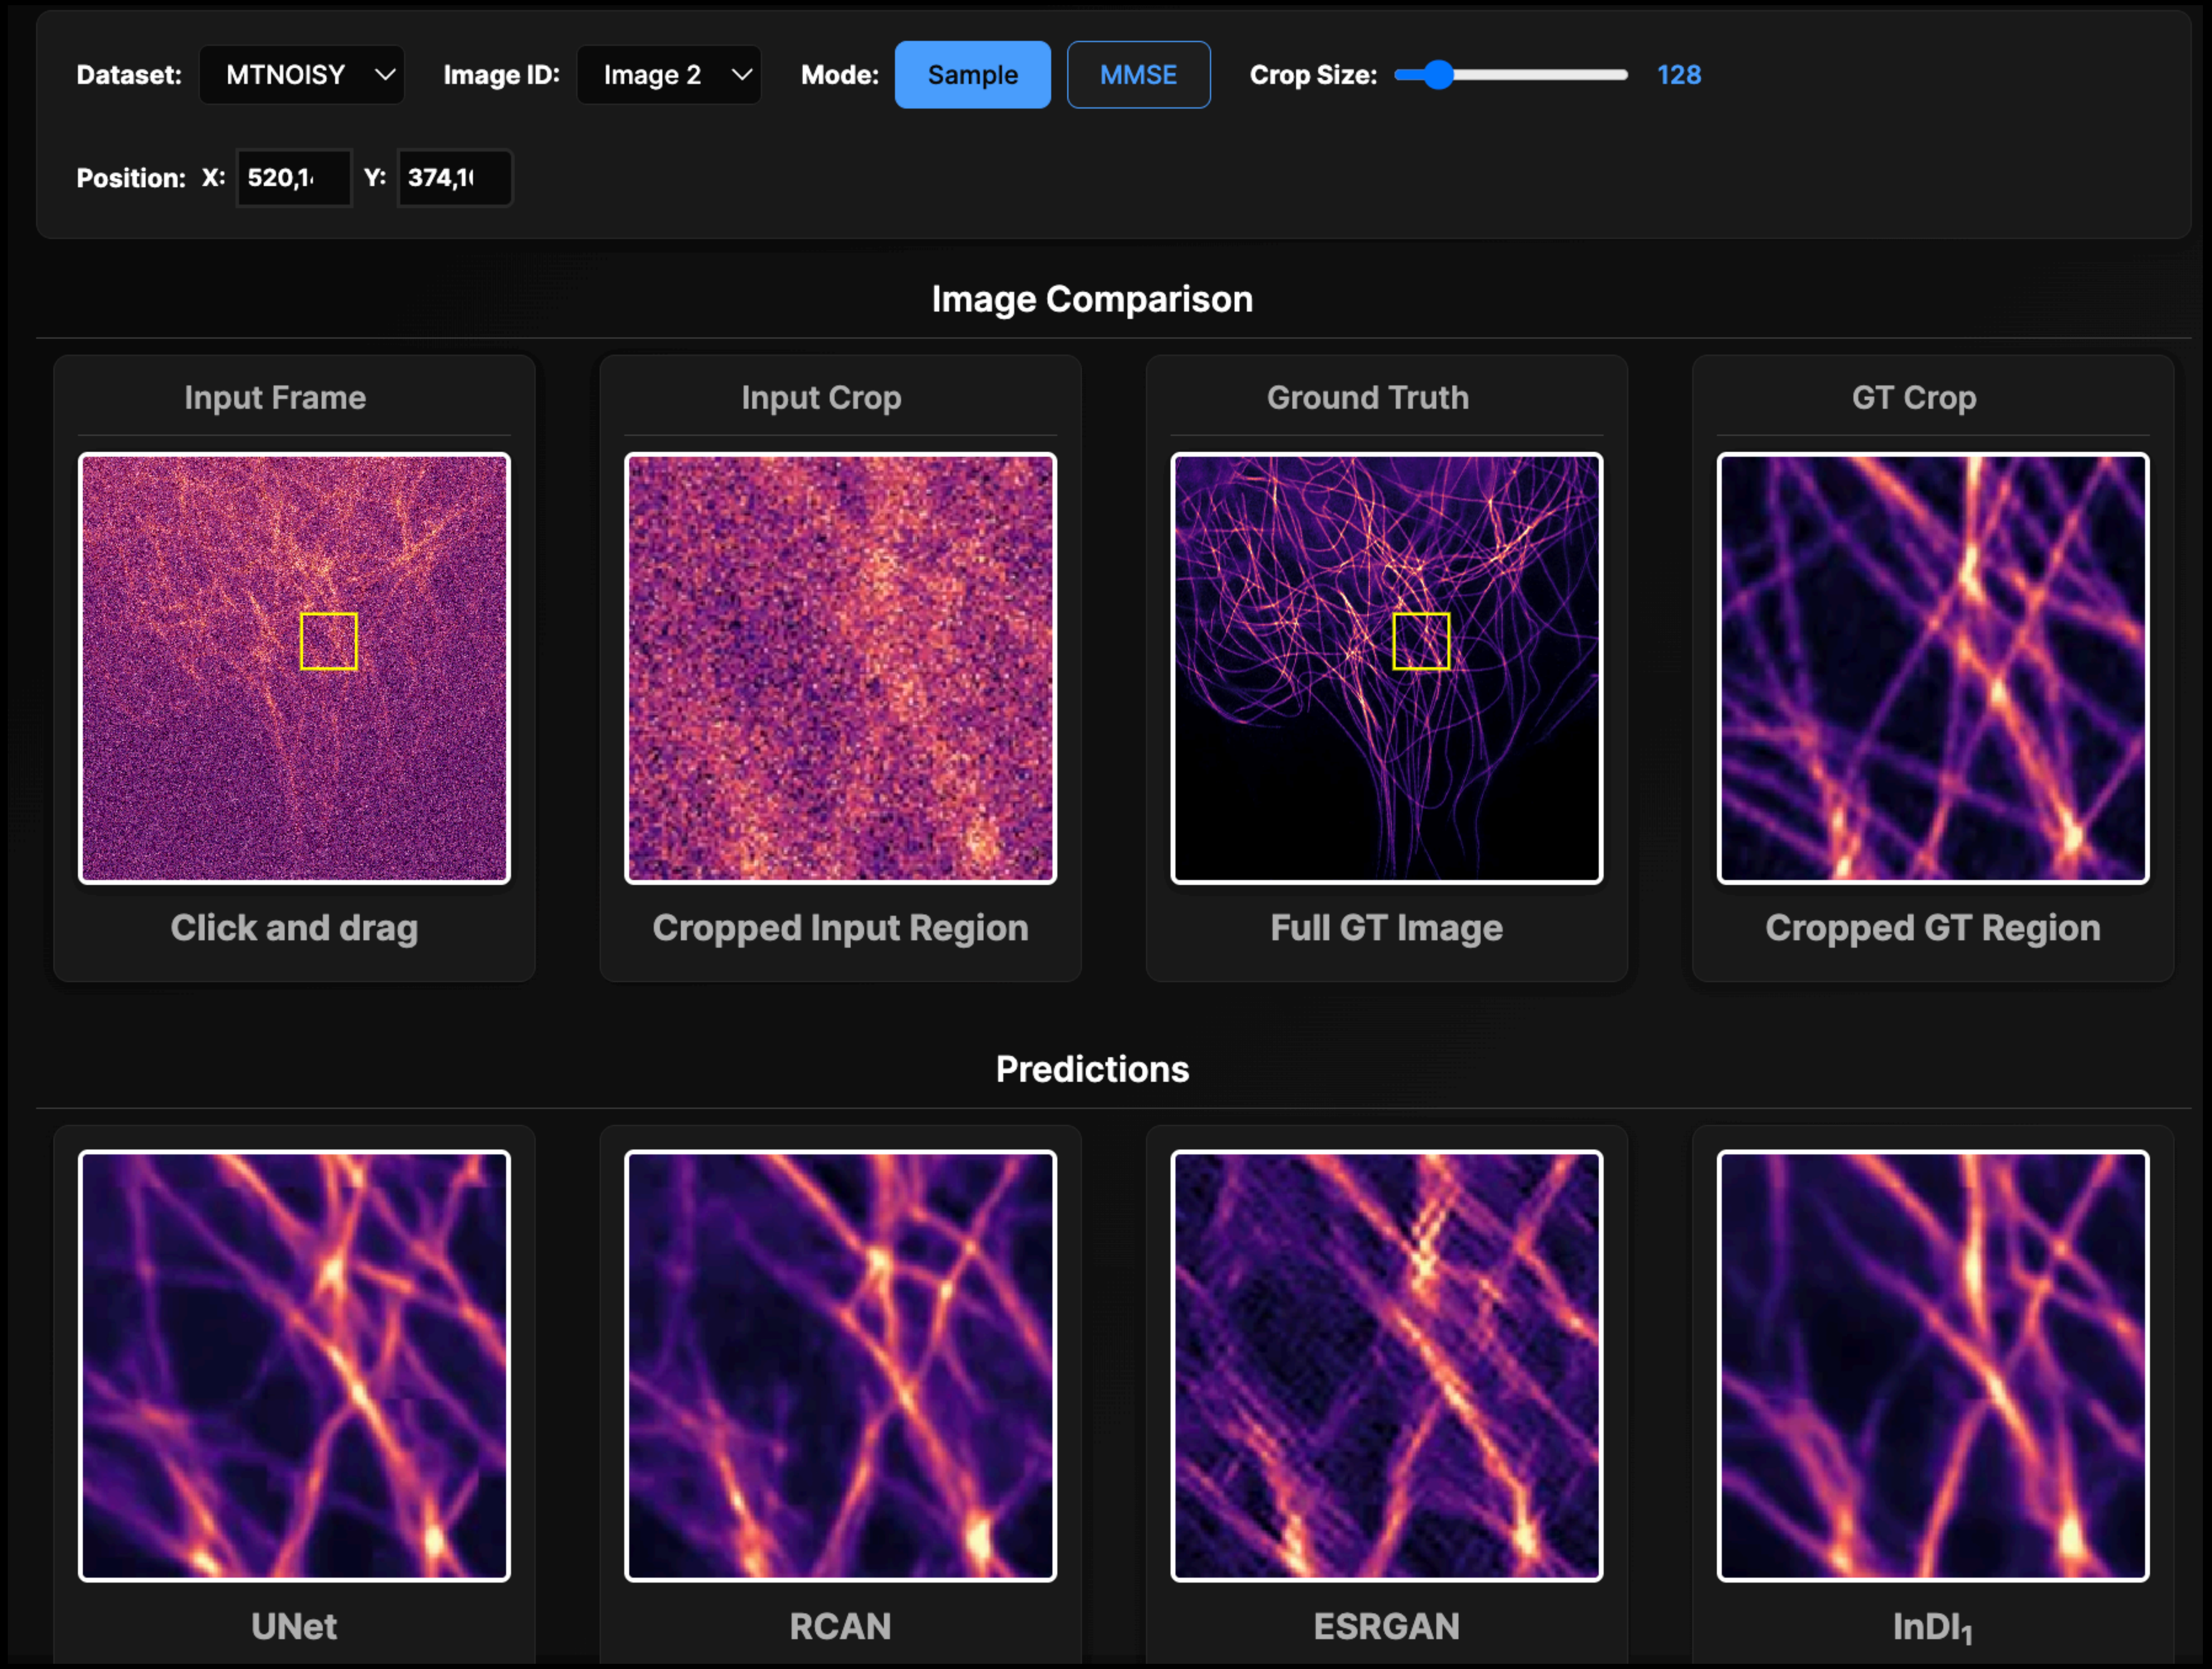

Wanna know more?

Find our poster! 😎

If I can find a spot