CRNDIFF: COUNT-NATIVE DIFFUSION FRAMEWORK VIA CHEMICAL REACTION NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Scientific measurements such as single-cell RNA (scRNA) sequencing often take the form of nonnegative integer counts, whereas continuous-state diffusion models approximate this discrete structure using continuous coordinates. Building on stochastic chemical reaction networks (CRNs), a class of count-native Markov jump processes, we introduce CRNDiff, a structured framework that combines count-space diffusion with inference-time conditioning on rare subpopulations. An independent birth–death instantiation yields a closed-form transition kernel for forward noising. This kernel enables reverse sampling via forward-filtering backward-sampling (FFBS) and supports data-driven selection of the terminal noising time, eliminating the need for a validation sweep. This tractability also lets us introduce tilted Feynman–Kac (FK) steering, a method for sampling target subpopulations from a frozen generator without retraining. By tilting posterior marginals before FK particle correction, steering mitigates importance-weight concentration when the target population is rare. Using scRNA-seq data from the human heart cell atlas, we test the ability of CRNDiff to generate cell-type-specific distributions. Across the three evaluated target populations, CRNDiff achieves the highest conditional fidelity among the evaluated generative models, with larger mean purity margins for rarer target populations. Generated cells preserve marker-level differential-expression structure. Replacing real training cells for the target classes with generated cells yields downstream classification performance approaching that of the real-data reference.

1 INTRODUCTION

Diffusion models generate data by reversing a stochastic process that progressively destroys its structure (Ho et al., 2020; Song et al., 2020). The forward process is therefore a modeling choice, and its design should reflect the state space and the meaningful local transformations of the data. We focus on nonnegative integer counts. In single-cell RNA sequencing, for example, each coordinate records a transcript count and lies on the nonnegative integer lattice (Lopez et al., 2018). While finite-state discrete diffusion models operate on a fixed, finite set of states (Austin et al., 2021; Campbell et al., 2022), nonnegative integer counts additionally have a natural ordering and no intrinsic upper bound. To accommodate these properties, recent count-native models formulate diffusion directly on the nonnegative integer lattice, a countably infinite state space. Blackout Diffusion and CountsDiff employ pure-death or thinning dynamics (Santos et al., 2023; Soatto et al., 2026), while Count Bridges and Count-FM use birth–death dynamics (Fishman et al., 2026; Wei & Pearson, 2026). However, existing count-native diffusion models are formulated around specific transition mechanisms. A more general, structured framework is needed to specify diverse admissible state transitions and their state-dependent rates.

Stochastic chemical reaction networks (CRNs) provide such a framework. Reaction channels specify admissible changes in count-valued states, while propensity functions determine their state-dependent rates; stochastic mass-action kinetics further arise naturally from combinatorial interactions among discrete reactant molecules (Gillespie, 1992). CRNs thus provide a compositional parameterization of count-space Markov dynamics together with useful analytic structure. Their master equations admit probability-generating-function representations (Doi, 1976; Peliti, 1985; Baez & Biamonte, 2018) that, when tractable, yield exact transition kernels and Markov bridges (Jahnke & Huisinga, 2007; Vastola, 2021). Building on this structure, we introduce CRNDiff and instantiate it

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

with an independent birth–death CRN. The forward process preserves empirical coordinate means and converges to a product-Poisson stationary distribution. Its closed-form transition kernels support bridge-based reverse sampling. The explicit factorial-cumulant dynamics yield a data-driven criterion for selecting the terminal time of the forward process based on excess-covariance decay, without requiring a validation sweep.

This tractable birth–death construction also supports tilted Feynman–Kac steering for inference-time conditioning of a frozen generator (Wu et al., 2023; Singhal et al., 2025). For rare targets where steering based on sequential Monte Carlo (SMC) can suffer from severe importance-weight concentration (Del Moral, 2004; Chopin et al., 2020), marginal tilt first adapts the reverse proposal toward the target marginals, and a Feynman–Kac correction of the residual discrepancy then accounts for the remaining joint mismatch along the reverse chain, without retraining the generator.

We evaluate CRNDiff on a controlled synthetic construction, which separates the effects of marginal tilt and residual correction, and on single-cell RNA sequencing data from an adult human heart cell atlas (Litviňuková et al., 2020), where one frozen generator serves target populations spanning a wide abundance range. CRNDiff outperforms the evaluated baselines across all conditional-fidelity metrics. Across the three evaluated target populations, its mean purity margin over the baseline with the highest mean purity increases as training abundance decreases. Generated cells preserve cell-type-specific differential-expression rankings and, when used for downstream cell-type classification, yield performance approaching that of a classifier trained entirely on real cells.

2 Background

2.1 Continuous-time diffusion models

Let $\mathcal{D}$ be a distribution on the count space $\mathcal{X}$, which in this work is a nonnegative integer lattice. Let $\overrightarrow{q}(x_t \mid x_s)$, $s < t$, denote the forward transition kernel of a continuous-time Markov jump process initialized from $q_0 = \mathcal{D}$. The corresponding single-time marginal is

$$q_t(x_t) = \sum_{x_0 \in \mathcal{X}} \overrightarrow{q}(x_t \mid x_0)\, \mathcal{D}(x_0).$$

We choose the forward process such that q_t converges to a tractable stationary distribution π.

Forward and reverse dynamics. With transition rates $w(x \to y)$, q_t satisfies

$$\partial_t q_t(x) = \sum_y \left[w(y \to x) q_t(y) - w(x \to y) q_t(x) \right]. \tag{1}$$

The corresponding generator is $(Lf)(x) = \sum_y w(x \to y)[f(y) - f(x)]$, so that $\partial_t \mathbb{E}[f(X_t)] = \mathbb{E}[(Lf)(X_t)]$. For reverse time $s \in [0, T]$, the time-reversed process has rates

$$\widetilde{w}_s(y \to x) = w(x \to y) \frac{q_{T-s}(x)}{q_{T-s}(y)}. \tag{2}$$

It evolves from q_T back to $q_0 = \mathcal{D}$; in practice q_T is replaced by π once the forward process has sufficiently relaxed.

Forward-filtering backward-sampling. For a tractable forward kernel, define the clean-state posterior $Q(x_0 \mid x_t, t)$. The corresponding Markov bridge $B(x_s \mid x_0, x_t)$ is the conditional distribution of the intermediate state x_s given the endpoints x_0 and x_t:

$$Q(x_0 \mid x_t, t) = \frac{\overrightarrow{q}(x_t \mid x_0)\mathcal{D}(x_0)}{\sum_{x'} \overrightarrow{q}(x_t \mid x')\mathcal{D}(x')}, \tag{3}$$

$$B(x_s \mid x_0, x_t) = \frac{\overrightarrow{q}(x_t \mid x_s)\, \overrightarrow{q}(x_s \mid x_0)}{\overrightarrow{q}(x_t \mid x_0)}. \tag{4}$$

Marginalizing over x_0 gives

$$\overleftarrow{q}(x_s \mid x_t) = \sum_{x_0} Q(x_0 \mid x_t, t) B(x_s \mid x_0, x_t) = \overrightarrow{q}(x_t \mid x_s) \frac{q_s(x_s)}{q_t(x_t)}. \tag{5}$$

Sampling this kernel backward on any time grid gives exact FFBS when the posterior and terminal distribution are exact (Carter & Kohn, 1994; Chopin et al., 2020). We approximate Q by a learned posterior Q_θ trained by the cross-entropy loss on $x_0 \sim \mathcal{D}$ and $x_t \sim \overrightarrow{q}(x_t \mid x_0)$.

2.1.1 CONDITIONAL GENERATION

Let P be the distribution at time 0 produced by a reverse sampling process, and P^{tar} the desired distribution. Assuming $P(x) > 0$ whenever $P^{\mathrm{tar}}(x) > 0$, define

$$H(x) := \frac{P^{\mathrm{tar}}(x)}{P(x)}. \tag{6}$$

Endpoint tilting. For the posterior associated with prior P, changing the endpoint distribution preserves the forward kernel and bridge and reweights the posterior,

$$Q^H(x_0 \mid x_t, t) = \frac{H(x_0)Q(x_0 \mid x_t, t)}{\sum_{x'} H(x')Q(x' \mid x_t, t)}. \tag{7}$$

Thus a learned posterior can be adapted to a new endpoint distribution without retraining the generator. While exact sampling also requires the target terminal marginal $\pi^T(x) = \sum_{x_0} \overrightarrow{q}(x_T = x|x_0)P^{\mathrm{tar}}(x_0)$, initialization from the common stationary distribution π is the usual large-T approximation.

Feynman–Kac steering. When P^{tar} concentrates on a region with little probability under P, applying H only at the endpoint can produce concentrated importance weights. Feynman–Kac (FK) steering mitigates this by distributing the same correction along the reverse trajectory (Del Moral, 2004; Wu et al., 2023; Singhal et al., 2025).

On $0 = t_0 < \cdots < t_K = T$, let $\mathbb{P}$ be the reverse proposal distribution over paths $x_{0:K} = (x_{t_0}, \ldots, x_{t_K})$. Choose positive potentials $\psi_j(x_{t_j})$ with $\psi_K \equiv 1$, $\psi_0 = H$. The FK path measure is

$$\mathbb{P}_{\mathrm{FK}}(dx_{0:K}) \propto \left(\prod_{j=1}^{K} G_j\right) \mathbb{P}(dx_{0:K}), \qquad \text{where} \qquad G_j = \frac{\psi_{j-1}(x_{t_{j-1}})}{\psi_j(x_{t_j})}. \tag{8}$$

Since the product telescopes to $H(x_0)$, its endpoint distribution is proportional to $P(x_0)H(x_0)$. The intermediate potentials ψ_j therefore affect sampling efficiency but not the desired endpoint correction. When ψ is not available in closed form, particle methods approximate the continuation potential $\psi_j(x_{t_j}) = \mathbb{E}[H(x_0) \mid x_{t_j}]$ on the fly, while propagating, reweighting, and resampling trajectories.

2.2 CHEMICAL REACTION NETWORKS

Notation. For any $S \times R$ matrix M, we use $M_{\cdot,r}$ to denote its r-th column, $M_{i,\cdot}$ its i-th row, and $M_{i,r}$ its (i,r)-th entry. For $\mathbf{x}, \mathbf{y} \in \mathbb{R}^S_{\geq 0}$, we also use the multi-index notation $\mathbf{x}^{\mathbf{y}} := \prod_{i=1}^{S} x_i^{y_i}$.

Definition. Chemical reaction networks (CRNs) provide a structured class of continuous-time Markov jump processes on the count-space $\mathcal{X} = \mathbb{Z}^S_{\geq 0}$. A CRN is specified by a set of species $\mathcal{S} = \{X_1, \ldots, X_S\}$, whose copy numbers define the state $n = (n_1, \ldots, n_S) \in \mathcal{X}$, together with R reactions. Let $\mathbf{S}^-, \mathbf{S}^+ \in \mathbb{Z}^{S \times R}_{\geq 0}$ denote the input and output stoichiometric matrices, and define the stoichiometric matrix $\mathbf{S} = \mathbf{S}^+ - \mathbf{S}^-$. Thus, $(\mathbf{S}^-)_{i,r}$ and $(\mathbf{S}^+)_{i,r}$ are the numbers of molecules of species X_i consumed and produced by reaction r, respectively, while $\mathbf{S}_{\cdot,r}$ is the corresponding change in the full count vector n. Reaction r is therefore written as

$$\sum_{i=1}^{S} (\mathbb{S}^-)_{i,r} X_i \longrightarrow \sum_{i=1}^{S} (\mathbb{S}^+)_{i,r} X_i, \qquad \mathbf{n} \longrightarrow \mathbf{n} + \mathbb{S}_{\cdot,r}. \tag{9}$$

Master equation and mass action. Each reaction is assigned a propensity $j_r(n)$, interpreted as its instantaneous firing rate in state n. Starting from a count vector, the system waits a random time, fires one of the available reactions, and updates the state by its stoichiometric change. Reactions whose required inputs are unavailable have zero propensity, so the dynamics remain in the nonnegative integer orthant. The CRN therefore defines transition rates $w(\mathbf{n} \to \mathbf{n} + \Delta) = \sum_{r:\, \mathbb{S}_{\cdot,r} = \Delta} j_r(\mathbf{n})$.

Substituting these rates into Eq. (1) gives the chemical master equation

$$\partial_t q_t(\mathbf{n}) = \sum_{r=1}^{R} \left[j_r(\mathbf{n} - \mathbb{S}_{\cdot,r}) q_t(\mathbf{n} - \mathbb{S}_{\cdot,r}) - j_r(\mathbf{n}) q_t(\mathbf{n}) \right], \tag{10}$$

where terms corresponding to states outside $\mathbb{Z}^S_{\geq 0}$ are understood to vanish.

For stochastic mass-action kinetics, the propensities are

$$j_r(\mathbf{n}) = k_r V^{1-|\mathbb{S}^-_{\cdot,r}|} \prod_i \frac{n_i!}{(n_i - \mathbb{S}^-_{i,r})!} \tag{11}$$

with rate constant k_r, system volume V, and reaction order $|\mathbb{S}^-_{\cdot,r}| = \sum_i \mathbb{S}^-_{i,r}$. Propensities vanish if any reactant count is insufficient. We set $V = 1$; Appendix A.5 gives the large-V diffusion approximation.

Generating functions. Define the probability generating function (PGF) $g(\mathbf{z}, t) = \mathbb{E}[\mathbf{z}^{\mathbf{n}(t)}] = \sum_{\mathbf{n} \geq 0} \mathbf{z}^{\mathbf{n}} q_t(\mathbf{n})$. Multiplication by z_i and differentiation by z_i implement the addition and removal of count units, giving the standard creation–annihilation representation (Doi, 1976; Peliti, 1985; Baez & Biamonte, 2018). For mass-action kinetics,

$$\partial_t g(\mathbf{z}, t) = \sum_r k_r V^{1-|\mathbb{S}^-_{\cdot,r}|} \left(\mathbf{z}^{\mathbb{S}^+_{\cdot,r}} - \mathbf{z}^{\mathbb{S}^-_{\cdot,r}} \right) \partial_{\mathbf{z}}^{\mathbb{S}^-_{\cdot,r}} g(\mathbf{z}, t). \tag{12}$$

Here $\partial_{\mathbf{z}}^{\mathbf{a}} = \prod_i \partial_{z_i}^{a_i}$; see Appendix A.2 for the derivation. With $g(\mathbf{z}, 0) = \mathbf{z}^{\mathbf{n}_0}$, coefficient extraction gives $q_t(\mathbf{n} \mid \mathbf{n}_0) = [\mathbf{z}^{\mathbf{n}}] g(\mathbf{z}, t)$, which supplies the transition kernel required by FFBS.

For $|\mathbb{S}^-_{\cdot,r}| \leq 1$, the PGF equation is first-order and can be treated by characteristics. This class includes births, deaths, conversions, and branching reactions, and is considerably richer than the independent example used below (Jahnke & Huisinga, 2007; Vastola, 2021).

3 CRN DIFFUSION MODEL

We instantiate the CRN framework with a simple birth–death process. We use these reactions purely as a noising mechanism, and they provide a natural way to progressively randomize count-valued data while remaining on the nonnegative integer lattice. Importantly, the resulting stochastic process is exactly solvable and gives closed-form transition probabilities. We also explicitly track how statistical structure inherited from the data—such as correlations and deviations from independent Poisson statistics—is gradually erased by the forward process, and use this decay to determine how long the data should be noised before reverse sampling begins. Finally, we take the endpoint distribution induced by the CRNDiff generator and show how it can be tilted toward a prescribed target distribution, enabling inference-time conditional generation without retraining the underlying model.

3.1 NOISING WITH A BIRTH–DEATH MODEL

Birth–death network. For each coordinate $i \in [S]$, consider the "birth" and "death" reactions that increase and decrease the count by one unit, respectively:

$$\varnothing \xrightarrow{V k_i^+} X_i, \qquad X_i \xrightarrow{k_i^- n_i} \varnothing. \tag{13}$$

We set $V = 1$, define $\mu_i = k_i^+ / k_i^-$ and $\beta_{t,i} = e^{-k_i^- t}$. Solving the PGF equation gives the exact transition kernel

$$\begin{aligned} \overrightarrow{q}(\mathbf{n}_t \mid \mathbf{n}_0) &= \prod_{i=1}^{S} \left[\mathrm{Bin}(n_{0,i}, \beta_{t,i}) * \mathrm{Pois}(\mu_i(1 - \beta_{t,i})) \right](n_{t,i}) \\ &= \prod_{i=1}^{S} \sum_{m=0}^{n_{t,i}} \mathrm{Bin}(n_{0,i}, \beta_{t,i})(m)\ \mathrm{Pois}(\mu_i(1 - \beta_{t,i}))\,(n_{t,i} - m) \end{aligned} \tag{14}$$

where $\mathrm{Bin}(N,p)$ denotes the binomial distribution $\mathrm{Bin}(N,p)(m) = \binom{N}{m}p^m(1-p)^{N-m}$, and $\mathrm{Pois}(\lambda)$ denotes the Poisson distribution $\mathrm{Pois}(\lambda)(m) = e^{-\lambda}\lambda^m/m!$.

Thus, conditional on $n_{0,i}$, the count $n_{t,i}$ is distributed as the sum of a binomial number of surviving initial count units and an independent Poisson number of newly created count units. Reverse sampling uses this analytic kernel as the bridge: it samples a clean-state proposal from the learned posterior, then an intermediate count state conditioned on that proposal and the current noisy state. As $t \to \infty$, $q_t \to \pi(\mathbf{n}) = \prod_i \mathrm{Pois}(\mu_i)(n_i)$. The PGF derivation is given in Appendix A.2.

Factorial cumulants. Let $\Phi_t = \log g_t$. The m-th factorial cumulant tensor is defined as $\kappa^{(m)}_{t,i_1\dots i_m} := \frac{\partial^m \Phi_t(\mathbf{z})}{\partial z_{i_1}\cdots\partial z_{i_m}}\Big|_{\mathbf{z}=\mathbf{1}}$. We denote the second factorial cumulant, also called the excess covariance, as $\Psi_{t,ij} := \kappa^{(2)}_{t,ij} = \mathrm{Cov}(n_i(t), n_j(t)) - \delta_{ij}\mathbb{E}[n_i(t)]$. The exact PGF solution then gives

$$\kappa^{(1)}_{t,i} = \mathbb{E}[n_i(t)] = \mu_i(1-\beta_{t,i}) + \beta_{t,i}\mathbb{E}[n_i(0)], \qquad \Psi_{t,ij} = \beta_{t,i}\beta_{t,j}\Psi_{0,ij}. \tag{15}$$

More generally, an order-m factorial cumulant is multiplied by the corresponding m survival factors, $\kappa^{(m)}_{t,i_1\dots i_m} = \beta_{t,i_1}\dots\beta_{t,i_m}\kappa^{(m)}_{0,i_1\dots i_m}$. We choose $k_i^- = 1$ and $\mu_i = k_i^+ = \mathbb{E}_{\mathcal{D}}[n_i]$. The means are then preserved, while every factorial cumulant of order $m \geq 2$ decays as e^{-mt}.

Terminal noising time. Let $\widehat{\boldsymbol{\mu}}_0 = \mathrm{diag}(\mu_1,\dots,\mu_S)$ and define the normalized excess-covariance signal along a direction $\mathbf{v} \in \mathbb{R}^S$ by

$$\Sigma_t(\mathbf{v}) = \frac{\mathbf{v}^\top \Psi_t \mathbf{v}}{\mathbf{v}^\top \widehat{\boldsymbol{\mu}}_0 \mathbf{v}}. \tag{16}$$

Let $\mathbf{w}_k$ be a generalized eigenvector of the pair $(\Psi_0, \widehat{\boldsymbol{\mu}}_0)$, with eigenvalue σ_k, so that $\Psi_0\mathbf{w}_k = \sigma_k\widehat{\boldsymbol{\mu}}_0\mathbf{w}_k$. Since $\Psi_t = e^{-2t}\Psi_0$, the signal along each generalized eigendirection decays as

$$\Sigma_t(\mathbf{w}_k) = e^{-2t}\sigma_k. \tag{17}$$

For a product-Poisson distribution, each coordinate has equal mean and variance and distinct coordinates are independent, so the population excess-covariance matrix vanishes, $\Psi = 0$. A finite sample, however, produces a nonzero apparent excess covariance through sampling fluctuations alone. We define σ_{noise} as the spectral level expected from these finite-sample fluctuations under the product-Poisson null. This threshold is determined in closed form by the number of species (genes) and observations (cells), without requiring Monte Carlo sampling; see Eq. (35).

Let σ_1 denote the largest generalized eigenvalue of $(\Psi_0, \widehat{\boldsymbol{\mu}}_0)$. We choose the terminal noising time as the time at which this strongest positive excess-covariance mode reaches the finite-sample Poisson noise floor,

$$T_O = \frac{1}{2}\log\frac{\sigma_1}{\sigma_{\mathrm{noise}}}, \qquad \sigma_1 > \sigma_{\mathrm{noise}}, \tag{18}$$

and set $T_O = 0$ otherwise. Thus, T_O is the earliest time at which the strongest positive second-order signal is no longer distinguishable from finite-sample fluctuations of the product-Poisson null. This is a second-order calibration criterion, not a guarantee of full distributional convergence. The spectral calculation and qualifications are given in Appendix A.3.

3.2 Conditional denoising with tilted FK steering

Let $\mathcal{D}_\theta$ be the endpoint distribution induced by the frozen CRNDiff generator and its unconditional reverse sampler, $\mathcal{D}_\theta^{\mathrm{tilt}}$ be its marginally tilted counterpart, and $\mathcal{D}_R$ be the target distribution. The two conditioning stages are

$$\mathcal{D}_\theta \xrightarrow{\text{marginal tilt}} \mathcal{D}_\theta^{\mathrm{tilt}} \xrightarrow{\text{residual correction}} \mathcal{D}_R.$$

The first stage uses the factorized posterior to adapt the proposal toward the target marginals; the second stage corrects the remaining joint discrepancy. We call this conditioning procedure tilted FK steering. Algorithm 2 summarizes the conditioning procedure, while Algorithm 1 describes generator training. The generator parameters remain frozen during conditioning. Target-specific preparation estimates $\mathcal{D}_{\theta,i}$ from an unconditional generation run, estimates $\mathcal{D}_{R,i}$ from target training cells, and fits the residual discriminator using target cells and marginally tilted proposal samples.

The final conditional run starts from the terminal distribution and interleaves tilted reverse transitions with FK weight updates and resampling.

Marginal tilting. For high-dimensional count data we sample the learned posterior coordinate-wise,

$$\widehat{Q}_\theta(\mathbf{n}_0 \mid \mathbf{n}_t, t) = \prod_{i=1}^{S} Q_{\theta,i}(n_{0,i} \mid \mathbf{n}_t, t). \tag{19}$$

Each posterior factor is conditioned on the full noisy count vector; coordinate-wise posterior sampling therefore does not imply a factorized endpoint distribution.

We define the marginally tilted posterior

$$\widehat{Q}^{\mathrm{tilt}}_\theta(\mathbf{n}_0 \mid \mathbf{n}_t, t) = \prod_i \frac{u_i(n_{0,i})\, Q_{\theta,i}(n_{0,i} \mid \mathbf{n}_t, t)}{\sum_m u_i(m)\, Q_{\theta,i}(m \mid \mathbf{n}_t, t)}, \tag{20}$$

where the coordinate-wise tilt is

$$u_i(n_i) := \left(\frac{\mathcal{D}_{R,i}(n_i)}{\mathcal{D}_{\theta,i}(n_i)}\right)^{\tau}, \qquad \tau \in (0, 1]. \tag{21}$$

Here $\mathcal{D}_{R,i}$ and $\mathcal{D}_{\theta,i}$ are the i-th coordinate marginals of the target and generator endpoint distributions, respectively. Because the endpoint weight $\prod_i u_i(n_{0,i})$ factorizes over coordinates, inserting it into Eq. (7) preserves the factorized posterior structure and tilts each coordinate separately. Combined with the original bridge, this defines a marginally tilted reverse process with endpoint distribution $\mathcal{D}^{\mathrm{tilt}}_\theta$.

Observe that u_i upweights count values that are more common under the target distribution and downweights those that are more common under the original sampler. The exponent τ controls the strength of the tilt: $\tau = 1$ applies the full marginal density ratio, while smaller values temper the correction toward the original posterior. Since these ratios are less reliably estimated for rare targets, we use stronger tempering for rarer target populations; see Appendix B.5. Importantly, tempering changes only the proposal $\mathcal{D}^{\mathrm{tilt}}_\theta$, not the target distribution $\mathcal{D}_R$.

Proposition 3.1 (Marginal information limitation)**.** *For a fixed τ, targets with identical one-dimensional marginals induce the same marginally tilted reverse process.*

The statement follows immediately because such targets give identical u_i. The tilt therefore adapts the proposal using marginal information, but cannot by itself distinguish targets that differ only in their joint dependence. Thus, the proposition concerns the information used by the marginal tilt ratios, not factorization of the endpoint law.

Residual correction. We correct the remaining discrepancy using the FK construction of Section 2.1.1. The discriminator $r_\phi(\mathbf{n})$ estimates the probability of the target class $\mathcal{D}_R$ versus the proposal class $\mathcal{D}^{\mathrm{tilt}}_\theta$. With equal class priors, the population-optimal discriminator under binary cross-entropy has odds (Goodfellow et al., 2014)

$$\rho^*(\mathbf{n}) = \frac{r^*(\mathbf{n})}{1 - r^*(\mathbf{n})} = \frac{\mathcal{D}_R(\mathbf{n})}{\mathcal{D}^{\mathrm{tilt}}_\theta(\mathbf{n})}. \tag{22}$$

We therefore use $\rho_\phi(\mathbf{n}) = r_\phi(\mathbf{n})/(1 - r_\phi(\mathbf{n}))$ as the FK endpoint weight. Applying the potentials of Section 2.1.1 to the marginally tilted reverse chain gives

$$\mathcal{D}_{\mathrm{FK}}(\mathbf{n}) \propto \mathcal{D}^{\mathrm{tilt}}_\theta(\mathbf{n})\rho_\phi(\mathbf{n}), \tag{23}$$

which equals $\mathcal{D}_R$ for the exact density ratio.

In practice, we estimate the intermediate FK potentials from predicted clean endpoints and resample when the effective sample size (ESS) becomes small; the estimator and particle construction are given in Appendix A.4.

Algorithm 1 CRNDiff training

Require: count data $\mathcal{D}$ with N cells and S genes; posterior network Q_θ

$\mu_i \leftarrow \mathbb{E}_{\mathcal{D}}[n_i]$; $k_i^- \leftarrow 1$, $k_i^+ \leftarrow \mu_i$ Eq. (13)
$\sigma_1 \leftarrow \max\{\sigma : \Psi_0 \mathbf{w} = \sigma \widehat{\boldsymbol{\mu}}_0 \mathbf{w}\}$ Eq. (17)
$\sigma_{\text{noise}} \leftarrow (1 + \sqrt{S/N})^2 - 1$ Eq. (35)
$T_O \leftarrow \max\{\frac{1}{2}\log(\sigma_1/\sigma_{\text{noise}}),\, 0\}$ Eq. (18)
while not converged **do**
  $\mathbf{n}_0 \sim \mathcal{D}$, $t \sim \mathcal{U}(0, T_O]$
  $\mathbf{n}_t \sim \overrightarrow{q}_t(\cdot \mid \mathbf{n}_0)$ Eq. (14)
  update θ on $-\sum_i \log Q_{\theta,i}(n_{0,i} \mid \mathbf{n}_t, t)$
end while
return Q_θ, T_O

Algorithm 2 Tilted FK steering

Require: Q_θ, T_O; target marginals $\mathcal{D}_{R,i}$, exponent τ; odds ρ_ϕ of Eq. (22); M particles; resampling steps $\mathcal{C}$, ESS threshold ϑ

$\mathcal{D}_{\theta,i} \leftarrow$ marginals of an unconditional run of Q_θ
$\widehat{Q}_\theta^{\text{tilt}} \leftarrow$ tilt Q_θ with $u_i = (\mathcal{D}_{R,i}/\mathcal{D}_{\theta,i})^\tau$ Eq. (20)
$\mathbf{n}_{t_K}^{(m)} \sim \pi$, $w^{(m)} \leftarrow 1$ for $m = 1, \ldots, M$
for $j = K, \ldots, 1$ **do**
  $\mathbf{n}_{t_{j-1}}^{(m)} \sim$ tilted reverse step, all m Eqs. (4),(20)
  $w^{(m)} \leftarrow w^{(m)}\, G_j^{(m)}$, $G_j = \widehat{\psi}_{j-1}/\widehat{\psi}_j$ Eq. (38)
  if $j{-}1 \in \mathcal{C}$, $\text{ESS} < \vartheta M$: resample_w; $w \leftarrow 1$
end for
return n_{out} of the $\mathbf{n}_0^{(m)}$, drawn $\propto w^{(m)}$

4 Related work

Discrete and count-native generative models. Discrete diffusion and its continuous-time formulations generate through Markov transitions on discrete states (Austin et al., 2021; Campbell et al., 2022); masked diffusion language models instantiate absorbing-state transitions (Sahoo et al., 2024; Shi et al., 2025), and discrete flow matching and generator matching separate the probability path from the parameterization of its dynamics (Gat et al., 2024; Holderrieth et al., 2025). On integer counts, Blackout Diffusion and CountsDiff use death or thinning dynamics (Santos et al., 2023; Soatto et al., 2026), JUMP connects generation to thinning and thickening (Chen & Zhou, 2023), and Count Bridges and count-FM employ birth–death jump processes (Fishman et al., 2026; Wei & Pearson, 2026). CRNDiff similarly instantiates a birth–death process as a tractable count-space noising mechanism. More generally, the CRN formulation treats this construction as one instance within a broader class of reaction-network-based forward processes.

Stochastic kinetics and generating functions. The operator representation used here originates in Doi–Peliti theory (Doi, 1976; Peliti, 1985); related work connects CRNs, stochastic Petri nets, creation–annihilation operators, and factorial-moment hierarchies (Baez & Biamonte, 2018; Smith & Krishnamurthy, 2017). Our kernels, bridges, and factorial-cumulant dynamics instantiate established exact solutions for monomolecular and first-order systems (Jahnke & Huisinga, 2007; Vastola, 2021).

Conditional generation and inference-time steering. Best-of-N selection ranks generated candidates by target score (Singhal et al., 2025), while value-guided resampling injects target information during reverse sampling (Li et al., 2024). Twisted diffusion sampling and FK steering use sequential Monte Carlo (SMC) potentials to steer a frozen generator (Wu et al., 2023; Singhal et al., 2025), but rare targets can cause concentrated importance weights and genealogical degeneracy (Del Moral, 2004; Chopin et al., 2020; Jacob et al., 2015). CRNDiff adapts the reverse proposal toward target marginals before applying the joint density-ratio correction.

Single-cell generative models. scVI, scANVI, and CFGen model single-cell counts through latent representations and count-valued decoders (Lopez et al., 2018; Xu et al., 2021; Palma et al., 2025); CRNDiff instead models count-valued dynamics directly and conditions target populations at sampling time.

5 Experiments

We evaluate CRNDiff on a synthetic count mixture and human-heart scRNA-seq data, assessing unconditional and conditional generation, genealogical diversity, and downstream cell-type classification. Appendices B.2 and B.3 provide experimental settings and metric definitions.

5.1 A count mixture with matched-marginal targets

We construct an eight-component count mixture with four abundant and four rare components (Figure 1). The target assigns equal mass to rare components c_1 and c_5, together 1% of the training

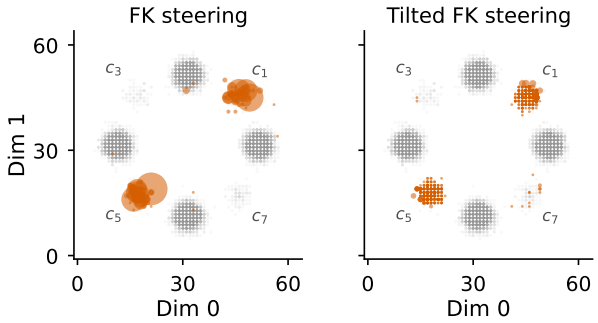

Figure 1: Generated samples from FK steering and tilted FK steering targeting c_1, c_5. The alternative pair c_3, c_7 has identical one-dimensional marginals. Gray markers show training samples; orange markers show one generated representative per initial ancestor, sized by output multiplicity.

mixture; the alternative pair c_3, c_7 has identical one-dimensional marginals but different joint structure.

On the same frozen generator, FK steering (Wu et al., 2023; Singhal et al., 2025) and tilted FK steering use untilted and marginally tilted proposals, respectively. Table 1 reports energy distance (ED), mean marginal total variation (mTV), and genealogical diversity (Anc.). Here Anc. is the number of distinct initial ancestors in the output divided by output sample size. The independent and identically distributed (i.i.d.) row is the target reference. Tilted FK steering reduces the genealogical concentration under FK steering and approaches the reference on all three measures. Appendix B.2 shows that marginal tilt alone cannot distinguish the matched-marginal component pairs and illustrates how marginal tilt and FK steering alter the reverse-chain particle distribution.

Table 1: Comparison on the two-dimensional count mixture; metrics are defined in Appendix B.2.

Method	ED ↓	mTV ↓	Anc. ↑
FK steering	0.255	0.352	0.116
Tilted FK	**0.036**	**0.117**	**0.875**
i.i.d. target	0.017	0.050	1.000

5.2 Generation and evaluation on single-cell data

We compare CRNDiff with scVI, scANVI, CFGen, and MDLM on scRNA-seq data from endothelial, myeloid, and neuronal cells in the adult human heart atlas (Litviňuková et al., 2020; Lopez et al., 2018; Xu et al., 2021; Palma et al., 2025; Sahoo et al., 2024). The CRNDiff generator is trained unconditionally; tilted FK steering conditions its samples on the target cell type. All fidelity evaluations use the held-out test set as the real-data reference. Appendix B.3 details preprocessing, model configurations, sampling, and evaluation.

Conditional fidelity and steering. Table 2 evaluates conditional fidelity using purity, mean per-gene Wasserstein-1 distance (W_1), squared maximum mean discrepancy (MMD^2), and Pearson correlation coefficient (PCC). Purity is the fraction of generated cells assigned to the target population by a fixed external CellTypist heart model (Domínguez Conde et al., 2022). W_1, MMD^2, and PCC compare generated and test cells at the per-gene, joint-distribution, and mean-expression levels, respectively. CRNDiff has the highest mean purity and PCC, and the lowest mean W_1 and MMD^2, for all three targets. Its mean purity margin over the strongest baseline increases as training abundance decreases, and neuronal mean purity approaches the real-cell reference.

To isolate conditioning effects from generator quality, Tables 6, 8, and 10 evaluate conditioning choices on the frozen CRNDiff generator (Appendices B.6 and B.5). Larger best-of-N pools and stronger value-guided resampling can increase purity and sometimes match or exceed tilted FK steering, often with higher W_1 and MMD^2. Ablations show that tilted FK steering preserves more distinct initial ancestors across targets, with metric-dependent distributional trade-offs.

Unconditional generation. While Table 2 reflects both generator quality and conditioning, Table 9, which evaluates generators without conditional steering, reports sliced Wasserstein-1 distance (sliced-W_1), median Fano factor, coefficients of variation for library size and detected-gene counts, and CellTypist-assigned target fractions. CRNDiff attains the lowest sliced-W_1 and near-reference detected-gene variability, whereas lower library-size variability and median Fano factor indicate generator-side dispersion limits.

Table 2: Conditional fidelity for endothelial, myeloid, and neuronal targets. Distributional metrics use held-out target cells as reference; purity uses fixed CellTypist target-label mappings. Results are mean ± sd over three training seeds. Boldface indicates the best mean among generative models.

Method	Purity ↑	W_1 ↓	MMD^2 ↓	PCC ↑
Endothelial *abundant*, 80,463 training cells				
Real cells	0.944	0.011	0.0001	1.000
CRNDiff	$\mathbf{0.906}_{\pm.005}$	$\mathbf{0.064}_{\pm.003}$	$\mathbf{0.0011}_{\pm.0001}$	$\mathbf{0.998}_{\pm.000}$
MDLM (Sahoo et al., 2024)	$0.898_{\pm.026}$	$0.078_{\pm.029}$	$0.0058_{\pm.0010}$	$0.991_{\pm.003}$
CFGen (Palma et al., 2025)	$0.797_{\pm.009}$	$0.102_{\pm.001}$	$0.0036_{\pm.0000}$	$0.993_{\pm.000}$
scANVI (Xu et al., 2021)	$0.857_{\pm.016}$	$0.079_{\pm.007}$	$0.0016_{\pm.0001}$	$0.997_{\pm.000}$
scVI (Lopez et al., 2018)	$0.423_{\pm.014}$	$0.241_{\pm.004}$	$0.0516_{\pm.0002}$	$0.862_{\pm.001}$
Myeloid 18,422 training cells				
Real cells	0.929	0.037	0.0003	0.999
CRNDiff	$\mathbf{0.886}_{\pm.014}$	$\mathbf{0.110}_{\pm.017}$	$\mathbf{0.0032}_{\pm.0010}$	$\mathbf{0.987}_{\pm.005}$
MDLM	$0.840_{\pm.072}$	$0.147_{\pm.012}$	$0.0060_{\pm.0027}$	$0.979_{\pm.014}$
CFGen	$0.625_{\pm.027}$	$0.217_{\pm.011}$	$0.0052_{\pm.0003}$	$0.982_{\pm.001}$
scANVI	$0.625_{\pm.018}$	$0.196_{\pm.009}$	$0.0056_{\pm.0008}$	$0.981_{\pm.004}$
scVI	$0.482_{\pm.040}$	$0.268_{\pm.002}$	$0.0190_{\pm.0004}$	$0.919_{\pm.002}$
Neuronal *rare*, 3,168 training cells				
Real cells	0.982	0.031	0.0019	0.993
CRNDiff	$\mathbf{0.971}_{\pm.012}$	$\mathbf{0.051}_{\pm.003}$	$\mathbf{0.0102}_{\pm.0018}$	$\mathbf{0.956}_{\pm.005}$
MDLM	$0.784_{\pm.063}$	$0.059_{\pm.011}$	$0.0120_{\pm.0030}$	$0.944_{\pm.019}$
CFGen	$0.734_{\pm.004}$	$0.063_{\pm.003}$	$0.0130_{\pm.0003}$	$0.943_{\pm.006}$
scANVI	$0.269_{\pm.019}$	$0.079_{\pm.003}$	$0.0478_{\pm.0027}$	$0.719_{\pm.021}$
scVI	$0.535_{\pm.073}$	$0.074_{\pm.003}$	$0.0394_{\pm.0062}$	$0.781_{\pm.037}$

Marker preservation and downstream utility. In addition, Table 4 compares differential-expression rankings from generated and real test cells; CRNDiff has the highest mean top-100 marker overlap and Spearman correlation for endothelial and myeloid cells, while neuronal results are metric-dependent. Table 5 evaluates generated cells as replacement training data for the three target classes in an eleven-class classifier. Class-specific F_1 is the harmonic mean of precision and recall, and macro F_1 is the unweighted mean across all eleven classes. CRNDiff achieves the highest mean macro F_1, myeloid F_1, and neuronal F_1.

6 Discussion

CRNDiff specifies count-space diffusion through reaction channels and propensity functions. In the birth–death instance studied here, closed-form transition kernels, Markov bridges, and factorial-cumulant dynamics support reverse sampling, terminal-time selection, and inference-time conditioning of a frozen generator. The heart-atlas results suggest that separating generator learning from sampler-level conditioning is useful for rare target populations. Across three target cell-types, CRNDiff improves conditional fidelity, preserves marker-level differential-expression structure, supports downstream classification, and reduces genealogical concentration relative to FK steering.

These conclusions have boundaries: terminal noising uses second-order excess covariance, higher-order dependence may require additional diagnostics, and finite-particle, density-ratio, or target-discriminator errors can shift rare-target purity–fidelity trade-offs. Lower library-size variability and median Fano factor indicate that steering selects target populations within the learned count distribution rather than repairing dispersion mismatch (Appendix B.6). Future work should test coupled reactions and structured-state extensions while preserving tractable kernels, bridges, and samplers (Danos et al., 2007; Behr et al., 2016).

Reproducibility Statement

Code, configuration files, and the metric summaries used to construct all tables are available at `https://anonymous.4open.science/r/CRNDiff-EDCA/`. Appendices B.2–B.6 document data preprocessing, model training, sampler settings, and evaluation protocols.

LARGE LANGUAGE MODEL STATEMENT

Large language models and associated AI tools were used for the following:

1. **Writing.** LLM-based tools were used to assist with the retrieval and discovery of related work and to polish portions of the manuscript for clarity and readability. All cited references and AI-assisted text were checked, reviewed, and revised by the authors.
2. **Code.** AI coding assistants were used to assist with portions of code implementation and debugging. All AI-assisted code was reviewed and tested by the authors.
3. **Not used.** LLMs were not used to develop the theoretical or conceptual framework, formulate or prove mathematical claims, propose the method or hypotheses, design the experiments, generate reported data, or interpret the experimental results.

The authors take full responsibility for the final content of this work, including all text, claims, code, and other artifacts produced with the aid of generative AI.

REFERENCES

Jacob Austin, Daniel D. Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. In *Advances in Neural Information Processing Systems*, volume 34, pp. 17981–17993, 2021.

John Baez and Jacob D Biamonte. *Quantum techniques in stochastic mechanics*. World Scientific, 2018.

Nicolas Behr, Vincent Danos, and Ilias Garnier. Stochastic mechanics of graph rewriting. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science*, pp. 46–55, 2016.

Andrew Campbell, Joe Benton, Valentin De Bortoli, Tom Rainforth, George Deligiannidis, and Arnaud Doucet. A continuous time framework for discrete denoising models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 28266–28279, 2022.

Chris K Carter and Robert Kohn. On gibbs sampling for state space models. *Biometrika*, 81(3): 541–553, 1994.

Tianqi Chen and Mingyuan Zhou. Learning to jump: Thinning and thickening latent counts for generative modeling. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 5367–5382. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/chen23ap.html.

Nicolas Chopin, Omiros Papaspiliopoulos, et al. *An introduction to sequential Monte Carlo*, volume 4. Springer, 2020.

Vincent Danos, Jérôme Feret, Walter Fontana, Russell Harmer, and Jean Krivine. Rule-based modelling of cellular signalling. In *International conference on concurrency theory*, pp. 17–41. Springer, 2007.

Pierre Del Moral. *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. Probability and Its Applications. Springer, New York, NY, 2004. doi: 10.1007/978-1-4684-9393-1.

Masao Doi. Second quantization representation for classical many-particle system. *Journal of Physics A: Mathematical and General*, 9(9):1465–1477, 1976.

Cecilia Domínguez Conde, Chuan Xu, Lorna B. Jarvis, David B. Rainbow, Sarah B. Wells, Tom Gomes, Sarah K. Howlett, Olga Suchanek, Krzysztof Polanski, Helena W. King, et al. Cross-tissue immune cell analysis reveals tissue-specific features in humans. *Science*, 376(6594): eabl5197, 2022. doi: 10.1126/science.abl5197.

Nic Fishman, Gokul Gowri, Tanush Kumar, Jiaqi Lu, Valentin de Bortoli, Jonathan S. Gootenberg, and Omar Abudayyeh. Count bridges enable modeling and deconvolving transcriptomic data, 2026. URL https://arxiv.org/abs/2603.04730.

Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky T. Q. Chen, Gabriel Synnaeve, Yossi Adi, and Yaron Lipman. Discrete flow matching, 2024. URL https://arxiv.org/abs/2407.15595.

Daniel T. Gillespie. A rigorous derivation of the chemical master equation. *Physica A: Statistical Mechanics and its Applications*, 188(1):404–425, 1992. ISSN 0378-4371. doi: https://doi.org/10.1016/0378-4371(92)90283-V. URL https://www.sciencedirect.com/science/article/pii/037843719290283V.

Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

Peter Holderrieth, Marton Havasi, Jason Yim, Neta Shaul, Itai Gat, Tommi Jaakkola, Brian Karrer, Ricky T. Q. Chen, and Yaron Lipman. Generator matching: Generative modeling with arbitrary markov processes, 2025. URL https://arxiv.org/abs/2410.20587.

Pierre E. Jacob, Lawrence M. Murray, and Sylvain Rubenthaler. Path storage in the particle filter. *Statistics and Computing*, 25(2):487–496, 2015. doi: 10.1007/s11222-013-9445-x.

Tobias Jahnke and Wilhelm Huisinga. Solving the chemical master equation for monomolecular reaction systems analytically. *Journal of Mathematical Biology*, 54(1):1–26, 2007. doi: 10.1007/s00285-006-0034-x.

Xiner Li, Yulai Zhao, Chenyu Wang, Gabriele Scalia, Gokcen Eraslan, Surag Nair, Tommaso Biancalani, Shuiwang Ji, Aviv Regev, Sergey Levine, and Masatoshi Uehara. Derivative-free guidance in continuous and discrete diffusion models with soft value-based decoding, 2024. URL https://arxiv.org/abs/2408.08252.

Monika Litviňuková, Carlos Talavera-López, Henrike Maatz, Daniel Reichart, Catherine L Worth, Eric L Lindberg, Masatoshi Kanda, Krzysztof Polanski, Matthias Heinig, Michael Lee, et al. Cells of the adult human heart. *Nature*, 588(7838):466–472, 2020.

Romain Lopez, Jeffrey Regier, Michael B Cole, Michael I Jordan, and Nir Yosef. Deep generative modeling for single-cell transcriptomics. *Nature methods*, 15(12):1053–1058, 2018.

V. A. Marčenko and L. A. Pastur. Distribution of eigenvalues for some sets of random matrices. *Mathematics of the USSR-Sbornik*, 1(4):457–483, 1967. doi: 10.1070/SM1967v001n04ABEH001994.

Alessandro Palma, Till Richter, Hanyi Zhang, Manuel Lubetzki, Alexander Tong, Andrea Dittadi, and Fabian Theis. Multi-modal and multi-attribute generation of single cells with cfgen, 2025. URL https://arxiv.org/abs/2407.11734.

Luca Peliti. Path integral approach to birth-death processes on a lattice. *Journal de Physique*, 46(9):1469–1483, 1985.

Subham Sekhar Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin T Chiu, Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language models. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 130136–130184. Curran Associates, Inc., 2024. doi: 10.52202/079017-4135. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/eb0b13cc515724ab8015bc978fdde0ad-Paper-Conference.pdf.

Javier E Santos, Zachary R. Fox, Nicholas Lubbers, and Yen Ting Lin. Blackout diffusion: Generative diffusion models in discrete-state spaces, 2023. URL https://arxiv.org/abs/2305.11089.

Dino Sejdinovic, Bharath Sriperumbudur, Arthur Gretton, and Kenji Fukumizu. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *The Annals of Statistics*, 41(5): 2263–2291, 2013. doi: 10.1214/13-AOS1140.

Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis K. Titsias. Simplified and generalized masked diffusion for discrete data, 2025. URL https://arxiv.org/abs/2406.04329.

Raghav Singhal, Zachary Horvitz, Ryan Teehan, Mengye Ren, Zhou Yu, Kathleen Mckeown, and Rajesh Ranganath. A general framework for inference-time scaling and steering of diffusion models. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 55810–55827. PMLR, 13–19 Jul 2025. URL https://proceedings.mlr.press/v267/singhal25b.html.

Eric Smith and Supriya Krishnamurthy. Flows, scaling, and the control of moment hierarchies for stochastic chemical reaction networks. *Physical review. E*, 96(6):062102, 2017.

Renzo G. Soatto, Anders Hoel, Greycen Ren, Shorna Alam, Stephen Bates, Nikolaos P. Daskalakis, Caroline Uhler, and Maria Skoularidou. Countsdiff: A diffusion model on the natural numbers for generation and imputation of count-based data, 2026. URL https://arxiv.org/abs/2604.03779.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.

Gábor J. Székely and Maria L. Rizzo. Energy statistics: A class of statistics based on distances. *Journal of Statistical Planning and Inference*, 143(8):1249–1272, 2013. doi: 10.1016/j.jspi.2013.03.018.

John J Vastola. Solving the chemical master equation for monomolecular reaction systems and beyond: a doi-peliti path integral view. *Journal of Mathematical Biology*, 83(5):48, 2021.

Ganchao Wei and John Pearson. Flow matching for count data, 2026. URL https://arxiv.org/abs/2605.07746.

Luhuan Wu, Brian L. Trippe, Christian A. Naesseth, David M. Blei, and John P. Cunningham. Practical and asymptotically exact conditional sampling in diffusion models. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Chenling Xu, Romain Lopez, Edouard Mehlman, Jeffrey Regier, Michael I Jordan, and Nir Yosef. Probabilistic harmonization and annotation of single-cell transcriptomics data with deep generative models. *Molecular systems biology*, 17(1):MSB20209620, 2021.

A Derivations

A.1 Reverse process, FFBS, and endpoint tilting

Reverse master equation.

Let $r_s(x) = q_{T-s}(x)$. Differentiating with respect to reverse time,

$$\begin{aligned}\partial_s r_s(x) &= \sum_y \left[w(x \to y) r_s(x) - w(y \to x) r_s(y)\right] \\ &= \sum_y \left[\widetilde{w}_s(y \to x) r_s(y) - \widetilde{w}_s(x \to y) r_s(x)\right], \end{aligned} \tag{24}$$

where

$$\widetilde{w}_s(y \to x) = w(x \to y) \frac{r_s(x)}{r_s(y)}.$$

This gives Eq. (2).

Writing $\mathbb{S}^r := \mathbb{S}_{\cdot,r}$, the reverse of CRN reaction r changes $\mathbf{n} \to \mathbf{n} - \mathbb{S}^r$, with propensity

$$\widetilde{j}_r(\mathbf{n}, s) = j_r(\mathbf{n} - \mathbb{S}^r) \frac{q_{T-s}(\mathbf{n} - \mathbb{S}^r)}{q_{T-s}(\mathbf{n})}. \tag{25}$$

Hence

$$\partial_s \widetilde{q}_s(\mathbf{n}) = \sum_r \left[\widetilde{j}_r(\mathbf{n} + \mathbb{S}^r, s) \widetilde{q}_s(\mathbf{n} + \mathbb{S}^r) - \widetilde{j}_r(\mathbf{n}, s) \widetilde{q}_s(\mathbf{n})\right], \qquad \widetilde{q}_s = q_{T-s}. \tag{26}$$

Reverse propensities generally need not be mass action even when the forward propensities are.

Exact backward sampling. Substituting Eqs. (3)–(4),

$$\begin{aligned}\sum_{x_0} Q(x_0 \mid x_t, t) B(x_s \mid x_0, x_t) &= \frac{\overrightarrow{q}(x_t \mid x_s)}{q_t(x_t)} \sum_{x_0} \overrightarrow{q}(x_s \mid x_0) \mathcal{D}(x_0) \\ &= \overrightarrow{q}(x_t \mid x_s) \frac{q_s(x_s)}{q_t(x_t)}. \end{aligned} \tag{27}$$

Consequently,

$$q_T(x_{t_K}) \prod_{j=1}^{K} \overleftarrow{q}(x_{t_{j-1}} \mid x_{t_j}) = \mathcal{D}(x_{t_0}) \prod_{j=1}^{K} \overrightarrow{q}(x_{t_j} \mid x_{t_{j-1}}), \tag{28}$$

which proves exact FFBS on any grid when initialized from q_T.

Endpoint tilting. If the endpoint distribution P is replaced by

$$P^H(x_0) = \frac{H(x_0) P(x_0)}{\sum_{x'} H(x') P(x')},$$

Bayes' rule immediately gives

$$Q^H(x_0 \mid x_t, t) = \frac{H(x_0) Q(x_0 \mid x_t, t)}{\sum_{x'} H(x') Q(x' \mid x_t, t)},$$

proving Eq. (7). The bridge is unchanged because the forward transition process is unchanged.

A.2 Generating functions and birth–death solution

Mass-action PGF equation. For $f_{\mathbf{z}}(\mathbf{n}) = \mathbf{z}^{\mathbf{n}}$, $\partial_t \mathbb{E}[f_{\mathbf{z}}(\mathbf{n}_t)] = \mathbb{E}[(L f_{\mathbf{z}})(\mathbf{n}_t)]$ gives

$$\begin{aligned}\partial_t g(\mathbf{z}, t) &= \sum_r \mathbb{E}\left[j_r(\mathbf{n}) \left(\mathbf{z}^{\mathbf{n} + \mathbb{S}^r} - \mathbf{z}^{\mathbf{n}}\right)\right] \\ &= \sum_r k_r V^{1 - |\mathbb{S}^-_{\cdot,r}|} \left(\mathbf{z}^{\mathbb{S}^+_{\cdot,r}} - \mathbf{z}^{\mathbb{S}^-_{\cdot,r}}\right) \mathbb{E}\left[\prod_i \frac{n_i!}{(n_i - \mathbb{S}^-_{i,r})!} \mathbf{z}^{\mathbf{n} - \mathbb{S}^-_{\cdot,r}}\right]. \end{aligned} \tag{29}$$

Since $\mathbb{E}\left[\prod_i \frac{n_i!}{(n_i-a_i)!}\mathbf{z}^{\mathbf{n}-\mathbf{a}}\right] = \partial_{\mathbf{z}}^{\mathbf{a}} g$, this proves Eq. (12).

Birth–death characteristics.

For one species,

$$\partial_t g = (z-1)(k^+ - k^- \partial_z) g, \tag{30}$$

with Lagrange–Charpit equations

$$dt = \frac{dz}{k^-(z-1)} = \frac{dg}{k^+(z-1)g}. \tag{31}$$

Writing $\beta_t = e^{-k^- t}$ and $\mu = k^+/k^-$, the characteristic through z_0 satisfies

$$z_0 = 1 + \beta_t(z-1), \qquad \frac{dg}{g} = \mu\, dz,$$

and therefore

$$g_t(z) = e^{\mu(1-\beta_t)(z-1)} g_0(1 + \beta_t(z-1))\,. \tag{32}$$

For independent birth–death coordinates this gives

$$g_t(\mathbf{z}) = \exp\left[\sum_i \mu_i(1-\beta_{t,i})(z_i-1)\right] g_0(\mathbf{1} + \boldsymbol{\beta}_t \odot (\mathbf{z}-\mathbf{1}))\,, \tag{33}$$

where $\odot$ is an element-wise or Hadamard product. No factorization of g_0 is required. Setting $g_0(\mathbf{z}) = \mathbf{z}^{\mathbf{n}_0}$ and extracting coefficients gives Eq. (14).

Taking $\Phi_t = \log g_t$ yields

$$\Phi_t(\mathbf{z}) = \sum_i \mu_i(1-\beta_{t,i})(z_i-1) + \Phi_0(\mathbf{1} + \boldsymbol{\beta}_t \odot (\mathbf{z}-\mathbf{1}))\,, \tag{34}$$

from which Eq. (15) and the general order-m factorial-cumulant decay follow by differentiation.

A.3 Spectral terminal-time criterion

Under our parameter choice, $\widehat{\boldsymbol{\mu}}_t = \widehat{\boldsymbol{\mu}}_0$ and $\Psi_t = e^{-2t}\Psi_0$. Hence

$$\Sigma_t(\mathbf{v}) = e^{-2t}\frac{\mathbf{v}^\top \Psi_0 \mathbf{v}}{\mathbf{v}^\top \widehat{\boldsymbol{\mu}}_0 \mathbf{v}}.$$

The stationary directions of this generalized Rayleigh quotient satisfy

$$\Psi_0 \mathbf{w}_k = \sigma_k \widehat{\boldsymbol{\mu}}_0 \mathbf{w}_k,$$

and therefore $\Sigma_t(\mathbf{w}_k) = e^{-2t}\sigma_k$. Equivalently, the σ_k are the eigenvalues of

$$C = \widehat{\boldsymbol{\mu}}_0^{-1/2} \Psi_0 \widehat{\boldsymbol{\mu}}_0^{-1/2}.$$

Thus the strongest positive second-order signal is $e^{-2t}\sigma_1$, and equating it to the empirical Poisson threshold σ_{noise} gives Eq. (18). Under the product-Poisson null the excess covariance vanishes in population, so σ_{noise} is the largest normalized eigenvalue that a sample of N cells over S genes produces by fluctuation alone. We take it to be the Marchenko–Pastur edge (Marčenko & Pastur, 1967)

$$\sigma_{\text{noise}} = \left(1 + \sqrt{S/N}\right)^2 - 1, \tag{35}$$

a closed-form function of S and N that involves no sampling. This criterion calibrates the decay of detectable positive excess-covariance structure; it does not by itself provide a bound on the discrepancy between q_{T_O} and π. For the independent birth–death model, factorial cumulants of order $m \geq 2$ decay as e^{-mt}, but the thresholding rule is tied to the strongest positive second-order mode.

A.4 CONDITIONAL SAMPLER DETAILS

Marginal information limitation. Suppose two target distributions P_1 and P_2 have identical coordinate marginals. They then give identical marginal ratios u_i in Eq. (21), and hence identical tilted posteriors in Eq. (20) at every reverse step. Their marginally tilted reverse processes are therefore identical, proving Proposition 3.1.

This does not imply that the resulting endpoint distribution is a product distribution: the factors $Q_{\theta,i}(n_{0,i} \mid \mathbf{n}_t, t)$ share the full conditioning variable $\mathbf{n}_t$, and marginalizing over $\mathbf{n}_t$ can induce dependence.

Residual density ratio. For balanced binary classification between $\mathcal{D}_R$ and $\mathcal{D}_\theta^{\text{tilt}}$, the population log-loss minimizer is

$$r^*(\mathbf{n}) = \frac{\mathcal{D}_R(\mathbf{n})}{\mathcal{D}_R(\mathbf{n}) + \mathcal{D}_\theta^{\text{tilt}}(\mathbf{n})}.$$

Its odds therefore satisfy

$$\frac{r^*(\mathbf{n})}{1 - r^*(\mathbf{n})} = \frac{\mathcal{D}_R(\mathbf{n})}{\mathcal{D}_\theta^{\text{tilt}}(\mathbf{n})},$$

which gives Eq. (22).

Intermediate FK potentials. For the marginally tilted reverse chain, an ideal continuation potential is

$$\psi_j(\mathbf{n}_{t_j}) = \mathbb{E}\left[\rho_\phi(\mathbf{n}_0) \mid \mathbf{n}_{t_j}\right], \qquad 0 < j < K. \tag{36}$$

We approximate it using J clean-state predictions,

$$\widehat{\psi}_j(\mathbf{n}_{t_j}) = \frac{1}{J}\sum_{\ell=1}^{J} \rho_\phi\left(\widetilde{\mathbf{n}}_0^{(\ell)}\right), \qquad \widetilde{\mathbf{n}}_0^{(\ell)} \sim \widehat{Q}_\theta^{\text{tilt}}(\cdot \mid \mathbf{n}_{t_j}, t_j). \tag{37}$$

With $\widehat{\psi}_K \equiv 1$ and $\widehat{\psi}_0(\mathbf{n}_0) = \rho_\phi(\mathbf{n}_0)$, the incremental weights are

$$G_j = \frac{\widehat{\psi}_{j-1}(\mathbf{n}_{t_{j-1}})}{\widehat{\psi}_j(\mathbf{n}_{t_j})}. \tag{38}$$

Along a particle lineage the same realized potential is carried forward, so the product telescopes:

$$\prod_{j=1}^{K} G_j = \rho_\phi(\mathbf{n}_0).$$

Consequently the weighted endpoint distribution is

$$\mathcal{D}_\theta^{\text{tilt}}(\mathbf{n})\rho_\phi(\mathbf{n}),$$

up to normalization, giving Eq. (23). Intermediate approximations affect weight variance and resampling efficiency; the terminal density-ratio correction determines the formal endpoint target.

Sampling configuration. For the single-cell experiments, the sampler uses $K = 32$ reverse steps, $M = 2n_{\text{out}}$ particles to generate n_{out} cells, and $J = 16$ endpoint draws per particle for the intermediate potentials, taken from the tilted posterior of the same step through a separate random stream so that the proposal trajectory is unchanged. Resampling is systematic and is triggered when $\text{ESS}/M < 0.5$ ($\vartheta = 0.5$ in Algorithm 2), subject to the two-step minimum gap listed in Table 3, with

$$\text{ESS} = \frac{\left(\sum_{m=1}^{M} w^{(m)}\right)^2}{\sum_{m=1}^{M}\left(w^{(m)}\right)^2}. \tag{39}$$

After all reverse steps, the reported n_{out} cells are selected from the M terminal particles by weighted systematic resampling using the final weights. This terminal weighted output selection is separate from the chain-internal resampling operations triggered by the ESS rule. Diagnostic resampling-event counts exclude the terminal output selection unless explicitly stated otherwise. The tempering exponent and the remaining settings are listed in Table 3.

A.5 CONNECTION TO DIFFUSION MODELS

The CRN construction contains ordinary continuous diffusion as its large-system-size approximation. Using $e^{\Delta\cdot\partial_n} f(n) \equiv f(n+\Delta)$, the chemical master equation can be written

$$\partial_t q_t(n) = \sum_r \left(e^{-\mathbb{S}^r\cdot\partial_n} - 1\right) j_r(n) q_t(n).$$

Set $x = n/V$, and let $p_t(x)$ denote its density in the continuous approximation. For mass action, $j_r(Vx) = V a_r(x) + O(1)$, with $a_r(x) = k_r x^{\mathbb{S}^{\cdot,r}}$. Expanding to second order in V^{-1} gives

$$\partial_t p_t = -\nabla\cdot(F p_t) + \frac{1}{2V}\sum_{i,j} \partial_{x_i}\partial_{x_j}(D_{ij} p_t) + O(V^{-2}), \tag{40}$$

where

$$F(x) = \sum_r \mathbb{S}^r a_r(x), \qquad D(x) = \sum_r \mathbb{S}^r (\mathbb{S}^r)^\top a_r(x).$$

Equivalently,

$$dX_t = F(X_t)\,dt + V^{-1/2}\sigma(X_t)\,dW_t, \qquad \sigma = \sum_r \mathbb{S}^r \sqrt{a_r(x)}. \tag{41}$$

Write $\rho_s(n) = q_{T-s}(n)$ for the reverse-time mass function, and use $\rho_s(x) = p_{T-s}(x)$ for its continuous density approximation. Expanding the reverse propensities

$$\widetilde{j}_r(n,s) = j_r(n-\mathbb{S}^r)\frac{\rho_s(n-\mathbb{S}^r)}{\rho_s(n)}$$

in the same way gives the reverse chemical Langevin equation

$$dX_s = \left[-F(X_s) + \frac{1}{V}\left(\nabla\cdot D(X_s) + D(X_s)\nabla\log\rho_s(X_s)\right)\right] ds + V^{-1/2}\sigma(X_s)\,dW_s, \tag{42}$$

where

$$\nabla f(X_s) := \nabla_x f(x)|_{x=X_s}, \qquad (\nabla\cdot D)_i := \sum_{j=1}^{S} \partial_{x_j} D_{ij}.$$

Thus the large-V reverse CRN has the familiar reverse-diffusion structure. At finite V, the corresponding reaction-resolved log-density difference is

$$f_r(n,s) = \log\frac{\rho_s(n-\mathbb{S}^r)}{\rho_s(n)}, \tag{43}$$

so $\widetilde{j}_r(n,s) = j_r(n-\mathbb{S}^r)e^{f_r(n,s)}$. Our sampler uses exact count kernels at $V=1$, so neither the forward nor the reverse diffusion approximation is required.

B ADDITIONAL EXPERIMENTAL RESULTS

B.1 TERMINAL-TIME CALIBRATION

Figure 2 illustrates the terminal-time criterion of Equation (18) on a one-dimensional example. The estimated excess-covariance signal reaches the empirical Poisson resolution floor at $T_O = 3.09$. The marginal reconstructed using terminal time $0.25T_O$ differs visibly from the data distribution. Those obtained using T_O and $1.5T_O$ are similar to the data and to each other. This example illustrates that the criterion selects a terminal time beyond which additional forward noising produces little visible change in the reconstructed marginal.

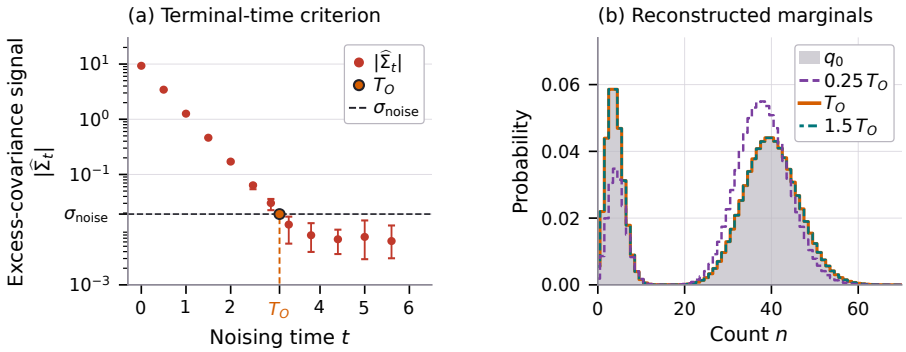

Figure 2: Terminal-time illustration. (a) The estimated excess-covariance signal reaches the empirical Poisson resolution floor at $T_O = 3.09$. (b) reconstructed marginals obtained from three terminal times of the forward process are compared with the data distribution in gray.

B.2 TWO-DIMENSIONAL COUNT MIXTURE

Construction. The eight-component mixture of Figure 3 is defined on the count lattice $\mathcal{G} = \{0, \dots, 63\}^2$. Each component is a discrete Gaussian kernel with

$$p_j(\mathbf{n}) \propto \exp\left(-\frac{\|\mathbf{n} - \boldsymbol{\mu}_j\|_2^2}{2\sigma^2}\right), \qquad \mathbf{n} \in \mathcal{G},$$

normalized over $\mathcal{G}$, with $\sigma = 1.8$ and

$$\boldsymbol{\mu}_j = (31.5, 31.5) + 20\big(\cos(j\pi/4), \sin(j\pi/4)\big), \qquad j = 0, \dots, 7.$$

The four axis components $\{c_0, c_2, c_4, c_6\}$ each have mixture mass 0.245, while the four diagonal components $\{c_1, c_3, c_5, c_7\}$ each have mixture mass 0.005. The target distribution assigns equal mass to c_1 and c_5, which together account for 1% of the training mixture mass. The pair $\{c_3, c_7\}$ has the same one-dimensional marginals as the target pair to numerical precision but a different joint distribution. This construction therefore tests both sampling from a rare target distribution and recovery of joint structure that is not determined by the marginals alone.

We compare three sampling configurations using the same frozen generator: marginal tilt, FK steering, and tilted FK steering. All configurations use $M = 10^4$ particles, output sample size $n_{\text{out}} = 10^3$, and $K = 32$ reverse steps. The two FK-steering configurations use $J = 16$ posterior draws per intermediate potential. FK steering and tilted FK steering resample when $\text{ESS}/M < 0.5$.

Distributional and genealogical metrics. For a generated distribution P and the target distribution P_R, let X and X' be independent samples from P, and let Y and Y' be independent samples from P_R, with all four variables mutually independent. The energy distance is

$$\mathcal{E}(P, P_R) = 2\mathbb{E}\|X - Y\| - \mathbb{E}\|X - X'\| - \mathbb{E}\|Y - Y'\|. \tag{44}$$

It measures discrepancy between joint distributions without requiring a kernel bandwidth and is related to a distance-induced MMD (Székely & Rizzo, 2013; Sejdinovic et al., 2013). The reported ED compares generated samples with an equally sized i.i.d. sample from the exact target distribution.

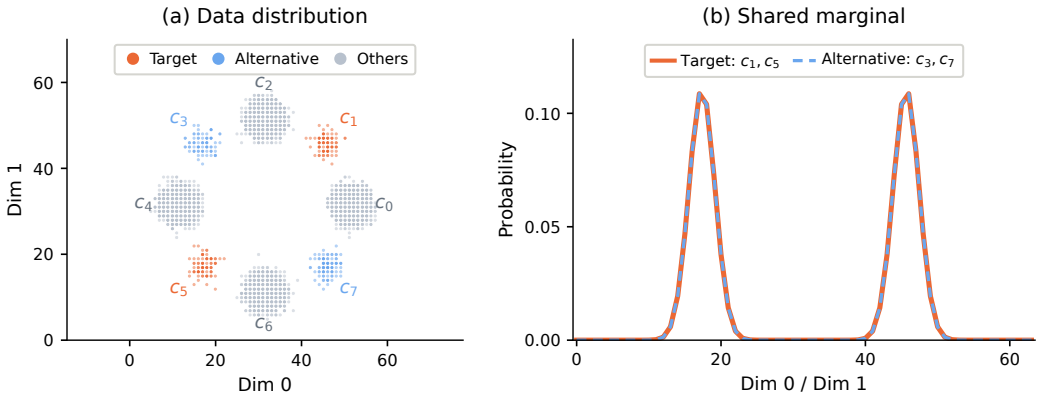

Figure 3: Two-dimensional count mixture and matched one-dimensional marginals. Components $\{c_1, c_5\}$ define the target distribution, and $\{c_3, c_7\}$ define an alternative distribution with identical coordinate marginals. The component pairs occupy different regions of the count space.

Mean marginal total variation (mTV) averages, across the two coordinates, the total-variation distance between the generated and target one-dimensional marginals. Because the target pair $\{c_1, c_5\}$ and the alternative pair $\{c_3, c_7\}$ have identical coordinate marginals, mTV alone cannot distinguish them. Anc. denotes genealogical diversity, defined as the number of distinct initial ancestors represented among the $n_{\text{out}} = 10^3$ generated samples divided by n_{out}. Higher Anc. therefore indicates that the final output sample descends from a larger set of initial particles. The i.i.d. reference provides a finite-sample calibration for ED and mTV.

Reverse-chain analysis. Figure 4 compares the three sampling configurations at matched reverse steps. Marginal tilt generates samples near all four rare components, without distinguishing the target pair $\{c_1, c_5\}$ from $\{c_3, c_7\}$, consistent with Proposition 3.1. FK steering uses an untilted proposal and progressively concentrates particles on the target pair, but its final output contains many descendants of relatively few initial particles. Tilted FK steering also concentrates on the target pair while retaining more distinct initial ancestors in the output. The intermediate snapshots and final genealogy therefore distinguish concentration on the target pair from genealogical diversity.

Figure 5 compares unweighted first-coordinate histograms of the endpoint particles before output sampling with those of the output samples. The endpoint particles from tilted FK steering are already closer to the target marginal before output sampling, whereas FK steering requires a larger redistribution of probability mass.

B.3 Single-cell experimental details

Experimental setup. The scRNA-seq dataset from the adult human heart is split into training, validation, and test sets in an 8:1:1 ratio, shared across methods. Checkpoints are selected by validation loss. Test cells serve as the real-data reference for distributional evaluation and are used to evaluate downstream classifiers. Generated-cell purity is scored from fixed CellTypist label histograms of each generated pool. Real-cell reference purity is computed over all held-out test cells in the target population after mapping CellTypist labels to the same fixed target label sets used for generated cells. All models use the same 2,000 highly variable genes. CRNDiff and MDLM cap model-input counts at 512, while scVI, scANVI, and CFGen use uncapped counts. Metric-specific preprocessing is shared across methods: W_1 uses raw counts, differential-expression analysis uses CP10K-normalized counts, and MMD^2, PCC, and downstream classification use log-CP10K values. CP10K normalizes each cell to total count 10^4 over the modeled genes; log-CP10K then applies $\log(1 + x)$. The endothelial, myeloid, and neuronal populations contain 80,463, 18,422, and 3,168 training cells, respectively. Generated and test samples are matched in size for distributional evaluation.

Generator. The posterior Q_θ is parameterized by a transformer with self-attention across gene tokens and a per-gene categorical count head (Table 3). Each token encodes the noisy count, gene

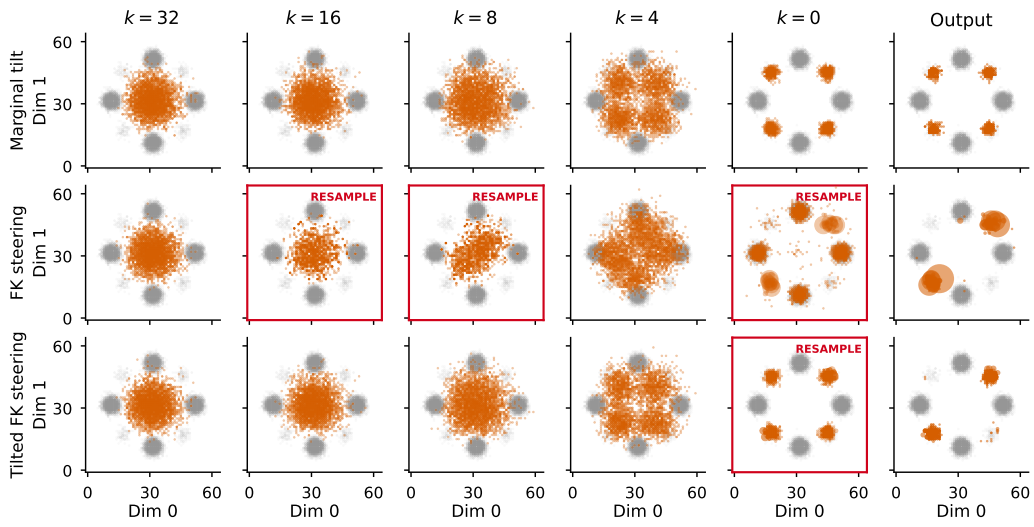

Figure 4: Reverse-chain snapshots for marginal tilt, FK steering, and tilted FK steering. Gray markers show training samples; orange markers show generated particles. In the intermediate columns, orange-marker area scales with normalized particle weight. In the output column, one marker per initial ancestor is placed at its last descendant in output-array order, with area proportional to that ancestor's number of output descendants. Red frames mark resampling events. FK steering also resamples at $k = 24$, which is not shown, and a final resampling is forced at $k = 0$ before output sampling.

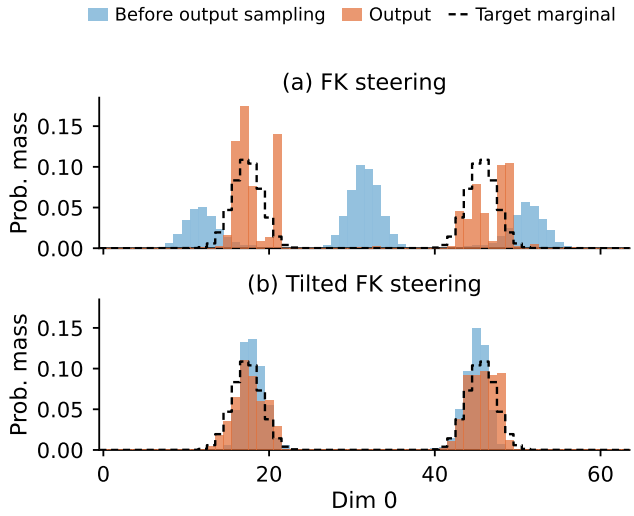

Figure 5: Unweighted first-coordinate histograms before and after output sampling for FK steering and tilted FK steering, with the exact target marginal as a reference. Before output sampling, histograms use the endpoint particles; Output denotes the final samples drawn using the terminal weights. The comparison measures marginal agreement; the target and alternative component pairs have identical one-dimensional marginals.

identity, and time. Training uses the cross-entropy objective of Section 2.1, with noisy samples drawn from the exact forward kernel in Equation (14). Three independent training seeds produce the models evaluated in Table 2. Conditioning comparisons and ablations use a single frozen model from this set as the reference CRNDiff generator.

Marginal-ratio estimation. For each frozen CRNDiff generator, per-gene histograms on $0, \ldots, 512$ are computed from all training cells of the target population and a fixed pool of $20{,}000$ unconditional samples from that generator. The estimated log-marginal ratio is $\log(\widehat{p}_i + 0.5/N_{\text{target}}) - \log(\widehat{q}_i + 0.5/20{,}000)$, clipped to $\pm\log(10^3)$, where $\widehat{p}_i$ and $\widehat{q}_i$ are normalized target and generator histograms, and N_{target} counts target training cells. Zero-frequency target bins retain the smoothed correction. These log-ratio estimates are reused across sampling runs. At each reverse step, we add τ times the estimated log-marginal ratio to the posterior logits and normalize the resulting posterior probabilities. We use $\tau = 0.4$ for endothelial cells and 0.2 for myeloid and neuronal cells.

Residual discriminator. For each target population, r_ϕ is a gradient-boosted tree classifier. Positive examples are real target cells from the training set; negative examples are cells generated by the marginally tilted proposal. Genes are ranked on log-CP10K fitting samples by

$$s_g = \frac{|\bar{z}_{P,g} - \bar{z}_{Q,g}|}{\sqrt{(v_{P,g} + v_{Q,g})/2 + \epsilon}},$$

Here P and Q denote target and proposal fitting samples; $\bar{z}_{P,g}, \bar{z}_{Q,g}$ and $v_{P,g}, v_{Q,g}$ are their means and empirical variances in log-CP10K space. The positive constant ϵ stabilizes the denominator. The top 256 genes are retained, except for the reference neuronal discriminator, which uses 192 genes. Training uses early stopping (Table 3). The odds $r_\phi/(1-r_\phi)$ provide a feature-based approximation to the residual density ratio in Equation (22). For FK steering, the discriminator is fitted using cells generated by the untilted proposal. The residual discriminator is distinct from the external CellTypist evaluation classifier.

Evaluation metrics. The metrics in Section 5.2 use the same preprocessing and scoring rules across methods. CellTypist labels are mapped to the three target cell types using fixed label sets. W_1 is computed on raw counts at matched sample sizes. MMD^2 uses the biased estimator in log-CP10K space, with a Gaussian kernel whose bandwidth is the median pairwise distance among test target cells. PCC uses the corresponding $2{,}000$-gene population mean vectors. Unconditional sliced-W_1 uses 512 random projections (Table 9) and is distinct from per-gene W_1.

Downstream classifier. Classifier hyperparameters for the downstream-utility evaluation are listed in Table 3; the training and test protocol is detailed in Appendix B.4.

Baselines. scVI, scANVI, and CFGen use the same data split and metric-specific preprocessing as CRNDiff. For conditional generation, scVI samples a standard-normal latent variable and decodes it with the target cell type fixed as a batch covariate. scANVI samples its latent variable from a prior conditioned on the target cell type, then decodes counts. CFGen fixes the target cell type during latent flow sampling and count decoding. For these three conditional baselines, library sizes are sampled with replacement from real test cells of the target population; Table 9 separately reports library-size sources for unconditional generation. Because the released MDLM implementation is designed for text data, we implement its published masked-diffusion objective using the same transformer backbone architecture as CRNDiff, representing counts as discrete tokens with an additional mask symbol. The conditional MDLM is initialized from an unconditional checkpoint and trained with a cell-type embedding added to the time embedding; sampling fixes this label throughout progressive unmasking.

Experimental repetitions. Table 2 reports three independently trained models per method, with one generation run from each trained model. Means and sample standard deviations are computed across these three model-level realizations. Tables 6–8 and Figure 9 use the frozen reference CRNDiff generator and repeat stochastic sampling three times. Differential-expression and downstream-classification results use five evaluation resamplings of fixed generated-cell pools; these repetitions quantify evaluation variation rather than generator-training variation, and their purity columns repeat the means from Table 2. For downstream classification, each evaluation resampling redraws training subsets without replacement from the fixed generated-cell and real training-cell pools, refits the logistic-regression classifier on those subsets, and evaluates it on the same held-out test set.

Table 3: Generator, sampling, residual-discriminator, and downstream-classifier settings used in the single-cell experiments.

Setting	Value
Generator	
Architecture	Transformer encoder; self-attention across genes
Input encoding	Scaled/log-count features, gene embeddings, time MLP
Layers / width / attention heads	3 / 128 / 4
Output head	Categorical over counts $0, \ldots, 512$
Optimizer	AdamW
Learning rate / weight decay	2×10^{-3} / 10^{-4}
Learning-rate schedule	500-step warm-up, then cosine decay
Batch size / training steps	384 / 20,000
EMA decay	0.999
Checkpoint selection	Lowest validation loss
Forward rates	$k_i^- = 1$, $\mu_i = k_i^+ = \widehat{\mathbb{E}}_{\mathcal{D}}[n_i]$, using training-set means
Terminal time	$T_O = 5.20$, determined by Equation (18)
Sampling	
Reverse steps K	32
Particles M	$2n_{\text{out}}$, where n_{out} is the output sample size
Endpoint draws per potential J	16
Resampling trigger	$\text{ESS}/M < 0.5$, minimum gap 2 steps
Terminal output	weighted draw from terminal particles
Tilt exponent τ	0.4 (endothelial), 0.2 (myeloid, neuronal)
Discriminator probability floor	10^{-8}, no upper clip on the odds
Residual discriminator (gradient-boosted trees)	
Input genes	256 selected genes (192 for the reference neuronal discriminator)
Learning rate / boosting rounds	0.06 / at most 180
Leaves / min. samples per leaf / L_2	31 / 30 / 1.0
Early stopping	15% internal validation split, patience 20
Downstream classifier (multinomial logistic regression)	
Penalty / regularization	L_2 / $C = 1$
Solver / maximum iterations	lbfgs / 2000
Class weights	balanced

B.4 Biological fidelity and downstream utility

We complement the conditional-fidelity metrics with analyses of low-dimensional localization, cell-type-specific differential-expression structure, and the utility of generated cells for downstream classification.

Fixed embeddings. Figure 6 compares representative generated cells from the five generative models in the fixed atlas PCA basis. Figure 7 shows CRNDiff-generated cells in the fixed atlas UMAP embedding. Both mappings place generated cells in the coordinate systems learned from real atlas cells, without refitting.

The PCA panels show that CRNDiff-generated neuronal cells concentrate near the real neuronal population, whereas generated cells from the baseline models show more dispersed projections. The CRNDiff UMAP panels provide a complementary view of target-population localization. These embeddings describe low-dimensional geometry and do not establish full count-space fidelity.

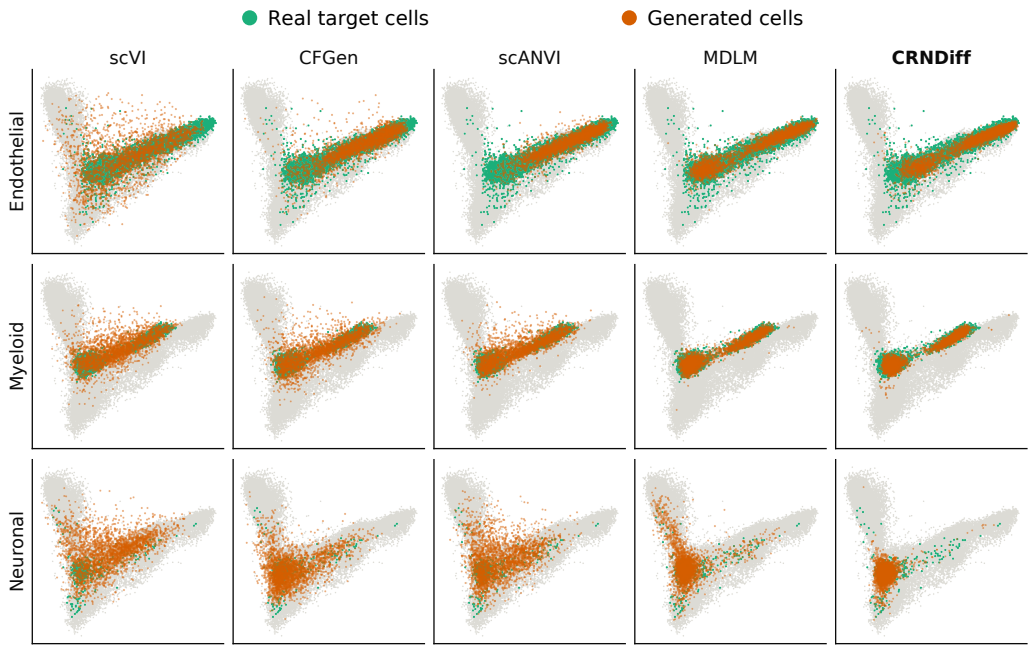

Figure 6: Representative generated cells in the fixed atlas PCA basis, with one column per generative model and one row per target population in decreasing training abundance. Gray points show all test cells, with test cells from the target population highlighted in green; orange points show generated cells. Generated cells are subsampled to 2,500 cells per panel; the PCA basis is applied without refitting.

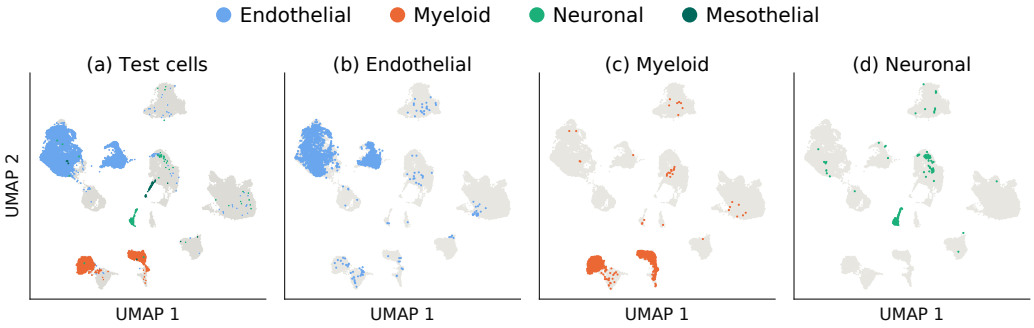

Figure 7: CRNDiff-generated cells in the fixed atlas UMAP embedding. (a) Test cells ($n = 45{,}145$); (b)–(d) 2,500 generated cells for endothelial, myeloid, and neuronal target populations. Mesothelial cells appear only in the test-set panel.

Differential-expression metrics. For each target cell type, we compare differential-expression rankings obtained from generated target cells and real test target cells. In both cases, the remaining cell types are represented by test cells. For each gene, log-fold change is the base-2 logarithm of the target-to-rest ratio of mean CP10K expression, with one added to each mean. Table 4 reports overlap among the 100 genes with the largest signed log-fold changes and Spearman correlation across all 2,000 gene rankings. Results are averaged over five evaluation resamplings of fixed generated-cell pools; cells are not regenerated for these repetitions, and the real reference and test split remain fixed. These metrics assess differential-expression agreement rather than cell-type purity.

Training on generated cells. An eleven-class multinomial logistic-regression cell-type classifier, with hyperparameters in Table 3, is trained on 2,500 cells per class in log-CP10K space and evaluated on all 45,145 held-out test cells. Generated cells replace real training cells for the three target classes; the remaining eight classes use real training cells. For each of the five evaluation resamplings, the

Table 4: Differential-expression agreement with test target cells. Marker overlap measures the fraction of shared genes among the top 100 ranked by signed target-versus-rest log-fold change; Spearman correlation uses all 2,000 genes. Marker overlap and Spearman correlation are averaged over five evaluation resamplings of fixed generated-cell pools; standard deviations are omitted. Purity values are reproduced from Table 2. Endo., Mye., and Neu. denote endothelial, myeloid, and neuronal cells.

	Purity ↑			Top-100 overlap ↑			Spearman ↑		
Model	Endo.	Mye.	Neu.	Endo.	Mye.	Neu.	Endo.	Mye.	Neu.
scVI	0.423	0.482	0.535	0.806	0.736	0.692	0.8923	0.9228	0.7234
CFGen	0.797	0.625	0.734	0.932	0.830	0.812	0.9812	0.9668	0.8735
scANVI	0.857	0.625	0.269	0.932	0.860	0.742	0.9772	0.9680	0.6930
MDLM	0.898	0.840	0.784	0.902	0.872	0.842	0.9767	0.9551	0.8003
CRNDiff	0.906	0.886	0.971	0.944	0.874	0.820	0.9852	0.9688	0.8355

training cells are redrawn without replacement from these fixed pools and the downstream classifier is refit; the generators are not retrained.

Table 5: Downstream classification on held-out real cells after replacing real training cells for three target classes with generated cells. Macro F_1 is the unweighted mean over all eleven atlas classes.

Training source	Macro F_1 ↑	Endo. F_1 ↑	Mye. F_1 ↑	Neu. F_1 ↑
scVI	0.915	0.947	0.879	0.800
CFGen	0.935	0.982	0.916	0.901
scANVI	0.895	0.970	0.940	0.457
MDLM	0.923	0.982	0.936	0.711
CRNDiff	0.941	0.980	0.943	0.915
Real cells	0.951	0.987	0.966	0.960

CRNDiff-generated training cells yield macro F_1 and target-class F_1 values close to the real-data reference. Differences among generative models are largest for neuronal F_1, while endothelial F_1 is similar for several methods. This comparison establishes utility for the evaluated classification task, rather than interchangeability between generated and real cells.

B.5 Sampler diagnostics and ablations

We next analyze the conditioning mechanism on a frozen CRNDiff generator, focusing on the contributions of marginal tilt and FK correction, genealogical diversity under resampling, and sensitivity to the marginal-tilt exponent.

Conditioning-component ablation. Table 6 and Figure 8 compare the conditioning components of Section 3.2 on one frozen CRNDiff generator. Marginal tilt adapts the posterior at every reverse step without particle resampling. FK steering uses the untilted proposal and a discriminator refitted against samples from that proposal. Tilted FK steering combines marginal tilt with particle weighting and resampling.

Marginal tilt yields dispersed PCA projections and low purity, particularly for the myeloid and neuronal target populations. These results indicate that adapting coordinate marginals is insufficient to recover the target population under the tested settings.

Compared with tilted FK steering, FK steering gives better endothelial distributional metrics, while the myeloid results are mixed. Tilted FK steering improves the neuronal distributional metrics (Table 6) and preserves greater genealogical diversity for every evaluated target population (Table 7), rather than uniformly improving every metric. Because the two procedures use separately fitted discriminators, their comparison also includes differences in density-ratio estimation.

Table 6: Ablation of marginal tilt and residual correction on the frozen reference CRNDiff generator, using the metrics of Table 2. The FK-steering discriminator is fitted against samples from the untilted proposal. Entries are means and standard deviations over three sampling runs.

Sampler	Purity ↑	W_1 ↓	MMD^2 ↓	PCC ↑
Endothelial *abundant*, 80,463 training cells				
Marginal tilt	0.682 ± 0.000	0.061 ± 0.002	0.0055 ± 0.0002	0.994 ± 0.000
FK steering	$\mathbf{0.909} \pm 0.007$	$\mathbf{0.043} \pm 0.009$	$\mathbf{0.0005} \pm 0.0003$	$\mathbf{0.999} \pm 0.000$
Tilted FK steering	0.906 ± 0.002	0.058 ± 0.002	0.0010 ± 0.0001	0.998 ± 0.000
Myeloid 18,422 training cells				
Marginal tilt	0.137 ± 0.004	0.297 ± 0.002	0.0504 ± 0.0011	0.764 ± 0.008
FK steering	$\mathbf{0.889} \pm 0.016$	$\mathbf{0.115} \pm 0.038$	0.0064 ± 0.0048	0.987 ± 0.007
Tilted FK steering	0.879 ± 0.019	0.131 ± 0.007	$\mathbf{0.0045} \pm 0.0029$	$\mathbf{0.991} \pm 0.008$
Neuronal *rare*, 3,168 training cells				
Marginal tilt	0.063 ± 0.005	0.131 ± 0.007	0.0588 ± 0.0003	0.629 ± 0.003
FK steering	$\mathbf{0.990} \pm 0.005$	0.062 ± 0.009	0.0179 ± 0.0105	0.923 ± 0.043
Tilted FK steering	0.984 ± 0.002	$\mathbf{0.053} \pm 0.001$	$\mathbf{0.0118} \pm 0.0003$	$\mathbf{0.954} \pm 0.001$

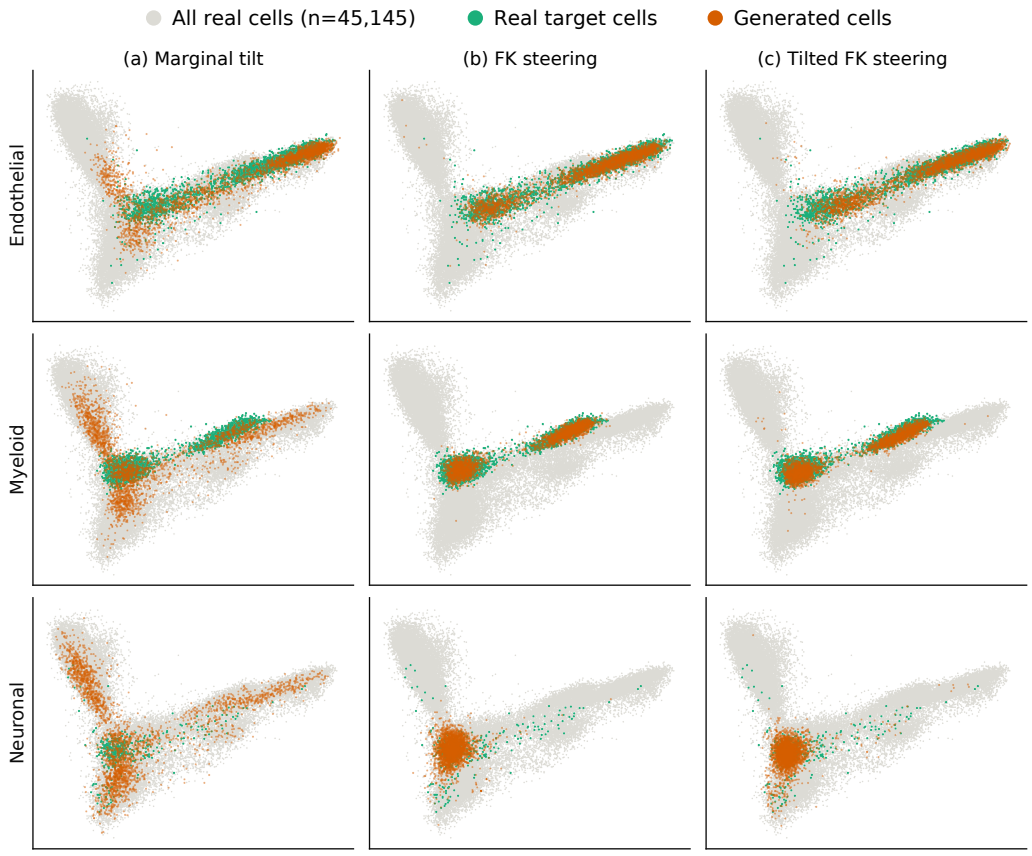

Figure 8: Three sampling methods across the three target populations in the fixed atlas PCA basis: (a) marginal tilt, (b) FK steering, and (c) tilted FK steering. Colors follow Figure 6. Each panel shows 2,500 generated cells subsampled from the combined outputs of three sampling runs.

Genealogical diversity. When particle weights are concentrated, resampling favors descendants of a few particles and can reduce the number of distinct initial ancestors represented in the output. Stochastic transitions can subsequently produce distinct count vectors from descendants of the same ancestor, so output uniqueness and genealogical diversity measure different properties (Jacob et al., 2015). Table 7 shows high output uniqueness alongside low genealogical diversity, particularly for neuronal FK steering. Distinct-ancestor counts diagnose shared particle histories but are not themselves estimates of effective sample size.

Table 7: Particle diagnostics on the frozen reference CRNDiff generator. Genealogical diversity (Anc.) is the number of distinct initial ancestors represented in the final output divided by n_{out}. Unique is the fraction of distinct output count vectors. $\text{ESS}_{\text{pre}}/M$ is the median normalized ESS at intermediate resampling events. Entries are means and standard deviations over three sampling runs.

Sampler	Resampling events	Anc. ↑	Unique ↑	$\text{ESS}_{\text{pre}}/M$ ↑
Endothelial *abundant*, 80,463 training cells				
FK steering	11.3 ± 0.6	0.300 ± 0.004	0.950 ± 0.023	0.379 ± 0.017
Tilted FK steering	6.0 ± 0.0	0.533 ± 0.005	0.959 ± 0.042	0.469 ± 0.010
Myeloid 18,422 training cells				
FK steering	11.3 ± 0.6	0.059 ± 0.005	0.920 ± 0.014	0.051 ± 0.007
Tilted FK steering	15.0 ± 0.0	0.112 ± 0.010	0.986 ± 0.024	0.147 ± 0.004
Neuronal *rare*, 3,168 training cells				
FK steering	14.3 ± 1.5	0.013 ± 0.003	0.936 ± 0.042	0.022 ± 0.009
Tilted FK steering	16.0 ± 0.0	0.045 ± 0.004	0.995 ± 0.002	0.053 ± 0.002

Despite more intermediate resampling events for myeloid and neuronal targets, tilted FK steering preserves greater final genealogical diversity (Table 7; Figure 9). Its higher median pre-resampling ESS is consistent with less concentrated weights at typical events. In these diagnostics, resampling-event counts and dashed markers refer to chain-internal resampling events and exclude the terminal weighted output selection described in Algorithm 2. Pre-resampling ESS is measured after the incremental-weight update and may fall below the resampling threshold.

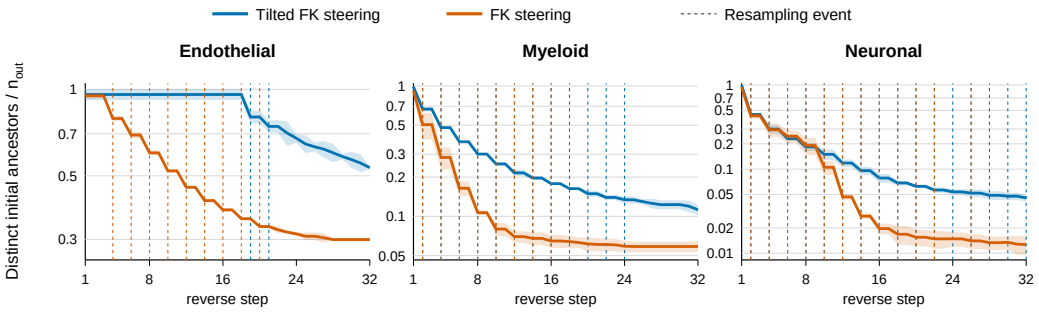

Figure 9: Number of distinct initial ancestors represented along the reverse chain, normalized by output sample size. On logarithmic y-axes, curves show means and standard-deviation bands over three sampling runs. Dashed lines mark chain-internal resampling events shared across the three sampling runs, excluding the terminal weighted output selection. Final values correspond to Table 7.

Tempering the marginal tilt. The exponent τ in Section 3.2 controls the strength of marginal proposal adaptation while leaving the target distribution unchanged. Table 8 evaluates the resulting purity and distributional fidelity on the frozen reference CRNDiff generator.

Table 8: Effect of the marginal tilt exponent, using the metrics of Table 2. The main experiments use $\tau = 0.4$ for endothelial cells and $\tau = 0.2$ for myeloid and neuronal cells. Entries are means and standard deviations over three sampling runs on the frozen reference CRNDiff generator.

τ	Purity ↑	W_1 ↓	MMD^2 ↓	PCC ↑
Endothelial *abundant*, 80,463 training cells				
0.2	0.906 ± 0.005	0.097 ± 0.002	0.0018 ± 0.0001	0.998 ± 0.000
0.4	0.906 ± 0.002	0.058 ± 0.002	0.0010 ± 0.0001	0.998 ± 0.000
1.0	0.933 ± 0.004	0.527 ± 0.015	0.1011 ± 0.0027	0.948 ± 0.001
Myeloid 18,422 training cells				
0.2	0.879 ± 0.019	0.131 ± 0.007	0.0045 ± 0.0029	0.991 ± 0.008
0.4	0.796 ± 0.006	0.189 ± 0.008	0.0084 ± 0.0010	0.964 ± 0.005
1.0	0.322 ± 0.036	0.839 ± 0.065	0.0923 ± 0.0111	0.914 ± 0.003
Neuronal *rare*, 3,168 training cells				
0.2	0.984 ± 0.002	0.053 ± 0.001	0.0118 ± 0.0003	0.954 ± 0.001
0.4	0.944 ± 0.024	0.059 ± 0.000	0.0106 ± 0.0011	0.958 ± 0.006
1.0	0.912 ± 0.045	0.399 ± 0.116	0.0407 ± 0.0092	0.882 ± 0.025

The tested tempered settings improve all three distributional metrics relative to $\tau = 1$ across the three target populations. Purity and distributional fidelity do not always improve together: endothelial purity is highest at $\tau = 1$, despite worse distributional agreement. The trade-off between $\tau = 0.2$ and $\tau = 0.4$ depends on the target population and metric, so these results do not identify a single exponent that optimizes every outcome.

B.6 Separating generator quality from conditioning

We separate generator quality from conditioning by first evaluating unconditional generation and then comparing conditioning strategies on a fixed CRNDiff generator.

Unconditional generation. For the 2,000 modeled genes, library size is the total count per cell, and detected-gene count is the number of nonzero entries. Their coefficients of variation (CVs) are sample standard deviations divided by means across cells. Median Fano is the median sample variance-to-mean ratio across genes with positive mean counts. These dispersion statistics use raw counts.

We evaluate unconditional fidelity and dispersion using the first 20,000 cells in each generated pool and an equally sized test sample (Table 9). The CRNDiff generator has the lowest sliced-W_1 and the closest detected-gene coefficient of variation to that of the test sample. Its median Fano factor and library-size coefficient of variation are lower than the test-sample values. scANVI most closely matches the median Fano factor, while CFGen most closely matches library-size variability.

CellTypist-assigned population fractions are computed over each full generated pool and assess composition separately from distributional fidelity. The CRNDiff generator most closely matches the test-sample endothelial and neuronal fractions, while MDLM is closest for myeloid cells. Neuronal cells constitute less than six percent of each generated pool; the number available for conditional selection therefore depends on both this fraction and the candidate budget.

Conditioning a frozen generator. Table 10 compares best-of-N selection, value-guided resampling (Li et al., 2024), and tilted FK steering on one frozen CRNDiff generator. Best-of-N selects n_{out} cells from N candidates ranked by target-class probabilities from a one-versus-rest gradient-boosted classifier trained on real training cells. A budget of $b\times$ uses bB_1 candidates, with $B_1 = 39{,}599$ for endothelial cells and 9,843 for each other target, calibrated from generator forward-pass counts for the reference sampler.

At each reverse step, value-guided resampling selects one of $J = 16$ candidates per particle with probability proportional to $\widehat{P}_\phi(y \mid \mathbf{n}_s, s)^\gamma$, where y is the target class and γ controls selection strength. The time-conditioned classifier is trained by cross-entropy on forward-noised real training cells. A separate baseline fine-tunes the same generator on labeled cells.

Table 9: Unconditional generation quality. For scVI, scANVI, and CFGen, total library sizes are drawn from the training data, whereas MDLM and CRNDiff generate library sizes. Fidelity and dispersion metrics use the first $20{,}000$ cells from each generated pool. The final columns report cell-type fractions over the full generated pool, assigned by the fixed external CellTypist heart model used for purity, with the test sample shown as a reference. Entries are means over three sampling runs per generator.

	Median	Libsize	Det. genes	sliced-W_1	CellTypist-assigned cell-type fraction		
Generator	Fano	CV	CV	↓	Endo.	Myeloid	Neuronal
Test sample	2.953	1.412	0.593	—	*0.233*	*0.055*	*0.021*
scVI	2.625	1.403	0.718	0.0758	0.194	0.040	0.056
scANVI	2.734	1.416	0.674	0.0345	0.221	0.042	0.029
CFGen	3.922	1.411	0.698	0.0359	0.212	0.044	0.040
MDLM	3.884	1.360	0.644	0.0723	0.254	0.062	0.027
CRNDiff	2.403	1.363	0.582	**0.0234**	0.229	0.041	0.025

For endothelial cells, intermediate selection or guidance settings improve individual metrics relative to tilted FK steering, while stronger selection increases purity and distributional distances. For myeloid cells, best-of-N at $8\times$ improves all three distributional metrics at similar purity; $16\times$ increases purity but worsens fidelity.

Rare neuronal target population. Tilted FK steering achieves higher purity and lower W_1 than every tested best-of-N and value-guided setting. Best-of-N retains a large purity gap even at the largest tested candidate-pool budget. Value-guided resampling narrows the gap, but increasing γ from 2 to 4 worsens all three distributional metrics. Some value-guided settings have lower MMD^2 and higher PCC than tilted FK steering. Fine-tuning further improves neuronal distributional fidelity, while tilted FK steering retains higher purity without retraining the generator.

Table 10: Conditioning strategies for the CRNDiff generator, using the metrics of Table 2. Entries are means over three sampling runs. The frozen-generator methods use the reference CRNDiff generator; † denotes its conditionally fine-tuned counterpart. Candidate-pool sizes and guidance strengths do not imply equal computational cost.

Conditioning	Purity ↑	W_1 ↓	MMD^2 ↓	PCC ↑
Endothelial *abundant*, 80,463 training cells; $n_{\text{out}} = 10{,}057$				
Best-of-N, 1×	0.830	0.046	0.0012	0.998
Best-of-N, 2×	0.949	0.051	0.0033	0.998
Best-of-N, 4×	0.958	0.078	0.0079	0.996
Value-guided, $\gamma = 1$	0.711	0.119	0.0117	0.992
Value-guided, $\gamma = 2$	0.889	0.029	0.0006	0.999
Value-guided, $\gamma = 4$	0.959	0.175	0.0169	0.994
Tilted FK steering	0.906	0.058	0.0010	0.998
Fine-tuned CRNDiff generator†	0.929	0.043	0.0004	0.999
Myeloid 18,422 training cells; $n_{\text{out}} = 2{,}500$				
Best-of-N, 1×	0.147	0.312	0.0577	0.690
Best-of-N, 2×	0.294	0.260	0.0343	0.878
Best-of-N, 4×	0.573	0.175	0.0104	0.976
Best-of-N, 8×	0.883	0.064	0.0012	0.996
Best-of-N, 16×	0.919	0.143	0.0093	0.988
Value-guided, $\gamma = 1$	0.510	0.240	0.0156	0.964
Value-guided, $\gamma = 2$	0.784	0.109	0.0019	0.992
Value-guided, $\gamma = 4$	0.882	0.155	0.0150	0.973
Tilted FK steering	0.879	0.131	0.0045	0.991
Fine-tuned CRNDiff generator†	0.897	0.106	0.0013	0.996
Neuronal *rare*, 3,168 training cells; $n_{\text{out}} = 2{,}500$				
Best-of-N, 1×	0.076	0.116	0.0612	0.610
Best-of-N, 2×	0.132	0.107	0.0542	0.664
Best-of-N, 4×	0.227	0.095	0.0457	0.731
Best-of-N, 8×	0.371	0.089	0.0366	0.795
Best-of-N, 16×	0.604	0.074	0.0230	0.883
Value-guided, $\gamma = 1$	0.577	0.094	0.0250	0.882
Value-guided, $\gamma = 2$	0.881	0.059	0.0077	0.965
Value-guided, $\gamma = 4$	0.968	0.088	0.0116	0.956
Tilted FK steering	0.984	0.053	0.0118	0.954
Fine-tuned CRNDiff generator†	0.961	0.032	0.0026	0.990