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

Anirban Ray[1,2], Vera Galinova[1], and Florian Jug[1]
[1]Human Technopole, Milan, Italy
[2]Technische Universität Dresden, Germany

IEEE ISBI 2026
2026 IEEE International Symposium on Biomedical Imaging
April 8–11, 2026 | London, UK

Abstract

Computational **S**uper-**R**esolution (**CSR**) fundamentally relies on learning a prior to extrapolate unobserved high-frequency details from low-resolution micrographs. With modern data-driven methods, stronger priors can be learned, improving CSR performance. We introduce **ResMatching**, a CSR approach based on guided conditional flow matching to learn such priors. Evaluated on **4** biological structures from the BioSR dataset against **7** baselines, **ResMatching** consistently achieves the best trade-off between data fidelity and perceptual realism. It is particularly effective in challenging settings with **noisy** low-resolution inputs, where learning strong priors is difficult. Furthermore, ResMatching enables sampling from an implicitly learned, well-calibrated posterior, providing pixel-wise uncertainty estimates to help identify unreliable predictions.

Background of the Problem

Super-Resolution (Inverse Problem)

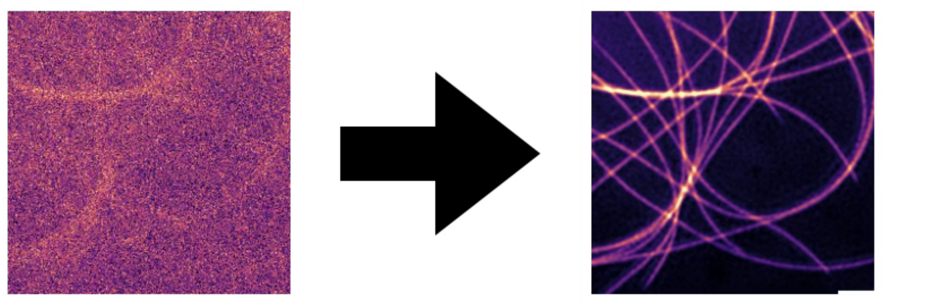

No Unique Solution (ill-posed)!

(Typical) one-shot restoration approaches

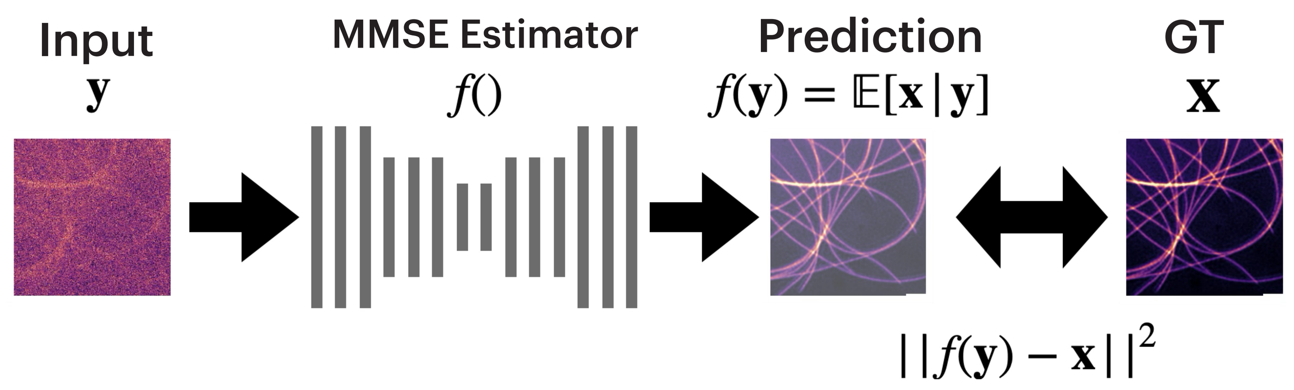

MSE Loss / L2 Loss / Quadratic Error

Regression-to-Mean Effect (Blurry Structures) [1]

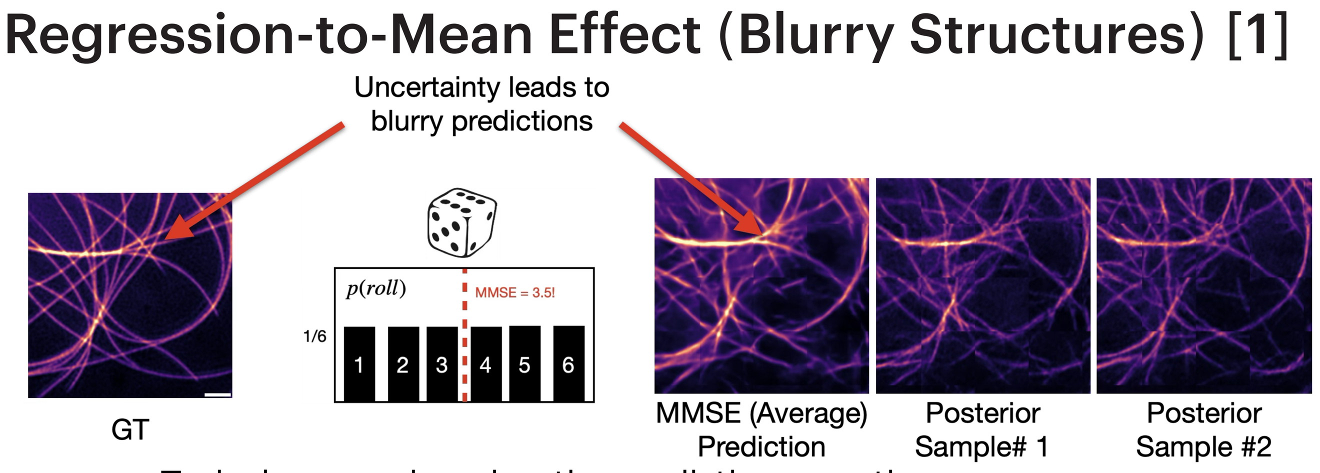

Typical approach makes the prediction smooth

Why it happens and how can we address this issue?

Many plausible solutions

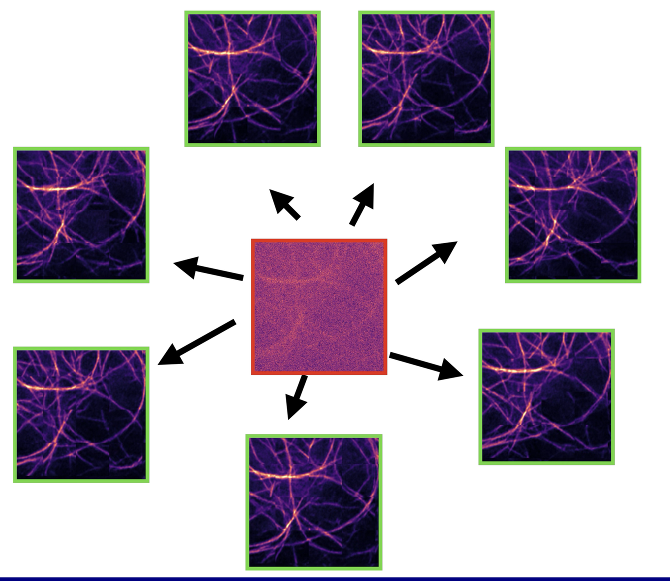

Perception-Distortion Tradeoff

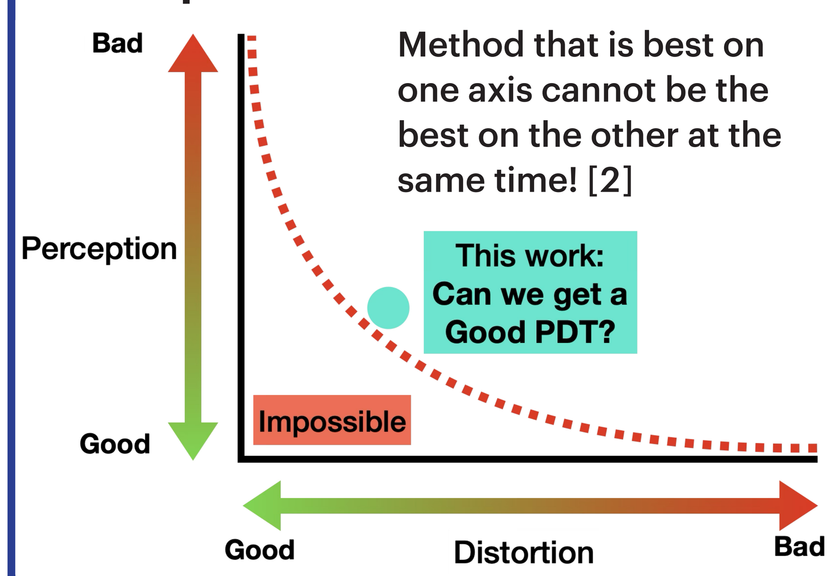

Restricting Averaging

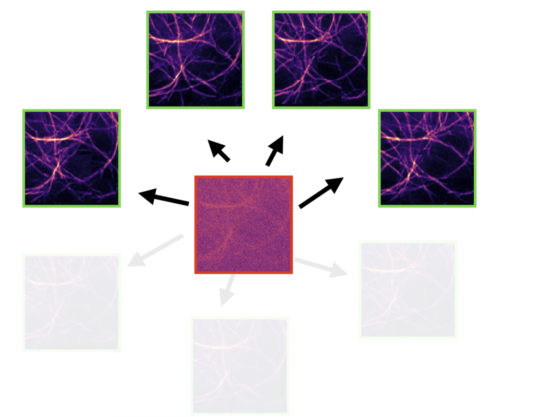

Can we restrict the averaging space?

Iterative Prediction reduces averaging effect

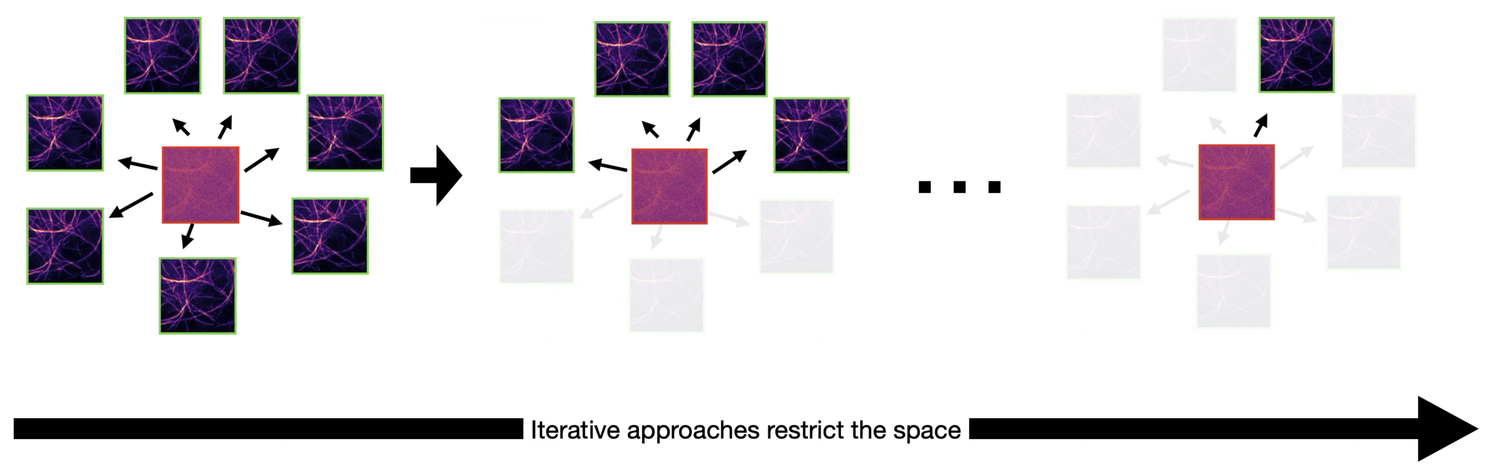

Iteratively predicting from less degraded images can be thought of as a step-wise reduction in uncertainty thereby restricting averaging!

Solution: *Guided Conditional Flow Matching [3]* — an ODE based iterative CSR approach

Training Data Generation

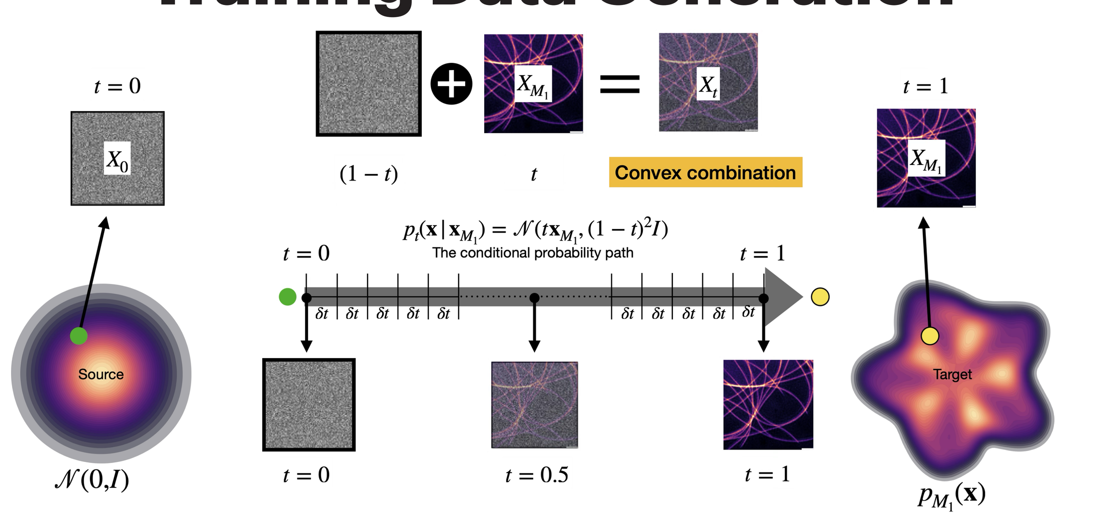

Training input data is a linear interpolation of noise and clean data

Training Scheme

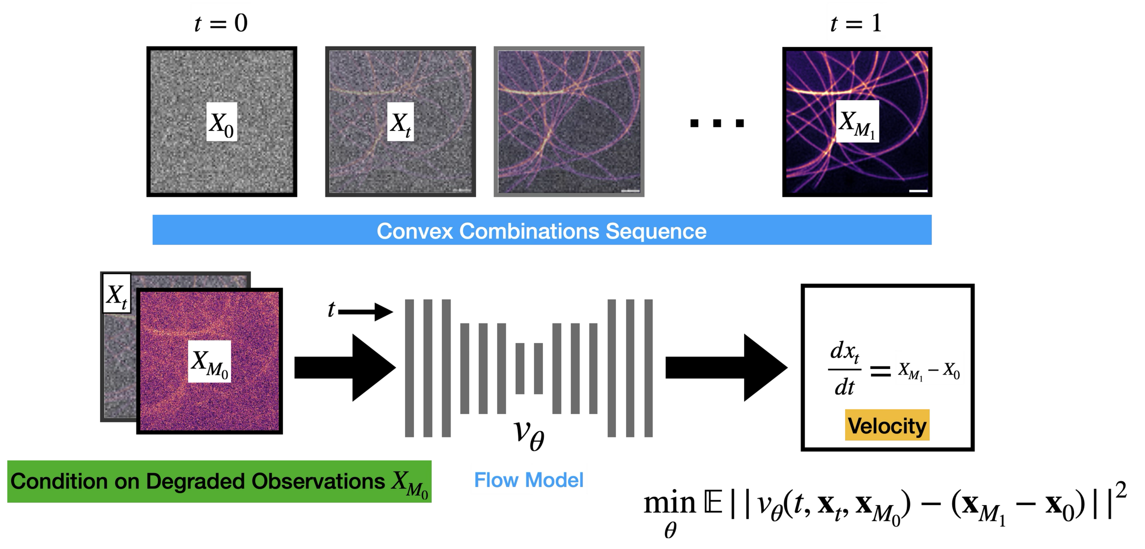

$$\min_{\theta} \mathbb{E}\,||\,v_\theta(t, \mathbf{x}_t, \mathbf{x}_{M_0}) - (\mathbf{x}_{M_1} - \mathbf{x}_0)\,||^2$$

Iterative Inference Scheme

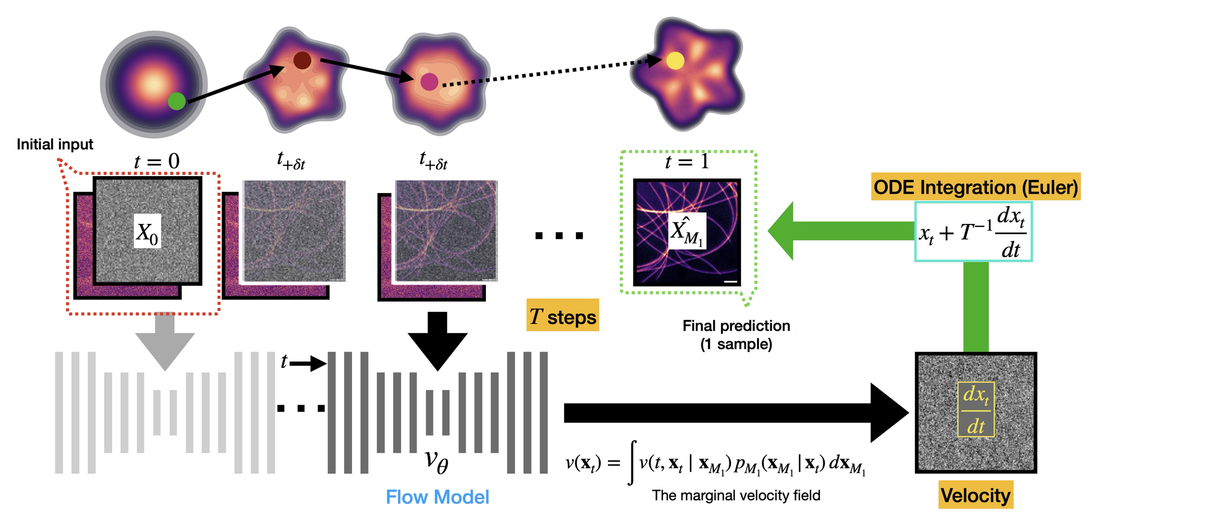

Our method iteratively produces crisp structures and allows for posterior sampling

CSR on BioSR Data [4] under strong noise

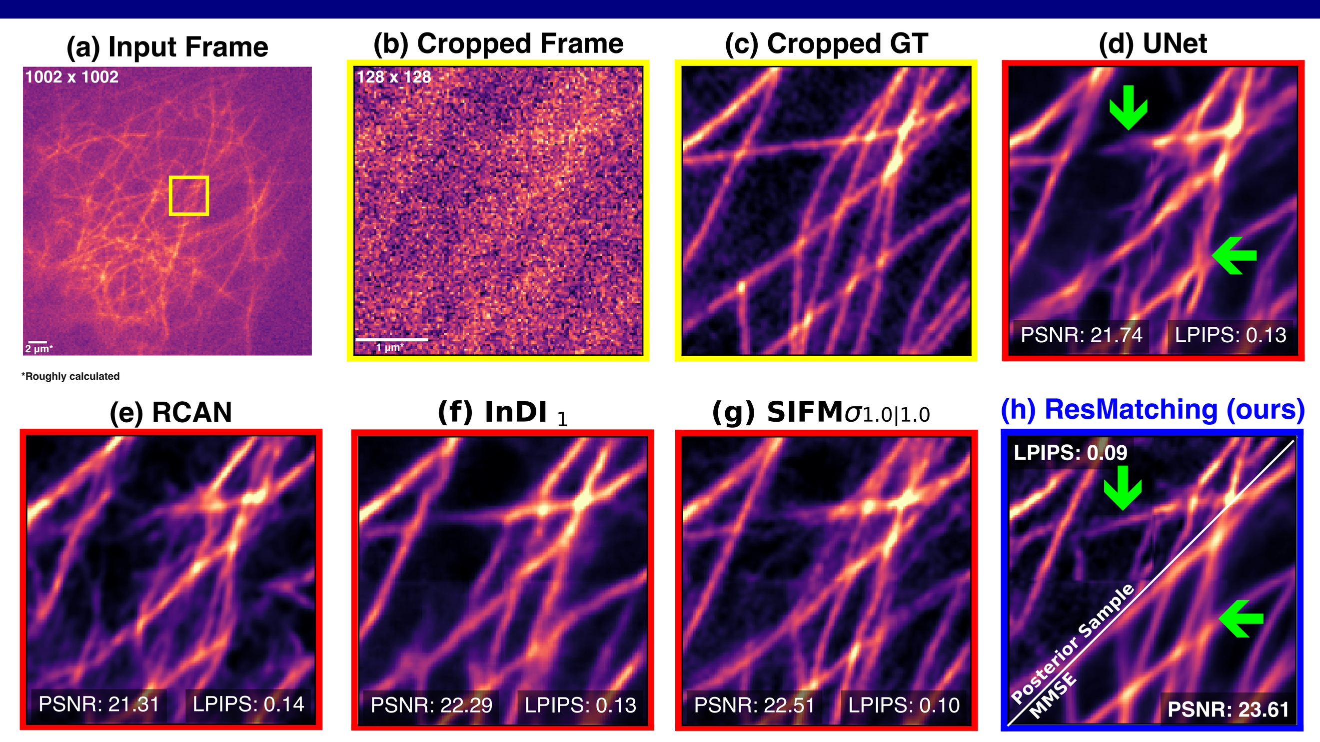

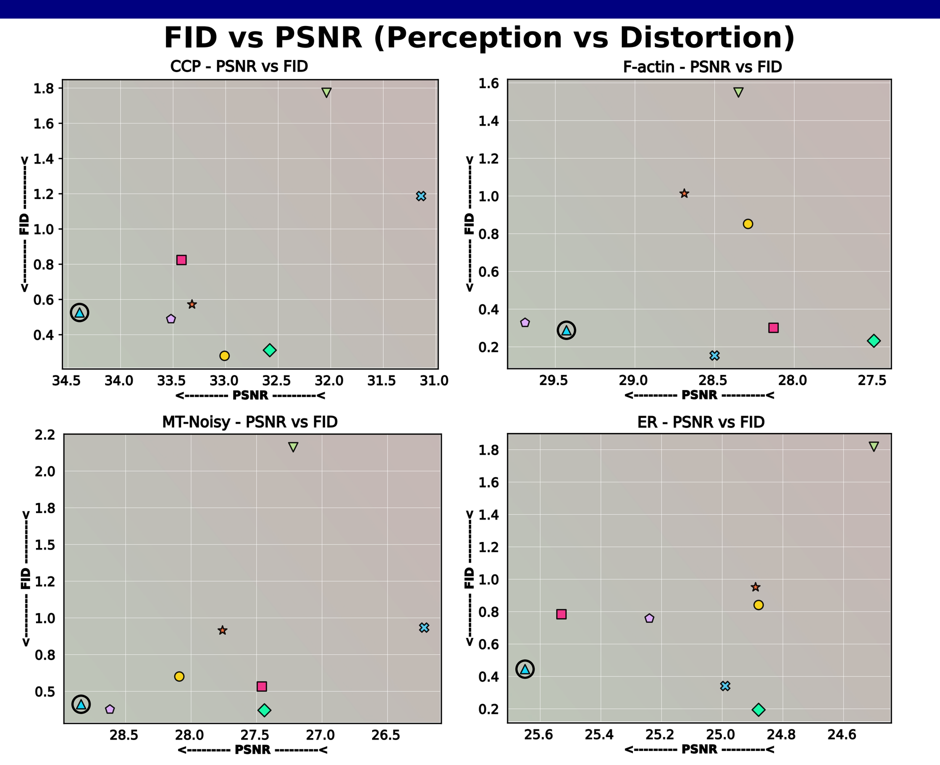

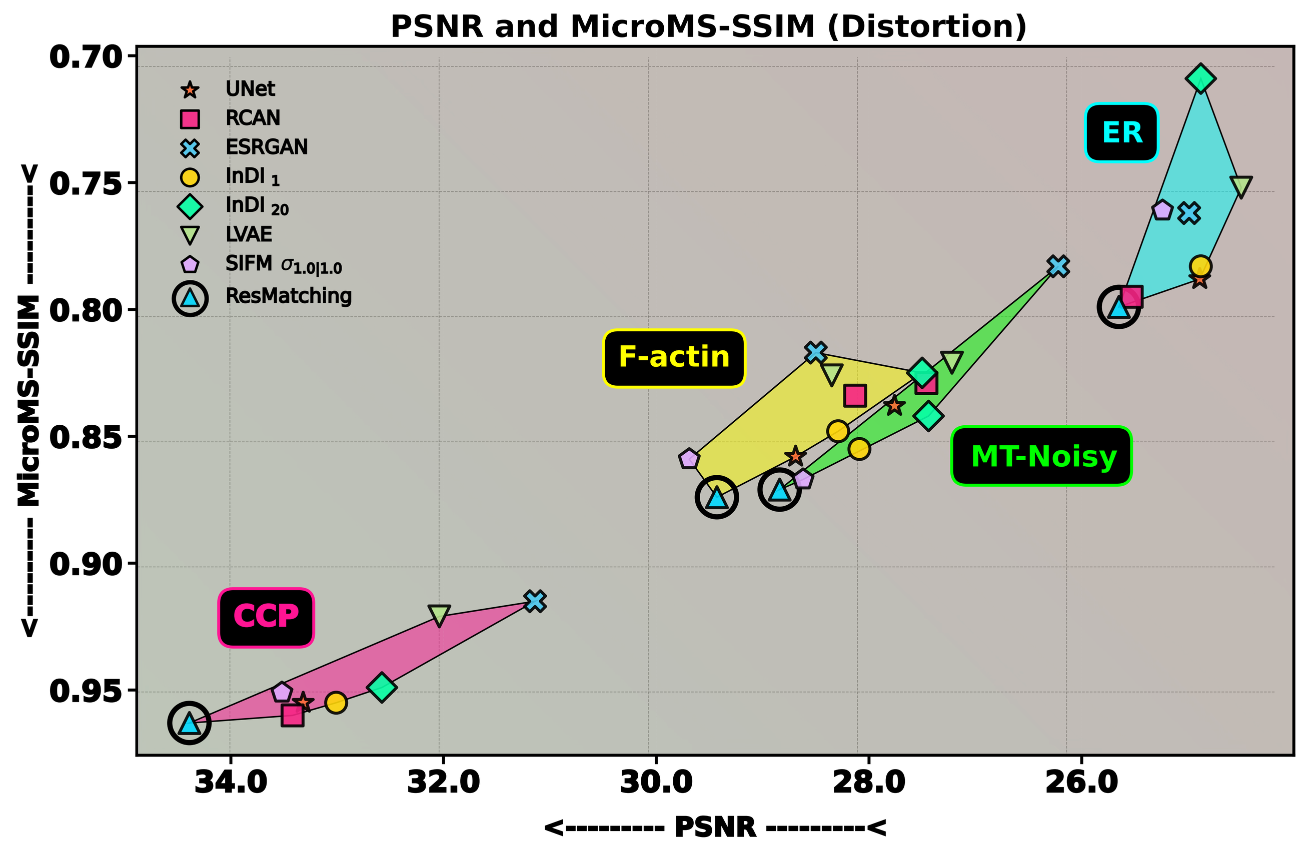

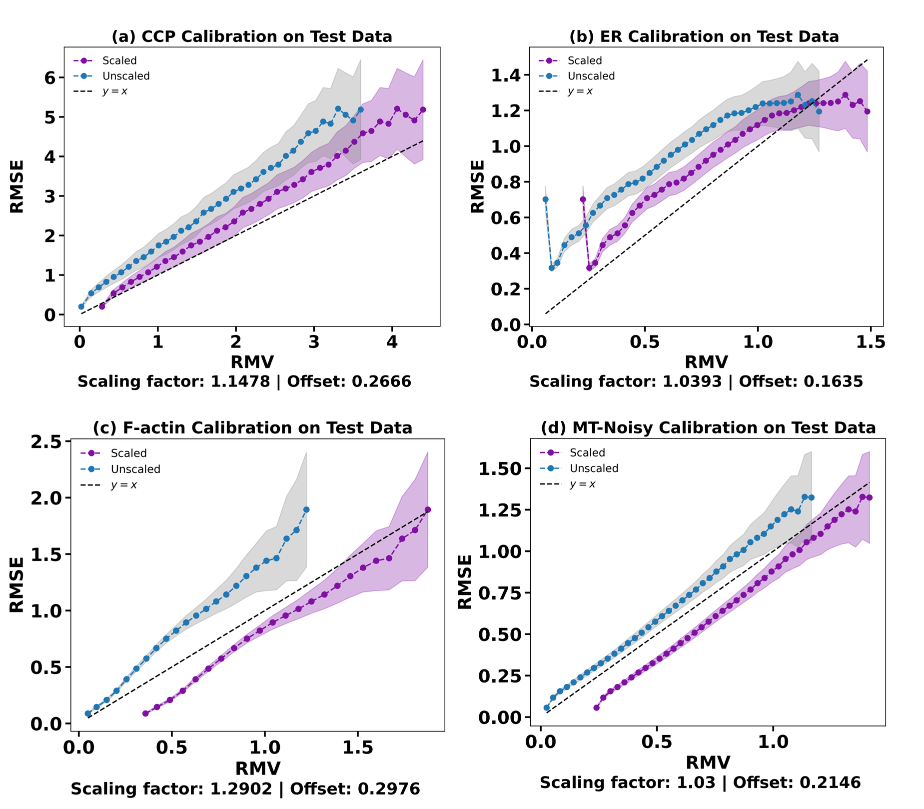

A Cautionary Note

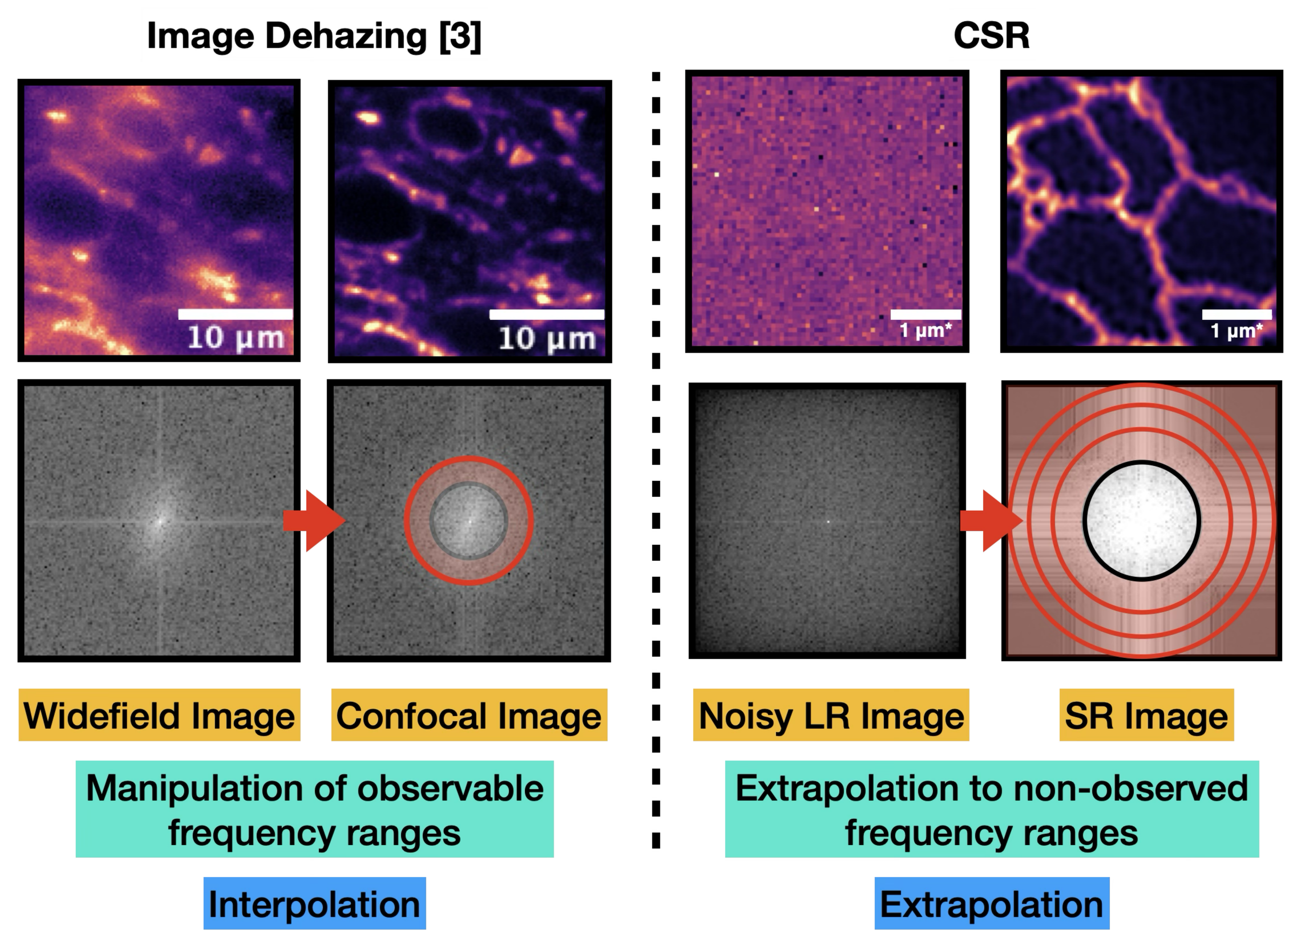

- CSR is an inverse problem that predicts high frequencies that go beyond the frequency range observable by a given microscope - **"extrapolation"**.
- The learned prior is the only source for these unobserved frequencies - it is critical for CSR to have the **best possible prior**.
- How can we know if the prior is suitable for a given input? Assessing a model's uncertainty is one solution, and Calibration plots show if these uncertainties scale with the true error. **This is an active field of research.**
- Hence, we believe that CSR should be **used with caution** and only if your project needs the gain in image resolution and quality enough to justify the uncertainty of sound predictions.

Conclusions

- We introduce ResMatching, an ODE-based iterative method for ill-posed computational super-resolution that successfully restricts the averaging space to navigate the perception-distortion tradeoff.
- Flow Matching models are especially useful where there is extreme uncertainty in the data such as heavy noise.
- This improved visual fidelity comes with the inherent challenge of higher computational costs compared to standard one-shot inference methods. Future work may optimize for throughput.
- Stitching artifacts continues to be a problem and additional method development will be required (a work in progress at JugLab!).

References

[1] Delbracio, M. and Milanfar, P. (2023). Inversion by direct iteration: An alternative to denoising diffusion for image restoration. TMLR 2023.
[2] Blau, Y. and Michaeli, T. (2018). The perception-distortion tradeoff (CVPR 2018)
[3] Ray, A, Ashesh A, and Jug F. (2026). HazeMatching: Dehazing Light Microscopy Images with Guided Conditional Flow Matching. CVPR 2026 (Findings).
[4] Qiao, C. et al. (2021). Evaluation and development of deep neural networks for image super-resolution in optical microscopy," Nature Methods