Towards Efficient Inference for Coupled Hidden Markov Models in Continuous Time and Discrete Space

Giosue Migliorini
Department of Statistics
University of California, Irvine

Padhraic Smyth
Department of Computer Science
University of California, Irvine

Abstract

Systems of interacting continuous time Markov chains are a powerful model class, but inference is typically intractable in high dimensional settings. Auxiliary information, such as noisy observations, is typically only available at discrete times, and incorporating it via a Doob's $h-$transform gives rise to an intractable posterior process that requires approximation. We introduce Hidden Interacting Particle Models (HIPMs), a model class parameterizing the generator of each Markov chain in the system. Our inference method involves estimating look-ahead functions (twist potentials) that anticipate future information, for which we introduce an efficient parameterization. We incorporate this approximation in a twisted Sequential Monte Carlo sampling scheme. We demonstrate the effectiveness of our approach on a challenging posterior inference task for a latent SIRS model on a graph, and benchmark different methods to approximate the twist function.

1 Introduction

Many real-world phenomena, from epidemics to wildfires, can be modeled as systems of interacting components evolving in continuous time, where the underlying dynamics are governed by discrete latent states. This approach builds upon concepts from continuous time hidden Markov models [1, 2] and extends them to spatially structured, high dimensional processes. Interacting particle systems (IPSs) [3, 4] provide a powerful mathematical framework for describing local propagation dynamics in discrete state spaces and continuous time, and are an important subset of the broader class of continuous time Markov chains (CTMCs). We formulate our goal as performing probabilistic inference on systems whose latent dynamics follow an IPS, given only incomplete or indirect information. We refer to this model class as Hidden Interacting Particle Models (HIPMs). Our model can be thought of as a discrete analogue to latent stochastic differential equations [5]. The core challenge within the HIPM framework lies in sampling from a smoothed path measure over latent trajectories. To address this problem, we make the following contributions:

- We propose a twisted sequential Monte Carlo (tSMC) scheme tailored to our model. Our approach involves learning a twist function [6], for which we extend the mass-covering loss proposed in [7]. A distinctive characteristic of our approach is that we do not require learning a separate proposal distribution as in [8**? ?**].
- We identify an invariance property that the twist function should satisfy, and use it to design an architecture with favorable inductive biases. We benchmark twist objectives and parameterizations on posterior path inference for a spatial SIRS model.

We provide an introduction to CTMCs and SMC, as well as references to related work, in A.

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: Frontiers in Probabilistic Inference: Sampling Meets Learning.

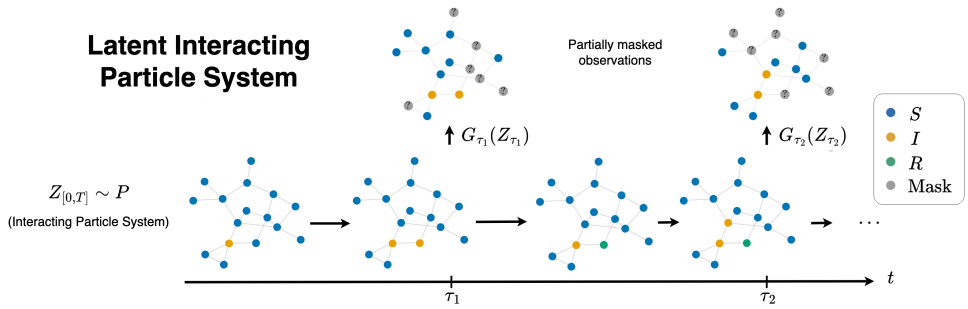

Figure 1: Example of an HIPM as a state space model. Here, the latent trajectory is sampled from a continuous time SIR model, and potentials are emission distributions of partially masked observations.

2 Background

2.1 Interacting Particle Systems

While CTMCs can be extended to high-dimensional systems, the size of the associated rate matrix increases exponentially with the number of dimensions, making exact inference intractable. We consider state spaces of the form $\mathcal{Z} = \mathcal{V}^d$, where $\mathcal{V} = \{1, \dots, V\}$ is a fixed *vocabulary*. In many real-world scenarios it is reasonable to assume that each transition only involves a jump in a single dimension. The rates for each dimension i can still depend on the global configuration Z_t, and this *spatial* dependence can be made explicit by considering a graph $\mathcal{G} = (\mathcal{I}, \mathcal{E})$, where $\mathcal{I}$ is a set of d indices and $\mathcal{E}$ is the set of edges. It is then natural to define the CTMC using local transition rates [9], corresponding to

$$r_t^i(v \mid z_t) = \lim_{\Delta_t \to 0} \frac{1}{\Delta_t} \mathbb{P}(Z_{t+\Delta_t}^i = v \mid Z_t^j = z_t^j : (i,j) \in \mathcal{E}), \quad v \neq z_t^i, \tag{1}$$

and $r_t^i(z_t^i \mid z_t) = -\sum_{v \in \mathcal{V}} r_t^i(v \mid z_t)$ for any $i \in \mathcal{I}$. The assumption that transitions occur for a single dimension at a time can be formalized by letting the rate matrix of the CTMC on the global state space $\mathcal{Z}$ be $R_t(z_{t+\Delta_t} \mid z_t) = \sum_{i \in \mathcal{I}} r_t^i(z_{t+\Delta_t}^i \mid z_t) \prod_{j \neq i} \delta_{z_t^j}(z_{t+\Delta_t}^j)$. CTMCs following these assumptions are known as IPSs on a finite state space and a finite graph [3, 4].

3 Twisted SMC for Hidden Interacting Particle Models

3.1 Hidden Interacting Particle Models

We are interested in scenarios where the underlying system dynamics follow an Interacting Particle System (IPS) $\{Z_t\}_{t \in [0,T]}$ with a prior path measure P, as described in Section 2.2. However, the inference task often involves conditioning this process on auxiliary information, which might arise from partial observations or specific desiderata on the process outcomes. We assume auxiliary information is present at discrete time points $\tau_1, \dots, \tau_K \in [0, T]$, and we represent it using non-negative *potential functions* $G_{\tau_k} : \mathcal{Z} \to \mathbb{R}_+$, for $k = 1, \dots, K$, which we assume to be tractable to evaluate pointwise. We let the target distribution of our **Hidden Interacting Particle Model (HIPM)** be the *smoothed path measure* $P^\star$, obtained by reweighting the original IPS path measure P according to the potentials:

$$P^\star(dZ_{[0,T]}) \propto \left(\prod_{k=1}^{K} G_{\tau_k}(Z_{\tau_k}) \right) P(dZ_{[0,T]}). \tag{2}$$

The construction in equation 2, for a general choice of prior dynamics, is commonly known as a Feynman-Kac (FK) path measure [10–12]. A crucial property of the HIPM target measure $P^\star$ is that the underlying stochastic process $\{Z_t^\star\}_{t \in [0,T]}$ it induces is, itself, an IPS. This insight stems from Doob's h-transform applied to the original IPS dynamics [13, 14]. To characterize the dynamics of this transformed process, we introduce the "look-ahead" function $h_t^\star : \mathcal{Z} \to \mathbb{R}_+$ for $t \in [0, T)$. At

times τ_k, $k \in [1:K]$ when a potential is observed, $h_t^\star$ undergoes discontinuities known as *reset conditions* [**?**]:

$$h_t^\star(z) \coloneqq \mathbb{E}_P\big[\prod_{k:\tau_k>t} G_{\tau_k}(Z_{\tau_k})\big|Z_t = z\big], \quad h_T^\star(z) = 1, \quad h_{\tau_k^-}^\star(z) \coloneqq \lim_{t\to\tau_k^-} h_t^\star(z) = G_{\tau_k}(z)h_{\tau_k}^\star(z). \tag{3}$$

The smoothed dynamics is then described by the following proposition:

Proposition 1 (Doob's h-transform of an IPS)**.** *Let P be the path measure of an IPS with initial distribution p_0 and local transition rates $r_t^i(v \mid z)$. The process $\{Z_t^\star\}_{t\in[0,T]}$ governed by the HIPM path measure $P^\star$ defined in equation 2 is also an IPS. Its initial distribution is $p_0^\star(z) \propto p_0(z)h_0^\star(z)$, and its time-dependent local transition rates $r_t^{\star,i}$ are given by:*

$$r_t^{\star,i}(v \mid z) \coloneqq \begin{cases} r_t^i(v \mid z)\frac{h_t^\star(z^{i\to v})}{h_t^\star(z)}, & v \neq z^i, \\ -\sum_{u\neq v} r_t^i(u \mid z)\frac{h_t^\star(z^{i\to u})}{h_t^\star(z)}, & v = z^i. \end{cases} \tag{4}$$

where $z^{i\to v}$ denotes the state z with its i-th dimension changed to state $v \in \mathcal{V}$, and $h_t^\star(z)$ is the look-ahead function defined in equation 3.

This result can be derived using conditioning properties for Markov processes [13, 14], and generalizes the posterior distribution used in reward-guided fine-tuning for discrete diffusion models [15–17]. We provide a complete proof of 1 in Appendix B.1. Simulating directly from these transformed rates $r_t^{\star,i}$ is generally intractable because computing the look-ahead function $h_t^\star(z)$ requires solving high-dimensional expectations over future trajectories under P. This computational challenge motivates the need for approximate inference techniques, such as the twisted Sequential Monte Carlo method we introduce next.

3.2 tSMC for HIPM

Classic design choices for SMC algorithms, such as the bootstrap particle filter (BPF) [18] notoriously display poor performance in continuous time problems with sparse potentials, as weights are uniform in between potential times, leading to degeneracy [19].

Twisted Targets. Twisted SMC (tSMC) addresses this limitation by modifying the intermediate target distributions of SMC [20, 6]. The core idea is to introduce a learnable function $h_t^\psi : \mathcal{Z} \to \mathbb{R}_+$, the *twist function*, approximating the look-ahead function $h_t^\star(z)$ from equation 3. The tSMC algorithm targets a sequence of *twisted* distributions P_t^ψ, defined by substituting $\prod_{k:\tau_k>t} G_{\tau_k}(z_{\tau_k})$ in equation 2 with $h_t^\psi(z_t)$. Following Proposition 1, this is also an IPS. We can then use the twist-induced proposal $\prod_{i\in\mathcal{I}} \delta_{z_s^i, z_t^i} + \Delta\, r_s^i(z_t^i \mid z_s)\frac{h_s^\psi(z_s^{i\to z_t^i})}{h_s^\psi(z_s)}$. The resampling procedure will correct errors induced by time discretization, and account for observed potentials.

Twist Learning. Prior work has explored approaches based on consistency inspired by optimal control theory [6, 8], and density ratio estimation [**? ?**]. Recently, [7] proposed to minimize the mass-covering forward KL divergence with respect to the true posterior in the context of autoregressive language models. We adapt this approach to our setting, and learn an amortized twist function by optimizing the following loss:

$$\mathcal{L}(\psi) = \mathbb{E}_{P_T^\star}\Big[-\sum_{i\in\mathcal{I}}\Big(\int_0^T r_t^{\psi,i}(Z_t^i \mid Z_t)dt + \sum_{t:Z_t^i\neq Z_{t^-}^i} \log r_t^{\psi,i}(Z_{t^-}^i \mid Z_t)\Big)\Big], \tag{5}$$

which can be approximated using tSMC weights, discussed in C.3. A similar approach has been proposed for learning the proposal distribution [21**?**]

Efficient Twist Parameterization. Parameterizing the twist-induced optimal proposal poses a notable challenge: at each sampling step, for a given starting state z_t, the proposal requires evaluating ratios of the twist for each value of $h_t^\psi(z^{i\to v})$, for $i \in \mathcal{I}, v \in \mathcal{V}$. We denote this mapping as $H_t^\psi : \mathcal{Z} \to \mathbb{R}_+^{d\times V}$, where $H_t^\psi(z)_{[i,z^i]} = h_t^\psi(z)$ for $i \in \mathcal{I}$. Note that, at the optimum, each element

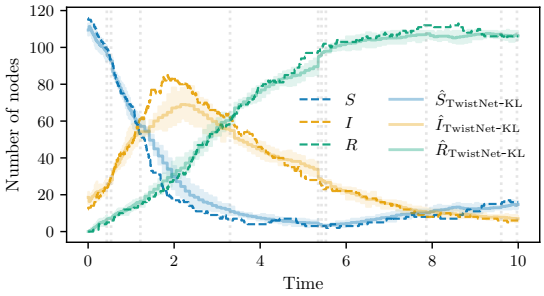

(a) Example of ground truth and generated latent trajectories, with counts of each state averaged over the graph at each timestep. The model is the SIRS IPS described in 4, with $d = 128$ nodes. Observation times are dashed gray lines; intervals are 95% quantiles over one SMC run with 25 particles.

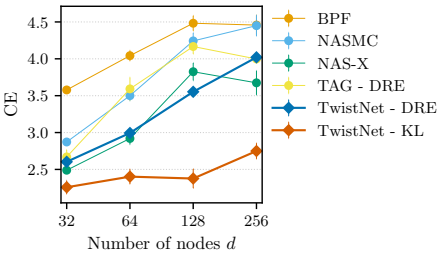

(b) Latent trajectory reconstruction, measured by NLL of ground truth trajectories with respect to the posterior approximations. Methods using TwistNet are highlighted by a thicker line. Error bars correspond to two standard errors.

in $H_t^\psi(z)$ should correspond to a conditional expectation where a single dimension of the input is swapped, bearing close resemblance to the numerator of the concrete score [22]. It follows that the learned mapping should satisfy the following invariance property:

$$h_t^\psi(z^{i\to v}) = H_t^\psi(z^{i\to u})_{[i,v]}, \quad u, v \in \mathcal{V},\ i \in \mathcal{I}. \tag{6}$$

In principle, this property can only be achieved by applying the twist function to each $z^{i\to v}$. However, it can be the case that we have access to highly informative contextual and positional information e_t, c_t. For instance, in SSM c_t could correspond to future observation states and times, i.e. $x_{\geq t}, \tau_{\geq t}$. We can leverage this by considering a *context head* $\Phi_t(e_t, c_t) \in \mathbb{R}^{V\times d\times m}$,

$$h_t^\psi(z) = \rho\left(\bigoplus_{i\in\mathcal{I}} \Phi_t(e_t, c_t)_{[i,z^i]}\right), \tag{7}$$

where $\rho : \mathbb{R}^m \to \mathbb{R}_+$ can be a learned function. When computing $\Psi_t(z)$, this formulation requires a single forward pass of the context head and the operation in equation 7 can be easily parallelized by pre-computing the sum $S_t(z) = \bigoplus_{i\in\mathcal{I}} \Phi_t(e_t, c_t)_{[i,z^i]}$ and letting $h_t^\psi(z^{i\to v}) = \rho\left(S_t(z) + \Phi_t(e_t, c_t)_{[i,v]} - \Phi_t(e_t, c_t)_{[i,z^i]}\right), \quad v \in \mathcal{V},\ i \in \mathcal{I}$. We refer to this parameterization as **TwistNet**. By [23, Thm. 9], this formulation characterizes an arbitrarily accurate function approximator, for suitable choices of Φ_t and ρ. Note that, while [23, Thm. 9] was introduced for permutation invariant functions, including positional information effectively bypasses this limitation. Another option is to use a discrete Taylor approximation, as we discuss in C.3.

4 Experiments

Latent SIRS model. We study posterior path inference in a latent spatial SIRS IPS on a graph with local state space $\mathcal{V} = \{S, I, R\}$ [4]. Observations arrive at $K = 10$ irregular *snapshots* $\tau_1 < \cdots < \tau_K$ and follow a state-space model: conditioned on the latent trajectory, the per-node emissions at each τ_k are independent across $i \in \mathcal{I}$. At each snapshot, node states are masked with 50% probability, and potentials are defined by the emission distribution of observations. We fix the parameters of the forward model P, and let the target of our inference procedure be the smoothed $P^\star$ from equation 2 for a prescribed test set of observations. We compare our approach to NASMC [21] and NAS-X [**?**]. For our approach, we experiment learning the twist using both our forward KL loss 5 and the DRE loss [**?**]. Experimental details are discussed in C.

5 Conclusion

We introduced an efficient posterior inference framework for systems of latent interacting CTMCs. Our method improves upon alternatives on a task of latent trajectory reconstruction, showing favorable scaling with dimensions. In future work, we plan on testing our approach on parameter learning tasks, and on including real world data such as WildFireSpreadTS [24], a dataset of wildfires trajectories.

References

[1] Leonard E Baum and Ted Petrie. Statistical inference for probabilistic functions of finite state markov chains. *The annals of mathematical statistics*, 37(6):1554–1563, 1966.

[2] Guy Leonard Kouemou and Dr Przemyslaw Dymarski. History and theoretical basics of hidden markov models. *Hidden Markov models, theory and applications*, 1, 2011.

[3] Thomas Milton Liggett. *Interacting particle systems*, volume 2. Springer, 1985.

[4] Nicolas Lanchier. *Stochastic interacting systems in life and social sciences*, volume 5. Walter de Gruyter GmbH & Co KG, 2024.

[5] Xuechen Li, Ting-Kam Leonard Wong, Ricky TQ Chen, and David Duvenaud. Scalable gradients for stochastic differential equations. In *International Conference on Artificial Intelligence and Statistics*, pages 3870–3882. PMLR, 2020.

[6] Jeremy Heng, Adrian N Bishop, George Deligiannidis, and Arnaud Doucet. Controlled sequential monte carlo. *The Annals of Statistics*, 48(5):2904–2929, 2020.

[7] Stephen Zhao, Rob Brekelmans, Alireza Makhzani, and Roger Grosse. Probabilistic Inference in Language Models via Twisted Sequential Monte Carlo, April 2024. URL `http://arxiv.org/abs/2404.17546`. arXiv:2404.17546 [cs].

[8] Dieterich Lawson, George Tucker, Christian A Naesseth, Chris Maddison, Ryan P Adams, and Yee Whye Teh. Twisted variational sequential monte carlo. In *Third workshop on Bayesian Deep Learning (NeurIPS)*, 2018.

[9] Nicolas Lanchier. *Stochastic modeling*. Springer, 2017.

[10] Nicolas Chopin and Omiros Papaspiliopoulos. *An Introduction to Sequential Monte Carlo*. Springer Series in Statistics. Springer International Publishing, Cham, 2020. ISBN 978-3-030-47844-5 978-3-030-47845-2. doi: 10.1007/978-3-030-47845-2. URL `https://link.springer.com/10.1007/978-3-030-47845-2`.

[11] Jianfeng Lu and Yuliang Wang. Guidance for twisted particle filter: a continuous-time perspective, September 2024. URL `http://arxiv.org/abs/2409.02399`. arXiv:2409.02399 [stat].

[12] Byoungwoo Park, Hyungi Lee, and Juho Lee. Amortized Control of Continuous State Space Feynman-Kac Model for Irregular Time Series, February 2025. URL `http://arxiv.org/abs/2410.05602`. arXiv:2410.05602 [stat].

[13] Marc Corstanje and Frank van der Meulen. Guided simulation of conditioned chemical reaction networks. *arXiv preprint arXiv:2312.04457*, 2023.

[14] Marc Corstanje, Frank van der Meulen, and Moritz Schauer. Conditioning continuous-time markov processes by guiding. *Stochastics*, 95(6):963–996, 2023.

[15] 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, October 2024. URL `http://arxiv.org/abs/2408.08252`. arXiv:2408.08252 [cs].

[16] Chenyu Wang, Masatoshi Uehara, Yichun He, Amy Wang, Tommaso Biancalani, Avantika Lal, Tommi Jaakkola, Sergey Levine, Hanchen Wang, and Aviv Regev. Fine-Tuning Discrete Diffusion Models via Reward Optimization with Applications to DNA and Protein Design, March 2025. URL `http://arxiv.org/abs/2410.13643`. arXiv:2410.13643 [cs].

[17] Cheuk Kit Lee, Paul Jeha, Jes Frellsen, Pietro Lio, Michael Samuel Albergo, and Francisco Vargas. Debiasing Guidance for Discrete Diffusion with Sequential Monte Carlo, February 2025. URL `http://arxiv.org/abs/2502.06079`. arXiv:2502.06079 [cs].

[18] Arnaud Doucet, Adam M Johansen, et al. A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of nonlinear filtering*, 12(656-704):3, 2009.

[19] Nicolas Chopin, Andras Fulop, Jeremy Heng, and Alexandre H. Thiery. Computational Doob's h-transforms for Online Filtering of Discretely Observed Diffusions, May 2023. URL http://arxiv.org/abs/2206.03369. arXiv:2206.03369 [stat].

[20] Pieralberto Guarniero, Adam M. Johansen, and Anthony Lee. The iterated auxiliary particle filter, June 2016. URL http://arxiv.org/abs/1511.06286. arXiv:1511.06286 [stat].

[21] Shixiang Shane Gu, Zoubin Ghahramani, and Richard E Turner. Neural adaptive sequential monte carlo. *Advances in neural information processing systems*, 28, 2015.

[22] Chenlin Meng, Kristy Choi, Jiaming Song, and Stefano Ermon. Concrete score matching: Generalized score matching for discrete data. *Advances in Neural Information Processing Systems*, 35:34532–34545, 2022.

[23] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan Salakhutdinov, and Alexander Smola. Deep Sets, April 2018. URL http://arxiv.org/abs/1703.06114. arXiv:1703.06114 [cs].

[24] Sebastian Gerard, Yu Zhao, and Josephine Sullivan. Wildfirespreadts: A dataset of multi-modal time series for wildfire spread prediction. *Advances in Neural Information Processing Systems*, 36:74515–74529, 2023.

[25] James R Norris. *Markov chains*. Cambridge university press, 1998.

[26] Pierre Del Moral and Spiridon Penev. *Stochastic Processes: From Applications to Theory*. Chapman and Hall/CRC, 2017.

[27] Christian A. Naesseth, Fredrik Lindsten, and Thomas B. Schön. Elements of Sequential Monte Carlo, December 2024. URL http://arxiv.org/abs/1903.04797. arXiv:1903.04797 [stat].

[28] Christopher Jackson. Multi-state models for panel data: the msm package for r. *Journal of statistical software*, 38:1–28, 2011.

[29] Mogens Bladt and Michael Sørensen. Statistical inference for discretely observed markov jump processes. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 67(3): 395–410, 2005.

[30] Robert T McGibbon and Vijay S Pande. Efficient maximum likelihood parameterization of continuous-time markov processes. *The Journal of chemical physics*, 143(3), 2015.

[31] Yu-Ying Liu, Shuang Li, Fuxin Li, Le Song, and James M Rehg. Efficient learning of continuous-time hidden markov models for disease progression. *Advances in neural information processing systems*, 28, 2015.

[32] Qingcan Wang. *Selected Topics in Deep Learning Theory and Continuous-Time Hidden Markov Models*. Princeton University, 2021.

[33] Richard J Boys, Darren J Wilkinson, and Thomas BL Kirkwood. Bayesian inference for a discretely observed stochastic kinetic model. *Statistics and Computing*, 18:125–135, 2008.

[34] Asger Hobolth and Eric A Stone. Simulation from endpoint-conditioned, continuous-time markov chains on a finite state space, with applications to molecular evolution. *The annals of applied statistics*, 3(3):1204, 2009.

[35] Vinayak Rao and Yee Whye Teh. Fast mcmc sampling for markov jump processes and extensions. *Journal of Machine Learning Research*, 14:3295–3320, 2013.

[36] Manfred Opper and Guido Sanguinetti. Variational inference for markov jump processes. *Advances in neural information processing systems*, 20, 2007.

[37] Ido Cohn, Tal El-Hay, Nir Friedman, and Raz Kupferman. Mean field variational approximation for continuous-time bayesian networks. *The Journal of Machine Learning Research*, 11: 2745–2783, 2010.

[38] Christian Wildner and Heinz Koeppl. Moment-based variational inference for markov jump processes. In *International Conference on Machine Learning*, pages 6766–6775. PMLR, 2019.

[39] Boqian Zhang, Jiangwei Pan, and Vinayak A Rao. Collapsed variational bayes for markov jump processes. *Advances in Neural Information Processing Systems*, 30, 2017.

[40] Lukas Köhs, Bastian Alt, and Heinz Koeppl. Variational inference for continuous-time switching dynamical systems. *Advances in Neural Information Processing Systems*, 34:20545–20557, 2021.

[41] Patrick Seifner and Ramsés J Sánchez. Neural markov jump processes. In *International Conference on Machine Learning*, pages 30523–30552. PMLR, 2023.

[42] David Berghaus, Kostadin Cvejoski, Patrick Seifner, Cesar Ojeda, and Ramses J Sanchez. Foundation inference models for markov jump processes. *arXiv preprint arXiv:2406.06419*, 2024.

[43] Andrew Golightly and Chris Sherlock. Efficient sampling of conditioned markov jump processes. *Statistics and Computing*, 29:1149–1163, 2019.

[44] Haoran Sun, Lijun Yu, Bo Dai, Dale Schuurmans, and Hanjun Dai. Score-based Continuous-time Discrete Diffusion Models, March 2023. URL `http://arxiv.org/abs/2211.16750`. arXiv:2211.16750 [cs].

[45] Andrew Campbell, Joe Benton, Valentin De Bortoli, Thomas Rainforth, George Deligiannidis, and Arnaud Doucet. A continuous time framework for discrete denoising models. *Advances in Neural Information Processing Systems*, 35:28266–28279, 2022.

[46] Ilia Igashov, Arne Schneuing, Marwin Segler, Michael Bronstein, and Bruno Correia. Retrobridge: Modeling retrosynthesis with markov bridges. *arXiv preprint arXiv:2308.16212*, 2023.

[47] Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion language modeling by estimating the ratios of the data distribution. *arXiv preprint arXiv:2310.16834*, 2023.

[48] Andrew Campbell, Jason Yim, Regina Barzilay, Tom Rainforth, and Tommi Jaakkola. Generative flows on discrete state-spaces: Enabling multimodal flows with applications to protein co-design. In *Forty-first International Conference on Machine Learning*, 2024.

[49] Jarrid Rector-Brooks, Mohsin Hasan, Zhangzhi Peng, Cheng-Hao Liu, Sarthak Mittal, Nouha Dziri, Michael M Bronstein, Pranam Chatterjee, Alexander Tong, and Joey Bose. Steering masked discrete diffusion models via discrete denoising posterior prediction. In *The Thirteenth International Conference on Learning Representations*, 2024.

[50] Masatoshi Uehara, Yulai Zhao, Chenyu Wang, Xiner Li, Aviv Regev, Sergey Levine, and Tommaso Biancalani. Inference-time alignment in diffusion models with reward-guided generation: Tutorial and review. *arXiv preprint arXiv:2501.09685*, 2025.

[51] Ricky TQ Chen, Brandon Amos, and Maximilian Nickel. Neural spatio-temporal point processes. In *International Conference on Learning Representations*, 2021.

[52] Hunter Nisonoff, Junhao Xiong, Stephan Allenspach, and Jennifer Listgarten. Unlocking Guidance for Discrete State-Space Diffusion and Flow Models, March 2025. URL `http://arxiv.org/abs/2406.01572`. arXiv:2406.01572 [cs].

[53] Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress: Discrete denoising diffusion for graph generation. *arXiv preprint arXiv:2209.14734*, 2022.

[54] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[55] Jörg Bornschein and Yoshua Bengio. Reweighted wake-sleep. *arXiv preprint arXiv:1406.2751*, 2014.

[56] Tuan Anh Le, Adam R. Kosiorek, N. Siddharth, Yee Whye Teh, and Frank Wood. Revisiting Reweighted Wake-Sleep for Models with Stochastic Control Flow, September 2019. URL http://arxiv.org/abs/1805.10469. arXiv:1805.10469 [stat].

[57] Declan McNamara, Jackson Loper, and Jeffrey Regier. Sequential monte carlo for inclusive kl minimization in amortized variational inference. In *International Conference on Artificial Intelligence and Statistics*, pages 4312–4320. PMLR, 2024.

A Background

A.1 Continuous time Markov chains.

Consider a Markov process $\{Z_t\}_{t\in[0,T]}$, taking value in a discrete state space $\mathcal{Z} = \{1, \dots, S\}$ over a finite, continuous time interval $[0, T]$. This system, known as a CTMC, can be fully characterized by an initial probability distribution $p_0 \in \mathcal{P}(\mathcal{Z})$, and a transition rate matrix $\mathbf{R} : \mathcal{Z}\times\mathcal{Z}\times[0,T] \to \mathbb{R}^{S\times S}$, where each entry represents an instantaneous probability rate defined as

$$R_t(z_{t+\Delta_t} \,|\, z_t) \coloneqq \lim_{\Delta t\to 0} \frac{1}{\Delta_t}\mathbb{P}(Z_{t+\Delta_t} = z_{t+\Delta_t} \,|\, Z_t = z_t), \quad z_{t+\Delta_t} \neq z_t,$$

and $R_t(z_t \,|\, z_t) = -\sum_{z\in\mathcal{Z}} R_t(z \,|\, z_t)$ for any z_t, $z_{t+\Delta_t} \in \mathcal{Z}$. Random trajectories $z : [0, T] \to \mathcal{Z}$ are càdlàg paths with a finite number of jumps, between which trajectories stay constant. We refer to a distribution of these a trajectories as *path measure*, denoted $P(dz_{[0,t]})$ for $t \in [0, T]$. A detailed introduction to the topic can be found in [25, 26].

A.2 Sequential Monte Carlo in Continous Time

Suppose we can define a sequence of *intermediate target distributions* $\{P_t^\circ\}_{t\in[0,T]}$, $P_t^\circ \in \mathcal{P}(\Omega_{[0,t]})$, such that $P_T^\circ = P^\star$. Moreover, we assume that $P_0^\circ \in \mathcal{P}(\mathcal{Z})$ is a distribution we can easily sample from and whose unnormalized density $\gamma_0^\circ(z)$ we can evaluate, and that we can evaluate unnormalized densities of transition probabilities $P^\circ_{t+\Delta_t\,|\,t}(z_t, dz_{t+1})$, denoted as $\gamma^\circ_{t+\Delta_t|t}(z_{t+\Delta_t} \,|\, z_t)$. Then, we can decompose the problem of sampling from $P^\star$ into a series of subproblems, where samples are evolved from one intermediate target to the next. Sequential Monte Carlo (SMC) methods are an ideal choice for approximating sequences of distributions of this kind [18, 27]. In an SMC algorithm, a set of $S \geq 1$ particles $z_0^{1:S} = (z_0^{(1)}, \dots, z_0^{(S)}) \in \mathcal{Z}^S$ is initialized as $z_0^{1:S} \overset{\text{iid}}{\sim} P_0^\circ$, and importance weights are initialized as $\bar{w}_0^s = w_0^s / \sum_{j=1}^S w_0^j$, where $w_0^s = \gamma_0^\circ(z_t^{(s)})/p_0(z_t^{(s)})$. Then, particles are evolved in a discrete time mesh

$$t_0 = 0 < t_1 < \cdots < t_n < t_{n+1} < \cdots < t_N = T. \tag{8}$$

This is achieved by iteratively performing the following steps for $n = 1, \dots, N-1$:

- Sample ancestry variables $A_{t_n}^{1:S} = (A_{t_n}^1, \dots, A_{t_n}^S)$ independently from a categorical distribution $A_{t_n}^{1:S} \overset{\text{iid}}{\sim} \text{Cat}(\bar{w}_{t_n}^{1:S})$ and reset weights to $w_{t_n}^s = 1/S$. In *adaptive* SMC [18], ,resampling is only performed if the effective sample size, defined as $ESS_{t_n} = \left(\sum_{s=1}^S (\bar{w}_{t_n}^s)^2\right)^{-1}$, falls below a pre-specified threshold. If this is not the case, weights are left unchanged.
- Propose new states $z_{t_{n+1}}^{(s)}$ at a future time t_{n+1} using a pre-specified proposal kernel $Q(z_{t_n}^{(A_{t_n}^s)}, dz_{t_{n+1}}^{(s)})$, whose density $q_{t_{n+1}|t_n}(z_{t_{n+1}}^{(s)} \mid z_{t_n}^{(A_{t_n}^s)})$ we assume we can evaluate.
- Update the trajectories by setting $z_{[t_0:t_{n+1}]}^{(s)} \leftarrow (z_{[t_0:t_n]}^{(A_{t_n}^s)}, z_{t_{n+1}}^{(s)})$
- Update the weights as

$$w_{t_{n+1}}^s = w_{t_n}^s \times \frac{\gamma^\circ_{t_{n+1}|t_n}(z_{t_{n+1}}^{(s)} \mid z_{t_n}^{(A_{t_n}^s)})}{q_{t_{n+1}|t_n}(z_{t_{n+1}}^{(s)} \mid z_{t_n}^{(A_{t_n}^s)})}, \quad \bar{w}_{t_{n+1}}^s = \frac{w_{t_{n+1}}^s}{\sum_{j=1}^S w_{t_{n+1}}^j}. \tag{9}$$

A.3 Related work

Inference for CTMCs. Inference methods for CTMCs have been extensively studied. Maximum likelihood estimation for time-homogeneous CTMCs is discussed in Jackson [28], Bladt and Sørensen [29], McGibbon and Pande [30]. Expectation-maximization techniques for continuous-time hidden Markov models have been developed in Liu et al. [31], and an overview of the topic can be found in Wang [32]. Bayesian approaches include Markov chain Monte Carlo methods [33–35] and variational methods. The latter include mean-field [36, 37], moment-based methods [38], combinations with MCMC [39], and extensions to hybrid processes [40]. Novel methods include black-box variational

inference with neural networks [41], foundation models [42], and expectation propagation [**?**]. We do not compare to these approaches, since to the best of our knowledge they can't easily be scaled to high dimensional systems of interacting CTMCs.

Another directly related line of research focuses on simulation methods for Markov bridges, notably [14, 13] and Golightly and Sherlock [43]. While less directly related, it is worth noting recent work discrete flow matching and diffusion methods [44, 22, 45–48]. In particular, reward-guided generation for discrete diffusion models targets a path measure that is a special case of 2 where a potential is only observed at the endpoint T [16, 49], see [50] for a review.

B Proofs

B.1 Proof of Proposition 1

Proof. We proceed by analyzing the transition kernel of $P^\star$ in an arbitrary interval $[s,t] \subseteq [0,T]$.

Let

$$M_T \coloneqq \frac{P^\star(dz_{[0,T]})}{P(dz_{[0,T]})} = \frac{\prod_{k=1}^K G_{\tau_k}(z_{\tau_k})}{\mathbb{E}_P[\prod_{k=1}^K G_{\tau_k}(z_{\tau_k})]}, \tag{10}$$

then, let the *filtered* path measure $P_t^\star$ be a restriction of $P^\star$ to the filtration $\mathcal{H}_t$, where $0 < t < T$, and denote

$$\frac{P_t^\star(dz_{[0,t]})}{P_t(dz_{[0,t]})} = M_t. \tag{11}$$

For an event $B \in \mathcal{H}_t$, by a simple application of the Radon-Nykodym theorem and the tower property we can write

$$P_t^\star(B) = P^\star(B) = \mathbb{E}_P[1_B M_T] = \mathbb{E}_P[1_B \mathbb{E}_P[M_T \,|\, \mathcal{H}_t]] = \mathbb{E}_{P_t}[1_B \mathbb{E}_{P_t}[M_T \,|\, \mathcal{H}_t]], \tag{12}$$

where the last step follows from $\mathbb{E}_P[M_T \,|\, \mathcal{H}_t]$ being measurable with respect to $\mathcal{H}_t$. Hence,

$$M_t = \mathbb{E}_{P_t}[M_T \,|\, \mathcal{H}_t] = \frac{1}{\mathbb{E}_P[\prod_{k=1}^K G_{\tau_k}(z_{\tau_k})]} \mathbb{E}_P\left[\prod_{k=1}^K G_{\tau_k}(Z_{\tau_k}) \middle| \mathcal{H}_t\right]. \tag{13}$$

By change of measure under conditional expectation, it follows that

$$\mathbb{E}_{P^\star}[f(Z_t) \mid Z_s = z] = \frac{\mathbb{E}_P[f(Z_t) M_t \mid Z_s = z]}{\mathbb{E}_P[M_t \mid Z_s = z]} = \frac{\mathbb{E}_P[f(Z_t) h_t(Z_t) \mid Z_s = z]}{h_s(z)}, \tag{14}$$

where h is defined as in equation 3 By definition, we can express the generator $\mathcal{L}_t^\star$ of $P_t^\star$ as

$$\mathcal{L}_t^\star(f)(z) = \lim_{\Delta_t \to 0} \frac{\mathbb{E}_{P^\star}[f(Z_{t+\Delta_t}) \mid Z_t = z] - f(z)}{\Delta_t} \tag{15}$$

$$= \lim_{\Delta_t \to 0} \frac{\mathbb{E}_P\left[f(Z_{t+\Delta_t}) \frac{h_{t+\Delta_t}(Z_{t+\Delta_t})}{h_t(z)} \mid Z_t = z\right] - f(z)}{\Delta_t} \tag{16}$$

Moreover, we can approximate $h_{t+\Delta_t}(z)$ for $t \in [\tau_k, \tau_{k+1} - \Delta_t)$, $k \in [1:K]$ using a Taylor expansion around time t

$$h_{t+\Delta_t}(z) = h_t(z) + \Delta_t \frac{\partial h_t(z)}{\partial t} + o(\Delta_t) \tag{17}$$

$$= h_t(z) - \Delta_t \sum_{i, v \neq z^i} r_t^i(v \,|\, z) \left[h_t(z^{i \to v}) - h_t(z)\right] + o(\Delta_t), \tag{18}$$

where the last line follows from Kolmogorov backward equation [25]:

$$\frac{\partial h_t(z)}{\partial t} = \mathcal{L}_t(h_t)(z) = \sum_{i, v \neq z^i} r_t^i(v \,|\, z) \left[h_t(z^{i \to v}) - h_t(z)\right].$$

For small Δ_t, we can express $\mathbb{E}_P\left[f(Z_{t+\Delta_t}) \frac{h_t(Z_{t+\Delta_t})}{h_t(z)} \mid Z_t = z\right]$ using the law of total expectation, where we split the expectation based on the number of jumps in the interval $[t, t+\Delta_t]$.

$$\mathbb{E}_P\left[f(Z_{t+\Delta_t})h_t(Z_{t+\Delta_t}) \mid Z_t = z\right] \tag{19}$$

$$= f(z)h_t(z)\left(1 - \Delta_t \sum_{i,v\neq z^i} r_t^i(v \mid z)\right) + \sum_{i,v\neq z^i} f(z^{i\to v})h_t(z^{i\to v})\Delta_t r_t^i(v \mid z) + o(\Delta_t) \tag{20}$$

$$= f(z)h_t(z) + \Delta_t \sum_{i,v\neq z^i} r_t^i(v \mid z)\left[f(z^{i\to v})h_t(z^{i\to v}) - f(z)h_t(z)\right] + o(\Delta_t). \tag{21}$$

Combining equation 18 and equation 21, we obtain

$$\mathbb{E}_P\left[f(Z_{t+\Delta_t})h_{t+\Delta_t}(Z_{t+\Delta_t}) \mid Z_t = z\right] \tag{22}$$

$$= f(z)h_t(z) + \Delta_t \sum_{i,v\neq z^i} r_t^i(v \mid z)h_t(z^{i\to v})\left[f(z^{i\to v}) - f(z)\right] + o(\Delta_t). \tag{23}$$

Then, plugging equation 23 back into equation 16, we get

$$\mathcal{L}_t^\star(f)(z) = \lim_{\Delta_t\to 0} \frac{\Delta_t \sum_{i,v\neq z^i} r_t^i(v \mid z)\frac{h_t(z^{i\to v})}{h_t(z)}\left[f(z^{i\to v}) - f(z)\right] + o(\Delta_t)}{\Delta_t} \tag{24}$$

$$= \sum_{i,v\neq z^i} r_t^i(v \mid z)\frac{h_t(z^{i\to v})}{h_t(z)}\left[f(z^{i\to v}) - f(z)\right] \tag{25}$$

By inspection, we recognize in equation 25 the generator of an IPS, with local rates

$$r_t^{\star,i}(v \mid z) \coloneqq \begin{cases} r_t^i(v \mid z)\frac{h_t(z^{i\to v})}{h_t(z)}, & v \neq z^i, \\ \sum_{u\neq v} r_t^i(u \mid z)\frac{h_t(z^{i\to u})}{h_t(z)}, & v = z^i. \end{cases} \tag{26}$$

The initial distribution follows from a simple application of the Bayes theorem, and is equal to

$$p_0^\star(z) = \frac{p_0(z)h_0(z)}{\mathbb{E}_{p_0}[h_0(Z_0)]}. \tag{27}$$

□

C Experimental details

C.1 Data generation

We simulate SIRS epidemics on undirected graphs with $|\mathcal{I}| = d$ nodes. Graphs are sampled from an expected-degree model. Each node i has a feature vector $\xi_i \in \mathbb{R}^p$ and we write $\hat{\xi}_i = \xi_i/\|\xi_i\|_2$. For any global state $z \in \mathcal{Z}$, define $S(z) = \{i : z^i = S\}$, $I(z) = \{i : z^i = I\}$, $R(z) = \{i : z^i = R\}$. The local SIRS rates are

$$r_t^i(I \mid z) = \left(\alpha_0 + \alpha_1 \sum_{j\neq i} a_{ij}\,\sigma(\langle\hat{\xi}_i, \hat{\xi}_j\rangle)\,\delta_{z^j,I}\right)\delta_{z^i,S}, \tag{28}$$

$$r_t^i(R \mid z) = \beta\,\delta_{z^i,I}, \qquad r_t^i(S \mid z) = \gamma\,\delta_{z^i,R}, \tag{29}$$

$$r_t^i(z^i \mid z) = -\sum_{v\neq z^i} r_t^i(v \mid z), \tag{30}$$

with adjacency a_{ij} and σ the logistic function to maintain positivity. We fix $(\alpha_0, \alpha_1, \beta, \gamma)$ throughout. Ground truth paths on $[0, T]$ for the test set are drawn using Gillespie's algorithm.

C.2 Observation model

At $K = 10$ snapshot times $\tau_1 < \cdots < \tau_K$, emissions are conditionally independent across nodes given Z_{τ_k}. For each node $i \in \mathcal{I}$ we include an explicit mask token $\varnothing$ and use a node-factorized emission

$$p(x_{\tau_k} \mid Z_{\tau_k} = z) \;=\; \prod_{i\in\mathcal{I}} g\big(x_{\tau_k}^i \mid z^i\big)\,,$$

with a masking probability set to $p_{\text{mask}} = \frac{1}{2}$ and small symmetric label noise $\delta > 0$, for numerical stability:

$$g(c \mid z^i) \;=\; \begin{cases} p_{\text{mask}}, & c = \varnothing, \\ (1 - p_{\text{mask}})\big[(1 - \delta(V-1))\,\mathbf{1}\{c = z^i\} + \delta\,\mathbf{1}\{c \neq z^i\}\big], & c \in \mathcal{V}, \end{cases}$$

where $V = |\mathcal{V}|$.

C.3 Methodology

We discretize each interval $[\tau_k, \tau_{k+1}]$ with a uniform grid of step Δ (rescaled inside each interval) following [51], and run twisted SMC with M particles. At each sub-step we simulate jumps in parallel for each dimension using the twisted transition implied by h_t^ψ (Eq. equation **??**), update importance weights with the incremental Radon–Nikodym factor between prior and twisted kernels, and, at τ_k, multiply by G_{τ_k}, i.e.

We resample using systematic resampling [10]. All methods share Δ, M, and the resampling rule.

Baselines.

- **Bootstrap particle filter** (BPF)[18]: this method is an SMC algorithm using the filtering distributions as intermediate targets, and the prior transition probabilities as a proposal distribution.
- **NASMC and NAS-X**[21**?**]: we learn a proposal distribution by fitting a free-form score network minimizing the forward KL loss in 5, i.e. we parameterize

  $$\log \frac{h_t(z^{i \to v})}{h_t(z)} \approx s_t^\psi(z)_{[i,v]},$$

  where $s_t^\psi : \mathcal{Z} \to \mathbb{R}^{d \times V}$ and $s_t^\psi(z)_{[i,z^i]} = 1$ for $i \in \mathcal{I}$. In [21], the intermediate targets are the filtering distributions. In [**?**], the intermediate targets are the twisted distributions, where the twist function is learned using a density ratio estimation loss.
- **Taylor-approximated guidance** (TAG): to overcome the inefficiency of computing $d \times (V-1) + 1$ forward passes of the twist function, [52] proposed to compute a first-order Taylor approximation of the log-twist evaluated at a specific value z, i.e.

  $$\log \psi_t(z_t) \approx \log \psi_t(\mathbf{z}) + \mathbf{z_t}^\top \nabla_{\mathbf{z}} \log \psi_t(\mathbf{z}) \tag{31}$$

  where $\mathbf{z}, \mathbf{z_t}$ are one-hot encoded versions of z, z_t, enabling a single forward pass at z of the twist function.

For all of our methods (except BPF), we also train a variational initial distribution $q_0(z_0)$ with a forward KL loss on generated trajectories. For simplicity, we let this be a different head of our twist network.

Twist training. We train ψ on trajectories simulated from the forward model, following [**? ?**]. We use the following losses for the twist:

$$\hat{\mathcal{L}}_{\text{KL-sleep}}(\psi) = \sum_{m=1}^{M} \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I}} \left(r_t^{\psi,i}(z_t^{i,(m)} \mid z_t^{(m)}) dt - \delta(z_t^i - z_{t-\Delta_t}^i) \log r_t^{\psi,i}(z_{t^-}^i \mid z_t) \right) \tag{32}$$

Note that an alternative approximation to 5 is given by the *reweighted wake loss*

$$\hat{\mathcal{L}}_{\text{KL-wake}}(\psi) = \sum_{m=1}^{M} \sum_{t \in \mathcal{T}} \bar{w}_t^m \sum_{i \in \mathcal{I}} \left(r_t^{\psi,i}(z_t^{i,(m)} \mid z_t^{(m)}) dt - \delta(z_t^i - z_{t-\Delta_t}^i) \log r_t^{\psi,i}(z_{t^-}^i \mid z_t) \right) \tag{33}$$

where $\{z_t^{(m)}, w_t^{(m)}\}_{t,m}$ are generated from previous runs of tSMC and stored in a buffer, and $\bar{w}_t^{(m)} = \texttt{stop\_grad}(w_t^{(m)})$.

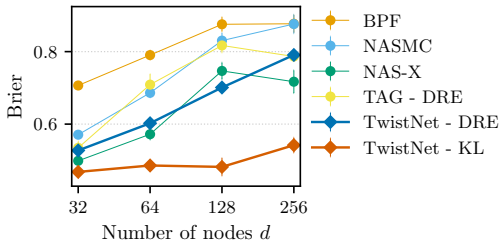

Figure 3: Latent trajectory reconstruction, measured by Brier score of ground truth trajectories with respect to the posterior approximations. See Figure 2b.

The density ratio estimation loss is:

$$\hat{\mathcal{L}}_{\text{DRE}}(\psi) = \sum_{t\in\mathcal{T}} \sum_{i\in\mathcal{I}} \log \sigma(\log h_t^{\psi}(z_t^{+}; x_{\geq t}, \tau_{\geq t})) + \log(1 - \sigma(\log h_t^{\psi}(z_t^{-}; x_{\geq t}, \tau_{\geq t})) \tag{34}$$

where $\sigma : \mathbb{R} \to [0, 1]$ is the logistic function. Positive samples z_t^{+} are generated by the forward model using ancestral sampling of $z_{[0,T]}^{+} \sim P$ first, and then $x_{1:K}, \tau_{1:K}$. Negative samplese $z_t^{-} \sim P$, and are hence uncoupled from $x_{1:K}, \tau_{1:K}$. Using this loss is equivalent to training a classifier to distinguish between coupled and uncoupled sampled.

We parameterize the score networks and the encoder of the TwistNet with a Graph Transformer (GT) from [53]. For all of our models we use 2 GT layers and 4 heads, with a node embedding dimensions of 128, an edge embedding dimension of 32, and global information embedding of 128. We did not tune any of these hyperparameters.

For any time $t \in [0, T]$, we feed as input the feature vector, future observation $x_{\geq t}$ and observation times $\tau_{\geq t}$, and graph statistics computed on the adjacency matrix as in [53]. The score network also takes as input the current state z_t, while in the TwistNet this is only considered when aggregating encoder outputs. For the TwistNet, we let the last layer be a two-layer MLP with $m = 512$. For optimization, we use the Adam optimizer [54] with learning rate 0.001 and the default PyTorch hyperparameters [1].

We do *not* perform parameter learning of $(\alpha_0, \alpha_1, \beta, \gamma)$ in these experiments, but rather consider them fixed at the ground truth. We are interested in inferring unknown latent trajectories for a given set of observations, i.e. inverting the forward model. Note that our procedure could be alternated with parameter fitting in a Monte Carlo EM algorithm, yielding a procedure similar to reweighted wake-sleep [55] or wake-wake [56], reweighted by SMC weights as in [57]. At the time of writing we are finalizing these experiments, and they will likely be included in future versions of this manuscript.

C.4 Evaluation metrics

We are interested in understanding whether, for a prescribed forward model, our method can be used to perform posterior inference given a set of observations.

From weighted particles we form per-time nodewise marginals $\hat{p}_t(z) = \sum_{m=1}^{M} \delta_{z_t^{(m)}}(z)$, and we add a small uniform weight to avoid numerical issues when support is scarce:

$$\tilde{p}_t \;=\; (1-\epsilon)\,\hat{p}_t \;+\; \epsilon\,\text{Unif}(\mathcal{V}).$$

Let $z_t^{\star} \in \mathcal{Z}$ be the ground-truth state of the latent trajectory at time $t \in [0, T]$. We report an average of the following metrics over a test set of 50 trajectories:

$$\text{NLL} \;=\; -\frac{1}{|\mathcal{T}|}\sum_{t\in\mathcal{T}} \log \tilde{p}_t(z_t^{\star}), \quad \text{Brier} \;=\; \frac{1}{|\mathcal{T}|}\sum_{t\in\mathcal{T}} \|\tilde{p}_t - \mathbf{z}_t^{\star}\|_2^2.$$

where $\mathcal{T}$ is the set of discretized time indices, and $\mathbf{z}_t$ is the one-hot encoded $z_t^{\star}$. For each method we consider 25 particles ($M = 25$), except for BPF for which we take 250 particles. The NLL over

[1] `https://docs.pytorch.org/docs/stable/generated/torch.optim.Adam.html`

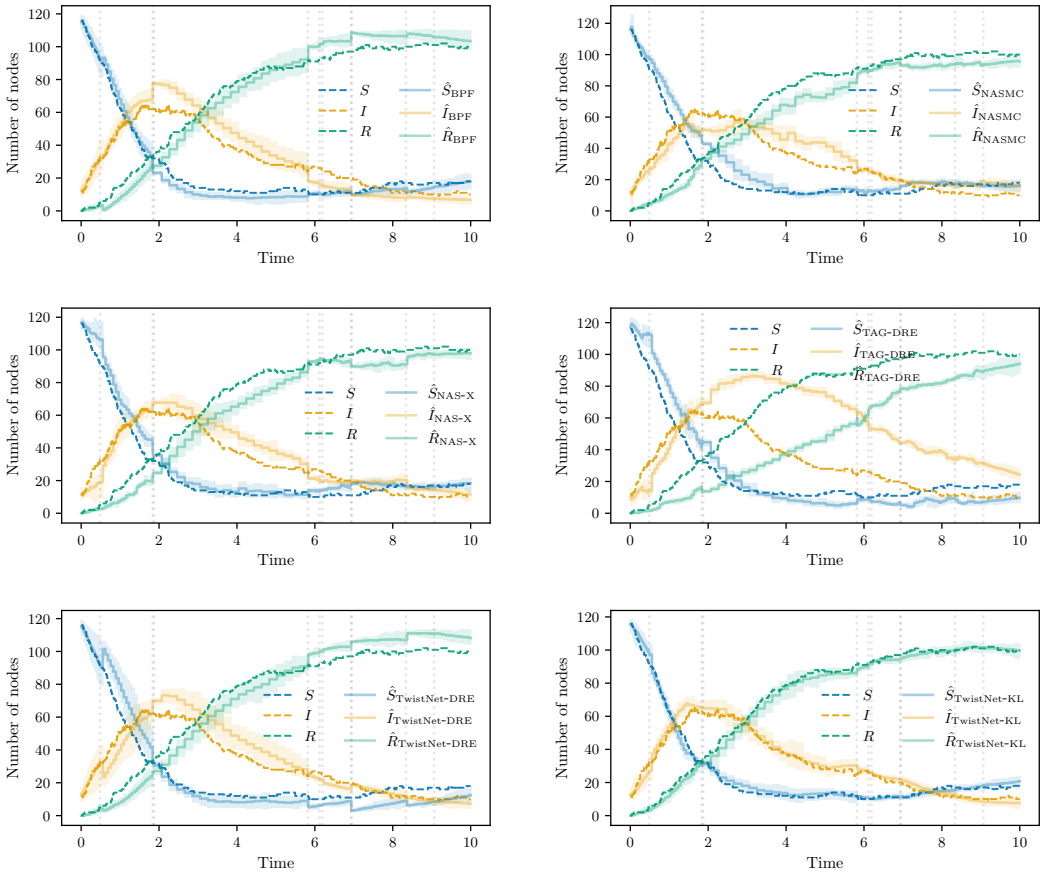

Figure 4: First example of latent trajectories, with counts of each state averaged over the graph at each timestep.

dimensions is displayed in Figure 2b, and the Brier score in Figure 3. In Figures 4 and 5, we show a summary of generated latent trajectories for each method for different graphs with $d = 128$ nodes. In line with the evaluation in Figures 2b and 3, the TwistNet-KL method seems to display much better performance than the alternatives in terms of closeness to the ground truth. In the future, we plan on expanding this evaluation to evaluate whether the model can recover multiple modes of the smoothing path measure.

While twisted SMC methods shows huge improvements in performance over traditional schemes for high dimensional problems, this comes at a cost: the need to perform a forward pass of the neural network at each timestep. This significantly increases the runtime of this family of methods, and this is possibly the biggest limitation of these methods when applied to large datasets.

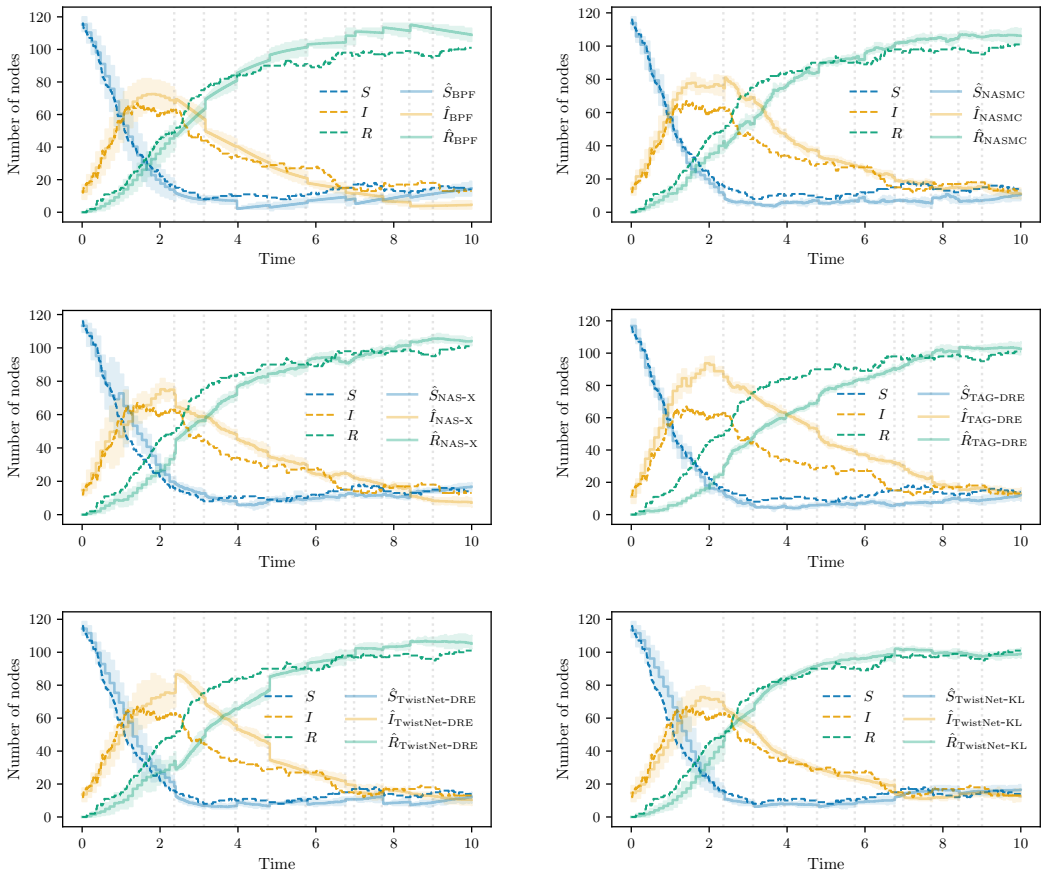

Figure 5: Second example of latent trajectories, with counts of each state averaged over the graph at each timestep.