

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



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

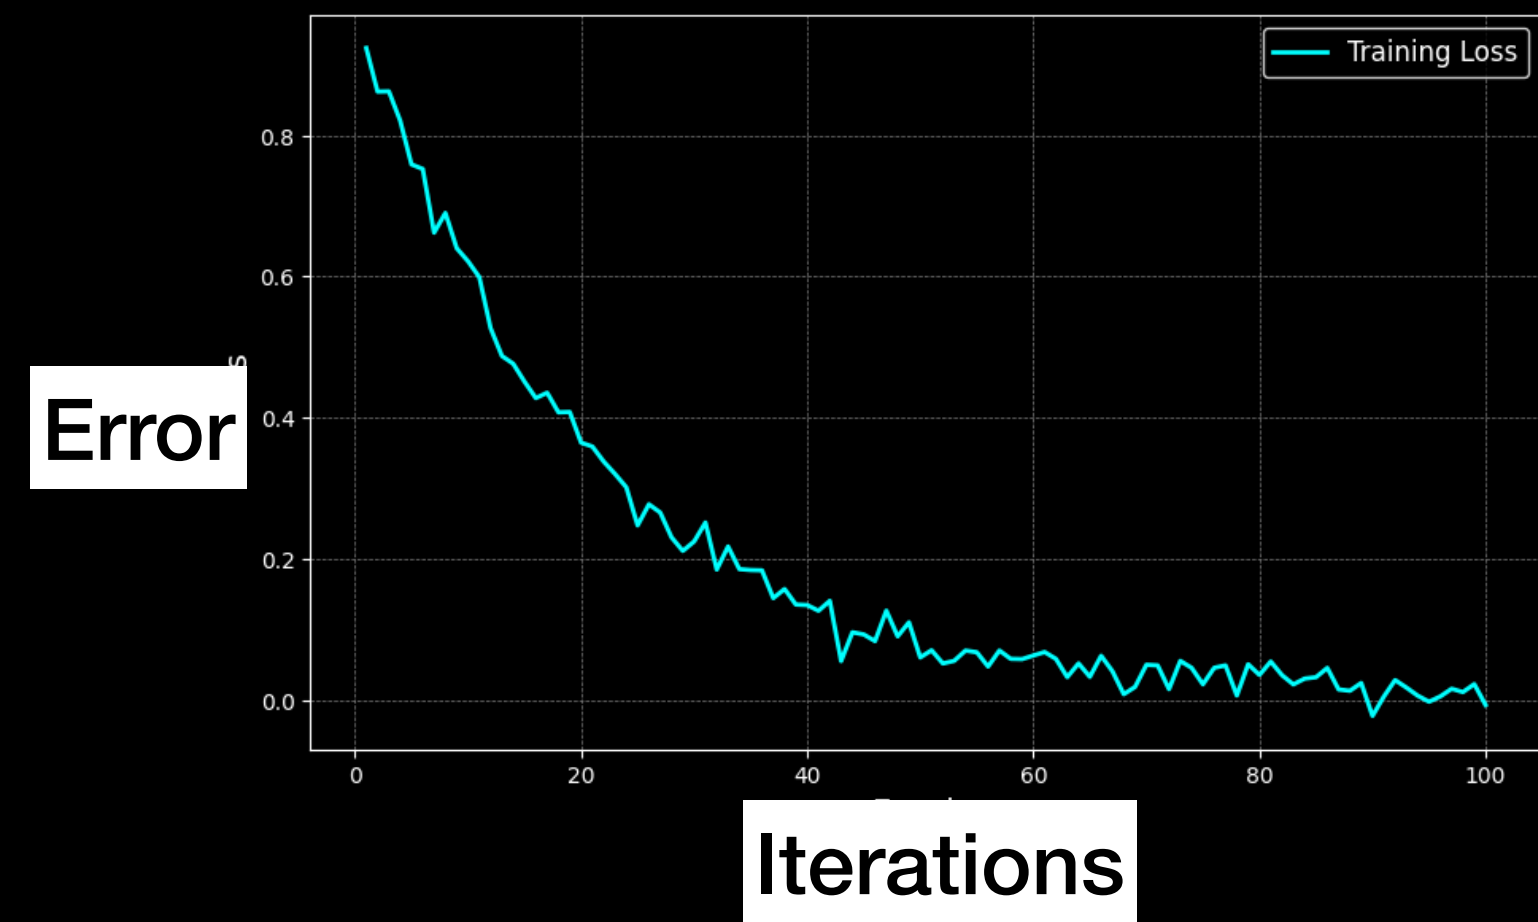
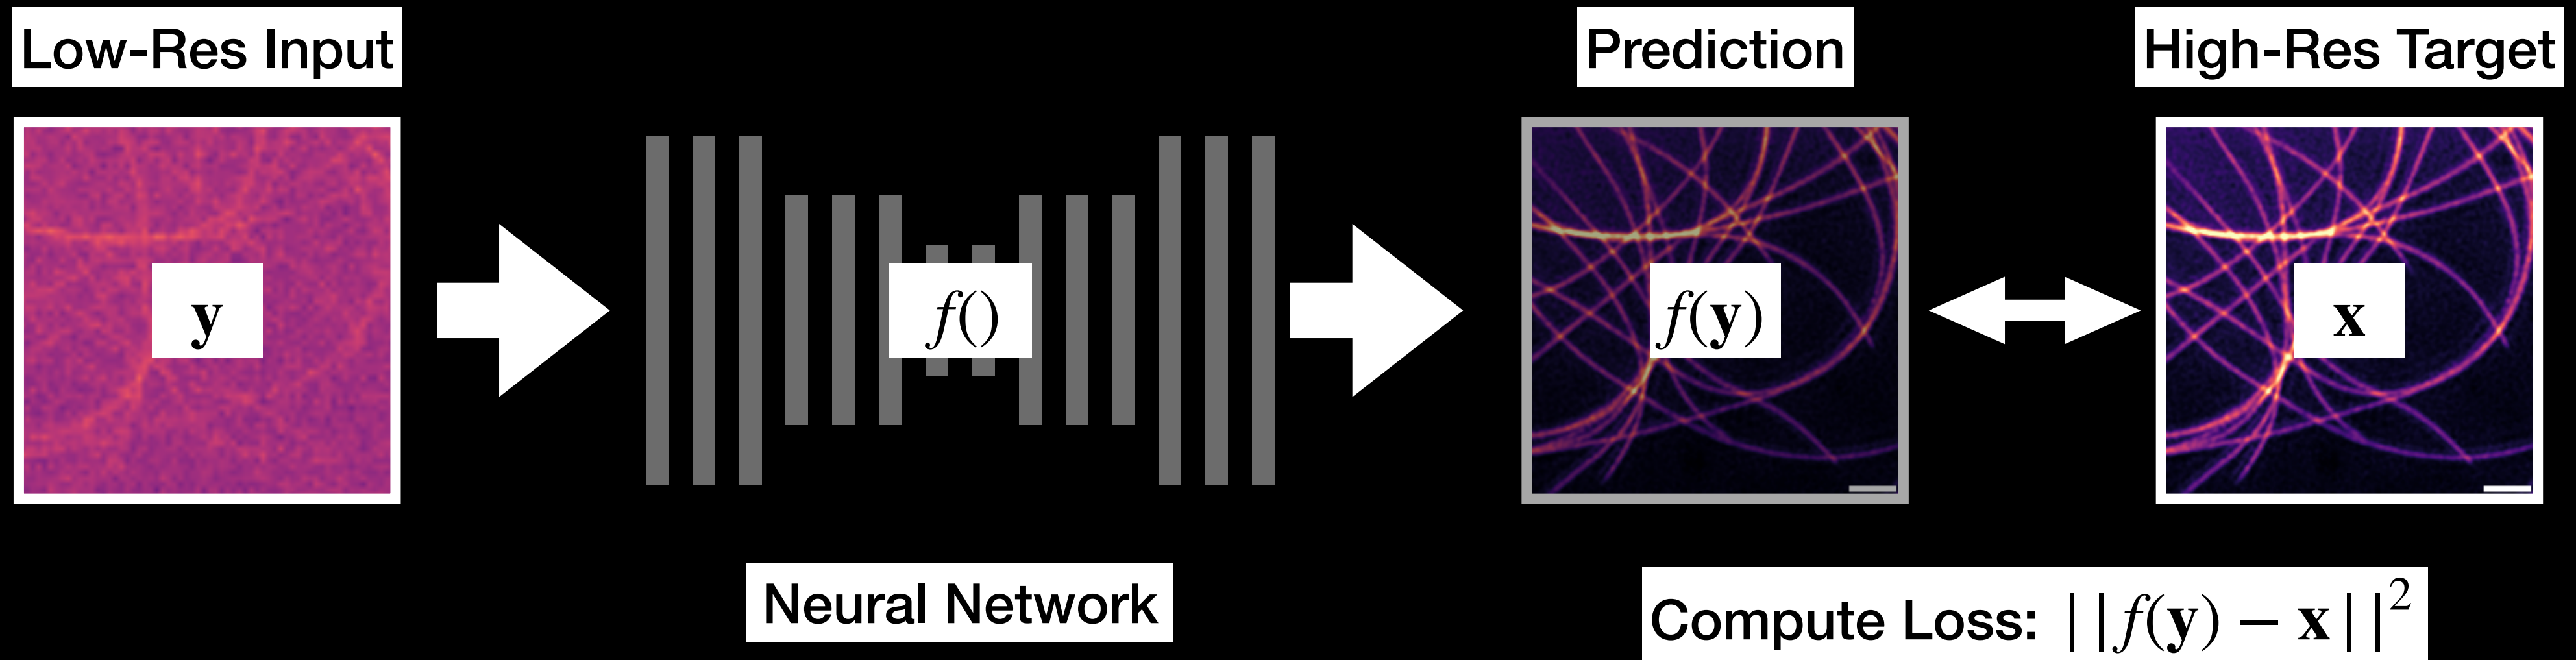
- Enable **imaging at low light** (less phototoxicity & photobleaching).
- Improve **temporal resolution** in live-cell imaging.
- ...

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





Neural Network


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

Various flavors


Low-Res Input



One-Shot Mapping!



Neural Network



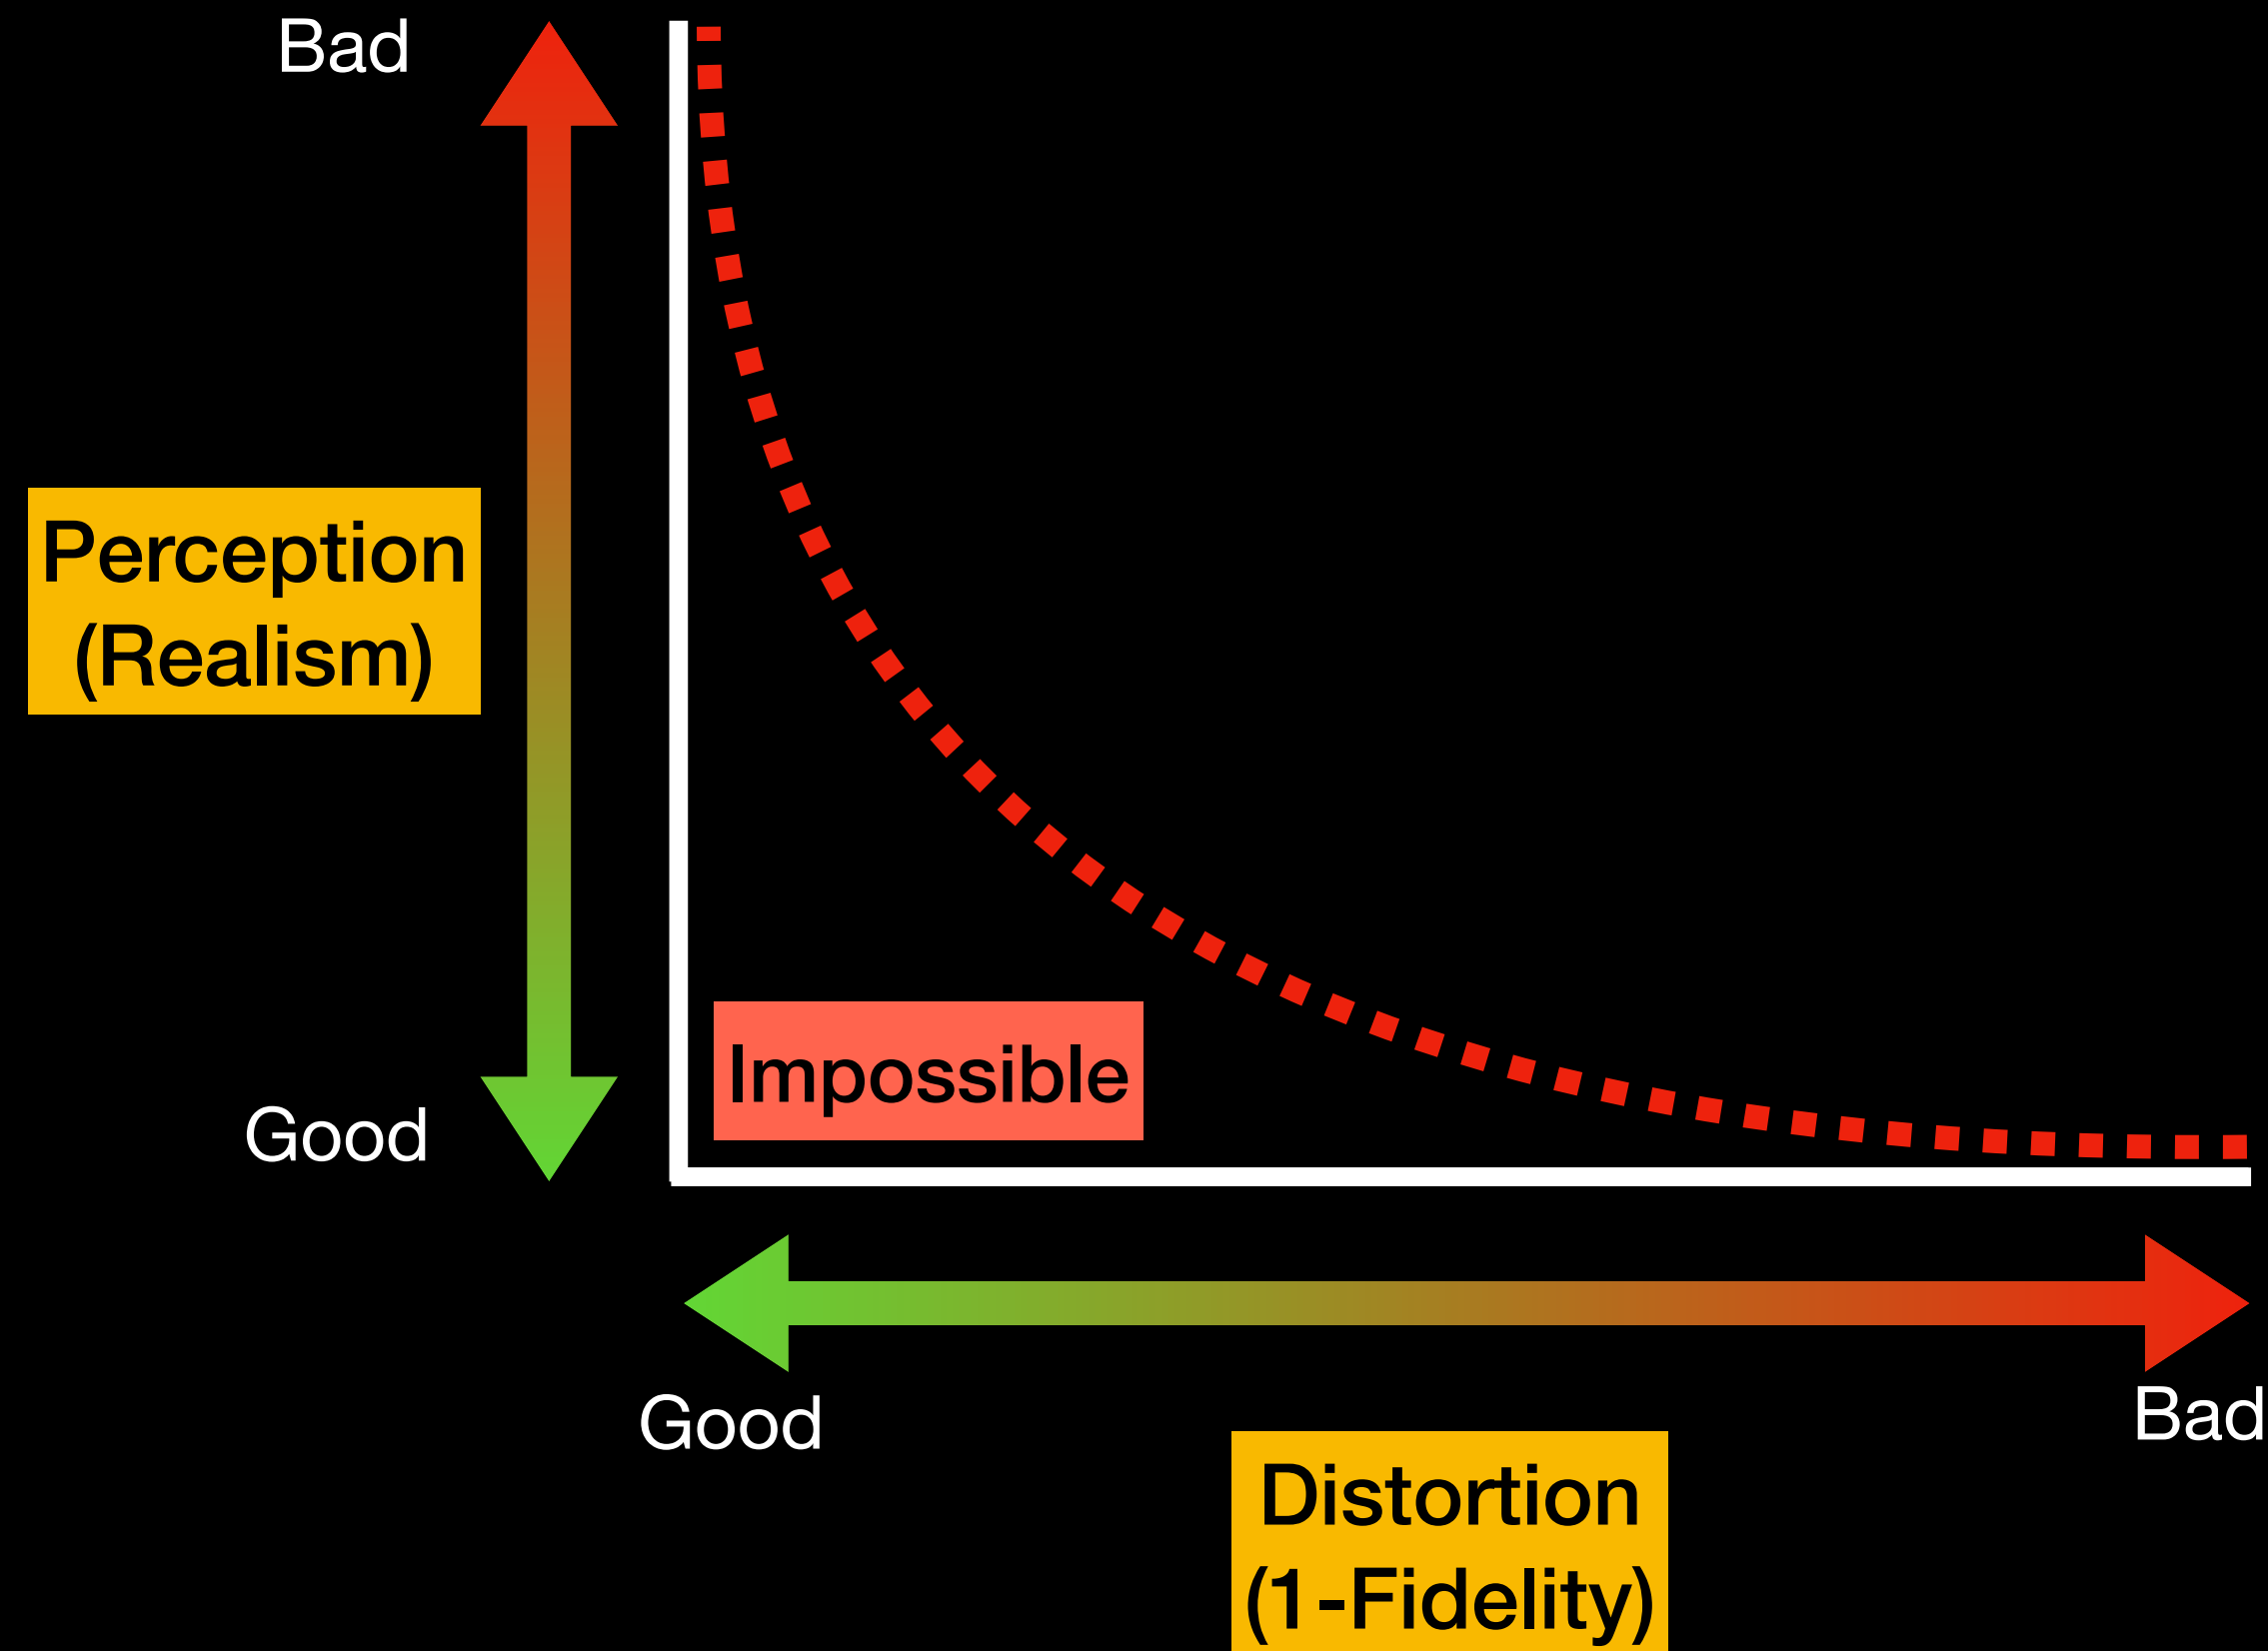
Prediction



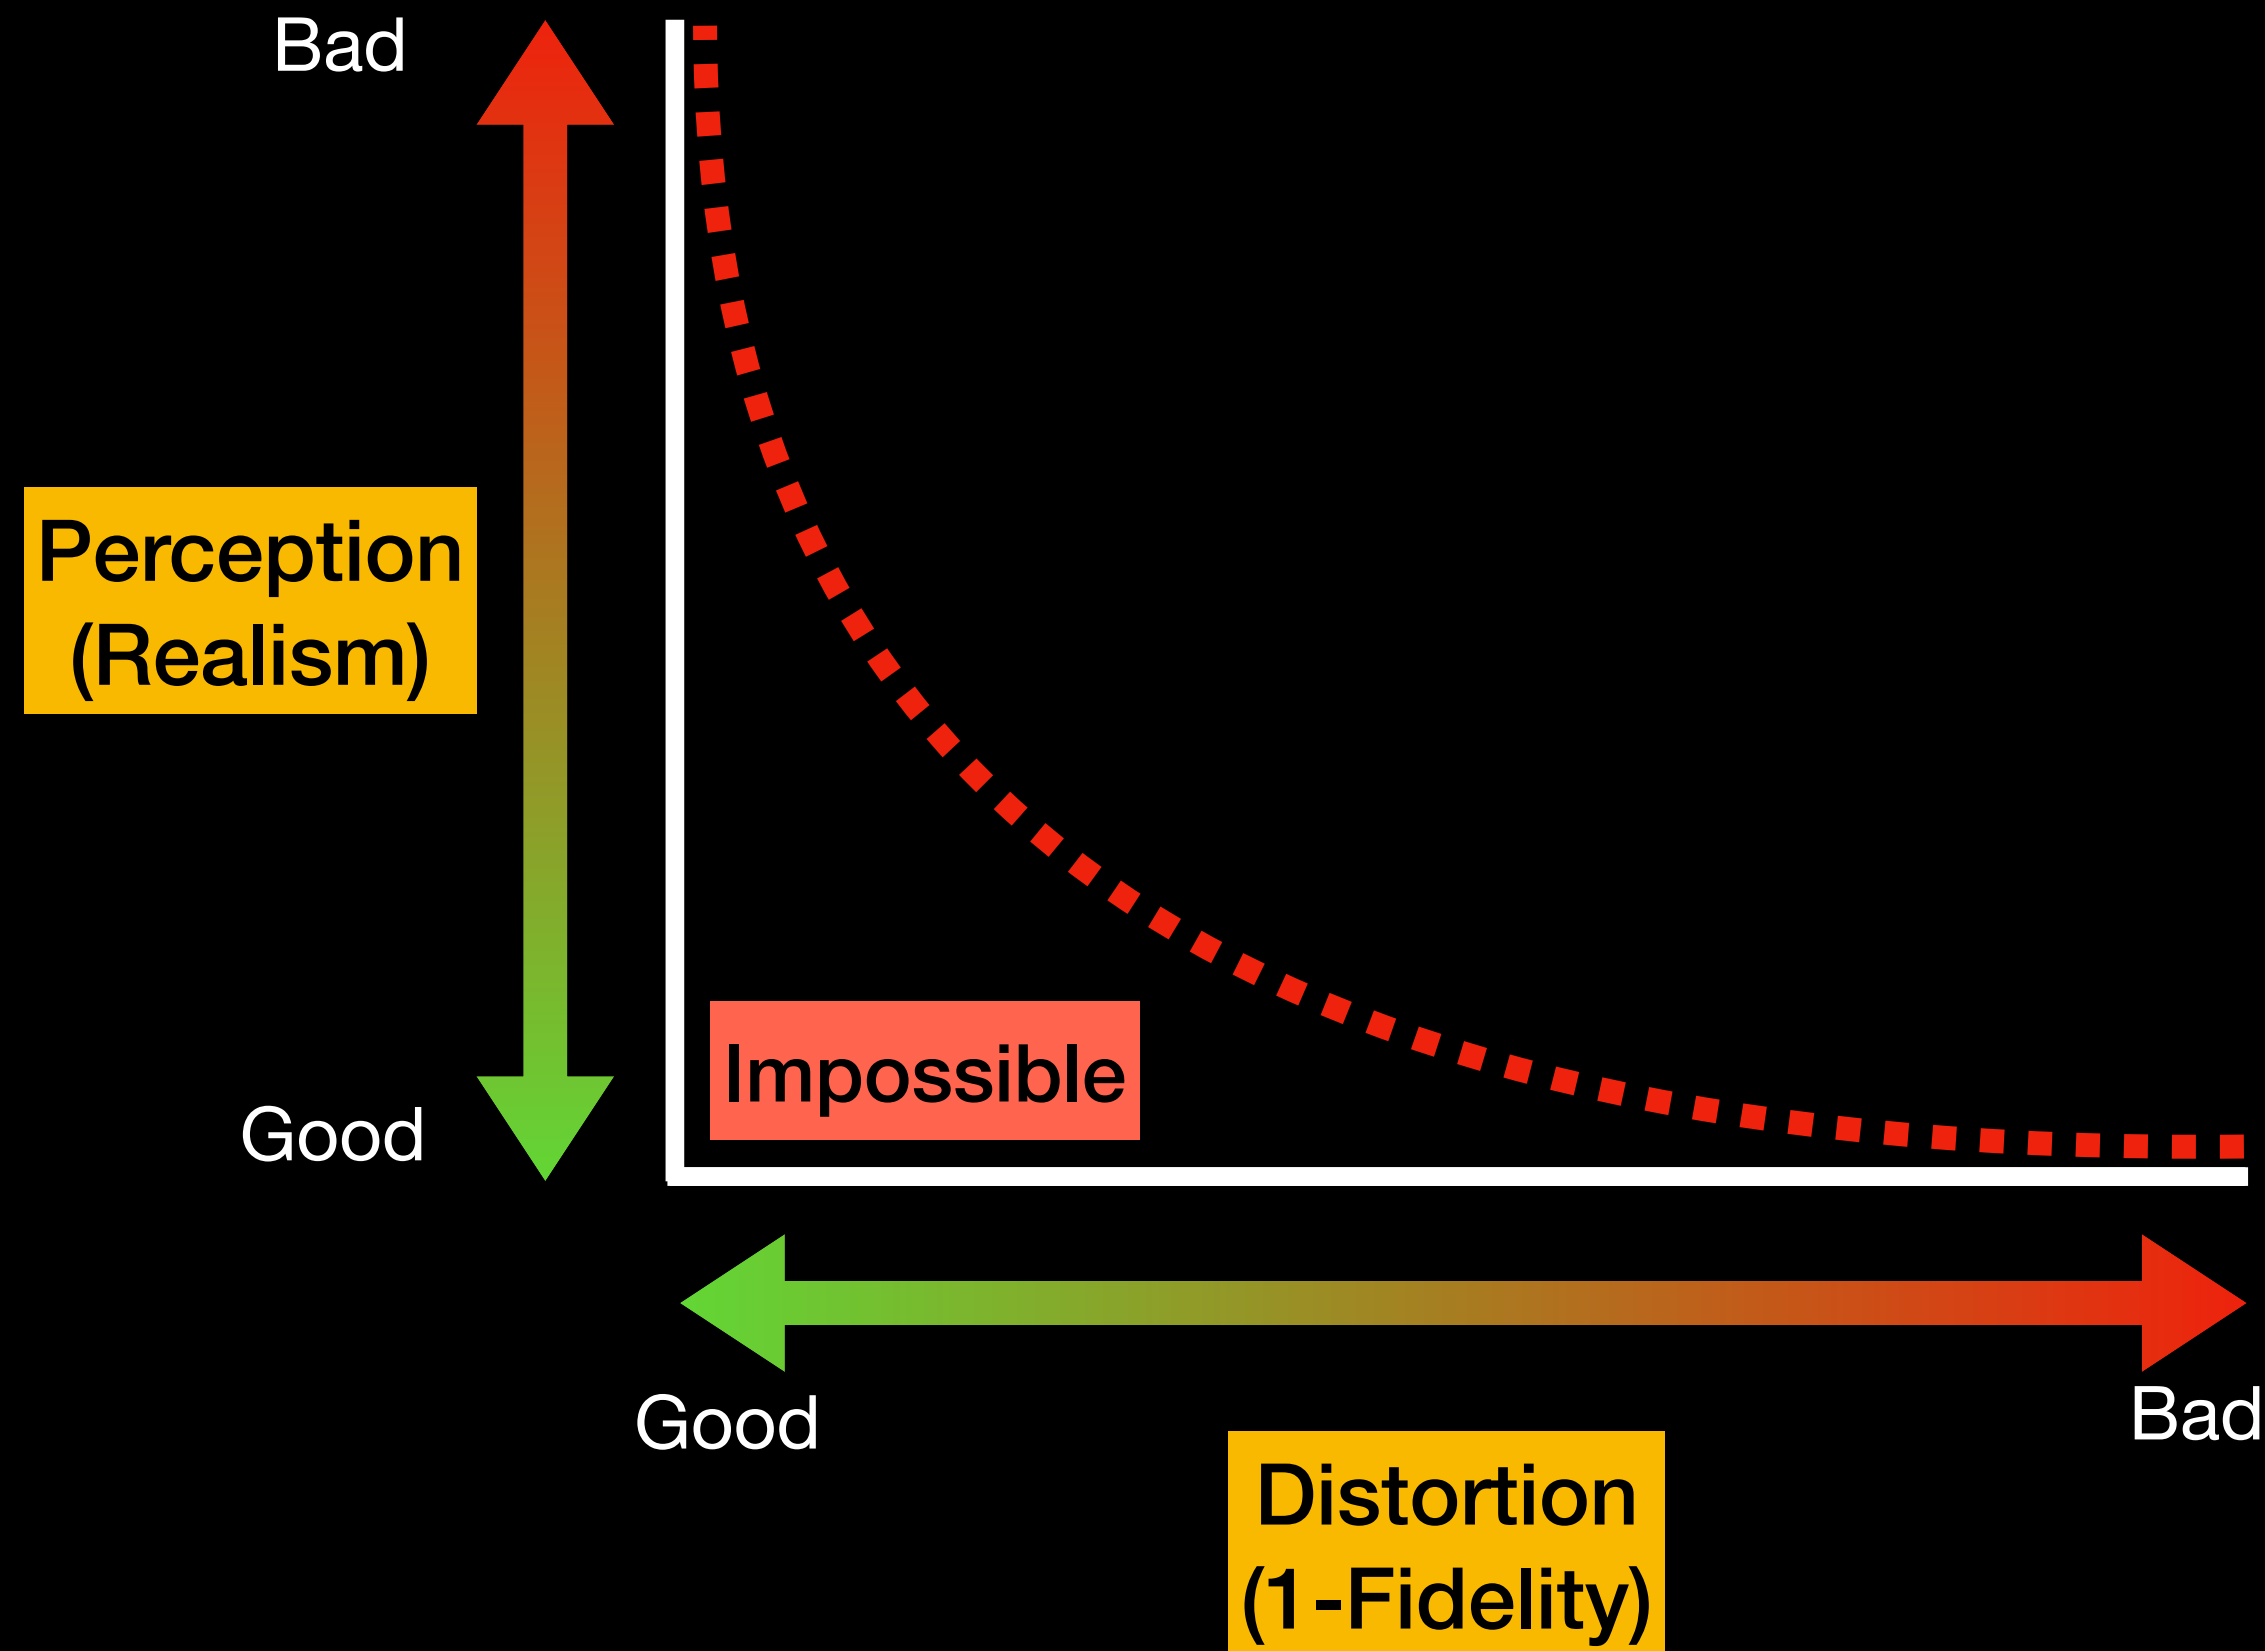
Prediction is usually done in a single step




Perception Distortion Tradeoff (PDT)



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!



Perception Distortion Tradeoff (PDT)



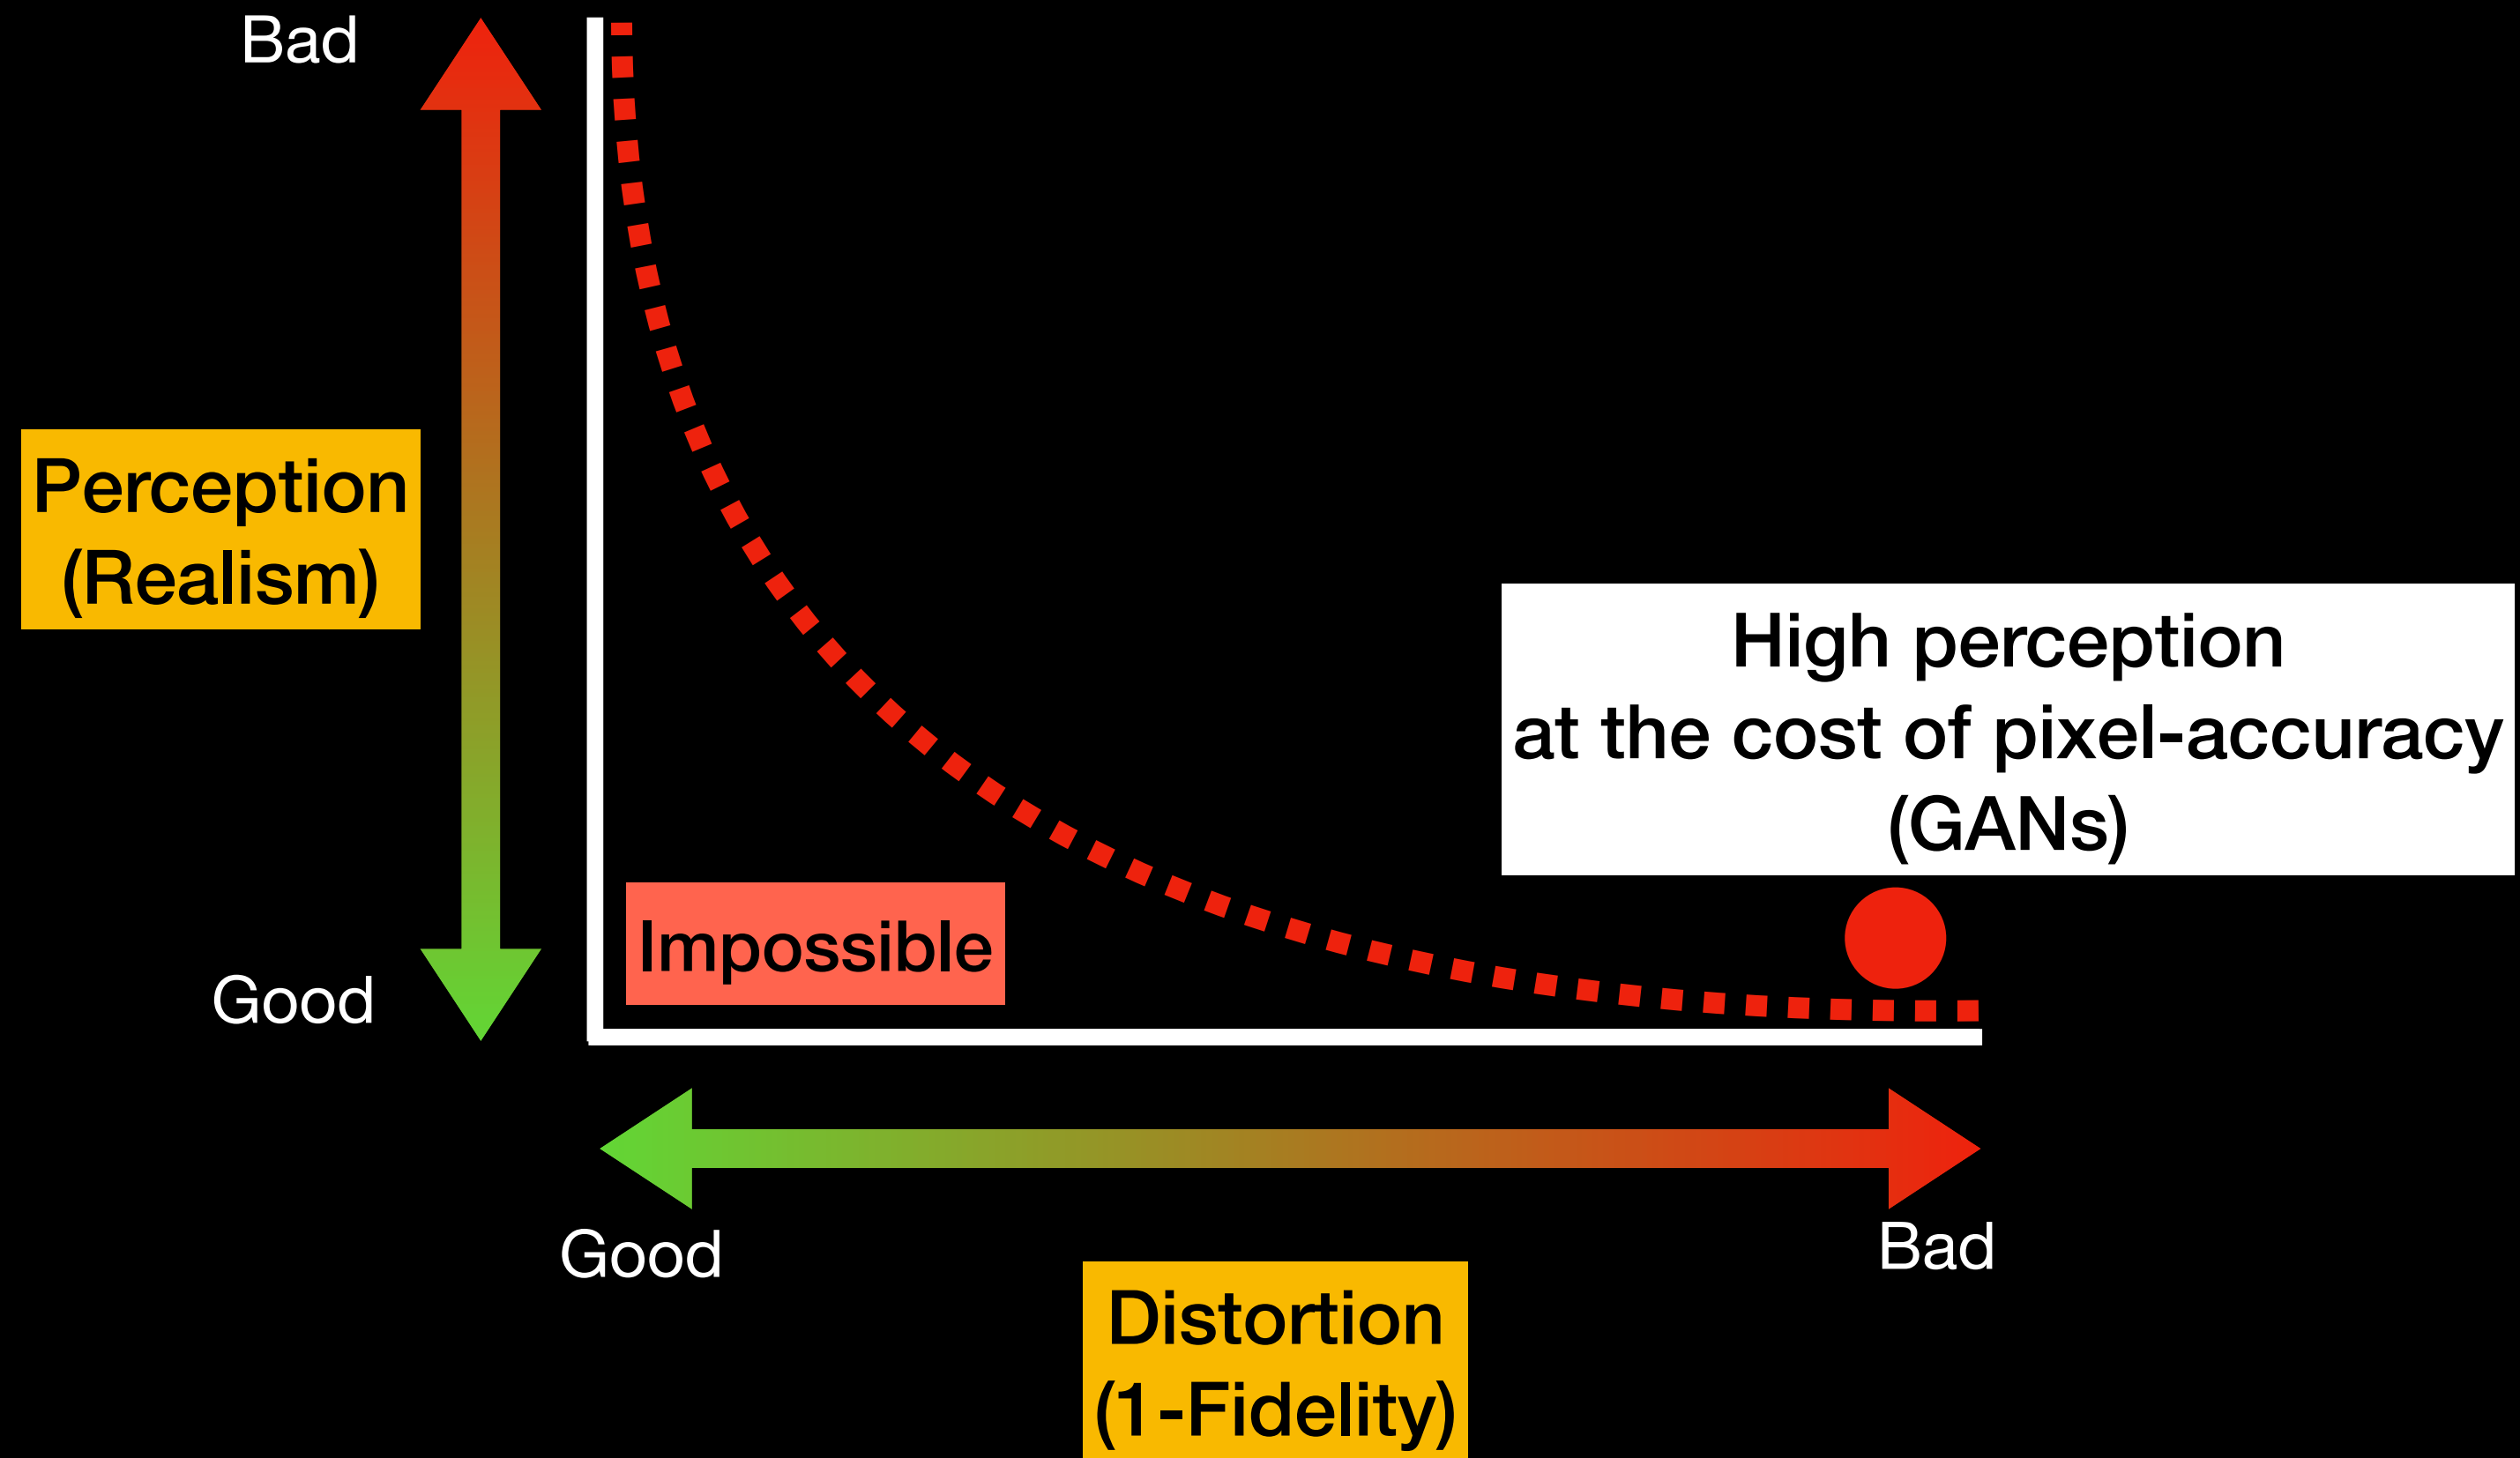
An model cannot be optimized for the *best perception* and *best distortion* at the same time!




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



Perception Distortion Tradeoff (PDT)



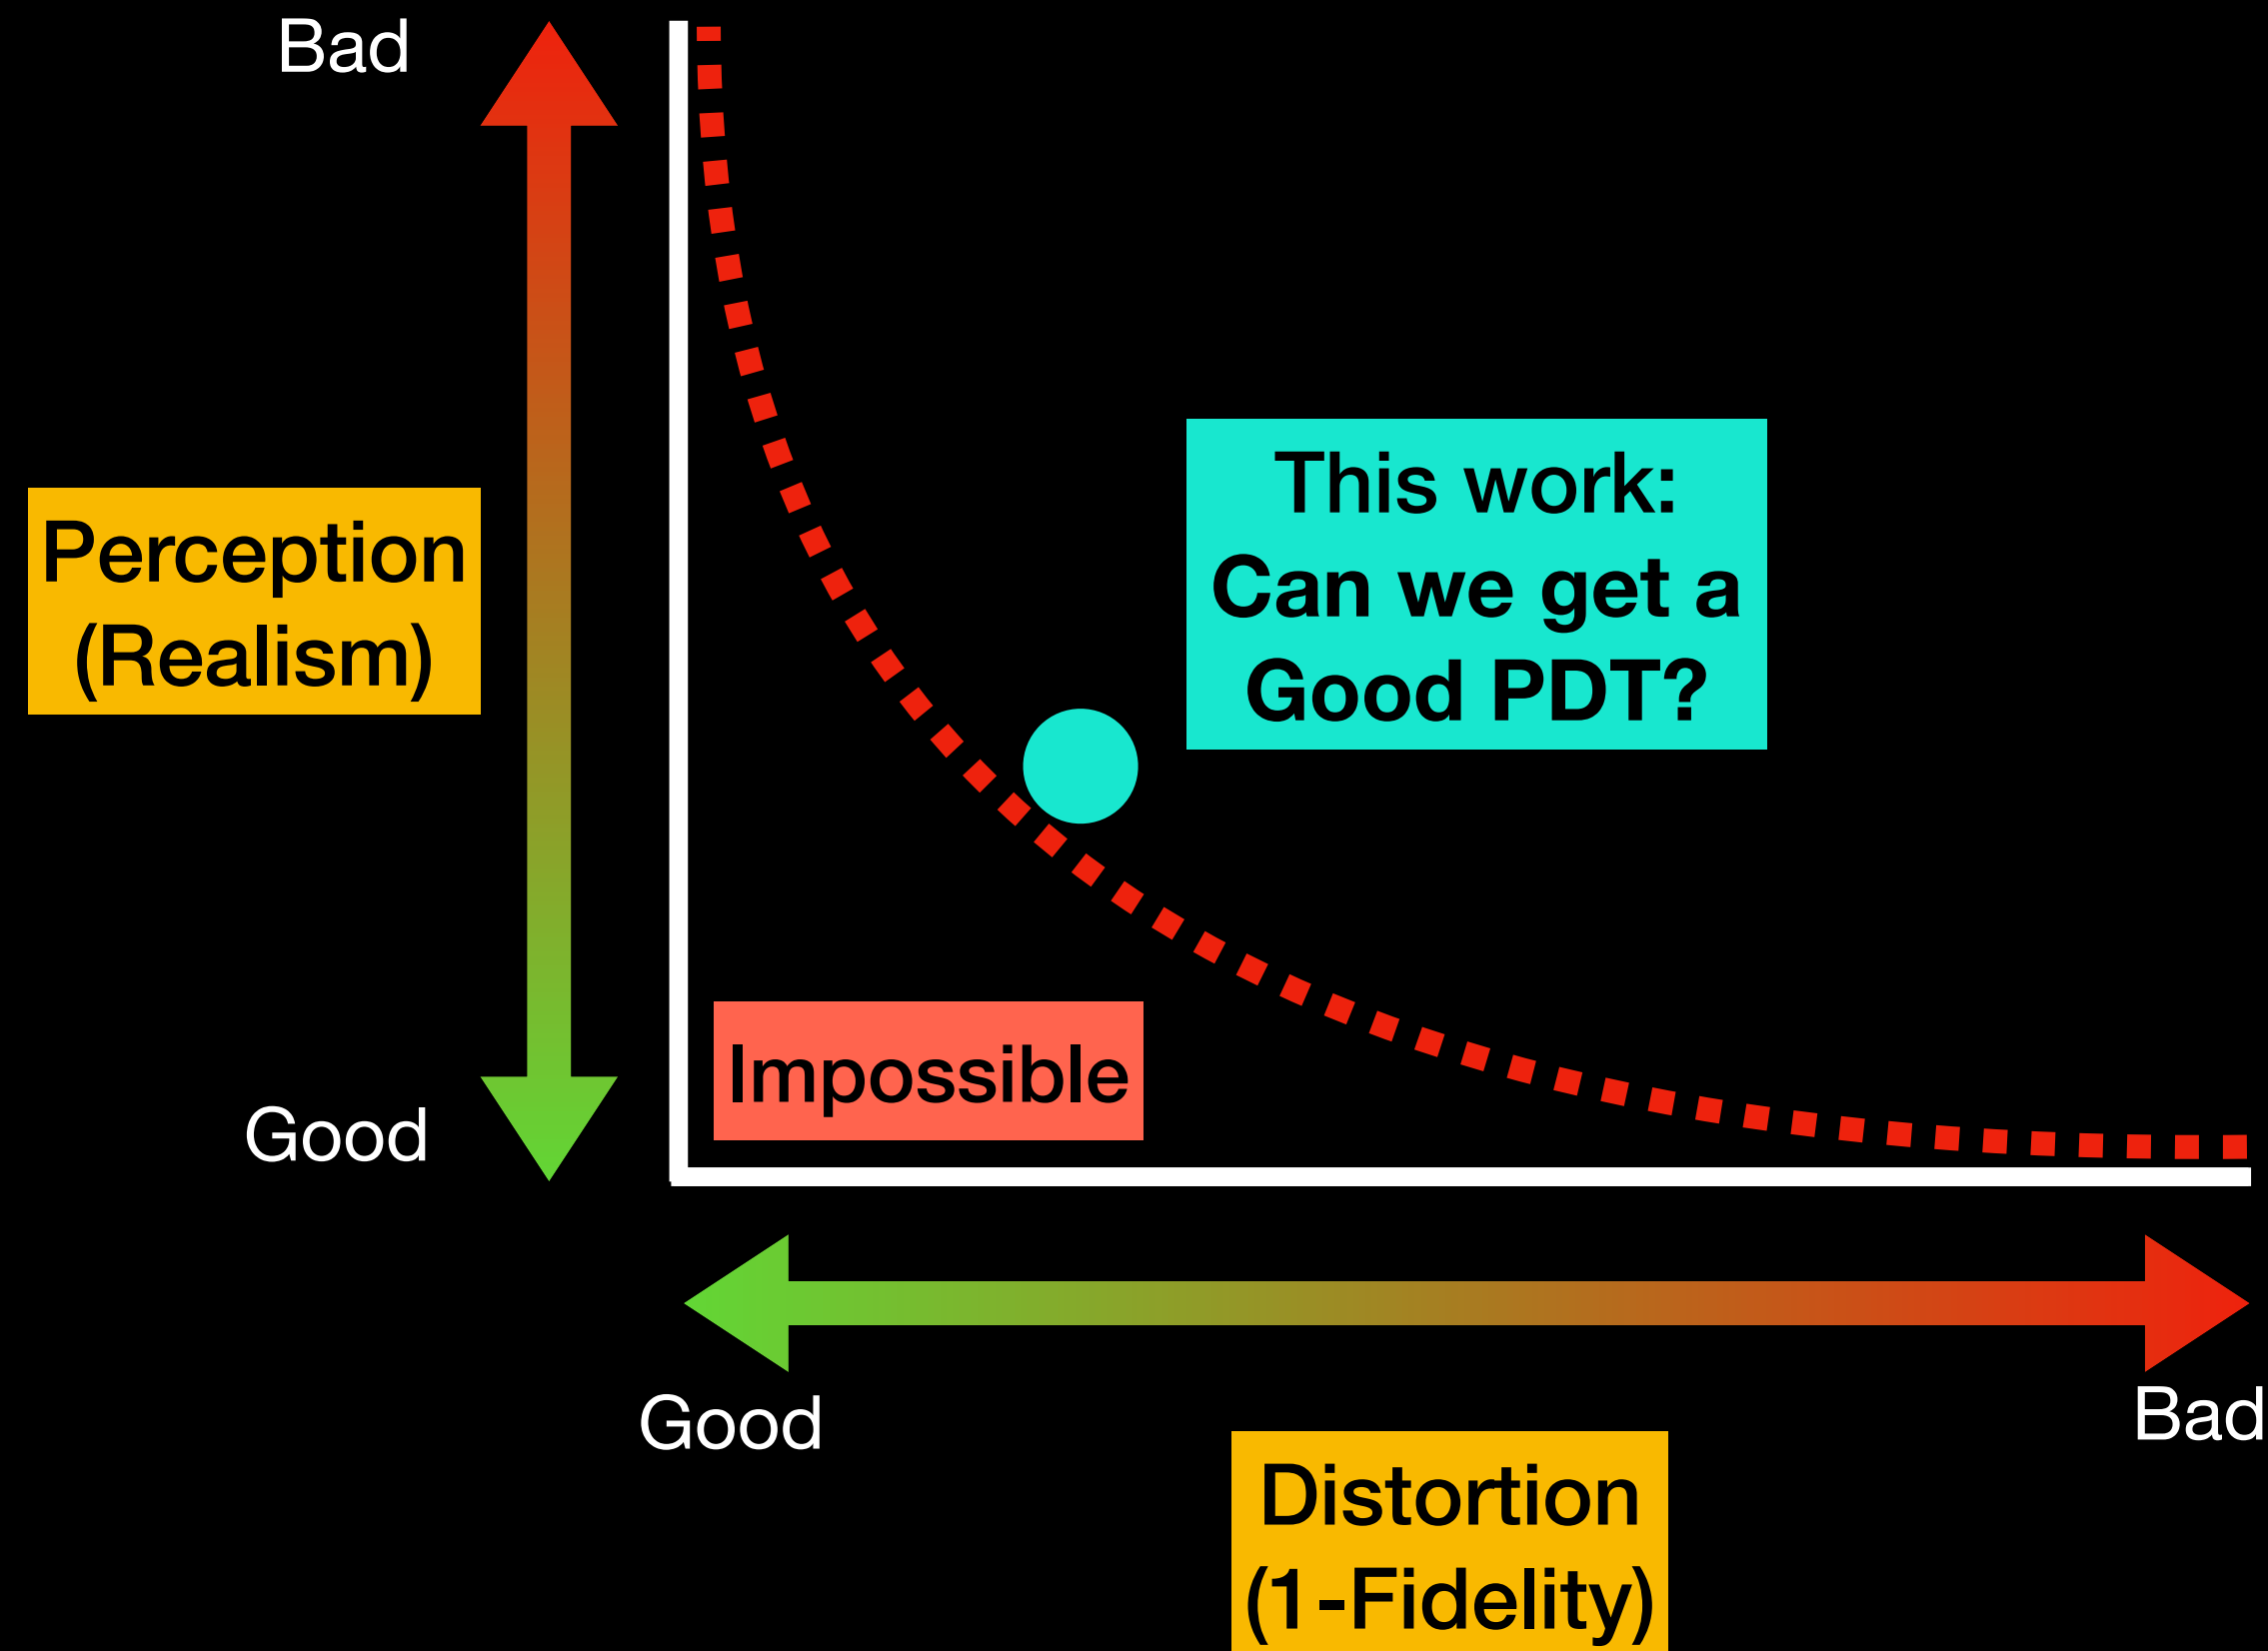
An model cannot be optimized for the *best perception* and *best distortion* at the same time!



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



Perception Distortion Tradeoff (PDT)



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)

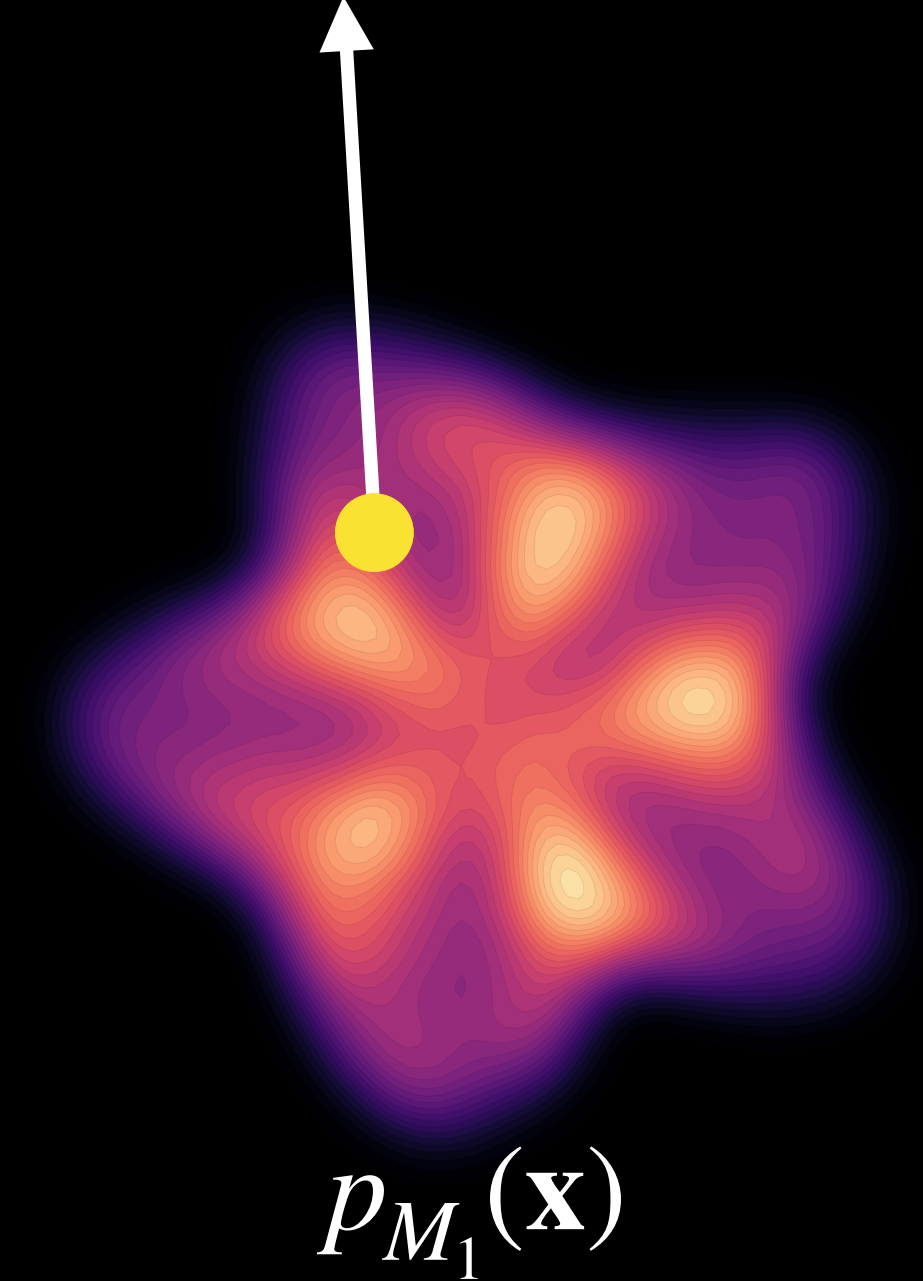


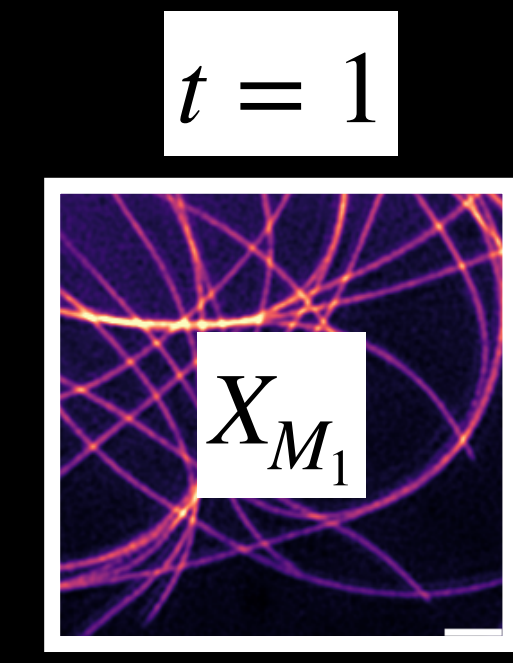



*HazeMatching, Ray et al. - CVPR 2026 (Findings)


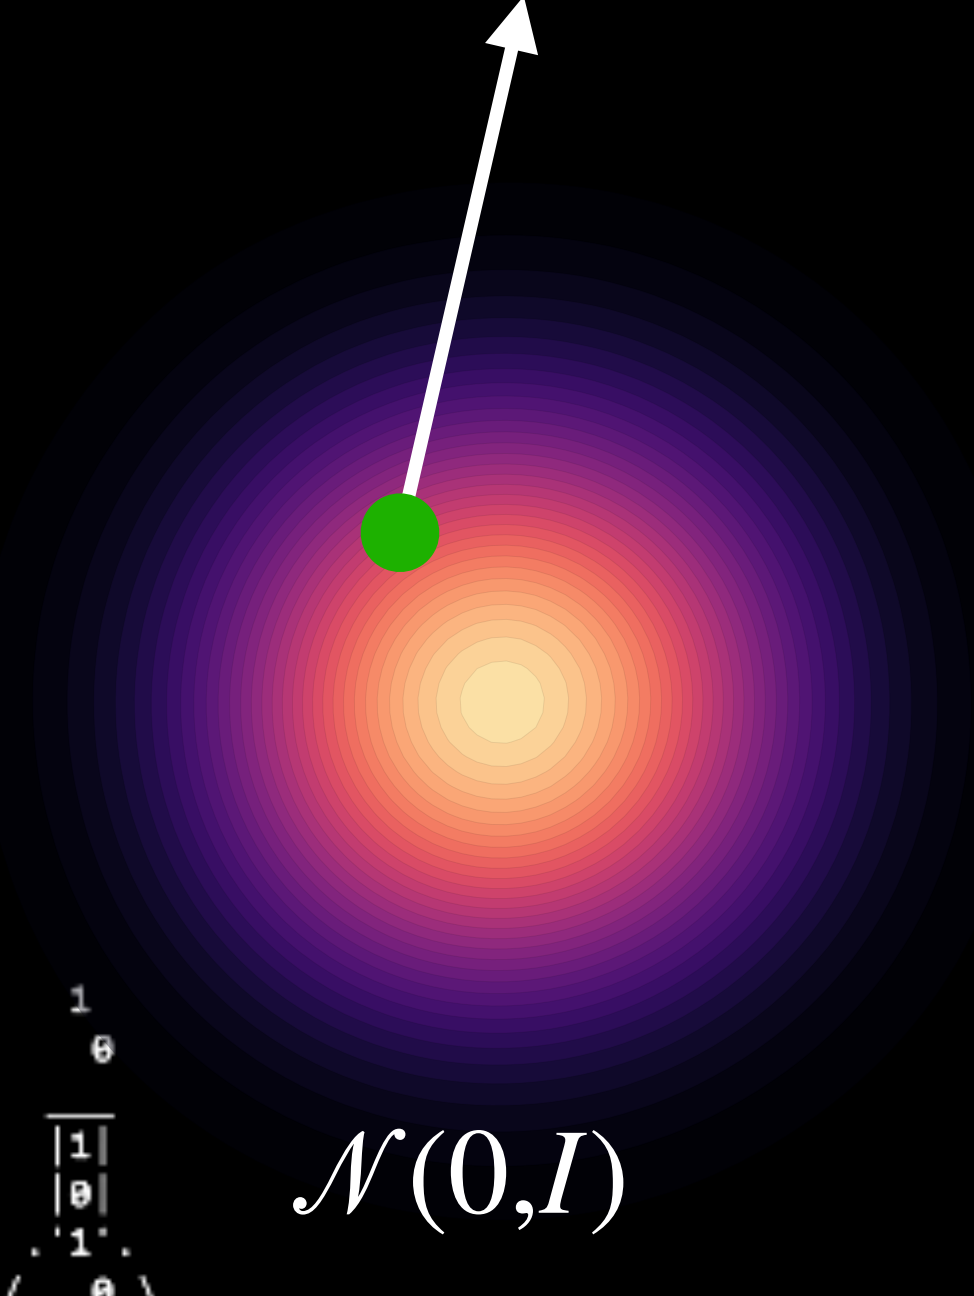


ResMatching - Training





ResMatching - Training



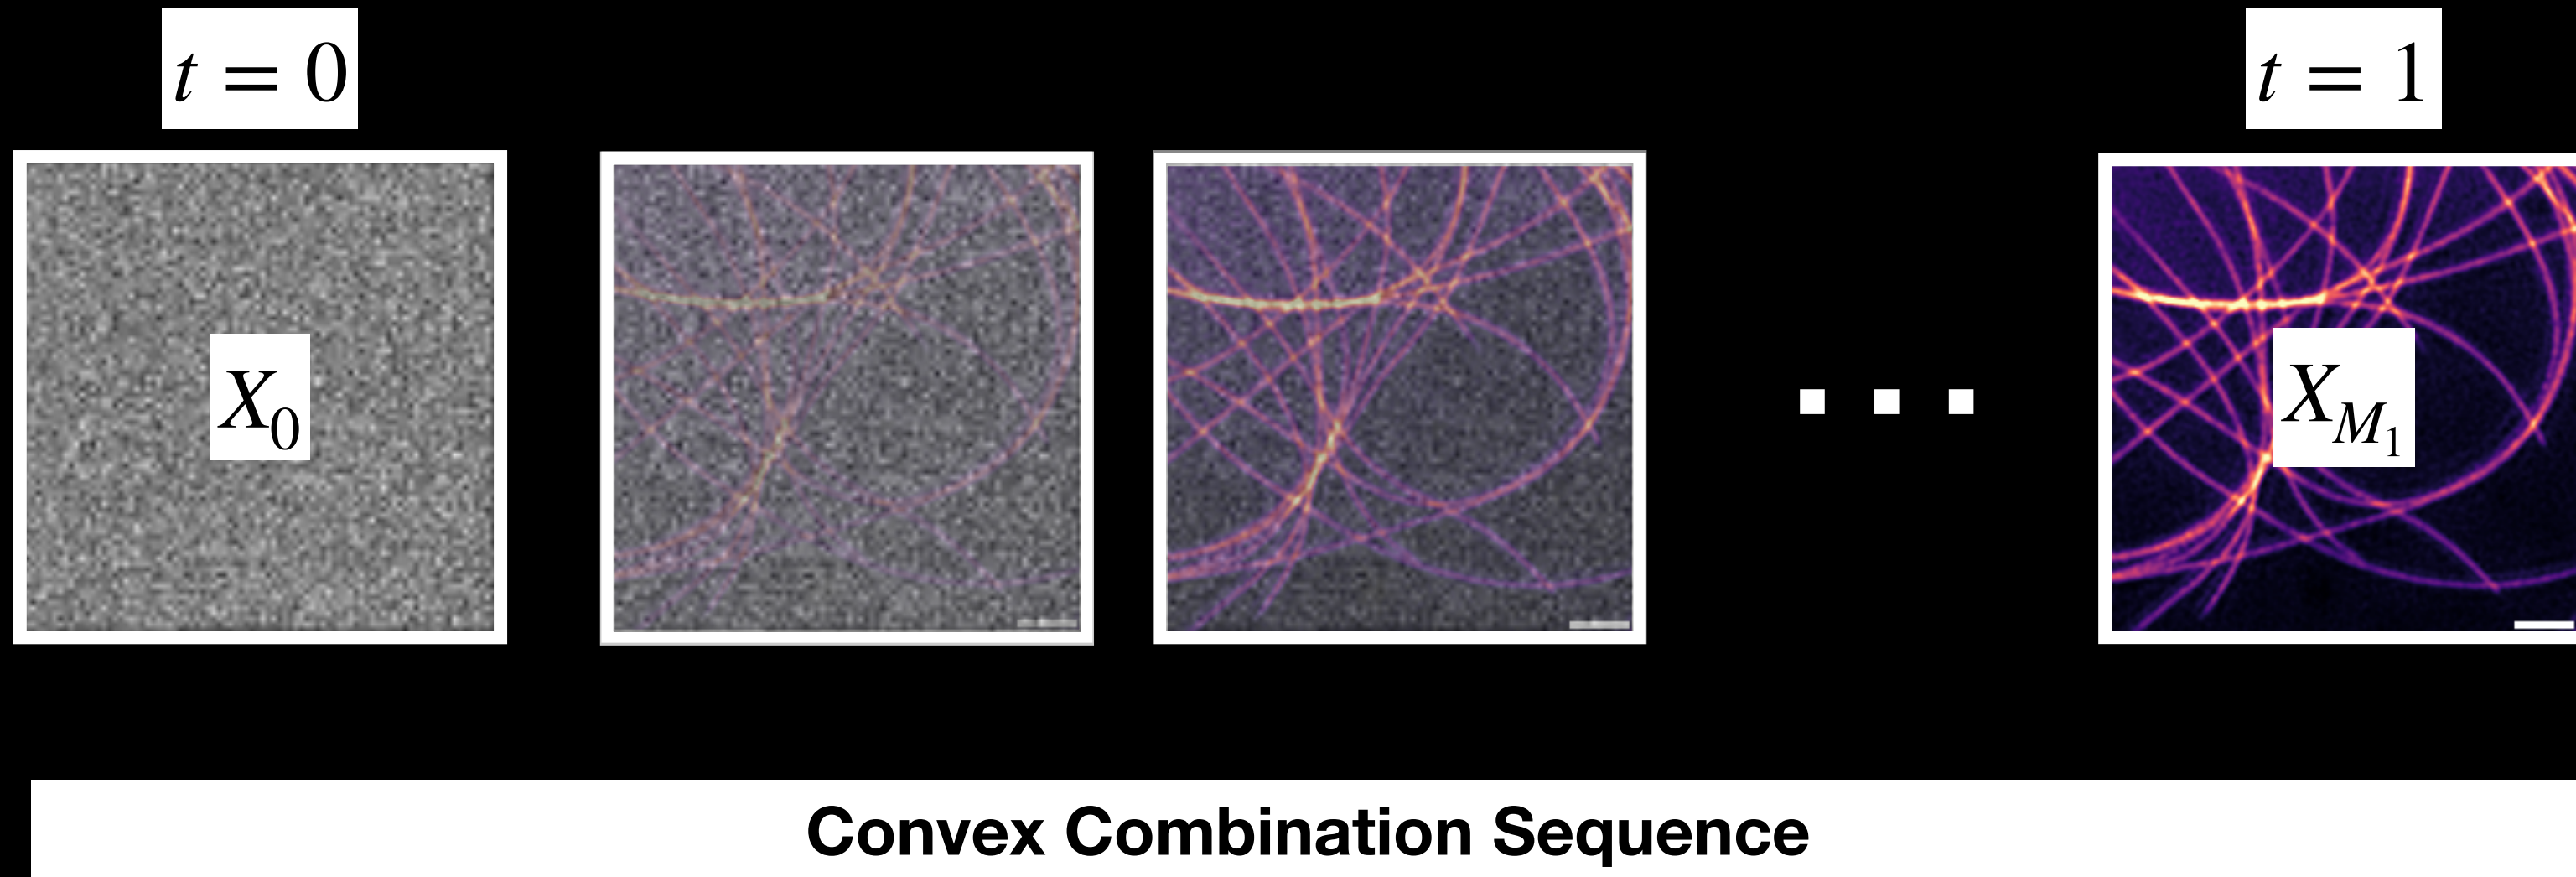
ResMatching - Training



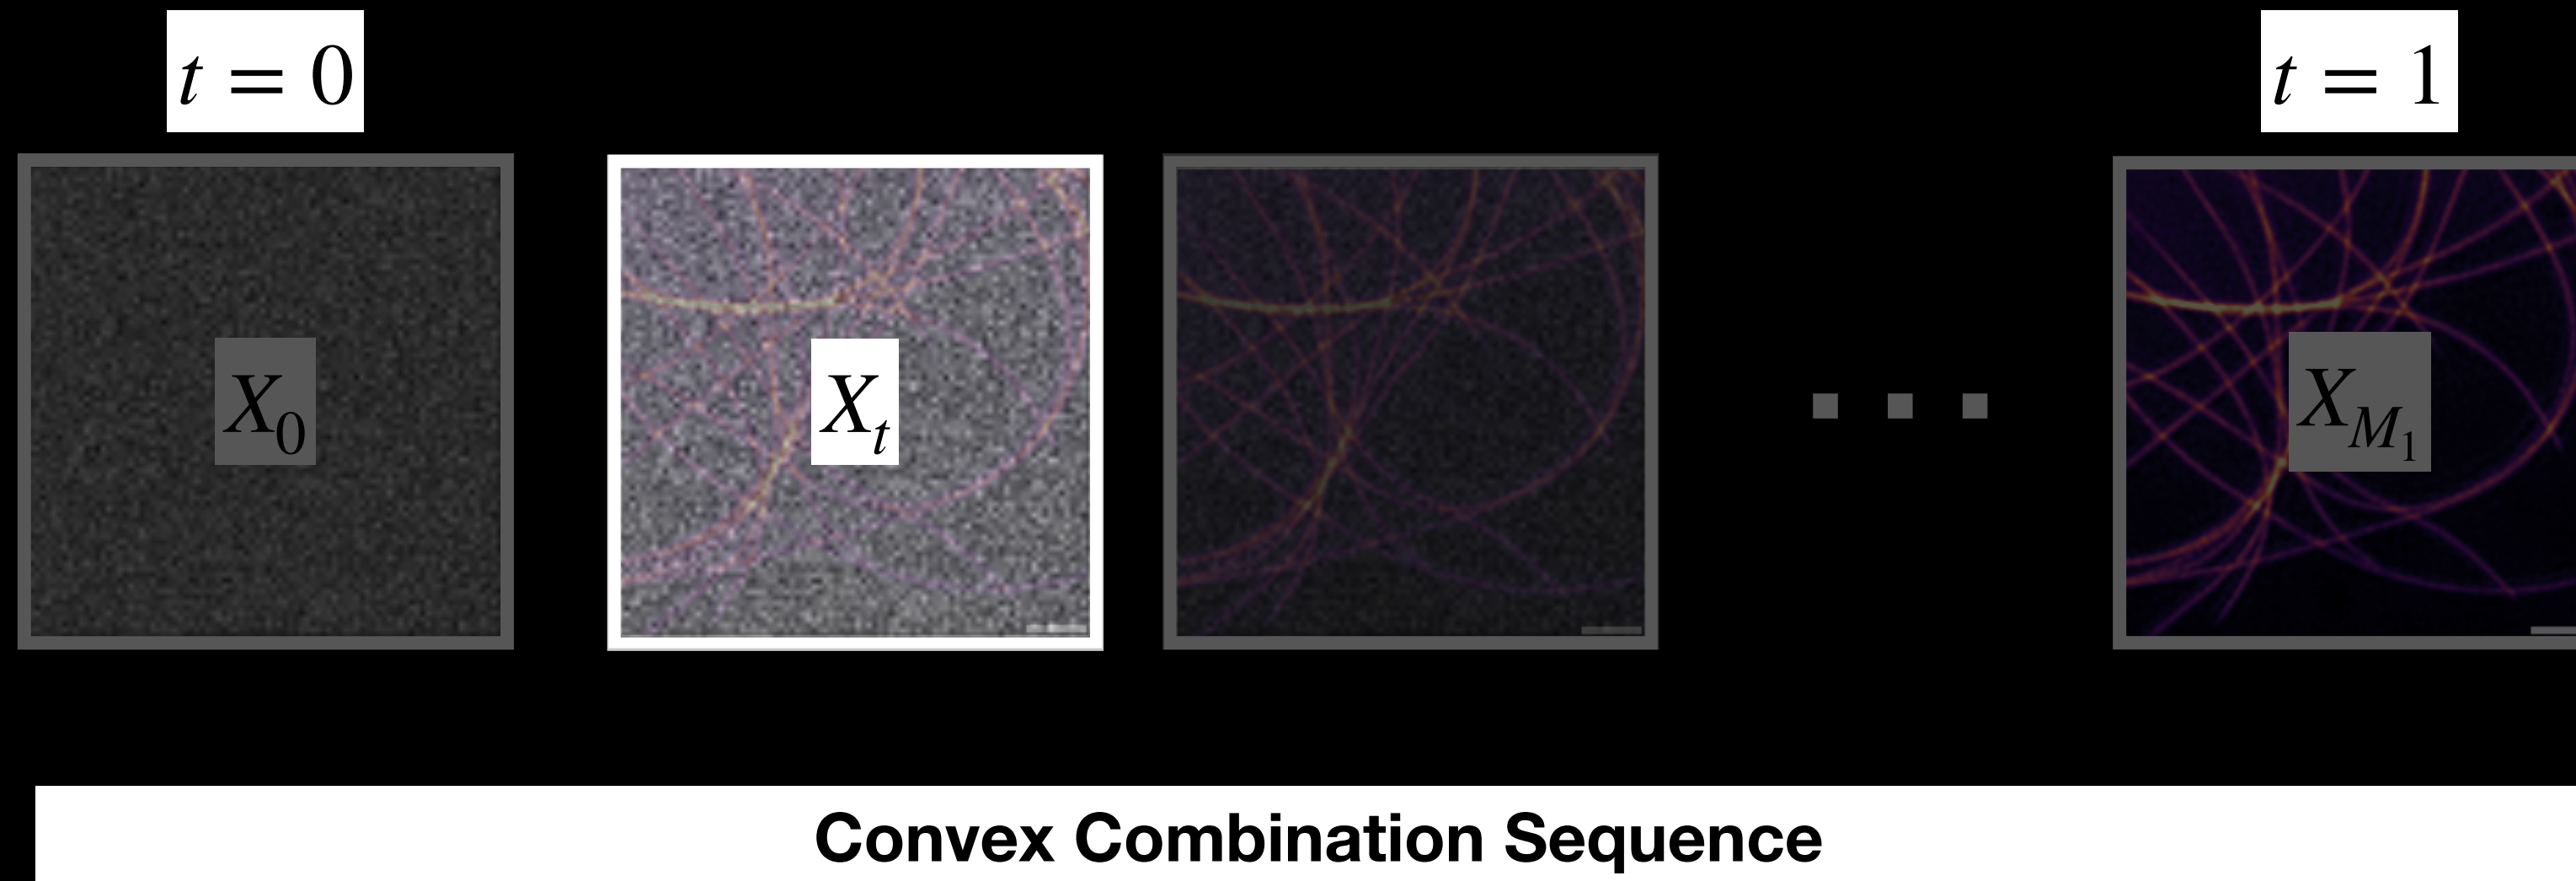
ResMatching - Training




ResMatching - Training




ResMatching - Training



ResMatching - Training



Convex Combination Sequence



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*




ResMatching - *Iterative Inference*




ResMatching - *Iterative Inference*



ResMatching - *Iterative Inference*



ResMatching - Posterior Sampling



Multiple inference cycles generates posterior samples

$\times K$

K posterior samples


$$\frac{1}{K} \sum^K \text{MMSE} \quad \text{Posterior variance} \quad \text{Var}(K)$$




BioSR Data (heavy Noise)

Input


CCP




ER







F-actin



MT-Noisy







GT



2x Upsampling under extreme uncertainty (noise)

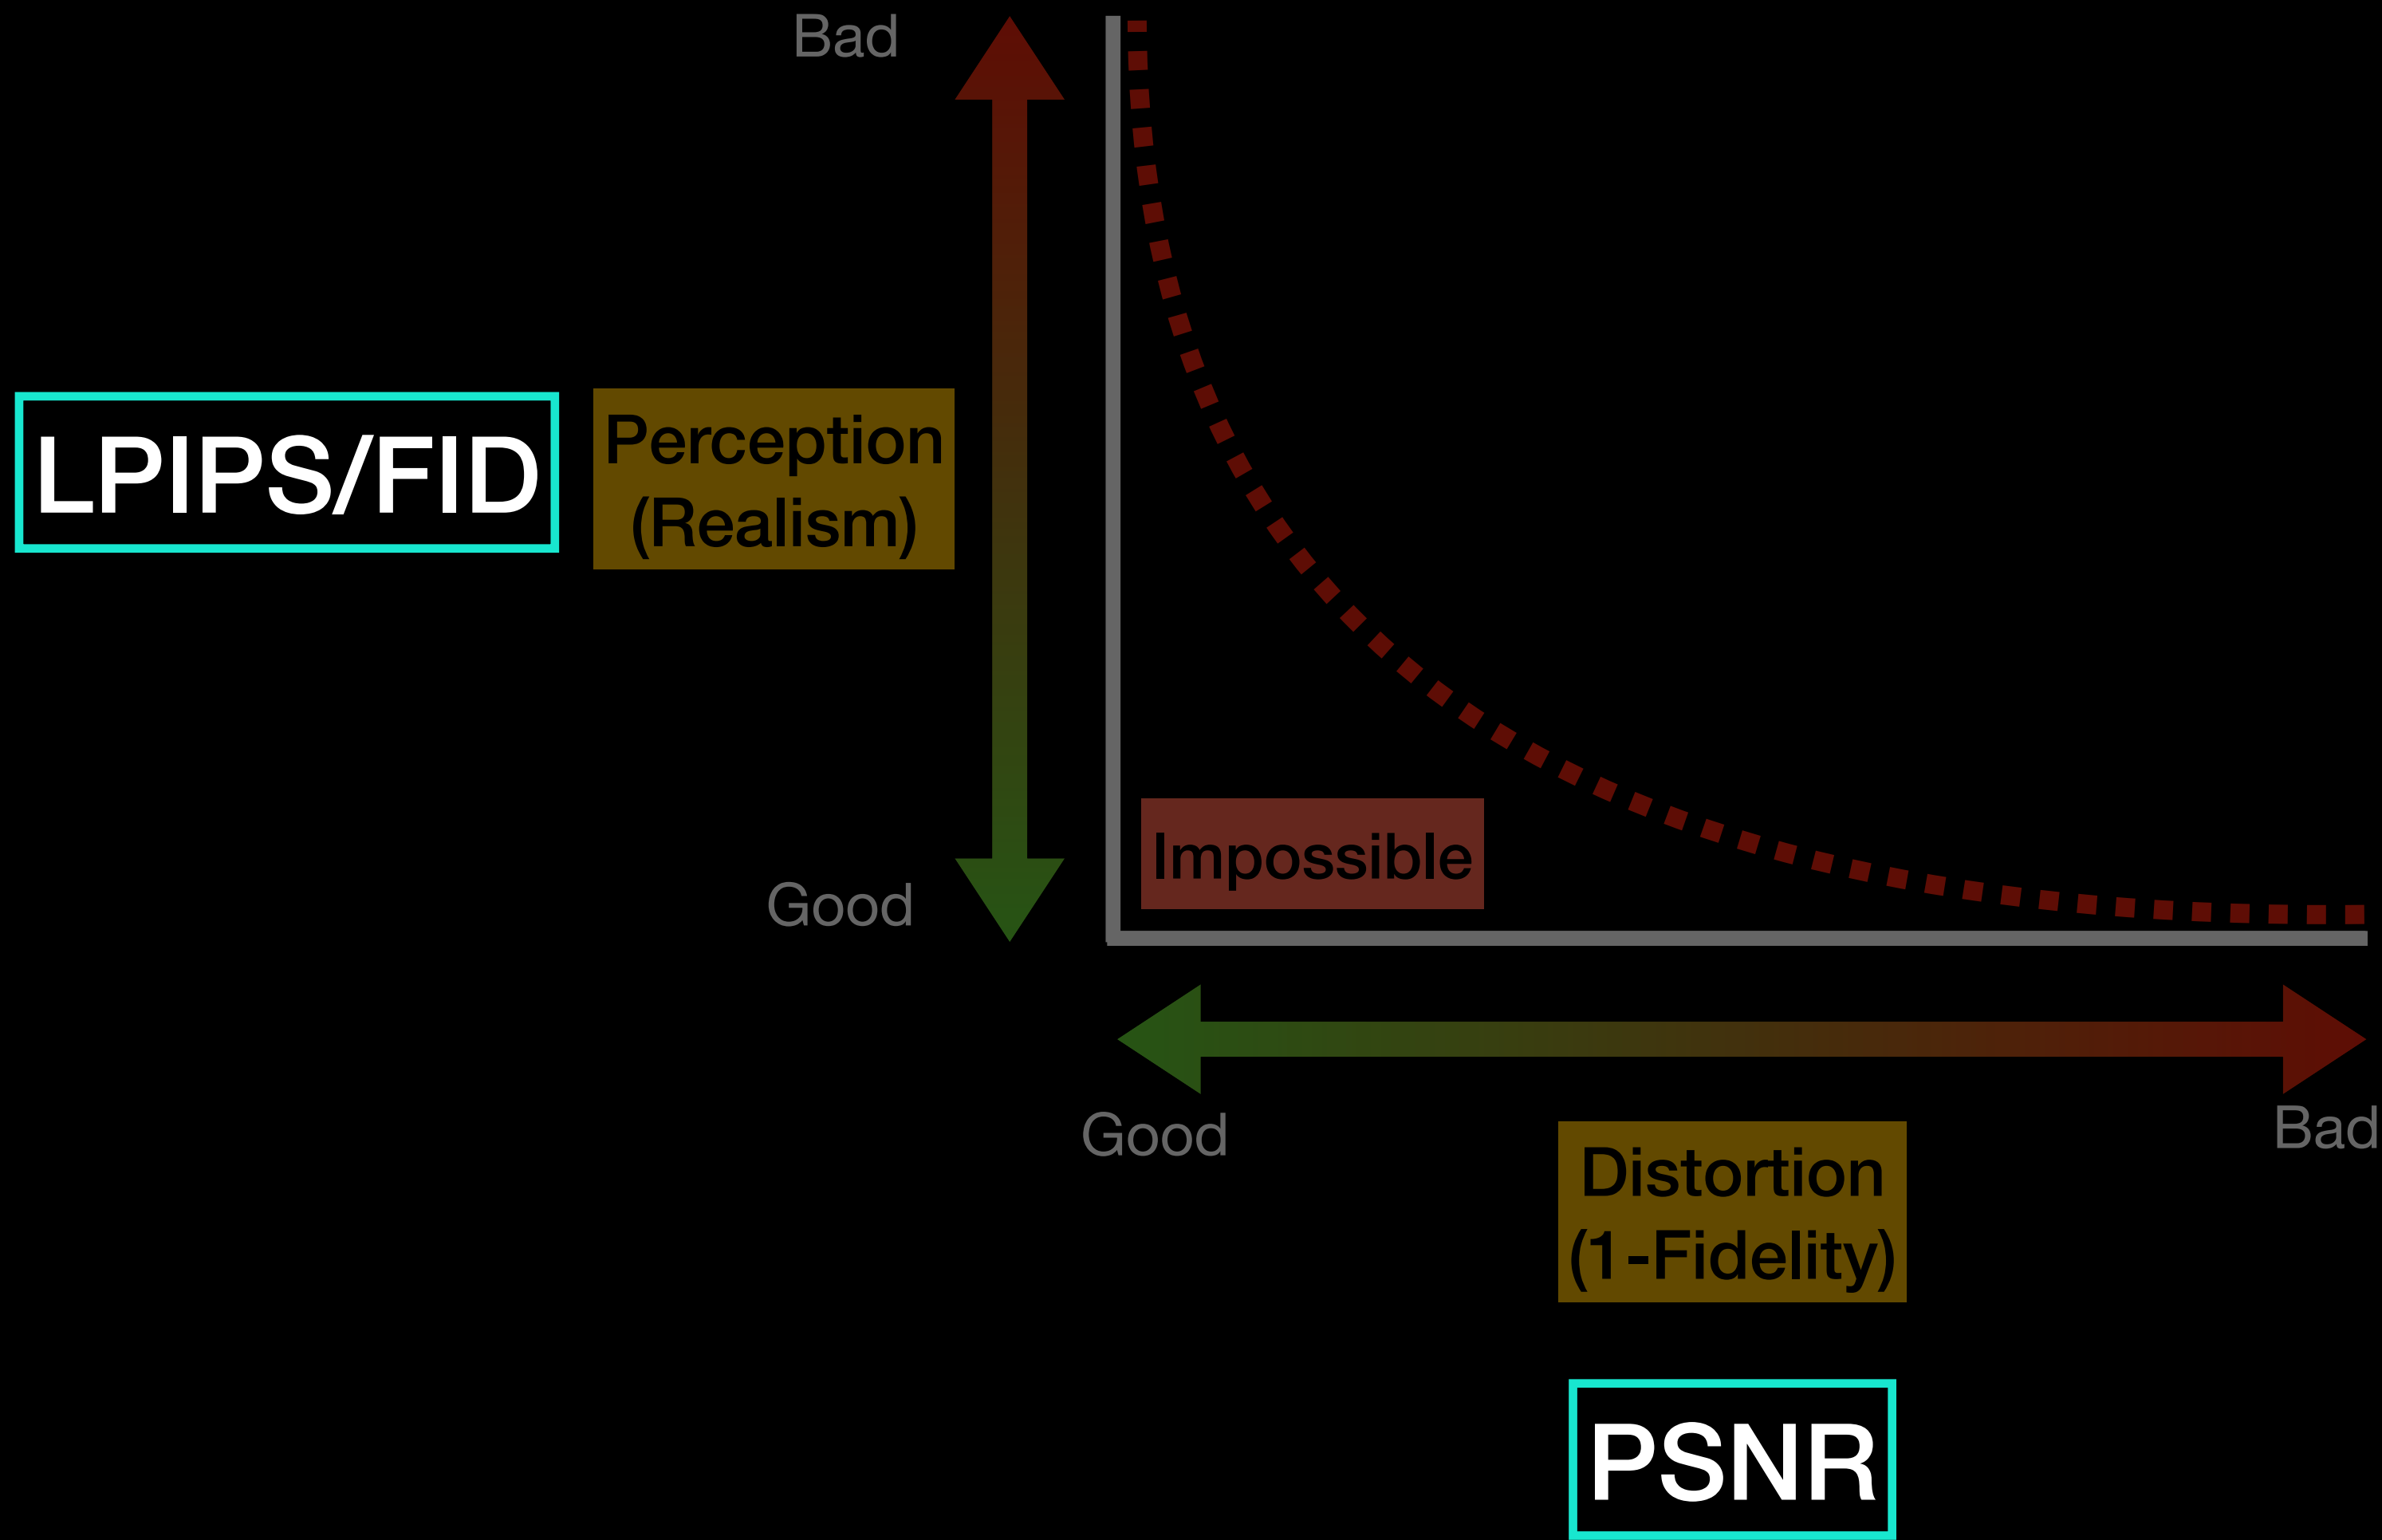


Baselines

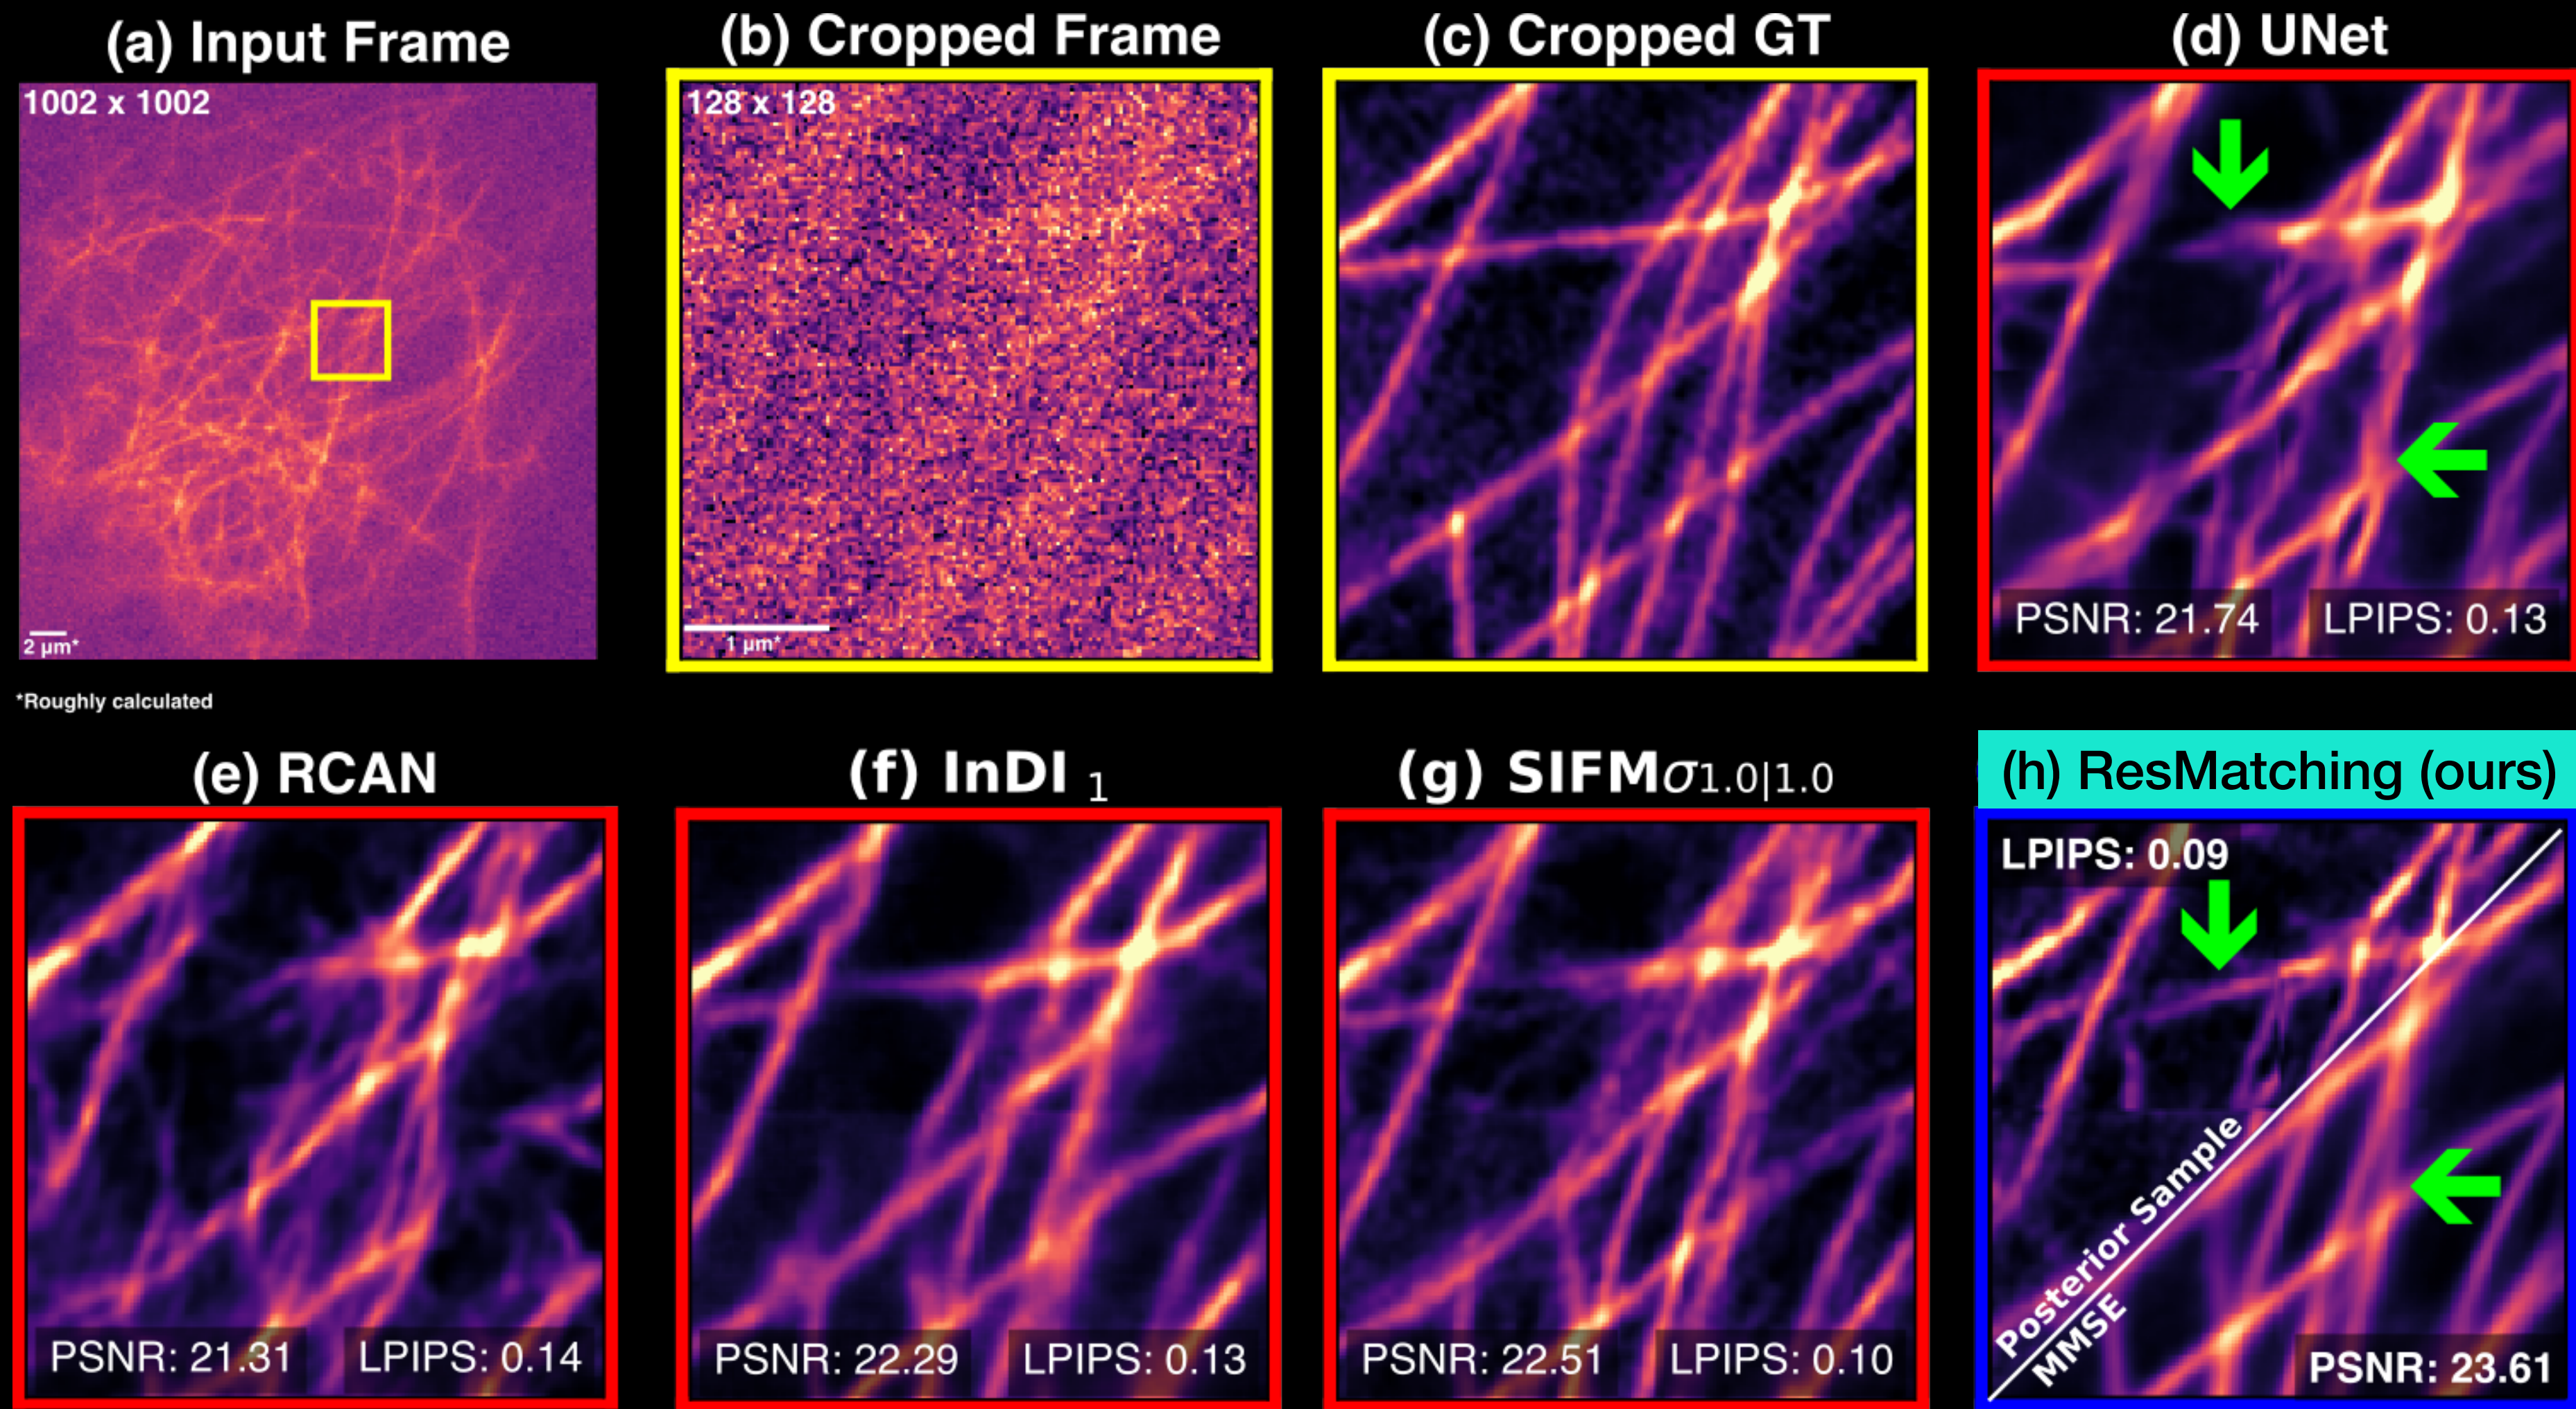
	Single Prediction	Posterior Sampling
One-shot	 UNet, RCAN, ESRGAN, InDI-1	 LVAE
Iterative	 InDI-20	 SIFM, ResMatching



Metrics



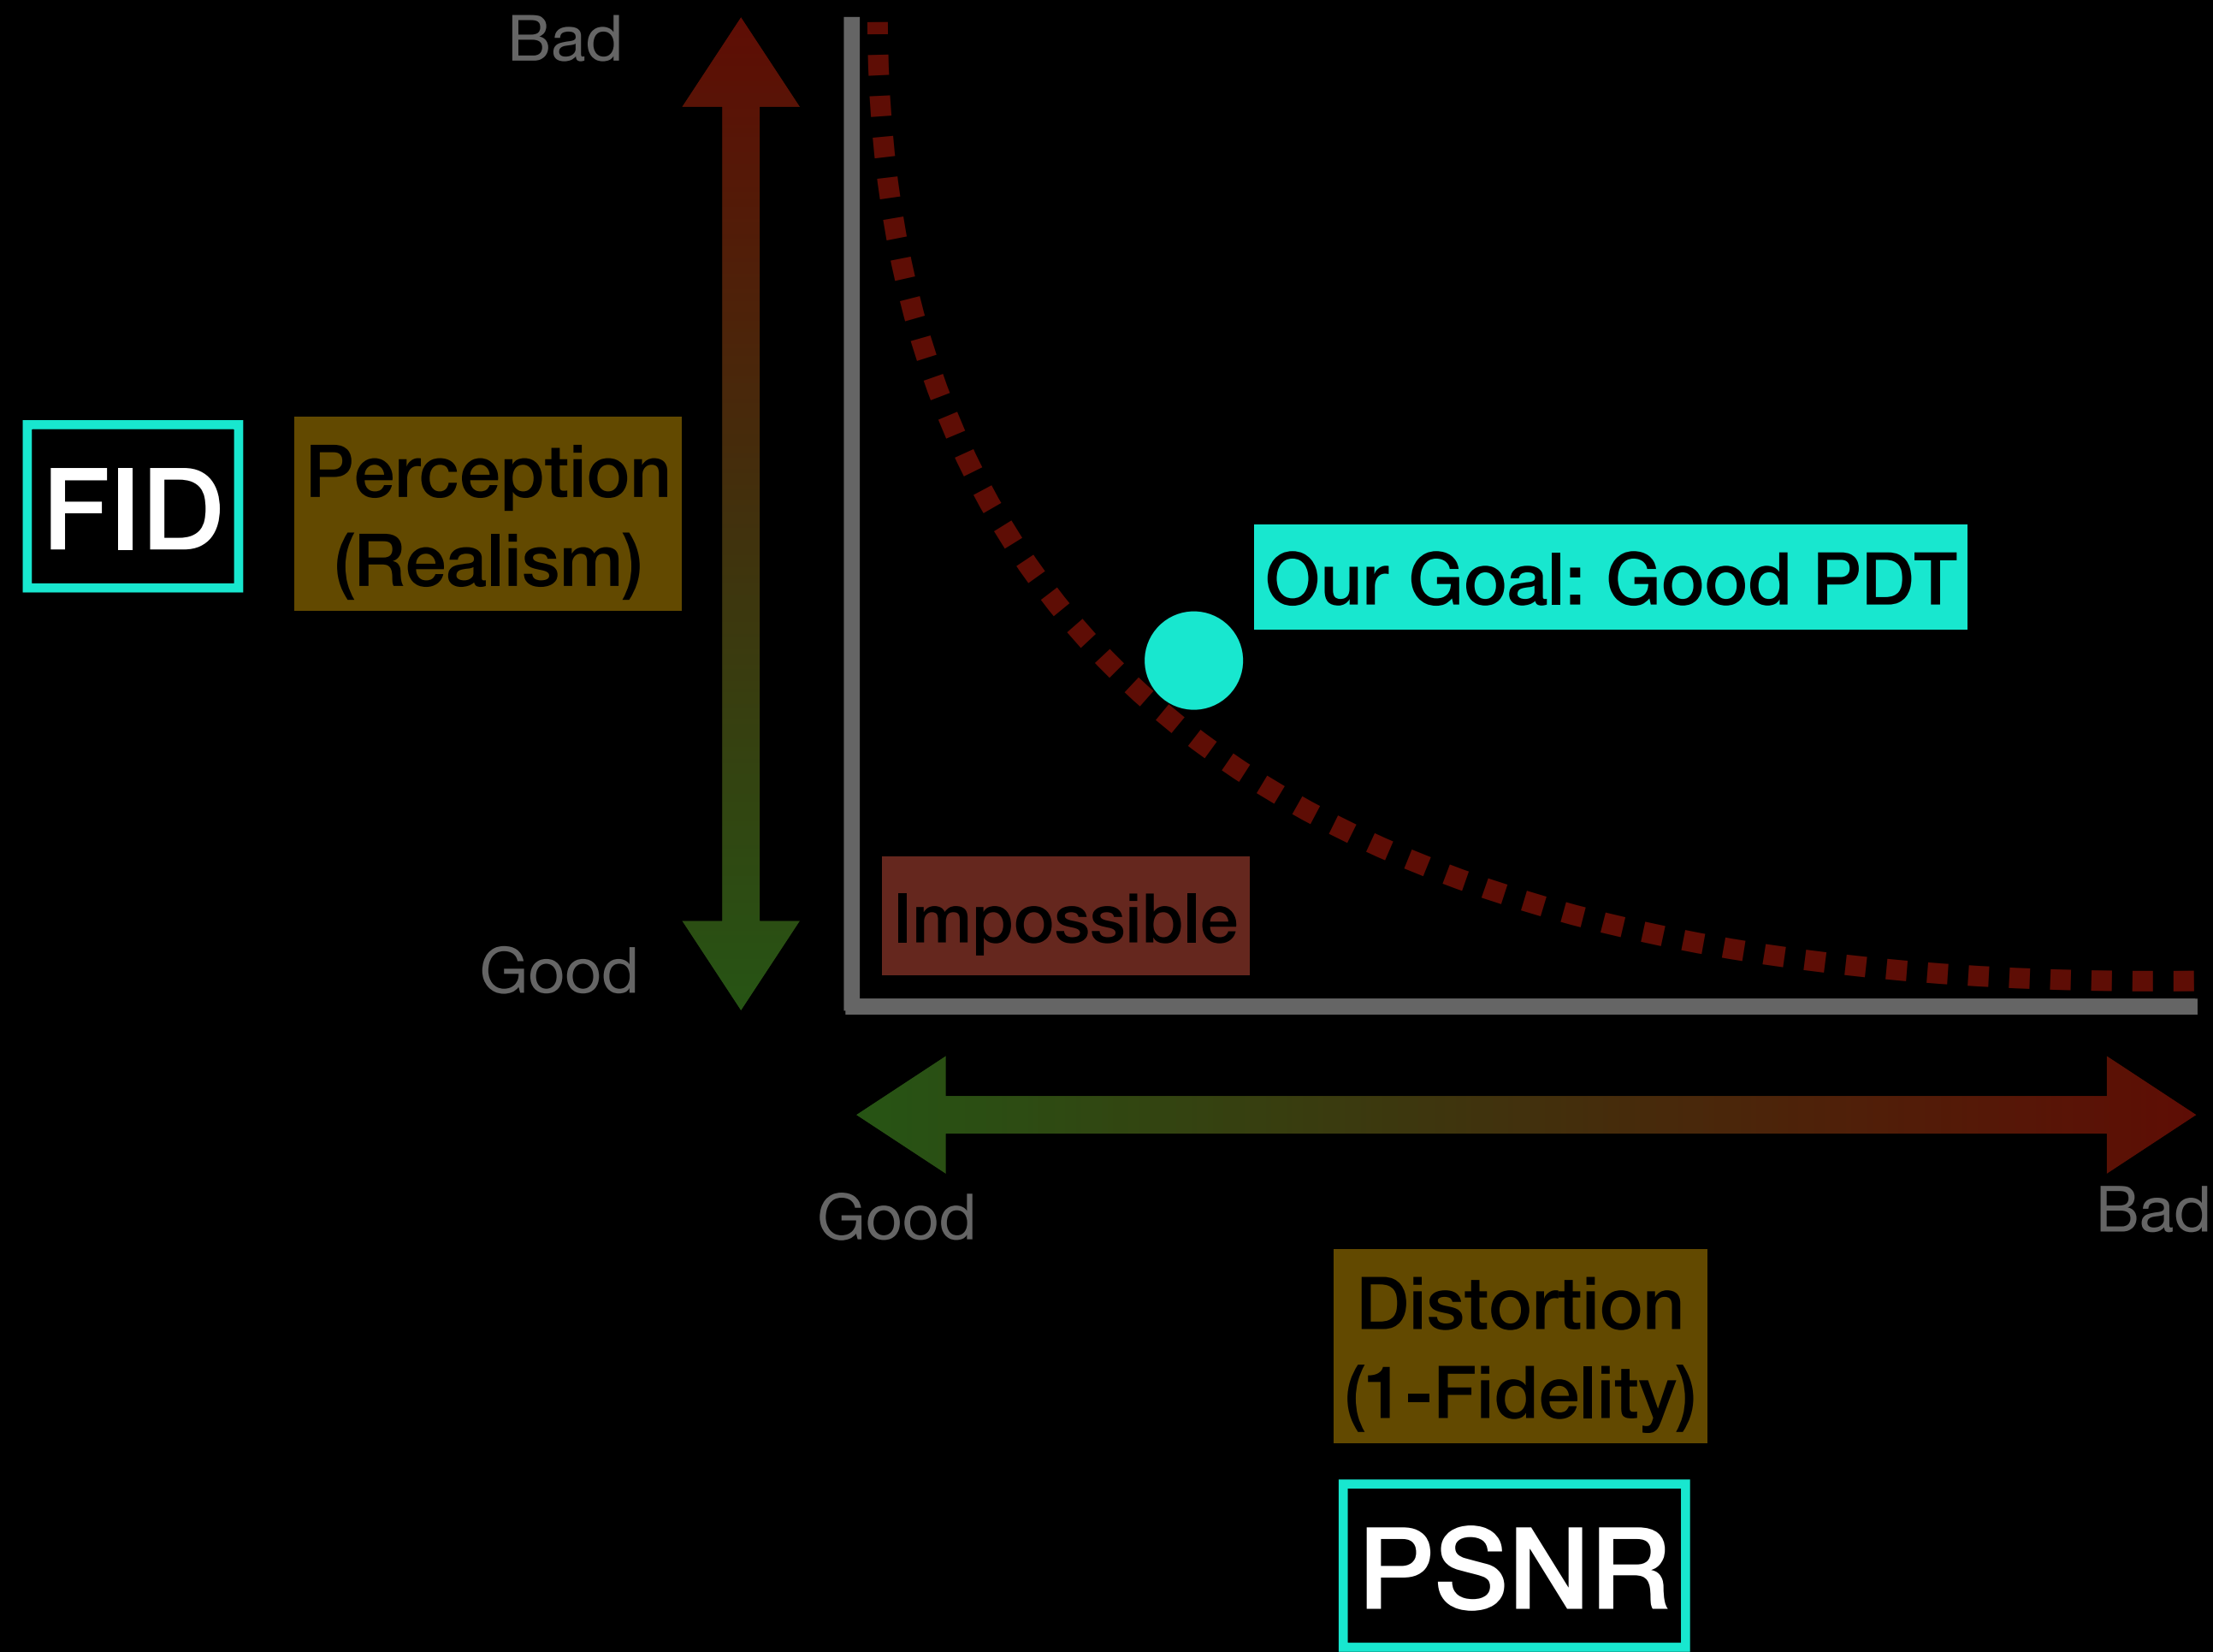
ResMatching Qualitative Results: MT-Noisy



*Roughly calculated






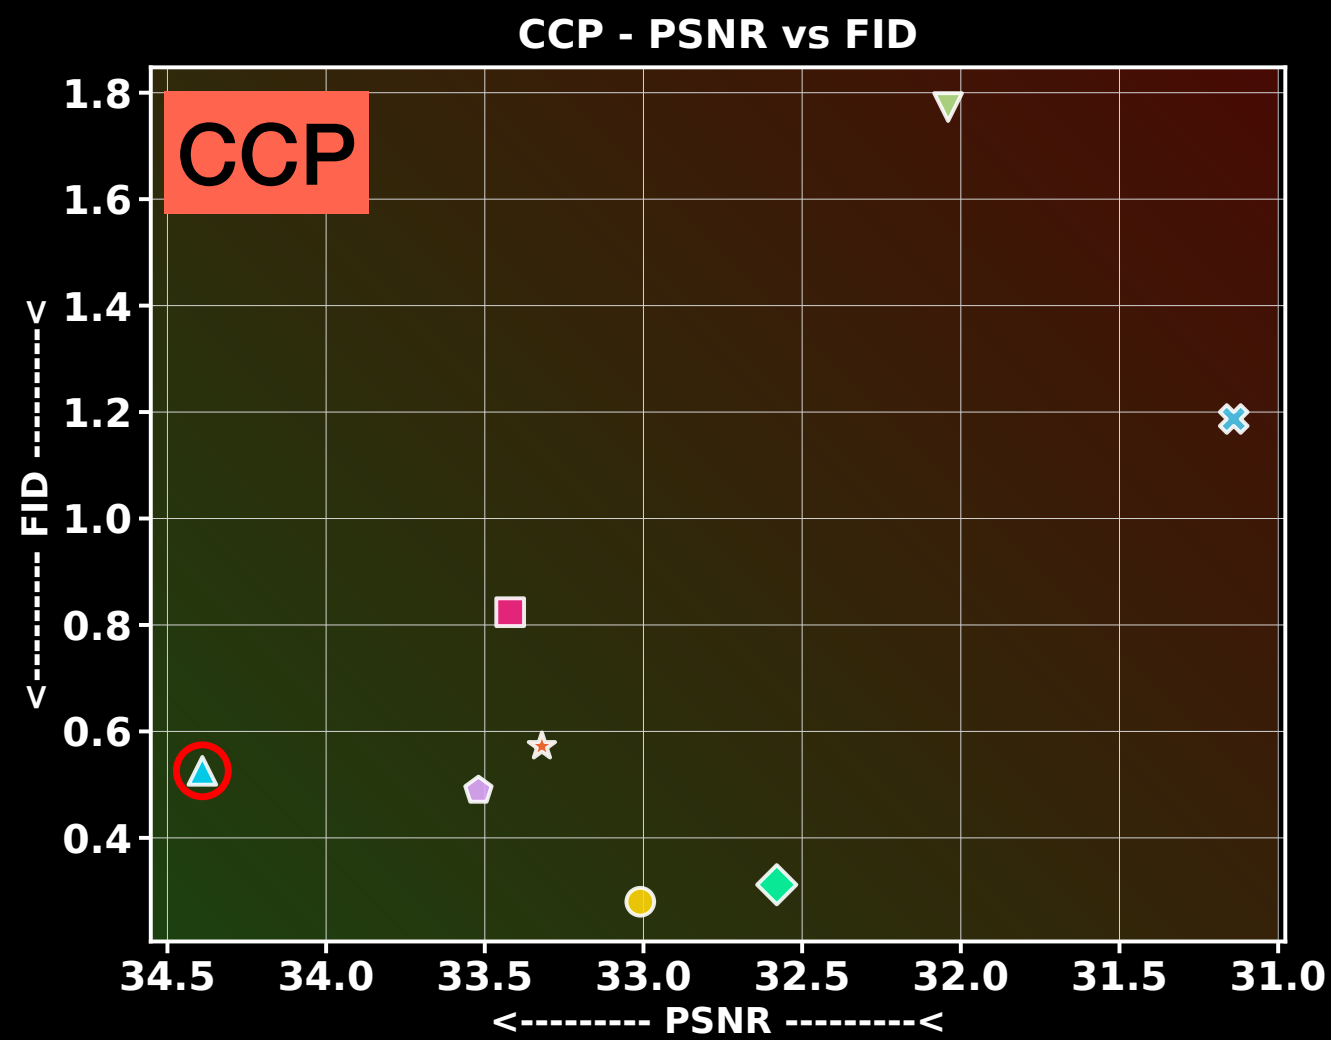
Quantitative: FID vs PSNR




Quantitative: FID vs PSNR

Worst

Best



- ★ UNet
- RCAN
- ✕ ESRGAN
- InDI₁
- ◆ InDI₂₀
- ▽ LVAE
- ◇ SIFM_{σ_{1.0}|1.0}
- ⊠ ResMatching





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?





MT Data GT

BioSR MT LR Data




ResMatching Pred



- ★ UNet
- RCAN
- ⊗ ESRGAN
- InDI₁
- ◆ InDI₂₀
- ▽ LVAE
- ⬠ SIFM $\sigma_{1.0|1.0}$
- ⊠ ResMatching




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




MT Data GT


BioSR MT LR Data




ResMatching Pred



MT-Noisy LR Data




Additional Synthetic Noise



- ★ UNet
- RCAN
- ⊗ ESRGAN
- InDI₁
- ◆ InDI₂₀
- ▽ LVAE
- ⬠ SIFM_{σ_{1.0}|1.0}
- ⊠ ResMatching




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




MT Data GT


BioSR MT LR Data



ResMatching Pred





MT-Noisy LR Data




Additional Synthetic Noise

ResMatching Pred




- ★ UNet
- RCAN
- ⊗ ESRGAN
- InDI₁
- ◆ InDI₂₀
- ▽ LVAE
- ⬠ SIFM_{σ_{1.0}|1.0}
- ⊙ ResMatching



Model Calibration



ResMatching - Model Calibration



$\times K$

K posterior samples


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

MMSE

$Var(K)$
Posterior variance



ResMatching - Model Calibration



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

K posterior samples

$$\frac{1}{K} \sum^K \text{MMSE} \quad \text{Posterior variance} \quad \text{Var}(K)$$



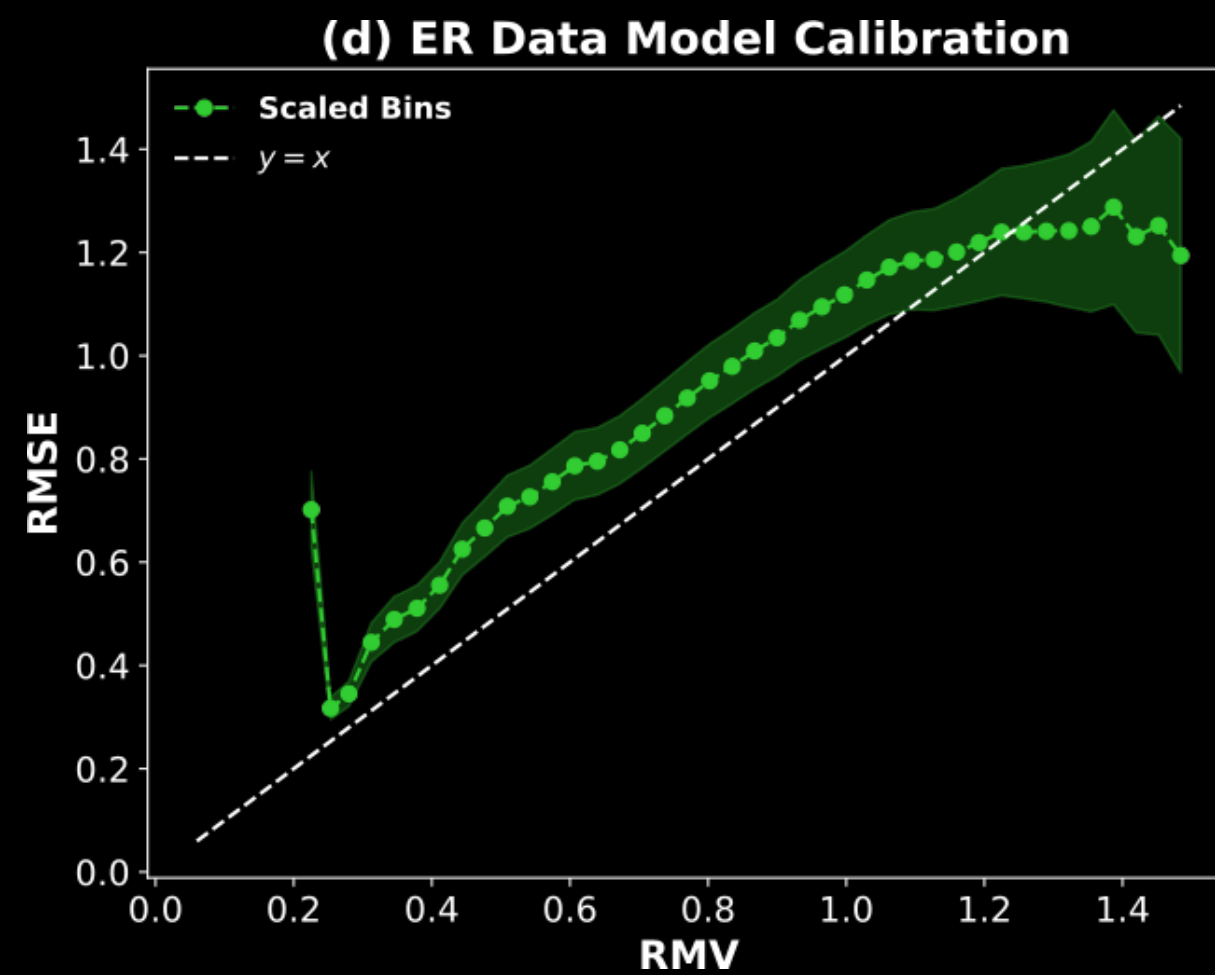



ResMatching - Model Calibration

Error

Posterior Variance



ResMatching - Model Calibration



Error

Posterior Variance


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



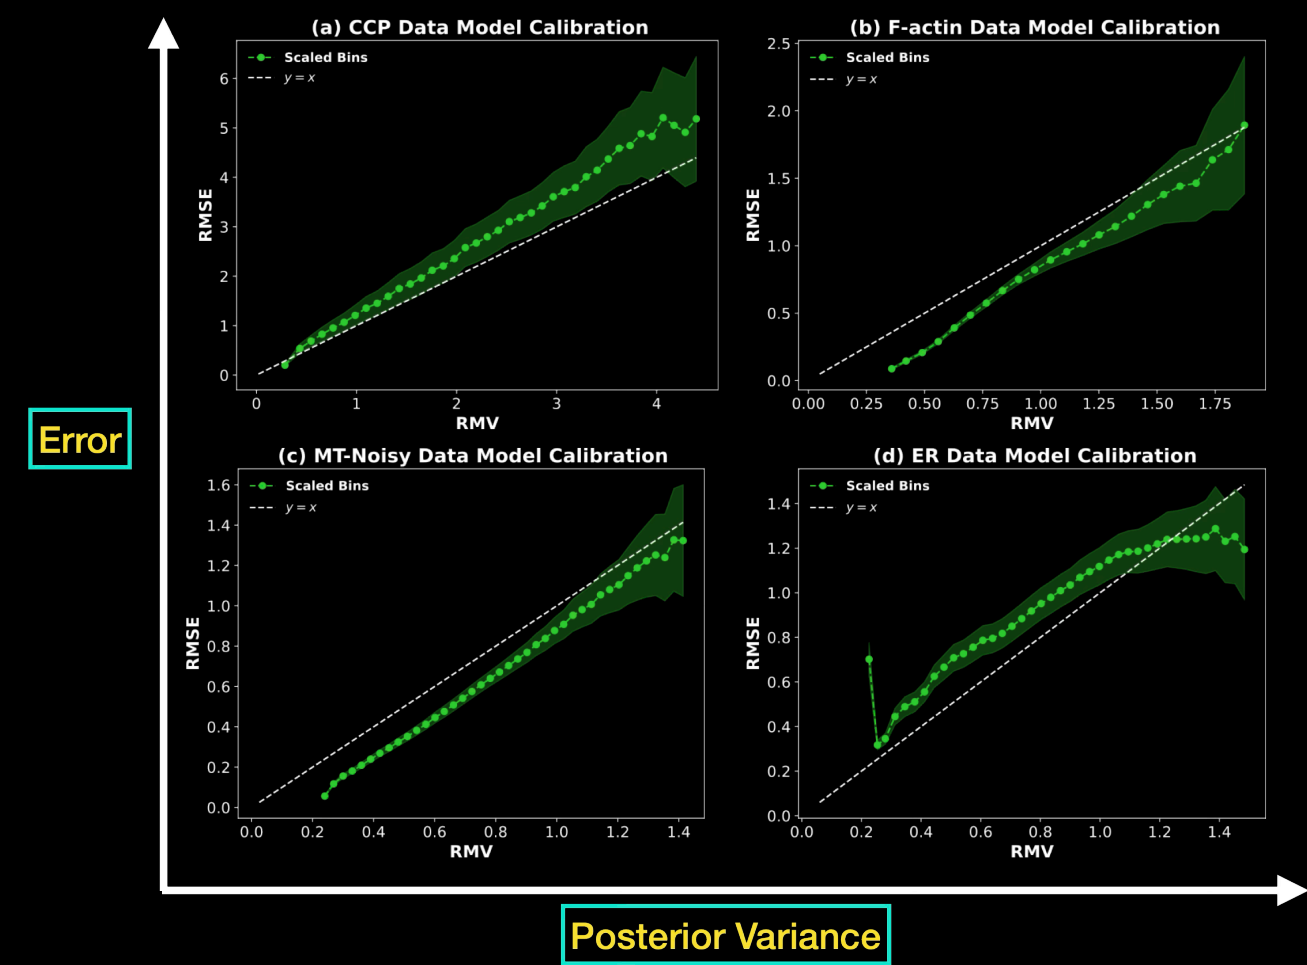
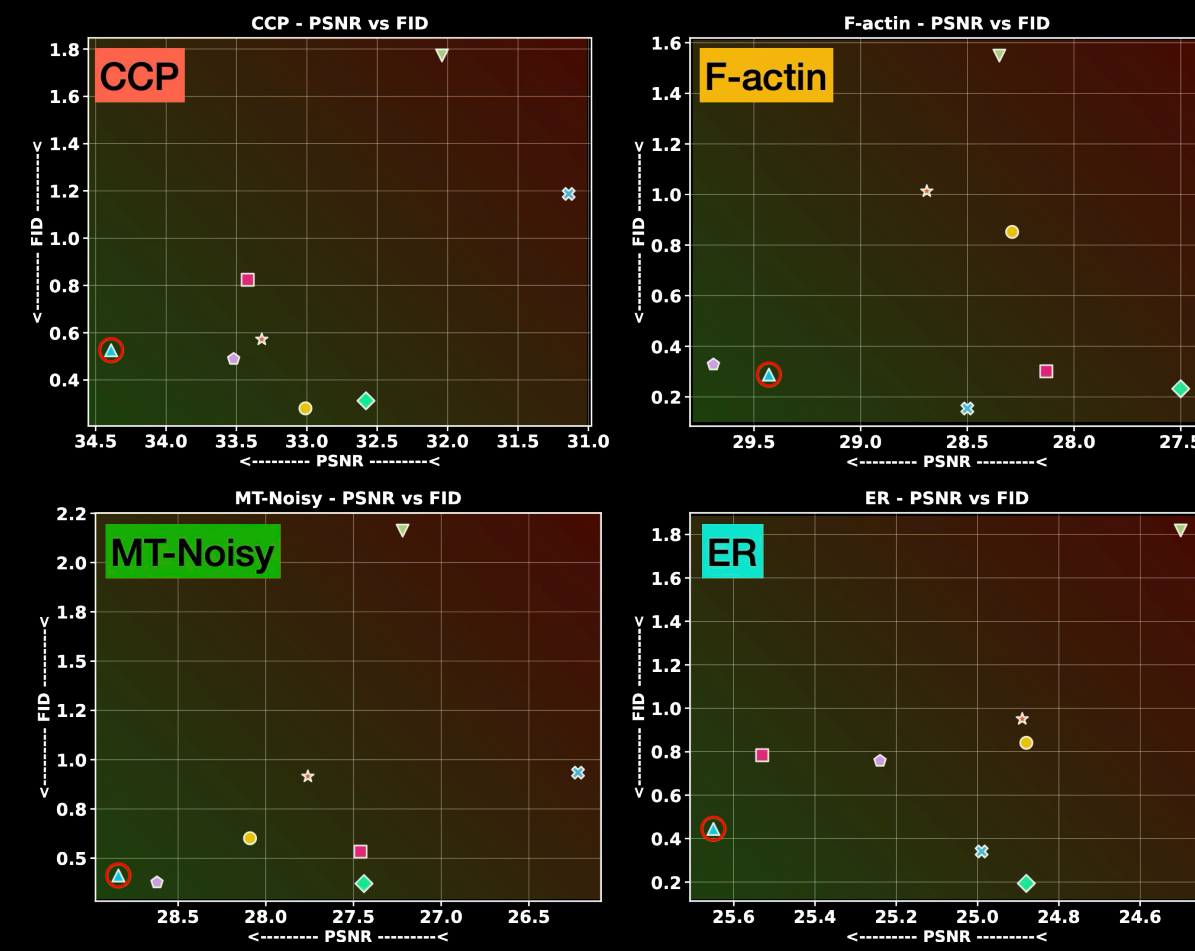

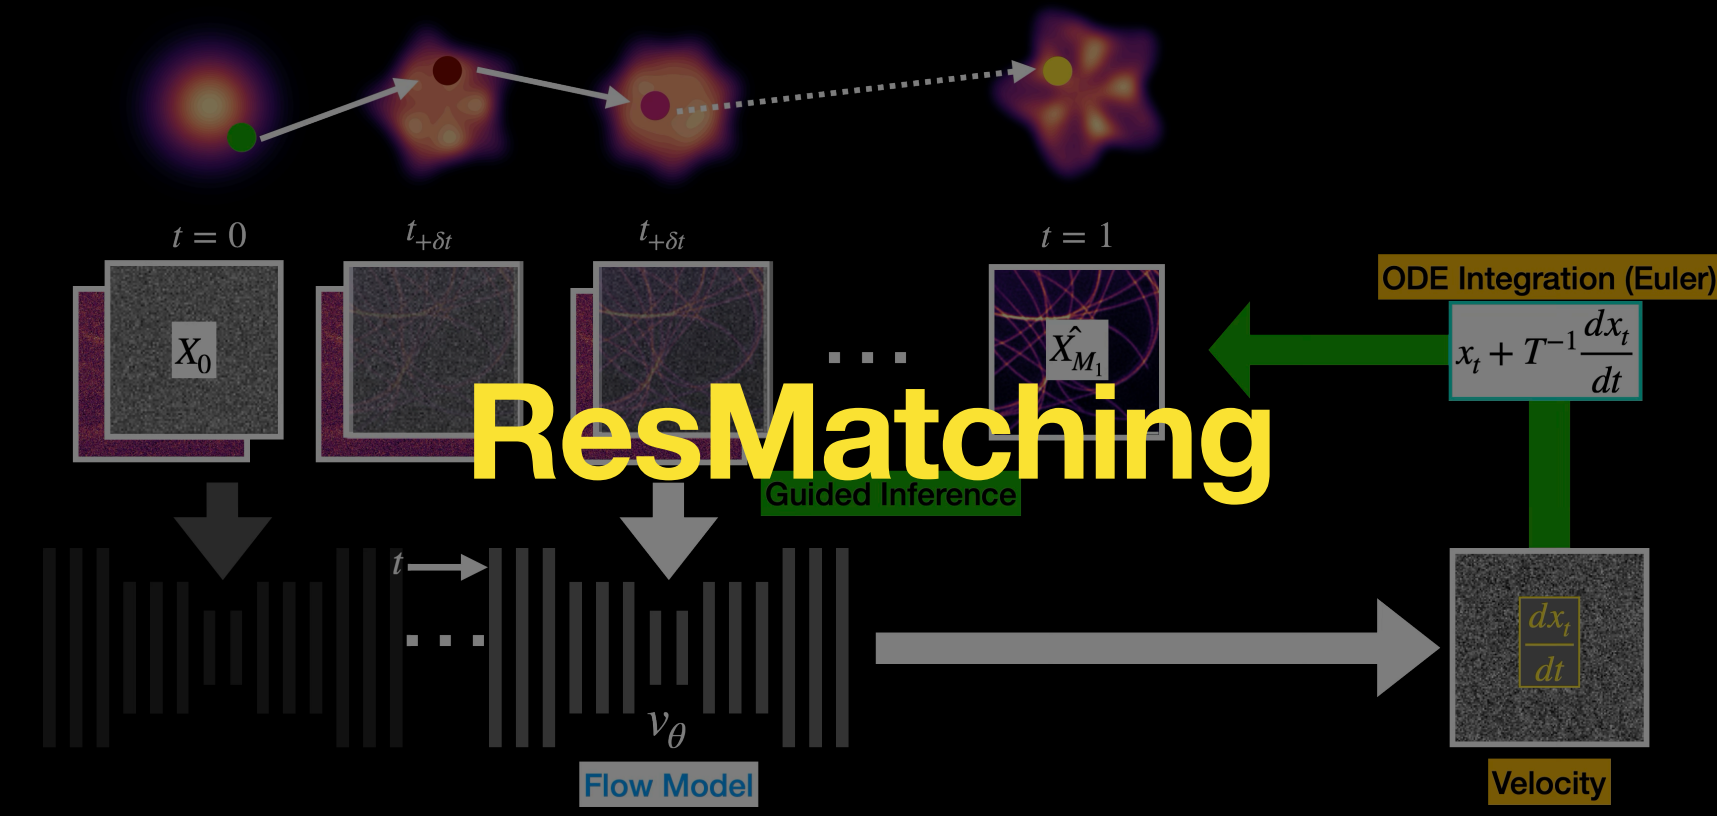

Summary



Summary





Summary



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*





Widefield Image

Confocal Image

Manipulation of observable frequency ranges

Interpolation

CSR



Noisy LR Image

SR Image

Extrapolation to non-observed frequency ranges

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).

Paper, more results and interactive tool



Wanna know more?

Find our poster! 😎

If I can find a spot

