



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

FreeBridge: Variational Schrödinger Bridges for Cellular Transition Dynamics

Xurui Wang^{1,2*}, Qin Ren^{1*}, Jun Ma³, Haibin Ling¹, and Chenyu You^{1**}

¹ Stony Brook University, Stony Brook, NY, USA
chenyu.you@stonybrook.edu

² University of Toronto, Toronto, ON, Canada

³ University Health Network, Toronto, ON, Canada

Abstract. High-content imaging assays quantify cellular responses to chemical and genetic perturbations, yet continuous trajectories of individual cells are unobservable because cells are chemically fixed at acquisition. Perturbation modeling therefore reduces to inferring stochastic transport between control and treated populations observed only as separate marginals. While recent generative models achieve strong endpoint alignment, boundary consistency does not determine intermediate evolution: multiple stochastic processes may connect identical marginals while traversing regions unsupported by observed single-cell morphologies. We introduce **FreeBridge**, a Schrödinger Bridge formulation for single-cell transition modeling under endpoint-only supervision. FreeBridge defines atomic states as instance-segmented single-cell representations, establishing a fixed cellular manifold, and learns stochastic transport constrained within this geometry via empirical latent support regularization. Across BBBC021, RxRx1, and JUMP, FreeBridge maintains competitive or improved endpoint fidelity under a unified evaluation protocol; on BBBC021, it further reduces intermediate support violations and improves mechanism-of-action retention. These findings highlight the importance of geometric grounding for biologically interpretable perturbation dynamics.

Keywords: High-Content Imaging · Cellular Morphology Modeling · Schrödinger Bridges · Perturbation Dynamics

1 Introduction

Building a virtual cell capable of simulating cellular responses *in silico* has long been a central aspiration in quantitative cell biology [19]. Image-based profiling enables large-scale measurement of morphological responses under chemical and genetic perturbations [5, 15]. The emergence of large perturbation atlases such as RxRx1 and JUMP [23, 6], together with recent efforts to integrate artificial intelligence into virtual cell construction [4], creates an opportunity to model perturbation dynamics directly from imaging data.

* Equal contribution.

** Corresponding author.




Fig. 1. Overview of FreeBridge. (a) Multi-cell microscopy images are segmented into atomic single-cell states, forming the State Unit. Each crop is embedded as $z = \phi(x)$, defining control and perturbed endpoint distributions. On this fixed latent geometry, FreeBridge learns a time-conditioned drift $u_\theta(z, t, c)$ to model Schrödinger Bridge transport, with an empirical support cost $V(z)$ regularizing intermediate states. (b) On BBBC021, FreeBridge achieves lower overall and condition-level FID and KID than PhenDiff, IMPA, and CellFlux under the unified single-cell protocol.

Recent generative models learn mappings between control and perturbed distributions. Diffusion models [9,20,17,22], flow matching [12,14,24], and morphology predictors including PhenDiff, IMPA, and CellFlux [3,16,25] achieve high endpoint fidelity by framing perturbation prediction as distribution-to-distribution transformation. However, high-content assays measure chemically fixed cells; individual trajectories are unobservable, and only marginal endpoint distributions are available. Infinitely many stochastic processes can connect identical marginals [7,11], so endpoint alignment does not determine intermediate evolution.

Most existing methods optimize distribution-level objectives to match control and perturbed populations [9,12,3,25,16]. While effective for endpoint realism, such formulations leave intermediate cellular structure underdetermined. The notion of an atomic cellular state is often implicit and entangled with population-level statistics, so intermediate samples may satisfy FID [8] or KID [2] while traversing regions weakly supported by observed single-cell morphologies.

A perturbation model therefore requires an explicit definition of admissible cellular states together with a transition rule constrained within that state space. We describe this structure as a *Cell Engine*, comprising a state unit and a transition mechanism. Atomic states are instantiated as instance-segmented single-cell representations [21], establishing a geometrically grounded manifold independent of population composition. On this manifold, perturbation dynamics are modeled through a Schrödinger Bridge-inspired stochastic transport process [7,11], regularized by empirical latent support constraints grounded in nonparametric density estimation [18,1].

Based on this formulation, we develop **FreeBridge**, a constrained stochastic transport framework for single-cell transition modeling under endpoint-only super-

vision. An overview of FreeBridge is shown in Fig. 1. By explicitly separating state specification from stochastic transport and grounding stochastic evolution within an empirically supported cellular manifold, FreeBridge constrains intermediate trajectories while preserving endpoint fidelity. Evaluations on BBBC021 [15], RxRx1 [23], and JUMP [6] show competitive or improved generative performance relative to prior distribution-alignment approaches [3, 16, 25]; on BBBC021, FreeBridge further reduces intermediate support violations.

2 Method

2.1 Overview: FreeBridge as *Cell Engine*

High-content screening provides endpoint populations of single cells under control and perturbed conditions. Because imaging requires chemical fixation, per-cell temporal trajectories are not observable. Perturbation modeling is therefore formulated as an endpoint-constrained stochastic transport problem between empirical single-cell distributions rather than as identity-preserving evolution.

As illustrated in Fig. 1(a), we develop **FreeBridge** under a *Cell Engine* formulation that separates cellular state specification from stochastic transport. The **State Unit** defines the admissible geometry of valid single-cell representations. The **Transition Rule** models entropy-regularized stochastic transport within that fixed geometry. This separation ensures that dynamical learning does not implicitly redefine the underlying notion of a biologically valid cellular state.

2.2 State Unit: Explicit Single-Cell Geometry

The State Unit specifies the atomic representation of a cellular state. We extract instance-segmented single-cell crops using Cellpose [21], ensuring that each observation corresponds to a physically identifiable cell rather than a field-level mixture. Each crop x is embedded using a fixed encoder $\phi : \mathcal{X} \rightarrow \mathbb{R}^d$, producing a latent representation: $z = \phi(x)$. The encoder is frozen during transport learning so that the latent geometry is defined prior to learning the Transition Rule. Let μ denote the distribution of control single-cell images and ν_c the distribution under perturbation condition c . Their pushforwards under ϕ are

$$\mu_z = \phi_{\#}\mu, \quad \nu_{c,z} = \phi_{\#}\nu_c. \quad (1)$$

To characterize the empirically supported cellular manifold, we construct a latent support bank

$$\mathcal{Z}_{\text{emp}} = \{\phi(x_i)\}_{i=1}^N. \quad (2)$$

Regions near $\mathcal{Z}_{\text{emp}}$ correspond to morphologies observed in real data.

2.3 Transition Rule: Schrödinger Bridge Transport

Given the fixed geometry defined by the State Unit, the Transition Rule specifies stochastic transport between μ_z and $\nu_{c,z}$. We model latent dynamics via the controlled stochastic differential equation:

$$dZ_t = u_t(Z_t) dt + \sigma dW_t, \quad Z_0 \sim \mu_z. \quad (3)$$

The objective is to identify a drift field u_t such that:

$$Z_1 \sim \nu_{c,z}. \quad (4)$$

Endpoint-only marginal matching is underdetermined before a reference process or path criterion is specified. A classical Schrödinger Bridge becomes well-defined once a Brownian reference process and boundary marginals are fixed:

$$\min_{\mathbb{P}^u} \text{KL}(\mathbb{P}^u \parallel \mathbb{P}^0) \quad \text{s.t. } Z_0 \sim \mu_z, Z_1 \sim \nu_{c,z}, \quad (5)$$

which is equivalent to the stochastic control objective [7]

$$\min_u \mathbb{E} \left[\frac{1}{2\sigma^2} \int_0^1 \|u_t(Z_t)\|^2 dt \right]. \quad (6)$$

FreeBridge follows this endpoint-constrained bridge view but augments it with a data-driven state cost that anchors intermediate paths to the empirical cellular morphology manifold, yielding a generalized Schrödinger Bridge formulation with task-specific state cost [13]. The stochastic term σdW_t captures intrinsic cellular heterogeneity. Deterministic transport yields a single trajectory per initial state, whereas the stochastic formulation permits condition-consistent variability.

2.4 Geometric Regularization of Intermediate Dynamics

Once the boundary marginals are fixed, the intermediate marginals of the bridge are determined by the chosen reference process; departing from a purely diffusive reference therefore requires an explicit path criterion. In cellular modeling, intermediate states correspond to putative morphologies along the perturbation response and should remain consistent with empirically observed single-cell structures. We therefore anchor stochastic transport to the geometry defined by the State Unit. For any latent state z , define the empirical support distance:

$$d_{\text{emp}}(z) = \min_{z_i \in \mathcal{Z}_{\text{emp}}} \|z - z_i\|. \quad (7)$$

This nearest-neighbor statistic [1] serves as a nonparametric proxy for density support [18]. Regions distant from $\mathcal{Z}_{\text{emp}}$ correspond to latent configurations not supported by observed cells. We penalize deviations along the trajectory:

$$\mathcal{L}_{\text{support}} = \mathbb{E} \left[\int_0^1 d_{\text{emp}}(Z_t)^2 dt \right]. \quad (8)$$

This regularization biases intermediate marginals toward empirically supported regions without imposing sample-level pairing constraints (Fig. 2).

Practical estimation. Direct evaluation of $d_{\text{emp}}(z)$ against the full support bank requires $O(N)$ comparisons per query. During training, we approximate the population support using the union of control and perturbed embeddings in the current minibatch, $\mathcal{Z}_B = \{z_0^{(i)}\}_{i=1}^B \cup \{z_1^{(i)}\}_{i=1}^B$. For each intermediate state $Z_t^{(i)}$, we compute

$$\hat{d}_{\text{emp}}(Z_t^{(i)}) = \min_{z_j \in \mathcal{Z}_B} \|Z_t^{(i)} - z_j\|_2. \quad (9)$$

This yields a stochastic minibatch approximation to the population state-cost term under i.i.d. sampling. The pairwise distances are computed with batched matrix operations (`torch.cdist`), resulting in $O(B^2)$ batch complexity. The nearest-neighbor index is treated as non-differentiable, while gradients propagate through the selected squared distance.

Interior Dynamics of Schrödinger Bridges. The Schrödinger Bridge constrains boundary marginals; given a fixed reference process, the interior marginals are then determined, so altering the empirical morphology of intermediate states requires an explicit state cost rather than endpoint supervision alone [7, 11]. Endpoint alignment therefore does not, by itself, fix the desired transition geometry on the data manifold. Figure 2 shows this effect on a synthetic manifold. Even when boundary distributions coincide, transport under a diffusive reference may traverse sparsely supported regions of latent space. Such trajectories preserve endpoint statistics yet alter interior structure. The support regularization restricts stochastic evolution to the empirical cellular manifold.

2.5 FreeBridge Objective

The FreeBridge objective is defined as a generalized Schrödinger Bridge with task-specific state cost $V_t(z) = \lambda d_{\text{emp}}(z)^2$:

$$\min_u \mathbb{E} \left[\frac{1}{2\sigma^2} \int_0^1 \|u_t(Z_t)\|^2 dt + \lambda \int_0^1 d_{\text{emp}}(Z_t)^2 dt \right] \quad \text{s.t. } Z_0 \sim \mu_z, Z_1 \sim \nu_{c,z}. \quad (10)$$

Intra-batch pairing for endpoint supervision. Endpoint populations are unpaired, but training requires batch-level tuples. For each condition c , we sample $\{z_0^{(i)}\}_{i=1}^B \sim \mu_z$ and $\{z_1^{(i)}\}_{i=1}^B \sim \nu_{c,z}$ independently and pair them by index after random shuffling. This stochastic coupling is used for optimization only and does not impose identity-level correspondence or constrain intermediate dynamics.

2.6 Numerical Approximation and Optimization

We parameterize the drift as $u_\theta(z, t, c)$ using a time-conditioned neural network. The continuous process is discretized via Euler–Maruyama:

$$Z_{t+\Delta t} = Z_t + u_\theta(Z_t, t, c)\Delta t + \sigma\sqrt{\Delta t}\epsilon, \quad \epsilon \sim \mathcal{N}(0, I). \quad (11)$$




Fig. 2. Intermediate trajectories under endpoint-constrained transport. Columns show increasing time along the stochastic evolution. Top (unconstrained): endpoints align, but intermediate states traverse low-support regions. Bottom (FreeBridge): support regularization prevents such excursions and preserves trajectory consistency.

With K Euler–Maruyama steps and $\Delta t = 1/K$, the training objective is approximated as

$$\hat{\mathcal{L}}(\theta) = \frac{1}{B} \sum_{i=1}^B \sum_{k=0}^{K-1} \left[\frac{\Delta t}{2\sigma^2} \|u_{\theta}(Z_{t_k}^{(i)}, t_k, c)\|_2^2 + \lambda \Delta t \hat{d}_{\text{emp}}(Z_{t_k}^{(i)})^2 \right] + \mathcal{L}_{\text{end}}, \quad (12)$$

where the endpoint term $\mathcal{L}_{\text{end}}$ aligns the terminal generated batch with samples from $\nu_{c,z}$ in the same latent metric used for evaluation. All endpoint samples are drawn independently; the random intra-batch pairing is used only as an optimization device and does not imply observed cell-level correspondences.

3 Experiments

3.1 Experimental Setup

Datasets and Baselines. We evaluate FreeBridge on three widely used high-content imaging benchmarks: BBBC021 [15], RxRx1 [23], and JUMP [6]. These datasets cover chemical and genetic perturbations across multiple experimental conditions. We compare against recent morphology prediction methods, including PhenDiff [3], IMPA [16], and CellFlux [25].

Implementation Details. All methods are evaluated under a unified single-cell pipeline. Images are instance-segmented and cropped using the same preprocessing procedure. This ensures that differences in performance arise from transition modeling rather than data preparation. The encoder ϕ is a frozen ResNet-50 backbone (latent dimension $d = 2048$), and latent embeddings are precomputed for all segmented single-cell crops (96×96). The drift network is a time-conditioned multilayer perceptron (hidden width 512, 6 layers) that takes (z, t, c) as input. We optimize with Adam [10] using batch size $B = 256$, learning rate 3×10^{-4} , and $K = 50$ Euler–Maruyama discretization steps during training, with the minibatch support bank described in Sec. 2. The support weight λ_{bank} , diffusion noise σ , and inference budget follow the default configuration; sensitivity to each is reported

Table 1. Quantitative comparison across datasets (5k generated samples per condition). We report overall and condition-level FID/KID together with the Support Violation Rate ($\mathcal{R}_{\text{viol}}$). All methods are evaluated under an identical single-cell preprocessing and inference protocol. FreeBridge achieves improved endpoint fidelity while maintaining lower intermediate support violations.

	BBBC021					RxRx1					JUMP				
Method	FID _o	FID _c	KID _o	KID _c	$\mathcal{R}_{\text{viol}}$	FID _o	FID _c	KID _o	KID _c	$\mathcal{R}_{\text{viol}}$	FID _o	FID _c	KID _o	KID _c	$\mathcal{R}_{\text{viol}}$
PhenDiff	48.25	107.50	3.08	3.26	0.38	67.12	175.62	5.28	5.34	48.57	126.52	5.13	5.23		
IMPA	33.91	75.76	2.74	2.79	0.34	39.94	163.91	2.89	2.87	14.81	101.03	1.10	1.00		
CellFlux	18.72	56.80	1.62	1.59	0.31	33.00	162.85	2.40	2.41	9.00	83.21	0.65	0.67		
FreeBridge	15.35	47.28	1.10	1.02	0.11	28.94	147.33	1.98	2.03	8.51	79.05	0.59	0.53		

in Fig. 4. All reported metrics use 5k generated samples per condition and are averaged over three random seeds. Endpoint fidelity is measured using Fréchet Inception Distance (FID) [8] and Kernel Inception Distance (KID) [2]. To assess trajectory consistency, we report the Support Violation Rate ($\mathcal{R}_{\text{viol}}$), defined as the proportion of intermediate states whose feature representations fall outside the empirical morphology support. Because fixed-cell assays do not provide observed single-cell trajectories, no metric can directly verify true temporal ordering at the individual-cell level; $\mathcal{R}_{\text{viol}}$ is therefore a support-feasibility metric and should not be interpreted as direct evidence of biological temporal correctness. Semantic retention is evaluated using Mechanism-of-Action (MoA) classification accuracy obtained from a linear probe trained on real perturbed samples. Unless otherwise stated, official hyperparameters are used for baseline methods.

3.2 Results

Quantitative and Qualitative Comparisons. Table 1 reports quantitative results across datasets. Fig. 1(b) provides a compact summary on BBBC021. FreeBridge achieves strongest endpoint fidelity on all three benchmarks under the matched evaluation protocol. In addition to improvements in FID and KID, FreeBridge exhibits substantially lower $\mathcal{R}_{\text{viol}}$ on BBBC021, indicating that intermediate states remain closer to empirically observed morphology regions. Figure 3 shows representative examples on BBBC021. The visual patterns align with quantitative results, with clearer preservation of compound-specific morphology. **Intermediate Trajectory Consistency.** As illustrated in Fig. 2, endpoint alignment alone does not constrain interior geometry. On real data, models without geometric anchoring may enter low-support regions of the single-cell manifold at intermediate times, producing distortions not captured by endpoint metrics. Table 2 quantifies this effect: increasing λ_{bank} reduces both the feasibility energy $\mathcal{E}_{\text{feas}}$ and the support violation rate $\mathcal{R}_{\text{viol}}$, confirming that the support prior improves intermediate validity. As noted above, this measures support feasibility rather than verified temporal ordering.

Semantic Retention. We further evaluate MoA classification accuracy to quantify preservation of condition-specific information. On BBBC021, removing

Table 2. Trajectory feasibility on BBBC021. All variants share identical architectures and differ only in the support weight λ_{bank} . Metrics are computed over $T = 11$ time points per trajectory and reported as the median over three seeds. Higher λ_{bank} improves geometric feasibility. Lower is better.

Variant	$\mathcal{E}_{\text{feas}} \downarrow$	$\mathcal{R}_{\text{viol}} \downarrow$
$\lambda_{\text{bank}} = 0$ (Unconstrained)	2.48	0.31
FreeBridge ($\lambda_{\text{bank}} = 0.3$)	1.85	0.19
FreeBridge (Default, $\lambda_{\text{bank}} = 0.5$)	1.32	0.11




Fig. 3. Qualitative comparison of endpoint morphology on BBBC021. Representative single-cell crops for two perturbations are shown. Columns display real targets and samples from IMPA, PhenDiff, CellFlux, and FreeBridge under the unified pipeline. FreeBridge preserves compound-specific structures (e.g., nuclear compaction under Demecolcine and cytoplasmic shrinkage under Mevinolin), whereas baselines show texture smoothing or incomplete perturbation patterns.

geometric regularization reduces MoA accuracy from 72.08% to 65.08%, indicating constraining intermediate evolution contributes to preserving mechanistic signal.

3.3 Ablation

Effect of Regularizations. Table 3 reports ablation results on BBBC021. Removing instance segmentation degrades both fidelity and MoA accuracy, suggesting that explicit single-cell state specification improves conditional modeling. Removing the support regularization increases $\mathcal{R}_{\text{viol}}$ and reduces semantic retention. These observations are consistent with the role of explicit single-cell geometry and geometric regularization in constraining intermediate evolution.

Hyperparameter Sensitivity. Fig. 4 shows sensitivity to the support weight λ_{bank} , diffusion noise σ , and inference steps. Performance remains stable across a range of settings, with moderate λ_{bank} values providing improved MoA retention while maintaining endpoint fidelity.

4 Conclusion

We introduced FreeBridge, a Schrödinger Bridge formulation for single-cell transition modeling under endpoint-only supervision. By coupling entropy-regularized

Table 3. Ablation study on BBBC021. We examine the effect of removing intra-batch pairing, support regularization, and instance segmentation. Both geometric anchoring and explicit single-cell state specification contribute to improved fidelity and semantic retention. Full-model results are from an independent ablation run and may differ marginally from Table 1.

Variant	FID _o ↓	FID _c ↓	MoA ↑	KID _o ↓
w/o intra-batch pairing	19.94	60.81	63.27	1.81
w/o support cost $V(\cdot)$	17.18	54.37	65.08	1.55
w/o instance segmentation	28.54	82.15	51.34	2.45
Full model	15.37	47.41	72.08	1.13




Fig. 4. Hyperparameter sensitivity on BBBC021. Performance is evaluated as a function of the support weight λ_{bank} , diffusion noise σ , and inference steps (NFE). Results indicate stable endpoint fidelity across a range of settings, with improved semantic retention at moderate regularization strength.

stochastic transport with explicit single-cell state geometry, FreeBridge links distribution-level alignment with trajectory-level consistency within empirically observed morphology space. Across multiple large-scale perturbation benchmarks, FreeBridge improves endpoint fidelity under a unified evaluation protocol; on BBBC021, it further reduces intermediate support violations. These findings suggest that modeling biological transitions requires not only boundary alignment but also explicit geometric grounding of intermediate evolution. Such a formulation offers a principled basis for virtual cell simulation and related transition modeling problems.

Limitations and cost. FreeBridge relies on a frozen encoder to define the latent cellular geometry; morphological factors absent from encoder pretraining may therefore not be recoverable through transport learning alone. The support-distance computation uses batched pairwise distances and scales as $O(B^2)$ with minibatch size. This cost is manageable at the batch sizes used here, but larger support banks may require approximate nearest-neighbor search or memory-aware batching.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Biau, G., Devroye, L.: Lectures on the nearest neighbor method, vol. 246. Springer (2015)
2. Bińkowski, M., Sutherland, D.J., Arbel, M., Gretton, A.: Demystifying mmd gans. arXiv preprint arXiv:1801.01401 (2018)
3. Bourou, A., Boyer, T., Gheisari, M., Daupin, K., Dubreuil, V., De Thonel, A., Mezger, V., Genovesio, A.: Phendiff: Revealing subtle phenotypes with diffusion models in real images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 358–367. Springer (2024)
4. Bunne, C., Roohani, Y., Rosen, Y., Gupta, A., Zhang, X., Roed, M., Alexandrov, T., AlQuraishi, M., Brennan, P., Burkhardt, D.B., et al.: How to build the virtual cell with artificial intelligence: Priorities and opportunities. *Cell* **187**(25), 7045–7063 (2024)
5. Carpenter, A.E.: Image-based chemical screening. *Nature Chemical Biology* **3**(8), 461–465 (2007)
6. Chandrasekaran, S.N., Ackerman, J., Alix, E., Ando, D.M., Arevalo, J., Bennion, M., Boisseau, N., Borowa, A., Boyd, J.D., Brino, L., et al.: Jump cell painting dataset: morphological impact of 136,000 chemical and genetic perturbations. *BioRxiv* pp. 2023–03 (2023)
7. Chen, Y., Georgiou, T.T., Pavon, M.: On the relation between optimal transport and schrödinger bridges: A stochastic control viewpoint. *Journal of Optimization Theory and Applications* **169**(2), 671–691 (2016)
8. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* **30** (2017)
9. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
10. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
11. Léonard, C.: A survey of the schrödinger problem and some of its connections with optimal transport. arXiv preprint arXiv:1308.0215 (2013)
12. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. arXiv preprint arXiv:2210.02747 (2022)
13. Liu, G.H., Lipman, Y., Nickel, M., Karrer, B., Theodorou, E., Chen, R.T.: Generalized schrödinger bridge matching. In: International Conference on Learning Representations. vol. 2024, pp. 14527–14552 (2024)
14. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. arXiv preprint arXiv:2209.03003 (2022)
15. Ljosa, V., Sokolnicki, K.L., Carpenter, A.E.: Annotated high-throughput microscopy image sets for validation. *Nature methods* **9**(7), 637 (2012)
16. Palma, A., Theis, F.J., Lotfollahi, M.: Predicting cell morphological responses to perturbations using generative modeling. *Nature Communications* **16**(1), 505 (2025)
17. Ren, Q., Wang, Y., Guo, L., Zhang, W., Fan, Z., You, C.: Scale where it matters: Training-free localized scaling for diffusion models. arXiv preprint arXiv:2511.19917 (2025)
18. Silverman, B.W.: Density estimation for statistics and data analysis. Routledge (2018)

19. Slepchenko, B.M., Schaff, J.C., Macara, I., Loew, L.M.: Quantitative cell biology with the virtual cell. *Trends in cell biology* **13**(11), 570–576 (2003)
20. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456* (2020)
21. Stringer, C., Wang, T., Michaelos, M., Pachitariu, M.: Cellpose: a generalist algorithm for cellular segmentation. *Nature methods* **18**(1), 100–106 (2021)
22. Sun, S., Wang, Y., Zhang, H., Xiong, Y., Ren, Q., Fang, R., Xie, X., You, C.: Ouroboros: Single-step diffusion models for cycle-consistent forward and inverse rendering. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2025)
23. Sypetkowski, M., Rezanejad, M., Saberian, S., Kraus, O., Urbanik, J., Taylor, J., Mabey, B., Victors, M., Yosinski, J., Sereshkeh, A.R., et al.: Rxxr1: A dataset for evaluating experimental batch correction methods. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (2023)
24. Wang, Y., Ma, Y., Li, W., You, C.: Let eeg models learn eeg. In: *International conference on machine learning* (2026)
25. Zhang, Y., Su, Y., Wang, C., Li, T., Wefers, Z., Nirschl, J., Burgess, J., Ding, D., Lozano, A., Lundberg, E., et al.: Cellflux: Simulating cellular morphology changes via flow matching. *arXiv preprint arXiv:2502.09775* (2025)