Variational Inference for Interacting Particle Systems with Discrete Latent States

Anonymous Author(s)
Affiliation
Address
email

Abstract

We present a novel Bayesian learning framework for interacting particle systems with discrete latent states, addressing the challenge of inferring dynamics from partial, noisy observations. Our approach learns a variational posterior path measure by parameterizing the generator of the underlying continuous-time Markov chain. We formulate the problem as a multi-marginal Schrödinger bridge with aligned samples, employing a two-stage learning procedure. Our method incorporates an emission distribution for decoding latent states and uses a scalable variational approximation.

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 paradigm extends the concept of hidden Markov models [Baum and Petrie, 1966, Kouemou and Dymarski, 2011] to spatially structured, continuous-time processes. Interacting particle systems (IPSs) [Liggett, 1985, Lanchier, 2024] offer a powerful mathematical framework for describing local propagation dynamics. However, inferring the rules governing these systems from partial, noisy observations remains a significant challenge. We propose a novel Bayesian approach that addresses this challenge by learning a variational posterior path measure on the space of IPS trajectories. Our approach parameterizes the generator of the continuous-time Markov chain (CTMC) of the latent IPS using neural networks and incorporates an emission model that can decode internal discrete states to continuous data and noisy observations. Key contributions of our approach include:

- Framing the problem as a multi-marginal Schrödinger bridge with aligned samples [Somnath et al., 2023], solved by a two-stage procedure: learning an endpoint-conditioned generator for trajectory reconstruction, followed by distillation to an unconditional generator for prediction.
- A scalable variational approximation using site-wise factorization of time-marginals and assuming independent particle evolution in infinitesimal time intervals conditionally on the present global configuration, enabling efficient learning for high-dimensional spatio-temporal processes.
- Flexibility in incorporating domain knowledge through through informative priors on rate matrix entries and neural architectures with desirable inductive biases.

We demonstrate preliminary results of our approach on two simulated datasets, for the following tasks: reconstructing the trajectory of an epidemic on a network and predicting wildfire spread on a lattice. For a description of the notation, see Appendix A. An overview of the relevant literature is presented in Appendix B, whereas proofs and other derivations are provided in Appendix C.

Submitted to Workshop on Bayesian Decision-making and Uncertainty, 38th Conference on Neural Information Processing Systems (BDU at NeurIPS 2024). Do not distribute.

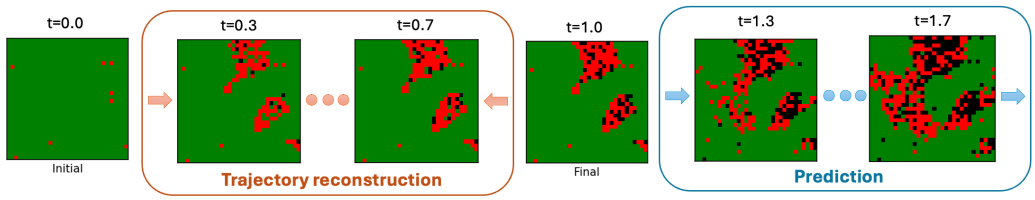

Figure 1: An illustration of our methodology on a simulated noiseless dataset of wildfire propagation. The first model approximates a Markovian bridge interpolating between the observations, enabling to reconstruct the unobserved trajectory. The second model, approximating the unconditional process, can predict beyond the last observation. Results shown for a held-out example.

2 Background

Interacting particle systems Consider a graph $\mathcal{G} = (V, E)$, and denote $i \sim j$ if there is an edge between nodes i, j, i.e., $\{i, j\} \in E$. Following Liggett [1985], we refer to vertices $i \in V$ as sites. For a countable state space S, consider the configuration space $\mathcal{X} := \{\mathbf{x} \mid \mathbf{x} : V \to S\}$. For our analysis, we assume both V and S are finite. An IPS adds a continuous-time dimension to this setting. Specifically, we obtain a CTMC $(\mathbf{x}_t)$ on $\mathcal{X}$ restricted to a time interval $[0, T]$, whose path space we denote $\Omega_{[0,T]}$. We define $\mathbf{x}_t(i)$ as the state of site i at time t. We consider a scenario where the dynamics of each site are described by local transition rates that depend on the graph's connectivity [Lanchier, 2017], corresponding to

$$\lambda_t^{s\to s'}(i, \mathbf{x}_t) \coloneqq \lim_{h\downarrow 0} h^{-1}\mathbb{P}\left(\mathbf{x}_{t+h}(i) = s' \mid \mathbf{x}_t(i) = s,\, \mathbf{x}_t(j) : i \sim j\right),$$

for s to $s' \neq s$ at site i and time $t \in [0, T]$, and set $\lambda_t^{s\to s}(i, \mathbf{x}_t) \coloneqq -\sum_{s'\neq s} \lambda_t^{s\to s'}(i, \mathbf{x}_t)$. The local transition rates are aggregated into a generator $\mathbf{\Lambda}_t(\mathbf{x}_t) = [\lambda_t^{s\to s'}(i, \mathbf{x}_t)]_{i\in V;,s,s'\in S}$. Note that this is a mapping $\mathbf{\Lambda}_\cdot(\cdot) : \mathcal{X} \times \mathbb{R}+ \to \mathbb{R}^{|V|\times|S|\times|S|}$. For brevity, we denote the collection of transition rates between two fixed configurations $\mathbf{x}$, $\mathbf{x}' \in \mathcal{X}$ as $\Lambda_t(\mathbf{x}' \mid \mathbf{x}) \coloneqq [\lambda_t^{\mathbf{x}(i)\to\mathbf{x}'(i)}(i, \mathbf{x})]_{i\in V}$. The path measure $\Pi \in \mathcal{M}(\Omega_{[0,T]})$ of a realization $(\mathbf{x}_t)_{t\in[0,T]}$ can then be described by the solution to an initial value problem with starting condition $\pi_t(\mathbf{x}_0) = \pi_0$, evolving according to the master equation

$$\partial_t \pi_t(\mathbf{x}_t) = \sum_{\mathbf{x}'\neq\mathbf{x}_t} \left(\Lambda_t(\mathbf{x}_t \mid \mathbf{x}')\pi_t(\mathbf{x}') - \Lambda_t(\mathbf{x}' \mid \mathbf{x}_t)\pi_t(\mathbf{x}_t)\right). \tag{1}$$

Markovian bridges Let $\{\mathbf{y}_k\}_{k\in[K]} \in \mathcal{Y}^K$ be a sequence recorded at times $\{t_k\}_{k\in[K]} \in \mathbb{R}^K$, where $0 = t_0 < \cdots < t_{K-1} = T$, and assume that each value in the sequence is independent conditionally on the Markov process $(\mathbf{x}_t)_{t\in[0,T]} \in \Omega_{[0,T]}$. The conditioned Markov process with path measure $\Pi_{\cdot|\{\mathbf{y}_k\}_{k\in[K]}} \in \mathcal{M}(\Omega_{[0,T]})$ is known as a Markovian bridge. At any time $t \in (t_k, t_{k+1}]$ and for states s and s' ($s \neq s'$), the transition rate is:

$$\lambda_t^{s\to s'}(i, \mathbf{x}_t \mid \{\mathbf{y}_k\}_{k\in[T]}) = \lambda_t^{s\to s'}(i, \mathbf{x}_t)\frac{\mathbb{P}(\mathbf{y}_{k+1} \mid \mathbf{x}_t = s')}{\mathbb{P}(\mathbf{y}_{k+1} \mid \mathbf{x}_t = s)}. \tag{2}$$

Hence, the conditional generator $\mathbf{\Lambda}_t(\mathbf{x}_t \mid \{\mathbf{y}_k\}_{k\in[T]})$ of $\Pi_{\cdot|\{\mathbf{y}_k\}_{k\in[K]}}$ is equivalent to the generator $\mathbf{\Lambda}_t(\mathbf{x}_t \mid \mathbf{y}_{t_{k+1}})$ of the Markovian bridge $\Pi_{\cdot|\mathbf{y}_{t_k},\mathbf{y}_{t_{k+1}}}$ at any time $t \in (t_k, t_{k+1}]$. See Appendix C.1 for an overview of Markovian bridges, and we refer the reader to Fitzsimmons et al. [1992] for a detailed construction.

3 Variational Discrete Interacting Particle Systems

We consider sequences of observations $\{\mathbf{y}_k^i\}_{k\in[K]} \in \mathcal{Y}^K$ recorded at times $\{t_k^i\}_{k\in[K]} \in \mathbb{R}^K$ and assumed to be independent conditionally on the interacting particle system $(\mathbf{x}_t^i)_{t\in[0,T^i]} \in \Omega_{[0,T^i]}$, for $i = 1, \ldots, N$. The discrete set of measurement times $0 = t_0^i < \cdots < t_{K-1}^i = T^i$ are allowed to be arbitrarily defined for each sequence, e.g., at random or regularly spaced. We assume that the graph determining the particles' dependence structure is fixed for each realization and directly deducible from the observed sequences, for example if the observations are noisy measurements of the IPS. We consider an emission distribution $p_t(\mathbf{y} \mid \mathbf{x}) \in \mathcal{P}(\mathcal{Y})$ for the observations, and a prior path measure $P \in \mathcal{M}(\Omega_{[0,T]})$ for the latent IPS. This can be specified directly on the entries of a prior generator, encoding possible constraints in the latent dynamics, and by an initial prior distribution. The marginal distribution of the data at any time $t \in [0,T]$ is denoted as $\pi_t \in \mathcal{P}(\mathcal{Y})$, and we are interested in parameterizing a variational path measure $Q \in \mathcal{M}(\Omega_{[0,T]})$ and an encoder $q_t(\mathbf{x} \mid \mathbf{y}) \in \mathcal{P}(\mathcal{X})$. To make inference tractable, we detail specific choices of our variational approximation in Section 3.1. Considering a single sequence $\{\mathbf{y}_k\}_{k\in[K]}$, our goal is twofold:

- **Trajectory reconstruction**, by learning the conditional generator $\mathbf{\Lambda}_t(\cdot \mid \{\mathbf{y}_k\}_{k\in[T]})$ of the Markovian bridge $Q_{\cdot|\{\mathbf{y}_k\}_{k\in[K]}}$;
- **Prediction**, by learning the generator $\mathbf{\Lambda}_t(\cdot)$ of Q, enabling extrapolation beyond an observed time window or with no past observations at all for a given graph.

Inference for both models can be conveniently amortized by parameterizing the generators with neural cellular automata, as detailed in Section 3.2. We show that the second goal can be achieved by distilling knowledge from a model trained for the first goal into a model that does not glance at future observations. Sampling is discussed in Appendix D.2.

3.1 Variational approximations

Contrary to previous work [Opper and Sanguinetti, 2007, Cohn et al., 2010, Seifner and Sánchez, 2023] we do not adopt a *mean-field* approximation, but rather the *star-approximation* introduced for continuous-time Bayesian networks in Linzner and Koeppl [2018]. While still assuming site-wise factorization of the time-marginals as $q_t(\mathbf{x}) = \prod_{i\in V} q_t^i(\mathbf{x}(i))$, we couple sites by letting the generator depend on the present global configuration $\mathbf{x}_t$. This is achieved by considering a system of marginal master equations

$$\partial_t q_t^i(\mathbf{x}_t(i)) = \sum_{s' \neq \mathbf{x}_t(i)} \left(\lambda^{s' \to \mathbf{x}_t(i)}(i, \mathbf{x}_t) q_t^i(s') - \lambda^{\mathbf{x}_t(i) \to s'}(i, \mathbf{x}_t) q_t^i(\mathbf{x}_t(i)) \right), \quad i \in V. \tag{3}$$

Under these assumptions, the explicit form for the Kullback-Leibler (KL) divergence of two CTMCs with path measures $Q,\ P \in \mathcal{M}(\Omega_{[0,T]})$ is tractable, see Appendix C.2 for a derivation.

3.2 Multi-marginal aligned Schrödinger bridges

Let $\mathrm{P} \in \mathcal{P}(\mathcal{Y}^K \times \Omega_{[0,T]})$ denote the reference measure constructed by gluing the prior and emission probabilities at each observed timestep. For a given sequence of distributions $\{\pi_{t_k}\}_{k\in[K]}$ on $\mathcal{P}(\mathcal{Y})$, we can express a multi-marginal Schrödinger bridge problem with noisy observations as

$$\mathrm{Q}^* := \operatorname*{arg\,min}_{\mathrm{Q}\in\mathcal{P}(\mathcal{Y}^K\times\Omega)} \{D_{\mathrm{KL}}(\mathrm{Q} \,||\, \mathrm{P}) \mid q_{t_k} = \pi_{t_k},\ k \in [K]\}, \tag{4}$$

where $q_{t_k} \in \mathcal{P}(\mathcal{Y})$ correspond to marginalizations of the variational distribution at each observed timepoint. Let $\pi_{t_{0:T-1}}$ denote the coupling solving the static version of (4), that is

$$\pi_{t_{0:T-1}} = \operatorname*{arg\,min}_{q_{t_{0:T-1}}\in\mathcal{P}(\mathcal{Y}^K)} \{D_{\mathrm{KL}}(q_{t_{0:T-1}} \,||\, p_{t_{0:T-1}}) \mid q_{t_k} = \pi_{t_k},\ k \in [K]\}, \tag{5}$$

where $p_{t_{0:T-1}} \in \mathcal{P}(\mathcal{Y}^K)$ is the marginal of the observed trajectories obtained from the reference measure P. Similarly to the setting considered in Somnath et al. [2023], we assume that our dataset is comprised of trajectories of *aligned* samples, in the sense that each observed trajectory $\{\mathbf{y}_k\}_{k\in[T]}$ is sampled from the coupling $\pi_{t_{0:T-1}}$. We denote its marginals at any pair of observed time points as $\pi_{t_k,t_{k+1}} \in \mathcal{P}(\mathcal{Y}\times\mathcal{Y})$, for $k \in [K-1]$.

Proposition 1 *Let (5) admit a solution $\pi_{t_{0:T-1}}$. Moreover, assume conditional independence of $\{\mathbf{y}_k\}_{k\in[K]}$ given $(\mathbf{x}_t) \in \Omega_{[0,T]}$, and let $P \in \mathcal{M}(\Omega_{[0,T]})$. Then, the problem in (4) reduces to*

$$\operatorname*{arg\,min}_{q_{\cdot|\mathbf{y}_0},\{Q_{\cdot|\mathbf{y}_k,\mathbf{y}_{k+1}}\}} \mathbb{E}_{\pi_{t_0}}\left[D_{\mathrm{KL}}(q_{\cdot|\mathbf{y}_0} \,||\, p_{\cdot|\mathbf{y}_0})\right] + \sum_{k\in[K]} \mathbb{E}_{\pi_{t_k,t_{k+1}}}\left[D_{\mathrm{KL}}(Q_{\cdot|\mathbf{y}_k,\mathbf{y}_{k+1}} \,||\, P_{\cdot|\mathbf{y}_k,\mathbf{y}_{k+1}})\right], \tag{6}$$

where $q_{\cdot|\mathbf{y}_0} \in \mathcal{P}(\mathcal{X})$ and $Q_{\cdot|\mathbf{y}_k,\mathbf{y}_{k+1}} \in \mathcal{M}(\Omega_{(t_k,t_{k+1}]})$ for $k \in [K-1]$.

Trajectory reconstruction As any conditional path measure can be fully characterized by its generator and an initial distribution, we leverage the correspondence explored in Section 2 to specify $Q_{\cdot|\{\mathbf{y}_k\}_{k\in[K]}}$ by a sequence of Markovian bridges $Q_{\cdot|\mathbf{y}_k,\mathbf{y}_{k+1}}$ and conditional distributions $q_{t_k}(\cdot \mid \mathbf{y}_k)$, for $k \in [K-1]$. We parameterize the former with a neural model $\mathbf{\Lambda}_t^\theta(\cdot \,|\, \mathbf{y}_{k+1})$, and the latter with a probabilistic encoder $q_t^{i,\theta}(\cdot \mid \mathbf{y}_k) = \text{Categorical}(g_t^\theta(i,\, \mathbf{y}_k))$, where $g^\theta \,:\, \mathbb{R}_+ \times V \times \mathcal{Y} \to \Delta_{|S|}$ and $\Delta_{|S|}$ denotes the $|S|-$dimensional simplex. Moreover, we parameterize the emission distribution with a probabilistic decoder $p_t^\xi(\mathbf{y} \mid \mathbf{x})$. We learn θ and ξ by minimizing an evidence lower bound derived from (6), given by

$$\mathcal{L}_1(\theta,\xi) := \sum_{k\in[K-1]} \mathbb{E}_{\pi_{t_k,t_{k+1}}(\mathbf{y}_k,\mathbf{y}_{k+1})}[\mathcal{L}_1^k(\theta,\xi,\mathbf{y}_k,\,\mathbf{y}_{k+1})] - \mathbb{E}_{q^\theta_{t_0}(\mathbf{x}_{t_0}|\boldsymbol{y}_0),\pi_{t_0}(\boldsymbol{y}_0)}[\log p^\xi_{t_0}(\mathbf{y}_0 \mid \mathbf{x}_{t_0})], \tag{7}$$

where

$$\mathcal{L}_1^k(\theta,\xi,\mathbf{y}_k,\,\mathbf{y}_{t_{k+1}}) := D_{\mathrm{KL}}(Q^\theta_{\cdot|\mathbf{y}_k,\,\mathbf{y}_{k+1}} \,||\, P) - \mathbb{E}_{q^\theta_{t_{k+1}}(\mathbf{x}_{t_{k+1}}|\mathbf{y}_k,\mathbf{y}_{k+1})}[\log p^\xi_{t_k}(\mathbf{y}_{k+1} \mid \mathbf{x}_{t_{k+1}})]. \tag{8}$$

See Appendix C.4 for a derivation, and Wildner and Koeppl [2019] for an alternative proof. A similar result is derived in Lavenant et al. [2021] for diffusion processes. At training time, we start by sampling a minibatch of pairs $\mathbf{y}_k,\ \mathbf{y}_{k+1}$. The sequence of variational distributions $\{(q_t^{i,\theta})_{t\in[t_k,t_{k+1}]}\}_{i\in V}$ is then obtained by numerically solving a system of marginal master equations as in (3), according to the neighborhood dynamic established by $\mathbf{\Lambda}_t^\theta(\mathbf{x}_t \,|\, \mathbf{y}_{t_{k+1}})$, for $t \in (t_k,\, t_{k+1}]$. This step can be achieved by using adaptive solvers, but for illustration purposes we remind the reader that an Euler update with step size $0 < \epsilon \ll 1$ can be performed as

$$q_{t+\epsilon}^{i,\theta} = q_t^{i,\theta} + \epsilon[\mathbf{\Lambda}_t^\theta(\mathbf{x}_t \,|\, \mathbf{y}_{t_{k+1}})]_i^\top q_t^{i,\theta}, \quad i \in V. \tag{9}$$

As we are interested in evolving the entire vector of probability mass functions at each timestep, we sample $\mathbf{x}_t$ for $t \in (t_k,\, t_{k+1}]$ from $q_t^{i,\theta}$ using the Gumbell-Softmax trick [Jang et al., 2017]. Additional details are reported in Appendix D.1, and considerations of computational cost are discussed in Appendix D.4.

Prediction We fix the conditional process Q^θ, and learn the unconditional IPS by distillation to a Markovian generator $\mathbf{\Lambda}_t^\phi(\mathbf{x}_t)$. This can be achieved by minimizing a mean squared error loss on the rates, given by

$$\mathcal{L}_2(\phi) := \sum_{k\in[K-1]} \mathbb{E}_{\pi_{t_k,t_{k+1}}(\mathbf{y}_k,\mathbf{y}_{k+1})}\left[\mathcal{L}_2^k(\phi,\mathbf{y}_k,\mathbf{y}_{t_{k+1}})\right], \tag{10}$$

where

$$\mathcal{L}_2^k(\phi,\mathbf{y}_k,\mathbf{y}_{t_{k+1}}) := \int_{t_k}^{t_{k+1}} \mathbb{E}_{q_t^\theta(\mathbf{x}_t|\mathbf{y}_k,\mathbf{y}_{k+1})} \sum_{\mathbf{x}'\neq\mathbf{x}_t} \left\|\mathbf{\Lambda}_t^\theta(\mathbf{x}_t \,|\, \mathbf{y}_{t_{k+1}}) - \mathbf{\Lambda}_t^\phi(\mathbf{x}_t)\right\|_2^2 dt. \tag{11}$$

If the conditional generator $\mathbf{\Lambda}_t^\theta$ is unbiased and $q_t^\theta \approx q_t^*$, the minimizer of (10) recovers the Markov process $Q \in \mathcal{P}(\Omega_{[0,T]})$ obtained by marginalizing Q^* in (4) to $\Omega_{[0,T]}$. A derivation is provided in Appendix C.5.

4 Experiments

We demonstrate our methodology on two simulated scenarios: epidemic trajectory inference on networks and wildfire spread prediction on lattices. We parameterize the neural models for the generators with a Vision Transformer Cellular Automata [Tesfaldet et al., 2022]. Results and details of the simulations are reported in Appendix E.

5 Conclusion

We introduce a variational inference method to fit partially observed trajectories whose dynamic can be modeled by a continuous-time latent process, parameterized to be an interacting particle system. Our solution is an approximation to a multi-marginal Schrödinger bridge, that we obtain by first fitting an endpoint-conditioned model and then distilling it into an unconditional one. This methodology enables both trajectory reconstruction and prediction of future states. In future work we aim at testing our models on real data, comparing with state-of-the-art methods.

References

Michael Samuel Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *The Eleventh International Conference on Learning Representations*, 2022.

David Aldous. Interacting particle systems as stochastic social dynamics. *Bernoulli*, pages 1122–1149, 2013.

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.

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.

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.

Nicholas M Boffi and Eric Vanden-Eijnden. Deep learning probability flows and entropy production rates in active matter. *Proceedings of the National Academy of Sciences*, 121(25):e2318106121, 2024.

Mattia Bongini, Massimo Fornasier, Markus Hansen, and Mauro Maggioni. Inferring interaction rules from observations of evolutive systems i: The variational approach. *Mathematical Models and Methods in Applied Sciences*, 27(05):909–951, 2017.

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.

Maury Bramson and David Griffeath. Asymptotics for interacting particle systems on z d. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 53(2):183–196, 1980.

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.

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.

Ricky T. Q. Chen. torchdiffeq, 2018. URL `https://github.com/rtqichen/torchdiffeq`.

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

Yifan Chen, Mark Goldstein, Mengjian Hua, Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Probabilistic forecasting with stochastic interpolants and f\" ollmer processes. *arXiv preprint arXiv:2403.13724*, 2024.

Yongxin Chen, Giovanni Conforti, Tryphon T Georgiou, and Luigia Ripani. Multi-marginal schrödinger bridges. In *International Conference on Geometric Science of Information*, pages 725–732. Springer, 2019.

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.

Armand Comas, Yilun Du, Christian Fernandez Lopez, Sandesh Ghimire, Mario Sznaier, Joshua B Tenenbaum, and Octavia Camps. Inferring relational potentials in interacting systems. In *International Conference on Machine Learning*, pages 6364–6383. PMLR, 2023.

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

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34:17695–17709, 2021.

Sergey Dolgov and Dmitry Savostyanov. Tensor product approach to modelling epidemics on networks. *Applied Mathematics and Computation*, 460:128290, 2024.

J.R. Dormand and P.J. Prince. A family of embedded Runge-Kutta formulae. *Journal of Computational and Applied Mathematics*, 6:19–26, 1980.

Rick Durrett. Ten lectures on particle systems. *Lectures on Probability Theory: Ecole d'Eté de Probabilités de Saint-Flour XXIII—1993*, pages 97–201, 2006.

Jinchao Feng, Mauro Maggioni, Patrick Martin, and Ming Zhong. Learning interaction variables and kernels from observations of agent-based systems. *IFAC-PapersOnLine*, 55(30):162–167, 2022.

Pat Fitzsimmons, Jim Pitman, and Marc Yor. Markovian bridges: construction, palm interpretation, and splicing. In *Seminar on Stochastic Processes, 1992*, pages 101–134. Springer, 1992.

Andrew Gelman, Aki Vehtari, Daniel Simpson, Charles C Margossian, Bob Carpenter, Yuling Yao, Lauren Kennedy, Jonah Gabry, Paul-Christian Bürkner, and Martin Modrák. Bayesian workflow. *arXiv preprint arXiv:2011.01808*, 2020.

Daniel T Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *The Journal of chemical physics*, 115(4):1716–1733, 2001.

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

Daniele Grattarola, Lorenzo Livi, and Cesare Alippi. Learning graph cellular automata. *Advances in Neural Information Processing Systems*, 34:20983–20994, 2021.

G Grinstein, C Jayaprakash, and Yu He. Statistical mechanics of probabilistic cellular automata. *Physical review letters*, 55(23):2527, 1985.

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.

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

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

Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparametrization with gumble-softmax. In *International Conference on Learning Representations (ICLR 2017)*, 2017.

John Kalkhof, Arlene Kühn, Yannik Frisch, and Anirban Mukhopadhyay. Frequency-time diffusion with neural cellular automata. *arXiv preprint arXiv:2401.06291*, 2024.

Beomseok Kang, Harshit Kumar, Minah Lee, Biswadeep Chakraborty, and Saibal Mukhopadhyay. Learning locally interacting discrete dynamical systems: Towards data-efficient and scalable prediction. In *Proceedings of the 6th Annual Learning for Dynamics and Control Conference*, volume 242, pages 1357–1369, 2024.

Matt J Keeling and Ken TD Eames. Networks and epidemic models. *Journal of the Royal Society Interface*, 2(4):295–307, 2005.

P Kidger. *On neural differential equations*. PhD thesis, University of Oxford, 2021.

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.

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

Christian Kümmerle, Mauro Maggioni, and Sui Tang. Learning transition operators from sparse space-time samples. *IEEE Transactions on Information Theory*, 2024.

Nicolas Lanchier. *Stochastic modeling*. Springer, 2017.

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

Quanjun Lang, Xiong Wang, Fei Lu, and Mauro Maggioni. Interacting particle systems on networks: joint inference of the network and the interaction kernel. *arXiv preprint arXiv:2402.08412*, 2024.

Hugo Lavenant, Stephen Zhang, Young-Heon Kim, and Geoffrey Schiebinger. Towards a mathematical theory of trajectory inference. *arXiv preprint arXiv:2102.09204*, 2021.

Christian Léonard. A survey of the schr\" odinger problem and some of its connections with optimal transport. *arXiv preprint arXiv:1308.0215*, 2013.

Zhongyang Li, Fei Lu, Mauro Maggioni, Sui Tang, and Cheng Zhang. On the identifiability of interaction functions in systems of interacting particles. *Stochastic Processes and their Applications*, 132:135–163, 2021.

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

Dominik Linzner and Heinz Koeppl. Cluster variational approximations for structure learning of continuous-time bayesian networks from incomplete data. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2022.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

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.

Yuxuan Liu, Scott G McCalla, and Hayden Schaeffer. Random feature models for learning interacting dynamical systems. *Proceedings of the Royal Society A*, 479(2275):20220835, 2023.

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.

Fei Lu, Mauro Maggioni, and Sui Tang. Learning interaction kernels in heterogeneous systems of agents from multiple trajectories. *Journal of Machine Learning Research*, 22(32):1–67, 2021.

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

Alexander Mordvintsev, Ettore Randazzo, Eyvind Niklasson, and Michael Levin. Growing neural cellular automata. *Distill*, 2020. URL `https://distill.pub/2020/growing-ca/`.

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

Rasmus Berg Palm, Miguel González Duque, Shyam Sudhakaran, and Sebastian Risi. Variational neural cellular automata. In *International Conference on Learning Representations*, 2022.

Philip E Paré, Carolyn L Beck, and Tamer Başar. Modeling, estimation, and analysis of epidemics over networks: An overview. *Annual Reviews in Control*, 50:345–360, 2020.

Stefano Peluchetti. Diffusion bridge mixture transports, schrödinger bridge problems and generative modeling. *Journal of Machine Learning Research*, 24(374):1–51, 2023.

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

L. C. G. Rogers and David Williams. *Diffusions, Markov Processes and Martingales*. Cambridge Mathematical Library. Cambridge University Press, 2 edition, 2000.

Salva Rühling Cachay, Bo Zhao, Hailey Joren, and Rose Yu. Dyffusion: A dynamics-informed diffusion model for spatiotemporal forecasting. *Advances in Neural Information Processing Systems*, 36, 2024.

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

Yunyi Shen, Renato Berlinghieri, and Tamara Broderick. Multi-marginal schr\" odinger bridges with iterative reference. *arXiv preprint arXiv:2408.06277*, 2024.

Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36, 2024.

Vignesh Ram Somnath, Matteo Pariset, Ya-Ping Hsieh, Maria Rodriguez Martinez, Andreas Krause, and Charlotte Bunne. Aligned diffusion schrödinger bridges. In *Uncertainty in Artificial Intelligence*, pages 1985–1995. PMLR, 2023.

Ella Tamir, Martin Trapp, and Arno Solin. Transport with support: Data-conditional diffusion bridges. *Transactions on Machine Learning Research*, 2023.

Mattie Tesfaldet, Derek Nowrouzezahrai, and Chris Pal. Attention-based neural cellular automata. *Advances in Neural Information Processing Systems*, 35:8174–8186, 2022.

Alexander Tong, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Kilian Fatras, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. In *ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems*, 2023.

Francisco Vargas, Pierre Thodoroff, Austen Lamacraft, and Neil Lawrence. Solving schrödinger bridges via maximum likelihood. *Entropy*, 23(9):1134, 2021.

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

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

Stephen Wolfram. Theory and applications of cellular automata. *World Scientific*, 1986.

N Wulff and J A Hertz. Learning cellular automaton dynamics with neural networks. *Advances in Neural Information Processing Systems*, 5, 1992.

Liu Yang, Constantinos Daskalakis, and George E Karniadakis. Generative ensemble regression: Learning particle dynamics from observations of ensembles with physics-informed deep generative models. *SIAM Journal on Scientific Computing*, 44(1):B80–B99, 2022.

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

A Notation

Let $\Omega_{[0,T]}$ be the space of $\mathcal{X}$-valued cadlag functions over a time interval $[0, T]$, and denote by $\mathcal{P}(\Omega_{[0,T]})$ the space of probability measure on the path space. For a path measure $Q \in \mathcal{P}(\Omega_{[0,T]})$, timesteps $v,\ t \in [0, T]$ s.t. $v > t$, and configurations $\mathbf{x},\ \mathbf{x}' \in \mathcal{X}$, we assume that time-marginal and transition probability measures are absolutely continuous w.r.t. the counting measure. Their Radon-Nikodym derivative can then be expressed by the probability mass function $q_t(\mathbf{x})$ and the transition probability $q_{s|t}(\mathbf{x}' \mid \mathbf{x})$. We denote by $\Omega_{(t_k,t_{k+1}]}$ time restrictions of $\Omega_{[0,T]}$ to $(t_k, t_{k+1}] \subseteq [0, T]$. We denote the index set of $N-$many elements as $[N] = 0, \ldots, N-1$. Denote the cartesian product $\times_{k\in[K]}\mathcal{Y}$ of observations at K times as $\mathcal{Y}^K$. Moreover, let $\mathcal{M}(\Omega_{[0,T]}) \subset \mathcal{P}(\Omega_{[0,T]})$ denote the space of Markovian probability measures on $\Omega_{[0,T]}$.

B Related work

B.1 Learning interacting particle systems

The dynamics of many physical systems can be described through the local interaction laws of their constituent components. This principle has inspired computational frameworks that directly parameterize these governing interactions, both deterministically and stochastically. A prime example is cellular automata [Wolfram, 1986, Grinstein et al., 1985]. Early developments focused on studying the emergence of global patterns from a fixed set of rules on the evolution of individual cells. The inverse problem —inferring such rules from observations— has been of historical interest in the machine learning community [Wulff and Hertz, 1992, Mordvintsev et al., 2020], with recent developments incorporating attention-based architectures, graph neural networks, and black-box variational inference [Tesfaldet et al., 2022, Kang et al., 2024, Grattarola et al., 2021, Palm et al., 2022]. Models that learn interaction rules find applications across many domains, including physical simulators, multi-agent dynamics, dynamic graphs, as well as deep generative modeling [Kalkhof et al., 2024].

Within this context, most existing methods have proposed iterative updating schemes by parameterizing transition rules in discrete time. Interacting particle systems (IPSs) offer an alternative mathematical formalism that extends cellular automata to continuous time. Interacting particle systems are structured CTMCs whose states evolve with dependence on neighbors within a topology, typically established by a graph. Lanchier [2017] provides a modern introduction to this field. Classical literature focused on systems with finite states and often countably many sites [Bramson and Griffeath, 1980, Liggett, 1985, Durrett, 2006], while more recent work has focused on systems with finitely many sites [Aldous, 2013]. These systems have found applications in multi-agent modeling [Comas et al., 2023] and have been extended to systems of coupled stochastic differential equations (SDEs) in Euclidean space. This extension has seen increased attention recently [Lu et al., 2021, Yang et al., 2022, Feng et al., 2022, Liu et al., 2023, Lang et al., 2024, Kümmerle et al., 2024, Boffi and Vanden-Eijnden, 2024]. The learnability and identifiability of interaction rules in these systems have also been explored [Bongini et al., 2017, Li et al., 2021].

B.2 Inference for CTMCs

Inference methods for Markov jump processes (MJPs) have been extensively studied. Maximum likelihood estimation for time-homogeneous MJPs is discussed in Jackson [2011], Bladt and Sørensen [2005], McGibbon and Pande [2015]. Expectation-maximization techniques for continuous-time hidden Markov models have been developed in Liu et al. [2015], and an overview of the topic can be found in Wang [2021]. Bayesian approaches include Markov chain Monte Carlo methods [Boys et al., 2008, Hobolth and Stone, 2009, Rao and Teh, 2013, Golightly and Sherlock, 2019] and variational methods. The latter include mean-field [Opper and Sanguinetti, 2007, Cohn et al., 2010], moment-based methods [Wildner and Koeppl, 2019], combinations with MCMC [Zhang et al., 2017], and extensions to hybrid processes [Köhs et al., 2021]. Novel methods include black-box variational inference with neural networks [Seifner and Sánchez, 2023] and foundation models (i.e., meta-learning) [Berghaus et al., 2024]. While less directly related, it's worth noting recent work on guidance and conditioning for Markovian bridges [Corstanje et al., 2023] and discrete flow matching and diffusion methods [Campbell et al., 2022, Igashov et al., 2023, Lou et al., 2023, Campbell et al., 2024].

B.3 Trajectory Inference

Trajectory inference is a crucial component of our work, with connections to several recent developments. The Schrödinger bridge (SB) problem with multi-marginal constraints has been explored by Chen et al. [2019], Lavenant et al. [2021]. Recent advances in SB methods with a source and a target are presented in Vargas et al. [2021] and De Bortoli et al. [2021], with extensions to the multi-marginal setting by Shen et al. [2024]. Our approach shares similarities with Somnath et al. [2023], Shi et al. [2024], and Peluchetti [2023] in that it relies on samples from couplings solving the static SB problem. However, our methodology differs in that we learn the Markovian bridge and recover the unconditional path measure by distillation, rather than relying on closed-form endpoint-conditioned diffusions. The concept of Markov bridge by interpolation with a fictitious dynamic, as proposed by Igashov et al. [2023], is related to stochastic interpolants [Albergo and Vanden-Eijnden, 2022, Tong et al., 2023, Lipman et al., 2022, Liu et al., 2022] for probabilistic forecasting [Chen et al., 2024]. Ad-hoc variants for dynamical systems have also been developed [Rühling Cachay et al., 2024]. Our methodology also shares connections with flow matching using Gaussian process and Kalman filter interpolants [Tamir et al., 2023], in the fact that we are interested in *model-based* interpolants in a Bayesian framework.

C Proofs

C.1 Markovian bridges

Consider a sequence of observations $\{\mathbf{y}_k\}_{k\in[K]} \in \mathcal{Y}^K$ recorded at times $\{t_k\}_{k\in[K]} \in \mathbb{R}^K$, and assume conditional independence with respect to a Markov process $(\mathbf{x}_t)_{t\in[0,T]}$. For $t \in [t_0, t_{K-1}]$, let $\mathbf{y}_{>t} = \{\mathbf{y}_k \,|\, t_k > t,\, k = 1,\dots,K\}$ and $\mathbf{y}_{\leq t} = \{\mathbf{y}_k \,|\, t_k \leq t,\, k = 1,\dots,K\}$. The next observation after t is at time $t' := \min\{t_k : t_k > t,\, k = 1,\dots,K\}$, and we assume $t + h < t'$ for $h \approx 0$, by right-continuity of the transition probabilities. We can then denote the conditional transition rates for $\mathbf{x}' \neq \mathbf{x}$ as

$$\begin{aligned}
\Lambda(\mathbf{x}' \mid \mathbf{x},\, \mathbf{y}_{0:K}) &= \lim_{h\downarrow 0} h^{-1} \left[\mathbb{P}(\mathbf{x}_{t+h} = \mathbf{x}' \mid \mathbf{x}_t = \mathbf{x}, \mathbf{y}_{0:K})\right] \\
&= \lim_{h\downarrow 0} h^{-1} \left[\frac{\mathbb{P}(\mathbf{x}_{t+h} = \mathbf{x}', \mathbf{x}_t = \mathbf{x}, \mathbf{y}_{>t} \mid \mathbf{y}_{\leq t})}{\mathbb{P}(\mathbf{x}_t = \mathbf{x}, \mathbf{y}_{>t} \mid \mathbf{y}_{\leq t})}\right] \\
&= \lim_{h\downarrow 0} h^{-1} \left[\frac{\mathbb{P}(\mathbf{y}_{>t+h} \mid \mathbf{x}_{t+h} = \mathbf{x}')\mathbb{P}(\mathbf{x}_{t+h} = \mathbf{x}' \mid \mathbf{x}_t = \mathbf{x})\mathbb{P}(\mathbf{x}_t = \mathbf{x} \mid \mathbf{y}_{\leq t})}{\mathbb{P}(\mathbf{y}_{>t} \mid \mathbf{x}_t = \mathbf{x})\mathbb{P}(\mathbf{x}_t = \mathbf{x} \mid \mathbf{y}_{\leq t})}\right] \\
&= \lim_{h\downarrow 0} h^{-1} \left[\frac{\mathbb{P}(\mathbf{y}_{t'} \mid \mathbf{x}_{t+h} = \mathbf{x}')\mathbb{P}(\mathbf{x}_{t+h} = \mathbf{x}' \mid \mathbf{x}_t = \mathbf{x})}{\mathbb{P}(\mathbf{y}_{t'} \mid \mathbf{x}_t = \mathbf{x})}\right] \\
&= \Lambda_t(\mathbf{x}' \mid \mathbf{x}, t)\frac{\mathbb{P}(\mathbf{y}_{t'} \mid \mathbf{x}_t = \mathbf{x}')}{\mathbb{P}(\mathbf{y}_{t'} \mid \mathbf{x}_t = \mathbf{x})},
\end{aligned}$$

and similarly

$$\Lambda(\mathbf{x} \mid \mathbf{x},\, \mathbf{y}_{0:K}) = -\sum_{\mathbf{x}'\neq\mathbf{x}} \Lambda_t(\mathbf{x}' \mid \mathbf{x}, t)\frac{\mathbb{P}(\mathbf{y}_{t'} \mid \mathbf{x}_t = \mathbf{x}')}{\mathbb{P}(\mathbf{y}_{t'} \mid \mathbf{x}_t = \mathbf{x})} = \Lambda_t(\mathbf{x} \mid \mathbf{x}, t)\frac{1 - \mathbb{P}(\mathbf{y}_{t'} \mid \mathbf{x}_t = \mathbf{x})}{\mathbb{P}(\mathbf{y}_{t'} \mid \mathbf{x}_t = \mathbf{x})}. \quad (12)$$

C.2 Derivation of $D_{\mathrm{KL}}(Q \,||\, P)$

Consider two CTMCs with path measures $Q,\, P \in \mathcal{P}(\Omega_{[0,T]})$, and denote their respective rate matrices entries with $\Lambda_t(\mathbf{x}' \mid \mathbf{x})$ and $\Psi_t(\mathbf{x}' \mid \mathbf{x})$ for $\mathbf{x},\, \mathbf{x}' \in \mathcal{X}$. Their KL divergence, as discussed in Opper and Sanguinetti [2007], Seifner and Sánchez [2023], can be derived from the limit of discrete-time transitions with step size $h \coloneqq T/K$ as

$$
\begin{aligned}
&D_{\mathrm{KL}}(Q||P) \\
&= \lim_{K\to\infty} \sum_{\mathbf{x}_{0:K}} q_0(\mathbf{x}_0) \prod_{k=0}^{K-1} q_{k+h|k}(\mathbf{x}_{k+h} \mid \mathbf{x}_k) \log \frac{q_0(\mathbf{x}_0)\prod_{k=0}^{K-1} q_{k+h|k}(\mathbf{x}_{k+h} \mid \mathbf{x}_k)}{p_0(\mathbf{x}_0)\prod_{k=0}^{K-1} p_{k+h|k}(\mathbf{x}_{k+h} \mid \mathbf{x}_k)} \\
&= \sum_{\mathbf{x}_0} q_0(\mathbf{x}_0) \log \frac{q_0(\mathbf{x}_0)}{p_0(\mathbf{x}_0)} + \lim_{K\to\infty} \sum_{k=0}^{K-1} \mathbb{E}_{q_t(\mathbf{x})} \left[\sum_{\mathbf{x}_{k+h}} q_{k+h|k}(\mathbf{x}_{k+h} \mid \mathbf{x}_k) \log \frac{q_{k+h|k}(\mathbf{x}_{k+h} \mid \mathbf{x}_k)}{p_{k+h|k}(\mathbf{x}_{k+h} \mid \mathbf{x}_k)} \right] \qquad (13) \\
&= D_{\mathrm{KL}}(q_0||p_0) + \int_0^T \mathbb{E}_{q_t(\mathbf{x})} \sum_{\mathbf{x}'\neq\mathbf{x}} \left\{ \Psi_t(\mathbf{x}' \mid \mathbf{x}) + \Lambda_t(\mathbf{x}' \mid \mathbf{x}) \left(\log \frac{\Lambda_t(\mathbf{x}' \mid \mathbf{x})}{\Psi_t(\mathbf{x}' \mid \mathbf{x})} - 1 \right) \right\} dt,
\end{aligned}
$$

where the last line follows from dividing and multiplying each summand in (13) by h, and substituting the transition probabilities with rates,

$$
\frac{q_{k+h|k}(\mathbf{x}_{k+h} \mid \mathbf{x}_k)}{h} \log \frac{q_{k+h|k}(\mathbf{x}_{k+h} \mid \mathbf{x}_k)}{p_{k+h|k}(\mathbf{x}_{k+h} \mid \mathbf{x}_k)} \xrightarrow{h\to 0} \begin{cases} \Lambda_t(\mathbf{x}_{k+h} \mid \mathbf{x}_k) \log \frac{\Lambda_t(\mathbf{x}_{k+h}|\mathbf{x}_k)}{\Psi_t(\mathbf{x}_{k+h}|\mathbf{x}_k)} & \mathbf{x}_{k+h} \neq \mathbf{x}_k, \\ \sum_{\mathbf{x}'\neq\mathbf{x}} [\Psi_t(\mathbf{x}' \mid \mathbf{x}) - \Lambda_t(\mathbf{x}' \mid \mathbf{x})] & \mathbf{x}_{k+h} = \mathbf{x}_k. \end{cases}
$$

By assuming:

- Site-wise factorization of the time marginals $q_t(\mathbf{x}) = \prod_{i\in V} q_t^i(\mathbf{x}(i))$,
- Coupled transitions $q_{t+h|t}(\mathbf{x}' \mid \mathbf{x}) = \prod_{i\in V} q_{t+h|t}(\mathbf{x}'(i) \mid \mathcal{N}_t^i(\mathbf{x}))$ where we define a neighborhood $\mathcal{N}_t^i(\mathbf{x}) \coloneqq \{\mathbf{x}(i), \mathbf{x}(j) : i \sim j\}$,

we can rewrite each summand in (13) as

$$
\begin{aligned}
&\mathbb{E}_{q_k(\mathbf{x}_k)} \left[\sum_{\mathbf{x}_{k+h}} q_{k+h|k}(\mathbf{x}_{k+h} \mid \mathbf{x}_k) \log \frac{q_{k+h|k}(\mathbf{x}_{k+h} \mid \mathbf{x}_k)}{p_{k+h|k}(\mathbf{x}_{k+h} \mid \mathbf{x}_k)} \right] \\
&\quad = \mathbb{E}_{q_k(\mathbf{x}_k)} \left[\sum_{i\in V} q^i_{k+h,\,k}(\mathbf{x}_{k+h}(i) \mid \mathcal{N}_t^i(\mathbf{x}_k)) \log \frac{q^i_{k+h,\,k}(\mathbf{x}_{k+h}(i) \mid \mathcal{N}_t^i(\mathbf{x}_k))}{p^i_{k+h,\,k}(\mathbf{x}_{k+h}(i) \mid \mathcal{N}_t^i(\mathbf{x}_k))} \right].
\end{aligned}
$$

Letting $K \to \infty$ and plugging (3), we get

$$
\begin{aligned}
&D_{\mathrm{KL}}(Q||P) \\
&= D_{\mathrm{KL}}(q_0||p_0) + \int_0^T \mathbb{E}_{q_t(\mathbf{x}_t)} \sum_{i\in V} \sum_{s\neq\mathbf{x}_t(i)} \Big\{ \psi_t^{\mathbf{x}_t(i)\to s}(i, \mathbf{x}_t) - \lambda_t^{\mathbf{x}_t(i)\to s}(i, \mathbf{x}_t) \\
&\qquad\qquad + \lambda_t^{\mathbf{x}_t(i)\to s}(i, \mathbf{x}_t) \left(\log \frac{\lambda_t^{\mathbf{x}_t(i)\to s}(i, \mathbf{x}_t)}{\psi_t^{\mathbf{x}_t(i)\to s}(i, \mathbf{x}_t)} \right) \Big\} dt. \qquad (14)
\end{aligned}
$$

Moreover, notice that if we let p_0 be a uniform distribution the KL between initial distributions reduces to $D_{\mathrm{KL}}(q_0||p_0) = H(q_0) - \log(|\mathcal{X}|)$, where $H(\cdot)$ is the entropy.

C.3 Proof of Proposition 1

The *additive property* of the KL divergence [Léonard, 2013] states that for a Polish space $\mathcal{Q} \coloneqq \mathcal{Y}^K \times \Omega_{[0,T]}$, the canonical projecting onto the trajectory coordinates $\phi : \mathcal{Q} \to \mathcal{Y}^K$, a measurable mapping $\mathbf{y}_{0:K-1} \in \mathcal{Y}^K \mapsto Q_{\cdot|\mathbf{y}_{0:K-1}} \in \mathcal{P}(\Omega_{[0,T]})$, and $\mathrm{Q}, \mathrm{P} \in \mathcal{P}(\mathcal{Q})$, we get

$$
D_{\mathrm{KL}}(\mathrm{Q} \,||\, \mathrm{P}) = D_{\mathrm{KL}}(\phi_\# \mathrm{Q} \,||\, \phi_\# \mathrm{P}) + \int_{\mathcal{Y}^K} D_{\mathrm{KL}}(Q_{\cdot|\mathbf{y}_{0:K-1}} \,||\, P_{\cdot|\mathbf{y}_{0:K-1}}) \phi_\# \mathrm{Q}(d\mathbf{y}_{0:K-1}), \qquad (15)
$$

where $\phi_{\#}\mathrm{Q}(A) = \mathrm{Q}(\phi^{-1}(A))$ for any Borel set $A \subseteq \mathcal{Y}^K$ denotes a pushforward measure. Denote by $q_{t_{0:K-1}}(\mathbf{y}_{0:K-1})$ the p.m.f. associated to $q_{t_{0:K-1}} \coloneqq \phi_{\#}\mathrm{Q}$, and $p_{t_{0:K-1}}(\mathbf{y}_{0:K-1})$ the p.m.f. for $p_{t_{0:K-1}} \coloneqq \phi_{\#}\mathrm{P}$. For clarity, we rewrite (15) as

$$D_{\mathrm{KL}}(\mathrm{Q}\,||\,\mathrm{P}) = \mathbb{E}_{q_{t_{0:K-1}}(\mathbf{y}_{0:K-1})}\left[\log\frac{q_{t_{0:K-1}}(\mathbf{y}_{0:K-1})}{p_{t_{0:K-1}}(\mathbf{y}_{0:K-1})} + D_{\mathrm{KL}}(Q_{\cdot|\mathbf{y}_{0:K-1}}\,||\,P_{\cdot|\mathbf{y}_{0:K-1}})\right]. \quad (16)$$

Now consider the canonical projection $\varphi : \mathcal{Q} \to \Omega_{[0,T]}$, and assume that $P \coloneqq \varphi_{\#}\mathrm{P} \in \mathcal{M}(\Omega_{[0,T]})$. It follows from Léonard [2013, Prop. 2.10] that $Q^* \coloneqq \varphi_{\#}\mathrm{Q}^* \in \mathcal{M}(\Omega_{[0,T]})$, where Q^* is the solution to (4). Hence, without loss of generality we restrict our analysis to measures Q s.t. $Q \coloneqq \varphi_{\#}\mathrm{Q} \in \mathcal{M}(\Omega_{[0,T]})$.

Moreover, consider the case where Q and P are fully specified by their $\mathcal{Y}^K-$projection and disintegrations, such that

$$d\mathrm{Q} = dq_{\cdot|\mathbf{y}_0}(\mathbf{x}_0)dQ_{\cdot|\mathbf{y}_0,\mathbf{y}_1}((\mathbf{x}_t)_{t\in(0,1]})\dots dQ_{\cdot|\mathbf{y}_{T-2},\mathbf{y}_{T-1}}((\mathbf{x}_t)_{t\in(K-2,K-1]})dq_{t_{0:K-1}}(\mathbf{y}_{0:K-1})$$

for any Borel set $B \in \mathcal{Q}$, and the same for P. This property can be described as conditional independence of $\mathbf{y}_{0:K-1}$ given $(\mathbf{x}_t) \in \Omega_{[0,T]}$. Then, we can decompose (16) as

$$D_{\mathrm{KL}}(\mathrm{Q}\,||\,\mathrm{P}) = J_1 + J_2,$$

where

$$J_1 = \mathbb{E}_{q_{t_{0:K-1}}(\mathbf{y}_{0:K-1})}\left[\log\frac{q_{t_{0:K-1}}(\mathbf{y}_{0:K-1})}{p_{t_{0:K-1}}(\mathbf{y}_{0:K-1})}\right],$$
$$J_2 = \mathbb{E}_{q_{t_0}}\left[D_{\mathrm{KL}}(q_{\cdot|\mathbf{y}_0}\,||\,p_{\cdot|\mathbf{y}_0})\right] + \sum_{k\in[K]}\mathbb{E}_{q_{t_k,t_{k+1}}}\left[D_{\mathrm{KL}}(Q_{\cdot|\mathbf{y}_k,\mathbf{y}_{k+1}}\,||\,P_{\cdot|\mathbf{y}_k,\mathbf{y}_{k+1}})\right].$$

If we have access to the coupling $\pi_{t_{0:K-1}}$ solving (5), then the only term left depending on Q is J_2, hence we recover (6).

C.4 Derivation of $\mathcal{L}_1$

For ease of illustration, we start by considering a loss $\mathcal{L}_1$ with a single component defined in the time frame $[0,T]$ between two observations $\mathbf{y}_0$ and $\mathbf{y}_T$. We might be interested in parameterizing emission distributions $p_{t|t}(\mathbf{y}\,|\,\mathbf{x})$, hence we denote them in short as $p_t(\mathbf{y}\,|\,\mathbf{x})$. We are interested in proving that

$$\begin{aligned}&D_{\mathrm{KL}}\left(Q\,||\,P_{\cdot|\mathbf{y}_0,\mathbf{y}_T}\right)\\&\quad= D_{\mathrm{KL}}\left(Q\,||\,P\right) - \mathbb{E}_{\mathbf{x}\sim q_0}\left[\log p_0(\mathbf{y}_0\mid\mathbf{x})\right] - \mathbb{E}_{\mathbf{x}\sim q_T}\left[\log p_T(\mathbf{y}_T\mid\mathbf{x})\right] + \log p_{0,T}(\mathbf{y}_0,\mathbf{y}_T).\end{aligned}$$

We now consider a time-discretization at $0 = \tau_0 < \cdots < \tau_{K-1} = T$ of the path measures Q and $P_{\cdot|\mathbf{y}_0,\mathbf{y}_T}$. For $\tau_{k+1},\,\tau_k \in [0,T]$ and $\mathbf{x}_{k+1},\mathbf{x}_k \in \mathcal{X}$, by the Markov property of $(\mathbf{x}_t)_{t\in[0,T]}$ under P and conditional independence of $\mathbf{y}_0,\,\mathbf{y}_T$ given $\mathbf{x}_k$, we can express the marginal and transition probability mass functions conditionally on $\mathbf{y}_0,\,\mathbf{y}_T$ as

$$\bar{p}_{\tau_k}(\mathbf{x}_k) = p_{\tau_k}(\mathbf{x}_k)\frac{p_{0|\tau_k}(\mathbf{y}_0\mid\mathbf{x}_k)p_{T|\tau_k}(\mathbf{y}_T\mid\mathbf{x}_k)}{p_{0,T}(\mathbf{y}_0,\mathbf{y}_T)}, \quad (17)$$

$$\bar{p}_{\tau_{k+1}|\tau_k}(\mathbf{x}_{k+1}\mid\mathbf{x}_k) = p_{\tau_{k+1}|\tau_k}(\mathbf{x}_{k+1}\mid\mathbf{x}_k)\frac{p_{T|\tau_{k+1}}(\mathbf{y}_T\mid\mathbf{x}_{k+1})}{p_{T|\tau_k}(\mathbf{y}_T\mid\mathbf{x}_k)}. \quad (18)$$

Notice that

$$\bar{p}_{\tau_0}(\mathbf{x}_0)\prod_{k=0}^{K-2}\bar{p}_{\tau_{k+1}|\tau_k}(\mathbf{x}_{k+1}\mid\mathbf{x}_k) = p_{\tau_0}(\mathbf{x}_0)\prod_{k=0}^{K-2}p_{\tau_{k+1}|\tau_k}(\mathbf{x}_{k+1}\mid\mathbf{x}_k)\frac{p_0(\mathbf{y}_0\mid\mathbf{x}_0)p_T(\mathbf{y}_T\mid\mathbf{x}_T)}{p_{0,T}(\mathbf{y}_0,\mathbf{y}_T)}.$$

Notice that every marginal and transition probability derived from the variational path measure $Q \in \mathcal{P}(\Omega_{[0,T]})$ can depend on the observations $\mathbf{y}_0,\,\mathbf{y}_1$, but we omit them from the notation for brevity.

We can then write $q_{\tau_0:\tau_T}(\mathbf{x}_{0:T}) = q_{\tau_0}(\mathbf{x}_0)\prod_{k=0}^{K-2} q_{\tau_{k+1}|\tau_k}(\mathbf{x}_{k+1} \mid \mathbf{x}_k)$, and it follows that

$$\begin{aligned}
&D_{\mathrm{KL}}\left(Q \,||\, P_{\cdot|\mathbf{y}_0,\mathbf{y}_T}\right) \\
&= \mathbb{E}_{q_{\tau_0:\tau_T}}\left[\log \frac{q_{\tau_0:\tau_K}(\mathbf{x}_{0:T})}{\bar{p}_{\tau_0}(\mathbf{x}_0)\prod_{k=0}^{K-2}\bar{p}_{\tau_{k+1}|\tau_k}(\mathbf{x}_{k+1} \mid \mathbf{x}_k)}\right] \\
&= \mathbb{E}_{q_{\tau_0:\tau_K}}\left[\log \frac{q_{\tau_0:\tau_K}(\mathbf{x}_{0:T})}{p_{\tau_0}(\mathbf{x}_0)\prod_{k=0}^{K-2}p_{\tau_{k+1}|\tau_k}(\mathbf{x}_{k+1} \mid \mathbf{x}_k)} - \log \frac{p_0(\mathbf{y}_0 \mid \mathbf{x}_0)p_T(\mathbf{y}_T \mid \mathbf{x}_T)}{p_{0,T}(\mathbf{y}_0,\mathbf{y}_T)}\right] \\
&= D_{\mathrm{KL}}\left(Q \,||\, P\right) - \mathbb{E}_{q_{\tau_0}(\mathbf{x}|\mathbf{y}_0)}\left[\log p_0(\mathbf{y}_0 \mid \mathbf{x}_0)\right] - \mathbb{E}_{q_{\tau_T}(\mathbf{x}|\mathbf{y}_T)}\left[\log p_T(\mathbf{y}_T \mid \mathbf{x}_T)\right] + \log p_{0,T}(\mathbf{y}_0,\mathbf{y}_T).
\end{aligned} \tag{19}$$

This formulation can trivially be extended for a sequence of observations $\{\mathbf{y}_k\}_{k\in[K]}$ at times $0 = t_1 < \cdots < t_{T-1} = T$, resulting in

$$D_{\mathrm{KL}}\left(Q \,||\, P_{\cdot|\{\mathbf{y}_k\}_{k\in[K]}}\right) = D_{\mathrm{KL}}\left(Q \,||\, P\right) - \sum_{k\in[K]} \mathbb{E}_{q_{t_k}(\mathbf{x}_{t_k}|\boldsymbol{y}_k),\pi_{t_k}(\boldsymbol{y}_k)}[\log p_{t_k}(\mathbf{y}_k \mid \mathbf{x}_{t_k})] + \log Z, \tag{20}$$

where $Z = p_{t_{0:K-1}}(\{\mathbf{y}_k\}_{k\in[K]})$. Taking the expectation w.r.t. $\{\mathbf{y}_k\}_{k\in[K]} \sim \pi_{t_{0:K-1}}$ on both sides recovers (6) on the LHS, and (7) on the RHS plus a term independent of Q. In practice, we might be interested in learning the emission distribution at the same time as we are learning the variational path measure, in which case we can interpret the objective in (7) as an evidence lower bound.

C.5 Derivation of $\mathcal{L}_2$

Once again, for ease of illustration we consider two observations $\mathbf{y}_0$, $\mathbf{y}_T$ at the endpoints of a time interval $[0, T]$. We are interested in learning the generator of a CTMC $(\mathbf{x}_t)_{t\in[0,T]}$ by approximating it with a neural model $\Lambda_t^\phi(\mathbf{x}' \mid \mathbf{x})$. We are only given access to its endpoint-conditioned generator $\Lambda_t(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T)$, and by Bayes rule we can recover the unconditional one as

$$\Lambda_t(\mathbf{x}' \mid \mathbf{x}) = \int_{\mathcal{Y}\times\mathcal{Y}} \Lambda_t(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T) q_{0,T|t}(\mathbf{y}_0, \mathbf{y}_T \mid \mathbf{x}) d\mathbf{y}_0 d\mathbf{y}_1 \tag{21}$$

$$= \int_{\mathcal{Y}\times\mathcal{Y}} \Lambda_t(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T) \frac{q_{t|0,T}(\mathbf{x} \mid \mathbf{y}_0, \mathbf{y}_T)\pi_{0,T}(\mathbf{y}_0, \mathbf{y}_T)}{q_t(\mathbf{x})} d\mathbf{y}_0 d\mathbf{y}_1. \tag{22}$$

First, we want to show that the intractable unconditional loss

$$\mathcal{L}_U(\phi) := \int_0^T \mathbb{E}_{q_t(\mathbf{x})} \sum_{\mathbf{x}'\neq\mathbf{x}} \left\|\Lambda_t(\mathbf{x}' \mid \mathbf{x}) - \Lambda_t^\phi(\mathbf{x}' \mid \mathbf{x})\right\|_2^2 dt \tag{23}$$

is equivalent to the tractable conditional one

$$\mathcal{L}_C(\phi) := \int_0^T \mathbb{E}_{\mathbf{y}_0,\mathbf{y}_T\sim\pi_{0,T},\, \mathbf{x}\sim q_{t|0,T}(\cdot|\mathbf{y}_0,\mathbf{y}_T)} \sum_{\mathbf{x}'\neq\mathbf{x}} \left\|\Lambda_t(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T) - \Lambda_t^\phi(\mathbf{x}' \mid \mathbf{x})\right\|_2^2 dt \tag{24}$$

up to a constant independent of ϕ. This proof is not novel, as it mirrors the proof of Lipman et al. [2022, Theorem 2] adapted to transition rate matrices in discrete spaces rather than vector fields in Euclidean spaces.

First, notice that each component $\mathbb{E}_{q_t(\mathbf{x})} \left\|\Lambda_t(\mathbf{x}' \mid \mathbf{x}) - \Lambda_t^\phi(\mathbf{x}' \mid \mathbf{x})\right\|_2^2$ for $\mathbf{x}' \neq \mathbf{x}$ can be expressed as

$$\mathbb{E}_{q_t(\mathbf{x})} \left\|\Lambda_t(\mathbf{x}' \mid \mathbf{x})\right\|_2^2 - 2\mathbb{E}_{q_t(\mathbf{x})} \left\langle \Lambda_t(\mathbf{x}' \mid \mathbf{x}), \Lambda_t^\phi(\mathbf{x}' \mid \mathbf{x}) \right\rangle + \mathbb{E}_{q_t(\mathbf{x})} \left\|\Lambda_t^\phi(\mathbf{x}' \mid \mathbf{x})\right\|_2^2. \tag{25}$$

Of the two summands that depend on ϕ, plugging (21) into the first term yields

$$\begin{aligned}
&\mathbb{E}_{q_t(\mathbf{x})}\left\langle \Lambda_t(\mathbf{x}' \mid \mathbf{x}), \Lambda_t^{\phi}(\mathbf{x}' \mid \mathbf{x})\right\rangle \\
&= \int_{\mathcal{X}} \left\langle \int_{\mathcal{Y}\times\mathcal{Y}} \Lambda_t(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T) \frac{q_{t|0,T}(\mathbf{x} \mid \mathbf{y}_0, \mathbf{y}_T)\pi_{0,T}(\mathbf{y}_0, \mathbf{y}_T)}{q_t(\mathbf{x})} d\mathbf{y}_0 d\mathbf{y}_1, \Lambda_t^{\phi}(\mathbf{x}' \mid \mathbf{x})\right\rangle q_t(\mathbf{x}) d\mathbf{x} \\
&= \int_{\mathcal{X}} \int_{\mathcal{Y}\times\mathcal{Y}} \left\langle \Lambda_t(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T) q_{t|0,T}(\mathbf{x} \mid \mathbf{y}_0, \mathbf{y}_T)\pi_{0,T}(\mathbf{y}_0, \mathbf{y}_T), \Lambda_t^{\phi}(\mathbf{x}' \mid \mathbf{x})\right\rangle d\mathbf{y}_0 d\mathbf{y}_1 d\mathbf{x} \\
&= \mathbb{E}_{\pi_{0,T}(\mathbf{y}_0,\mathbf{y}_T),\, q_{t|0,T}(\mathbf{x}|\mathbf{y}_0,\mathbf{y}_T)} \left\langle \Lambda_t(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T), \Lambda_t^{\phi}(\mathbf{x}' \mid \mathbf{x})\right\rangle .
\end{aligned}$$

Next, it follows by the law of total expectation that the second term is

$$\mathbb{E}_{q_t(\mathbf{x})} \left\| \Lambda_t^{\phi}(\mathbf{x}' \mid \mathbf{x}) \right\|_2^2 = \mathbb{E}_{\pi_{0,T}(\mathbf{y}_0,\mathbf{y}_T),\, q_{t|0,T}(\mathbf{x}|\mathbf{y}_0,\mathbf{y}_T)} \left\| \Lambda_t^{\phi}(\mathbf{x}' \mid \mathbf{x}) \right\|_2^2 .$$

Combining the terms, it follows from the linearity of expectation that $\nabla_\phi \mathcal{L}_U(\phi) = \nabla_\phi \mathcal{L}_C(\phi)$. Next, we want to show that if we approximate the true conditional generator with an estimator $\Lambda_t^{\theta}(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T)$ such that $\mathbb{E}_{\pi_{0,T}(\mathbf{y}_0,\mathbf{y}_T),\, q_{t|0,T}(\mathbf{x}|\mathbf{y}_0,\mathbf{y}_T)}[\Lambda_t^{\theta}(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T)] = \Lambda_t(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T)$, we recover $\mathcal{L}_2$ as specified in (10) and a component independent of ϕ. To show this, add and subtract $\Lambda_t^{\theta}(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T)$ from $\mathcal{L}_C(\phi)$ and complete the square to get

$$\begin{aligned}
&\mathcal{L}_C(\phi) \\
&= \int_0^T \mathbb{E}_{\pi_{0,T},\, q_{t|0,T}} \sum_{\mathbf{x}' \neq \mathbf{x}} \left\| \Lambda_t^{\theta}(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T) - \Lambda_t^{\phi}(\mathbf{x}' \mid \mathbf{x}) \right\|_2^2 + \left\| \Lambda_t(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T) - \Lambda_t^{\theta}(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T) \right\|_2^2 dt \\
&= \int_0^T \mathbb{E}_{\pi_{0,T},\, q_{t|0,T}} \sum_{\mathbf{x}' \neq \mathbf{x}} \left\| \Lambda_t^{\theta}(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T) - \Lambda_t^{\phi}(\mathbf{x}' \mid \mathbf{x}) \right\|_2^2 dt + K_\theta .
\end{aligned}$$

In order to retrieve $\mathcal{L}_2(\phi)$, we perform a change of measure by importance sampling with proposal $q_t^{\theta}(\mathbf{x} \mid \mathbf{y}_0, \mathbf{y}_T)$, and approximate the importance weights $q_t/q_t^{\theta} \approx 1$ to get

$$\begin{aligned}
\mathcal{L}_C(\phi) &\propto \int_0^T \mathbb{E}_{\pi_{0,T},\, q_{t|0,T}} \sum_{\mathbf{x}' \neq \mathbf{x}} \left\| \Lambda_t^{\theta}(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T) - \Lambda_t^{\phi}(\mathbf{x}' \mid \mathbf{x}) \right\|_2^2 dt \\
&\approx \int_0^T \mathbb{E}_{\pi_{0,T},\, q^{\theta}_{t|0,T}} \sum_{\mathbf{x}' \neq \mathbf{x}} \left\| \Lambda_t^{\theta}(\mathbf{x}' \mid \mathbf{x}, \mathbf{y}_T) - \Lambda_t^{\phi}(\mathbf{x}' \mid \mathbf{x}) \right\|_2^2 dt = \mathcal{L}_2(\phi).
\end{aligned}$$

D Implementation details

D.1 Training

The training algorithms for the trajectory reconstruction and prediction tasks are reported in Algorithm 1 and Algorithm 2 respectively. Notice that it is also possible to learn the unconditional generator at the same time as the unconditional one, by freezing the gradients of θ before updating the $\mathcal{L}_2$ loss. For illustration purposes we do not explicit the numerical solver we are using, but for our experiments we use a Dormand-Prince solver of order 5 [Dormand and Prince, 1980] from the `torchdiffeq` library [Chen, 2018]. While all datapoints in a batch are processed in parallel, we might need to evolve the solver through different time points for each batch. This is feasible by applying the tricks for parallel solving of neural ODEs with varying time-intervals presented in Chen et al. [2021].

Overall, we found training of the unconditional model quite straightforward to implement. On the other hand, training the unconditional model seems to be quite challenging, mainly due to vanishing gradients when backpropagating through the solver. A trick we found quite helpful in addressing this problem is annealing the time discretization grid, from very coarse to a finer and finer one.

Moreover, if the prior tends to push the model towards "high activity", we found the states generated by the conditional model to converge to the next observed state pretty quickly. This can hamper the training of the unconditional model, as trajectories generated by the conditional model will converge very quickly and then stay still for a long amount of time. This biases the distribution of samples seen at training time by the conditional model, that might then experience ”mode collapse” and predict all of the transition rates to be zero. We found that choosing priors that bias the conditional model towards performing fewer transitions helps addressing this issue.

Algorithm 1 Training for trajectory reconstruction

Require: Dataset $\mathcal{D} = \{\{t_k, \mathbf{y}^i_k, C_{\mathcal{G}^i_k}\}_{k\in[K_i]}\}_{i\in[N]}$, prior path measure P, OPTIMIZER, step size ϵ
Ensure: Learned parameters θ, ξ, for conditional generator $\mathbf{\Lambda}^\theta_t$ and emission distribution p^ξ_t

```
Initialize parameters θ, ξ and loss L_1 = 0
while not converged do
    Sample minibatch of pairs (y_k, y_{k+1}) from D
    for each pair (y_k, y_{k+1}) do
        Encode x_{t_k} ~ q^θ_{t_k}(·|y_k) = Categorical(t_k, g^θ_{t_k}(i, y_k))
        for t ∈ (t_k, t_{k+1}] do
            Sample x_t from q^{i,θ}_t using Gumbel-Softmax trick
            Evolve q^{i,θ}_t using Λ^θ_t(x_t|y_{k+1}) as in (9)
        end for
        Compute L_1 ← L_1 + L^k_1(θ, ξ, y_k, y_{k+1}) using (8)
    end for
    Update θ, ξ ← OPTIMIZER(∇L_1(θ, ξ)).
end while
```

Algorithm 2 Training for prediction

Require: Dataset $\mathcal{D} = \{\{t_k, \mathbf{y}^i_k, C_{\mathcal{G}^i_k}\}_{k\in[K_i]}\}_{i\in[N]}$, conditional generator $\mathbf{\Lambda}^\theta_t$, OPTIMIZER, step size ϵ
Ensure: Learned parameters ϕ for unconditional generator $\mathbf{\Lambda}^\phi_t$

```
Initialize parameters φ and loss L_1 = 0
while not converged do
    Sample minibatch of pairs (y_k, y_{k+1}) from D
    for each pair (y_k, y_{k+1}) do
        Encode x_{t_k} ~ q^θ_{t_k}(·|y_k) = Categorical(t_k, g^θ_{t_k}(i, y_k))
        for t ∈ (t_k, t_{k+1}] do
            Sample x_t from q^{i,θ}_t using Gumbel-Softmax trick
            Evolve q^{i,θ}_t using Λ^θ_t(x_t|y_{k+1}) as in (9)
            Compute L_2 ← L_2 + ||Λ^θ_t(x_t|y_{k+1}) − Λ^φ_t(x_t)||_2^2
        end for
    end for
    Update φ ← OPTIMIZER(∇L_2(φ)).
end while
```

D.2 Sampling

For sampling, we compute transitions directly in sample space, in order to respect possible constraints encoded in the generator. Denoting a neural encoder as q_t and the entries of a generator $\mathbf{\Lambda}_t$ as Λ_t (regardless of it being conditional or unconditional), we initialize $\mathbf{x}_0 \sim q_0(\cdot \mid \mathbf{y}_0)$ and employ a first-order approximation of the transition probability as

$$q_{t+\epsilon|t}(\mathbf{x}' \mid \mathbf{x}) \approx \begin{cases} \epsilon\Lambda_t(\mathbf{x}' \mid \mathbf{x}), & \mathbf{x}' \neq \mathbf{x}, \\ 1 - \epsilon\sum_{\mathbf{x}'\neq\mathbf{x}} \Lambda_t(\mathbf{x}), & \mathbf{x}' = \mathbf{x}. \end{cases} \tag{26}$$

written in short as $\mathbf{x}_{t+\epsilon} \sim \text{Categorical}\left(\delta_{\mathbf{x}_t} + \epsilon\mathbf{\Lambda}_t(\mathbf{x}_t)\right)$. This method, despite being a crude approximation, is typically employed in discrete flow matching [Campbell et al., 2024] for its high scalability.

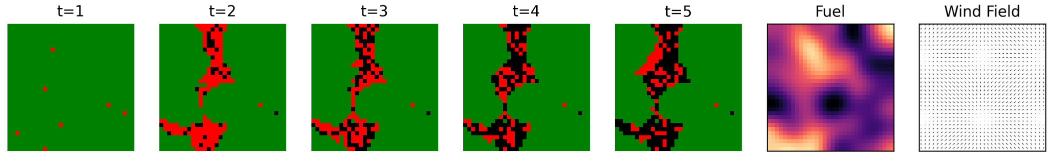

Figure 2: First 5 observations in time of a sequence from the wildfires dataset, with the corresponding covariates.

D.3 Alternative parameterization

The analytical form of the Markovian bridge in (2) suggests a potential unification of our two-stage approach into a single stage. This would involve fitting the conditional generator by minimizing the evidence lower bound in (7) and specifying the model as:

$$\mathbf{\Lambda}_t^\theta(\mathbf{x}_t \,|\, \mathbf{y}_{t_{k+1}}) = \mathbf{\Lambda}_t^\theta(\mathbf{x}_t)\boldsymbol{H}_t^\theta(\mathbf{y}_{t_{k+1}}, \mathbf{x}_t). \tag{27}$$

This approach is analogous to that taken by Somnath et al. [2023] for diffusion Schrödinger bridges, where they learn to approximate a target comprising a guidance term that ensures the diffusion process reaches the prescribed $\mathbf{y}_{k+1}$ at time t_{k+1}. However, their method relies on the availability of a closed-form for the endpoint-conditioned prior processes, which is not available in our setting. Both guidance approaches can be viewed as instances of Doob's h-transform [Rogers and Williams, 2000], and we refer the reader to Corstanje et al. [2023] for a detailed discussion. However, our preliminary experiments suggest that this single-stage approach may yield inferior unconditional models compared to the two-stage method, possibly due to identifiability issues. We leave a thorough investigation of single-stage methods to future work.

D.4 Computational considerations

Our method is not simulation-free, in the sense that learning is made possible by backpropagating through a solver. In doing so, a practitioner can incur in two fundamental problems, inaccurate gradients and memory-intensive training steps. The choice of a backpropagation technique can trade off one disadvantage for the other. In our experiments we use continuous adjoint methods, that provide memory-efficient numerical solutions (constant w.r.t. the time discretization grid) at the cost of incurring numerical errors that accumulate into potentially inaccurate gradient estimates. An overview of other possible approaches is presented in [Kidger, 2021].

E Experiments

E.1 Datasets

Epidemics The dataset is comprised of a collection of 250 random graphs with 128 nodes each and a given expected degree of 3, where edges are generated at random. Two covariates $\mathbf{c}_1^i$, $\mathbf{c}_2^i$ are generated for each node $i \in V$ by sampling from a standard normal distribution. An epidemic is then spread according to a Susceptible-Infected-Recovered (SIR) model [Keeling and Eames, 2005, Paré et al., 2020, Dolgov and Savostyanov, 2024]. Initially, all nodes are set to be susceptible (S) with the exception of p_0 nodes set to be infected (I) at random. Each graph in the dataset is evolved in the continuous-time interval $[0, 19]$, where a time-homogeneous functional form for the local transition rates from S to I and from I to recovered (R) is specified as

$$\begin{aligned} \lambda^{S\to I}(i, \mathbf{x}) &= \beta \exp\left(\sin(\mathbf{c}_1^i) + \cos(\mathbf{c}_2^i)\right) \left|\mathcal{N}_i^I\right|, \\ \lambda^{I\to R}(i, \mathbf{x}) &= \gamma, \end{aligned}$$

where $\mathcal{N}_i^I \coloneqq \{j \in V \,|\, \mathbf{x}(j) = I, j \sim i\}$, $\beta = 6$ and $\gamma = 0.2$. These parameters do not correspond to physically meaningful quantities, and adjusting them to reflect real-world spread dynamics remains an interesting avenue for future work. Each graph is observed at $K = 20$ regularly spaced time points, with no observation noise (i.e., $\mathcal{X} \equiv \mathcal{Y}$). The data is simulated using $\tau-$leaping [Gillespie, 2001], with $\tau = 1 \times 10^{-2}$. A sample observed in its first 5 time steps is displayed in Figure 3.

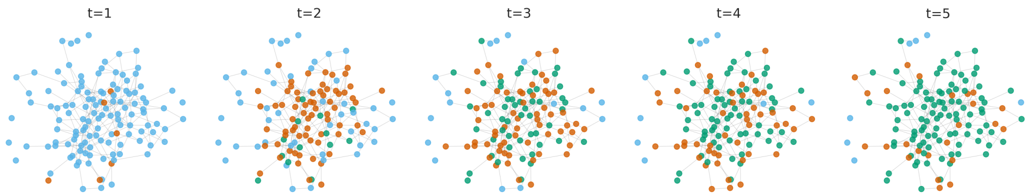

Figure 3: First 5 observations in time of a sequence from the epidemics dataset.

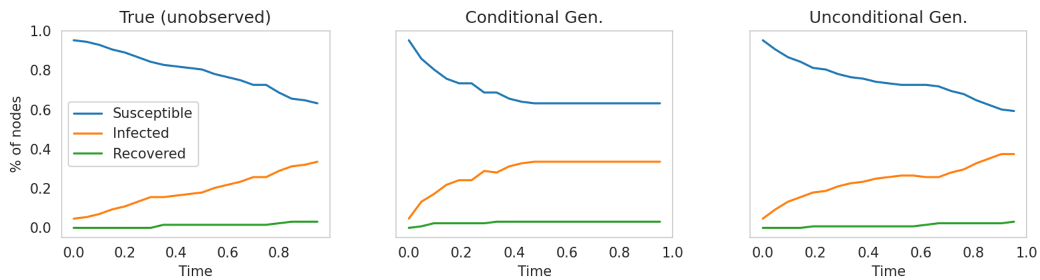

Figure 4: True and generated SIR curves in a time interval observed only at the two endpoints, in an held-out graph of 128 nodes.

Wildfires We consider 250 observations of 32^2-dimensional lattice-valued data represented as images, where each pixel can take three possible values: unburned (U), burning (B), or extinguished (E). Spatially structured covariates corresponding to wind fields $\mathbf{w}$ and ground-level fuel $\mathbf{f}$ are generated at the same resolution. At time zero, each pixel is set to B with a probability $p_0^B = 0.005$ (i.e., we expect 5 pixels to be burning), while all the others are set to U. The dynamic is then evolved in the continuous-time interval $[0, 19]$ by local transition rates with time-homogeneous functional forms

$$\lambda^{U \to B}(i, \mathbf{x}) = \text{ReLU}(a_0 + a_1 \mathbf{f}^i) \times \text{ReLU}\left(b_0 + b_1 \sum_{j \in \mathcal{N}_i^B} \mathbf{a}^{ij}\right),$$

$$\lambda^{E \to B}(i, \mathbf{x}) = \text{ReLU}(c_0 + c_1 \mathbf{f}^i) \times \text{ReLU}\left(d_0 + d_1 \sum_{j \in \mathcal{N}_i^B} \mathbf{a}^{ij}\right),$$

$$\lambda^{B \to E}(i, \mathbf{x}) = \gamma,$$

where $\mathcal{N}_i^B := \{j \in V \,|\, \mathbf{x}(j) = I, j \sim i\}$, and $\mathbf{a}^{ij}$ is a *wind alignment* value obtained by the dot product between the relative position of the neighbor j w.r.t. i and the value of the wind field at j. For our simulation, we set $a_0 = b_0 = c_0 = d_0 = 0.1$, $a_1 = 5$, $b_1 = 1$, $c_1 = d_1 = 0.01$, and $\gamma = 0.5$. Similarly to the first setting, each wildfire is observed at $K = 20$ regularly spaced time points with no observation noise. A sample observed in its first 5 time steps, as well as the related covariates, is displayed in Figure 2.

E.2 Model

Since there is no observation noise, all we need to parameterize in our experiments are the conditional and unconditional generators. Both can be thought of as mappings $\mathcal{X} \to \mathbb{R}_{\geq 0}^{|V| \times |S| \times |S|}$, i.e. the output shall be a local transition rate matrix at each site $i \in V$. In order to constrain the dependence structure of each local transition rate to the neighborhood of that site, we use the attention-based neural cellular automata presented in Tesfaldet et al. [2022]. This can be thought of as a depth-one vision transformer whose attention matrix is (efficiently) masked to attend only to a neighborhood of embedded input sites. For the wildfires experiment we simply consider a 3×3 Moore neighborhood, whereas for the epidemics we mask the attention matrix with the adjacency matrix of each observation. We obtain 512-dimensional input embeddings by a two-layer MLP of width 512. The

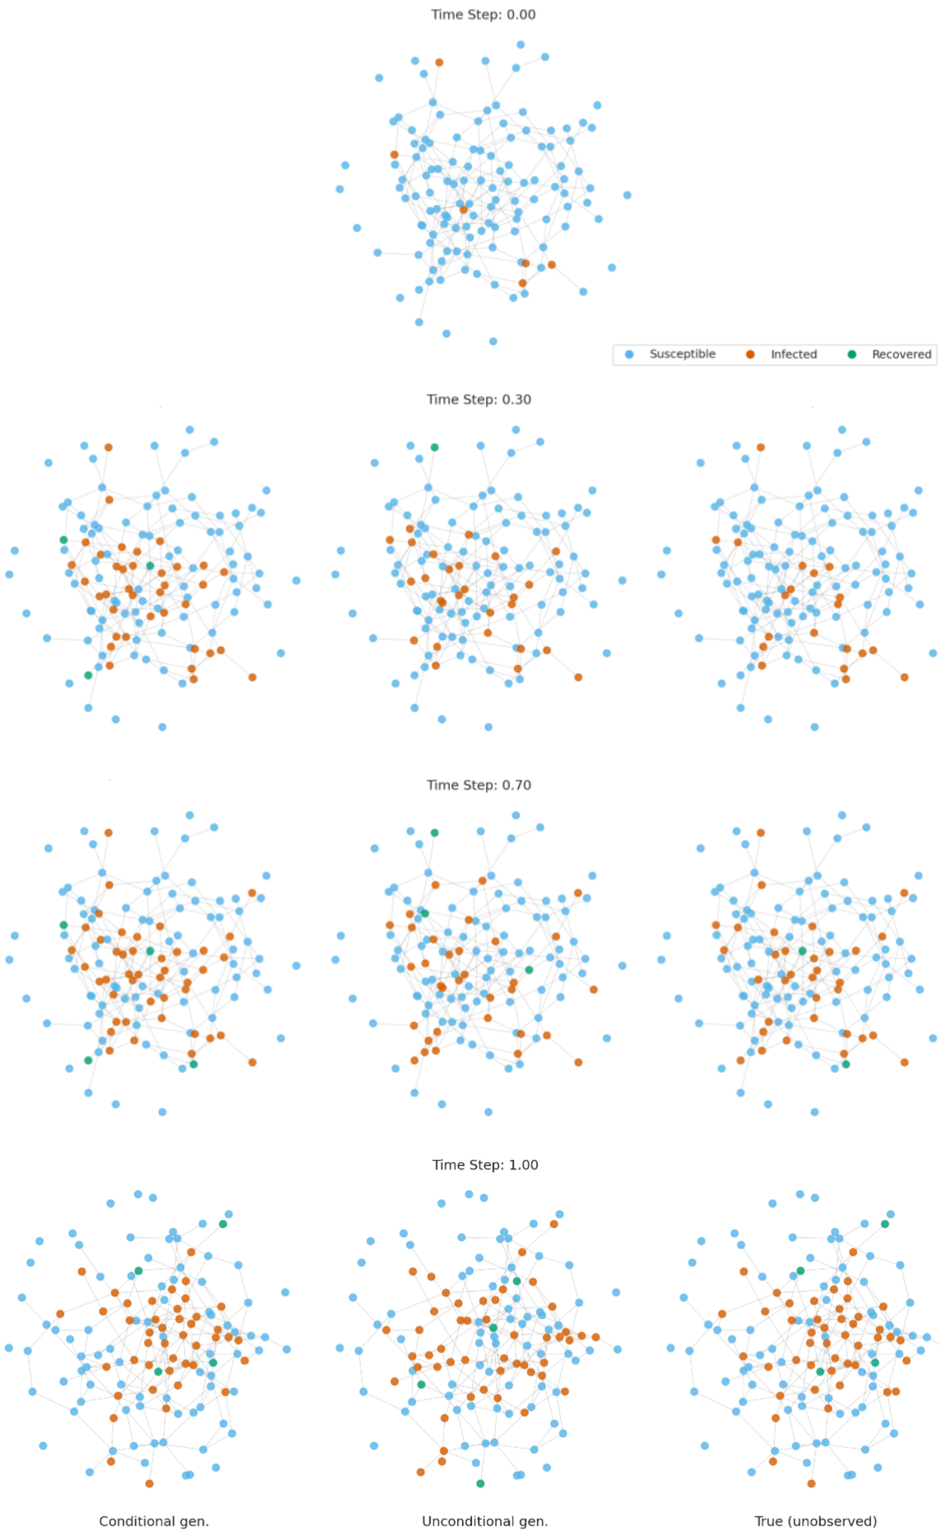

Figure 5: Evolution of an epidemic on an held-out graph. Endpoint-conditioned generation (left), unconditional generation (center), trajectory observed only at the endpoints (right).

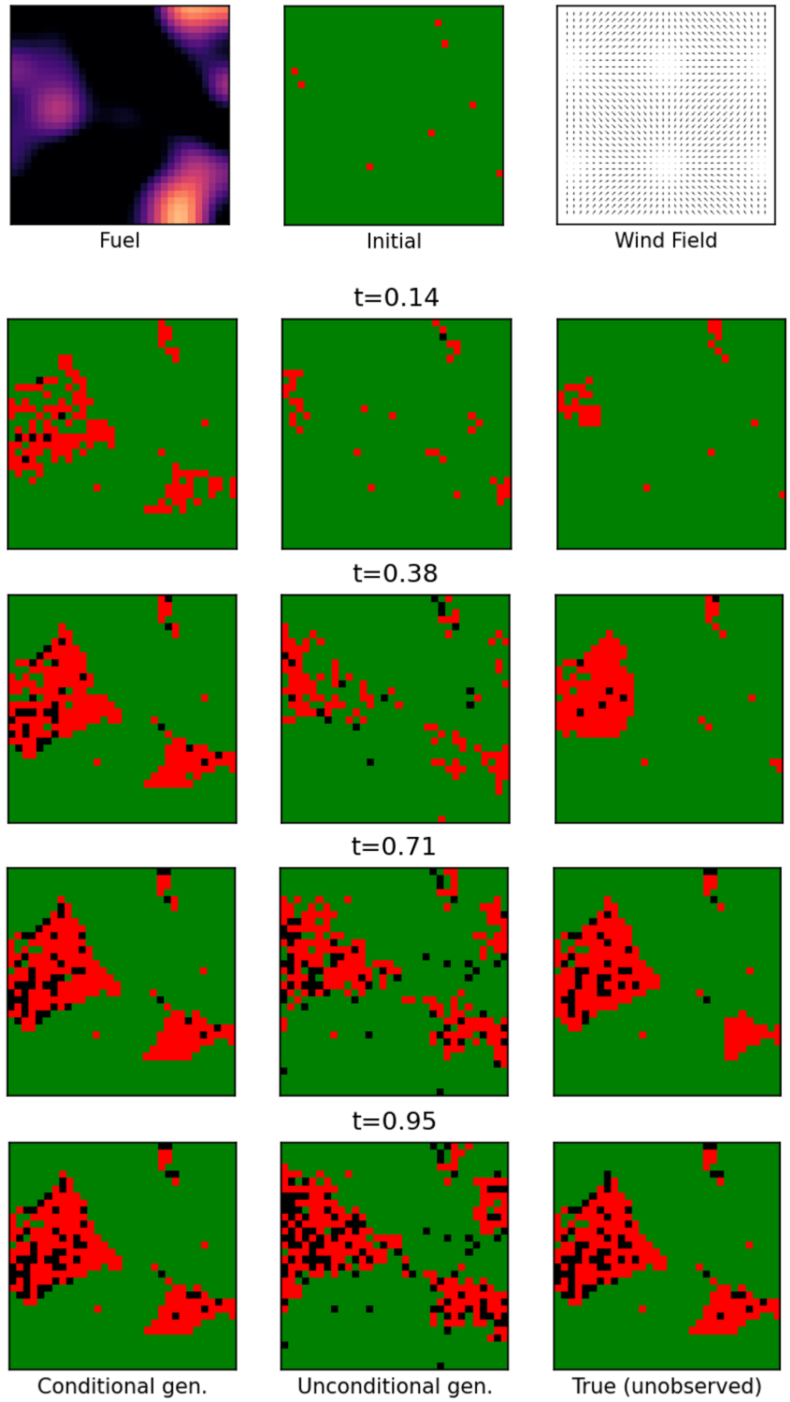

Figure 6: Initial conditions (top) and generated trajectories from the conditional (left) and unconditional (center) models, and true sequence observed only at the endpoints(right). Results shown for an held-out example.

embedding is then split into 4 attention heads, combined by another two-layer MLP with 512 hidden units that returns the off-diagonal values of the local transition rate matrices. We constrain the output to be positive by applying a ReLU function. We specify the prior path measure by a prior rate matrix, where we set to zero physically impossible transitions (e.g. $U \to E$ for wildfires, or $S \to R$ for epidemics) and the remaining off-diagonal elements to a constant value c. More complex functional forms are possible, and shall be chosen for example by simulating from the prior predictive distribution [Gelman et al., 2020].

E.3 Results

We provide a qualitative overview of the results we have obtained so far. These shall be considered preliminary, and a quantitative comparison with other baselines (e.g. the mean-field approximation from Seifner and Sánchez [2023]) will be carried out in future work. For the epidemics dataset, we display generated trajectories on an held-out graph in Figure 5, as well as the aggregated SIR curves for the same example in Figure 4. Notice how the conditional model tends to converge quickly to the end solution, while the unconditional model mirrors the true unobserved trajectory more closely. For the wildfires experiments, we display results on held-out examples in Figure 6 and Figure 7. Despite the lack of information at the initial time, the unconditional model can still predict an evolution very close to the ground truth final configuration.

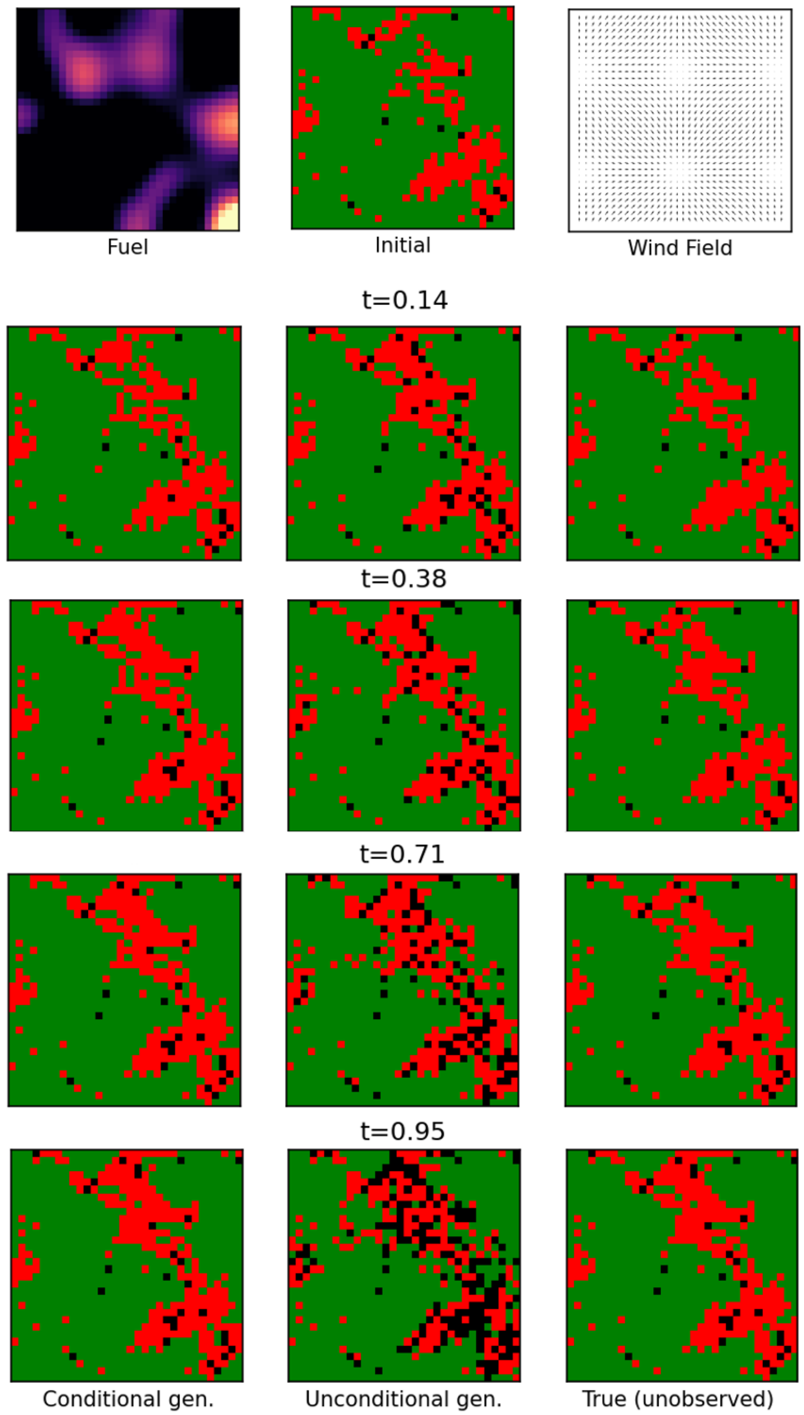

Figure 7: Same as Figure 6 but at a different stage of the simulated wildfire propagation, results shown for an held-out example.