

# RESMATCHING: NOISE-RESILIENT COMPUTATIONAL SUPER-RESOLUTION VIA GUIDED CONDITIONAL FLOW MATCHING

Anirban Ray<sup>1,2</sup>, Vera Galinova<sup>1</sup>, and Florian Jug<sup>1</sup>

<sup>1</sup>Human Technopole, Milan, Italy    <sup>2</sup>Technische Universität Dresden, Germany

{anirban.ray, vera.galinova, florian.jug}@fht.org

## ABSTRACT


Computational Super-Resolution (CSR) in fluorescence microscopy has, despite being an ill-posed problem, a long history [1]. At its very core, CSR is about finding a prior that can be used to extrapolate frequencies in a micrograph that have never been imaged by the image-generating microscope. It stands to reason that, with the advent of better data-driven machine learning techniques, stronger prior can be learned and hence CSR can lead to better results. Here, we present **RESMATCHING**, a novel CSR method that uses guided conditional flow matching to learn such improved data-priors. We evaluate RESMATCHING on 4 diverse biological structures from the BiOSR dataset [2] and compare its results against 7 baselines. RESMATCHING consistently achieves competitive results, demonstrating in all cases the best trade-off between data fidelity and perceptual realism [3]. We observe that CSR using RESMATCHING is particularly effective in cases where a strong prior is hard to learn, *e.g.* when the given low-resolution images contain a lot of noise. Additionally, we show that RESMATCHING can be used to sample from an implicitly learned posterior distribution and that this distribution is calibrated [4] for all tested use-cases, enabling our method to deliver a pixel-wise data-uncertainty term that can guide future users to reject uncertain predictions [5, 6, 7].

**Index Terms**— computational super-resolution, conditional flow matching, uncertainty estimation.

## 1. INTRODUCTION

In the context of fluorescence microscopy, many approaches for computational super-resolution (CSR) have been proposed over the years [1]. Classical CSR algorithms can be understood as attempts to invert the low-pass filtering imposed by the microscope’s optical transfer function (OTF) and thereby to extrapolate high-frequency information lost during image formation. Because this inverse problem is ill-posed, priors are essential for constraining the space of feasible solutions [1]. Early methods introduced explicitly designed priors that capture simple image characteristics, such as smoothness or sparsity. Prominent examples include Tikhonov [8] and total-variation regularization [9], which penalize high gradients to suppress noise while preserving edges, and iterative deconvolution schemes such as the Richardson–Lucy algorithm [10], which exploit a Poisson noise model and an implicit positivity constraint to recover plausible high-frequency detail.

With the advent of deep learning, hand-crafted regularization terms have been replaced by learned, data-driven priors. Early convolutional architectures such as the U-NET [11, 12] and RCAN [13] established strong baselines for microscopy restoration by learning hierarchical feature representations and attention-based refinement, respectively. However, their deterministic nature limits their ability to



**Fig. 1.** RESMATCHING, using a conditional flow matching approach, leads to best-in-class computational super-resolution results even in severely noisy microscopy data.

model the inherent uncertainty in fluorescence imaging, especially under high noise conditions. These models, trained directly on suitable microscopy datasets, capture the statistics of biological structures directly from the data and enable more powerful solutions to inverse problems such as denoising, deblurring, or super-resolution [14]. In particular, convolutional and adversarial networks [15] have demonstrated that structural priors learned from fluorescence micrographs can substantially improve resolution restoration, even under severe photon limitations. From this perspective, modern CSR can be viewed as the process of learning a *structural prior* that constrains the predictor toward biologically meaningful high-frequency content.

To overcome these limitations, probabilistic frameworks such as hierarchical variational autoencoders (HVAE [16]) and implicit diffusion inference models [17] have introduced stochastic latent variables to capture spatial uncertainty and diverse, plausible reconstructions. Recent progress in generative modeling has further expanded this view. Models such as variational autoencoders [16], diffusion processes (INDI [18]), and flow matching networks [19, 20] provide a formal framework to approximate the distribution of high-resolution images conditioned on low-resolution measurements. Such generative approaches also offer a flexible way to sample from an explicit or implicit posterior over possible reconstructions, which, in theory, enables uncertainty quantification.

In this work, we investigate whether conditional flow matching, a recent generative modeling paradigm, can yield improved priors for CSR in fluorescence microscopy. Rather than aiming for a general-purpose foundation model, we introduce **RESMATCHING**, which (i) learns a detailed, data-specific structural prior from fluorescence microscopy images, (ii) achieves a good trade-off between data fidelity and perceptual realism [3], (iii) supports posterior sampling to visualize the diversity of plausible reconstructions, and (iv) provides well-calibrated uncertainty estimates [4]. Together, these features make RESMATCHING a conceptually novel and practically powerful method for uncertainty-aware computational super-resolution.

## 2. METHOD

### 2.1. Problem definition

Computational super-resolution (CSR) in fluorescence microscopy seeks to recover high-resolution (HR) structures from low-resolution (LR) acquisitions corrupted by optical blur and imaging noise. We model the LR image as  $\mathbf{x}_{M_0} = H_{M_0}(s) + \eta(s)$ , where  $H_{M_0}$  represents the unknown degradation operator of microscope  $M_0$ ,  $s$  is the underlying biological structure, and  $\eta(s)$  is the signal-dependent noise. The corresponding HR image from a higher resolution microscope  $M_1$  is given by  $\mathbf{x}_{M_1} = H_{M_1}(s) + \eta(s)$ . Given paired LR–HR observations  $(\mathbf{x}_{M_0}, \mathbf{x}_{M_1})$  of the same sample  $s_i \in \mathcal{S}$ , our goal is to learn a function  $\hat{\mathbf{x}}_{M_1} = S(\mathbf{x}_{M_0})$  that reconstructs HR-like images  $\hat{\mathbf{x}}_{M_1}$  that are faithful to the unknown  $\mathbf{x}_{M_1}$ .

### 2.2. Conditional flow matching (CFM)

Flow matching unifies diffusion processes and normalizing flows through continuous-time dynamics [19, 20]. Instead of incrementally denoising a sample through discrete diffusion steps, flow-matching models learn a vector field that continuously transports samples from a simple base distribution to the target data distribution. In its conditional form, the model learns a velocity field  $v_\theta(t, \mathbf{x}_t, \mathbf{x}_{M_0})$  such that a clean HR image  $\mathbf{x}_{M_1}$  can be recovered from its degraded counterpart  $\mathbf{x}_{M_0}$  by integrating the associated ODE,  $\frac{d\mathbf{x}_t}{dt} = v_\theta(t, \mathbf{x}_t, \mathbf{x}_{M_0})$ , with  $\mathbf{x}_0 = \mathbf{x}_{t=0} \sim \mathcal{N}(0, I)$  and  $\hat{\mathbf{x}}_{M_1} = \mathbf{x}_{t=1}$  corresponding to the reconstructed HR image.

In fluorescence microscopy, the degradation operator  $H_{M_0}$  is typically unknown and highly variable across microscopes, while the acquisition conditions and noise statistics  $\eta(s)$  depend on the true signal intensity. Standard conditional flow-matching methods [19] assume that  $H_{M_0}$  and the noise level are known, an unrealistic assumption in our setting. Recent work such as HAZEMATCHING [20] addressed this limitation by introducing a guided variant of CFM that implicitly learns to invert complex degradations without access to  $H_{M_0}$  or explicit noise parameters.

---

#### Algorithm 1 RESMATCHING Training Procedure

---

**Require:** velocity model  $v_\theta$ ; learning rate  $\gamma$ ; steps  $T$

- 1: **repeat**
- 2:    $s \sim \mathcal{S}$ ;  $t \sim \mathcal{U}([i/T]_{i=0}^T)$ ;  $\mathbf{x}_0 \sim \mathcal{N}(0, I)$ ;
- 3:    $\mathbf{x}_{M_0} \sim p_{M_0}(\mathbf{x}|s)$ ;  $\mathbf{x}_{M_1} \sim p_{M_1}(\mathbf{x}|s)$
- 4:    $\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\mathbf{x}_{M_1}$
- 5:    $\mathcal{L} = \|v_\theta(t, \mathbf{x}_t, \mathbf{x}_{M_0}) - (\mathbf{x}_{M_1} - \mathbf{x}_0)\|^2$
- 6:    $\theta \leftarrow \theta - \gamma \nabla_\theta \mathcal{L}$
- 7: **until** converged
- 8: **return**  $v_\theta$

---



---

#### Algorithm 2 RESMATCHING Inference Procedure

---

**Require:** velocity model  $v_\theta$ , steps  $T$ , input  $\mathbf{x}_{M_0}$  to condition on

- 1:  $\delta = 1/T$
- 2:  $\mathbf{x}_0 \sim \mathcal{N}(0, I)$
- 3: **for**  $t = 1$  to  $T$  **do**
- 4:    $\mathbf{x}_t = \mathbf{x}_{t-1} + \delta v_\theta(\delta \cdot (t-1), \mathbf{x}_t, \mathbf{x}_{M_0})$
- 5: **end for**
- 6: **return**  $\mathbf{x}_T$   $\triangleright$  called  $\hat{\mathbf{x}}_{M_1}$  in Section 2

---

### 2.3. Guided conditional flow matching for CSR

To learn  $S$ , we extend the conditional flow matching framework proposed in HAZEMATCHING [20] to the CSR setting by guiding the conditional velocity field with the observed LR image  $\mathbf{x}_{M_0}$ . The model learns a continuous transport that maps samples drawn from a base Gaussian distribution to samples from the HR image distribution, conditioned on the LR input. This yields a generative formulation in which HR reconstructions are obtained by integrating the learned conditional velocity field from  $t = 0$  to  $t = 1$ .

### 2.4. Marginal probability path

We consider three relevant distributions: a base distribution  $\mathbf{x}_0 \sim \mathcal{N}(0, I)$ , the LR (source) distribution  $\mathbf{x}_{M_0} \sim p_{M_0}(\mathbf{x}|s_i)$ , and the HR (target) distribution  $\mathbf{x}_{M_1} \sim p_{M_1}(\mathbf{x}|s_i)$ . We define the interpolant between  $\mathbf{x}_0$  and  $\mathbf{x}_{M_1}$  as  $\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\mathbf{x}_{M_1}$  for  $t \in [0, 1]$ , leading to the conditional path distribution

$$p_t(\mathbf{x}|\mathbf{x}_{M_1}) = \mathcal{N}(t\mathbf{x}_{M_1}, (1-t)^2 I). \quad (1)$$

### 2.5. Marginal velocity field for CSR

Our velocity field  $v(t, \mathbf{x}_t, \mathbf{x}_{M_0}|\mathbf{x}_{M_1})$  depends on the interpolation time step  $t \in [0, 1]$ , the interpolant  $\mathbf{x}_t$ , and the guiding LR image  $\mathbf{x}_{M_0}|\mathbf{x}_{M_1}$  on which we condition. The marginal velocity field is then defined as

$$v(t, \mathbf{x}_t, \mathbf{x}_{M_0}) = \int v(t, \mathbf{x}_t|\mathbf{x}_{M_1}, \mathbf{x}_{M_0}) p_{M_1}(\mathbf{x}_{M_1}|\mathbf{x}_t, \mathbf{x}_{M_0}) d\mathbf{x}_{M_1}. \quad (2)$$

We train a neural network  $v_\theta$  to approximate this velocity field by minimizing

$$\min_{\theta} \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_{M_0}, \mathbf{x}_{M_1}), t} \|v_\theta(t, \mathbf{x}_t, \mathbf{x}_{M_0}) - (\mathbf{x}_{M_1} - \mathbf{x}_0)\|^2, \quad (3)$$

where  $(\mathbf{x}_0, \mathbf{x}_{M_0}, \mathbf{x}_{M_1}) \sim p(\mathbf{x}_0, \mathbf{x}_{M_0}, \mathbf{x}_{M_1}|s_i)$  and  $\mathbf{x}_t \sim p_t(\mathbf{x}|\mathbf{x}_{M_1})$ . The network learns to predict the conditional velocity that transports a noisy sample  $\mathbf{x}_0$  toward  $\mathbf{x}_{M_1}$ , guided by the LR input  $\mathbf{x}_{M_0}$ .

### 2.6. Inference and posterior sampling


At inference time, the CSR operator  $S(\mathbf{x}_{M_0})$  integrates the learned velocity field from  $t = 0$  to 1 via an ODE solver,

$$\frac{d\mathbf{x}_t}{dt} = v_\theta(t, \mathbf{x}_t, \mathbf{x}_{M_0}), \quad \mathbf{x}_0 \sim \mathcal{N}(0, I). \quad (4)$$

The final state  $\mathbf{x}_1$  corresponds to the CSR prediction  $\hat{\mathbf{x}}_{M_1}$ . Multiple predictions from an implicit posterior are possible by repeating the inference from different  $\mathbf{x}_0^j \sim \mathcal{N}(0, I)$ . Details of training and inference are similar to what was discussed in [20] and are shown here in detail in Algorithms 1 and 2, respectively.

### 2.7. Uncertainty calibration

Being capable of sampling solutions from an implicit posterior, see previous paragraph, we can adopt the uncertainty calibration strategy from [7, 21, 20]. Once posterior samples  $\{\hat{\mathbf{x}}_{M_1}^{(k)}\}_{k=1}^K$  were generated for all  $K$  LR input images, we computed the pixel-wise standard deviation over those samples per image and combined all these values to estimate the model uncertainty via the root mean variance (RMV). We additionally fit a linear calibration model between the RMV and the true prediction error, quantified via the root mean squared error (RMSE) w.r.t. a required set of ground truth HR images, by



**Fig. 2. Quantitative results – data fidelity and realism.** Each row corresponds to results obtained on each data-subset we trained on (see Section 3). We show PSNR vs. MicroMS-SSIM (left, note inverted axis for more consistent plots) and LPIPS vs. FID (center), capturing the trade-off between pixel-level fidelity and perceptual quality. RESMATCHING is highlighted with an additional red circle. Results tables (right) show the same data.

minimizing  $\min_{\alpha, \beta} \|\text{RMSE} - (\alpha \text{RMV} + \beta)\|^2$ , and report both calibrated and uncalibrated reliability plots to evaluate prediction trustworthiness (see Figure 4).

### 3. EXPERIMENTS


**Datasets:** We evaluate RESMATCHING on four representative biological structures from the BioSR dataset: Clathrin-Coated Pits (CCP), Endoplasmic Reticulum (ER), F-actin, and Microtubule-Noisy (MT-Noisy), where we add additional noise to BioSR’s MT data. The training sets contain 39, 53, 35, and 40 raw images of size  $1004 \times 1004$ , from which we crop 3120, 4240, 2800, and 3200  $128 \times 128$  patches, respectively. Each training data subset and hence experiment further includes 5 validation images and 10 test images.

**Training and Evaluation:** We train our networks according to [20], taking into account the adaptations for the CSR task we described in Section 2 and Algorithm 1 using a step size according to  $T = 20$ . Note that we feed the interpolation time step  $t$  to the velocity field  $v_\theta$ , after positionally encoding it, to all residual blocks by adding


it to their respective output tensor [22]. After training, to evaluate the quality of the trained models, each test image is processed according to Algorithm 2, using inner tiling with 50% overlap [6], *i.e.* from each  $128 \times 128$  tile only the central  $64 \times 64$  region is used to stitch the final prediction. Distortion metrics (PSNR, MicroMS-SSIM [23]) are then computed over the entire stitched image. Because perceptual metrics (LPIPS, FID) are sensitive to tiling boundaries, we compute them on individual  $64 \times 64$  inner tiling crops. Inference for a full test image requires  $3.976 \pm 0.003$  sec. on an NVIDIA V100 GPU.

**MMSE estimate:** Since RESMATCHING can sample from an implicit posterior, we also compute an approximate MMSE estimate by pixel-wise averaging 50 posterior samples [5].

**Baselines:** We compare RESMATCHING against point-predicting and variational/generative baselines. As point predictors, we chose a vanilla U-NET [11] and an RCAN [13]. Among generative approaches, we evaluate against the adversarially trained ESRGAN [15], INDI [18] (both with 1-step and 20-step inference), and a hierarchical variational autoencoder (HVAE) [16]. We also compare against SIFM [19], setting  $\sigma=1.0$  following the optimal



**Fig. 3. Qualitative results.** Representative predictions are shown for all data subsets (experiments). For each result (I–IV), we show: (a) the full low-resolution noisy input with a highlighted crop region (yellow box), (b) the cropped input region, (c) the high-resolution ground truth, (d–j) predictions of baseline methods (see Section 3), (k) a single posterior sample generated by RESMATCHING, and (l) the MMSE estimate computed over 50 posterior samples.



**Fig. 4. Model calibration.** We show root mean variance (RMV) versus root mean square error (RMSE) as described in Section 2. Each plot ((a)–(d)) corresponds to one experiment we conducted. The dashed line indicates  $y = x$ .

configuration reported in [20].

#### 4. RESULTS

Figure 2 shows the PSNR–MicroMS-SSIM and LPIPS–FID trade-offs for each training set described above. Across all experiments, RESMATCHING achieves competitive fidelity scores (PSNR, MicroMS-SSIM) while simultaneously maintaining very good perceptual quality (LPIPS, FID). These results demonstrate that our conditional flow-matching approach learns a structural prior strong enough to outperform all baselines. As indicated in [20, 18], the perception–distortion trade-off can additionally be modulated by varying the  $T$  during inference.

In Figure 3, we show representative crops for each dataset, com-

paring all methods with one another. Note that for RESMATCHING we show a single posterior sample (prediction) as well as the approximate MMSE. While baselines such as RCAN and ESRGAN tend to over-smooth or hallucinate high-frequency details, RESMATCHING preserves filament continuity and realistic texture.

**Model Calibration:** We plot the root-mean-variance (RMV) against the root-mean-squared-error (RMSE) for all datasets (Fig. 4). After linear rescaling, the calibration curves remain close to the ideal identity line, indicating that all trained networks are well calibrated [4].

#### 5. DISCUSSION AND CONCLUSION

Computational super-resolution (CSR) ultimately depends on the strength and validity of the prior used to infer frequencies that were never measured by the microscope. RESMATCHING contributes to this ongoing pursuit by demonstrating that guided conditional flow matching can learn expressive, biologically grounded priors that unify denoising and resolution enhancement in a single generative process. At the same time, our findings highlight the intrinsic difficulty of CSR: predicting unseen high-frequency structures is fundamentally uncertain, and visually plausible results do not necessarily imply physical correctness [14, 1]. By explicitly modeling uncertainty and calibrating the learned posterior, RESMATCHING offers a principled way to quantify and communicate this ambiguity.

We hope that this work not only advances the technical state of CSR but also encourages a more careful interpretation of its outputs—treating super-resolved reconstructions as informed hypotheses rather than as ground truth.

#### 6. COMPLIANCE WITH ETHICAL STANDARDS

This research was conducted using publicly available microscopy datasets that do not contain any human or animal subjects. No ethical approval was required.

## 7. ACKNOWLEDGMENTS

This work was supported by the European Union through the Horizon Europe program (IMAGINE project, grant agreement 101094250-IMAGINE and AI4Life project, grant agreement 101057970-AI4LIFE) and the generous core funding of Human Technopole. We also thank Ashesh for the fruitful discussions.

## 8. REFERENCES

- [1] Wenfeng Tian, Riwang Chen, and Liangyi Chen, “Computational super-resolution: An odyssey in harnessing priors to enhance optical microscopy resolution,” *Anal. Chem.*, vol. 97, no. 9, pp. 4763–4792, Mar. 2025.
- [2] Cheng Qiao, Dan Li, Yanan Guo, Shuo Liu, Tingting Jiang, Qionghai Dai, Yibo Chen, and Xiangyu Fan, “Evaluation and development of deep neural networks for image super-resolution in optical microscopy,” *Nature Methods*, vol. 18, pp. 194–202, 2021.
- [3] Yochai Blau and Tomer Michaeli, “The perception-distortion tradeoff,” in *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, 2018, pp. 6228–6237, Computer Vision Foundation / IEEE Computer Society.
- [4] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger, “On calibration of modern neural networks,” in *Proceedings of the 34th International Conference on Machine Learning*, 2017, vol. 70 of *Proceedings of Machine Learning Research*, pp. 1321–1330, PMLR.
- [5] Mangal Prakash, Mauricio Delbracio, Peyman Milanfar, and Florian Jug, “Interpretable unsupervised diversity denoising and artefact removal,” in *International Conference on Learning Representations*, 2022.
- [6] Ashesh Ashesh, Alexander Krull, Moises Di Sante, Francesco Pasqualini, and Florian Jug, “usplit: Image decomposition for fluorescence microscopy,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 21219–21229.
- [7] Ashesh Ashesh and Florian Jug, “denoisplit: A method for joint microscopy image splitting and unsupervised denoising,” in *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part XXIV*, Ales Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol, Eds. 2024, vol. 15082 of *Lecture Notes in Computer Science*, pp. 222–237, Springer.
- [8] Andrey N. Tikhonov and Vasilii Y. Arsenin, *Solutions of Ill-posed Problems*, Winston and Sons, Washington, D.C., 1977, Translated from the Russian by John C. Brown.
- [9] Leonid I. Rudin, Stanley Osher, and Emad Fatemi, “Nonlinear total variation based noise removal algorithms,” *Physica D: Nonlinear Phenomena*, vol. 60, no. 1–4, pp. 259–268, 1992.
- [10] L. B. Lucy, “An iterative technique for the rectification of observed distributions,” *The Astronomical Journal*, vol. 79, pp. 745, 1974.
- [11] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Proc. MICCAI*, 2015, pp. 234–241.
- [12] Martin Weigert, Uwe Schmidt, Tobias Boothe, Andreas Müller, Alexandr Dibrov, Akanksha Jain, Benjamin Wilhelm, Deborah Schmidt, Coleman Broadbudd, Siân Culley, Mauricio Rocha-Martins, Fabián Segovia-Miranda, Caren Norden, Ricardo Henriques, Marino Zerial, Michele Solimena, Jochen Rink, Pavel Tomancak, Loic Royer, Florian Jug, and Eugene W. Myers, “Content-aware image restoration: Pushing the limits of fluorescence microscopy,” *bioRxiv*, 2018.
- [13] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu, “Image super-resolution using very deep residual channel attention networks,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 286–301.
- [14] Hari Shroff, Iliaria Testa, Florian Jug, and Suliana Manley, “Live-cell imaging powered by computation,” *Nat. Rev. Mol. Cell Biol.*, Feb. 2024.
- [15] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy, “Esrgan: Enhanced super-resolution generative adversarial networks,” in *Proc. ECCV Workshops*, 2018, pp. 63–79.
- [16] Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther, “Ladder variational autoencoders,” in *Proc. NeurIPS*, 2016, pp. 3738–3746.
- [17] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi, “Image super-resolution via iterative refinement,” *arXiv:2104.07636*, 2021.
- [18] Mauricio Delbracio and Peyman Milanfar, “Inversion by direct iteration: An alternative to denoising diffusion for image restoration,” *Trans. Mach. Learn. Res.*, vol. 2023, 2023.
- [19] Michael Samuel Albergo, Mark Goldstein, Nicholas Matthew Boffi, Rajesh Ranganath, and Eric Vanden-Eijnden, “Stochastic interpolants with data-dependent couplings,” in *Proc. ICML*, 2024, pp. 921–937.
- [20] Anirban Ray, Ashesh, and Florian Jug, “Hazematching: Dehazing light microscopy images with guided conditional flow matching,” 2025.
- [21] Ashesh Ashesh, Federico Carrara, Igor Zubarev, Vera Galinova, Melisande Croft, Melissa Pezzotti, Daozheng Gong, Francesca Casagrande, Elisa Colombo, Stefania Giussani, Elena Restelli, Eugenia Cammarota, Juan Manuel Battagliotti, Nikolai Klena, Moises Di Sante, Gaia Pigino, Elena Taverna, Oliver Harschnitz, Nicola Maghelli, Norbert Scherer, Damian Edward Dalle Nogare, Joran Deschamps, Francesco Pasqualini, and Florian Jug, “Microsplit: Semantic unmixing of fluorescent microscopy data,” *bioRxiv*, 2025.
- [22] Prafulla Dhariwal and Alex Nichol, “Diffusion models beat gans on image synthesis,” 2021.
- [23] Ashesh Ashesh, Joran Deschamps, and Florian Jug, “Microsim: Improved structural similarity for comparing microscopy data,” 2024.