Markov-bridge generation of transition paths and its application to cell-fate choice

Guillaume Le Treut [1,*] Sarah Ancheta [1] Greg Huber [1,†] Henri Orland [2] and David Yllanes [1,3,4,5]

[1] *Chan Zuckerberg Biohub—SF, 499 Illinois Street, San Francisco, California 94158, USA*

[2] *Institut de Physique Théorique, CEA, CNRS, Université Paris-Saclay, 91191 Gif-sur-Yvette, France*

[3] *Instituto de Biocomputación y Física de Sistemas Complejos (BIFI), 50018 Zaragoza, Spain*

[4] *Fundación ARAID, Diputación General de Aragón, 50018 Zaragoza, Spain*

[5] *Zaragoza Scientific Center for Advanced Modeling (ZCAM), 50018 Zaragoza, Spain*

(Received 24 January 2024; accepted 18 October 2024; published 6 January 2025)

We present a method to sample Markov-chain trajectories constrained to both the initial and final conditions, which we term Markov bridges. The trajectories are conditioned to end in a specific state at a given time. We derive the master equation for Markov bridges, which exhibits the original transition rates scaled by a time-dependent factor. Trajectories can then be generated using a refined version of the Gillespie algorithm. We illustrate the benefits of our method by sampling trajectories in the Müller-Brown potential. This allows us to generate transition paths which would otherwise be obtained at a high computational cost with standard kinetic Monte Carlo methods because commitment to a transition path is essentially a rare event. We then apply our method to a single-cell RNA sequencing dataset from mouse pancreatic cells to investigate the cell differentiation pathways of endocrine-cell precursors. By sampling Markov bridges for a specific differentiation pathway, we obtain a time-resolved dynamics that can reveal features such as cell types which behave as bottlenecks. The ensemble of trajectories also gives information about the fluctuations around the most likely path. For example, we quantify the statistical weights of different branches in the differentiation pathway to alpha cells.

DOI: 10.1103/PhysRevResearch.7.013010

I. INTRODUCTION

In a world where computational capabilities are continually expanding, stochastic simulations have become indispensable tools, shedding light on myriad phenomena across physics, chemistry, and biology. These simulations are crucial for exploring processes such as the spontaneous folding and unfolding of proteins [1], allosteric transitions [2], glassy dynamics [3,4], the binding of molecules [5], and, more generally, for computing physical observables of complex systems [6]. However, sampling realizations of stochastic processes is fraught with challenges, particularly when it involves trajectories that explore rare states.

For instance, in protein folding, although the total folding time may be of the order of seconds, the time during which the system effectively jumps from an unfolded state to the folded state can be of the order of microseconds [7]. This most interesting part of the trajectory, during which the system effectively evolves from an unfolded state to the folded state, is called the transition path. While the typical time between folding-unfolding events is given by the Kramers time [8], which is exponential in the barrier height (also known as activation energy), the duration of the folding process, denoted as the transition-path time (TPT), is logarithmic in the barrier height [9,10]. Computationally, sampling such transitions is therefore very expensive since the majority of a simulation is spent waiting for the rare event of interest to occur. And yet, knowledge of the transition path ensemble is crucial to the exploration of the dynamics of the system during the TPT, including the monitoring of large conformational changes undergone by the system [11,12].

Langevin bridges have been proposed as an elegant solution to address such challenges [13–16]. They deal with the problem of sampling realizations of a reference Langevin dynamics (which is a diffusion process), starting in a certain known configuration [e.g., $\mathbf{x}(t_i) = \mathbf{x}_i$] and conditioned to end in a known final state (or final states sampled according a known distribution) in a given time t_f [e.g., $\mathbf{x}(t_f) = \mathbf{x}_f$]. In fact, a Langevin bridge is a Doob transform of the reference diffusion process. This transform was originally introduced by Doob to study Brownian bridges [17], and later adapted to deal with the conditioning of other stochastic processes (see [18] and references therein), including Schrödinger bridges [19–21]. The use of Doob's transform with diffusion processes has been receiving renewed attention in recent years among the machine-learning community since it is foundational to diffusion-based generative models [22–24].

While Langevin bridges are associated to a reference stochastic process of a continuous variable, namely, a diffusion process, in this work, we are interested in the sampling of stochastic bridges associated to a reference process of a

*Contact author: guillaume.letreut@gmail.com

†Contact author: gerghuber@gmail.com

discrete variable, namely, a discrete jump Markov process. We call them Markov bridges. Comparatively, the sampling of Markov bridges has received limited attention [25,26]. We derive Markov bridges from a master equation, which requires that transition rates between discrete states be given. We then show how Markov bridges can be efficiently sampled with kinetic Monte Carlo. The main difficulty in the practical implementation of our approach resides in the calculation of the transition probability $\mathbb{P}(\alpha_\mathrm{f}, t_\mathrm{f}|\alpha, t)$ from a state α towards the final state α_f over the duration $t_\mathrm{f} - t$, which requires the evaluation of a matrix exponential. Yet we will show through several examples that this problem can be overcome. In particular, in the context of developmental biology, we will use Markov bridges to investigate the dynamics of cell differentiation.

This article is organized as follows. In Sec. II, we introduce our framework and we derive a modified Gillespie algorithm to sample Markov bridges. In Sec. III, we show several applications of our method. Specifically, in Sec. III A, we validate our method by applying it to the case of diffusion on the one-dimensional (1D) lattice, for which analytical results are available. In Sec. III B, we illustrate the benefits of Markov bridges for rare-event statistics by sampling transition paths on a widely used benchmark potential. Finally, in Sec. III C, we apply our method to the investigation of cell-fate dynamics by sampling developmental trajectories for the differentiation of pancreatic endocrine-cell precursors. We conclude by discussing the strengths and limitations of our approach, and we put our work in context with respect to the existing body of literature in Sec. IV. The appendices provide additional details on the numerical algorithms, some analytical computations and clarifications on the statistical treatment of trajectories and of the transition-path time, and the procedure to compute transition rates in the application of our method to cell-fate dynamics.

II. MODEL

A. Master equation of a Markov bridge

Let us consider a system with N discrete states, namely, $[\![1, N]\!] = \{1, 2, \ldots, N\}$. For any pair of states $(\alpha, \beta) \in [\![1, N]\!]^2$, we introduce the transition rate $W_{\alpha\beta}$, such that the probability to jump from state β to state α during an infinitesimal time interval is $\mathbb{P}(\beta, t \to \alpha, t + \mathrm{d}t) = W_{\alpha\beta}\mathrm{d}t$. The time evolution of the probability to be in state α at time t conditioned to the initial state α_i at $t = 0$, namely, $P_\alpha(t) = \mathbb{P}(\alpha, t \mid \alpha_\mathrm{i}, 0)$, is described by the master equation [27]

$$\frac{\mathrm{d}P_\alpha(t)}{\mathrm{d}t} = \sum_\beta [W_{\alpha\beta}P_\beta(t) - W_{\beta\alpha}P_\alpha(t)]. \tag{1}$$

Let us now introduce the probability to be in the final state α_f at time t_f conditioned to being in state α at time t, namely, $Q_\alpha(t) = \mathbb{P}(\alpha_\mathrm{f}, t_\mathrm{f} \mid \alpha, t)$. Because of the Markovian nature of Eq. (1), the following relation is satisfied: $\sum_\alpha P_\alpha(t)Q_\alpha(t) = \mathbb{P}(\alpha_\mathrm{f}, t_\mathrm{f}|\alpha_\mathrm{i}, 0)$. This is called the Chapman-Kolmogorov identity (CKI) [27]. Differentiating the CKI with respect to t yields zero on the right-hand side, and $\mathrm{d}P_\alpha/\mathrm{d}t$ can be eliminated using Eq. (1). The resulting sum over states α and β can be relabeled and factorized to give the backward master equation,

$$\frac{\mathrm{d}Q_\alpha(t)}{\mathrm{d}t} = \sum_\beta W_{\beta\alpha}[Q_\alpha(t) - Q_\beta(t)]. \tag{2}$$

Both Eqs. (1) and (2) can be solved, and the formal solutions are given by

$$\begin{aligned} P(t) &= \mathrm{e}^{t\mathbb{W}}P(0), \\ Q(t) &= \mathrm{e}^{(t_\mathrm{f}-t)\mathbb{W}^T}Q(t_\mathrm{f}), \end{aligned} \tag{3}$$

where we introduced the $\mathbb{W}$ matrix [27] with entries,

$$\mathbb{W}_{\alpha\beta} = W_{\alpha\beta} - \delta_{\alpha\beta}\sum_\gamma W_{\gamma\alpha}. \tag{4}$$

B. Bridge master equation

We aim to generate bridges between states $(\alpha_\mathrm{i}, 0)$ and $(\alpha_\mathrm{f}, t_\mathrm{f})$. We therefore express the conditional probability to be in state (α, t) given the aforementioned initial and final states,

$$\begin{aligned} R_\alpha(t) &= \frac{\mathbb{P}(\alpha_\mathrm{f}, t_\mathrm{f} \mid \alpha, t)\mathbb{P}(\alpha, t \mid \alpha_\mathrm{i}, 0)}{\mathbb{P}(\alpha_\mathrm{f}, t_\mathrm{f} \mid \alpha_\mathrm{i}, 0)} \\ &= \frac{1}{Z}P_\alpha(t)Q_\alpha(t), \end{aligned} \tag{5}$$

where Z is a normalization factor ensuring that $\sum_\alpha R_\alpha(t) = 1$. Differentiating the previous equation, and using Eqs. (1) and (2), we obtain the master equation satisfied by $R_\alpha(t)$,

$$\begin{aligned} \frac{\mathrm{d}R_\alpha(t)}{\mathrm{d}t} &= \sum_\beta [V_{\alpha\beta}(t)R_\beta(t) - V_{\beta\alpha}(t)R_\alpha(t)], \\ V_{\alpha\beta}(t) &= W_{\alpha\beta}\frac{Q_\alpha(t)}{Q_\beta(t)}. \end{aligned} \tag{6}$$

Equation (6) is very similar to the master equation in Eq. (1), except that the original transition rates are scaled by a time-dependent factor $Q_\alpha(t)/Q_\beta(t)$. When $t_\mathrm{f} - t \gg 1$, $Q_\alpha(t) \to P_{\alpha_\mathrm{f}}(\infty) = \pi_{\alpha_\mathrm{f}}$ (the stationary distribution), and we recover the nonbridge transition rates $V_{\alpha\beta}(t) \to W_{\alpha\beta}$. On the other hand, when $t \to t_\mathrm{f}$, we have $Q_\alpha(t) \to \delta_{\alpha\alpha_\mathrm{f}}$ and thus for all $\alpha \neq \alpha_\mathrm{f}$ the incoming transition rate vanishes, $V_{\alpha\alpha_\mathrm{f}}(t) \to 0$, while the transition rate towards the final state diverges, $V_{\alpha_\mathrm{f}\alpha}(t) \to +\infty$. In between the aforementioned limiting cases, the evaluation of the bridge transition rates requires knowledge of $Q(t)$.

The stochastic process defined through Eq. (6) is akin to a Doob h-transform [17,18] of the original stochastic process from Eq. (1) with the time-dependent function $h(t, \alpha) = Q_\alpha(t)$. The resulting stochastic process is called a driven process, and it can be shown that its path measure only differs from the path measure of the original stochastic process by a multiplicative constant, namely, $Z = \mathbb{P}(\alpha_\mathrm{f}, t_\mathrm{f}|\alpha_\mathrm{i}, 0)$. In other words, the relative statistics of the paths is preserved. When considering a stochastic process of a continuous state variable (namely, a diffusion process), instead of Eq. (6), the Doob h-transform results in the addition of a term proportional to $\nabla \ln h(x, t)$ to the drift, where $h(x, t) = \mathbb{P}(x_\mathrm{f}, t_\mathrm{f}|x, t)$ [13,18,21,23,28]. We also note that an expression similar to Eq. (6) was proposed earlier [25,26], although they instead considered that the driven process evolves from the final state

towards the initial state, i.e., they used effective backward rates (in time) instead of the effective forward rates $V_{\alpha\beta}(t)$.

C. Kinetic Monte Carlo implementation

To sample bridge trajectories according to Eq. (6), we can adapt the original Gillespie algorithm [29] to the case of a master equation with time-dependent transition rates [30]. Given a current state (α, t), the probability to leave this state in the infinitesimal time interval $[t, t + \mathrm{d}t]$ is

$$\Gamma_\alpha(t)\mathrm{d}t = \sum_{\beta\neq\alpha} V_{\beta\alpha}(t)\mathrm{d}t. \tag{7}$$

It is convenient to introduce a random variable T representing the dwell time in state α. We define the cumulative distribution function $F_t(\tau) = \mathbb{P}(T < \tau)$ and the survival probability $G_t(\tau) = 1 - F_t(\tau)$. The latter function satisfies the ordinary differential equation (ODE),

$$\frac{\mathrm{d}G_t}{\mathrm{d}\tau}(\tau) = -\Gamma_\alpha(t+\tau)G_t(\tau), \quad G_t(0) = 1. \tag{8}$$

The distribution of T is therefore determined,

$$F_t(\tau) = 1 - \exp\left[-\int_0^\tau \mathrm{d}s\, \Gamma_\alpha(t+s)\right]. \tag{9}$$

We can generate T through the operation $T = F_t^{-1}(U)$, where $U \equiv \mathcal{U}(0, 1)$. This can be shown by noting that $F_t(T)$ is a uniform random variable between 0 and 1 [e.g., through its cumulative distribution function $\mathbb{P}(F_t(T) < u) = \mathbb{P}(T < F_t^{-1}(u)) = u$]. In practice, F_t^{-1} is not tractable; therefore, we solve for τ given a realization u,

$$\int_0^\tau \mathrm{d}s\, \Gamma_\alpha(t+s) + \ln(1-u) = 0. \tag{10}$$

The left-hand side of Eq. (10) is strictly increasing with τ so its root can be easily found with $0 < \tau < t_\mathrm{f} - t$. Since $V_{\alpha_\mathrm{f}\alpha}(t)$ diverges as $t \to t_\mathrm{f}$, $\Gamma_\alpha(t)$ also diverges as $t \to t_\mathrm{f}$. This ensures that there is always a solution. Eventually, we can sample Markov bridges using a modified version of the Gillespie algorithm in which jump events are determined by solving Eq. (10) (see Appendix A and Algorithm 1 for a pseudocode implementation).

So far, we have omitted giving any explanation on how to compute the time-dependent transition rates $V_{\alpha\beta}(t)$, yet they are the defining parameters of the bridge master equation, given by Eq. (6). They require evaluating $Q(t)$, whose formal solution given in Eq. (3) relies on the evaluation of a matrix exponential, which can be computationally expensive for large N. In this work, we decided to use the eigenvalue decomposition of $\mathbb{W}$ to evaluate this matrix exponential. While costly to compute, the eigenvalue decomposition $\mathbb{W} = -U\,\mathrm{Diag}(\lambda_1, \ldots, \lambda_N)U^{-1}$ $(0 = \lambda_1 < \lambda_2 \leqslant \cdots \leqslant \lambda_N)$ is computed only once and subsequently used every time $Q(t)$ needs to be evaluated, as

$$Q_\alpha(t) = \sum_k U_{\alpha_\mathrm{f}k}(U^{-1})_{k\alpha}\mathrm{e}^{-\lambda_k(t_\mathrm{f}-t)}. \tag{11}$$

The Perron-Frobenius theorem ensures that $\lambda_1 = 0$ is a unique eigenvalue and that the associated eigenvector is non-negative: $\pi_\alpha = U_{\alpha 1} \geqslant 0$. The vector π is the stationary distribution of the Markov process represented in Eq. (1). In general, $\mathbf{W}$ is a nonsymmetric matrix with real non-negative entries, and the nonsymmetric eigenvalue problem must be solved. When detailed balance is satisfied, namely, $W_{\alpha\beta}\pi_\beta = W_{\beta\alpha}\pi_\alpha$, the matrix $\pi_\alpha^{-1/2}W_{\alpha\beta}\pi_\beta^{1/2}$ is symmetric and can therefore be efficiently diagonalized, along with $\mathbb{W}$. For large systems, the eigenvalue decomposition of $\mathbb{W}$ (and most other matrix decompositions of the form $A = UBU^{-1}$) might be difficult to obtain since the usual methods require $\mathrm{O}(N^3)$ computing time. For such systems, one can still sample Markov bridges by first evaluating $Q(t)$ at discrete-time points through numerical integration of Eq. (2) and then using a standard method to sample Markov chains (see Appendix A).

In this work, we have restrained ourselves to Markov bridges connecting one initial state $(\alpha_\mathrm{i}, 0)$ to one final state $(\alpha_\mathrm{f}, t_\mathrm{f})$. Yet, our method can easily be extended to final states sampled according to a probability distribution Ψ. For example, to generate bridges with equal probabilities to end in any of L competing final states $\{\alpha_\mathrm{f}^{(l)}\}_{l=1,\ldots L}$, Eq. (3) needs to be solved with

$$\begin{aligned} Q_\alpha(t_\mathrm{f}) &= \Psi_\alpha \\ &= \frac{1}{L}\sum_{l=1}^{L}\delta_{\alpha\alpha_\mathrm{f}^{(l)}}, \end{aligned} \tag{12}$$

and Eq. (11) must be modified to

$$Q_\alpha(t) = \sum_\beta\sum_k \Psi_\beta U_{\beta k}(U^{-1})_{k\alpha}\mathrm{e}^{-\lambda_k(t_\mathrm{f}-t)}. \tag{13}$$

III. RESULTS

We now proceed to apply our method to several examples. We start with diffusion on the 1D lattice. Although trivial, since $P(t)$, $Q(t)$, and $R(t)$ can be computed exactly, it allows us to validate our model by comparing the results of the simulations to the theory. Second, we apply our method to sample bridge trajectories to study the transition between two wells in the Müller-Brown potential, which is widely used in physical chemistry as a benchmark potential to investigate transition paths in thermally activated processes. Theoretical results are not available, yet this example will help us illustrate the benefits of our method to sample rare events. Third, we illustrate the interest of our method in studying cell fate. A longstanding challenge in developmental biology is to reconstruct the temporal sequence of cell differentiation (i.e., cell fate) to understand commitment to competing cell types and to understand how robust differentiation pathways are to fluctuations. In recent years, single-cell RNA sequencing (scRNA-seq) has emerged as a pivotal technique giving molecular information about cell differentiation via high-throughput quantification of RNA abundance within individual cells [31–35]. More recently, RNA velocity methods [36–38], which estimate the time derivative of RNA abundance, have become instrumental in analyzing the direction of dynamic developmental processes involving cell-fate decisions. We leverage recent computational tools [37,39,40] to compute effective transition rates between cell "microstates" from RNA velocities, and we

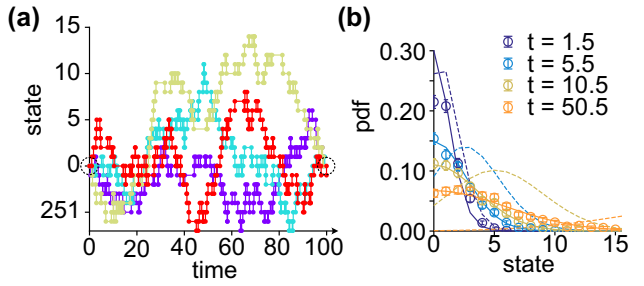

FIG. 1. Markov bridges for diffusion on the 1D lattice, with $N = 256$, $\alpha = 1$, $\beta = 0.5$, $n_\mathrm{i} = n_\mathrm{f} = 0$, and $t_\mathrm{f} = 100$. (a) Four selected trajectories. (b) The computed estimate for $R(t)$ (symbols) agrees with the theoretical expression (solid lines), but is markedly different from $P(t)$ that would be obtained with the nonbridge dynamics (dashed lines). We show the distributions at times $t = 1.5,\ 5.5,\ 10.5,\ 50.5$.

sample Markov bridges as putative single-cell developmental trajectories in pancreatic endocrinogenesis.

A. Validation with 1D diffusion

We first set out to validate our method by applying it to diffusion on the 1D lattice, for which $P(t)$ and $Q(t)$ can be expressed analytically. For this process, the master equation from Eq. (1) reads

$$\frac{\mathrm{d}P_n}{\mathrm{d}t} = \alpha P_{n-1} + \beta P_{n+1} - (\alpha+\beta)P_n, \quad \forall n \in [\![1, N]\!], \quad (14)$$

where $n \in [\![1, N]\!]$ denotes the state of a diffusing particle being at position $x_n = n\Delta x$ (Δx is the lattice site size), and α (β) is the forward (backward) transition rate. The diffusion process is biased when $\alpha \neq \beta$. We consider periodic boundary conditions, i.e., $x_{N+1} = x_1$. We generated 1024 trajectories using Algorithm 1, for which we used the following parameters: $N = 256$, $\alpha = 1$, $\beta = 0.5$ (hence the diffusion is biased), $n_\mathrm{i} = n_\mathrm{f} = 0$ and $t_\mathrm{f} = 100$. In Fig. 1(a), we show four of those trajectories.

Since $P(t)$ and $Q(t)$ can be expressed analytically, so can the bridge probability distribution $R(t)$ (see Appendix B). We can therefore compare this theoretical distribution to its estimate obtained by averaging over the sampled trajectories. For this purpose, we defined a time subdivision $t_j = j\Delta t$ for $j \in [\![0, M]\!]$, such that $t_M = t_\mathrm{f}$. The estimate of $R(t_j)$ from simulations is then

$$R_n^{\mathrm{sim}}(t_j) = \frac{1}{Z_j} \sum_s \sum_{(n_s, t_s)} \delta_{n_s,n} \mathbb{1}_{[t_j, t_{j+1}[}(t_s), \quad (15)$$

where the sum runs over all states (n_s, t_s) sampled in each bridge trajectory with index s, and Z_j is a normalization factor ensuring that $\sum_n R_n^{\mathrm{sim}}(t_j) = 1$. In Fig. 1(b), we computed R^{sim} at $t_j = 1.5,\ 5,\ 5,\ 10.5,\ 50.5$, with a subdivision size $\Delta t = 1$. We find that it agrees with the theoretical values of $R(t)$ (solid lines). For comparison, we also show the theoretical values of $P(t)$ (dashed lines), which are drastically different. Therefore, the method used to sample Markov bridges successfully recovers expected theoretical results.

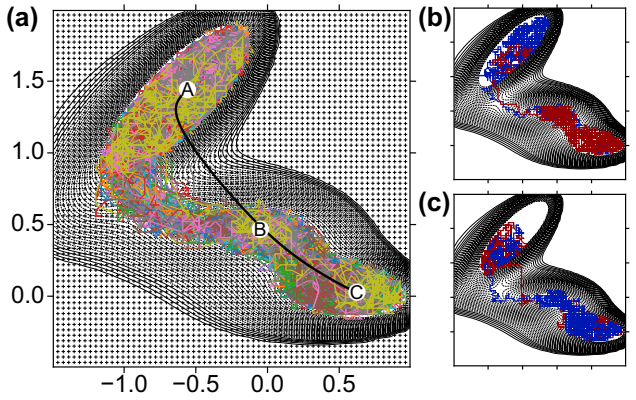

FIG. 2. (a) Markov bridges for the Müller-Brown potential, with $N = 4225$, $\alpha_\mathrm{i} = A$, $\alpha_\mathrm{f} = C$, and $t_\mathrm{f} = 1000$. The 2D grid of microstates is represented with black crosses (+). The level lines of the Müller-Brown potential are shown (thin solid black lines). We also give the mean trajectory (thick black line). We show a selection of trajectories going through (b) low- and (c) high-energy barriers. See also Movie 1 [44].

B. Application to the Müller-Brown potential

The Müller-Brown potential [41,42] is a standard benchmark potential used to check the validity of methods for generating transition paths. It is a two-dimensional (2D) potential given by

$$U(x, y) = \sum_{j=1}^{4} d_j \mathrm{e}^{a_j(x-x_j^0)^2 + b_j(x-x_j^0)(y-y_j^0) + c_j(y-y_j^0)^2}, \quad (16)$$

with

$$\begin{aligned} \mathbf{a} &= [-1, -1, -6.5, 0.7], \\ \mathbf{b} &= [0, 0, 11, 0.6], \\ \mathbf{c} &= [-10, -10, -6.5, 0.7], \\ \mathbf{d} &= [-200, -100, -170, 15], \\ \mathbf{x}^0 &= [1, 0, -0.5, -1], \\ \mathbf{y}^0 &= [0, 0.5, 1.5, 1]. \end{aligned} \quad (17)$$

It has three minima denoted by A $(-0.558, 1.442)$, B $(-0.05, 0.467)$, and C $(0.623, 0.028)$. As a second application of our method, we seek to sample Markov bridges from A to C. For that matter, we considered a 2D grid of points in $[-1.5, 1] \times [-0.5, 2]$, which defines $N = 4225$ states. We defined the transition rates as

$$W_{\alpha\beta} = \mathrm{e}^{-\frac{U_\alpha - U_\beta}{2k_\mathrm{B}T}}. \quad (18)$$

so that detailed balance is satisfied [43], namely, $W_{\alpha\beta}\pi_\beta = W_{\beta\alpha}\pi_\alpha$. This ensures that the stationary distribution is the Boltzmann distribution, $\pi_\alpha \propto \exp(-U_\alpha/k_\mathrm{B}T)$.

To accommodate the accuracy of double-precision computation arithmetic, we truncated the potential so that $\min U(x, y) = U_\mathrm{min} = -75\, k_\mathrm{B}T_0$ and $\max U(x, y) = U_\mathrm{max} = 15\, k_\mathrm{B}T_0$, and we chose $T/T_0 = 3.2$, so that the probability of the least-likely state is $\exp(-U_\mathrm{max}/k_\mathrm{B}T)/Z \approx 10^{-15}$. In Fig. 2, we show 128 Markov bridges generated with our method. We also represented the mean trajectory (details in

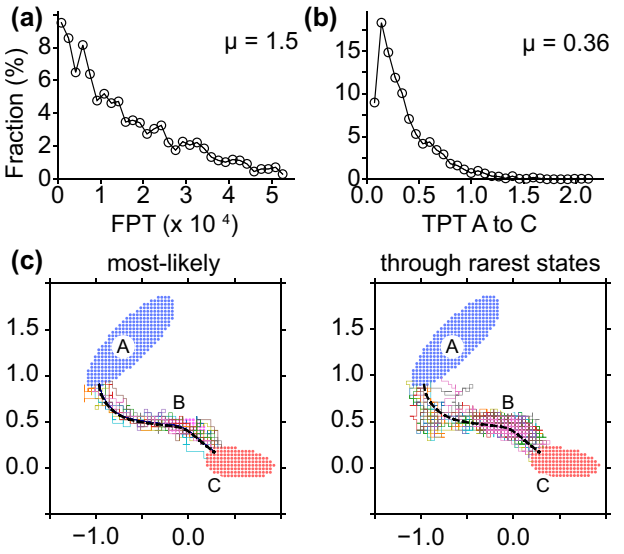

FIG. 3. Histograms for the (a) first-passage times and (b) transition-path times obtained by sampling 2048 trajectories with standard kinetic Monte Carlo [Eq. (1)] in the Müller-Brown potential, with $N = 4225$ and $\alpha_i = A$. (c) Transition paths for a selection of the most-likely trajectories (left) and trajectories visiting the rarest states (right). The filled circle indicates the bottom of each well $[U(x, y) = U_A, U_B, U_C]$. The thick dashed line denotes the mean transition path (see Appendix C). The FPT and TPT are given in units of Brownian time $\tau_B = \Delta x^{-2}$, where Δx is the lattice site size.

Appendix C). As can be expected, the most likely trajectories take paths going through low-energy barriers [Fig. 2(b)]. Yet some rare trajectories use paths which go through high-energy barriers [Fig. 2(c)]. While upon first inspection the mean trajectory is markedly different from the most likely trajectories, this is due to the inhomogeneous kinetics of the bridge trajectories. Indeed, each trajectory spends a varying amount of time sampling states in the potential well A before actually committing to transitioning towards states B and C. This asynchrony in the timing of the transition $A \to C$ causes the mean trajectory to deviate from the most likely one.

This example also emphasizes the difference between first-passage time (FPT) and TPT. To be more quantitative, we considered the nonbridge dynamics described in Eq. (1), and we sampled 2048 trajectories using the Gillespie algorithm. For each trajectory, we defined the last-passage time (LPT) as the time corresponding to the last occurrence of a state belonging to the bottom of well A [filled circles in Fig. 3(c)], and the FPT as the time corresponding to the first occurrence of a state belonging to well C. The transition-path time is then simply given by TPT = FPT − LPT. We checked that the distribution of the TPT agrees with previously reported theoretical results (see Appendix D). Figures 3(a) and 3(b) show that the mean TPT is about four orders of magnitude smaller than the mean FPT. Therefore, to study the transition $A \to C$ with standard kinetic Monte Carlo, one has to typically simulate trajectories for a duration of time of the order of 10^4 time units, whereas the portion of interest in which the transition $A \to C$ actually happens represents less than 0.01 % of the trajectory. Computationally, such an approach is very inefficient. This emphasizes the benefits of our method for rare-event statistics, which allows one to specifically sample the transition paths. That also means that the value for t_f should be chosen to be slightly larger than the mean TPT.

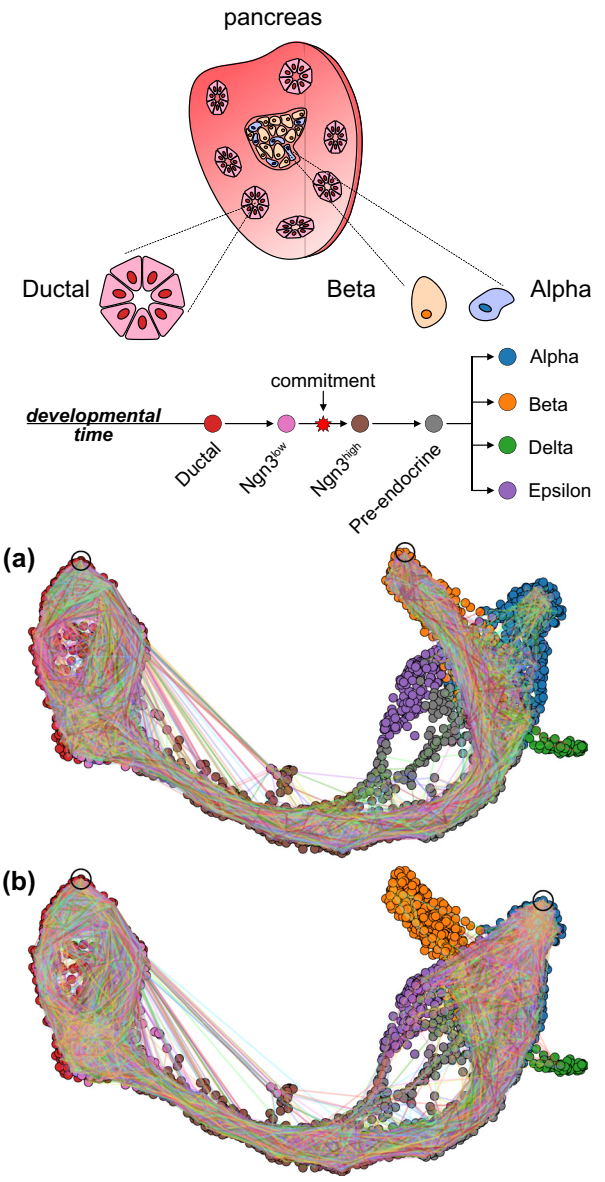

FIG. 4. Markov bridges for the pancreas dataset, with $N = 3696$, $\alpha_i = 1103$ (ductal state) and $t_f = 100$. We study the transition towards the final state (a) $\alpha_f = 2379$ (beta-cell type) and (b) $\alpha_f = 3437$ (alpha-cell type). We generated 512 bridge trajectories in each case. The state coordinates are the first two components of a UMAP analysis.

C. Inference of developmental trajectories

For a more concrete application of our method, we decided to use it to infer single-cell developmental trajectories in pancreatic endocrinogenesis. Development from endocrine progenitor cells (EPs) to fully differentiated cell types occurs in a stepwise manner, with each state having a characteristic genetic expression profile (Fig. 4). Starting from the

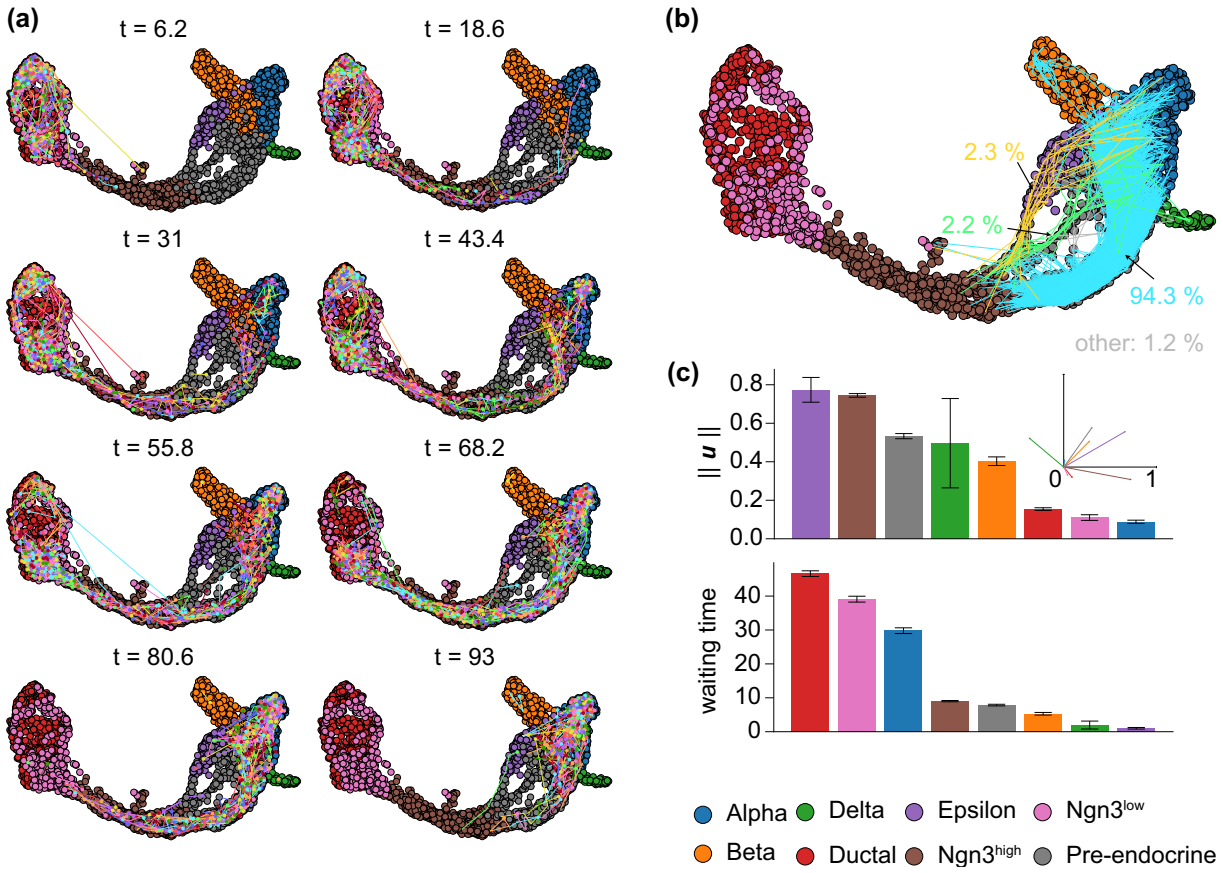

FIG. 5. Insights for the differentiation of pancreatic endocrine-cell precursors to alpha cells obtained from the bridge trajectories shown in Fig. 4. (a) Time lapse of the bridge dynamics. Current states are represented with filled circles and the tails show a few of the previous states. See also Movie 2 [44]. (b) Focus on the part of the bridge trajectories which goes from the $Ngn3^{high}$ state to the alpha-cell state. (c) Norm of the mean direction ($\|\mathbf{u}\|$) and waiting time (LPT − FPT) for each cell type. Inset: The mean direction vector $\mathbf{u}$ for the eight different cell types.

pool of bipotent ductal/EP cells physically located in the pancreatic epithelium, cells transition to a state expressing the transcription factor Neurogenin-3 at a low level (hereafter called $Ngn3^{low}$) and assume the identity of endocrine progenitors. These endocrine progenitors are still noncommitted to a specific cell fate and can, for example, revert to the ductal cell fate. Upon receiving adequate signals, endocrine progenitors can increase the expression levels of Ngn3 and transition to the $Ngn3^{high}$ state, thus assuming the identity of endocrine precursors, which are now committed to an endocrine cell fate. From the endocrine precursor state, cells transition to the pre-endocrine state. In this state, cells start to aggregate into proto-islet clusters (the future islets of Langerhans) physically located in the pancreatic mesenchyme. From the pre-endocrine state, cells differentiate into the different endocrine subtypes: glucagon-producing alpha cells, insulin-producing beta cells, somatostatin-producing delta cells, ghrelin-producing epsilon cells, and pancreatic polypeptide-producing (PP) cells. To analyze this developmental transition, we used a mouse pancreatic scRNA-seq dataset [45]. We used SCVELO [37] to compute transition rates between $N = 3696$ single cells (see Appendix E). With the purpose of defining a discrete jump Markov process, we will assimilate a single cell to a microstate in "cell-fate space," a high-dimensional space in which each microstate is a P-dimensional vector of mRNA transcript levels and where P is the number of genes resolved with scRNA-seq. Yet it can be visualized by applying dimension-reduction techniques to the gene expression vectors. The microstates' coordinates can be projected onto the first two components obtained by the uniform manifold approximation and projection (UMAP) method [46,47], as shown in Fig. 4. By contrast with the Müller-Brown potential application, the transition-rate matrix obtained with SCVELO did not satisfy detailed balance. Yet Markov bridges can still be generated in the same manner to sample putative single-cell developmental trajectories. Here, the time series of microstates represents the different genetic states a cell has to go through during the differentiation process. We generated 512 Markov bridges to study the transition from the ductal/EPs states to the beta-cell type [Fig. 4(a)] and from the ductal/EPs states to the alpha-cell type [Fig. 4(b)]. Qualitatively, we observe that many of the trajectories bound to the beta-cell fate also visit alpha-cell states [Fig. 4(a)]. By contrast, very few of the trajectories bound to the alpha-cell fate visit beta-cell states [Fig. 4(b)]. This suggests that the alpha-cell fate is specified before the beta-cell fate along the developmental path, in agreement with previous reports [37,45].

As a more quantitative analysis, we decided to focus on the differentiation into the alpha-cell type. The time lapse in Fig. 5(a) shows that trajectories spend a variable but significant amount of time in the ductal and $Ngn3^{low}$ states.

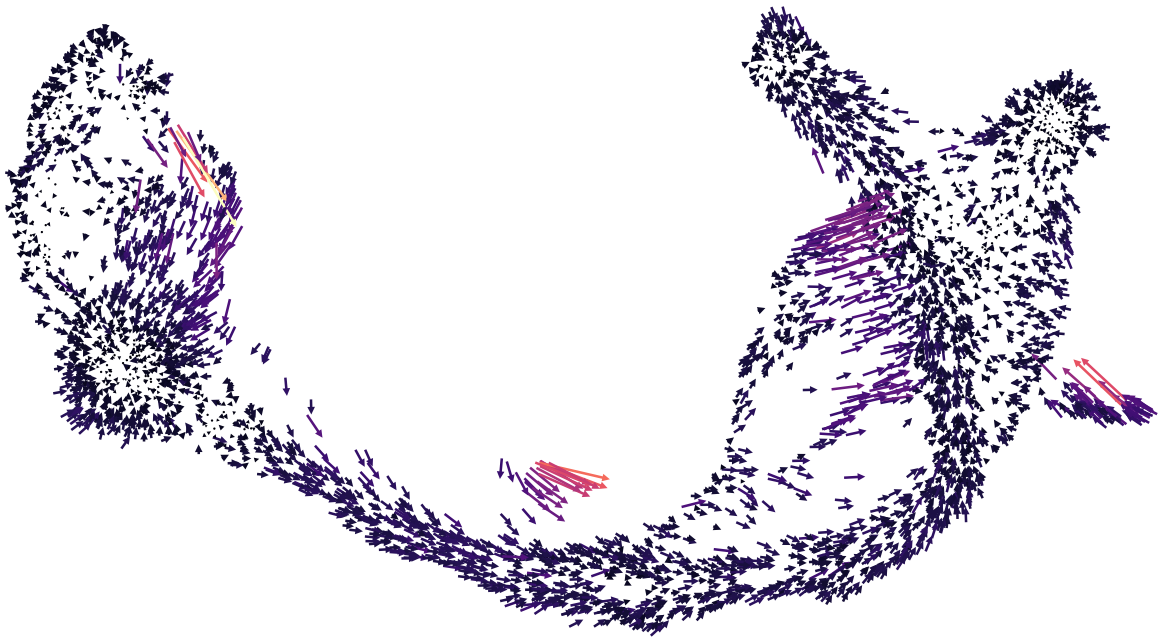

FIG. 6. Velocity field for the pancreas dataset of Fig. 4, computed as $\mathbf{v}_\alpha = \sum_{\beta\neq\alpha}(\mathbf{x}_\beta - \mathbf{x}_\alpha)W_{\beta\alpha}$. Both the lengths and the colors of the arrows scale linearly with $\|\mathbf{v}_\alpha\|$.

Quantitatively, this is confirmed by noting that the average waiting times in the ductal/Ngn3$^{\text{low}}$ states are the longest [Fig. 5(c)]. Here the waiting time was defined as the time of the last occurrence of a cell state minus its first occurrence, i.e., LPT − FPT. We note that the unit of time is set by the matrix of transition rates. Previous studies found that the expression levels of several cell-cycle inhibitors increased during commitment from the Ngn3$^{\text{low}}$ state to the Ngn3$^{\text{high}}$ state [45]. Analysis of RNA velocities then suggested that cycles and backflows visualized in cell-fate space could be associated to the cell-cycle process in the ductal/Ngn3$^{\text{low}}$ states [37]. Interestingly, the extensive sampling of microstates from the ductal/Ngn3$^{\text{low}}$ states by the bridge trajectories seems to corroborate these findings. The increased waiting time in these states suggests that the commitment to the Ngn3$^{\text{high}}$ state is the main bottleneck in this differentiation route.

We then computed the mean direction of motion by averaging the displacements (namely, $\mathbf{u}_t = \mathbf{v}_t/\|\mathbf{v}_t\|$, where $\mathbf{v}_t = \mathbf{x}_t - \mathbf{x}_{t-1}$) over trajectories in a given cell state. We found that the cell states exhibiting the most robust direction of motion are the Ngn3$^{\text{high}}$, pre-endocrine and epsilon-cell states [Fig. 5(c)], again corroborating RNA velocity analysis that suggested a strong directional flow towards the pre-endocrine state once committed to the Ngn3$^{\text{high}}$ state, and from the pre-endocrine state towards the endocrine subtypes [45].

Novel insights can be gained by using the fact that the sampling of the transition-path ensemble with Markov bridges is statistically exact. This property can be used to identify rare developmental trajectories, to compute effective transition rates between cell states, and to compute the statistical weights of competing developmental routes. For instance, in Fig. 5(b), we represented the part of the trajectories which is between the last occurrence of the Ngn3$^{\text{high}}$ state and the first occurrence of the alpha-cell type. Trajectories can be broadly divided into three groups, going through different branches in cell-fate space, in effect, defining a trifurcation. It appears that the vast majority of trajectories, about 94 %, uses a specific branch and the remainder is roughly equally split between the other two branches. Interestingly, the epsilon-cell type and the delta-cell type seem to only be accessible from these less likely routes. By comparison, RNA velocities cannot really be used to quantify the statistics of competing developmental routes. We speculate that our approach can be used to more broadly investigate branching occurring in cell-fate dynamics.

We note that velocities can be directly obtained from the matrix of transition rates, $\mathbf{W}$,

$$\mathbf{v}_\alpha = \sum_{\beta\neq\alpha}(\mathbf{x}_\beta - \mathbf{x}_\alpha)W_{\beta\alpha}. \tag{19}$$

As emphasized above, while the resulting velocity field (Fig. 6) is consistent with the results obtained by sampling Markov bridges, it is not, in and of itself, enough to get a temporal resolution of the dynamics or to analyze the role of fluctuations, such as the statistical weights of competing developmental routes. We also point out that the velocities shown in Eq. (19) are low-dimensional counterparts of RNA velocities. As such, their quantitative interpretation should be carried out with caution because they depend on the 2D cell-fate coordinates. By contrast, sampling Markov-bridge trajectories does not depend on the dimension-reduction method (namely, the UMAP) since the only input is the matrix of transition rates.

IV. DISCUSSION

A. Practical considerations

The evaluation of the bridge transition rates requires the evaluation of $Q(t)$, which, as already mentioned, involves a matrix exponential [Eq. (3)]. In this work, we have made

the choice to use the eigenvalue decomposition of $\mathbb{W}$ to evaluate $Q(t)$ as shown in Eq. (11). However, in general, for any matrix $A = UDU^{-1}$ ($D = \mathrm{Diag}(\lambda_1, \lambda_2, \ldots, \lambda_N)$), e^{tA} is not well approximated (numerically) by $U\mathrm{e}^{tD}U^{-1}$ when A is nonnormal (see, for example, [49] and sec. 11.3 of [50]). Unfortunately, $\mathbb{W}$ is seldom normal. Therefore, an improvement to our method would be to find more reliable ways to compute $Q(t)$. In practice, our method also requires choosing a duration of time for the transition, t_f. Although figuring out an appropriate value might be difficult in general, this can be done heuristically. Concretely, since the conditioned jump process in Eq. (6) converges towards the unconstrained jump process in Eq. (1) as $t_\mathrm{f} \to +\infty$, we can increase t_f until some dwelling is observed around both the initial and final states, which then means that t_f has reached a value slightly greater than the typical TPT.

B. Alternative methods to sample transition-path ensembles

Other methods can be used to sample transition paths and rare-event statistics, including forward flux sampling (FFS) [51], nonequilibrium umbrella sampling (NEUS) [52,53], and the weighted ensemble (WE) strategy [54]. FFS relies on the definition of (open) interfaces in phase space between initial and final states, while NEUS relies on a tiling of phase space into bins. The WE method also requires dividing phase space into bins. By assigning weights to trajectories and implementing a spawning/merging procedure to maintain a prescribed population of trajectories per bin, the WE method manages to repurpose most of the sampling effort towards the tail (or rarer events) of the distribution being sampled. In effect, these methods harvest segments of true trajectories in a piecewise fashion, which requires careful consideration for patching the segments into complete transition paths and for computing quantities of interest such as transition rates between initial and final states (e.g., tracking boundary crossings, reweighing bins), which generates most of the computational complexity associated with these methods. By contrast, the sampling of Markov bridges is agnostic to any partition of phase space. Its computational complexity lies in the evaluation of $Q(t)$, which effectively constrains the size of phase space that can be tackled. We also note that sampling Markov bridges generates uncorrelated trajectories while phase-space methods that spawn/merge trajectories, such as in the WE method and FFS, tend to produce correlated trajectories in each bin. Aguilar and co-workers proposed a method similar to ours to sample Markov bridges [25], although, as mentioned earlier, they used the backward generator instead of the forward generator of the reference process to express the bridge transition rates in Eq. (6). A more substantial difference between our approaches lies in the actual application of our method. In practice, they sampled Markov bridges using approximations for the bridge transition rates. Specifically, the probability distribution entering in the definition of the bridge transition rates [$Q(t)$ in Eq. (6)] was approximated with its time-independent stationary solution. By contrast, in this work, we used the eigenvalue decomposition of $\mathbb{W}$ to evaluate $Q(t)$, allowing us to compute the exact time-dependent bridge transition rates. In a more recent work, Aguilar and Gatto [26] numerically integrated the master equation of the reference process [Eq. (1)] to compute the bridge transition rates, somewhat similar to our Algorithm 2 (see Appendix A). However, their application was limited to discrete-time jump processes with a binomial tree structure (like random walks). By contrast, we sampled Markov bridges in continuous time and our implementation applies, in principle, to any transition-rate matrix.

C. Inference of developmental trajectories

Reconstructing a time-ordered sequence of states representing the progression of a single cell in a dynamic process such as cell differentiation constitutes a major challenge in developmental biology. While RNA velocity methods (e.g., VELOCYTO [36], SCVELO [37], CELLRANK [39]) enable one to visualize the "flow" of cell dynamics in cell-fate space, for instance, through the representation of streamlines of the RNA velocity vector field, they do not, strictly speaking, return cell trajectories. There are several trajectory inference methods that take advantage of the high-throughput transcriptomic data coming from scRNA-seq experiments [55]. Among those, one could distinguish methods inferring cell trajectories using only scRNA-seq data (e.g., SLINGSHOT [56]) from methods also using RNA velocity information (e.g., DIRECTED-PAGA [57], CELLRANK [39], VETRA [58], CELLPATH [48]). In most of these methods, trajectories are inferred using graph-theoretic methods such as shortest-path algorithms [48]. Furthermore, these methods typically return a small number of dominant trajectories. By contrast with such methods, we presented an approach to sample from an ensemble of trajectories. The dominance of specific transition paths and branching events can be inferred directly from the statistics of sampled trajectories. Crucially, our methods relies on estimates of cell-to-cell transition rates from RNA velocities. While the accuracy and robustness of such computational methods is under active investigation, it is not the ambition of this article to evaluate them or to take sides in the ongoing debate [59]. In fact, one may argue that snapshot techniques, such as scRNA-seq, which are destructive by nature and thus only reveal static snapshots of cellular states, do not capture all degrees of freedom in developmental trajectories, and that they only partially resolve the temporal evolution of developmental trajectories in cell-fate space. By contrast, live-cell imaging is a lower-throughput modality, yet unequivocally resolving the dynamics of a single cell, and it has been shown to identify developmental trajectories not identified with scRNA-seq data [60]. The combination of several modalities (e.g., morphological features with genetic expression data) holds tantalizing promises for a high-throughput, yet accurately time-resolved cell-fate space [61].

V. CONCLUSION

Given a system with discrete states, whose dynamics can be described by a master equation as in Eq. (1), we have presented a method to generate the subset of possible trajectories that go from one initial state α_i to a final state α_f in a time t_f. This method is especially valuable for the sampling of rare events, such as when the final state has a low probability to occur or when the initial state and the final state are separated by high-energy barriers (in the sense of Kramers' theory). Stochastic realizations of such processes

might exhibit a large FPT and would be sampled with very low efficiency using standard kinetic Monte Carlo methods (e.g., Gillespie algorithm) because most of the time would be spent waiting for a rare event to happen. A benefit of our method is therefore to discard the uninteresting waiting time. Although we mentioned some lines of improvement, our approach can be applied in its current state to study diverse phenomena, including cell-fate dynamics, and yield unique insights.

The code used to sample Markov bridges is implemented as a PYTHON package called MARKOV-BRIDGES available at Ref. [62].

ACKNOWLEDGMENTS

G.L.T. wishes to thank Loïc Royer for very useful discussions, and Alejandro Granados for introducing him to single-cell RNA sequencing and planting the seed for the application of Markov bridges to cell-fate dynamics. D.Y. was partially supported by Ministerio de Ciencia, Innovación y Universidades (Spain), Agencia Estatal de Investigación (AEI, Spain, 10.13039/501100011033), and European Regional Development Fund (ERDF, A way of making Europe) through Grant PID2022-136374NB-C21.

APPENDIX A: PSEUDOCODE IMPLEMENTATIONS

Algorithm 1 is the implementation we have used in this work to sample Markov bridges. However, as mentioned in the main text, for large systems it might not be adequate due to the computational cost of the eigenvalue decomposition of the $\mathbb{W}$ matrix. We therefore give Algorithm 2, which is an implementation not relying on a matrix decomposition of the form $e^{t\mathbb{W}} = Ue^{-tD}U^{-1}$, to evaluate $Q(t)$. In this approach, $Q(t)$ can be evaluated at specified discrete times $t_j = j\Delta t$ (with $t_M = t_f$) by integrating Eq. (2). We note here that matrix-vector multiplications can be performed very quickly with modern GPUs, which can be leveraged in an integration scheme. The time resolution Δt should be chosen so that $\Delta t < \Delta t^* = 1/\max_\alpha(\Gamma_\alpha)$ [Γ_α is the escape rate as defined in Eq. (7)].

ALGORITHM 1. Gillespie sampling for Markov bridges.

Require: $(\alpha_i, 0)$, rate matrix $\mathbf{W}$
  Set $(\alpha, t) \leftarrow (\alpha_i, 0)$
  print(α, t).
  while $t < t_f$ **do**
    Draw u_1 and u_2 as $\mathcal{U}(0, 1)$.
    $\tau \leftarrow$ Solve Eq. (10) with u_1. $Q(t)$ is computed through Eq. (11).
    Take γ so that $\sum_{\beta=1}^{\gamma-1} V_{\beta\alpha}(t+\tau) < u_2\Gamma_\alpha(t+\tau) \leqslant \sum_{\beta=1}^{\gamma} V_{\beta\alpha}(t+\tau)$.
    Update $(\alpha, t) \leftarrow (\gamma, t+\tau)$.
    print (α, t).
  end while

APPENDIX B: ANALYTICAL TREATMENT OF DIFFUSION ON THE 1D LATTICE

To solve Eq. (14), we introduce the discrete Fourier transform $\tilde{P}_k$, such that

$$\tilde{P}_k = \sum_{n=1}^{N} P_n e^{-i\frac{2\pi}{N}nk},$$
$$P_n = \frac{1}{N}\sum_{k=1}^{N} \tilde{P}_k e^{i\frac{2\pi}{N}nk}. \tag{B1}$$

Taking the time derivative of $\tilde{P}_k$ and using Eq. (14), we obtain the following system of ODE:

$$\frac{d\tilde{P}_k}{dt} = -\tilde{\omega}_k \tilde{P}_k, \quad \forall k \in [\![1, N]\!], \tag{B2}$$

where we introduced

$$\tilde{\omega}_k = -2i\sin\left(\frac{\pi k}{N}\right)\left(\beta e^{i\frac{\pi k}{N}} - \alpha e^{-i\frac{\pi k}{N}}\right),$$
$$\omega_n = (\alpha+\beta)\delta_n - \alpha\delta_{n-1} - \beta\delta_{n+1}, \tag{B3}$$

where δ_n is the Kronecker delta. In general, $\tilde{\omega}_k$ has a nonzero imaginary part. However, when $\alpha = \beta$, it reduces to $\tilde{\omega}_k = 4\alpha\sin^2(\frac{\pi k}{N})$.

Equation (B2) is straightforwardly integrated and yields

$$\tilde{P}_k(t) = e^{-\tilde{\omega}_k t} e^{-i\frac{2\pi}{N}n_i k}, \tag{B4}$$

where we used the initial condition $P_n(0) = \delta_{n,n_i}$.

For the backward master equation, we introduce $Q_n(t) = \mathbb{P}(n_f, t_f | n, t)$. Equation (2) reads

$$\frac{dQ_n}{dt} = -\alpha Q_{n+1} - \beta Q_{n-1} + (\alpha+\beta)Q_n, \quad \forall n \in [\![1, N]\!], \tag{B5}$$

and its discrete Fourier transform satisfies the ODE,

$$\frac{d\tilde{Q}_k}{dt} = \tilde{\omega}_k^* \tilde{Q}_k, \quad \forall k \in [\![1, N]\!], \tag{B6}$$

where the superscript $*$ denotes complex conjugation. Equation (B6) is again straightforwardly integrated and

ALGORITHM 2. Alternative sampling of Markov bridges.

Require: $(\alpha_i, 0)$, rate matrix $\mathbf{W}$, time resolution $\Delta t = t_f/M$
  Integrate Eq. (2) and evaluate $Q(t_j)$, $\forall j \in [\![0, M]\!]$.
  Set $(\alpha, t) \leftarrow (\alpha_i, 0)$
  print (α, t).
  while $t < t_f$ **do**
    Draw u_1 and u_2 as $\mathcal{U}(0, 1)$.
    $\tau \leftarrow j\Delta t$, where j is the smallest integer such that Eq. (10) is positive (with u_1).
    Take γ so that $\sum_{\beta=1}^{\gamma-1} V_{\beta\alpha}(t+\tau) < u_2\Gamma_\alpha(t+\tau) \leqslant \sum_{\beta=1}^{\gamma} V_{\beta\alpha}(t+\tau)$.
    Update $(\alpha, t) \leftarrow (\gamma, t+\tau)$.
    print (α, t).
  end while

yields

$$\tilde{Q}_k(t) = \mathrm{e}^{-\tilde{\omega}_k^*(t_\mathrm{f}-t)}\mathrm{e}^{-\mathrm{i}\frac{2\pi}{N}n_\mathrm{f}k}, \quad \forall k \in [\![1, N]\!], \tag{B7}$$

where we used the final condition $Q_n(t_\mathrm{f}) = \delta_{n,n_\mathrm{f}}$.

We can now compute the joint probability $R_n(t) \propto \mathbb{P}(n_\mathrm{f}, t_\mathrm{f}|n, t)\mathbb{P}(n, t|n_\mathrm{i}, 0)$,

$$R_n(t) = \frac{1}{Z}P_n(t)Q_n(t), \quad \forall n \in [\![1, N]\!],$$
$$\tilde{R}_k = \frac{\tilde{P} * \tilde{Q}_k}{\tilde{P} * \tilde{Q}_0}, \quad \forall k \in [\![1, N]\!], \tag{B8}$$

where $Z = \sum_n P_n(t)Q_n(t) = P_{n_\mathrm{f}}(t_\mathrm{f}) = Q_{n_\mathrm{i}}(0)$, and $*$ in this context is the convolution operator ($U * V_n = \sum_m U_{n-m}V_m$). Therefore, $R_n(t)$ can be exactly computed from Eqs. (B4), (B7), and (B8).

By differentiating $R_n(t)$ and making use of the forward [Eq. (14)] and backward [Eq. (B5)] master equations, we obtain the bridge master equation,

$$\frac{\mathrm{d}R_n}{\mathrm{d}t} = \alpha\left(\frac{Q_n}{Q_{n-1}}R_{n-1} - \frac{Q_{n+1}}{Q_n}R_n\right) + \beta\left(\frac{Q_n}{Q_{n+1}}R_{n+1} - \frac{Q_{n-1}}{Q_n}R_n\right), \quad \forall n \in [\![1, N]\!]. \tag{B9}$$

The bridge master equation (B9) has the same form as Eq. (6), where the forward and backward rates are now time dependent,

$$V_{n+1,n} = \alpha\frac{Q_{n+1}}{Q_n}, \quad V_{n+p,n} = 0 \quad \text{if} \quad p > 1,$$
$$V_{n-1,n} = \beta\frac{Q_{n-1}}{Q_n}, \quad V_{n-p,n} = 0 \quad \text{if} \quad p > 1. \tag{B10}$$

APPENDIX C: MEAN AND MOST LIKELY TRAJECTORIES

The mean trajectory shown in Fig. 2 was calculated by averaging trajectories after time registration. First, we defined a suitable subdivision of times $t_j = j\Delta t$. For a given trajectory $\{(t_k, \alpha_k)\}$, the state at time t_j, namely, α_j, is such that

$$\alpha_j := \alpha_{k_j} \quad \text{with} \quad t_{k_j} \leqslant t_j < t_{k_j+1}. \tag{C1}$$

The mean trajectory was obtained by averaging over the coordinates of the S trajectories,

$$\bar{\mathbf{x}}_j = \frac{1}{S}\sum_{s=1}^{S}\mathbf{x}_{\alpha_j^{(s)}}, \tag{C2}$$

where $\mathbf{x}_\alpha$ is the coordinates' vector associated with state α and $\alpha_j^{(s)}$ is the state of trajectory s at time t_j after registration as defined in Eq. (C1).

The mean transition path shown in Fig. 3 was obtained by setting the origin of time at the LPT in well A for each trajectory, and then averaging the coordinates as described above.

APPENDIX D: DISTRIBUTION OF THE TRANSITION-PATH TIME

We used theoretical results published in [10], investigating the TPT for the symmetric one-dimensional quartic potential.

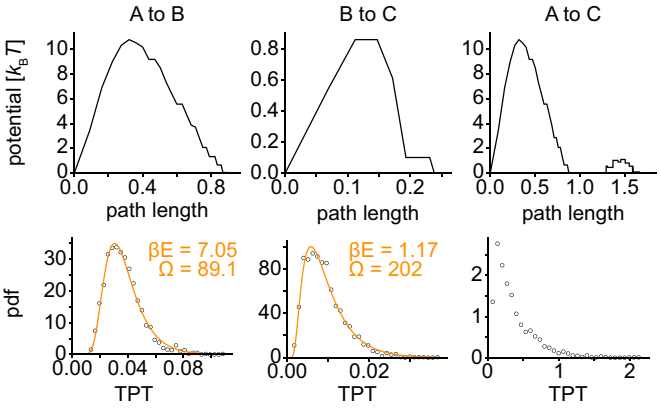

FIG. 7. Distribution of the TPT from A to B and from B to C for the Müller-Brown potential. The solid lines denote the fit to the theoretical prediction from [10].

For the overdamped Langevin dynamics, the probability distribution function of the TPT reads

$$p_\mathrm{TPT}(t) = \sqrt{\frac{\beta E}{\pi}}\frac{2\Omega}{1-\mathrm{erf}(\sqrt{\beta E})}\mathrm{e}^{-\Omega t}(1-\mathrm{e}^{-\Omega t})^{-\frac{3}{2}} \times (1+\mathrm{e}^{-\Omega t})^{-\frac{1}{2}}\exp\left(-\beta E\frac{1+\mathrm{e}^{-\Omega t}}{1-\mathrm{e}^{-\Omega t}}\right), \tag{D1}$$

where βE denotes the height of the potential barrier and Ω is proportional to the curvature at the top of the potential barrier. The previous equation makes use of the error function,

$$\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}}\int_0^x \mathrm{d}u\,\mathrm{e}^{-u^2}. \tag{D2}$$

The fits shown in Fig. 7 were performed by fitting βE and Ω using the Levenberg-Marquardt algorithm.

APPENDIX E: TRANSITION RATES FOR PANCREATIC CELLS

We used a mouse pancreatic development dataset [45] and the RNA velocity method SCVELO [37] to obtain a transition matrix. We normalized the raw counts data utilizing the default procedures from SCVELO. Each cell is size normalized to the median of total molecules, and 2000 top highly variable genes are chosen (with a minimum of 20 expressed counts for spliced and unspliced mRNA). The SCVELO pipeline calculates a nearest-neighbor graph using Euclidean distance utilizing 30 principal components on logarithmic spliced counts, and the mean and variance for each cell is calculated across its 30 nearest neighbors. Below we give a PYTHON implementation using the SCVELO package (`scv`), where `adata` is the annotated pancreas dataset from [45]:

```
scv.pp.filter_genes(adata, min_shared_counts=20)
scv.pp.normalize_per_cell(adata)
scv.pp.filter_genes_dispersion(adata, n_top_genes=2000)
scv.pp.log1p(adata)
scv.pp.moments(adata, n_pcs=30, n_neighbors=30).
```

We then compute the velocity and transition matrix using SCVELO-dynamical. We used the following functions to recover the splicing kinetics of each gene, and the cell-specific latent time, which are estimated using expectation maximization. We computed the transition matrix based on the velocity graph, using the SCVELO function with default parameters:

```
scv.tl.recover_dynamics(adata)
scv.tl.velocity(adata, mode="dynamical")
scv.tl.velocity_graph(adata)
scv.utils.get_transition_matrix(adata, 'velocity').
```


[1] D. E. Shaw, M. M. Deneroff, R. O. Dror, J. S. Kuskin, R. H. Larson, J. K. Salmon, C. Young, B. Batson, K. J. Bowers, J. C. Chao, M. P. Eastwood, J. Gagliardo, J. P. Grossman, C. R. Ho, D. J. Ierardi, I. Kolossváry, J. L. Klepeis, T. Layman, C. McLeavey, M. A. Moraes, R. Mueller, E. C. Priest, Y. Shan, J. Spengler, M. Theobald, B. Towles, and S. C. Wang, Anton, a special-purpose machine for molecular dynamics simulation, Commun. ACM **51**, 91 (2008).

[2] R. Elber, Simulations of allosteric transitions, Curr. Opin. Struct. Biol. **21**, 167 (2011).

[3] A. Cavagna, Supercooled liquids for pedestrians, Phys. Rep. **476**, 51 (2009).

[4] P. Charbonneau, E. Marinari, G. Parisi, F. Ricci-tersenghi, G. Sicuro, F. Zamponi, and M. Mezard, *Spin Glass Theory and Far Beyond: Replica Symmetry Breaking after 40 Years* (World Scientific, Singapore, 2023).

[5] Y. Shan, E. T. Kim, M. P. Eastwood, R. O. Dror, M. A. Seeliger, and D. E. Shaw, How does a drug molecule find its target binding site? J. Am. Chem. Soc. **133**, 9181 (2011).

[6] X. Wang, S. Ramírez-Hinestrosa, and D. Frenkel, Using molecular simulation to compute transport coefficients of molecular gases, J. Phys. Chem. B **124**, 7636 (2020).

[7] C. Hartmann, R. Banisch, M. Sarich, T. Badowski, and C. Schütte, Characterization of rare events in molecular dynamics, Entropy **16**, 350 (2014).

[8] P. Hänggi, P. Talkner, and M. Borkovec, Reaction-rate theory: Fifty years after Kramers, Rev. Mod. Phys. **62**, 251 (1990).

[9] B. W. Zhang, D. Jasnow, and D. M. Zuckerman, Transition-event durations in one-dimensional activated processes, J. Chem. Phys. **126**, 074504 (2007).

[10] M. Laleman, E. Carlon, and H. Orland, Transition path time distributions, J. Chem. Phys. **147**, 214103 (2017).

[11] K. Neupane, D. B. Ritchie, H. Yu, D. A. N. Foster, F. Wang, and M. T. Woodside, Transition path times for nucleic acid folding determined from energy-landscape analysis of single-molecule trajectories, Phys. Rev. Lett. **109**, 068102 (2012).

[12] H. S. Chung and W. A. Eaton, Protein folding transition path times from single molecule FRET, Curr. Opin. Struct. Biol. **48**, 30 (2018).

[13] H. Orland, Generating transition paths by Langevin bridges, J. Chem. Phys. **134**, 174114 (2011).

[14] M. Delarue, P. Koehl, and H. Orland, *Ab initio* sampling of transition paths by conditioned Langevin dynamics, J. Chem. Phys. **147**, 152703 (2017).

[15] R. Elber, D. E. Makarov, and H. Orland, *Molecular Kinetics in Condensed Phases: Theory, Simulation, and Analysis* (Wiley, New York, 2020).

[16] P. Koehl and H. Orland, Sampling constrained stochastic trajectories using Brownian bridges, J. Chem. Phys. **157**, 054105 (2022).

[17] J. L. Doob, Conditional Brownian motion and the boundary limits of harmonic functions, Bul. Soc. Math. France **79**, 431 (1951).

[18] R. Chetrite and H. Touchette, Nonequilibrium Markov processes conditioned on large deviations, Ann. Henri Poincaré **16**, 2005 (2015).

[19] E. Schrödinger, Über die Umkehrung der Naturgesetze, Sitzber, Preuss. Akad. Wiss. Phys.-Math. **K1**, 144 (1931).

[20] E. Schrödinger, Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique, Ann. Inst. Henri Poincare **2**, 269 (1932).

[21] B. Jamison, The Markov processes of Schrödinger, Z. Wahrscheinlichkeitstheor. Gebiete **32**, 323 (1975).

[22] J. Sohl-Dickstein, E. A. Weiss, N. Maheswaranathan, and S. Ganguli, Deep unsupervised learning using nonequilibrium thermodynamics, in *Proceedings of the 32nd International Conference on Machine Learning*, Vol. 37 (PMLR, Lille, 2015), pp. 2256–2265.

[23] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, Score-based generative modeling through stochastic differential equations, in *Proceedings of the 9th International Conference on Learning Representations* (2021).

[24] V. De Bortoli, J. Thornton, J. Heng, and A. Doucet, Diffusion Schrödinger bridge with applications to score-based generative modeling, in *Proceedings of the 35th Annual Conference on Neural Information Processing Systems* (2021).

[25] J. Aguilar, J. W. Baron, T. Galla, and R. Toral, Sampling rare trajectories using stochastic bridges, Phys. Rev. E **105**, 064138 (2022).

[26] J. Aguilar and R. Gatto, Unified perspective on exponential tilt and bridge algorithms for rare trajectories of discrete Markov processes, Phys. Rev. E **109**, 034113 (2024).

[27] N. G. Van Kampen, *Stochastic Processes in Physics and Chemistry* (Elsevier, Amsterdam, 1992), Vol. 1.

[28] J. Heng, V. De Bortoli, A. Doucet, and J. Thornton, Simulating diffusion bridges with score matching, 2023, https://hal.science/hal-03945487v1.

[29] D. T. Gillespie, A general method for numerically simulating the stochastic time evolution of coupled chemical reactions, J. Comput. Phys. **22**, 403 (1976).

[30] D. F. Anderson, A modified next reaction method for simulating chemical systems with time dependent propensities and delays, J. Chem. Phys. **127**, 214107 (2007).

[31] R. Sandberg, Entering the era of single-cell transcriptomics in biology and medicine, Nat. Methods **11**, 22 (2014).

[32] C. Gawad, W. Koh, and S. R. Quake, Single-cell genome sequencing: Current state of the science, Nat. Rev. Genet. **17**, 175 (2016).

[33] The Tabula Muris Consortium, Single-cell transcriptomics of 20 mouse organs creates a *Tabula Muris*, Nature **562**, 367 (2018).

[34] The Tabula Sapiens Consortium, The *Tabula Sapiens*: A multiple-organ, single-cell transcriptomic atlas of humans, Science **376**, eabl4896 (2022).

[35] M. Lange *et al.*, A multimodal zebrafish developmental atlas reveals the state-transition dynamics of late-vertebrate pluripotent axial progenitors, Cell **187**, 6742 (2024).

[36] G. La Manno, R. Soldatov, A. Zeisel, E. Braun, H. Hochgerner, V. Petukhov, K. Lidschreiber, M. E. Kastriti, P. Lönnerberg, A. Furlan, J. Fan, L. E. Borm, Z. Liu, D. van Bruggen, J. Guo, X. He, R. Barker, E. Sundström, G. Castelo-Branco, P. Cramer, I. Adameyko, S. Linnarsson, and P. V. Kharchenko, RNA velocity of single cells, Nature **560**, 494 (2018).

[37] V. Bergen, M. Lange, S. Peidli, F. A. Wolf, and F. J. Theis, Generalizing RNA velocity to transient cell states through dynamical modeling, Nat. Biotechnol. **38**, 1408 (2020).

[38] X. Qiu, Y. Zhang, J. D. Martin-Rufino, C. Weng, S. Hosseinzadeh, D. Yang, A. N. Pogson, M. Y. Hein, K. Hoi (Joseph) Min, L. Wang, E. I. Grody, M. J. Shurtleff, R. Yuan, S. Xu, Y. Ma, J. M. Replogle, E. S. Lander, S. Darmanis, I. Bahar, V. G. Sankaran, J. Xing, and J. S. Weissman, Mapping transcriptomic vector fields of single cells, Cell **185**, 690 (2022).

[39] M. Lange, V. Bergen, M. Klein, M. Setty, B. Reuter, M. Bakhti, H. Lickert, M. Ansari, J. Schniering, H. B. Schiller, D. Pe'er, and F. J. Theis, CellRank for directed single-cell fate mapping, Nat. Methods **19**, 159 (2022).

[40] P. Weiler, M. Lange, M. Klein, D. Pe'er, and F. Theis, CellRank 2: unified fate mapping in multiview single-cell data, Nat Methods **21**, 1196 (2024).

[41] K. Müller and L. D. Brown, Location of saddle points and minimum energy paths by a constrained simplex optimization procedure, Theor. Chim. Acta **53**, 75 (1979).

[42] K. Müller, Reaction Paths on Multidimensional Energy Hypersurfaces, Angew. Chem., Intl. Ed. Engl. **19**, 1 (1980).

[43] K. Binder and D. W. Heermann, *Monte Carlo Simulation in Statistical Physics* (Springer, Berlin, 2010).

[44] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevResearch.7.013010 for movies 1 and 2.

[45] A. Bastidas-Ponce, S. Tritschler, L. Dony, K. Scheibner, M. Tarquis-Medina, C. Salinno, S. Schirge, I. Burtscher, A. Böttcher, F. J. Theis, H. Lickert, and M. Bakhti, Comprehensive single cell mRNA profiling reveals a detailed roadmap for pancreatic endocrinogenesis, Development **146**, dev173849 (2019).

[46] L. McInnes, J. Healy, and J. Melville, UMAP: Uniform manifold approximation and projection for dimension reduction, arXiv:1802.03426.

[47] E. Becht, L. McInnes, J. Healy, C.-A. Dutertre, I. W. H. Kwok, L. G. Ng, F. Ginhoux, and E. W. Newell, Dimensionality reduction for visualizing single-cell data using UMAP, Nat. Biotechnol. **37**, 38 (2019).

[48] Z. Zhang and X. Zhang, Inference of high-resolution trajectories in single-cell RNA-seq data by using RNA velocity, Cell Rep. Methods **1**, 100095 (2021).

[49] C. Moler and C. Van Loan, Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later, SIAM Rev. **45**, 3 (2003).

[50] G. H. Golub and C. F. Van Loan, *Matrix Computations* (JHU Press, Baltimore, 2013).

[51] R. J. Allen, P. B. Warren, and P. R. ten Wolde, Sampling rare switching events in biochemical networks, Phys. Rev. Lett. **94**, 018104 (2005).

[52] A. Warmflash, P. Bhimalapuram, and A. R. Dinner, Umbrella sampling for nonequilibrium processes, J. Chem. Phys. **127**, 154112 (2007).

[53] A. Dickson, A. Warmflash, and A. R. Dinner, Separating forward and backward pathways in nonequilibrium umbrella sampling, J. Chem. Phys. **131**, 154104 (2009).

[54] R. M. Donovan, A. J. Sedgewick, J. R. Faeder, and D. M. Zuckerman, Efficient stochastic simulation of chemical kinetics networks using a weighted ensemble of trajectories, J. Chem. Phys. **139**, 115105 (2013).

[55] W. Saelens, R. Cannoodt, H. Todorov, and Y. Saeys, A comparison of single-cell trajectory inference methods, Nat. Biotechnol. **37**, 547 (2019).

[56] K. Street, D. Risso, R. B. Fletcher, D. Das, J. Ngai, N. Yosef, E. Purdom, and S. Dudoit, Slingshot: Cell lineage and pseudotime inference for single-cell transcriptomics, BMC Genomics **19**, 477 (2018).

[57] F. A. Wolf, F. K. Hamey, M. Plass, J. Solana, J. S. Dahlin, B. Göttgens, N. Rajewsky, L. Simon, and F. J. Theis, PAGA: Graph abstraction reconciles clustering with trajectory inference through a topology preserving map of single cells, Genome Biol. **20**, 59 (2019).

[58] G. Weng, J. Kim, and K. J. Won, VeTra: A tool for trajectory inference based on RNA velocity, Bioinformatics **37**, 3509 (2021).

[59] G. Gorin, M. Fang, T. Chari, and L. Pachter, RNA velocity unraveled, PLoS Comput. Biol. **18**, e1010492 (2022).

[60] W. Wang, D. Douglas, J. Zhang, S. Kumari, M. S. Enuameh, Y. Dai, C. T. Wallace, S. C. Watkins, W. Shu, and J. Xing, Live-cell imaging and analysis reveal cell phenotypic transition dynamics inherently missing in snapshot data, Sci. Adv. **6**, eaba9319 (2020).

[61] J. Copperman, I. C. Mclean, S. M. Gross, J. Singh, Y. H. Chang, D. M. Zuckerman, and L. M. Heiser, Single-cell morphodynamical trajectories enable prediction of gene expression accompanying cell state change, bioRxiv 2024.01.18.576248 (2024).

[62] https://github.com/czbiohub-sf/markov-bridges.