From Geometry to Dynamics: Learning Overdamped Langevin Dynamics from Sparse Observations with Geometric Constraints

Dimitra Maoutsa [1]

Abstract

How can we learn the laws underlying the dynamics of stochastic systems when their trajectories are sampled sparsely in time? Existing methods either require temporally resolved high-frequency observations, or rely on geometric arguments that apply only to conservative systems, limiting the range of dynamics they can recover. Here, we present a new framework that reconciles these two perspectives by reformulating inference as a stochastic control problem. Our method uses geometry-driven path augmentation, guided by the geometry in the system's invariant density to reconstruct likely trajectories and infer the underlying dynamics without assuming specific parametric models. Applied to overdamped Langevin systems, our approach accurately recovers stochastic dynamics even from extremely undersampled data, outperforming existing methods in synthetic benchmarks. This work demonstrates the effectiveness of incorporating geometric inductive biases into stochastic system identification methods.

1. Introduction

How can we discover the underlying driving forces that govern the behaviour of complex, stochastic systems when we only measure their state at discrete time points? From pollen motion in a liquid medium (Einstein, 1905) and chemical reactions (Li, 2020) to population dynamics (Silva-Dias and López-Castillo, 2018; Fisher and Mehta, 2014) and cell growth (Alonso et al., 2014), many natural processes evolve following stochastic dynamics, best described by Langevin or stochastic differential equations (SDEs) of the form

$$\mathrm{d}\mathbf{X}_t = \mathbf{f}(\mathbf{X}_t)\,\mathrm{d}t + \boldsymbol{\sigma}\,\mathrm{d}\mathbf{W}_t. \tag{1}$$

[1]Technical University of Berlin, Berlin, Germany. Correspondence to: Dimitra Maoutsa <dimitra.maoutsa@gmail.com>.

Proceedings of the 43rd International Conference on Machine Learning, Seoul, South Korea. PMLR 306, 2026. Copyright 2026 by the author(s).

Under this formalism, the deterministic part of the equation $\mathbf{f}(\cdot) : \mathbb{R}^d \to \mathbb{R}^d$, the *drift* function, captures the long-term evolution of the state variables $\mathbf{X}_t \in \mathbb{R}^d$, while the stochastic part $\boldsymbol{\sigma} : \mathbb{R}^{d\times d}$, the *diffusion*, accounts for the contribution of unresolved degrees of freedom.

A common challenge for most inference methods is the assumption of fine temporal resolution of observations. Most approaches estimate $\mathbf{f}$ from short-time increments or by approximating transition densities over a small time interval (Sec E). However, when the sampling interval τ is large, these approximations break down, and the inverse problem becomes poorly constrained: multiple drifts can induce similar transition statistics between sparse observations. Thus, in such a setting, to accurately recover $\mathbf{f}$ requires additional assumptions or inductive biases consistent with the data.

Two dominant perspectives for stochastic inference. Existing data-driven approaches for learning dynamical systems provide powerful tools in the deterministic setting and for densely sampled stochastic dynamics (Sec E). Broadly, data-driven system identification for stochastic systems largely follows two tracks: ***Temporal methods*** (Fig. 1**A.**) rely on the **temporal ordering** of measurements, regressing state increments against states to estimate the drift. This works well when the inter-observation interval (τ) is small (Ruttor et al., 2013; Friedrich and Peinke, 1997; Friedrich et al., 2011; Ragwitz and Kantz, 2001; Jacobs et al., 2023). ***Geometric methods*** on the other hand, approximate the **invariant density** (Batz et al., 2016; Gu et al., 2021) or eigenstructure of the infinitesimal generator of the diffusion process (Singer and Coifman, 2008; Nüske et al., 2021; Ionides et al., 2006; Talmon and Coifman, 2015; Dsilva et al., 2016; Berry and Harlim, 2018)) (Fig. 1**B.**). Each perspective has limitations: temporal approaches deteriorate with increasing inter-observation intervals (Fig. 1**C.**), whereas geometric methods are restricted to systems with conservative forces (Berry and Harlim, 2015; Batz et al., 2016) or decoupled state variables (Singer and Coifman, 2008).

A unifying perspective: reconcile temporal and geometric methods by constraining with most prob-

able paths extracted from the invariant density. Here, we recast inference into a stochastic control problem and introduce **geometry-aware path augmentation**. Our method follows a simple premise that incorporates **geometric inductive biases** informed by the system's *invariant density* into dynamical inference: we postulate that the augmented paths should lie **in the vicinity** of **geodesic curves** (Fig. 1**F.** middle, magenta line) that connect consecutive measurements on the **empirical manifold** induced by the observations. To achieve this, **(i)** we approximate the Riemannian metric induced by the observations (Fig. 1**F.**) without the need to predefine the dimensionality of the empirical manifold, **(ii)** compute geodesics between consecutive observations through nonparametric approximation of shortest path distances between consecutive observations according to the approximated metric, and **(iii)** estimate the unobserved path between consecutive observations by generating **geometrically constrained diffusion bridges** that both respect temporal order and are guided toward identified geodesics (Fig. 1 **F.**). Nonparametric estimation of the drift function based on the augmented paths within an Expectation Maximisation framework (**E.M.**) (Dempster et al., 1977) results in accurate approximations of the underlying stochastic dynamics. Extensive numerical experiments demonstrate the effectiveness of our proposed method in recovering the ground truth stochastic dynamics, even in challenging scenarios where existing approaches fail.

Our contributions in this work are the following:
1. We introduce a geometry-guided augmentation framework for inference of sparsely observed SDEs that reconciles temporal and geometric methods.
2. We formulate the geometry-constrained augmentation in terms of a stochastic control problem that can be solved independently across inter-observation intervals.
3. We provide theoretical support showing that for long inter-observation intervals, the bias of short-time approximations is controlled by higher-order terms involving the curvature of the vector field, explaining why purely temporal methods degrade as τ increases (Sec. H.2).
4. We integrate this augmentation into a model-independent drift estimation framework (using a nonparametric function approximation), and demonstrate improved recovery of dynamics in challenging large-τ settings.

2. Setup and background

Setting. We consider a system whose state evolves according to Eq. 1. We observe the system state at discrete time points $t_k = k\tau$ at **inter-observation intervals** of τ time units, obtaining a time-ordered set of observations $\{\mathcal{O}_k \doteq \mathbf{X}_{t_k}\}_{k=1}^{K}$.

Goal. Given $\{\mathcal{O}_k\}_{k=1}^{K}$, estimate the drift $\mathbf{f}$.
Background. Many drift-learning approaches (Batz et al., 2018; Friedrich and Peinke, 1997) rely on short-time approximations: for a small τ they assume

$$\mathbf{X}_{t+\tau} \mid \mathbf{X}_t \approx \mathcal{N}\big(\mathbf{X}_t + \mathbf{f}(\mathbf{X}_t)\tau,\ \boldsymbol{\sigma}\boldsymbol{\sigma}^\top \tau\big), \tag{2}$$

which leads to a quadratic objective in $\mathbf{f}$ and motivates nonparametric estimators such as Gaussian-process drift models (Ruttor et al., 2013; Hostettler et al., 2018; Zhao et al., 2020). However, in our considered setting the sampling interval τ is **not** small: the Gaussian increment model of Eq. 2 (and Euler–Maruyama-based likelihoods) becomes inaccurate because the underlying transition density is generally non-Gaussian (see more details in Sec A).

Common drift inferences approaches rely on *small* inter-observation intervals τ (Batz et al., 2018). As τ increases, the EuM approximation becomes inaccurate: transition densities are not Gaussian, and higher-order remainder terms related to the curvature of the flow field become important (see further theoretical analysis in Sec. H.2 and c.f. Fig. 8). Attempts to mitigate this problem by introducing bridge sampling to infer the unobserved path between observations (Batz et al., 2018; Sermaidis et al., 2013) provide small improvements, because these methods rely on linearised or otherwise simplified bridge dynamics that do not match the true transition densities (c.f. Sec. E).

Our approach. Here, we target this large inter-observation interval setting by merging insights from both temporal and geometric perspectives. Specifically, our approach combines **nonlinear** bridge sampling with **a geometric approximation of the system's invariant density** as detailed in the following. This introduces an implicit bias in our approach: among the alternative potential hypotheses regarding the underlying drift dynamics compatible with the data, we favour those whose likely paths follow high-density regions on the invariant density.

3. Methodology

Core idea. The invariant density of the observed system imposes a low-dimensional structure on the state space, within which the observations are confined. We propose that this low-dimensional structure is well approximated by a Riemannian manifold $\mathcal{M}_\infty \in \mathbb{R}^{m \leq d}$ in the ambient space (see Sec. G), and that the ensemble of observations $\{\mathcal{O}_k\}_{k=1}^{K}$ offers a reliable discrete approximation to $\mathcal{M}_\infty$. We term this observation-based approximation the *empirical manifold* $\mathcal{M}$. The central premise of our approach is that **unobserved paths between successive observations will be lying either *on* or *in the vicinity* of the empirical manifold** $\mathcal{M}$. In particular, we postulate that unobserved paths should lie **in the vicinity of geodesics that connect consecutive observations** on $\mathcal{M}$.

However, while this view of a lower dimensional manifold

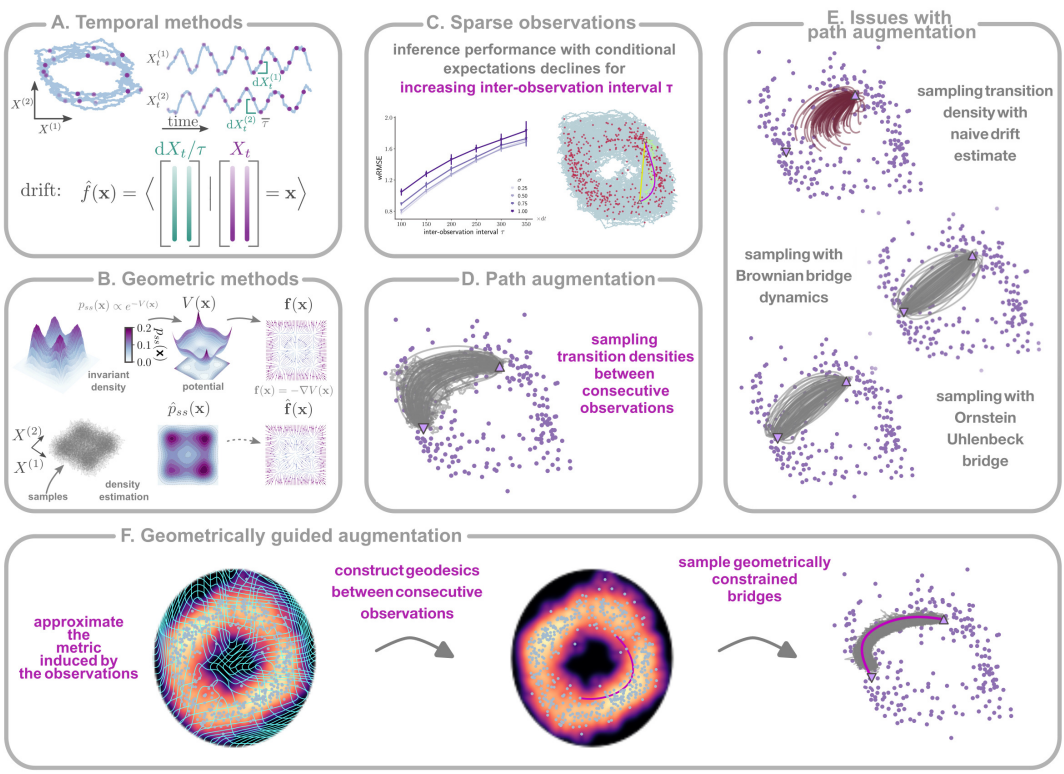

Figure 1

Temporal and geometric perspectives for discovering stochastic dynamics and proposed inference with geometrically guided augmentation. (**A.**) Temporal methods consider the time-ordering of observations $\{\mathcal{O}_k\}_{k=1}^K$ (*purple dots*) to approximate the drift $\mathbf{f}(\mathbf{x})$ with conditional rescaled state increments $\hat{\mathbf{f}}(\mathbf{x}) = \langle \frac{\mathrm{d}\mathbf{X}_t}{\tau} | \mathbf{X}_t = \mathbf{x} \rangle$. (**B.**) Geometric methods assume a conservative drift $\mathbf{f}(\mathbf{x}) = -\nabla V(\mathbf{x})$ as the gradient of a potential. (**C.**) With increasing inter-observation interval τ performance of temporal methods degrades because Euclidean distances ignore the curvature of the latent continuous path between consecutive observations. (**D.**) Path augmentation alternates between state estimation - by sampling diffusion bridges for each inter-observation interval - and drift inference. (**E.**) Commonly used path augmentation methods employ Brownian or Ornstein-Uhlenbeck bridges that increasingly deviate from the unobserved path as τ grows. (lower) Illustration of the ground truth (*neon green*) and geodesic (*magenta*) continuous path between two observations and of that assumed during inference with Gaussian likelihood (*yellow line*). (**F.**) Geometrically guided augmentation approximates first the metric induced by the invariant density, constructs geodesics connecting consecutive observations, and samples geometrically constrained diffusion bridges.

embedded in a higher dimensional ambient space helps to build intuition, for practical purposes we adopt a complementary view of the low dimensional manifold inspired by (Fröhlich et al., 2021). According to this view, we consider the entire observation space $\mathbb{R}^d$ as a smooth Riemannian manifold, $\mathcal{M} \doteq \mathbb{R}^d$, characterised by a Riemannian metric $\mathfrak{h}$. The effect of the nonlinear geometry of the observations is then captured by the metric $\mathfrak{h}$. Thus to approximate the geometric structure of the system's invariant density, we learn the Riemannian metric tensor $H : \mathbb{R}^d \to \mathbb{R}^{d \times d}$ and compute the geodesics between consecutive observations according to the learned metric. Intuitively, according to this view the observations $\{\mathcal{O}_k\}_{k=1}^K$ introduce distortions in the way we compute distances in the state space. The advantage of this approach is that we do not need to estimate the dimensionality of the empirical manifold, which would have been difficult due to the presence of fluctuations in the system's dynamics. Instead, we still operate in the original space, while the empirical manifold introduces distortions in the estimated metric (Fig. 1**F.i.**).

Inference framework. Our approach comprises three steps: **(α.)** Approximation of the geometric structure of the system's invariant density with metric learning, **(β.)** estimation of the (latent) system state between consecutive observations guided by the invariant density (**path augmentation**), and **(γ.)** data-driven estimation of the drift function (Fig. 1). We perform the two final steps in an iterative manner within an Expectation Maximisation (**E.M.**) framework (Dempster et al., 1977).

(α.) Approximating the Riemannian geometry induced by the observations. Although there are many methods for approximating Riemannian manifolds (Tenenbaum et al., 2000; Balasubramanian and Schwartz, 2002; Mead,

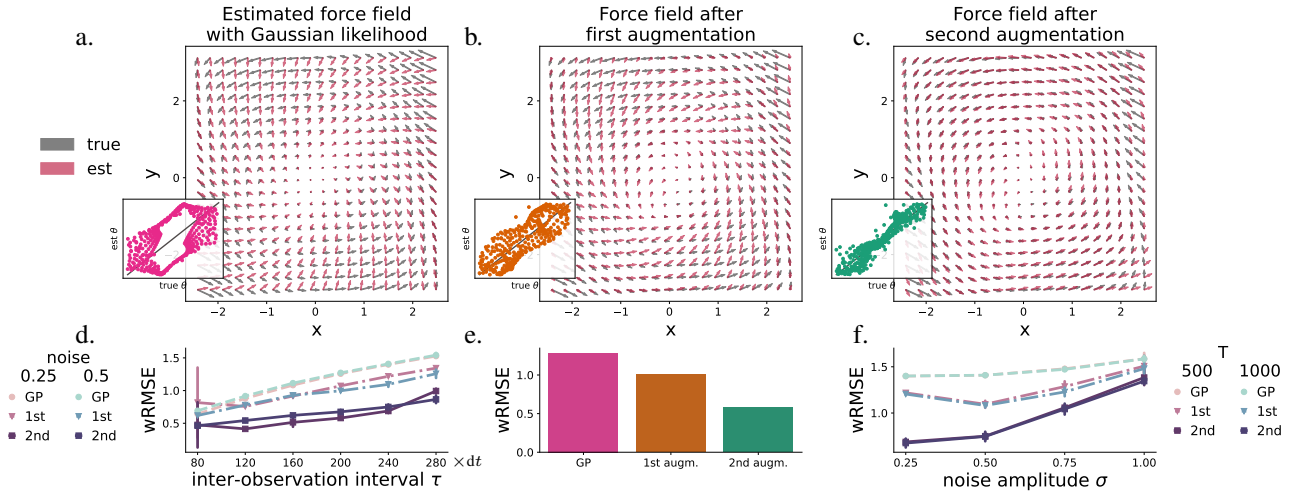

Figure 2

Geometry-aware path augmentation improves drift inference after two iterations. Estimated (*red*) vs. ground truth (*grey*) force field with **a.)** Gaussian likelihood, **b.)** after one, and **c.)** after two augmentations. (**Insets**) Ground truth vs. estimated angles at grid points. **d.)** Weighted (by observation density) root mean square error (wRMSE) vs. inter-observation interval τ for different noise levels $\sigma = \{0.25, 0.5\}$ for drift estimated with a Gaussian likelihood (*gaus*-circles), after first augmentation (*1st*-triangles), and after second augmentation (*2nd*-squares) for $T = 500$ (time units). **e.)** wRMSE across iterations for the presented example. **f.)** wRMSE vs. noise amplitude σ for different trajectory durations $T = \{500, 1000\}$ (time units) for inter-observation interval $\tau = 240$ (dt). Markers in **d.)** and **f.)** indicate augmentation steps. Error bars: one standard deviation over five independent runs. ($\mathrm{d}t = 0.01$)

1992; Roweis and Saul, 2000), our objective is to obtain a representation that acts as a *local* constraint for subsequent state estimation between consecutive observations. We achieve this in two steps: **(i.)** We approximate in the ambient space $\mathbb{R}^d$ the metric $\mathfrak{h}$ induced by the observations (see Fig. 1**F.i.**). This identifies regions of the state space with high observation density (represented with small metric values). **(ii.)** We construct geodesics between consecutive observations on the empirical manifold $(\mathcal{M} \doteq \mathbb{R}^d, \mathfrak{h})$ (see Fig. 1**F.ii.**). The geodesics provide geometry-informed reference curves between consecutive observations that identify high-density paths through the empirical geometry. Each such path subsequently functions as a soft proximity constraint during the latent state estimation.

(i.) Approximation of the invariant metric. To approximate the (local) metric $\mathfrak{h}$ in a nonparametric form at locations $\mathbf{x}$ of the state space, we follow (Arvanitidis et al., 2019), and consider the inverse of the weighted local diagonal covariance computed on the K observations as

$$H_{dd}(\mathbf{x}) = \left(\sum_{k=1}^{K} w_k(\mathbf{x}) \left(\mathcal{O}_k^{(d)} - x^{(d)} \right)^2 + \epsilon \right)^{-1}, \quad (3)$$

with weights $w_k(\mathbf{x}) = \exp\left(-\frac{\|\mathcal{O}_k - \mathbf{x}\|_2^2}{2\sigma_{\mathcal{M}}^2}\right)$, and $A^{(d)}$ denoting the d-th dimensional component of the vector $\mathbf{A}$ for $\mathbf{A} \in \{\mathbf{x}, \mathcal{O}_k\}$. The parameter $\epsilon > 0$ is a small value ensuring non-zero diagonals of the weighted covariance matrix, while $\sigma_{\mathcal{M}}$ is a hyper-parameter characterising the curvature of the approximated manifold.

(ii.) Constructing geodesics between consecutive observations. To compute the geodesic curves connecting consecutive observations on the empirical manifold, we employ the approximated metric tensor $\mathbf{H}(\mathbf{x})$. We identify the geodesic curve $\gamma_{t'}^k$ between $\mathcal{O}_k$ and $\mathcal{O}_{k+1}$ as the curve with minimum energy that connects these two points, i.e., as the minimiser of the kinetic energy functional $\mathcal{E}(\gamma_{t'}^k) = \int_0^1 L_{\mathcal{M}}(\gamma_{t'}^k, \dot{\gamma}_{t'}^k)\, \mathrm{d}t'$

$$\gamma_{t'}^{k*} = \underset{\substack{\gamma_{t'}^k, \\ \gamma_0^k = \mathcal{O}_k, \gamma_1^k = \mathcal{O}_{k+1}}}{\arg\min} \int_0^1 L_{\mathcal{M}}(\gamma_{t'}^k, \dot{\gamma}_{t'}^k) \mathrm{d}t', \quad (4)$$

$$\text{with} \quad \int_0^1 L_{\mathcal{M}}(\gamma_{t'}^k, \dot{\gamma}_{t'}^k)\, \mathrm{d}t' = \frac{1}{2} \int_0^1 \|\dot{\gamma}_{t'}^k\|_{\mathfrak{h}}^2 \mathrm{d}t',$$

where $L_{\mathcal{M}}(\gamma_{t'}^k, \dot{\gamma}_{t'}^k)$ is an appropriately constructed Lagrangian. The minimising curve of this functional is the same as the minimiser of the curve length functional $\ell(\gamma_{t'})$ (c.f. Eq. 31), i.e., the geodesic (Do Carmo, 1992). This results in a system of second order differential equations (Eq. 34) (Arvanitidis et al., 2017; Do Carmo, 1992) (Sec. A.3.2) with boundary conditions $\gamma_0^k = \mathcal{O}_k$ and $\gamma_1^k = \mathcal{O}_{k+1}$ that we solve with a probabilistic differential equation solver as in (Arvanitidis et al., 2019).

(β.) Latent state estimation: Geometry-guided augmentation. To estimate the unobserved system state between

wRMSE ↓	total duration T	$\tau = 80 \times \mathrm{d}t$	$\tau = 120 \times \mathrm{d}t$	$\tau = 160 \times \mathrm{d}t$	$\tau = 200 \times \mathrm{d}t$	$\tau = 240 \times \mathrm{d}t$	$\tau = 280 \times \mathrm{d}t$
				Van der Pol			
$\boldsymbol{\sigma = 0.25}$							
GP	500	0.642 ± 0.006	0.879 ± 0.005	1.083 ± 0.015	1.258 ± 0.011	1.399 ± 0.003	1.528 ± 0.0153
SVISE	500	1.465 ± 0.009	0.857 ± 0.021	0.740 ± 0.072	0.592 ± 0.026	$\mathbf{0.587 \pm 0.112}$	$\mathbf{0.824 \pm 0.003}$
KM-basis	500	$\mathbf{0.368 \pm 0.054}$	0.452 ± 0.011	0.671 ± 0.023	1.588 ± 0.021	1.751 ± 0.008	1.735 ± 0.020
LatentSDE	500	1.091 ± 0.316	1.091 ± 0.039	1.098 ± 0.023	1.089 ± 0.036	1.088 ± 0.038	1.091 ± 0.039
LatentSDE+GP-pre	500	1.095 ± 0.038	1.085 ± 0.039	1.101 ± 0.034	1.089 ± 0.038	1.106 ± 0.045	1.102 ± 0.039
GSBM	500	1.518 ± 0.033	1.435 ± 0.055	-	-	-	-
[SF]2M	1500	1.741 ± 0.304	1.801 ± 0.226	1.745 ± 0.322	1.583 ± 0.132	1.816 ± 0.228	1.721 ± 0.094
$\mathrm{MFM}_{\mathrm{LAND}}$	1500	1.463 ± 0.007	1.469 ± 0.005	1.469 ± 0.012	1.469 ± 0.008	1.469 ± 0.006	1.467 ± 0.008
Geometric (our)	500	0.474 ± 0.034	$\mathbf{0.413 \pm 0.016}$	$\mathbf{0.514 \pm 0.068}$	$\mathbf{0.578 \pm 0.022}$	0.687 ± 0.032	0.993 ± 0.037
$\boldsymbol{\sigma = 0.50}$							
GP	500	0.691 ± 0.029	0.916 ± 0.014	1.114 ± 0.15	1.272 ± 0.030	1.409 ± 0.019	1.542 ± 0.044
SVISE	500	1.235 ± 0.083	0.9935 ± 0.015	0.7505 ± 0.052	0.736 ± 0.072	1.3565 ± 0.278	1.425 ± 0.086
KM-basis	500	0.495 ± 0.010	0.727 ± 0.008	0.890 ± 0.024	1.683 ± 0.020	1.744 ± 0.038	1.732 ± 0.065
LatentSDE	500	1.158 ± 0.036	1.151 ± 0.045	1.160 ± 0.032	1.151 ± 0.036	1.146 ± 0.033	1.176 ± 0.046
LatentSDE+GP-pre	500	1.158 ± 0.045	1.159 ± 0.034	1.159 ± 0.027	1.151 ± 0.034	1.150 ± 0.028	1.191 ± 0.052
GSBM	500	6.106 ± 2.988	4.818 ± 3.060	4.738 ± 3.304	4.875 ± 3.222	9.076 ± 1.451	26.187 ± 18.804
[SF]2M	1500	1.869 ± 0.482	1.813 ± 0.286	1.484 ± 0.096	1.876 ± 0.247	1.753 ± 0.158	1.707 ± 0.233
$\mathrm{MFM}_{\mathrm{LAND}}$	1500	1.517 ± 0.011	1.526 ± 0.006	1.536 ± 0.009	1.537 ± 0.017	1.528 ± 0.015	1.545 ± 0.019
Geometric (our)	500	$\mathbf{0.462 \pm 0.019}$	$\mathbf{0.541 \pm 0.023}$	$\mathbf{0.621 \pm 0.012}$	$\mathbf{0.675 \pm 0.030}$	$\mathbf{0.750 \pm 0.038}$	$\mathbf{0.865 \pm 0.057}$

Table 1

Performance comparison in terms of weighted root mean square error (wRMSE) of considered frameworks for different noise conditions σ and inter-observation intervals τ for the Van der Pol system. For all settings $\mathrm{d}t = 0.01$.

consecutive observations $\boldsymbol{\mathcal{O}}_k$ and $\boldsymbol{\mathcal{O}}_{k+1}$, we perform variational inference (Beal, 2003)(see Sec. A.3). Given a prior diffusion process with drift $\hat{\mathbf{f}}(\cdot) : \mathbb{R}^d \to \mathbb{R}^d$ and diffusion σ, we construct an **approximating process** conditioned **i.)** to pass through the observations, and **ii.)** to respect the local geometry of the invariant density as it is represented by the geodesics. The conditioned process is also a diffusion process with the same diffusion constant and an effective drift function $\mathbf{g}(\mathbf{x}, t)$ (Chetrite and Touchette, 2015; Majumdar and Orland, 2015). The path probability measure $Q_X(\mathbf{X}_{0:T})$ induced by the approximating process

$$Q_X(\mathbf{X}_{0:T}) : \quad \mathrm{d}\mathbf{X}_t = \mathbf{g}\left(\mathbf{X}_t, t\right) \mathrm{d}t + \sigma \mathrm{d}\bar{\mathbf{W}}_t \tag{5}$$
$$= \left(\widehat{\mathbf{f}}(\mathbf{X}_t) + \mathbf{u}(\mathbf{X}_t, t)\right) \mathrm{d}t + \sigma \mathrm{d}\bar{\mathbf{W}}_t,$$

provides an approximation to the unobserved continuous system state. In Eq. 5, $\mathbf{u}(\cdot, \cdot) : \mathbb{R}^d \times \mathbb{R}^+ \to \mathbb{R}^d$ is a time-dependent control term that guides the approximating path distribution through the observations, while staying in the vicinity of the corresponding geodesics between them.

More precisely, we obtain the controlled drift $\mathbf{g}\left(\mathbf{X}_t, t\right)$ from the solution of the variational problem of minimising the functional (see Sec. A.3.1)

$$\mathcal{F}[Q_X] = \mathcal{KL}\Big(Q_X(\mathbf{X}_{0:T}) || \mathrm{P}(\mathbf{X}_{0:T} \mid \hat{\mathbf{f}})\Big) \tag{6}$$
$$- \sum_{k=1}^{K} \Big\langle \ln \mathrm{P}(\boldsymbol{\mathcal{O}}_k \mid \mathbf{X}_{t_k}) \Big\rangle_Q + \Big\langle \|\boldsymbol{\Gamma}_t - \mathbf{X}_{0:T}\|^2 \Big\rangle_Q$$
$$= \tfrac{1}{2} \int_0^T \int \Big[\|\mathbf{g}(\mathbf{x},t) - \hat{\mathbf{f}}(\mathbf{x})\|^2_{\mathbf{D}} + U_{\mathcal{O}}(\mathbf{x},t) + \beta U_{\mathcal{G}}(\mathbf{x},t) \Big] q_t(\mathbf{x}) \mathrm{d}\mathbf{x} \mathrm{d}t,$$

where $\boldsymbol{\Gamma}_t$ denotes the sequence of K geodesics indexed by time t, $\boldsymbol{\Gamma}_t \doteq \{\boldsymbol{\gamma}^k_{t'}\}_{t=(k-1)\tau + t'\tau}$, where $\boldsymbol{\gamma}^k_{t'}$ is the geodesic connecting $\boldsymbol{\mathcal{O}}_k$ and $\boldsymbol{\mathcal{O}}_{k+1}$, and $t' \in [0, 1]$ denotes a rescaled time variable, and β is a weighting term. In Eq. 6, the term $U_{\mathcal{O}}(\mathbf{x}, t) = -\sum_{t_k} \ln \mathrm{P}(\boldsymbol{\mathcal{O}}_k \mid \mathbf{x})\, \delta(t - t_k)$ **forces the augmentation to pass through the observations at each bridge boundary**, while $U_{\mathcal{G}}(\mathbf{x}, t) \doteq \|\boldsymbol{\Gamma}_t - \mathbf{x}\|^2$ **guides the latent path towards the identified geodesics**. This geodesic term should be considered as a soft penalty rather than a hard constraint: the controlled process is allowed to deviate from $(\boldsymbol{\Gamma}_t)$, with the strength of the geometric bias determined by the scalar weight β (Fig. 6).

This minimisation can be construed as a stochastic control problem (Opper, 2019) with the objective to identify a time-dependent drift adjustment $\mathbf{u}(\mathbf{x}, t) := \mathbf{g}(\mathbf{x}, t) - \hat{\mathbf{f}}(\mathbf{x})$ for the system with drift $\hat{\mathbf{f}}(\mathbf{x})$ so that the controlled dynam-

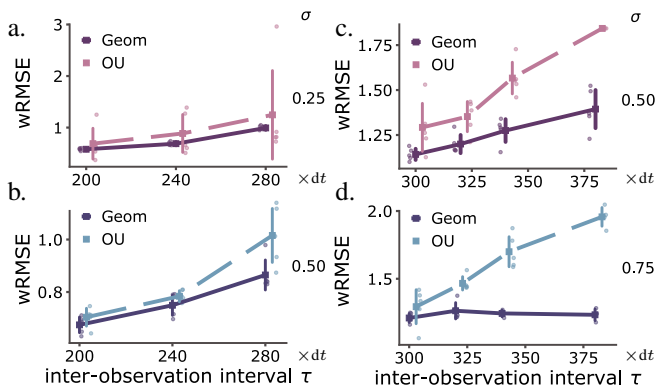

Figure 3

Comparison of geometry-aware inference against inference with Ornstein-Uhlenbeck augmentation. Weighted root mean square error (wRMSE) vs. different inter-observation intervals τ for different noise amplitudes for moderate inter-observation intervals with **a.)** $\sigma = 0.25$ and **b.)** $\sigma = 0.50$, and for large inter-observation intervals with **c.)** $\sigma = 0.50$ and **d.)** $\sigma = 0.75$, where only one observation per oscillation period is available. Error bars indicate one standard deviation over five independent runs. ($\mathrm{d}t = 0.01$)

ics fulfil the path constraints $U_{\mathcal{O}}(\mathbf{x}, t)$ and $U_{\mathcal{G}}(\mathbf{x}, t)$. The first term, $\|g(x,t) - \hat{f}(x)\|_D^2$, is the quadratic control cost: it penalises changes of the drift away from the current estimate $\hat{f}$. The potential $U_O(x,t)$ enforces consistency with the observed endpoints, while $U_G(x,t)$ biases the bridge toward the geodesic.

The optimal time-dependent control for the interval between $\boldsymbol{\mathcal{O}}_k$ and $\boldsymbol{\mathcal{O}}_{k+1}$ results from the solution of the backward equation (Kappen, 2005a; Maoutsa and Opper, 2022)

$$\frac{\partial \phi_t(\mathbf{x})}{\partial t} = -\mathcal{L}_{\hat{\mathbf{f}}}^{\dagger}\phi_t(\mathbf{x}) + \beta\, U_{\mathcal{G}}(\mathbf{x}, t)\phi_t(\mathbf{x}), \tag{7}$$

with terminal condition $\phi_{t_{k+1}}(\mathbf{x}) = \chi(\mathbf{x}) = \delta(\mathbf{x} - \boldsymbol{\mathcal{O}}_{k+1})$ and with $\mathcal{L}_{\hat{f}}^{\dagger}$ denoting the adjoint Fokker-Planck operator for the process of Eq. 24. As shown in (Maoutsa and Opper, 2022) the optimal drift adjustment $\mathbf{u}(\mathbf{x}, t)$ can be expressed in terms of the difference of the logarithmic gradients of two probability flows

$$\mathbf{u}^*(\mathbf{x}, t) = \mathbf{D}\Big(\nabla \ln q_{T-t}(\mathbf{x}) - \nabla \ln \rho_t(\mathbf{x})\Big), \tag{8}$$

where ρ_t fulfils the forward (filtering) partial differential equation (PDE)

$$\frac{\partial \rho_t(\mathbf{x})}{\partial t} = \mathcal{L}_{\hat{\mathbf{f}}}\rho_t(\mathbf{x}) - \beta\, U_{\mathcal{G}}(\mathbf{x}, t)\rho_t(\mathbf{x}), \tag{9}$$

while q_t is the solution of a time-reversed PDE with initial condition $q_0(\mathbf{x}) \propto \rho_T(\mathbf{x})\chi(\mathbf{x})$

$$\frac{\partial q_t(\mathbf{x})}{\partial t} = -\nabla\cdot\Big[\Big(\mathbf{D}\,\nabla \ln \rho_{T-t}(\mathbf{x}) - \hat{\mathbf{f}}(\mathbf{x}, T-t)\Big)q_t(\mathbf{x})\Big] + \frac{\mathbf{D}}{2}\nabla^2 q_t(\mathbf{x}).$$

The density ρ_t is therefore a forward filtering density: it propagates probability mass under the current drift estimate $\hat{\mathbf{f}}$, while down-weighting paths that accumulate large geometric cost. The terminal factor $\chi(\mathbf{x})$ imposes the endpoint condition at $\mathcal{O}_{k+1}$. The density q_t is then evolved backward in time from $q_0(\mathbf{x}) \propto \rho_T(\mathbf{x})\chi(\mathbf{x})$, and carries the information required to steer the process toward the endpoint.

For each interval $[\boldsymbol{\mathcal{O}}_k, \boldsymbol{\mathcal{O}}_{k+1}]$ we identify the posterior path measure (minimiser of Eq. 35) by solving such a stochastic control problem for the time-varying control $\mathbf{u}(\mathbf{x}, t)$ of Eq. 8. This results in a set of $K - 1$ *independent* optimal control problems, that are solved in parallel for efficiency.

(γ.) Estimating the drift. We approximate the drift function in a model independent framework by imposing a Gaussian process prior on the function values $\mathbf{f} \sim \mathrm{P}_o(\mathbf{f}) = \mathcal{GP}(\mathbf{m}^f, k^f)$, where $\mathbf{m}^f$ and k^f denote the mean and covariance function of the Gaussian process. The optimal measure for the drift Q_f is a Gaussian process given by (Batz et al., 2018)

$$Q_f \propto \mathrm{P}_o \exp\left(-\frac{1}{2}\int \|\mathbf{f}(\mathbf{x})\|_D^2 A(\mathbf{x}) - 2\langle \mathbf{f}(\mathbf{x}), B(\mathbf{x})\rangle_D \mathrm{d}\mathbf{x}\right), \tag{10}$$

with $A(\mathbf{x}) \doteq \int_0^T q_t(\mathbf{x})\mathrm{d}t$ and $B(\mathbf{x}) \doteq \int_0^T q_t(\mathbf{x})\mathbf{g}(\mathbf{x}, t)\mathrm{d}t$, where $q_t(\mathbf{x})$ denotes the marginal density of the constrained process' state obtained by the state estimation. The function $\mathbf{g}(\mathbf{x}, t)$ denotes the effective (time-dependent) drift of the constrained process (Eq. 5), resulting from the solution of the individual control problems accounting for the observations and the invariant geometry.

4. Results

Revealing stochastic dynamics in model systems. To demonstrate the effectiveness of our approach, we inferred the stochastic dynamics of model systems, and compared the resulting estimates to those obtained from: **(i.)** Gaussian process regression without state estimation (**GP**), **(ii.)** path augmentation with Ornstein-Uhlenbeck dynamics (**OU**) (Batz et al., 2018), **(iii.)** sparse variational inference with state estimation (**SVISE**) (Course and Nair, 2023a), **(iv.)** basis function approximation of Kramers-Moyal coefficients, i.e. the drift function (**KM-basis**) (Nabeel et al., 2025), and **(v.)** latent SDE inference with amortized reparameterization with (**LatentSDE+GP-pre**) and without pre-training (**LatentSDE**) (Course and Nair, 2023b), **(vi.)** metric flow matching (**MFM**) (Kapusniak et al., 2024)(with RBF (Arvanitidis et al., 2021) and LAND metric (Arvanitidis et al., 2019) metric approximations - see Sec I.5), **(vii.)** generalized Schrödinger bridge matching **(GSBM)** (Liu et al., 2024), **(viii.)** simulation-free Schrödinger bridges via score and flow matching (**[SF]^{2}M**) (Tong et al., 2024) (Sec. J.2). We tested our method on non-conservative systems inducing diverse types of invariant geometries: **(a.)** a Van der Pol system, **(b.)** a process with harmonic trapping and circulation and a Gaussian re-

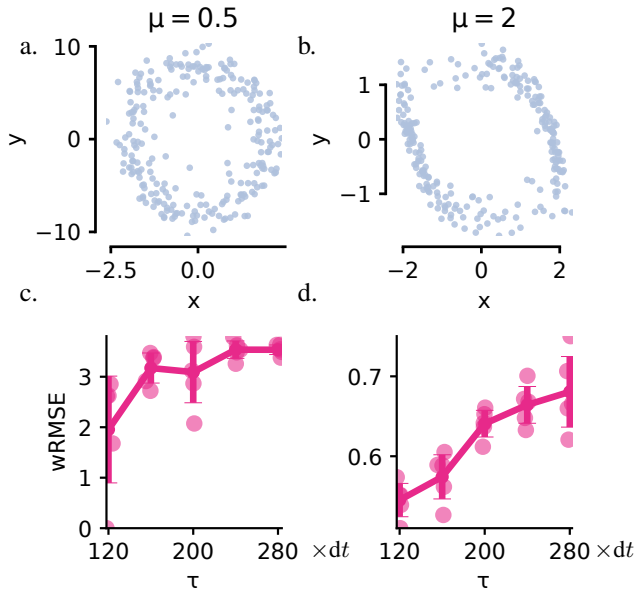

Figure 4

Geometry-aware inference provides accurate drift estimation for different empirical manifold geometries resulting from different parameter regimes of the Van der Pol system. a.-b.) Empirical manifold for the Van der Pol system with different μ parameters. Notice the different scales on the axes. **c.-d.)** Inference performance of the proposed framework against inter-observation interval τ. Error bars indicate one standard deviation over five independent runs. ($dt = 0.01$)

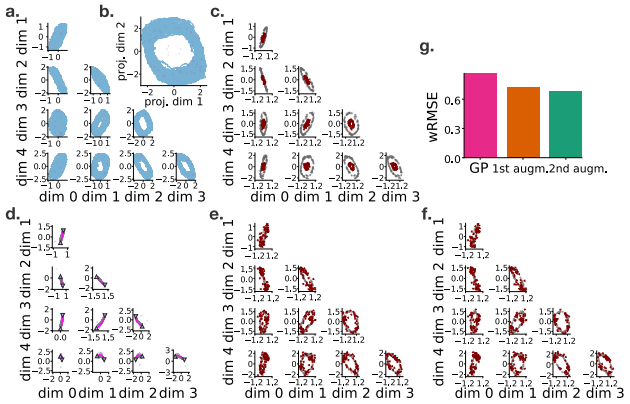

Figure 5

Van der Pol system embedded in a 5D ambient space with isotropic 5D dynamical noise and a mean restoring force orthogonal to the embedding subspace. a.) Sparse observations (purple) together with the underlying unobserved path (light blue). **b.)** Trajectory of the Van der Pol system in its intrinsic 2D coordinates. **c.)** Observations together with the flow field estimated from sparse data by Gaussian process regression (maroon), compared to the ground-truth flow field (grey). **d.)** Example of a geometry-guided augmentation step: constructed geodesic (magenta) and sampled bridge trajectory (grey). **e., f.)** Inferred flow field after the first and second augmentation steps, respectively. **g.)** Weighted root mean squared error (wRMSE) of the inferred flow field at each stage of inference: initial GP estimate, after the first augmentation, and after the second augmentation.

pulsive obstacle in the centre introduced in (Frishman and Ronceray, 2020), **(c.)** a Hopf system, **(d.)** a Selkov glycolysis model (Selkov, 1968) (see Sec. J), **(e.)** a Repressilator system (see Fig. 9), **(f.)** as well as Van der Pol dynamics embedded in higher dimensional spaces (5D and 8D) (see Fig. 5 and Fig. 10). For most settings, the proposed framework outperformed existing methods, especially for large inter-observation intervals (Table 2 and 1).

We quantified the quality of the inference in terms of weighted root mean square error (**wRMSE**) between the estimated and ground truth drift functions evaluated on a $d-$dimensional grid spanning the state space volume of the observations. The weights for each grid point were obtained from a kernel density estimation of the observations. Thus misalignment of ground truth and estimated dynamics were penalised stronger for regions of the state space visited more frequently by the observed process.

For a system with a drift function following Van der Pol dynamics, we found that only after two E.M. iterations, the estimated force field (red arrows) was well aligned to the ground truth force field that generated the observations (grey arrows) (Fig. 2a.). For comparison we demonstrate also the result of a naive GP estimation, which resulted in a flow fields orthogonal to the ground truth ones.
We performed systematic estimations for this system under different noise conditions σ, observed at different inter-observation intervals τ, for different trajectory lengths T (Sec. J). For the examined noise amplitudes (Fig. 2 f.), the proposed algorithm improved the naive estimation with Gaussian assumptions within two iterations (Fig. 2). For increasing noise the improvement contributed by our method decreased (Fig. 2f.), as the invariant geometry is less well defined for large noise, but was still considerable.

Impact of the geometry of empirical manifold. We performed inference for different parameter values of the Van der Pol system ($\mu = 1$ (as above) and $\mu = 0.5$ and $\mu = 2$), that resulted in asymmetries of the invariant density (Fig. 4). We observed that the performance of all inference frameworks deteriorated for increasing asymmetry (larger dynamic range along one dimension), yet our method still delivered more accurate predictions compared to the other considered frameworks. Moreover, employing an alternative metric-learning algorithm to approximate the invariant geometry resulted in comparable performance. Specifically, $\text{Geometric}_{\text{RBF}}$ in Table 3 uses the metric of (Arvanitidis et al., 2021), further developed in (Kapusniak et al., 2024), in which a diagonal metric is represented as a positive linear combination of Gaussian RBFs centred at selected cluster centres.

Impact of noise amplitude. For systems with small dynamical noise (small σ), geodesics approximate the manifold structure better, however the path integral control is

wRMSE ↓	Harmonic trapping system			Hopf			Selkov	
	$\tau = 150$	$\tau = 200$	$\tau = 250 \times \mathrm{d}t$	$\tau = 200$	$\tau = 300$	$\tau = 400 \times \mathrm{d}t$	$\tau = 100$	$\tau = 200 \times \mathrm{d}t$
GP	2.632 ± 0.007	3.387 ± 0.012	3.733 ± 0.011	0.781 ± 0.006	0.969 ± 0.015	1.069 ± 0.006	0.550 ± 0.021	0.682 ± 0.040
SVISE	35.204 ± 39.888	3.462 ± 0.129	7.540 ± 7.602	2.113 ± 0.658	4.960 ± 2.687	3.936 ± 1.063	5.793 ± 0.028	2.028 ± 0.045
LatentSDE	**2.348** ± 0.032	**2.340** ± 0.047	**2.356** ± 0.042	1.168 ± 0.052	1.161 ± 0.053	1.173 ± 0.046	0.742 ± 0.022	0.747 ± 0.021
Geometric (ours)	2.762 ± 0.132	3.034 ± 0.143	2.693 ± 0.992	**0.210** ± 0.013	**0.237** ± 0.010	**0.255** ± 0.028	**0.414** ± 0.245	**0.682** ± 0.071

Table 2

Performance comparison in terms of wRMSE for the considered frameworks for three different nonlinear dynamical systems and for increasing inter-observation interval τ. Numbers indicate mean wRMSE and standard deviation of five independent runs for each setting. For all settings $\mathrm{d}t = 0.01$.

limited by the control costs proportional to inverse noise covariance. Our framework had comparable accuracy for all inter-observation lengths, but improvement was small for small lengths since in that setting the estimation with Gaussian likelihood already provided a good approximation of the ground truth drift.

We compared our method to the approach proposed in (Batz et al., 2018). In this work, the authors perform augmentation with Ornstein-Uhlenbeck bridges, i.e. assuming linear underlying dynamics. We found that our approach delivered more accurate estimates for larger inter-observation intervals. For inter-observation intervals with only one observation per oscillation period (Fig. 3c.,d.), our approach delivered better results by considering additionally knowledge of the direction of movement in the state space (Sec. J). The variance of estimates of the proposed method was smaller compared to Batz et al. due to the consistency imposed by conditioning on the invariant geometry of the system. Predictions improved with longer observation intervals T, and for decreasing noise amplitude σ. In both settings the invariant geometry was more well approximated by the empirical manifold. Overall, state estimation with Ornstein-Uhlenbeck dynamics (Batz et al., 2018), was increasingly less capable of correctly estimating the latent system state and subsequently correctly approximating the unknown drift function especially as the length of the inter-observation interval τ increased.

Effects of noise miss-estimation. We further investigated the impact of noise misestimation on the accuracy of drift inference (S.I. Fig. 11). Our findings indicate that after two augmentations conditioned on the invariant geometry, small inaccuracies in the employed dynamical noise during the simulation of augmented paths had a negligible effect on the overall accuracy of the inferred drift. In particular, for small inter-observation intervals, our approach remained robust to misestimated noise amplitudes. As the inter-observation intervals increased, the effect of noise deviations from ground truth on the performance remained small, provided the noise used in the augmentation deviated by at most ± 0.1 from the ground truth noise amplitude. Thus, stochastic dynamics may still be identified with our method even with inaccurate or moderately misestimated diffusion constants.

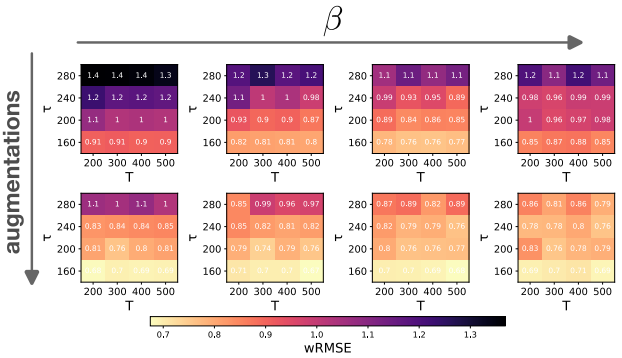

Figure 6

Sensitivity analysis with respect to the parameter β, which controls the relative weight of the geodesic control term. Weighted root mean square error values between ground truth and estimated drift, averaged over 3 independent runs, for different total trajectory lengths T and inter-observation intervals τ with noise amplitude $\sigma = 0.5$. Upper row shows results after the first iteration, while second row shows results after the second one. Values of β from left to right $[0.1, 0.5, 1.0, 2.0]$.

Additional experiments with three- and eight-dimensional systems are provided in Fig. 9 and Fig. 10. We further provide in the Supplement more ablations with respect to **(i)** the metric-learning algorithm (Sec. I.5), **(ii)** and noise misspecification (Sec. I.3).[1]

5. Discussion

Discovering unknown driving forces from sparse observations of stochastic systems remains still challenging: when the sampling interval is large, short-time transition approximations break down and the inverse problem of identifying the drift function becomes challenging. Our results show that combining the **temporal ordering** of the observations while accounting for the **geometry of the invariant density** can substantially improve drift recovery in nonlinear stochastic systems in non-conservative settings with steady-state probability currents.

Geometric inductive bias. We introduced **geometric inductive biases** into inference of stochastic systems by treating the deterministic flow field as a scaffold upon which system states fluctuate. We approximated this scaffold in terms of **distortions of a metric induced by the**

[1]Code is available at: https://github.com/dimitra-maoutsa/From_geometry_to_dynamics.

system's measurements. This approach effectively approximates the low-dimensional invariant density (empirical manifold) without the need to project to a lower dimensional space, whose dimensionality would be hard to estimate due to the presence of fluctuations.

The key insight is that **geodesics** computed on the empirical manifold with respect to the approximated metric provide geometry-informed reference curves for latent path augmentation. These curves are motivated by the **most probable paths** of the unknown system between consecutive observations in the Onsager-Machlup sense, but are not explicitly assumed to be the true unobserved paths. They rather serve as soft constraints induced by the approximated invariant geometry. Using these **geodesics as control constraints**, we formulated a path-augmentation scheme that bridges sparse observations with trajectories consistent with both the temporal order and the geometry of the data.

Why standard path augmentation fails at large τ. Widely used inference methods, predominantly developed within the statistics community, often employ path (*data*) augmentation to approximate transition densities between successive observations. However, this approach suffers from several challenges: **1.)** The unobserved information between successive observations is an infinite-dimensional object, requiring the solution of a complex and computationally intensive problem (bridge sampling) (Beskos et al., 2006a). We addressed this challenging problem using the computationally efficient framework developed in (Maoutsa and Opper, 2022). **2.)** Direct drift estimation from sparse observations results in estimated dynamics that significantly deviate from the ground truth. Thereby consecutive observations of the system have small probability under the law of the estimated SDE. This discrepancy, in turn, leads to several computational difficulties: **i)** Most bridge sampling schemes become computationally demanding, or fail, when attempting to generate transition densities between atypical states for the considered stochastic dynamics. For instance, the method of (Maoutsa and Opper, 2021) successfully generates transition densities between atypical states only for conservative systems through a reweighting with Brownian bridge dynamics. Alternatively, an exceedingly large number of samples would be required for accurate numerical approximation. **ii)** Iterative algorithms, such as Expectation Maximisation, which exhibit only local convergence (Romero et al., 2019), may converge to inaccurate solutions, when the initial estimation significantly deviates from the ground truth.

Relation to Schrödinger bridges and controlled transport. The framework we employed for the augmentation relies on a deterministic particle formulation of the path integral control formalism (Kappen, 2005b). This framework can be connected to the dynamic Schrödinger bridge problem, if we consider transferring probability mass between two Dirac measures or very narrow Gaussians that sit on each observation, considering additionally a potential that constraints the intermittent dynamics similar to the one considered in (Neklyudov et al., 2023a). Thus, in principle, one can employ one of the recently developed alternative frameworks that solve the dynamic Schrödinger bridge problem for path augmentation. Recent bridge- and flow-matching methods (Lipman et al., 2023; Albergo et al., 2025; Shi et al., 2023; Liu et al., 2024) solve related controlled-transport problems, but typically without the path constraint we employed here. In contrast, the Generalised Schrödinger Bridge Matching (GSBM) framework (Liu et al., 2024) uses a cost functional that is equivalent to the controlled cost we employ to construct our augmentations (SI Eq. 30). In this setting, the penalty term corresponds to the geodesic proximity constraint used in our framework. The GSBM could, in principle, replace the particle-based framework we use. We nevertheless adopted a particle-based formulation, because the resulting particles provide direct samples from the controlled path measures, and can therefore be reused to construct Monte Carlo approximations of the integrals entering the Gaussian-process drift inference (Eq. 40). Yet, the Gaussian variant of the GSBM framework that incorporates time-dependent penalty constraints (analogous to our geodesic constraints), might be an interesting avenue to explore for potential incorporation in our framework.

Limitations and outlook. The proposed approach relies on the geometric characterisation of the invariant density of the system's dynamics. This requires sufficiently long observation windows to accurately characterise said density and correctly approximate the unobserved paths with geodesic curves. Thus, our approach is limited to systems where the invariant density can be approximated by a manifold where we can identify geodesics. An alternative method worth exploring would consider the learned invariant metric directly in the dynamics of the augmented process. Our evaluation was restricted to synthetic systems with known ground-truth dynamics. While this allows us to isolate the effect of sparse sampling and directly quantify drift recovery, validating the method on real datasets with naturally sparse observations remains an important direction for future work. Moreover, we have considered here inference of stochastic differential equations with known state independent diffusion. While this approach might seem limited, several processes with state dependent diffusions can be transformed into processes with state independent diffusions (Beskos et al., 2006a; Roberts and Stramer, 2001) through the Lamperti transform if they fulfil the appropriate conditions for the drift function.

Acknowledgements

The author is indebted to Manfred Opper for prompting them to work on this problem, Nina Miolane for early guidance on extracting geodesics from variational autoencoder based approximation of data manifolds, Stefan Sommer for answering questions on diffusions on manifolds, and to Georgios Arvanitidis for maintaining an open repository with his previous work on approximating geodesics from data manifolds. The author gratefully acknowledges the anonymous reviewers for their careful reading of the manuscript and for their constructive feedback, which contributed to improving the final version of this manuscript. An earlier account of this work has been presented at the NeurIPS 2022 workshop *Machine Learning and the Physical Sciences* (Maoutsa, 2022), in ICLR 2023 workshop *Physics for Machine Learning* (**?**), and in ICLR 2026 workshop *Geometry-grounded Representation Learning and Generative Modeling* (Maoutsa). The author further acknowledges that previous work from the Python (Van Rossum and Drake Jr, 1995), numpy (Harris et al., 2020), scipy (Virtanen et al., 2020), matplotlib (Hunter, 2007), seaborn (Waskom, 2021), GPflow (Matthews et al., 2017), pyEMD (Pele and Werman, 2009), and pytorch (Paszke et al., 2017) communities facilitated the implementation of the computational part of this work.

Impact Statement

The aim of this work is to advance the field of dynamical inference for stochastic systems. While we do not foresee any direct societal consequences directly impinging from our work, we recognize that stochastic systems could be applied in military contexts, financial engineering, or more recently in machine learning for data (such as image, audio, video) generation. Still, the proposed method does not propose interventions that might lead to unfavourable societal outcomes. Overdamped Langevin systems are widespread in areas such as physics, biology, neuroscience, and ecology. We anticipate that our contributions will thus help these disciplines by offering a tool to identify and further study relevant systems.

Our contribution emphasises the importance of incorporating concepts from the evolving field of geometric statistics into system identification methods for stochastic systems. Although geometric and topological properties of invariant densities have been extensively studied in the context of deterministic systems, comparable attention is lacking for their stochastic counterparts. Our work further highlights that in settings where the amount of augmented data exceeds the number of observations, data augmentation frameworks can enhance inference accuracy by incorporating domain knowledge or other relevant information, such as the geometry of the system's invariant density we consider here. Many algorithms used for data augmentation, including the expectation maximisation algorithm employed in our work (Romero et al., 2019), show only **local convergence**. As a result, when the initial estimate deviates significantly from the ground truth value, naive data augmentation methods may converge to suboptimal solutions that fail to accurately identify the underlying system.

References

Albert Einstein. Über die von der molekularkinetischen Theorie der wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen. *Annalen der Physik*, 4, 1905. *(cited on page: 1)*

Tao Li. Chemical Langevin Equation for Complex Reactions. *The Journal of Physical Chemistry A*, 124(5):810–816, 2020. *(cited on page: 1)*

L Silva-Dias and A López-Castillo. Spontaneous Symmetry Breaking of Population: Stochastic Lotka–Volterra Model for Competition among two similar Preys and Predators. *Mathematical Biosciences*, 300:36–46, 2018. *(cited on page: 1)*

Charles K Fisher and Pankaj Mehta. The Transition between the Niche and Neutral Regimes in Ecology. *Proceedings of the National Academy of Sciences*, 111(36): 13111–13116, 2014. *(cited on page: 1)*

Antonio A Alonso, Ignacio Molina, and Constantinos Theodoropoulos. Modeling bacterial population growth from stochastic single-cell dynamics. *Applied and Environmental Microbiology*, 80(17):5241–5253, 2014. *(cited on page: 1)*

Andreas Ruttor, Philipp Batz, and Manfred Opper. Approximate Gaussian Process inference for the Drift Function in Stochastic Differential Equations. *Advances in Neural Information Processing Systems (NeurIPS 2013)*, 26, 2013. *(cited on pages: 1, 2, 23, 28, 29, 30, 32, and 39)*

Rudolf Friedrich and Joachim Peinke. Description of a Turbulent Cascade by a Fokker-Planck Equation. *Physical Review Letters*, 78(5):863, 1997. *(cited on pages: 1, 2, 22, 23, 32, and 39)*

Rudolf Friedrich, Joachim Peinke, Muhammad Sahimi, and M Reza Rahimi Tabar. Approaching Complexity by Stochastic Methods: From Biological Systems to Turbulence. *Physics Reports*, 506(5):87–162, 2011. *(cited on page: 1)*

Mario Ragwitz and Holger Kantz. Indispensable Finite Time Corrections for Fokker-Planck Equations from Time Series Data. *Physical Review Letters*, 87(25): 254501, 2001. *(cited on pages: 1, 22, 32, and 39)*

Mozes Jacobs, Bingni W Brunton, Steven L Brunton, J Nathan Kutz, and Ryan V Raut. Hypersindy: Deep Generative Modeling of Nonlinear Stochastic Governing Equations. *arXiv preprint arXiv:2310.04832*, 2023. *(cited on page: 1)*

Philipp Batz, Andreas Ruttor, and Manfred Opper. Variational Estimation of the Drift for Stochastic Differential Equations from the Empirical Density. *Journal of Statistical Mechanics: Theory and Experiment*, 2016(8): 083404, 2016. *(cited on pages: 1 and 32)*

Yiqi Gu, John Harlim, Senwei Liang, and Haizhao Yang. Stationary Density Estimation of Itô Diffusions using Deep Learning, 2021. *(cited on page: 1)*

Amit Singer and Ronald R Coifman. Non-linear Independent Component Analysis with Diffusion Maps. *Applied and Computational Harmonic Analysis*, 25(2):226–239, 2008. *(cited on pages: 1 and 33)*

Feliks Nüske, Péter Koltai, Lorenzo Boninsegna, and Cecilia Clementi. Spectral Properties of Effective Dynamics from Conditional Expectations. *Entropy*, 23(2):134, 2021. *(cited on page: 1)*

Edward L Ionides, Carles Bretó, and Aaron A King. Inference for Nonlinear Dynamical Systems. *Proceedings of the National Academy of Sciences*, 103(49):18438–18443, 2006. *(cited on page: 1)*

Ronen Talmon and Ronald R Coifman. Intrinsic Modeling of Stochastic Dynamical Systems using Empirical Geometry. *Applied and Computational Harmonic Analysis*, 39(1):138–160, 2015. *(cited on pages: 1 and 33)*

Carmeline J Dsilva, Ronen Talmon, C William Gear, Ronald R Coifman, and Ioannis G Kevrekidis. Data-driven Reduction for a Class of Multiscale fast-slow Stochastic Dynamical Systems. *SIAM Journal on Applied Dynamical Systems*, 15(3):1327–1351, 2016. *(cited on page: 1)*

Tyrus Berry and John Harlim. Iterated Diffusion Maps for Feature Identification. *Applied and Computational Harmonic Analysis*, 45(1):84–119, 2018. *(cited on page: 1)*

Tyrus Berry and John Harlim. Nonparametric Uncertainty Quantification for Stochastic Gradient Flows. *SIAM/ASA Journal on Uncertainty Quantification*, 3(1):484–508, 2015. *(cited on page: 1)*

Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22, 1977. *(cited on pages: 2 and 3)*

Philipp Batz, Andreas Ruttor, and Manfred Opper. Approximate Bayes Learning of Stochastic Differential Equations. *Physical Review E*, 98(2):022109, 2018. *(cited on pages: 2, 6, 8, 22, 24, 32, and 56)*

Roland Hostettler, Filip Tronarp, and Simo Särkkä. Modeling the Drift Function in Stochastic Differential Equations using Reduced Rank Gaussian Processes. *IFAC-PapersOnLine*, 51(15):778–783, 2018. *(cited on pages: 2 and 32)*

Zheng Zhao, Filip Tronarp, Roland Hostettler, and Simo Särkkä. State-space Gaussian process for drift estimation in stochastic differential equations. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5295–5299. IEEE, 2020. *(cited on pages: 2 and 32)*

Giorgos Sermaidis, Omiros Papaspiliopoulos, Gareth O Roberts, Alexandros Beskos, and Paul Fearnhead. Markov chain Monte Carlo for Exact Inference for Diffusions. *Scandinavian Journal of Statistics*, 40(2):294–321, 2013. *(cited on pages: 2, 22, 23, and 32)*

Christian Fröhlich, Alexandra Gessner, Philipp Hennig, Bernhard Schölkopf, and Georgios Arvanitidis. Bayesian Quadrature on Riemannian Data Manifolds. *International Conference on Machine Learning (ICML 2021)*, pages 3459–3468, 2021. doi: https://doi.org/10.48550/arXiv.2102.06645. *(cited on pages: 3 and 26)*

Joshua B Tenenbaum, Vin de Silva, and John C Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science*, 290(5500):2319–2323, 2000. *(cited on page: 3)*

Mukund Balasubramanian and Eric L Schwartz. The Isomap Algorithm and Topological Stability. *Science*, 295(5552):7–7, 2002. *(cited on page: 3)*

AL Mead. Review of the Development of Multidimensional Scaling Methods. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 41(1):27–39, 1992. *(cited on page: 3)*

Sam T Roweis and Lawrence K Saul. Nonlinear Dimensionality Reduction by Locally Linear Embedding. *Science*, 290(5500):2323–2326, 2000. *(cited on page: 4)*

Georgios Arvanitidis, Soren Hauberg, Philipp Hennig, and Michael Schober. Fast and Robust Shortest Paths on Manifolds Learned from Data. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1506–1515. PMLR, 2019. *(cited on pages: 4, 6, 26, 27, 51, 56, and 59)*

Manfredo Perdigao Do Carmo. *Riemannian Geometry*, volume 6. Springer, 1992. *(cited on pages: 4, 25, and 27)*

Georgios Arvanitidis, Lars Kai Hansen, and Søren Hauberg. Latent Space Oddity: on the Curvature of Deep Generative Models. In *International Conference on Learning Representations (ICLR 2017)*, 2017. *(cited on pages: 4, 27, 34, and 51)*

Matthew James Beal. *Variational Algorithms for Approximate Bayesian Inference*. University of London, University College London (United Kingdom), 2003. *(cited on pages: 5 and 24)*

Raphaël Chetrite and Hugo Touchette. Variational and Optimal Control Representations of Conditioned and Driven Processes. *Journal of Statistical Mechanics: Theory and Experiment*, 2015(12):P12001, 2015. *(cited on pages: 5 and 24)*

Satya N Majumdar and Henri Orland. Effective Langevin Equations for Constrained Stochastic Processes. *Journal of Statistical Mechanics: Theory and Experiment*, 2015 (6):P06039, 2015. *(cited on pages: 5 and 24)*

Manfred Opper. Variational Inference for Stochastic Differential Equations. *Annalen der Physik*, 531(3): 1800233, 2019. *(cited on pages: 5, 24, 25, 27, and 32)*

Hilbert J Kappen. Path Integrals and Symmetry Breaking for Optimal Control Theory. *Journal of Statistical Mechanics: theory and experiment*, 2005(11):P11011, 2005a. *(cited on page: 6)*

Dimitra Maoutsa and Manfred Opper. Deterministic Particle Flows for Constraining Stochastic Nonlinear Systems. *Phys. Rev. Research*, 4:043035, Oct 2022. *(cited on pages: 6, 9, 22, 27, 28, and 54)*

Kevin Course and Prasanth B Nair. State Estimation of a Physical System with unknown Governing Equations. *Nature*, 622(7982):261–267, 2023a. *(cited on pages: 6 and 56)*

Arshed Nabeel, Ashwin Karichannavar, Shuaib Palathingal, Jitesh Jhawar, David B Brückner, Danny Raj M, and Vishwesha Guttal. Discovering Stochastic Dynamical Equations from Ecological Time Series Data. *The American Naturalist*, 205(4):E100–E117, 2025. *(cited on pages: 6, 32, and 56)*

Kevin Course and Prasanth Nair. Amortized Reparametrization: Efficient and Scalable Variational Inference for Latent SDEs. *Advances in Neural Information Processing Systems (NeurIPS 2023)*, 36: 78296–78318, 2023b. *(cited on pages: 6, 33, and 56)*

Kacper Kapusniak, Peter Potaptchik, Teodora Reu, Leo Zhang, Alexander Tong, Michael Bronstein, Joey Bose, and Francesco Di Giovanni. Metric Flow Matching for Smooth Interpolations on the Data Manifold. *Advances in Neural Information Processing Systems (NeurIPS 2024)*, 37:135011–135042, 2024. *(cited on pages: 6, 7, 34, 51, and 56)*

Georgios Arvanitidis, Soren Hauberg, and Bernhard Schölkopf. Geometrically Enriched Latent Spaces. In *International Conference on Artificial Intelligence and Statistics*, pages 631–639. PMLR, 2021. *(cited on pages: 6 and 7)*

Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos A Theodorou, and Ricky TQ Chen. Generalized Schrödinger Bridge Matching. *International Conference for Learning Representations*, 2024. *(cited on pages: 6, 9, 33, 34, and 56)*

Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free Schrödinger Bridges via Score and Flow Matching. *International Conference for Learning Representations (ICLR 2024)*, 2024. *(cited on pages: 6, 33, and 56)*

Anna Frishman and Pierre Ronceray. Learning Force Fields from Stochastic Trajectories. *Physical Review X*, 10(2): 021009, 2020. *(cited on pages: 7, 32, and 55)*

Evgeny Evgenievich Selkov. Self-oscillations in Glycolysis 1. A Simple Kinetic Model. *European Journal of Biochemistry*, 4(1):79–86, 1968. *(cited on pages: 7 and 55)*

Alexandros Beskos, Omiros Papaspiliopoulos, Gareth O Roberts, and Paul Fearnhead. Exact and Computationally Efficient Likelihood-based Estimation for Discretely Observed Diffusion Processes (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):333–382, 2006a. *(cited on page: 9)*

Dimitra Maoutsa and Manfred Opper. Deterministic Particle Flows for Constraining SDEs. *Machine Learning and the Physical Sciences, Workshop at the 35th Conference on Neural Information Processing Systems (NeurIPS 2021), arXiv preprint arXiv:2110.13020*, 2021. *(cited on pages: 9, 22, 27, and 59)*

Orlando Romero, Sarthak Chatterjee, and Sérgio Pequito. Convergence of the Expectation-Maximization Algorithm through Discrete-time Lyapunov Stability Theory. In *2019 American Control Conference (ACC)*, pages 163–168. IEEE, 2019. *(cited on pages: 9 and 10)*

Hilbert J Kappen. Linear Theory for Control of Nonlinear Stochastic Systems. *Physical Review Letters*, 95(20): 200201, 2005b. *(cited on page: 9)*

Kirill Neklyudov, Rob Brekelmans, Alexander Tong, Lazar Atanackovic, Qiang Liu, and Alireza Makhzani. A

Computational Framework for Solving Wasserstein Lagrangian Flows. In *Forty-first International Conference on Machine Learning (ICML 2023)*, 2023a. *(cited on page: 9)*

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow Matching for Generative Modeling. *International Conference for Learning Representations (ICLR 2023)*, 2023. *(cited on pages: 9 and 33)*

Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic Interpolants: A Unifying Framework for Flows and Diffusions. *Journal of Machine Learning Research*, 26(209):1–80, 2025. *(cited on pages: 9 and 33)*

Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion Schrödinger Bridge Matching. *Advances in Neural Information Processing Systems (NeurIPS 2023)*, 36:62183–62223, 2023. *(cited on page: 9)*

Gareth O Roberts and Osnat Stramer. On Inference for Partially Observed Nonlinear Diffusion Models using the Metropolis–Hastings Algorithm. *Biometrika*, 88(3): 603–621, 2001. *(cited on page: 9)*

Dimitra Maoutsa. Geometric Path Augmentation for Inference of Sparsely Observed Stochastic Nonlinear Systems. *Machine Learning and the Physical Sciences workshop (NeurIPS 2022)*, 2022. *(cited on page: 10)*

Dimitra Maoutsa. Flow Curvature Explains Failed SDE Drift Estimation under Sparse Sampling. In *ICLR 2026 Workshop on Geometry-grounded Representation Learning and Generative Modeling*. *(cited on page: 10)*

Guido Van Rossum and Fred L Drake Jr. *Python Tutorial*, volume 620. Centrum voor Wiskunde en Informatica Amsterdam, The Netherlands, 1995. *(cited on page: 10)*

Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array Programming with NumPy. *Nature*, 585(7825):357–362, September 2020. *(cited on page: 10)*

Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. *(cited on page: 10)*

J. D. Hunter. Matplotlib: A 2D Graphics Environment. *Computing In Science & Engineering*, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55. *(cited on page: 10)*

Michael L Waskom. Seaborn: Statistical Data Visualization. *Journal of Open Source Software*, 6(60):3021, 2021. *(cited on page: 10)*

Alexander G. de G. Matthews, Mark van der Wilk, Tom Nickson, Keisuke. Fujii, Alexis Boukouvalas, Pablo León-Villagrá, Zoubin Ghahramani, and James Hensman. GPflow: A Gaussian Process Library using TensorFlow. *Journal of Machine Learning Research*, 18(40): 1–6, apr 2017. *(cited on page: 10)*

Ofir Pele and Michael Werman. Fast and Robust Earth Mover's Distances. In *2009 IEEE 12th International Conference on Computer Vision*, pages 460–467. IEEE, September 2009. *(cited on pages: 10 and 28)*

Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic Differentiation in PYtorch. 2017. *(cited on page: 10)*

Andrew Golightly and Darren J Wilkinson. Bayesian Inference for Nonlinear Multivariate Diffusion Models Observed with Error. *Computational Statistics & Data Analysis*, 52(3):1674–1693, 2008. *(cited on pages: 22 and 23)*

Omiros Papaspiliopoulos, Yvo Pokern, Gareth O Roberts, and Andrew M Stuart. Nonparametric Estimation of Diffusions: a Differential Equations Approach. *Biometrika*, 99(3):511–531, 2012. *(cited on pages: 22 and 23)*

Alexandros Beskos, Omiros Papaspiliopoulos, and Gareth O Roberts. Retrospective Exact Simulation of Diffusion Sample Paths with Applications. *Bernoulli*, 12(6):1077–1098, 2006b. *(cited on pages: 22 and 23)*

Siddhartha Chib, Michael K Pitt, and Neil Shephard. Likelihood Based Inference for Diffusion Driven State Space Models. *Por Clasificar*, pages 1–33, 2006. *(cited on pages: 22, 23, and 32)*

Valentin De Bortoli, Arnaud Doucet, Jeremy Heng, and James Thornton. Simulating Diffusion Bridges with Score Matching. *arXiv preprint arXiv:2111.07243*, 2021. *(cited on page: 22)*

Ge Liu, Peter F Craigmile, and Radu Herbei. A Study of the Data Augmentation Strategy for Stochastic Differential Equations. *Journal of Statistical Computation and Simulation*, 90(10):1753–1772, 2020. *(cited on page: 22)*

Robert S Liptser and Albert N Shiryaev. *Statistics of Random Processes II: Applications*, volume 6. Springer Science & Business Media, 2013. *(cited on page: 22)*

Carl Edward Rasmussen. Gaussian Processes in Machine Learning. In *Summer School on Machine Learning*, pages 63–71. Springer-Verlag, 2003. *(cited on pages: 23 and 29)*

Steven J Lade. Finite Sampling Interval Effects in Kramers–Moyal Analysis. *Physics Letters A*, 373(41): 3705–3709, 2009. *(cited on pages: 23 and 35)*

Christoph Honisch and Rudolf Friedrich. Estimation of Kramers-Moyal coefficients at low sampling rates. *Physical Review E*, 83(6):066701, 2011. *(cited on page: 23)*

John M Lee. *Introduction to Riemannian Manifolds*, volume 176. Springer, 2018. *(cited on page: 25)*

Stephen Wiggins. *Normally Hyperbolic Invariant Manifolds in Dynamical Systems*, volume 105. Springer Science & Business Media, 1994. *(cited on page: 26)*

Salah-Eldin A Mohammed and Michael KR Scheutzow. The Stable Manifold Theorem for Stochastic Differential Equations. *Annals of Probability*, pages 615–652, 1999. *(cited on page: 26)*

TV Girya and Igor Dmitrievich Chueshov. Inertial Manifolds and Stationary Measures for Stochastically Perturbed Dissipative Dynamical Systems. *Sbornik: Mathematics*, 186(1):29–45, 1995. *(cited on page: 26)*

Neil Fenichel. Persistence and Smoothness of Invariant Manifolds for Flows. *Indiana University Mathematics Journal*, 21(3):193–226, 1971. *(cited on page: 26)*

Ludwig Arnold. Stochastic Differential Equations as Dynamical Systems. In *Realization and Modelling in System Theory*, pages 489–495. Springer, 1990. *(cited on page: 26)*

Andrew Carverhill. Flows of Stochastic Dynamical Systems: Ergodic Theory. *Stochastics: An International Journal of Probability and Stochastic Processes*, 14(4): 273–317, 1985. *(cited on page: 26)*

Sebastian Reich. A Nonparametric Ensemble Transform Method for Bayesian Inference. *SIAM Journal on Scientific Computing*, 35(4):A2013–A2024, 2013. *(cited on page: 28)*

Michalis Titsias. Variational Learning of Inducing Variables in Sparse Gaussian Processes. *Artificial Intelligence and Statistics*, pages 567–574, 2009. *(cited on page: 29)*

Lehel Csató and Manfred Opper. Sparse on-line Gaussian Processes. *Neural computation*, 14(3):641–668, 2002. *(cited on page: 29)*

Lars Onsager and Stefan Machlup. Fluctuations and Irreversible Processes. *Physical Review*, 91(6):1505, 1953. *(cited on page: 30)*

Tooru Taniguchi and EGD Cohen. Onsager-Machlup Theory for Nonequilibrium Steady States and Fluctuation theorems. *Journal of Statistical Physics*, 126(1):1–41, 2007. *(cited on page: 31)*

Artur B Adib. Stochastic Actions for Diffusive Dynamics: Reweighting, Sampling, and Minimization. *The Journal of Physical Chemistry B*, 112(19):5910–5916, 2008. *(cited on page: 31)*

Robert Graham. Path Integral Formulation of General Diffusion Processes. *Zeitschrift für Physik B Condensed Matter*, 26(3):281–290, 1977. *(cited on page: 31)*

Ruslan Leontievich Stratonovich. On the Probability Functional of Diffusion Processes. *Selected Translations in Mathematical Statistics and Probability*, 10:273–286, 1971. *(cited on page: 31)*

Detlef Dürr and Alexander Bach. The Onsager-Machlup Function as Lagrangian for the Most Probable Path of a Diffusion Process. *Communications in Mathematical Physics*, 60(2):153–170, 1978. *(cited on page: 31)*

Y Takahashi and S Watanabe. The Probability Functionals (Onsager-Machlup Functions) of Diffusion Processes. In *Stochastic Integrals*, pages 433–463. Springer, 1981. *(cited on page: 31)*

Robert Graham. Onsager-Machlup Function of Nonlinear Non-Equilibrium Thermodynamics. In *Functional Integration*, pages 263–280. Springer, 1980. *(cited on page: 31)*

Erlend Grong and Stefan Sommer. Most Probable Flows for Kunita SDEs. *arXiv preprint arXiv:2209.03868*, 2022. *(cited on page: 31)*

Mireille Capitaine. On the Onsager-Machlup Functional for Elliptic Diffusion Processes. *Séminaire de Probabilités XXXIV*, pages 313–328, 2000. *(cited on page: 31)*

Bernt Øksendal. Stochastic Differential Equations. In *Stochastic Differential Equations*, pages 65–84. Springer, 2003. *(cited on page: 31)*

J Cremers and A Hübler. Construction of Differential Equations from Experimental Data. *Zeitschrift für Naturforschung A*, 42(8):797–802, 1987. *(cited on page: 32)*

Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering Governing Equations from Data by Sparse Identification of Nonlinear Dynamical Systems. *Proceedings of the National Academy of Sciences*, 113(15): 3932–3937, 2016. *(cited on page: 32)*

Bryan C Daniels and Ilya Nemenman. Automated Adaptive Inference of Phenomenological Dynamical Models. *Nature Communications*, 6(1):1–8, 2015. *(cited on page: 32)*

Jose Casadiego, Dimitra Maoutsa, and Marc Timme. Inferring Network Connectivity from Event Timing Patterns. *Physical Review Letters*, 121(5):054101, 2018. *(cited on page: 32)*

Kevin McGoff, Sayan Mukherjee, and Natesh Pillai. Statistical Inference for Dynamical Systems: A Review. *Statistics Surveys*, 9:209–252, 2015. *(cited on page: 32)*

Holger Kantz and Thomas Schreiber. *Nonlinear Time Series analysis*, volume 7. Cambridge University Press, 2004. *(cited on page: 32)*

Michael Schmidt and Hod Lipson. Distilling Free-form Natural Laws from Experimental Data. *Science*, 324 (5923):81–85, 2009. *(cited on page: 32)*

Andonis Gerardos and Pierre Ronceray. Principled Model Selection for Stochastic Dynamics. *Physical Review Letters*, 135(16):167401, 2025. *(cited on page: 32)*

Norman H Packard, James P Crutchfield, J Doyne Farmer, and Robert S Shaw. Geometry from a Time Series. *Physical review letters*, 45(9):712, 1980. *(cited on page: 32)*

J Crutchfield and B McNamara. Equations of Motion from a Data Series. *Complex systems*, 1:417–452, 1987. *(cited on page: 32)*

Felipe Miguel Aparicio Acosta. Radial Basis Function and Related Models: an Overview. *Signal Processing*, 45(1): 37–58, 1995. *(cited on page: 32)*

Michael Small and Chi Kong Tse. Minimum Description Length Neural Networks for Time Series Prediction. *Physical Review E*, 66(6):066701, 2002. *(cited on page: 32)*

Kevin Judd and Alistair Mees. On Selecting Models for Nonlinear Time Series. *Physica D: Nonlinear Phenomena*, 82(4):426–444, 1995. *(cited on page: 32)*

Michael Small and Kevin Judd. Comparisons of new Nonlinear Modeling Techniques with Applications to Infant Respiration. *Physica D: Nonlinear Phenomena*, 117(1-4):283–298, 1998. *(cited on page: 32)*

David B Brückner, Pierre Ronceray, and Chase P Broedersz. Inferring the dynamics of underdamped stochastic systems. *Physical Review Letters*, 125(5):058103, 2020. *(cited on page: 32)*

Eurika Kaiser, J Nathan Kutz, and Steven L Brunton. Sparse identification of nonlinear dynamics for model predictive control in the low-data limit. *Proceedings of the Royal Society A*, 474(2219):20180335, 2018. *(cited on page: 32)*

Josh Bongard and Hod Lipson. Automated reverse engineering of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 104(24):9943–9948, 2007. *(cited on page: 32)*

Ioannis G Kevrekidis, C William Gear, James M Hyman, Panagiotis G Kevrekidis, Olof Runborg, Constantinos Theodoropoulos, et al. Equation-free, coarse-grained multiscale computation: enabling microscopic simulators to perform system-level analysis. *Commun. Math. Sci*, 1(4):715–762, 2003. *(cited on page: 32)*

Constantinos Theodoropoulos, Yue-Hong Qian, and Ioannis G Kevrekidis. "Coarse" stability and bifurcation analysis using time-steppers: A reaction-diffusion example. *Proceedings of the National Academy of Sciences*, 97(18):9840–9843, 2000. *(cited on page: 32)*

Mauricio Álvarez, David Luengo, and Neil D. Lawrence. Latent Force Models. In David van Dyk and Max Welling, editors, *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5 of *Proceedings of Machine Learning Research*, pages 9–16, Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA, 16–18 Apr 2009. PMLR. *(cited on page: 32)*

Guido Sanguinetti, Neil D Lawrence, and Magnus Rattray. Probabilistic inference of transcription factor concentrations and gene-specific regulatory activities. *Bioinformatics*, 22(22):2775–2781, 2006. *(cited on page: 32)*

Simo Särkkä. The Use of Gaussian Processes in System Identification. *arXiv preprint arXiv:1907.06066*, 2019. *(cited on page: 32)*

Peiyuan Teng. Machine-learning quantum mechanics: Solving quantum mechanics problems using radial basis function networks. *Physical Review E*, 98(3):033305, 2018. *(cited on page: 32)*

Ravinder Bhattoo, Sayan Ranu, and NM Krishnan. Learning the Dynamics of Particle-based Systems with Lagrangian Graph Neural Networks. *arXiv preprint arXiv:2209.01476*, 2022. *(cited on page: 32)*

Thomas Jüngling, Thomas Lymburn, Thomas Stemler, Débora Corrêa, David Walker, and Michael Small. Reconstruction of complex dynamical systems from time series using reservoir computing. In *2019 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5. IEEE, 2019. *(cited on page: 32)*

Cagatay Yildiz, Markus Heinonen, Jukka Intosalmi, Henrik Mannerstrom, and Harri Lahdesmaki. Learning Stochastic Differential Equations with Gaussian Processes without Gradient Matching. In *2018 IEEE 28th International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6. IEEE, 2018. *(cited on page: 32)*

J Peinke, R Friedrich, and Antoon Naert. A new approach to characterize disordered structures. *Zeitschrift für Naturforschung A*, 52(8-9):588–592, 1997. *(cited on page: 32)*

Ting-Ting Gao, Baruch Barzel, and Gang Yan. Learning Interpretable Dynamics of Stochastic Complex Systems from Experimental Data. *Nature Communications*, 15 (1):6029, 2024. *(cited on page: 32)*

Rudolf Friedrich, Silke Siegert, Joachim Peinke, Marcus Siefert, Michael Lindemann, Jan Raethjen, Güntner Deuschl, Gerhard Pfister, et al. Extracting model equations from experimental data. *Physics Letters A*, 271(3): 217–222, 2000. *(cited on page: 32)*

Federica Ferretti, Victor Chardès, Thierry Mora, Aleksandra M Walczak, and Irene Giardina. Building General Langevin Models from Discrete Datasets. *Physical Review X*, 10(3):031018, 2020. *(cited on page: 32)*

Lorenzo Boninsegna, Feliks Nüske, and Cecilia Clementi. Sparse learning of stochastic dynamical equations. *The Journal of chemical physics*, 148(24):241723, 2018. *(cited on page: 32)*

Yunfei Huang, Youssef Mabrouk, Gerhard Gompper, and Benedikt Sabass. Sparse inference and active learning of stochastic differential equations from data. *arXiv preprint arXiv:2203.11010*, 2022. *(cited on page: 32)*

David Lamouroux and Klaus Lehnertz. Kernel-based regression of drift and diffusion coefficients of stochastic processes. *Physics Letters A*, 373(39):3507–3512, 2009. *(cited on page: 32)*

George J Jiang and John L Knight. A nonparametric approach to the estimation of diffusion processes, with an application to a short-term interest rate model. *Econometric Theory*, 13(5):615–645, 1997. *(cited on page: 32)*

Stefan Klus, Feliks Nüske, Sebastian Peitz, Jan-Hendrik Niemann, Cecilia Clementi, and Christof Schütte. Data-driven approximation of the Koopman generator: Model reduction, system identification, and control. *Physica D: Nonlinear Phenomena*, 406:132416, 2020. *(cited on page: 32)*

Richard Stanton. A nonparametric model of term structure dynamics and the market price of interest rate risk. *The Journal of Finance*, 52(5):1973–2002, 1997. *(cited on page: 32)*

Youngkyoung Bae, Seungwoong Ha, and Hawoong Jeong. Inferring the Langevin Equation with Uncertainty via Bayesian Neural Networks. *Chaos, Solitons & Fractals*, 197:116440, 2025. *(cited on page: 32)*

Lars P Hansen and Jose A Scheinkman. Back to the future: Generating moment implications for continuous-time Markov processes, 1993. *(cited on page: 32)*

Mario Ragwitz and Holger Kantz. Ragwitz and Kantz reply. *Physical Review Letters*, 89(14):149402, 2002. *(cited on page: 32)*

D Kleinhans, R Friedrich, A Nawroth, and J Peinke. An Iterative Procedure for the Estimation of Drift and Diffusion Coefficients of Langevin Processes. *Physics Letters A*, 346(1-3):42–46, 2005. *(cited on page: 32)*

David Kleinhans and Rudolf Friedrich. Maximum Likelihood Estimation of Drift and Diffusion Functions. *Physics Letters A*, 368(3-4):194–198, 2007. *(cited on page: 32)*

Bjørn Eraker. MCMC analysis of diffusion models with application to finance. *Journal of Business & Economic Statistics*, 19(2):177–191, 2001. *(cited on page: 32)*

Monica Billio, Alain Monfort, and Christian Robert. The simulated likelihood ratio method. *Institut National de la Statistique et des Etudes Economiques*, 1998. *(cited on page: 32)*

Lea Duncker, Gergo Bohner, Julien Boussard, and Maneesh Sahani. Learning interpretable continuous-time models of latent stochastic dynamical systems. In *International conference on machine learning*, pages 1726–1734. PMLR, 2019. *(cited on pages: 32 and 33)*

Prakhar Verma, Vincent Adam, and Arno Solin. Variational Gaussian process diffusion processes. In *International Conference on Artificial Intelligence and Statistics*, pages 1909–1917. PMLR, 2024. *(cited on page: 32)*

Cédric Archambeau, Manfred Opper, Yuan Shen, Dan Cornford, and John Shawe-Taylor. Variational Inference for Diffusion Processes. *Advances in Neural Information Processing Systems (NeurIPS 2007)*, 20:17–24, 2007. *(cited on page: 33)*

Xuechen Li, Ting-Kam Leonard Wong, Ricky T. Q. Chen, and David Duvenaud. Scalable Gradients for Stochastic Differential Equations, 2020. *(cited on page: 33)*

Yury A Kutoyants. *Statistical Inference for Ergodic Diffusion Processes*. Springer Science & Business Media, 2004. *(cited on page: 33)*

Ronald R Coifman, Stephane Lafon, Ann B Lee, Mauro Maggioni, Boaz Nadler, Frederick Warner, and Steven W Zucker. Geometric Diffusions as a Tool for Harmonic Analysis and Structure Definition of Data: Diffusion Maps. *Proceedings of the National Academy of Sciences of the United States of America*, 102(21): 7426–7431, 2005. *(cited on page: 33)*

Boaz Nadler, Stéphane Lafon, Ronald R Coifman, and Ioannis G Kevrekidis. Diffusion Maps, Spectral Clustering and Reaction Coordinates of Dynamical Systems. *Applied and Computational Harmonic Analysis*, 21(1): 113–127, 2006. *(cited on page: 33)*

Dimitrios Giannakis. Data-driven Spectral Decomposition and Forecasting of Ergodic Dynamical Systems. *Applied and Computational Harmonic Analysis*, 47(2):338–396, 2019. *(cited on pages: 33 and 35)*

Andrew L Ferguson, Athanassios Z Panagiotopoulos, Ioannis G Kevrekidis, and Pablo G Debenedetti. Nonlinear dimensionality reduction in molecular simulation: The diffusion map approach. *Chemical Physics Letters*, 509 (1-3):1–11, 2011. *(cited on page: 33)*

Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, and Ricky TQ Chen. Multisample Flow Matching: Straightening Flows with Minibatch Couplings. In *International Conference on Machine Learning*, 2023. *(cited on page: 33)*

Michael S Albergo and Eric Vanden-Eijnden. Building Normalizing Flows with Stochastic Interpolants. *International Conference for Learning Representations (ICLR 2023)*, 2023. *(cited on page: 33)*

Stephen Zhang and Michael Stumpf. Learning non-equilibrium Diffusions with Schrödinger Bridges: from exactly solvable to simulation-free. *Advances in Neural Information Processing Systems*, 38:119861–119898, 2026. *(cited on page: 33)*

Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023. *(cited on page: 33)*

Kirill Neklyudov, Rob Brekelmans, Daniel Severo, and Alireza Makhzani. Action Matching: Learning Stochastic Dynamics from Samples. In *International Conference on Machine Learning (ICML 2023)*, pages 25858–25889. PMLR, 2023b. *(cited on page: 33)*

Andres F Duque, Sacha Morin, Guy Wolf, and Kevin R Moon. Geometry regularized Autoencoders. *IEEE transactions on pattern analysis and machine intelligence*, 45(6):7381–7394, 2022. *(cited on page: 34)*

Dimitris Kalatzis, David Eklund, Georgios Arvanitidis, and Søren Hauberg. Variational Autoencoders with Riemannian Brownian Motion Priors. In *37th International Conference on Machine Learning (ICML 2020)*, pages 5020–5034. Curran Associates, Inc., 2021. *(cited on page: 34)*

Emile Mathieu and Maximilian Nickel. Riemannian Continuous Normalizing Flows. *Advances in neural information processing systems*, 33:2503–2515, 2020. *(cited on page: 34)*

Valentin De Bortoli, Emile Mathieu, Michael Hutchinson, James Thornton, Yee Whye Teh, and Arnaud Doucet. Riemannian Score-based Generative Modelling. *Advances in Neural Information Processing Systems (NeurIPS 2022)*, 35:2406–2422, 2022. *(cited on page: 34)*

Yunyi Shen, Renato Berlinghieri, and Tamara Broderick. Multi-marginal Schr\" odinger Bridges with Iterative Reference Refinement. *International Conference on Artificial Intelligence and Statistics, PMLR*, 2025. *(cited on page: 34)*

Nina Miolane, Nicolas Guigui, Alice Le Brigant, Johan Mathe, Benjamin Hou, Yann Thanwerdas, Stefan Heyder, Olivier Peltre, Niklas Koep, Hadi Zaatiti, et al. Geomstats: a Python Package for Riemannian Geometry in Machine Learning. *Journal of Machine Learning Research*, 21(223):1–9, 2020. *(cited on page: 35)*

Stefan Sommer. Probabilistic Approaches to Geometric Statistics: Stochastic Processes, Transition Distributions, and Fiber Bundle Geometry. In *Riemannian Geometric Statistics in Medical Image Analysis*, pages 377–416. Elsevier, 2020. *(cited on page: 35)*

Charles Fefferman, Sanjoy Mitter, and Hariharan Narayanan. Testing the Manifold Hypothesis. *Journal*

of the American Mathematical Society, 29(4):983–1049, 2016. *(cited on page: 35)*

Tal Shnitzer, Ronen Talmon, and Jean-Jacques Slotine. Manifold Learning for Data-Driven Dynamical System Analysis. In *The Koopman Operator in Systems and Control*, pages 359–382. Springer, 2020. *(cited on page: 35)*

Jared L Callaham, J-C Loiseau, Georgios Rigas, and Steven L Brunton. Nonlinear Stochastic Modelling with Langevin Regression. *Proceedings of the Royal Society A*, 477(2250):20210092, 2021. *(cited on page: 35)*

Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. TrajectoryNet: A Dynamic Optimal Transport Network for Modeling Cellular Dynamics. In *International Conference on Machine Learning (ICML 2020)*, pages 9526–9536. PMLR, 2020. *(cited on page: 35)*

Tal Shnitzer, Ronen Talmon, and Jean-Jacques Slotine. Manifold Learning with Contracting Observers for Data-driven Time-series Analysis. *IEEE Transactions on Signal Processing*, 65(4):904–918, 2016. *(cited on page: 35)*

Nelida Črnjarić-Žic, Senka Maćešić, and Igor Mezić. Koopman Operator Spectrum for Random Dynamical Systems. *Journal of Nonlinear Science*, 30(5):2007–2056, 2020. *(cited on page: 35)*

Dorothy Maharam. Invariant Measures and Radon-Nikodym Derivatives. *Transactions of the American Mathematical Society*, 135:223–248, 1969. *(cited on page: 35)*

David Ruelle. Ergodic theory of differentiable dynamical systems. *Publications Mathématiques de l'Institut des Hautes Études Scientifiques*, 50(1):27–58, 1979. *(cited on page: 35)*

David Ruelle. *Chaotic Evolution and Strange Attractors*, volume 1. Cambridge University Press, 1989. *(cited on page: 36)*

Lai-Sang Young. What are SRB Measures, and which Dynamical Systems have them? *Journal of Statistical Physics*, 108(5):733–754, 2002. *(cited on page: 36)*

Edward Ott. *Chaos in Dynamical Systems*. Cambridge University Press, 2002. *(cited on page: 36)*

Lai-Sang Young. Dimension, Entropy and Lyapunov Exponents. *Ergodic theory and dynamical systems*, 2(1): 109–124, 1982. *(cited on page: 36)*

François Ledrappier and Lai-Sang Young. The metric entropy of diffeomorphisms: part II: relations between entropy, exponents and dimension. *Annals of Mathematics*, 122(3):540–574, 1985. *(cited on page: 36)*

Valery Iustinovich Oseledets. A multiplicative ergodic theorem. Characteristic Ljapunov, exponents of dynamical systems. *Trudy Moskovskogo Matematicheskogo Obshchestva*, 19:179–210, 1968. *(cited on page: 36)*

Balázs Bárány and Antti Käenmäki. Ledrappier–Young Formula and Exact Dimensionality of Self-affine Measures. *Advances in Mathematics*, 318:88–129, 2017. *(cited on page: 36)*

David Ruelle. An Inequality for the Entropy of Differentiable Maps. *Boletim da Sociedade Brasileira de Matemática-Bulletin/Brazilian Mathematical Society*, 9 (1):83–87, 1978. *(cited on page: 36)*

Hannes Risken. *The Fokker-Planck Equation: Methods of Solution and Applications*. Springer-Verlag, 2nd edition, 1996. *(cited on page: 37)*

Lars Hörmander. Hypoelliptic Second Order Differential Equations. *Acta Mathematica*, 1967. *(cited on page: 37)*

Yuri Kifer. Random perturbations of dynamical systems. *Nonlinear Problems in Future Particle Accelerators*, 189, 1988. *(cited on page: 37)*

Ludwig Arnold. Random Dynamical Systems. In *Dynamical Systems: Lectures Given at the 2nd Session of the Centro Internazionale Matematico Estivo (CIME) held in Montecatini Terme, Italy, June 13–22, 1994*, pages 1–43. Springer, 2006. *(cited on page: 37)*

Arnulf Jentzen and Peter E Kloeden. *Taylor Approximations for Stochastic Partial Differential Equations*. SIAM, 2011. *(cited on page: 41)*

Wolfgang Kühnel. Differential Geometry: Curves - Surfaces - Manifolds, Second edition. *American Mathematical Society, Providence, RI*, 2, 2002. *(cited on page: 43)*

Christopher Scarvelis and Justin Solomon. Riemannian Metric Learning via Optimal Transport. In *International Conference on Learning Representations (ICLR 2023)*, 2023. *(cited on page: 51)*

Søren Hauberg, Oren Freifeld, and Michael Black. A Geometric Take on Metric Learning. *Advances in Neural Information Processing Systems (NeurIPS 2012)*, 25, 2012. *(cited on page: 51)*

Arnab Barua, Haralampos Hatzikirou, and Sumiyoshi Abe. Geodesic Learning. *Physica A: Statistical Mechanics and its Applications*, 669:130539, 2025. *(cited on page: 51)*

Samuel Gruffaz and Josua Sassen. Riemannian Metric Learning: Closer to you than you imagine. *arXiv preprint arXiv:2503.05321*, 2025. *(cited on page: 51)*

Evelyn Fix. *Discriminatory analysis: nonparametric discrimination, consistency properties*, volume 1. USAF School of Aviation Medicine, 1985. *(cited on page: 57)*

Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE transactions on information theory*, 13(1):21–27, 1967. *(cited on page: 57)*

Dimitra Maoutsa, Sebastian Reich, and Manfred Opper. Interacting Particle Solutions of Fokker–Planck Equations through Gradient–log–density Estimation. *Entropy*, 22 (8):802, 2020. *(cited on page: 59)*

Supplementary Information

A Drift inference for high and low frequency observations **21**

A.1 High frequency observations . . . 22

A.2 Low frequency observations . . . 23

A.3 Approximate posterior over paths. . . . 24

A.3.1 Approximate posterior over paths without geometric constraints . . . 24

A.3.2 Approximate posterior over paths with geometric constraints . . . 25

Riemannian geometry. . . . 25

Riemannian geometry of observations. . . . 26

Extended free energy functional. . . . 27

A.4 Approximate posterior over drift functions. . . . 28

B Sparse Gaussian process estimation **29**

C Theoretical evidence that supports the use of geodesics as geometric constraints representing most probable paths between observations **30**

D Does the proposed approach invalidate the Markovian property of the diffusion process? **31**

E Related work and positioning of the present work **32**

▷ Modelling general SDEs from state observations. . . . 32

▷ Modelling SDEs from population level snapshots/boundary conditions. . . . 33

Geometry aware generative methods. . . . 34

Approximating observation geometry in the ambient space. . . . 34

Positioning of the present work. . . . 34

F Geometric constraints on inference **35**

G Theoretical justification for Riemannian manifold approximation of the invariant density **35**

G.1 Dimensionality of invariant measures induced by deterministic dynamics . . . 35

G.2 Dimensionality of invariant measures induced by stochastic dynamics . . . 37

H Theoretical justification of geometric augmentation for large inter-observation intervals **39**

H.1 Inference performance deteriorates with increasing inter-observation interval for existing frameworks . . . 39

H.2 Inference based on Euler-Maruyama discretisation does not account for the curvature of the trajectories in the state space . . . 40

H.2.1 First remainder term $R_{1,a}$ . . . 41

First component $R^1_{1,a}$ of remainder term $R_{1,a}$: Flow curvature term. . . . 42

Second component $R^2_{1,a}$ of remainder term $R_{1,a}$. . . . 44

Third component $R^3_{1,a}$ of remainder term $R_{1,a}$. 45
H.2.2 Second remainder term $R_{1,b}$ 45
H.2.3 Third remainder term $R_{1,c}$ 46
H.2.4 Fourth remainder term $R_{1,d}$ 46

I Additional numerical results **47**
I.1 Inference for a 3-dimensional system 47
I.2 Inference for 8-dimensional system 48
I.3 Inference with noise miss-estimation 49
I.4 Convergence of inference over iterations 50
I.5 Ablations with respect to metric learning algorithm 51
I.6 Comparison of geodesic curves with Onsager–Machlup minimizers 51
Discrete OM action. 52
Numerical computation of the OM-minimiser. 53
Computation of action gap and relative action gap. 53

J Details on numerical experiments **54**
J.1 On computation of geodesic curves 56
J.2 Details on baseline methods 56

K Computational complexity of the proposed method **57**
Metric estimation. 58
Geodesic construction. 58
Geometry-guided bridge augmentation. 58
Sparse Gaussian-process drift estimation. 58
Overall cost. 59

L Algorithmic details **59**

A. Drift inference for high and low frequency observations

Effective dynamics of systems with many degrees of freedom or inherently stochastic are often described in terms of a stochastic differential equation (SDE)

$$\mathrm{d}\mathbf{X}_t = \mathbf{f}(\mathbf{X}_t)\mathrm{d}t + \boldsymbol{\eta}(t)\mathrm{d}t = \mathbf{f}(\mathbf{X}_t)\mathrm{d}t + \boldsymbol{\sigma}\mathrm{d}\mathbf{W}_t, \tag{11}$$

where the drift $\mathbf{f}(\cdot) : \mathbb{R}^d \to \mathbb{R}^d$ describes the deterministic forces acting on the system, while the delta-correlated Gaussian white noise term $\boldsymbol{\eta}(t)$, $\langle \boldsymbol{\eta}(t)\boldsymbol{\eta}(t')\rangle = \boldsymbol{\sigma}\boldsymbol{\sigma}^\top \delta(t-t')$ describes the effect of stochastic forces as a product of a diffusion matrix (or constant) $\boldsymbol{\sigma} : \mathbb{R}^{d\times d}$ that accounts for the magnitude of the stochastic forces acting on the system, and a d-dimensional Wiener process $\mathbf{W}_t$ that contributes random influences.

Often the detailed equation that governs the evolution of the state of the system is unknown. Therefore, understanding a system of interest often requires identification from time series observations of its state. In more practical terms, given some **prior probability** for the drift function, we want to compute the **posterior probability** $\mathrm{P}(\mathbf{f}|\{\mathcal{O}_k\}_{k=1}^K)$ that identifies

the unknown drift function of Eq. 11 that most likely gave rise to the observations of the system state $\{\boldsymbol{\mathcal{O}}_k\}_{k=1}^K$. The exact relationship between the observations and the system state will be defined more precisely in the following.

When a system is observed nearly continuously (inter-observation interval length τ much smaller than the characteristic time scale of the system $\tau \ll \tau_{\text{char}}$), temporal methods regress the system state $\mathbf{X}_t$ against the state increments $\mathbf{Y}_t \doteq \frac{\mathbf{X}_{t+\tau}-\mathbf{X}_t}{\tau}$ to identify the drift function (Friedrich and Peinke, 1997; Ragwitz and Kantz, 2001). In a Bayesian framework, this corresponds to Gaussian process regression with a Gaussian likelihood (SI A.1). However, for large inter-observation intervals τ, these methods fail (Batz et al., 2018), as the Gaussian likelihood assumption is invalid for general nonlinear systems with sparse observations (Fig.1**C.**). In such cases, the likelihood is a *path integral* over continuous trajectories of the unobserved process (SI A.2), making Gaussian-based estimates inaccurate (Fig. 1**C.**).

This underwhelming performance has motivated the development of methods that combine state estimation (or **path augmentation**) and dynamical inference. These methods reconstruct continuous paths to approximate transition densities between observations, enabling inference by estimating the system's state between observations. However, for large time intervals, transition densities are usually analytically intractable, except in a few trivial cases of scalar or linear processes. As a result, the prevailing strategy is to approximate transition densities by sampling marginal distributions of **diffusion bridges**, which are diffusion processes constrained by their initial and terminal states (Golightly and Wilkinson, 2008; Papaspiliopoulos et al., 2012; Sermaidis et al., 2013; Beskos et al., 2006b; Chib et al., 2006). Yet, existing methods employ path augmentation with simplified bridge dynamics (e.g., Brownian (Chib et al., 2006; Golightly and Wilkinson, 2008) or Ornstein-Uhlenbeck bridges (Batz et al., 2018)) that do not accurately reflect the underlying transition densities for nonlinear systems (Fig. 1**E.**).

An alternative path augmentation strategy would obtain a coarse drift estimate, typically achieved by assuming a Gaussian likelihood between observations (see SI Eq. 14), and would subsequently employ a stochastic bridge sampler (De Bortoli et al., 2021; Maoutsa and Opper, 2022; 2021) to construct stochastic bridges using the coarsely estimated nonlinear drift. However, for large inter-observation intervals, the coarsely estimated drift function often deviates significantly from the true function that generated the observations. Consequently, the observations frequently fall into low-probability regions of the estimated diffusion dynamics (Fig. 1 **E.**), rendering the construction of diffusion bridges either too computationally demanding or impossible (Liu et al., 2020).

A.1. High frequency observations

In an optimal but rather practically unrealistic scenario, we would observe the system (path) $\mathbf{X}_{0:T}$ in (nearly) continuous time, and thus we would try to identify the drift from $\mathrm{P}(\mathbf{f}|X_{0:T})$. In such a case, the infinitesimal transition probabilities of the diffusion process between consecutive time-points are Gaussian, i.e.,

$$\mathrm{P}_f(\mathbf{X}_{0:T} \mid \mathbf{f}) \propto \exp\left(-\frac{1}{2\mathrm{d}t}\sum_t \|\mathbf{X}_{t+\mathrm{d}t} - \mathbf{X}_t - \mathbf{f}(\mathbf{X}_t)\mathrm{d}t\|_D^2\right). \tag{12}$$

Here we have introduced the weighted norm $\|\mathbf{u}\|_D \doteq \mathbf{u}^\top \cdot \mathbf{D}^{-1} \cdot \mathbf{u}$, with $\mathbf{D} \doteq \boldsymbol{\sigma}\boldsymbol{\sigma}^\top$ indicating the noise covariance.

In turn, the transition probabilities of a discretised drift-less process (a Wiener path) $P_{\mathcal{W}}(\mathbf{X}_{0:T})$ with same diffusion σ is

$$\mathrm{P}_{\mathcal{W}}(\mathbf{X}_{0:T}) = \exp\left(-\frac{1}{2\mathrm{d}t}\sum_t \|\mathbf{X}_{t+\mathrm{d}t} - \mathbf{X}_t\|_D^2\right). \tag{13}$$

We can thus express the likelihood for the drift f as the likelihood ratio between the transition probabilities of Eq. 12 and Eq. 13, which for diffusion processes is expressed by the Radon-Nykodym derivative between $\mathrm{P}_f(\mathbf{X}_{0:T}|f)$ and $\mathrm{P}_{\mathcal{W}}(\mathbf{X}_{0:T})$ for paths $\mathbf{X}_{0:T}$ within the time interval $[0,\, T]$ (Liptser and Shiryaev, 2013)

$$\mathcal{L}(\mathbf{X}_{0:T} \mid \mathbf{f}) = \exp\left(-\frac{1}{2}\sum_t \|\mathbf{f}(\mathbf{X}_t)\|_D^2 \mathrm{d}t + \sum_t \langle \mathbf{f}(\mathbf{X}_t), \mathbf{X}_{t+\mathrm{d}t} - \mathbf{X}_t\rangle_D\right), \tag{14}$$

where for brevity we have introduced the notation $\langle \mathbf{u}, \mathbf{v}\rangle_D \doteq \mathbf{u}^\top \cdot \mathbf{D}^{-1} \cdot \mathbf{v}$ for the weighted inner product with respect to the inverse noise covariance $\mathbf{D}^{-1}$. This expression results from applying the Girsanov theorem on the path measures induced by a process with drift $\mathbf{f}$ and a Wiener process, with same diffusion σ, and employing an Euler-Maruyama discretisation on the continuous path $\mathbf{X}_{0:T}$.

The likelihood of a continuously observed path of the SDE (Eq. 14) has a quadratic form in terms of the drift function. Therefore a Gaussian measure over function values (Gaussian process) is a natural conjugate prior for this likelihood. Thus, to identify the drift in a non-parametric form, we assume a Gaussian process prior for the function values $\mathbf{f} \sim \mathrm{P}_0(\mathbf{f}) = \mathcal{GP}(\mathbf{m}^f, k^f)$, where $\mathbf{m}^f$ and k^f denote the mean and covariance function of the Gaussian process (Ruttor et al., 2013). The prior measure can be written as

$$\mathrm{P}_0(\mathbf{f}) = \exp\left(-\frac{1}{2}\int\int \mathbf{f}(\mathbf{x})\left(k^f(\mathbf{X},\mathbf{X}')\right)^{-1}\mathbf{f}(\mathbf{X}')\mathrm{d}\mathbf{X}\mathrm{d}\mathbf{X}'\right), \tag{15}$$

if we consider a zero mean Gaussian process $\mathbf{m}^f = \mathbf{0}$.

Bayesian inference for the drift function $\mathbf{f}$ requires the computation of a probability distribution in the function space, the posterior probability distribution $\mathrm{P}_f(\mathbf{f} \mid \mathbf{X}_{0:T})$. From the Bayes' rule the posterior can be written as

$$\mathrm{P}_f(\mathbf{f} \mid \mathbf{X}_{0:T}) = \frac{\mathrm{P}_0(\mathbf{f})\mathcal{L}(\mathbf{X}_{0:T} \mid \mathbf{f})}{Z} \propto \mathrm{P}_0(\mathbf{f})\mathcal{L}(\mathbf{X}_{0:T} \mid \mathbf{f}), \tag{16}$$

where Z denotes a normalising factor defined as

$$Z = \int \mathrm{P}_0(\mathbf{f})\mathcal{L}(\mathbf{X}_{0:T} \mid \mathbf{f})\mathcal{D}\mathbf{f}, \tag{17}$$

where $\mathcal{D}\mathbf{f}$ denotes integration over the Hilbert space $\mathbf{f} : H_0[\mathbf{f}] < \infty$. Here we have expressed the prior probability over functions as $\mathrm{P}_0(\mathbf{f}) = e^{-H_0[\mathbf{f}]}$. In (Ruttor et al., 2013) the authors show that in this continuous-time setting, nonparametric estimation of the drift can be attained through a Gaussian process regression (Rasmussen, 2003) with the objective to identify the mapping from the system state $\mathbf{X}_t$ to state increments $\mathrm{d}\mathbf{X}_t$. More precisely, we consider as the regressor the N observations of the system state $\mathbf{X}_t$ and as the associated response variables the state increments

$$\mathbf{Y}_t = \frac{\mathbf{X}_{t+\mathrm{d}t} - \mathbf{X}_t}{\mathrm{d}t}, \tag{18}$$

and select the kernel function of the Gaussian process as $k^f(\mathbf{X},\mathbf{X}')$.

If we denote with $\mathcal{X} = \{\mathbf{X}_t\}_{t=0}^{T-\mathrm{d}t}$ and $\mathcal{Y} = \{\mathbf{Y}_t\}_{t=0}^{T-\mathrm{d}t}$ the set of state observations and observation increments, the mean of the posterior process over drift functions $\mathbf{f}$ can be expressed as

$$\bar{\mathbf{f}}(\mathbf{x}) = k^f(\mathbf{x},\mathcal{X})^\top\left(\mathcal{K} + \frac{\mathbf{D}}{\mathrm{d}t}I_N\right)^{-1}\mathcal{Y}, \tag{19}$$

where we abused the notation and denoted with $k^f(\mathbf{x},\mathcal{X})$ the vector resulting from evaluating the kernel k^f at points $\mathbf{x}$ and $\{\mathcal{O}_k\}_{k=1}^{K-1}$. Similarly $\mathcal{K} = k^f(\mathcal{X},\mathcal{X})$ stands for the $(K-1)\times(K-1)$ matrix resulting from evaluation of the kernel on all observation pairs. In a similar vein, the posterior variance can be written as

$$\Sigma^2(\mathbf{x}) = k^f(\mathbf{x},\mathbf{x}) - k^f(\mathbf{x},\mathcal{X})^\top\left(\mathcal{K} + \frac{\mathbf{D}}{\mathrm{d}t}\right)^{-1}k^f(\mathbf{x},\mathcal{X}), \tag{20}$$

where the term $\mathbf{D}/\mathrm{d}t$ plays the role of observation noise.

A.2. Low frequency observations

As the inter-observation interval increases (*low frequency observations*), the validity of the Gaussian likelihood used in drift estimation diminishes as the transition density is no longer Gaussian. Consequently, methods for drift estimation with Gaussian assumptions (Friedrich and Peinke, 1997; Ruttor et al., 2013) become increasingly inaccurate. To discount the effects of low frequency sampling, Lade (Lade, 2009) proposed a method to compute finite-time corrections for drift estimates, which has been mainly applied to one-dimensional problems (Honisch and Friedrich, 2011). In parallel, the statistics community has proposed path augmentation techniques that involve sampling with a simplified system's dynamics between time-consecutive observations to augment the observed trajectory to a nearly continuous-time path (Golightly and Wilkinson, 2008; Papaspiliopoulos et al., 2012; Sermaidis et al., 2013; Beskos et al., 2006b; Chib et al., 2006). However,

for large inter-observation intervals and nonlinear systems, the augmented trajectories match poorly the underlying path statistics and these methods often exhibit poor convergence rates or fail to identify the correct dynamics (Figure 1 c. and d.). We note that path augmentation using Ornstein-Uhlenbeck bridges and local linearisation of the **ground truth** dynamics provides a reasonable approximation of the underlying transition density up to a certain inter-observation interval. Nevertheless, during inference, the ground truth dynamics is unknown, and the proposed local linearisations based on inaccurate drift estimates (Batz et al., 2018) perform poorly in this sparsely sampled regime.

As the inter-observation interval τ increases, if the system is nonlinear, the likelihood assumed between two consecutive observations is no longer Gaussian, but is rather expressed as a *path integral*

$$\mathrm{P}(\boldsymbol{\mathcal{O}}_{1:K} \mid \mathbf{f}) = \int \mathrm{P}(\boldsymbol{\mathcal{O}}_{1:K} \mid \mathbf{X}_{0:T})\mathrm{P}(\mathbf{X}_{0:T} \mid \mathbf{f})\mathcal{D}(\mathbf{X}_{0:T}), \tag{21}$$

where $\boldsymbol{\mathcal{O}}_{1:K} \doteq \{\boldsymbol{\mathcal{O}}_k\}_{k=1}^{K}$ identifies the set of K observations collected within the interval $[0,\, T]$, $\mathrm{P}(\mathbf{X}_{0:T} \mid \mathbf{f})$ the prior path probability resulting from a diffusion process with drift $\mathbf{f}(\mathbf{x})$, $\mathcal{D}(\mathbf{X}_{0:T})$ identifies the formal volume element on the path space, and $\mathrm{P}(\boldsymbol{\mathcal{O}}_{1:K} \mid \mathbf{X}_{0:T})$ stands for the likelihood of observations given the latent path $\mathbf{X}_{0:T}$.

However, the path integral of Eq. 21 is in general intractable for nonlinear systems. thus we need to simultaneously estimate the drift and latent state of the diffusion process, i.e., to approximate the joint posterior measure of latent paths and drift functions $\mathrm{P}(\mathbf{X}_{0:T}, \mathbf{f} \mid \boldsymbol{\mathcal{O}}_{1:K})$. Therefore we consider the unobserved continuous path $\mathbf{X}_{0:T}$ as latent random variables and employ an Expectation Maximisation (E.M.) algorithm to identify a maximum a posteriori estimate for the drift function. More precisely, we follow an iterative algorithm, where at each iteration n we alternate between the two following steps: An **Expectation** step, where given a drift estimate $\hat{\mathbf{f}}^n(\mathbf{x})$ we construct an approximate posterior over the latent variables $Q(\mathbf{X}_{0:T}) \approx \mathrm{P}(\mathbf{X}_{0:T} \mid \boldsymbol{\mathcal{O}}_{1:K}, \hat{\mathbf{f}}^n(\mathbf{x}))$, and compute the expected log-likelihood of the augmented path

$$\mathfrak{L}\big(\hat{\mathbf{f}}^n(\mathbf{x}), Q\big) = \mathbb{E}_Q\Big[\ln \mathscr{L}\big(\mathbf{X}_{0:T}, \boldsymbol{\mathcal{O}}_{1:K} \mid \hat{\mathbf{f}}^n(\mathbf{x})\big)\Big]. \tag{22}$$

A **Maximisation** step, where we update the drift estimation by maximising the expected log likelihood

$$\mathbf{f}^{n+1}(\mathbf{x}) = \arg\max_{f} \Big[\mathfrak{L}\big(\mathbf{f}^n(\mathbf{x}), Q\big) - \ln \mathrm{P}_0\big(\mathbf{f}^n(\mathbf{x})\big)\Big]. \tag{23}$$

In Eq. 23, P_0 denotes the Gaussian process prior over function values.

A.3. Approximate posterior over paths.

To obtain an approximate posterior over the latent paths we perform **variational inference** (Beal, 2003). In this section, we first formulate the approximate posterior over paths (conditional distribution for the path given the observations) by considering only individual observations as constraints (Section A.3.1). However, this approach results computationally taxing calculations during path augmentation, since the observations are atypical states of the initially estimated drift. To overcome this issue, we subsequently extend the formalism (Section A.3.2) to incorporate constraints that consider also the local geometry of the observations.

A.3.1. Approximate posterior over paths <u>without</u> geometric constraints

Given a drift function (or a drift estimate) $\hat{\mathbf{f}}(\mathbf{x})$ we can apply variational techniques to approximate the posterior measure over the latent path conditioned on the observations $\boldsymbol{\mathcal{O}}_{1:K}$. We consider that the **prior process** (the process without considering the observations $\boldsymbol{\mathcal{O}}_{1:K}$) is described by the equation

$$\mathrm{P}(\mathbf{X}_{0:T} \mid \hat{\mathbf{f}}): \qquad \mathrm{d}\mathbf{X}_t = \hat{\mathbf{f}}(\mathbf{X}_t)\mathrm{d}t + \sigma \mathrm{d}\mathbf{W}_t. \tag{24}$$

We will define an approximating (posterior) process that is conditioned on the observations. The conditioned process is also a diffusion process with the same diffusion as Eq. 24 but with a modified, time-dependent drift $g(x, t)$ that accounts for the observations (Chetrite and Touchette, 2015; Majumdar and Orland, 2015). We identify the approximate posterior measure Q with the posterior measure induced by an approximating process that is conditioned by the observations $\boldsymbol{\mathcal{O}}_{1:K}$ (Opper, 2019), with governing equation

$$Q(\mathbf{X}_{0:T}): \qquad \mathrm{d}\mathbf{X}_t = \mathbf{g}(\mathbf{X}_t, t)\mathrm{d}t + \sigma\, \mathrm{d}W_t = \Big(\hat{\mathbf{f}}(\mathbf{X}_t) + \mathbf{u}(\mathbf{X}_t, t)\Big)\, \mathrm{d}t + \sigma\, \mathrm{d}\mathbf{W}_t. \tag{25}$$

The effective drift $\mathbf{g}(\mathbf{X}_t, t)$ of Eq. 25 may be obtained from the solution of the variational problem of minimising the free energy

$$\mathcal{F}[Q] = \mathcal{KL}\Big(Q(\mathbf{X}_{0:T})||\mathrm{P}(\mathbf{X}_{0:T} \mid \hat{\mathbf{f}})\Big) - \sum_{k=1}^{K} \Big\langle \ln \mathrm{P}(\boldsymbol{\mathcal{O}}_k \mid \mathbf{X}_{t_k}) \Big\rangle_Q. \tag{26}$$

By applying the Cameron-Girsanov-Martin theorem we can express the Kullback-Leibler divergence between the two path measures induced by the diffusions with drift $\hat{\mathbf{f}}(\mathbf{x})$ and $\mathbf{g}(\mathbf{x}, t)$ as

$$\begin{aligned}
\mathcal{KL}\Big(Q(\mathbf{X}_{0:T})||\mathrm{P}(\mathbf{X}_{0:T}|\hat{\mathbf{f}})\Big) &= \left\langle \ln \left(\frac{\mathrm{d}Q(\mathbf{X}_{0:T})}{\mathrm{d}P\left(\mathbf{X}_{0:T}|\hat{\mathbf{f}}\right)} \right) \right\rangle_Q && (27)\\
&= \left\langle \left(-\frac{1}{2}\int_0^T \|\hat{\mathbf{f}}(\mathbf{X}_t) - \mathbf{g}(\mathbf{X}_t, t)\|^2_{\mathbf{D}} \mathrm{d}t + \int_0^T \frac{\hat{\mathbf{f}}(\mathbf{X}_t) - \mathbf{g}(\mathbf{X}_t, t)}{\mathbf{D}} \mathrm{d}\mathbf{W}_t \right) \right\rangle_Q \\
&= \left\langle \left(-\frac{1}{2}\int_0^T \|\hat{\mathbf{f}}(\mathbf{X}_t) - \mathbf{g}(\mathbf{X}_t, t)\|^2_{\mathbf{D}} \mathrm{d}t + V_T \right) \right\rangle_Q && (28)\\
&= \frac{1}{2}\int_0^T \int \|\mathbf{g}(\mathbf{x}, t) - \hat{\mathbf{f}}(\mathbf{x})\|^2_{\mathbf{D}}\, q_t(\mathbf{x})\, \mathrm{d}\mathbf{x}\, \mathrm{d}t + \mathfrak{C}, && (29)
\end{aligned}$$

where $q_t(\mathbf{x})$ stands for the marginal density for $\mathbf{X}_t$ of the approximate process. In the third line we have introduced the random variable $V_T = \int_0^T \frac{\hat{\mathbf{f}}(\mathbf{X}_t) - \mathbf{g}(\mathbf{X}_t, t)}{\mathbf{D}} \mathrm{d}\mathbf{W}_t$. Under the assumption that the function $\ell(\mathbf{X}_t) = \hat{\mathbf{f}}(\mathbf{X}_t) - \mathbf{g}(\mathbf{X}_t, t)$ is bounded, piece-wise continuous, and in $L^2[0, \infty)$, V_T follows the distribution $\mathcal{N}\left(V_T \mid 0, \int_0^T \ell^2(s)\mathrm{d}s\right)$, which for a given T will result into a constant $\mathfrak{C}$. Thus the second term in Eq. 29 is not relevant for the minimisation of the free energy and will be omitted.

We can thus express the free energy of Eq. 26 as (Opper, 2019)

$$\mathcal{F}[Q] = \frac{1}{2}\int_0^T \int \Big[\|\mathbf{g}(\mathbf{x}, t) - \hat{\mathbf{f}}(\mathbf{x})\|^2_{\mathbf{D}} + U(\mathbf{x}, t)\Big]\, q_t(\mathbf{x})\, \mathrm{d}\mathbf{x}\, \mathrm{d}t, \tag{30}$$

where the term $U(\mathbf{x}, t)$ accounts for the observations $U(\mathbf{x}, t) = -\sum_{t_k} \ln \mathrm{P}(\boldsymbol{\mathcal{O}}_k \mid \mathbf{x})\, \delta(t - t_k)$.

The minimisation of the functional of the free energy can be construed as a stochastic control problem (Opper, 2019) with the objective to identify a time-dependent drift adjustment $\mathbf{u}(\mathbf{x}, t) := \mathbf{g}(\mathbf{x}, t) - \hat{\mathbf{f}}(\mathbf{x})$ for the system with drift $\hat{\mathbf{f}}(\mathbf{x})$ so that the controlled dynamics fulfil the constraints imposed by the observations.

A.3.2. Approximate posterior over paths with geometric constraints

The previously described construction of the approximate measure in terms of stochastic bridges is relevant when the observations have non vanishing probability under the law of the prior diffusion process of Eq. 24. However, when the prior process (with the estimated drift $\hat{f}$) differs considerably from the process that generated the observations, such a construction might either provide a bad approximation of the underlying path measure, or show slow numerical convergence in the construction of the diffusion bridges. To overcome this issue, we consider here additional constraints for the posterior process that force the paths of the posterior measure to respect the local geometry of the observations. In the following we provide a brief introduction on the basics of Riemannian geometry and consequently continue with the geometric considerations of the proposed method.

Riemannian geometry. A d-dimensional **Riemannian manifold** (Do Carmo, 1992; Lee, 2018) $(\mathcal{M}, \mathfrak{h})$ embedded in a d-dimensional ambient space $\mathcal{X} = \mathbb{R}^d$ is a smooth curved d-dimensional surface endowed with a smoothly varying inner product (Riemannian) **metric** $\mathfrak{h} : \mathbf{x} \to \langle \cdot | \cdot \rangle_{\mathbf{x}}$ on $\mathcal{T}_{\mathbf{x}}\mathcal{M}$. A tangent space $\mathcal{T}_{\mathbf{x}}\mathcal{M}$ is defined at each point $\mathbf{x} \in \mathcal{M}$. The Riemannian metric $\mathfrak{h}$ defines a canonical volume measure on the manifold $\mathcal{M}$. Intuitively this characterises how to

compute inner products locally between points on the tangent space of the manifold $\mathcal{M}$, and therefore determines also how to compute norms and thus distances between points on $\mathcal{M}$.

A **coordinate chart** (G, ϕ) provides the mapping from an open set G on $\mathcal{M}$ to an open set V in the Euclidean space. The dimensionality of the manifold is d if for each point $\mathbf{x} \in \mathcal{M}$ there exists a local neighborhood $G \subset \mathbb{R}^d$. We can represent the metric $\mathfrak{h}$ on the local chart (G, ϕ) by the positive definite matrix (**metric tensor**) $H(\mathbf{x}) = (\mathfrak{h}_{i,j})_{\mathbf{x}, 0 \leq i,j, \leq d} = \left(\langle \frac{\partial}{\partial x_i} | \frac{\partial}{\partial x_j} \rangle_{\mathbf{x}} \right)_{0 \leq i,j, \leq d}$ at each point $\mathbf{x} \in G$.

For $\mathbf{v}, \mathbf{w} \in \mathcal{T}_{\mathbf{x}}\mathcal{M}$ and $\mathbf{x} \in G$, their inner product can be expressed in terms of the matrix representation of the metric $\mathfrak{h}$ on the tangent space $\mathcal{T}_{\mathbf{x}}\mathcal{M}$ as $\langle \mathbf{v} | \mathbf{w} \rangle_{\mathbf{x}} = \mathbf{v}^\top H(\mathbf{x}) \mathbf{w}$, where $H(\mathbf{x}) \in \mathbb{R}^{d \times d}$.

The **length of a curve** $\gamma : [0, 1] \to \mathcal{M}$ on the manifold is defined as the integral of the norm of the tangent vector

$$\ell(\gamma_{t'}) = \int_0^1 \|\dot{\gamma}_{t'}\|_{\mathfrak{g}} \mathrm{d}t' = \int_0^1 \sqrt{\dot{\gamma}_{t'}^\top H(\gamma_{t'}) \dot{\gamma}_{t'}} \mathrm{d}t', \tag{31}$$

where the dotted letter indicates the velocity of the curve $\dot{\gamma}_{t'} = \partial_{t'} \gamma_{t'}$. A **geodesic curve** is a locally length minimising smooth curve that connects two given points on the manifold.

Riemannian geometry of observations. For approximating the posterior over paths we take into account the geometry of the invariant density as it is represented by the observations. To that end, we consider systems whose dynamics induce invariant (inertial) manifolds that contain the global attractor of the system and on which system trajectories concentrate (Wiggins, 1994; Mohammed and Scheutzow, 1999; Girya and Chueshov, 1995; Fenichel, 1971; Arnold, 1990; Carverhill, 1985). We assume thus that the continuous-time trajectories $\mathbf{X}_{0:T} \in \mathbb{R}^d$ of the underlying system concentrates on an invariant manifold $\mathcal{M} \in \mathbb{R}^{m \leq d}$ of dimensionality m (possibly) smaller than d. The discrete-time observations $\boldsymbol{\mathcal{O}}_k$ are thus samples of the manifold $\mathcal{M}$. The central premise of our approach is that **unobserved paths between successive observations will be lying either *on* or *in the vicinity* of the manifold** $\mathcal{M}$. In particular, we postulate that unobserved paths should lie **in the vicinity of geodesics that connect consecutive observations** on $\mathcal{M}$. To that end we propose a path augmentation framework that constraints the augmented paths to lie in the vicinity of identified geodesics between consecutive observations.

However, while this view of a lower dimensional manifold embedded in a higher dimensional ambient space helps to build our intuition for the proposed method, for computational purposes we adopt a complementary view inspired by the discussion in (Fröhlich et al., 2021). According to this view, we consider the entire observation space $\mathbb{R}^d$ as a smooth Riemannian manifold, $\mathcal{M} \doteq \mathbb{R}^d$, characterised by a Riemannian metric $\mathfrak{h}$. The effect of the nonlinear geometry of the observations is then captured by the metric $\mathfrak{h}$. Thus to approximate the geometric structure of the system's invariant density, we learn the Riemannian metric tensor $H : \mathbb{R}^d \to \mathbb{R}^{d \times d}$ and compute the geodesics between consecutive observations according to the learned metric. Intuitively according to this view the observations $\{\boldsymbol{\mathcal{O}}_k\}_{k=1}^K$ introduce distortions in the way we compute distances on the state space.

In effect this approach does not reduce the dimensionality of the space we operate, but changes the way we compute inner products and thus distances, lengths, and geodesic curves on $\mathcal{M}$. The alternative perspective of working on a lower dimensional manifold would strongly depend on the correct assessment of the dimensionality of said manifold. For example, one could use a Variational Autoencoder to approximate the observation manifold and subsequently obtain the Riemannian metric from the embedding of the manifold mediated by the decoder. However, our preliminary results of such an approach revealed that such a method requires considerable fine tuning to adapt to the characteristics of each dynamical system and is sensitive to the estimation of the dimensionality of the approximated manifold.

To learn the Riemannian metric and compute the geodesics we follow the framework proposed by Arvanitidis et al. in (Arvanitidis et al., 2019). In particular, we approximate the local metric induced by the observations at location $\mathbf{x}$ of the state space, in a non-parametric form by the inverse of the weighted local diagonal covariance computed on the observations as (Arvanitidis et al., 2019)

$$H_{dd}(\mathbf{x}) = \left(\sum_{i=1}^K w_i(\mathbf{x}) \left(x_i^{(d)} - x^{(d)} \right)^2 + \epsilon \right)^{-1}, \tag{32}$$

with weights $w_i(\mathbf{x}) = \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}\|_2^2}{2\sigma_{\mathcal{M}}^2} \right)$, and $x^{(d)}$ denoting the d-th dimensional component of the vector $\mathbf{x}$. The parameter $\epsilon > 0$ ensures non-zero diagonals of the weighted covariance matrix, while $\sigma_{\mathcal{M}}$ characterises the curvature of the manifold.

Between consecutive observations for each interval $[\boldsymbol{\mathcal{O}}_k, \boldsymbol{\mathcal{O}}_{k+1}]$, we identify the geodesic $\boldsymbol{\gamma}_{t'}^k$ as the energy minimising curve, i.e., as the minimiser of the kinetic energy functional $\mathcal{E}(\boldsymbol{\gamma}_{t'}^k) = \int_0^1 L_{\mathcal{M}}(\boldsymbol{\gamma}_{t'}^k, \dot{\boldsymbol{\gamma}}_{t'}^k)\,\mathrm{d}t'$

$$\boldsymbol{\gamma}_{t'}^{k*} = \underset{\boldsymbol{\gamma}_{t'}^k, \boldsymbol{\gamma}_0^k = \boldsymbol{\mathcal{O}}_k, \boldsymbol{\gamma}_1^k = \boldsymbol{\mathcal{O}}_{k+1}}{\arg\min} \int_0^1 L_{\mathcal{M}}(\boldsymbol{\gamma}_{t'}^k, \dot{\boldsymbol{\gamma}}_{t'}^k)\,\mathrm{d}t',$$

$$\text{with} \quad \int_0^1 L_{\mathcal{M}}(\gamma_{t'}^k, \dot{\gamma}_{t'}^k)\mathrm{d}t' = \frac{1}{2}\int_0^1 \|\dot{\gamma}_{t'}^k\|_{\mathfrak{h}}^2, \tag{33}$$

where $L_{\mathcal{M}}(\boldsymbol{\gamma}_{t'}^k, \dot{\boldsymbol{\gamma}}_{t'}^k)$ denotes the Lagrangian. The minimising curve of this functional is the same as the minimiser of the curve length functional $\ell(\boldsymbol{\gamma}_{t'})$ (Eq. 31), i.e., the geodesic (Do Carmo, 1992).

By applying calculus of variations, the minimising curve of the functional $\mathcal{E}(\boldsymbol{\gamma}_{t'}^k)$ can be obtained from the Euler-Lagrange equations, resulting in the following system of second order differential equations (Arvanitidis et al., 2017; Do Carmo, 1992)

$$\ddot{\boldsymbol{\gamma}}_t{}^k = -\frac{1}{2}H(\boldsymbol{\gamma}_t^k)^{-1}\left(2\left(I \otimes (\dot{\boldsymbol{\gamma}}_t{}^k)^\top\right)\frac{\partial \mathrm{vec}[H(\boldsymbol{\gamma}_t^k)]}{\partial \boldsymbol{\gamma}_t^k}\dot{\boldsymbol{\gamma}}_t{}^k - \frac{\partial \mathrm{vec}[H(\boldsymbol{\gamma}_t^k)]^\top}{\partial \boldsymbol{\gamma}_t^k}\left(\dot{\boldsymbol{\gamma}}_t{}^k \otimes \dot{\boldsymbol{\gamma}}_t{}^k\right)\right), \tag{34}$$

with boundary conditions $\boldsymbol{\gamma}_0^k = \boldsymbol{\mathcal{O}}_k$ and $\boldsymbol{\gamma}_1^k = \boldsymbol{\mathcal{O}}_{k+1}$, where $\otimes$ stands for the Kroenecker product, and $\mathrm{vec}[A]$ denotes the vectorisation operation of matrix A through stacking the columns of A into a vector. We follow (Arvanitidis et al., 2019) and obtain the geodesics by approximating the solution of the boundary value problem of Eq. 34 with a probabilistic differential equation solver.

Extended free energy functional. We denote the collection of individual geodesics by $\boldsymbol{\Gamma}_t \doteq \{\boldsymbol{\gamma}_{t'}^k\}_{t=(k-1)\tau+t'\tau}$, where $\boldsymbol{\gamma}_{t'}^k$ is the geodesic connecting $\boldsymbol{\mathcal{O}}_k$ and $\boldsymbol{\mathcal{O}}_{k+1}$, and $t' \in [0,1]$ denotes a rescaled time variable. Additional to the constraints imposed in the previously explained setting (Sec A.3.1), here we add an extra term in the free energy $U_{\mathcal{G}}(\mathbf{x}, t) \doteq \|\boldsymbol{\Gamma}_t - \mathbf{x}\|^2$ that accounts for the local geometry of the invariant density, and guides the latent path towards the geodesic curves $\gamma_{t'}^k$ that connect consecutive observations

$$\mathcal{F}[Q] = \frac{1}{2}\int_0^T\int \left[\|g(\mathbf{x}, t) - \hat{f}(\mathbf{x})\|_D^2 + U_{\mathcal{O}}(\mathbf{x}, t) + \beta U_{\mathcal{G}}(\mathbf{x}, t)\right] q_t(\mathbf{x})\, d\mathbf{x}\, \mathrm{d}t. \tag{35}$$

Here we denote the observation term by $U_{\mathcal{O}}(\mathbf{x}, t) \doteq -\sum_{t_k} \ln \mathrm{P}(\boldsymbol{\mathcal{O}}_k|\mathbf{x})\delta(t - t_k)$, while β stands for a weighting constant that determines the relative weight of the geometric term in the control objective.

Eq. 35 can be construed as a least-action principle for the augmented path measure. The first term is the energetic cost of changing the drift from ($\hat{f}$) to ($g = \hat{f} + u$). The observation potential ($U_{\mathcal{O}}$) imposes the bridge boundary conditions at the sparse measurement times. The geometric potential ($U_{\mathcal{G}}$) does not force the path to equal the geodesic, but penalises deviations from it, with strength controlled by (β). Therefore, the optimal-control formulation provides a principled way to combine three requirements: dynamical consistency with the current drift estimate, interpolation through the observations, and regularisation by the invariant-density geometry.

Following (Opper, 2019), for each inter-observation interval $[\boldsymbol{\mathcal{O}}_k, \boldsymbol{\mathcal{O}}_{k+1}]$ we identify the posterior path measure (minimiser of Eq. 35) by the solution of a stochastic optimal control problem (Maoutsa and Opper, 2022) with the objective to obtain a time-dependent drift adjustment $\mathbf{u}(\mathbf{x}, t) := \mathbf{g}(\mathbf{x}, t) - \hat{\mathbf{f}}(\mathbf{x})$ for the system with drift $\hat{\mathbf{f}}(\mathbf{x})$ with initial and terminal constraints defined by $U_{\mathcal{O}}(\mathbf{x}, t)$, and additional path constraints $U_{\mathcal{G}}(\mathbf{x}, t)$.

For the case of exact observations, i.e., for an observation process $\boldsymbol{\psi}(\mathbf{x}) = \mathbf{x}$, we can compute the drift adjustment for each of the $K - 1$ inter-observation intervals independently. Thus for each interval between consecutive observations, we identify the optimal control $\mathbf{u}(\mathbf{x}, t)$ required to construct a stochastic bridge following the dynamics of Eq. 24 with initial and terminal states the respective observations $\boldsymbol{\mathcal{O}}_k$ and $\boldsymbol{\mathcal{O}}_{k+1}$.

The optimal drift adjustment for such a stochastic control problem for the inter-observation interval between $\boldsymbol{\mathcal{O}}_k$ and $\boldsymbol{\mathcal{O}}_{k+1}$ can be obtained from the solution of the backward equation (see (Maoutsa and Opper, 2022; 2021))

$$\frac{\partial \phi_t(\mathbf{x})}{\partial t} = -\mathcal{L}_{\hat{f}}^\dagger \phi_t(\mathbf{x}) + U_{\mathcal{G}}(\mathbf{x}, t)\phi_t(\mathbf{x}), \tag{36}$$

with terminal condition $\phi_T(\mathbf{x}) = \chi(\mathbf{x}) = \delta(\mathbf{x} - \boldsymbol{\mathcal{O}}_{k+1})$ and with $\mathcal{L}^{\dagger}_{\hat{f}}$ denoting the adjoint Fokker-Planck operator for the process of Eq. 24. As shown in (Maoutsa and Opper, 2022) the optimal drift adjustment $\mathbf{u}(\mathbf{x}, t)$ can be expressed in terms of the difference of the logarithmic gradients of two probability flows

$$\mathbf{u}^*(\mathbf{x}, t) = D\Big(\nabla \ln q_{T-t}(\mathbf{x}) - \nabla \ln \rho_t(\mathbf{x})\Big), \tag{37}$$

where ρ_t fulfils the forward (filtering) partial differential equation (PDE)

$$\frac{\partial \rho_t(\mathbf{x})}{\partial t} = \mathcal{L}_{\hat{f}} \rho_t(\mathbf{x}) - U_{\mathcal{G}}(\mathbf{x}, t)\rho_t(\mathbf{x}), \tag{38}$$

while q_t is the solution of a time-reversed PDE that depends on the logarithmic gradient of $\rho_t(\mathbf{x})$

$$\frac{\partial q_t(\mathbf{x})}{\partial t} = -\nabla \cdot \left[\Big(\sigma^2 \nabla \ln \rho_{T-t}(\mathbf{x}) - \mathbf{f}(\mathbf{x}, T-t)\Big) q_t(\mathbf{x})\right] + \frac{\sigma^2}{2} \nabla^2 q_t(\mathbf{x}), \tag{39}$$

with initial condition $q_0(\mathbf{x}) \propto \rho_T(\mathbf{x})\chi(\mathbf{x})$.

We obtain the optimal control by solving the backward equation associated with this controlled diffusion problem. The solution defines the change of path measure required to condition the prior process on the endpoint and path potentials. The resulting control can be expressed as a logarithmic gradient of the corresponding desirability or path-density functions, resulting in a time-dependent correction ($u^*(x, t)$). In the implementation, we apply this construction independently to each interval ($[\mathcal{O}_k, \mathcal{O}_{k+1}]$), obtaining one controlled bridge per pair of consecutive observations.

For the numerical solution of the control problem we use the numerical framework accompanying (Maoutsa and Opper, 2022), where the path constraints associated with the geodesic curves are imposed through the two staged process for particle propagation described in the paper for path constraints, with the particle reweighting being performed through optimal transport implemented using the PyEMD python toolbox (Pele and Werman, 2009).

More precisely, according to this framework we propagate a particle representation of the probability density $\rho_t(\mathbf{x})$ according to the filtering equation of Eq. 38. This follows the dynamics of the uncontrolled process with drift $\hat{\mathbf{f}}$ and particle reweighting at each time step as determined by the path constrained (potential) $U_{\mathcal{G}}(\mathbf{x}, t)$, that quantifies the proximity to the geodesic at each time point. In the particle representation we apply this reweighting in the form of a deterministic optimal transportation of the particles (Reich, 2013).

A.4. Approximate posterior over drift functions.

For a fixed path measure Q, the optimal measure for the drift Q_f is a Gaussian process given by

$$Q_f \propto \mathrm{P}_f \exp\left(-\frac{1}{2}\int \|\mathbf{f}(\mathbf{x})\|_D^2 A(\mathbf{x}) - 2\langle \mathbf{f}(\mathbf{x}), B(\mathbf{x})\rangle_D \mathrm{d}\mathbf{x}\right), \tag{40}$$

with

$$A(\mathbf{x}) \doteq \int_0^T q_t(\mathbf{x}) \mathrm{d}t,$$

and

$$B(\mathbf{x}) \doteq \int_0^T q_t(\mathbf{x}) g(\mathbf{x}, t) \mathrm{d}t,$$

where $q_t(\mathbf{x})$ denotes the marginal constrained density of the state $\mathbf{X}_t$. The function $\mathbf{g}(\mathbf{x}, t)$ denotes the effective drift.

We assume a Gaussian process prior for the unknown function $\mathbf{f}$, i.e., $\mathbf{f} \sim \mathrm{P}_0(\mathbf{f}) = \mathcal{GP}(\boldsymbol{m}^f, k^f)$ where m^f and k^f denote the mean and covariance function of the Gaussian process. Following Ruttor *et al.* (Ruttor et al., 2013), we employ a sparse kernel approximation for the drift f by optimising the function values over a sparse set of S inducing points $\{Z_i\}_{i=1}^S$. We obtain the resulting drift from

$$\hat{\mathbf{f}}_S(\mathbf{x}) = k^{\mathbf{f}}(\mathbf{x}, \mathcal{Z}) \left(I + \Lambda \mathcal{K}_S\right)^{-1} \mathbf{d}, \tag{41}$$

where we have defined introduced the notation $\mathcal{K}_S \doteq k^f(\mathcal{Z}, \mathcal{Z})$

$$\Lambda = \frac{1}{\sigma^2}\mathcal{K}_S^{-1}\left(\int k^f(\mathcal{Z}, \mathbf{x})A(\mathbf{x})k^f(\mathbf{x}, \mathcal{Z})\mathrm{d}\mathbf{x}\right)\mathcal{K}_S^{-1}. \tag{42}$$

$$\mathbf{d} = \frac{1}{\sigma^2}\mathcal{K}_S^{-1}\left(\int k^f(\mathcal{Z}, \mathbf{x})B(\mathbf{x})\mathrm{d}\mathbf{x}\right)\mathcal{K}_S^{-1}, \tag{43}$$

The associated variance results similarly from the equation

$$\Sigma_S^2(\mathbf{x}) = k^f(\mathbf{x}, \mathbf{x}) - k^f(\mathbf{x}, \mathcal{Z})\left(I + \Lambda\,\mathcal{K}_S\right)^{-1}\Lambda\, k^f(\mathcal{Z}, \mathbf{x}). \tag{44}$$

We employ a sample based approximation of the densities in Eq. 40 resulting from the particle sampling of the path measure Q resulting from the geometric augmentation, i.e. the integrals over $\int q_t(\mathbf{x})$ are over the samples of the augmented paths. Thus by representing the densities by samples, we can rewrite the density $p_t(x)$ in terms of a sum of Dirac delta functions centered around the particles positions

$$p_t(\mathbf{x}) \approx \frac{1}{N}\sum_{j=1}^{N}\delta(\mathbf{x} - \mathbf{X}_j(t)),$$

and replace the Riemannian integrals with summation over particles, i.e. perform a Monte Carlo integration. Here $\mathbf{X}_j(t)$ represents the position of the j-th particle at time point t.

B. Sparse Gaussian process estimation

Since the amount of required observations for accurate drift estimation is generally large for systems with nonlinear dynamics, regular Gaussian process regression becomes computationally intensive. Its computational complexity scales as $\mathcal{O}(N^3)$ with the number of observations N due to the $N \times N$ kernel matrix inversions required for inference (c.f. Eq.20 and (Rasmussen, 2003)). Therefore, (Ruttor et al., 2013) employ the sparse (low dimensional approximation) counterpart of Gaussian process regression (Titsias, 2009; Csató and Opper, 2002) that reduces significantly the computation time by reducing the computational complexity to $\mathcal{O}(NM^2)$, where $M \ll N$ denotes the number of selected sparse (inducing) points. Here we present briefly the derivation.

For sparse Gaussian process drift inference, we augment the distributions with M inducing points $\mathbf{z} = [z_1, \ldots, z_M]$ with inducing values $\mathbf{u} = [\mathbf{f}(z_m)]_{m=1}^M$ that are jointly Gaussian distributed with the latent function values $\{\mathbf{f}(\mathbf{X}_t)\}_{t=0}^T$.

As demonstrated previously the true posterior for function values $\mathbf{f}$ is expressed as a product

$$P_f(\mathbf{f}) = \frac{1}{Z}P_o(\mathbf{f})e^{-\mathcal{A}(\mathbf{f})}, \tag{45}$$

where Z a normalisation constant, $\mathcal{A}(\mathbf{f}) = \frac{1}{2}\mathbf{f}^T\Lambda\mathbf{f} - \mathbf{a}^T\mathbf{f}$ a quadratic form of $\mathbf{f}$ (see Eq. 14), while $P_o(\mathbf{f})$ denotes a prior Gaussian measure. Thus the posterior $P_f(\mathbf{f})$ is also Gaussian.

To employ sparse Gaussian process inference, we approximate P_f with $Q_f = \mathcal{GP}\left(m^q(\cdot), k^q(\cdot,\cdot)\right)$, with mean and variance functions to be calculated, depending only on the *smaller* subset ($M \ll N$) of inducing function values $\mathbf{u}$,

$$Q_f(\mathbf{f}) \propto R(\mathbf{u})P_o(\mathbf{f}). \tag{46}$$

The effective likelihood $R(\mathbf{u})$ is chosen as the minimiser of the Kullback-Leibler divergence $\mathcal{KL}\left(Q_f||P_f\right)$.

We may now express the prior $P_o(\mathbf{f})$ and the approximate marginal $Q_f(\mathbf{f})$ in terms of the inducing points

$$P_o(\mathbf{f}) = P_o(\mathbf{f}|\mathbf{u})P_o(\mathbf{u}), \tag{47}$$

and

$$Q_f(\mathbf{f}) = Q_f(\mathbf{f}|\mathbf{u})Q_f(\mathbf{u}) = P_o(\mathbf{f}|\mathbf{u})Q_f(\mathbf{u}), \tag{48}$$

under the assumption that the posterior conditional $Q_f(\mathbf{f}|\mathbf{u})$ matches the prior conditional $P_o(\mathbf{f}|\mathbf{u})$.

We select the effective likelihood $R(u)$ as the minimiser of the relative entropy between Q_f and P_f

$$\begin{aligned}
\mathcal{KL}\left(Q_f||P_f\right) &= \int Q_f(\mathbf{f})\,\ln\frac{Q_f(\mathbf{f})}{P_f(\mathbf{f})}\mathrm{d}\mathbf{f} \\
&= \int P_o(\mathbf{f}|\mathbf{u})Q_f(\mathbf{u})\,\ln\frac{P_o(\mathbf{f})R(\mathbf{u})}{\frac{1}{Z}P_o(\mathbf{f})e^{-\mathcal{A}(\mathbf{f})}}\mathrm{d}\mathbf{f}\mathrm{d}\mathbf{u} \\
&= \int P_o(\mathbf{f}|\mathbf{u})Q_f(\mathbf{u})\,\ln\frac{P_o(\mathbf{f})R(\mathbf{u})}{\frac{1}{Z}P_o(\mathbf{f}|\mathbf{u})e^{-\mathcal{A}(\mathbf{f}|\mathbf{u})}P_o(\mathbf{u})}\mathrm{d}\mathbf{f}\mathrm{d}\mathbf{u} \\
&= \int P_o(\mathbf{f}|\mathbf{u})Q_f(\mathbf{u})\,\ln\frac{P_o(\mathbf{u})R(\mathbf{u})}{\frac{1}{Z}e^{-\mathcal{A}(\mathbf{f}|\mathbf{u})}P_o(\mathbf{u})}\mathrm{d}\mathbf{f}\mathrm{d}\mathbf{u} \\
&= \int P_o(\mathbf{f}|\mathbf{u})Q_f(\mathbf{u})\,\ln\frac{R(\mathbf{u})}{\frac{1}{Z}e^{-\mathcal{A}(\mathbf{f}|\mathbf{u})}}\mathrm{d}\mathbf{f}\mathrm{d}\mathbf{u} \\
&= \ln Z + \int Q_f(\mathbf{u})\,\ln\left(\frac{e^{\ln R(\mathbf{u})}}{e^{-\mathbb{E}_o[\mathcal{A}(\mathbf{f}|\mathbf{u})]}}\right)\mathrm{d}\mathbf{u}.
\end{aligned} \tag{49}$$

In Eq. 49 in the second line, we have introduced Eq. 45-Eq. 48. In the third line we have introduced $\frac{P_0(\mathbf{f})}{P_0(\mathbf{f}|\mathbf{u})} = P_0(\mathbf{u})$ from Eq. 47. In the final line we rearranged the terms that do not depend on $\mathbf{f}$ outside of the integral over $\mathbf{f}$, moved the $\ln Z$ term out of the integration over $\mathbf{u}$, and denoted $\langle\cdot\rangle_0 = \int P_0(\mathbf{f}|\mathbf{u})\,\mathrm{d}\mathbf{f}$.

To minimise the relative entropy $\mathcal{KL}\left[Q_f||P_f\right]$ we conclude that the optimal choice for the effective likelihood $R(\mathbf{u})$ is

$$R(\mathbf{u}) \propto e^{-\langle\mathcal{A}(\mathbf{f}|\mathbf{u})\rangle_o}. \tag{50}$$

Given the quadratic form of $A(\mathbf{f})$ we may write the conditional expectation in Eq. 50 as a quadratic form too

$$\begin{aligned}
\langle\mathcal{A}(\mathbf{f}|\mathbf{u})\rangle_o &= \frac{1}{2}\langle\mathbf{f}|\mathbf{u}\rangle_o^\top\,\Lambda\,\langle\mathbf{f}|\mathbf{u}\rangle_o + \frac{1}{2}\mathrm{Tr}\left(\mathrm{Cov}_o[\mathbf{f}|\mathbf{u}]\Lambda\right) - a^\top\langle\mathbf{f}|\mathbf{u}\rangle_o \\
&= \frac{1}{2}\langle\mathbf{f}|\mathbf{u}\rangle_o^\top\,\Lambda\,\langle\mathbf{f}|\mathbf{u}\rangle_o - a^\top\langle\mathbf{f}|\mathbf{u}\rangle_o + \text{const.},
\end{aligned} \tag{51}$$

where in the last line we take into account that the term $\mathrm{Tr}\left(\mathrm{Cov}_o[\mathbf{f}|\mathbf{u}]\Lambda\right)$ is independent of the sparse function values $\mathbf{u}$ (c.f. (Ruttor et al., 2013)). In Eq. 51 $\Lambda \doteq \mathrm{diag}\big[\Delta t\,D^{-1},\ldots,\Delta t\,D^{-1}\big]$.

In particular, the conditional expectation of function values f conditioned on the inducing point function values $\mathbf{u}\equiv\mathcal{U}$ at inducing point locations $\mathbf{z}\equiv\mathcal{Z}$ equals

$$\bar{f}^s(\mathbf{x}) = \langle f|\mathbf{u}\rangle_o = k(\mathbf{x},\mathcal{Z})k(\mathcal{Z},\mathcal{Z})^{-1}\mathcal{U}, \tag{52}$$

while the covariance equals

$$(\Sigma^s)^2(\mathbf{x}) = k(\mathbf{x},\mathbf{x}) - k(\mathbf{x},\mathcal{Z})k(\mathcal{Z},\mathcal{Z})^{-1}k(\mathcal{Z},\mathbf{x}), \tag{53}$$

where we have employed similar notation for the kernel functions as in Eqs. 19-20.

C. Theoretical evidence that supports the use of geodesics as geometric constraints representing most probable paths between observations

The Onsager-Machlup functional for diffusion processes has been known in theoretical physics as a characteriser of the most probable path (MPP) between two pre-defined states of the process. In (Onsager and Machlup, 1953), Onsager and Machlup used the thermal fluctuations of a diffusion process to show that the probability density of a path $\boldsymbol{\gamma}\in C^1\left([0,T],\mathbb{R}^d\right)$ in $\mathbb{R}^d$ over finite interval can be expressed as a Boltzmann factor

$$\mathrm{P}(\boldsymbol{\gamma}) \sim \exp\left[-\int_0^T L(\boldsymbol{\gamma}(t),\dot{\boldsymbol{\gamma}}(t))dt\right], \tag{54}$$

where

$$L(\boldsymbol{\gamma}(t), \dot{\boldsymbol{\gamma}}(t)) = \frac{1}{2}\|\dot{\boldsymbol{\gamma}}(t) - \mathbf{f}(\boldsymbol{\gamma}(t))\|^2_{\mathbf{D}^{-1}} + \frac{1}{2}\nabla \cdot \mathbf{f}(\boldsymbol{\gamma}(t)).^{2} \tag{55}$$

The function $L(\boldsymbol{\gamma}(t), \dot{\boldsymbol{\gamma}}(t))$ is known as the **Onsager-Machlup** function (action), while its integral over time is known as Onsager-Machlup action functional. It has been used as Lagrangian in Euler-Lagrange minimisation schemes to identify the **most probable path (MPP)** of a diffusion process between two given points in the state space (Graham, 1977; Stratonovich, 1971).

Stratonovich (Stratonovich, 1971) considered the probability that a sample of a multidimensional diffusion process will lie in the vicinity of (within a tube of infinitesimal thickness around) an idealised smooth path in the state space. To compute this probability he constructed a probability functional which is identical to the Onsager-Machlup functional considered as Lagrangian for the diffusion process. Duerr et al. (Dürr and Bach, 1978) considered scalar diffusion processes and constructed the Onsager-Machlup function from the asymptotic limit of the transition probability between the starting and end state of the path using a Girsanov transformation.

Considering Brownian motions defined on a Riemannian manifold $(\mathcal{M}, \mathfrak{g})$ with associated Riemannian metric $\mathfrak{g}$, the Onsager-Machlup functional can be expressed as the integral over the Lagrangian (Takahashi and Watanabe, 1981; Graham, 1980; Grong and Sommer, 2022)

$$L(\boldsymbol{\gamma}, \dot{\boldsymbol{\gamma}}) = \frac{1}{2}\|\dot{\boldsymbol{\gamma}}(t)\|^2_{\mathfrak{g}} - \frac{1}{12}S(\boldsymbol{\gamma}(t)), \tag{56}$$

where $\|\cdot\|_{\mathfrak{g}}$ denotes the Riemannian norm on the tangent space $\mathcal{T}_X\mathcal{M}$ of the manifold with respect to the metric $\mathfrak{g}$, and $S(\cdot)$ stands for the scalar curvature of the manifold at each point. The first term is the Lagrangian used to identify geodesic curves on manifolds (c.f. A.3.2)

In our proposed formalism, for computational purposes we have assumed the entire $\mathbb{R}^d$ as smooth manifold. We can identify the first term of Eq. 56 with the Lagrangian we optimised for computing the geodesics on the manifold $(\mathbb{R}^d, \mathfrak{g})$, where $\mathfrak{g}$ is the metric learned from the observations.

However the system we observed was a diffusion process defined in $\mathbb{R}^d$ with an Euclidean metric. Constructing a path augmentation scheme that guides the augmented paths towards the geodesics of a diffusion defined with respect to a different metric raises questions about the validity of our approach. Here we should note that diffusions with a general state dependent diffusion coefficient $\boldsymbol{\sigma} \in \mathbb{R}^{d\times m}$, and m-dimensional Brownian motion, can be considered as evolving on the manifold $\left(\mathbb{R}^d, \mathfrak{g}\right)$, with the associated metric $\mathfrak{g} = \left(\boldsymbol{\sigma}\boldsymbol{\sigma}^\top\right)^{-1}$ (Capitaine, 2000). Thus it may be possible to associate the metric learned from the data with the metric arising from a state dependent diffusion by applying a transformation akin to an inverse Lamperti transform (Øksendal, 2003) to transform our learned SDE to one that would have induced the learned metric due to the state dependent diffusion. The existence of such a transformation would justify the proposed method. Our empirical results demonstrate that such a transformation may be possible.

D. Does the proposed approach invalidate the Markovian property of the diffusion process?

The proposed path augmentation seemingly invalidates the Markovian property of the diffusion process. According to the Markov property of the diffusion of Eq. 1, the system state $\mathbf{X}_{k\tau+\delta t}$ should depend only the state $\mathbf{X}_{k\tau}$, i.e., the observation $\boldsymbol{\mathcal{O}}_k$. The proposed augmentation makes the state $\mathbf{X}_{k\tau+\delta t}$ depending not only on the next observation $\boldsymbol{\mathcal{O}}_{k+1} = \mathbf{X}_{(k+1)\tau}$, but also on past and future states that lie in the vicinity of these observations.

We effectively construct the augmented paths to compute the likelihood of a drift estimate. To compute this likelihood we require to evaluate the transition probabilities between consecutive observations. Since for general nonlinear systems the transition probabilities are in general intractable, we have to resort to numerical approximations. Ideally we would approximate the transition density with a bridge sampler that would consider the nonlinear estimated SDE conditioned to pass though consecutive observations. However for coarse drift estimates, the observations have zero probability under the law of the estimated SDE, and construction of those bridges would result either in very taxing computations or would fail altogether. Instead, here, we compute the likelihood of a "corrected" estimate (the correction resulting from the invariant density) under which the observations have non-zero probability, and subsequently re-estimate the drift on the

[2] Onsager and Machlup's initial work concentrated around linear processes and therefore the functional initially introduced by the did not include the second term with the divergence of $\mathbf{f}$ as this is a constant for linear $\mathbf{f}$. It was later added to the OM function to account for trajectory entropy corrections (Taniguchi and Cohen, 2007; Adib, 2008)

augmented path with this "corrected" estimate. By taking into account the local geometry of the observations, we provide systematic corrections for the misestimated drift function to generate the augmented paths. This effectively nudges the augmentation process towards the second observation of each inter-observation interval through the path constraint that forces the augmented paths towards the geodesics.

E. Related work and positioning of the present work

Here, we briefly review further related work on inference or modelling of SDEs and position our work further with respect to the existing literature.

Recent advances in dynamical system inference have delivered valuable tools for identifying continuous-time *deterministic* systems from observations (Cremers and Hübler, 1987; Brunton et al., 2016; Daniels and Nemenman, 2015; Casadiego et al., 2018; McGoff et al., 2015; Kantz and Schreiber, 2004; Schmidt and Lipson, 2009; Gerardos and Ronceray, 2025; Packard et al., 1980; Crutchfield and McNamara, 1987). **Data-driven** (or **nonparametric**, or **equation-free**) approaches seek to reconstruct the governing equations of observed systems directly from state observations, without imposing explicit assumptions or inductive biases about the underlying dynamical models. They rely on function approximation to infer the system's structure from observations, such as basis functions (Acosta, 1995; Small and Tse, 2002; Judd and Mees, 1995; Small and Judd, 1998; Brückner et al., 2020; Frishman and Ronceray, 2020), symbolic regression (Kaiser et al., 2018; Brunton et al., 2016; Bongard and Lipson, 2007; Daniels and Nemenman, 2015), spectral approximations (Kevrekidis et al., 2003; Theodoropoulos et al., 2000), Gaussian processes (Álvarez et al., 2009; Sanguinetti et al., 2006; Särkkä, 2019), or neural networks (Teng, 2018; Bhattoo et al., 2022; Jüngling et al., 2019). However, extending these methods to *stochastic* systems remains difficult. In this setting, inference must disentangle the influence of underlying deterministic forces from random fluctuations, a task that is particularly difficult for low sampling rates.

▷ **Modelling general SDEs from state observations.** As already mentioned in the Introduction and in Sec. A existing inference methods for SDEs can be broadly clustered in temporal and geometric methods, where the former accounts for the temporal order of the observations, while the latter approximate the invariant system density and discard any time information.

Temporal methods rely on the Euler-Maruyama discretisation of the SDE paths approximating conditional expectations of state increments (i.e. the Krammers Moyal coefficients). They model the drift either in terms of Gaussian processes (Ruttor et al., 2013; Batz et al., 2018; Hostettler et al., 2018; Zhao et al., 2020; Yildiz et al., 2018), basis functions (Nabeel et al., 2025; Ragwitz and Kantz, 2001; Friedrich and Peinke, 1997; Peinke et al., 1997; Gao et al., 2024; Friedrich et al., 2000; Ferretti et al., 2020) or libraries of functions (Boninsegna et al., 2018; Huang et al., 2022), kernel regression (Lamouroux and Lehnertz, 2009; Jiang and Knight, 1997), dynamic mode decomposition to learn the eigenfunctions of the Koopman operator (Klus et al., 2020), by approximating the central moments of the transition densities (Stanton, 1997), Bayesian neural networks (Bae et al., 2025), or by applying generalised methods of moments (Hansen and Scheinkman, 1993).

As explicitly detailed in Sec. A, most temporal methods do not provide accurate drift estimates when the interval between observations is large. The two prevailing approaches to mitigate this finite-sampling rate effects is to either account for the systematic bias introduced by the finite sampling rate by estimating an explicit correction term for the inferred drift (Ragwitz and Kantz, 2001; 2002; Kleinhans et al., 2005; Kleinhans and Friedrich, 2007), or by performing state estimation for the unobserved paths (also known as path or data augmentation) and then estimating the drift from the continuous paths.

The former approach works only for scalar systems, while the latter employs simplified bridge dynamics (e.g., Brownian (Chib et al., 2006; Eraker, 2001; Sermaidis et al., 2013) or Ornstein Uhlenbeck (Batz et al., 2018; Billio et al., 1998) bridges) that are analytically tractable or computationally non-demanding. However, for large τ and for nonlinear systems, these simplified bridge dynamics match poorly the underlying path statistics. (Fig. 1 **D.**). It is important to mention here, that path augmentation with Ornstein Uhlenbeck bridges similar to (Batz et al., 2018) provides a good approximation of the underlying transition density, when the underlying linear process employed for each bridge has a drift that comes from the local linearisation of the **ground truth** drift function. However, during inference the ground truth dynamics are unknown and the local linearisations on inaccurate drift estimates employed in (Batz et al., 2018) provide imprecise approximations for large τ.

Alternative methods, employ variational inference (Batz et al., 2016; Opper, 2019; Duncker et al., 2019; Verma et al., 2024) and approximate the posterior path measure with a tractable Gaussian process induced by a time–varying linear SDE. This

results in ODEs for the posterior mean and covariance matrix and an ELBO that is optimized directly (Archambeau et al., 2007; Duncker et al., 2019).

Building on the prolific line of work on neural ODEs, neural SDEs (Li et al., 2020) employ gradient-based stochastic variational inference and the stochastic adjoint sensitivity method to compute gradients of solutions of stochastic equations with respect to their parameters. Building on these methods, (Course and Nair, 2023b) remove the need for adjoint-based gradient computations by combining amortized inference with a reparametrization of the ELBO by assuming a latent linear process that generates the latent path.

Geometric approaches on the other hand, discard the temporal structure of the observations, and treat them as samples of the invariant density. Thereby these methods either employ density estimation to identify the drift as the gradient of a potential (Kutoyants, 2004), or resort to spectral approximations of the generator of the diffusion process through manifold learning.

Manifold learning methods employ often the *diffusion maps* algorithm, introduced by Coiffman and colleagues (Singer and Coifman, 2008), to learn the dominant part of the spectrum of the transfer operator of the observed diffusion process (Coifman et al., 2005; Nadler et al., 2006; Giannakis, 2019; Ferguson et al., 2011; Talmon and Coifman, 2015). In essence, these methods, learn from the data the few leading eigenfunctions of the Laplace–Beltrami operator that captures the Riemannian geometry of the observations, and consider them as a parametrisation of the manifold representing the invariant density.

▷ **Modelling SDEs from population level snapshots/boundary conditions.** With recent computational advances in optimal transport, a growing body of work focuses on the implementation of Schrödinger bridge sampling methods, including formulations with additional path constraints. These mostly generative methods aim to transport the data distribution from some initial boundary condition to a terminal one, typically by learning the underlying stochastic equation to perform this transport through Schrödinger bridge sampling (Lipman et al., 2023; Pooladian et al., 2023; Albergo et al., 2025; Albergo and Vanden-Eijnden, 2023; Zhang and Stumpf, 2026). Flow matching (Lipman et al., 2023) identifies the probability flow ODE that pushes forward an initial Gaussian density to a target one by solving a regression problem. The method relies on analytically tractable probability paths that provide closed-form regression targets for learning the velocity field, resulting in simulation-free training of deterministic flows. However, the framework is restricted to Gaussian distributions since the employed objective becomes intractable for general source distributions. Conditional flow matching (CFM) (Tong et al., 2023) generalizes flow matching by introducing conditional probability paths between paired samples, allowing the marginal velocity field to be learned with regression without requiring explicit evaluation of the marginal densities or restrictive assumptions on the source distribution. Generalized Schrödinger Bridge Matching (GSBM) (Liu et al., 2024) follows an alternating optimisation scheme that learns both drift and marginals. Given prescribed boundary conditions for initial and terminal densities, the framework minimises a kinetic energy term, and formulates the resulting problem in terms of a stochastic optimal control problem conditioned on the boundary conditions and a path cost that accounts for additional constraints. In the multi-marginal setting, **?** consider the problem of inferring unobserved population trajectories from marginal samples observed at several time points. Their method alternates between learning piecewise Schrödinger bridges between adjacent population snapshots and refining a reference SDE within a selected parametric family using trajectories sampled from the current bridge estimate. This addresses the limitation of standard multi-marginal Schrödinger bridges, that although they enforce multiple marginal constraints, they can reduce to independent adjacent-time bridge problems and therefore fail to identify the long-range temporal structure of the underlying dynamics. The proposed iterative approach introduces a global dynamical bias across time, but the method remains a snapshot-based bridge construction whose output is a path measure compatible with prescribed marginals. Action matching (Neklyudov et al., 2023b) introduces a simulation-free variational objective that identifies a time-dependent scalar potential (entropic action) s_t, whose gradient ∇s_t transports the densities from the initial to the boundary condition through the continuity equation. In its entropic formulation the ∇s_t can be considered as the drift of the underlying SDE, whose marginals match the boundary conditions. However, by construction, the framework can recover only gradient drifts and is therefore not suitable for identifying general stochastic systems with stationary probability currents. In contrast, simulation-free score and flow matching ([SF]^2M) (Tong et al., 2024) jointly learns the probability-flow ODE and the score function by regressing against closed-form quantities derived from conditional Brownian bridge paths, facilitating simulation-free identification of general Schrödinger bridge dynamics with non-gradient drifts. Zhang and Stumpf (2026) extend this simulation-free Schrödinger bridge perspective to non-equilibrium reference dynamics, by replacing the Brownian reference process with a multivariate Ornstein–Uhlenbeck process. This provides the added flexibility of employing reference dynamics of non gradient-type

with steady state probability currents, extending thereby the simulation-free Schrödinger bridge matching framework to non-conservative, non-equilibrium stochastic systems.

Geometry aware generative methods. Metric flow matching (MFM) generalizes CFM by learning interpolants that account for the geometry of the data. However, MFM does not assume a stochastic underlying process, as our framework does, only a deterministic interpolation (transport) that respects the data manifold. However, by assuming a specific noise amplitude for the underlying SDE, one can consider the flow field as generated by the effective drift of a probability flow ODE associated with the considered SDE and make inferences about the underlying drift function. This is the approach we followed when comparing the performance of MFM to our framework in Table 1.

Approximating observation geometry in the ambient space. In our work, we approximate the geometry induced by the observations by endowing the ambient space $\mathbb{R}^d$ with an observation-dependent Riemannian metric $H(\mathbf{x})$ (Eq. 3) that encodes the local anisotropy of the data distribution. In our framework this metric acts as a constraint for data-augmentation and as a geometric inductive bias for drift function inference: augmented paths are encouraged to remain in regions where the metric $H(\mathbf{x})$ induces smaller distances, i.e. in the vicinity of geodesics computed with respect to this metric, thereby aligning the augmented paths with the empirical observation geometry.

This perspective connects to a growing body of work that **approximates Riemannian metrics directly in the ambient space** as a proxy for the unknown curved low-dimensional data manifold, instead of first estimating its intrinsic dimensionality and then constructing explicit low-dimensional embeddings.

In parallel, an increasing body of literature focuses on endowing generative models with geometric constraints or inductive biases. While most methods function in an autoencoder-like setting, by learning an embedding function for projecting to a lower dimensional space that respects prescribed or learned geometric constraints (Duque et al., 2022; Kalatzis et al., 2021; Arvanitidis et al., 2017) geometry, "Riemannian" methods, similar to our proposed method, operate in the ambient space by directly a Riemannian geometry embedded there and define normalizing flows or other generative processes directly on the manifold of interest. (Mathieu and Nickel, 2020) introduce a framework for continuous normalizing flows defined in the ambient space, respecting a prescribed Riemannian geometry. Similarly, (De Bortoli et al., 2022) proposed a score-based generative model that models target densities with support on prescribed Riemannian manifolds in terms of a time-reversal of Langevin dynamics.

Metric flow matching (Kapusniak et al., 2024) interpolates data distributions that respect the geodesic interpolants computed according to the metric induced by the observations. The method employs a data-adapted metric in the ambient space to design interpolants (geodesic curves) with low kinetic energy under the approximated geometry, and constrains the generative paths to respect manifold induced by the data samples. Our construction is conceptually similar with these approaches, in that we also avoid explicit low-dimensional embeddings and instead approximate the observation manifold through a Riemannian metric living in the ambient space. However, in contrast to methods focused on deterministic transport or simulation-free matching, we use the learned metric to regularise continuous-time diffusion bridges and drift inference, through the stochastic controlled geometric augmentation, so that the recovered stochastic dynamics are geometrically consistent with the geometry of the observation-induced invariant measure.

Positioning of the present work. Our approach combines the nonparametric flexibility of Gaussian-process–based drift inference from time-series data with recent geometric ideas for population-level SDE modelling. Similar to Metric Flow Matching (Kapusniak et al., 2024), we posit that augmented trajectories should remain on the manifold induced by the observations: both frameworks estimate a data-adapted Riemannian metric and construct interpolants (geodesics and bridges) that respect this geometry. MFM learns the underlying ODE necessary to transport an initial distribution to a target one under the data-adapted metric, while our framework assumes underlying stochastic dynamics. Nevertheless, once the diffusion is known or coarsely estimated, one can interpret the inferred ODE as a probability flow ODE and make inferences about the underlying drift function of a stochastic system. The GSBM framework (Liu et al., 2024) employs a stochastic control objective that is similar to the objective we consider for constructing the augmented paths. In its standard formulation, GSBM does not impose the geometric constraints on augmented paths that we propose in our framework. For the baseline comparison presented here, we therefore employed a modified variant of GSBM in which we consider the same geometric path constraint term used in our method to the GSBM loss, while keeping all other components of GSBM unchanged. Finally, whereas these methods typically learn a drift that transports a single source distribution to a single terminal snapshot, yielding thus a **locally valid dynamics**, our method, akin to multi-marginal bridge sampling (Shen et al.,

2025), fits a sequence of bridges across multiple time points to recover a **single global drift** consistent with the underlying drift dynamics.

F. Geometric constraints on inference

Our method bridges the gap between approaches that rely only on the temporal structure of observations and those that approximate the invariant density, while ignoring temporal order. Motivated by advances in geometric statistics (Miolane et al., 2020; Sommer, 2020), and the growing interest on the concept of manifold hypothesis (Fefferman et al., 2016; Shnitzer et al., 2020), i.e., the consideration that the state of multi-dimensional dynamical systems often resides in low-dimensional regions of the state space, several recent methods integrate geometric and temporal constraints in stochastic system identification. In *Langevin regression* framework (Callaham et al., 2021), the Kramers-Moyal (KM) coefficients are estimated and low sampling effects are accounted for by solving an adjoint Fokker-Planck equation, with regularisation via moment matching (Lade, 2009).(Tong et al., 2020) consider the manifold of the observations for inference of cellular dynamics. Their method employs dynamic optimal transport to interpolate between measured distributions constrained to lie in the vicinity of the observations. While sharing similar intuitions with our method, Tong et al. do not employ SDE modelling for inherently stochastic cellular dynamics and do not consider the underlying geometry of the observations, relying solely on constraints penalizing pairwise distances between them. Shnitzer et al. (Shnitzer et al., 2020; 2016) employ diffusion maps to approximate the eigenfunctions of the backward Kolmogorov operator (the generator of the stochastic Koopman operator (Giannakis, 2019; Črnjarić-Žic et al., 2020)). By evolving the dominant operator eigenspectrum with a Kalman filter, they account for the temporal order of observations. However, their approach is limited to conservative systems and requires the presence of a spectral gap in the approximated operator's spectrum.

G. Theoretical justification for Riemannian manifold approximation of the invariant density

Our method is based on the argument that the invariant density[3] of the observed system imposes a low-dimensional structure on the state space, within which the observations are confined. We propose that this low-dimensional structure is well approximated by a Riemannian manifold $\mathcal{M}_\infty \in \mathbb{R}^{m \leq d}$ and that the observations $\{\boldsymbol{\mathcal{O}}_k\}_{k=1}^K$ offer a reliable discrete approximation to $\mathcal{M}_\infty$.

We employ the notion of a "low-dimensional structure" as a concise way to refer to the fact that for many dissipative dynamical systems, the invariant measure has support on a subset of the state space with dimension smaller than the ambient space dimension. This phenomenon arises due to the dissipative nature of these systems, which causes volume contraction in the state space, resulting in trajectories concentrating asymptotically on attractors of lower dimension than the state space dimension. To provide further justification on this, in the following section, we start by building intuition from deterministic dynamical systems and then generalise to stochastic dynamics.

G.1. Dimensionality of invariant measures induced by deterministic dynamics

We consider a dissipative deterministic dynamical system of the form

$$\dot{\mathbf{x}}_t = \mathbf{f}(\mathbf{x}_t), \qquad \mathbf{x}_t \in \mathbb{R}^d, \tag{57}$$

generating a semiflow $(\Phi^t)_{t\geq 0}$. Under standard assumptions, the dynamics admit an invariant probability measure μ describing the distribution of states along long-term typical trajectories. From an ergodic perspective, μ is the natural object characterising the asymptotic behaviour of the system. For almost every initial condition in μ, the empirical measure

$$\frac{1}{T}\int_0^T \delta_{\mathbf{x}_t}\, \mathrm{d}t \tag{58}$$

converges (in the weak sense) to μ.

For dissipative systems, phase–space volumes contract along typical trajectories, so the Lebesgue measure is not invariant under the dynamics, i.e. state space volume is not preserved when pushed forward through the flow (Ruelle, 1979). This

[3] In the following the discussion concentrates around invariant measures. We point out here that the invariant density is the Radon-Nikodym derivative of the invariant measure with respect to some reference measure, often the Lebesgue measure if it exists (Maharam, 1969).

implies that the system state does not explore the ambient space uniformly. Instead, trajectories concentrate asymptotically on subsets of state space of vanishing Lebesgue measure. In fact, this concentration phenomenon persists even in chaotic systems, where, although trajectories separate exponentially along unstable directions, contraction along stable directions dominates the evolution of infinitesimal volumes in the state space.

The resulting invariant measure μ typically has an **effective dimension** smaller than the ambient space dimension. To quantify this, we require a notion of dimensionality that remains meaningful when the Lebesgue measure vanishes. The Hausdorff dimension (Ruelle, 1989; Young, 2002; Ott, 2002) lends itself for such a purpose since it naturally extends from sets to probability measures (Young, 1982). More precisely, the Hausdorff dimension of an invariant measure μ is defined as the smallest Hausdorff dimension among all measurable sets containing μ

$$\dim_H(\mu) = \inf\big\{\dim_H(\mathcal{A}) \ : \ \mu(\mathcal{A}) = 1\big\}. \tag{59}$$

A useful aspect of this formulation is its local interpretation. Under mild regularity assumptions, $\dim_H(\mu)$ can be characterised by the scaling of probability mass around typical points under μ. If, for almost every $\mathbf{x}$,

$$\mu(B_\varepsilon(\mathbf{x})) \sim \varepsilon^{\,d_\mu} \qquad \text{as } \varepsilon \to 0, \tag{60}$$

then $d_\mu = \dim_H(\mu)$. Thus, this dimension reflects how probability mass concentrates across scales.

In (smooth) deterministic dynamical systems, the interplay between expansion and contraction along different directions governs this local scaling behaviour. This is well characterised by Lyapunov exponents that quantify the exponential deformation of infinitesimal neighbourhoods, while the metric (Kolmogorov-Sinai) entropy h_μ quantifies the rate at which trajectories generate information. Well known results in ergodic theory (Ledrappier and Young, 1985) show that the Hausdorff dimension of an invariant measure can be expressed directly in terms of these quantities, and **is strictly smaller than the ambient space dimension d in dissipative systems** with non-trivial Lyapunov exponents, i.e. both positive and negative exponents.

More precisely, according to the Oseledets' theorem (Oseledets, 1968), the system has a Lyapunov spectrum $\lambda_1 \geq \cdots \geq \lambda_d$, and dissipativity implies on average volume contraction, i.e.

$$\sum_{i=1}^{d} \lambda_i < 0. \tag{61}$$

(Ledrappier and Young, 1985) formulate an expression for the Hausdorff dimension of the invariant measure μ in terms of the Lyapunov exponents $\{\lambda_i\}_{i=1}^d$ and the Kolmogorov-Sinai entropy h_μ (Bárány and Käenmäki, 2017)

$$\dim_H(\mu) = k + \frac{h_\mu - \sum_{i=1}^{k} \lambda_i}{|\lambda_{k+1}|}, \tag{62}$$

where k is the largest integer for which $\sum_{i=1}^{k} \lambda_i \geq h_\mu$. This relation holds under standard smoothness and hyperbolicity assumptions (for instance for $C^{1+\alpha}$ systems with non-zero Lyapunov exponents almost everywhere). Intuitively, k here quantifies the number of expanding dimensions needed to characterise the system's entropy.

Since the sum of all Lyapunov exponents is negative (Eq. 61), and the metric entropy is bounded by the sum of positive Lyapunov exponents (Ruelle, 1978)

$$0 \leq h_\mu \leq \sum_{\lambda_i > 0} \lambda_i, \tag{63}$$

the equality of Eq. 62 implies

$$\dim_H(\mu) < d, \tag{64}$$

indicating that the invariant measure concentrates on a subset of the state space, whose Hausdorff dimension is strictly smaller than the ambient space dimension d.

G.2. Dimensionality of invariant measures induced by stochastic dynamics

We now consider stochastic dynamical systems of the form

$$\mathrm{d}\mathbf{X}_t = \mathbf{f}(\mathbf{X}_t)\mathrm{d}t + \sigma\,\mathrm{d}\mathbf{W}_t, \tag{65}$$

similar to the systems we discuss in the main text. Under mild conditions on $\mathbf{f}$ and σ, the corresponding Markov semi-group admits a unique invariant probability measure μ_σ, which coincides with the stationary solution of the associated Fokker–Planck equation (Risken, 1996).

The additive noise regularises the deterministic invariant measure, nevertheless its density concentrates exponentially around $\mathcal{A}$ as $\sigma \to 0$. For non-degenerate noise σ, the Hörmander condition ensures that μ_σ is absolutely continuous with respect to the Lebesgue measure, and thus possesses a *smooth* invariant density (Hörmander, 1967). However, the invariant measure μ_σ of the stochastic system of Eq. 65 satisfies the following exponential concentration inequality around the deterministic attractor $\mathcal{A}$ **for sufficiently small noise amplitude** σ

$$\mu_\sigma\big(\{\mathbf{x} \in \mathbb{R}^d : \mathrm{dist}(\mathbf{x}, \mathcal{A}) > \delta\}\big) \leq C(\delta)\exp\left(-\frac{c(\delta)}{\sigma^2}\right), \tag{66}$$

for all $\delta > 0$, where $C(\delta), c(\delta) > 0$ denote δ-dependent constants, that are nevertheless independent of noise amplitude σ (see Theorem 4.2.1 (Kifer, 1988)). This exponential concentration indicates that, although μ_σ is absolutely continuous with respect to the Lebesgue measure for $\sigma > 0$, it becomes increasingly confined near $\mathcal{A}$ as $\sigma \to 0$. Thus the effective dimension of μ_σ approaches that of the invariant measure of the deterministic system μ_0, while remaining bounded above by the ambient dimension d (Kifer, 1988; Arnold, 2006).

In this sense, stochasticity does not destroy the low-dimensional structure induced by the deterministic dynamics, but thickens the invariant measure around the deterministic attractor. Thus even though μ_σ is smooth, its effective dimension can still be low-dimensional in the sense of mass being tightly concentrated near a lower-dimensional skeleton determined by the underlying deterministic dynamics.

To demonstrate the intuition of behind the low dimensional manifold approximation mentioned in the main text, we embedded the 2 dimensional Van der Pol drift into an D-dimensional ($D \in [3, 10]$ space via a linear projection, add a stable linear force in the subspace orthogonal to the embedding, adding also D-dimensional isotropic noise, and simulated the resulting system. The simulated systems were full D-dimensional stochastic dynamical systems. Subsequently, we estimated the effective dimensionality of the observation manifold by computing the participation ratio of the covariance spectrum of the observations. Let $\lambda_1, \ldots, \lambda_D$ denote the eigenvalues of the empirical covariance matrix of the observed states. The participation ratio is defined as

$$\mathrm{PR} = \frac{\left(\sum_{i=1}^{D} \lambda_i\right)^2}{\sum_{i=1}^{D} \lambda_i^2}.$$

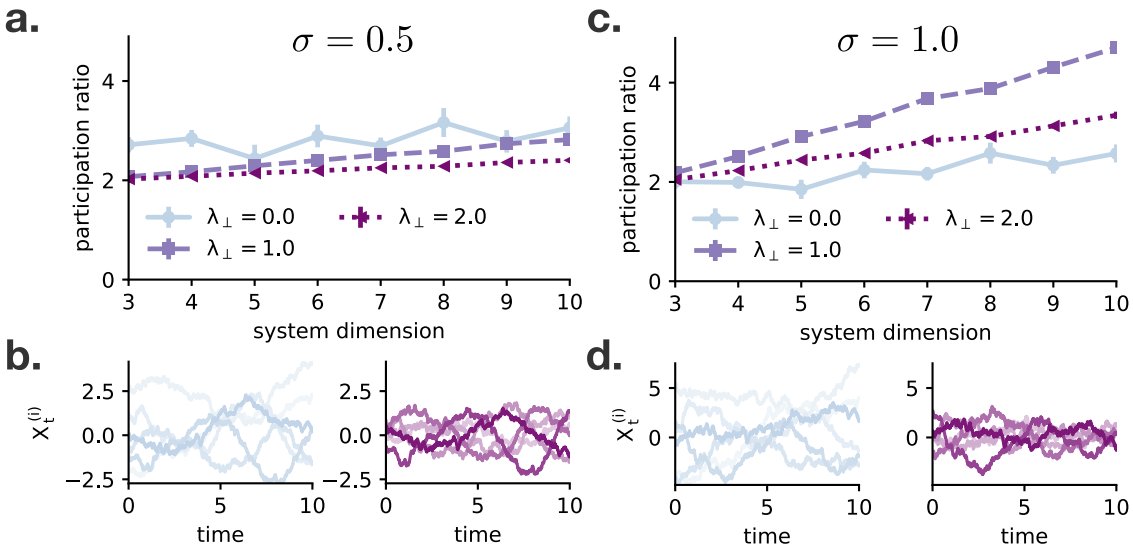

Figure 7

Effective dimensionality of the invariant density as a function of ambient system dimension for a D-dimensional system obtained by linearly embedding a 2-dimensional Van der Pol dynamics into spaces of dimension $D \in \{3, \ldots, 10\}$, with isotropic D-dimensional noise and an additional restoring drift acting in directions perpendicular to the embedded subspace. Effective dimensionality is quantified by the participation ratio for noise amplitudes **a.)** $\sigma = 0.5$ and **c.)** $\sigma = 1.0$. Indicated values are averages over 10 independent runs for each setting. Different curves correspond to different values of $\lambda_\perp$, which controls the strength of the restoring drift in directions orthogonal to the embedded 2-dim. subspace: $\lambda_\perp = 0$ corresponds to the absence of perpendicular confinement, whereas larger values of $\lambda_\perp$ increasingly pull trajectories toward the embedded manifold. **b., d.)** Example sample trajectories for systems of dimension $D = 6$, shown for two values of $\lambda_\perp = 0.0$ (*left*) and $\lambda_\perp = 2.0$ (*right*).

H. Theoretical justification of geometric augmentation for large inter-observation intervals

In the following sections we provide a theoretical analysis justifying our choice to employ geometric path augmentation to improve inference in the large inter-observation limit. In particular, in Sec H.1, we revisit the fact that inference starting from the Euler-Maruyama discretisation deteriorates for increasing inter-observation interval. Then we study the terms in the remainder of the discretisation that become important when the time step (or inter-observation interval) is large, and connect these terms with the geometry of the unknown vector field. We show that for non-linear systems the remainder contains terms related to the curvature of the flow, and that neglecting these terms amounts to assuming a vector field with straight flow-lines in-between observations. This introduces a bias in inference that is linear in the step size. By approximating the curvature by means of controlled path augmentation with reference the geodesic curves of the invariant manifold, our method partially accounts for these remainder terms.

H.1. Inference performance deteriorates with increasing inter-observation interval for existing frameworks

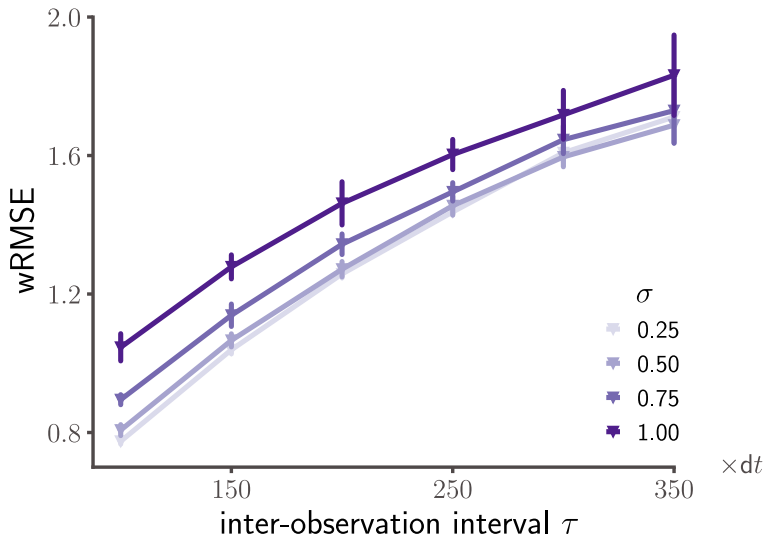

Figure 8
Increasing observation interval between successive observations τ deteriorates performance quantified by increasing weighted root mean squared error (wRMSE) for Gaussian process-based inference. Weighted root mean square error between estimated and ground truth drift vector fields for increasing observation interval τ between subsequent observations for different noise conditions (indicated by different hues). Observations were collected from a Van der Pol oscillator system simulated with $dt = 0.01$ for $T = 500$ time units. Error bars indicate one standard deviation over ten independent realizations.

We computed the weighted root mean square error (wRMSE) between ground truth flow fields and estimated ones for several commonly applied inference frameworks. We observed that the performance of all of them deteriorates once the inter-observation interval becomes large.

We started with the method that motivated our research, approximating drift functions through Gaussian processes, the method outlined in (Ruttor et al., 2013). The method approximates the drift functions with Gaussian process regression, using the system state $\mathbf{X}_t$ as the regressor and state increments as the response variable $\mathbf{Y}_t \doteq \frac{\mathbf{X}_{t+\tau} - \mathbf{X}_t}{\tau}$. This is the Bayesian counterpart of earlier methods encountered in physics literature (Friedrich and Peinke, 1997; Ragwitz and Kantz, 2001), providing additionally uncertainty estimation through the Gaussian process approximation.

As is evident from Figure 8 the discrepancy between ground truth and estimated vector fields increases for increasing temporal distance between successive observations. This should be understood, under the consideration that inference of the drift based on regression on state increments results from an approximation relying on a truncated Ito-Taylor expansion. This is also the starting point of the Euler Maruyama discretisation. As the time interval between successive steps of this approximation increases, the truncated approximation does not longer hold, and higher order terms should be considered.

H.2. Inference based on Euler-Maruyama discretisation does not account for the curvature of the trajectories in the state space

To be more precise, a general SDE of the form

$$\mathrm{d}\mathbf{X}_t = \mathbf{f}(\mathbf{X}_t, t)\mathrm{d}t + \boldsymbol{\sigma}(\mathbf{X}_t, t)\mathrm{d}\mathbf{W}_t. \tag{67}$$

is a shorthand for the integral equation

$$\mathbf{X}_t = \mathbf{X}_{t_0} + \int_{t_0}^{t} \mathbf{f}(\mathbf{X}_s, s)\,\mathrm{d}s + \int_{t_0}^{t} \boldsymbol{\sigma}(\mathbf{X}_s, s)\,\mathrm{d}\mathbf{W}_s, \tag{68}$$

where as previously in this manuscript, we consider the stochastic integrals in the **Itô sense**. (Here we start from a more general formulation of the stochastic equation with both diffusion and drift terms being state- and time-dependent to highlight that also for more general SDEs our geometric argument is valid.)

Applying the Itô formula on each integrand, and integrating from t_0 to t, we obtain the Itô-Taylor expansion of Eq. 67

$$\begin{aligned}
\mathbf{f}(\mathbf{X}_t, t) =& \mathbf{f}(\mathbf{X}_{t_0}, t_0) + \int_{t_0}^{t} \frac{\partial \mathbf{f}(\mathbf{X}_s, s)}{\partial s}\,\mathrm{d}s + \int_{t_0}^{t} \sum_u \frac{\partial \mathbf{f}(\mathbf{X}_s, s)}{\partial X^{(u)}} f_u(\mathbf{X}_s, s)\,\mathrm{d}s \\
&+ \int_{t_0}^{t} \sum_u \frac{\partial \mathbf{f}(\mathbf{X}_s, s)}{\partial X^{(u)}} \left[\boldsymbol{\sigma}(\mathbf{X}_s, s)\,\mathrm{d}\mathbf{W}_s\right]_u + \int_{t_0}^{t} \frac{1}{2} \sum_{u,v} \frac{\partial^2 \mathbf{f}(\mathbf{X}_s, s)}{\partial X^{(u)} \partial X^{(v)}} \left[\boldsymbol{\sigma}(\mathbf{X}_s, s)\,\boldsymbol{\sigma}^\top(\mathbf{X}_s, s)\right]_{uv} \mathrm{d}s \\
=& \mathbf{f}(\mathbf{X}_{t_0}, t_0) + \int_{t_0}^{t} \mathcal{L}_s^\dagger \mathbf{f}(\mathbf{X}_s, s)\,\mathrm{d}s + \sum_\nu \int_{t_0}^{t} \mathcal{L}_{W,\nu} \mathbf{f}(\mathbf{X}_s, s)\,\mathrm{d}W_s^{(\nu)},
\end{aligned} \tag{69}$$

and

$$\begin{aligned}
\boldsymbol{\sigma}(\mathbf{X}_t, t) =& \boldsymbol{\sigma}(\mathbf{X}_{t_0}, t_0) + \int_{t_0}^{t} \frac{\partial \boldsymbol{\sigma}(\mathbf{X}_s, s)}{\partial s}\,\mathrm{d}s + \int_{t_0}^{t} \sum_u \frac{\partial \boldsymbol{\sigma}(\mathbf{X}_s, s)}{\partial X^{(u)}} f_u(\mathbf{X}_s, s)\,\mathrm{d}s \\
&+ \int_{t_0}^{t} \sum_u \frac{\partial \boldsymbol{\sigma}(\mathbf{X}_s, s)}{\partial X^{(u)}} \left[\boldsymbol{\sigma}(\mathbf{X}_s, s)\,\mathrm{d}\mathbf{W}_s\right]_u + \int_{t_0}^{t} \frac{1}{2} \sum_{u,v} \frac{\partial^2 \boldsymbol{\sigma}(\mathbf{X}_s, s)}{\partial X^{(u)} \partial X^{(v)}} \left[\boldsymbol{\sigma}(\mathbf{X}_s, s)\,\boldsymbol{\sigma}^\top(\mathbf{X}_s, s)\right]_{uv} \mathrm{d}s \\
=& \boldsymbol{\sigma}(\mathbf{X}_{t_0}, t_0) + \int_{t_0}^{t} \mathcal{L}_s^\dagger \boldsymbol{\sigma}(\mathbf{X}_s, s)\,\mathrm{d}s + \sum_\nu \int_{t_0}^{t} \mathcal{L}_{W,\nu} \boldsymbol{\sigma}(\mathbf{X}_s, s)\,\mathrm{d}W_s^{(\nu)},
\end{aligned} \tag{70}$$

where we have used the fact that the product of stochastic differentials due to the Ito isometry and multiplication rules equals the noise covariance times the time step

$$dX_t^{(u)} dX_t^{(v)} = \left[\boldsymbol{\sigma}\boldsymbol{\sigma}^\top\right]_{uv} dt,$$

where

$$dX_s^{(u)} = f_u\, ds + \sum_{j=1}^{m} \sigma_{uj}\, dW_s^{(j)},$$

while the superscripts/subscripts u, v indicate dimensional components.

In the above equations, we have introduced the operators acting on test-functions $\mathbf{h} : \mathbb{R}^D \to \mathbb{R}^D$

$$\mathcal{L}_t^\dagger \mathbf{h} = \frac{\partial \mathbf{h}}{\partial t} + \sum_u \frac{\partial \mathbf{h}}{\partial X^{(u)}} f_u + \frac{1}{2} \sum_{u,v} \frac{\partial^2 \mathbf{h}}{\partial X^{(u)} \partial X^{(v)}} \left[\boldsymbol{\sigma}(\mathbf{X}_s, s)\,\boldsymbol{\sigma}^\top(\mathbf{X}_s, s)\right]_{uv} \tag{71}$$

and

$$\mathcal{L}_{W,v} \mathbf{h} = \sum_u \frac{\partial \mathbf{h}}{\partial X^{(u)}}\, \boldsymbol{\sigma}_{uv}, \qquad \text{for } v = 1, \ldots, n. \tag{72}$$

With these expressions, the original integral equation for $\mathbf{X}_t$ can be written as

$$\mathbf{X}_t = \mathbf{X}_{t_0} + \mathbf{f}(\mathbf{X}_{t_0}, t_0)(t - t_0) + \boldsymbol{\sigma}(\mathbf{X}_{t_0}, t_0)\,(\mathbf{W}_t - \mathbf{W}_{t_0}) + \tag{73}$$

$$R_1 = \begin{cases} + \displaystyle\int_{t_0}^{t}\int_{t_0}^{s} \mathcal{L}_u^\dagger \mathbf{f}(\mathbf{X}_u, u)\,\mathrm{d}u\,\mathrm{d}s + \sum_\nu \int_{t_0}^{t}\int_{t_0}^{s} \mathcal{L}_{W,\nu}\mathbf{f}(\mathbf{X}_u, u)\,\mathrm{d}W_u^{(\nu)}\,\mathrm{d}s \\ + \displaystyle\int_{t_0}^{t}\int_{t_0}^{s} \mathcal{L}_u^\dagger \boldsymbol{\sigma}(\mathbf{X}_u, u)\,\mathrm{d}u\,\mathrm{d}\mathbf{W}_s + \sum_\nu \int_{t_0}^{t}\int_{t_0}^{s} \mathcal{L}_{W,\nu}\boldsymbol{\sigma}(\mathbf{X}_u, u)\,\mathrm{d}W_u^{(\nu)}\,\mathrm{d}\mathbf{W}_s. \end{cases}$$

In the last equation, dropping the terms in the remainder R_1 results in the Euler–Maruyama integration scheme (Jentzen and Kloeden, 2011). Introducing the discrete time and noise increments

$$\Delta t_n = t_{n+1} - t_n = \int_{t_n}^{t_{n+1}} \mathrm{d}s, \quad \Delta\mathbf{W}_n = \mathbf{W}_{t_{n+1}} - \mathbf{W}_{t_n} = \int_{t_n}^{t_{n+1}} \mathrm{d}\mathbf{W}_s, \tag{74}$$

we result in the discrete time equation commonly used for numerical integration of SDEs

$$\mathbf{X}_{n+1} = \mathbf{X}_n + \mathbf{f}(\mathbf{X}_n, t_n)\,\Delta t_n + \boldsymbol{\sigma}\,\Delta\mathbf{W}_n. \tag{75}$$

This is also the starting point of most inference methods that employ the regression scheme mentioned above by approximating the drift as

$$\hat{\mathbf{f}}(\mathbf{X}_n, t_n) \approx \frac{\mathbf{X}_{n+1} - \mathbf{X}_n}{\Delta t} \sim \mathcal{N}\left(\hat{\mathbf{f}}(\mathbf{X}_n, t_n), \frac{\boldsymbol{\sigma}\,\boldsymbol{\sigma}^\top}{\Delta t}\right). \tag{76}$$

This discretisation is a zero-order approximation of the ground truth dynamics, and assumes that $\mathbf{f}(\cdot)$ remains constant throughout the interval Δt, i.e. throughout the inter-observation interval τ in the inference setting. However as τ increases, higher-order terms in the remainder R_1 of the Itô-Taylor expansion become significant, since the assumption that the drift is approximately constant over τ does not hold.

We can glean onto the terms that become important once the inter-observation interval becomes large, by applying the Itô formula onto each one of the integrands in R_1 separately **for the specific setting we consider in this manuscript**, i.e. that of time-independent drift function $\mathbf{f}(\mathbf{x})$ and constant diffusion matrix $\boldsymbol{\sigma}$. In the following, we demonstrate that the leading-order error of this approximation is governed by the intrinsic geometry of the drift vector field. This provides further insight and a geometric explanation for the deterioration of inference methods for increasing inter-observation interval τ.

In short we show that, inference methods based on the Euler-Maruyama discretisation-based inference effectively assume that the vector field between consecutive observations $\mathbf{X}_n$ and $\mathbf{X}_{n+1}$ does not change. Our analysis shows this is equivalent to assuming trajectories are straight lines ($\mathbf{J}_f\mathbf{f} \parallel \mathbf{f}$) and the Itô correction is constant. In reality, trajectories curve ($\mathbf{J}_f\mathbf{f}$ has also a perpendicular component), and this curvature itself changes along the vector field. The Euler-Maruyama discretisation-based inference scheme systematically misses these higher-order geometric features, leading to biased drift estimates.

H.2.1. First remainder term $R_{1,a}$

We denote the first term of the reminder by $R_{1,a}$

$$R_{1,a} = \int_{t_0}^{t}\int_{t_0}^{s} \mathcal{L}_u^\dagger \mathbf{f}(\mathbf{X}_u)\,\mathrm{d}u\,\mathrm{d}s. \tag{77}$$

Applying Itô's formula to the integrand $\mathcal{L}_t^\dagger \mathbf{f}(\mathbf{X}_u, u)$, we get

$$\mathrm{d}\mathcal{L}_u^\dagger \mathbf{f}(\mathbf{X}_u) = \frac{\partial}{\partial u}\mathcal{L}_u^\dagger \mathbf{f}(\mathbf{X}_u)\,\mathrm{d}u + \sum_{j=1}^{d} \frac{\partial \mathcal{L}_u^\dagger \mathbf{f}}{\partial X^{(j)}}(\mathbf{X}_u)\,\mathrm{d}X_u^{(j)} + \frac{1}{2}\sum_{j,k=1}^{d} \frac{\partial^2 \mathcal{L}_u^\dagger \mathbf{f}}{\partial X^{(j)}\partial X^{(k)}}(\mathbf{X}_u)\,\left[\boldsymbol{\sigma}\boldsymbol{\sigma}^\top\right]_{jk}\,\mathrm{d}u. \tag{78}$$

Plugging in the original equation $\mathrm{d}X_u^{(j)} = f_j\,\mathrm{d}u + \sum_{\nu=1}^m \sigma_{j\nu}\,\mathrm{d}W_u^{(\nu)}$, and integrating over the time from t_0 to u

$$\mathcal{L}_u^\dagger \mathbf{f}(\mathbf{X}_u) = \mathcal{L}_{t_0}^\dagger \mathbf{f}(\mathbf{X}_{t_0}) + \int_{t_0}^u \left(\frac{\partial}{\partial w}(\mathcal{L}_w^\dagger \mathbf{f}(\mathbf{X}_w)) + \sum_j \frac{\partial(\mathcal{L}_w^\dagger \mathbf{f})}{\partial X^{(j)}} f_j + \frac{1}{2}\sum_{j,k} \frac{\partial^2(\mathcal{L}_w^\dagger \mathbf{f})}{\partial X^{(j)}\partial X^{(k)}} [\boldsymbol{\sigma}\boldsymbol{\sigma}^\top]_{jk} \right) \mathrm{d}w$$
$$+ \int_{t_0}^u \sum_j \frac{\partial(\mathcal{L}_w^\dagger \mathbf{f})}{\partial X^{(j)}} [\boldsymbol{\sigma}\mathrm{d}\mathbf{W}_w]_j. \tag{79}$$

Applying Fubini's theorem in the original double integral, we change the order of integration

$$\int_{t_0}^t \int_{t_0}^s \phi(u)\,\mathrm{d}u\,\mathrm{d}s = \int_{t_0}^t (t-u)\,\phi(u)\,\mathrm{d}u, \tag{80}$$

and applying it twice we obtain

$$R_{1,a} = \int_{t_0}^t \int_{t_0}^s \mathcal{L}_u^\dagger \mathbf{f}(\mathbf{X}_u)\,\mathrm{d}u\,\mathrm{d}s = \int_{t_0}^t \frac{(t-u)^2}{2} \left[\underbrace{\sum_j \frac{\partial \mathcal{L}_u^\dagger \mathbf{f}}{\partial X^{(j)}} f_j}_{R_{1,a}^1} + \underbrace{\frac{1}{2}\sum_{j,k} \frac{\partial^2 \mathcal{L}_u^\dagger \mathbf{f}}{\partial X^{(j)}\partial X^{(k)}} [\boldsymbol{\sigma}\boldsymbol{\sigma}^\top]_{jk}}_{R_{1,a}^2} \right] \mathrm{d}u$$
$$+ \int_{t_0}^t \frac{(t-u)^2}{2} \underbrace{\sum_j \frac{\partial \mathcal{L}_u^\dagger \mathbf{f}}{\partial X^{(j)}} [\boldsymbol{\sigma}\,\mathrm{d}\mathbf{W}_u]_j}_{R_{1,a}^3} + \frac{\tau^2}{2}\mathcal{L}_t^\dagger \mathbf{f}(\mathbf{X}_{t_0}). \tag{81}$$

In the previous equation we have dropped the term $\frac{\partial}{\partial w}\left(\mathcal{L}_w^\dagger\,\mathbf{f}(\mathbf{X}_w)\right)$ that is equal to zero and that would require the drift $\mathbf{f}$ to be time-dependent to be non-negligible.

First component $R_{1,a}^1$ of remainder term $R_{1,a}$: Flow curvature term. The Backward Kolmogorov generator applied to a vector field $\mathbf{f}$ can be written as

$$\mathcal{L}^\dagger \mathbf{f} \;=\; \mathbf{J}_f\,\mathbf{f} \;+\; \frac{1}{2}\,\Delta_D \mathbf{f}. \tag{82}$$

In Eq. 82, $\mathbf{J}_f \doteq \nabla \mathbf{f}$ denotes the Jacobian of $\mathbf{f}$, $\mathbf{D} \doteq \boldsymbol{\sigma}\boldsymbol{\sigma}^\top$ the noise covariance, and $\Delta_\mathbf{D} \doteq \sum_{j,k} \mathbf{D}_{jk}\,\partial^2_{X^{(j)}X^{(k)}}$ is the noise-covariance weighted Laplacian operator. Thus each component of $\mathcal{L}^\dagger \mathbf{f}$ comprises the directional derivative of the drift $\mathbf{J}_f \mathbf{f}$ plus an anisotropic/noise-covariance weighted Laplacian of $\mathbf{f}$, which in component-wise form is expressed as

$$\left[\mathcal{L}^\dagger \mathbf{f}\right]_i = \sum_k \frac{\partial f_i}{\partial X^{(k)}}\,f_k \;+\; \frac{1}{2}\sum_{k,\ell} \mathbf{D}_{k\ell}\,\frac{\partial^2 f_i}{\partial X^{(k)}\partial X^{(\ell)}}. \tag{83}$$

Differentiating wrt to $X^{(j)}$ yields

$$\frac{\partial}{\partial X^{(j)}}\left[\mathcal{L}^\dagger \mathbf{f}\right]_i = \sum_k \frac{\partial^2 f_i}{\partial X^{(j)}\partial X^{(k)}}\,f_k \;+\; \sum_k \frac{\partial f_i}{\partial X^{(k)}}\,\frac{\partial f_k}{\partial X^{(j)}} \;+\; \frac{1}{2}\sum_{k,\ell} \mathbf{D}_{k\ell}\,\frac{\partial^3 f_i}{\partial X^{(j)}\partial X^{(k)}\partial X^{(\ell)}}, \tag{84}$$

and thus we rewrite the i-th component of the term $R_{1,a}^1$ as

$$\left[R_{1,a}^1\right]_i = \int_{t_0}^t \frac{(t-u)^2}{2} \left[\sum_{j,k} \frac{\partial^2 f_i}{\partial X^{(j)}\partial X^{(k)}}\,f_k\,f_j + \sum_{j,k} \frac{\partial f_i}{\partial X^{(k)}}\,\frac{\partial f_k}{\partial X^{(j)}}\,f_j + \frac{1}{2}\sum_{j,k,\ell} \mathbf{D}_{k\ell}\,\frac{\partial^3 f_i}{\partial X^{(j)}\partial X^{(k)}\partial X^{(\ell)}}\,f_j \right]_i \mathrm{d}u. \tag{85}$$

The third-order state-derivative in the last summand of Eq. 85, indicates that this last term is inactive for linear or quadratic drift functions $\mathbf{f}$.

We re-write again this part of the remainder in a more compact vector notation in terms of the directional derivative of $(\mathbf{J}_f\mathbf{f})$ and $\frac{1}{2}\,\Delta_D\mathbf{f}$ along the vector field as

$$R^1_{1,a} = \int_{t_0}^{t} \frac{(t-u)^2}{2} \Big[\underbrace{\nabla(\mathbf{J}_f\mathbf{f})\cdot\mathbf{f}}_{\text{flow curvature}} + \underbrace{\nabla\big(\tfrac{1}{2}\,\Delta_D\mathbf{f}\big)\cdot\mathbf{f}}_{\text{diffusive term along the flow}} \Big]\, \mathrm{d}u. \tag{86}$$

This part of the remainder captures two geometric effects that standard inference methods neglect: the **intrinsic curvature of deterministic flow trajectories in state space**, and the **systematic bias introduced by the spatial variation of both drift and diffusion** along these trajectories, when both drift and diffusion are assumed as constant between inter-observation intervals.

- To understand the **first term**, $\nabla(\mathbf{J}_f\mathbf{f})\cdot\mathbf{f}$, from a geometric perspective, let us consider a deterministic dynamical system with dynamics $\dot{\mathbf{x}}_t = \mathbf{f}(\mathbf{x}_t)$. A trajectory initiated from an initial condition $\mathbf{x}_0$ traces a streamline in the state space $\mathbb{R}^d$. We express the acceleration of this trajectory in terms of the directional derivative

$$\ddot{\mathbf{x}}_t = \frac{\mathrm{d}}{\mathrm{d}t}\mathbf{f}(\mathbf{x}_t) = \mathbf{J}_f(\mathbf{x}_t)\cdot\mathbf{f}(\mathbf{x}_t) = \mathbf{J}_f\cdot\mathbf{f}. \tag{87}$$

  The acceleration vector admits a natural orthogonal decomposition comprising a component parallel to the vector field $\mathbf{f}$ and an orthogonal component to $\mathbf{f}$

$$\mathbf{J}_f\cdot\mathbf{f} = P_{\|}(\mathbf{f})\,\mathbf{J}_f\cdot\mathbf{f} + P_{\perp}(\mathbf{f})\,\mathbf{J}_f\cdot\mathbf{f}. \tag{88}$$

  Here $P_{\|}(\mathbf{f}(\mathbf{x})) = \frac{\mathbf{f}(\mathbf{x})\mathbf{f}^\top(\mathbf{x})}{\|\mathbf{f}(\mathbf{x})\|^2}$ and $P_{\perp}(\mathbf{f}(\mathbf{x})) = \mathbb{I} - P_{\|}(\mathbf{f}(\mathbf{x}))$ stand for the parallel and orthogonal projectors. The parallel component quantifies the rate of change of speed along the trajectory (tangential acceleration), whilst the perpendicular component defines the **curvature vector** $\kappa_{\text{flow}}(x)$ (Kühnel, 2002), which quantifies the bending of the trajectories

$$\boldsymbol{\kappa}_{\text{flow}}(\mathbf{x}) = \frac{P_{\perp}(\mathbf{f}(\mathbf{x}))\mathbf{J}_f(\mathbf{x})\mathbf{f}(\mathbf{x})}{\|\mathbf{f}(\mathbf{x})\|^2}. \tag{89}$$

  When $\boldsymbol{\kappa}_{\text{flow}} = 0$, the trajectories are straight lines in the state space, while when $\|\boldsymbol{\kappa}_{\text{flow}}\| > 0$ they are curved.

  The term $\nabla(\mathbf{J}_f\mathbf{f})\cdot\mathbf{f}$ quantifies the **evolution of the trajectory curvature** [4] as the system moves along the flow field. From Eq. 85 we have for each dimensional component i of this term

$$\begin{aligned}[\nabla(\mathbf{J}_f\mathbf{f})\cdot\mathbf{f}]_i &= \sum_{j,k} \frac{\partial^2 f_i}{\partial X^{(j)}\partial X^{(k)}} f_k f_j + \sum_{j,k} \frac{\partial f_i}{\partial X^{(k)}}\frac{\partial f_k}{\partial X^{(j)}} f_j \\ &= [\mathbf{f}^\top(\nabla^2 f_i)\mathbf{f}] + [\mathbf{J}_f^2\mathbf{f}]_i.\end{aligned} \tag{90}$$

  We observe that this term captures the effects of how both second-order spatial variation of the flow field (the Hessian $\nabla^2 f_i$) and the Jacobian of the acceleration ($\mathbf{J}_f^2\mathbf{f}$) influence the evolution of trajectories.

  - In Eq. 90, the **first sub-term**, $\mathbf{f}^\top(\nabla^2 f_i)\mathbf{f}$, represents the **second directional derivative** (or quadratic variation) of f_i along the flow direction $\mathbf{f}$. It measures the curvature or second-order spatial variation of the i-th component of $\mathbf{f}$ in the direction $\mathbf{f}$. In regions where the Hessian $\nabla^2\mathbf{f}$ is large (as is for the case of a highly nonlinear drift with curving or bending behaviour), this term becomes significant, and it vanishes for linear or constant drift $\mathbf{f}$. Neglecting this term corresponds to approximating the flow by its linearisation.
  - The **second sub-term**, $\mathbf{J}_f^2\mathbf{f} = \mathbf{J}_f(\mathbf{J}_f\mathbf{f})$, of Eq. 90 represents the action of the Jacobian operator on the acceleration vector. Geometrically, it describes how the local linearised field acts on the acceleration as we move an infinitesimal step along the flow field, or in other words how the linear approximation changes when following the flow direction $\mathbf{f}$.

[4] More precisely the directional derivative of the acceleration, $\mathbf{J}_f(\mathbf{x})\cdot\mathbf{f}$ along the flow direction, or the **rate at which the acceleration changes along the flow, or a measure of how the local curvature of f as a vector field influences trajectory evolution**.

By temporal integration we have

$$R^1_{1,a} = \int_{t_0}^{t} \frac{(t-u)^2}{2} \Big(\nabla(\mathbf{J}_f \mathbf{f}) \cdot \mathbf{f} + \nabla\big(\tfrac{1}{2}\,\Delta_D \mathbf{f}\big) \cdot \mathbf{f} \Big)\, \mathrm{d}u \sim \frac{\tau^2}{2} \Big(\nabla(\mathbf{J}_f \mathbf{f}) \cdot \mathbf{f} + \nabla\big(\tfrac{1}{2}\,\Delta_D \mathbf{f}\big) \cdot \mathbf{f} \Big), \tag{91}$$

indicating that the evolution of trajectory curvature introduces an $O(\tau^3)$ correction to the transition density.

Drift inference based on Euler–Maruyama–type discretisation ignores between others the term $R^1_{1,a}$ introducing thereby a mean bias at each point $\mathbf{x}$ in the state space,

$$\beta^1_{1,a}(\mathbf{x}) = \frac{1}{\tau}\, R^1_{1,a} \;\approx\; \frac{\tau}{2} \left[\nabla(\mathbf{J}_f \mathbf{f}) \cdot \mathbf{f} + \nabla\big(\tfrac{1}{2}\,\Delta_D \mathbf{f}\big) \cdot \mathbf{f} \right](\mathbf{x}). \tag{92}$$

This bias induces a mean error in drift estimate, when using Euler–Maruyama-based inference, leading to under- or over-estimation of the ground truth drift at state $\mathbf{x}$. This error scales as $\mathcal{O}(\tau^2)$ with the interval τ.

Let us now consider the temporal rate of change experienced by a particle travelling along the flow field. The instantaneous speed of the particle at location $\mathbf{x}$ is $\|\mathbf{f}(\mathbf{x})\|$. The quantity in the brackets in Eq. 92, $\nabla(\mathbf{J}_f \mathbf{f}) \cdot \mathbf{f} + \nabla\left(\frac{1}{2}\Delta_D \mathbf{f}\right) \cdot \mathbf{f}$, is a spatial derivative measuring how quickly the curvature and diffusion variation change as the particle moves in space. The rate of change of this variation per unit of time is expressed as

$$\frac{\left\| \nabla(\mathbf{J}_f \mathbf{f}) \cdot \mathbf{f} + \nabla\left(\frac{1}{2}\Delta_D \mathbf{f}\right) \cdot \mathbf{f} \right\| (\mathbf{x})}{\|\mathbf{f}(\mathbf{x})\|} \doteq \tau^{-2}_{\mathrm{curv}}(\mathbf{x}). \tag{93}$$

In the last equation we have introduced the time scale of change τ_{curv} as the inverse of the rate of change, which captures the characteristic time it takes for the curvature/diffusion variation to change significantly along the particles trajectory. Then the relative magnitude error in the Euler-Maruyama-based drift estimate satisfies

$$\frac{\|\beta^1_{1,a}(\mathbf{x})\|}{\|\mathbf{f}(\mathbf{x})\|} = \frac{\tau^2}{6\,\tau^2_{\mathrm{curv}}(\mathbf{x})}, \tag{94}$$

implying that the estimate is reliable only when the inter-observation interval $\tau \ll \tau_{\mathrm{curv}}(\mathbf{x})$.

- The **second term** in Eq.86, $\nabla(\frac{1}{2}\Delta_D \mathbf{f}) \cdot \mathbf{f}$, accounts for the diffusion part of the backward generator acting on the vector field $\mathbf{f}$. The anisotropic Laplacian $\Delta_D \mathbf{f}$ quantifies the **diffusion–weighted second-order spatial variation of the vector field**

$$[\Delta_D \mathbf{f}]_i = \sum_{j,k} D_{jk} \frac{\partial^2 f_i}{\partial X^{(j)} \partial X^{(k)}} = \nabla \cdot \big(\mathbf{D}\, \nabla f_i \big). \tag{95}$$

The directional derivative quantifies how this term evolves along the flow field

$$\left[\nabla\big(\tfrac{1}{2}\Delta_D \mathbf{f}\big) \cdot \mathbf{f} \right]_i = \frac{1}{2} \sum_{j,k,\ell} D_{k\ell} \frac{\partial^3 f_i}{\partial X^{(j)} \partial X^{(k)} \partial X^{(\ell)}} f_j. \tag{96}$$

This term captures how the diffusion-weighted spatial variation of the flow field varies across the state space. As trajectories traverse regions of varying drift curvature, the effective Itô correction itself changes, introducing systematic bias in inference methods that assume that drift is piece-wise constant in-between observations.

Second component $R^2_{1,a}$ of remainder term $R_{1,a}$. The second component of the remainder term $R_{1,a}$ reads

$$R^2_{1,a} = \int_{t_0}^{t} (t-u)\, \frac{1}{2} \sum_{j,k} \frac{\partial^2 \big(\mathcal{L}^\dagger_u \mathbf{f}\big)}{\partial X^{(j)} \partial X^{(k)}} \left[\boldsymbol{\sigma}\boldsymbol{\sigma}^\top \right]_{jk} \mathrm{d}u. \tag{97}$$

For the i-th dimensional component we have

$$\begin{aligned}\frac{\partial^2}{\partial X^{(h)}\partial X^{(j)}}\left[\mathcal{L}_u^\dagger f\right]_i = &\sum_k \frac{\partial^3 f_i}{\partial X^{(h)}\partial X^{(j)}\partial X^{(k)}}\, f_k + \sum_k \frac{\partial^2 f_i}{\partial X^{(j)}\partial X^{(k)}}\,\frac{\partial f_k}{\partial X^{(h)}} \\ &+ \sum_k \frac{\partial^2 f_i}{\partial X^{(h)}\partial X^{(k)}}\,\frac{\partial f_k}{\partial X^{(j)}} + \sum_k \frac{\partial f_i}{\partial X^{(k)}}\,\frac{\partial^2 f_k}{\partial X^{(h)}\partial X^{(j)}} \\ &+ \frac{1}{2}\sum_{k,\ell}\mathbf{D}_{k\ell}\,\frac{\partial^4 f_i}{\partial X^{(h)}\partial X^{(j)}\partial X^{(k)}\partial X^{(\ell)}}.\end{aligned} \tag{98}$$

For this remainder term, we have for each dimensional component i

$$\left[R^2_{1,a}\right]_i = \int_{t_0}^{t}(t-u)\,\frac{1}{2}\sum_{j,k}\mathbf{D}_{jk}\,\left[\frac{\partial^2}{\partial X^{(k)}\partial X^{(j)}}\left[\mathcal{L}_u^\dagger \mathbf{f}\right]_i\right]\,\mathrm{d}u\,. \tag{99}$$

Geometrically, $R^2_{1,a}$ captures the **diffusion-weighted second-order spatial variation** of the generator $\mathcal{L}_u^\dagger\mathbf{f}$ across the $\sqrt{\tau}$-sized ellipsoid set by $\mathbf{D}$, i.e. the anisotropic Laplacian $\Delta_D(\mathcal{L}_u^\dagger\mathbf{f})$, the diffusion-weighted second spatial variation of the drift along the flow. Dropping this term in inference amounts to assuming $\mathcal{L}_u^\dagger\mathbf{f}$ is locally flat and results in an $O(\tau)$ drift bias of size $\beta^2_{1,a} \approx (\tau/4)\,\Delta_D(\mathcal{L}_u^\dagger\mathbf{f})$, underestimating anisotropy and the evolution of curvature of the flow field, so inferred flow-lines appear too straight.

Third component $R^3_{1,a}$ of remainder term $R_{1,a}$.

$$R^3_{1,a} = \int_{t_0}^{t}(t-u)\,\sum_j \frac{\partial\,\mathcal{L}_u^\dagger\mathbf{f}}{\partial X^{(j)}}\,[\sigma\,\mathrm{d}\mathbf{W}_u]_j, \tag{100}$$

$$\left[R^3_{1,a}\right]_i = \int_{t_0}^{t}(t-u)\,\sum_{j,m}\frac{\partial}{\partial X^{(j)}}\left[\mathcal{L}_u^\dagger\mathbf{f}\right]_i\,\sigma_{jm}\,\mathrm{d}\mathbf{W}_u^{(m)}, \tag{101}$$

This is a martingale term capturing the stochastic coupling between diffusion and the spatial inhomogeneity of the generator. In inference, this term doesn't introduce bias, since $\langle R^3_{1,a}\rangle = 0$. However, neglecting this term, ignores a second–order variance contribution with $\mathrm{Var}(R^3_{1,a}/\tau) = O(\tau)$.

H.2.2. Second remainder term $R_{1,b}$

We denote the second term of the reminder by $R_{1,b}$

$$R_{1,b} = \sum_{\nu=1}^{n}\int_{t_0}^{t}\int_{t_0}^{s}\mathcal{L}_{W,\nu}\,\mathbf{f}\,\mathrm{d}W_u^{(\nu)}\,\mathrm{d}s. \tag{102}$$

Applying Fubini's theorem again to change the order of integration, we re-write $R_{1,b}$ in the form of a stochastic integral

$$R_{1,b} = \sum_{\nu=1}^{n}\int_{t_0}^{t}(t-u)\,\mathcal{L}_{W,\nu}\,\mathbf{f}\,\mathrm{d}W_u^{(\nu)}. \tag{103}$$

Substituting the operator results in an expression for each dimensional component i

$$[R_{1,b}]_i = \sum_{\nu=1}^{n}\int_{t_0}^{t}(t-u)\left(\sum_{j=1}^{D}\frac{\partial f_i}{\partial X^{(j)}}\sigma_{j\nu}\right)\mathrm{d}W_u^{(\nu)}, \quad \text{for } i = 1,\ldots,D. \tag{104}$$

In matrix notation, this corresponds to

$$R_{1,b} = \int_{t_0}^{t}(t-u)\,\mathbf{J}_f\,\boldsymbol{\sigma}\,\mathrm{d}\mathbf{W}_u. \tag{105}$$

The remainder $R_{1,b}$ is a stochastic integral with zero mean, but non-zero covariance, given by

$$\mathrm{Cov}(R_{1,b}) = \langle R_{1,b}\, R_{1,b}^\top \rangle = \int_{t_0}^{t} (t-u)^2\, \mathbf{J}_f\, \boldsymbol{\sigma}\boldsymbol{\sigma}^\top\, \mathbf{J}_f^\top\, \mathrm{d}u. \tag{106}$$

For sufficiently smooth $\mathbf{J}_f$ and small time step $\tau = t - t_0$, this covariance scales on the order of τ^3.

The term $R_{1,b}$ quantifies the contribution to the remainder arising from stochastic fluctuations of the noise acting through the spatial derivatives of the drift $\mathbf{f}$. It does not contribute to additional systematic bias, but introduces variance in the drift estimator, especially when $\boldsymbol{\sigma}$ or $\mathbf{J}_f$ are large.

H.2.3. Third remainder term $R_{1,c}$

We denote the third remainder term by $R_{1,c}$ and re-write here for convenience

$$R_{1,c} = \int_{t_0}^{t} \int_{t_0}^{s} \mathcal{L}_u^\dagger \boldsymbol{\sigma}(\mathbf{X}_u, u)\, \mathrm{d}u\, \mathrm{d}\mathbf{W}_s. \tag{107}$$

In the general case of time- and state- dependent diffusion the integrand of this term would be expressed for the i-th row and ℓ-th column component of $\boldsymbol{\sigma}$ as follows

$$\left[\mathcal{L}_u^\dagger \boldsymbol{\sigma}(\mathbf{X}_u, u)\right]_{i\ell} = \frac{\partial}{\partial u}\sigma_{i\ell}(\mathbf{X}_u, u) + \sum_{j=1}^{D} \frac{\partial \sigma_{i\ell}}{\partial X^{(j)}}(\mathbf{X}_u, u) f_j(\mathbf{X}_u, u) \tag{108}$$

$$+ \frac{1}{2}\sum_{j,k=1}^{D} \frac{\partial^2 \sigma_{i\ell}}{\partial X^{(j)} \partial X^{(k)}}(\mathbf{X}_u, u)[\boldsymbol{\sigma}\boldsymbol{\sigma}^\top]_{jk}(\mathbf{X}_u, u). \tag{109}$$

However, in our setting we consider state- and time-independent diffusion matrix, and thus $\mathcal{L}_u^\dagger \boldsymbol{\sigma}(\mathbf{X}_u, u) = \mathbf{0}$, and by consequence $R_{1,c} = \mathbf{0}$

H.2.4. Fourth remainder term $R_{1,d}$

The fourth remainder term is

$$R_{1,d} = \sum_{\nu=1}^{n} \int_{t_0}^{t} \int_{t_0}^{s} \mathcal{L}_{W,\nu}\, \boldsymbol{\sigma}\, \mathrm{d}W_u^{(\nu)}\, \mathrm{d}\mathbf{W}_s. \tag{110}$$

For each component (i, ℓ) of $\boldsymbol{\sigma}$

$$[\mathcal{L}_{W,\nu}\boldsymbol{\sigma}]_{i\ell} = \sum_{j=1}^{D} \frac{\partial \sigma_{i\ell}}{\partial X^{(j)}}\, \sigma_{j\nu} \;=\; \mathbf{0}. \tag{111}$$

Thus, the omission of this remainder term does not contribute any bias or variance to the EuM-based drift estimator.

I. Additional numerical results

I.1. Inference for a 3-dimensional system

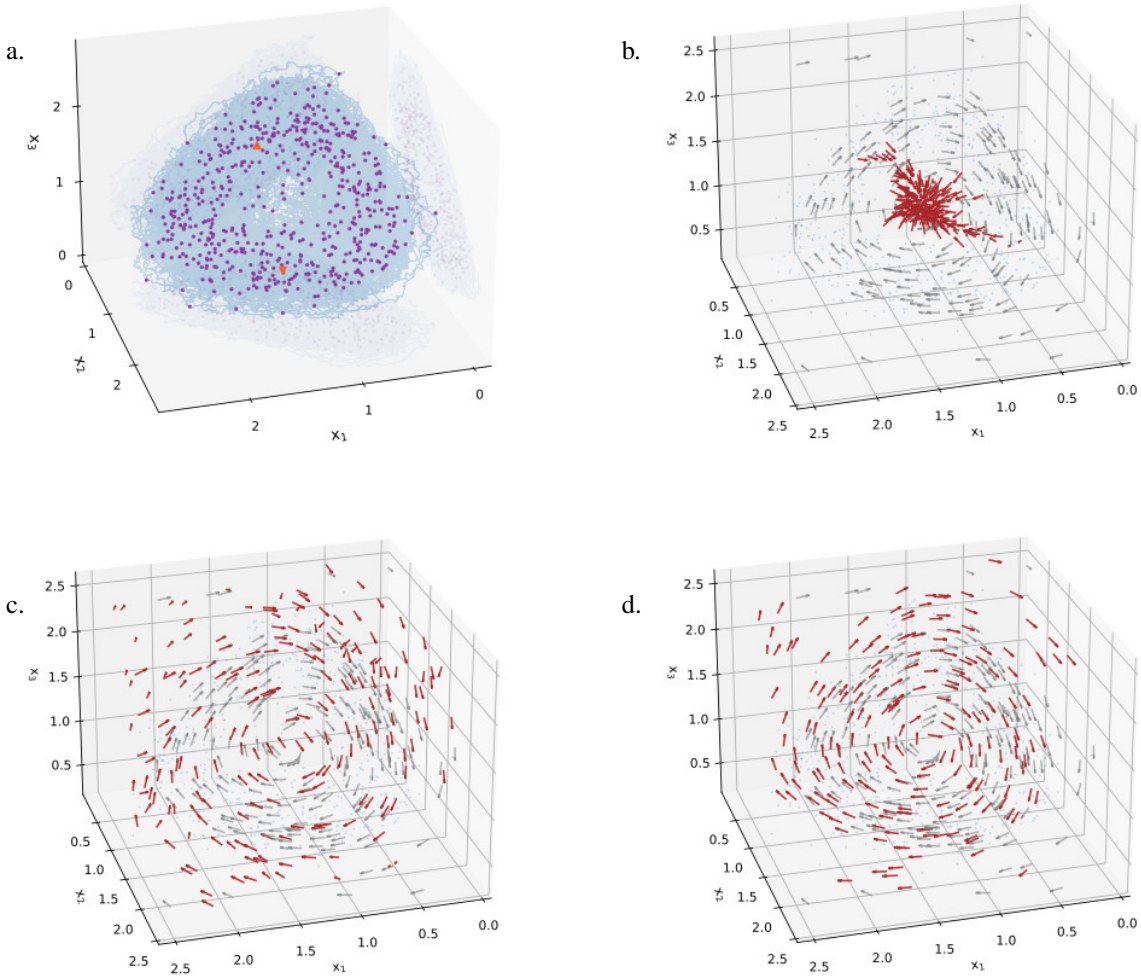

Figure 9

Inference for a three dimensional system following Repressilator dynamics. a.) Empirical manifold of the Repressilator system. Purple dots indicate the observations and blue line denotes the continuous unobserved trajectory. Orange triangles indicate two consecutive observations. **b.)** Ground truth (*grey*) and estimated (*maroon*) flow field with Gaussian process regression without augmentation **c.-d.)** Ground truth (*grey*) and estimated (*maroon*) flow field with the proposed framework after c.) one and d.) two augmentations.

I.2. Inference for 8-dimensional system

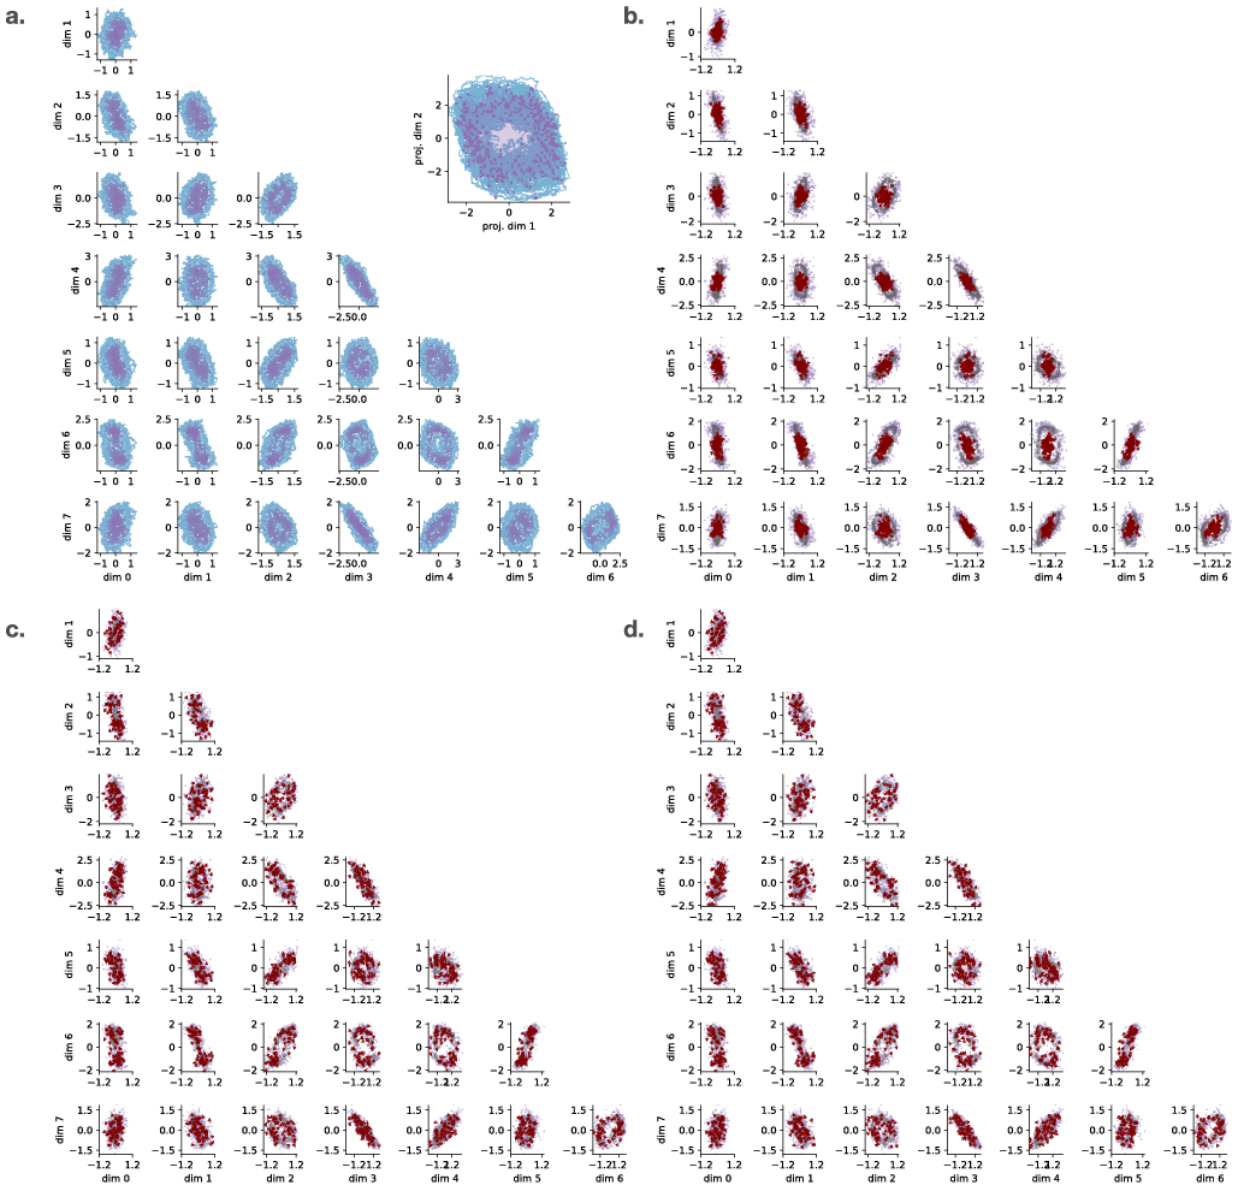

Figure 10

Example inference of an 8-dimensional system. **a.)** Observations (*purple*) and unobserved path (*light blue*) across all dimensions, **b.)** Inferred drift (*maroon*) with Gaussian process regression without augmentation against ground truth drift (*grey*). The vector fields are plotted with streamlines initiated from 60 random trajectories. Identified vector field after **c.)** one and **d.)** two augmentations.

I.3. Inference with noise miss-estimation

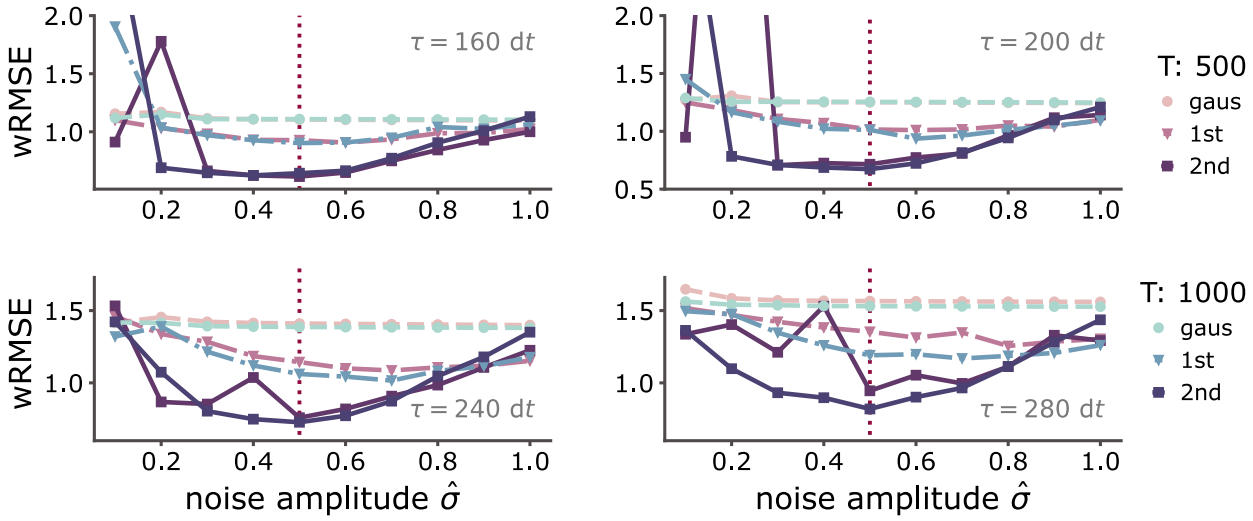

Figure 11

Small noise misestimation has small impact on estimation accuracy. Weighted root mean square error (wRMSE) vs. noise amplitude σ employed in the augmentation for different inter-observation intervals with **a.)** $\tau = 160\,dt$ **b.)** $\tau = 200\,dt$, **c.)** $\tau = 240\,dt$ **d.)** $\tau = 280\,dt$. Pink-purple lines correspond to estimation with total simulation length $T = 500$ time units, and blue markers correspond to total simulation length of $T = 1000$ time units. Red dotted line identifies the noise amplitude employed in the simulation of the observations.

I.4. Convergence of inference over iterations

In the main text we have reported performance results of the proposed approach for inferring the underlying stochastic dynamics after two augmentations. This was motivated by the observation that already after two augmentations the method outperformed existing approaches, and was a trade-off between compute requirements and performance. Here we provide demonstration that our approach converges both in terms of wRMSE and negative log likelihood over multiple augmentations for a Van der Pol system.

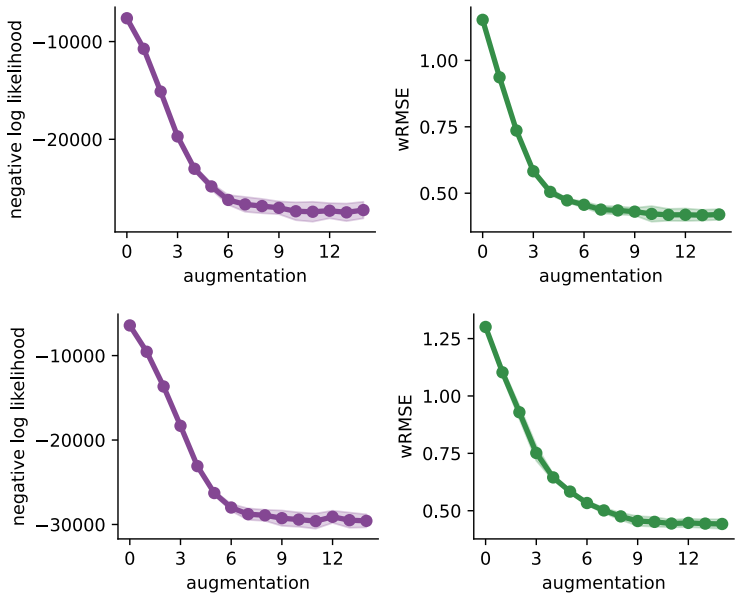

Figure 12
Convergence of the E.M. optimisation over augmentations both in terms of the negative log likelihood and weighted root mean squared error for inter-observation interval $\tau = 240\,\mathrm{d}t$ (*upper*) and $\tau = 280\,\mathrm{d}t$ (*lower*) for noise amplitude $\sigma = 0.25$ and total trajectory length $T = 500$ time units. Points indicate average over 5 (*upper*) and 3 (*lower*) independent runs.

	Van der Pol						
wRMSE ↓	total duration T	$\tau = 80 \times \mathrm{d}t$	$\tau = 120 \times \mathrm{d}t$	$\tau = 160 \times \mathrm{d}t$	$\tau = 200 \times \mathrm{d}t$	$\tau = 240 \times \mathrm{d}t$	$\tau = 280 \times \mathrm{d}t$
$\boldsymbol{\sigma = 0.25}$							
$\mathrm{MFM}_{\mathrm{RBF}}$	1500	1.462 ± 0.007	1.469 ± 0.005	1.470 ± 0.012	1.469 ± 0.008	1.469 ± 0.006	1.466 ± 0.008
$\mathrm{MFM}_{\mathrm{LAND}}$	1500	1.463 ± 0.007	1.469 ± 0.005	1.469 ± 0.012	1.469 ± 0.008	1.469 ± 0.006	1.467 ± 0.008
Geometric$_{\mathrm{RBF}}$ (our)	500	0.419 ± 0.052	0.458 ± 0.063	0.493 ± 0.031	0.517 ± 0.022	0.657 ± 0.040	1.001 ± 0.077
Geometric$_{\mathrm{LAND}}$ (our)	500	0.474 ± 0.034	$\mathbf{0.413 \pm 0.016}$	$\mathbf{0.514 \pm 0.068}$	$\mathbf{0.578 \pm 0.022}$	0.687 ± 0.032	0.993 ± 0.037
$\boldsymbol{\sigma = 0.50}$							
$\mathrm{MFM}_{\mathrm{RBF}}$	1500	1.516 ± 0.011	1.525 ± 0.006	1.538 ± 0.009	1.537 ± 0.017	1.528 ± 0.015	1.544 ± 0.019
$\mathrm{MFM}_{\mathrm{LAND}}$	1500	1.517 ± 0.011	1.526 ± 0.006	1.536 ± 0.009	1.537 ± 0.017	1.528 ± 0.015	1.545 ± 0.019
Geometric$_{\mathrm{RBF}}$ (our)	500	0.653 ± 0.014	0.690 ± 0.026	0.694 ± 0.026	0.761 ± 0.050	0.798 ± 0.047	0.933 ± 0.160
Geometric$_{\mathrm{LAND}}$ (our)	500	$\mathbf{0.462 \pm 0.019}$	$\mathbf{0.541 \pm 0.023}$	$\mathbf{0.621 \pm 0.012}$	$\mathbf{0.675 \pm 0.030}$	$\mathbf{0.750 \pm 0.038}$	$\mathbf{0.865 \pm 0.057}$

Table 3

Performance comparison in terms of weighted root mean square error (wRMSE) of proposed framework and Metric flow matching framework (Kapusniak et al., 2024) employing the LAND (Arvanitidis et al., 2019) and RBF (Arvanitidis et al., 2017) metric approximation algorithm for different noise conditions σ and inter-observation intervals τ for the Van der Pol system.

I.5. Ablations with respect to metric learning algorithm

For approximating the metric induced by the observations, we employed the framework of (Arvanitidis et al., 2019), while we could have employed alternative metric learning approaches (Scarvelis and Solomon, 2023; Hauberg et al., 2012; Barua et al., 2025; Gruffaz and Sassen, 2025). However, the framework of (Arvanitidis et al., 2019) perfectly fitted the purposes of our work, because it employs a non-parametric (kernel) estimation for approximating the metric and computes the geodesics through GP regression. This allowed us to evaluate the geodesic equation at different increments, that is necessary for imposing the time dependent geodesic constraint. A similar metric approximation has been recently employed in (Kapusniak et al., 2024) for metric flow matching, i.e., for augmentation that respects the geometry of the dataset. While our approach has a similar flavour to this work, our framework additionally requires the augmented data to be temporary ordered and to respect the stochastic flow of the estimated system. This results in learning a global drift that approximates the underlying stochastic dynamics, instead of learning a local drift that transports a snapshot of states from some initial to a final configuration.

To probe the robustness of our framework, when we employ a different approach to estimate the metric, following (Kapusniak et al., 2024) we tested our method when we employ a radial based function approximation to estimate the diagonal metric, similar to (Arvanitidis et al., 2017). In the table 3 we report the performance of our method when we employ the locally adaptive normal distribution framework (Geometric (our)) for approximation the metric (Arvanitidis et al., 2019) and when we employ the radial basis function variant of the metric approximation (Geometric$_{\mathrm{RBF}}$ (our)) for the Van der Pol system for different inter-observation intervals and noise conditions. We observe that the resulting drift is accurate also with the RBF method for estimating the metric, yet the method proposed in the main text performs slight better across all inter-observation intervals and noise conditions.

I.6. Comparison of geodesic curves with Onsager–Machlup minimizers

The **Onsager–Machlup** (OM) functional, we introduced in Section C, provides a variational characterisation of the **most probable path** (MPP) between two system states over a prescribed time interval. In our framework, the geodesic curves, we employed to guide bridge augmentation, are meant to serve as coarse empirical approximations of the OM-minimisers, i.e. of the most probable paths between successive observations. A natural question, therefore, is how accurately these inferred geodesics capture the corresponding OM-minimisers. To address this question, we compare the identified geodesic curves to the ground-truth OM-minimisers, and compare their **distance** and **action gap** to the true OM-minimisers with that of a simple straight interpolant between consecutive observations. We perform these comparisons for different inter-observation intervals $\tau \in \{200, 240, 280\} \times \mathrm{d}t$, total trajectory lengths $T \in \{200, 300, 400, 500\}$ time units, values of the hyper-parameter $\sigma_{\mathcal{M}} \in \{0.1, 0.2, 0.3, 0.4, 0.5, 1.0\}$ controlling the width of the kernel employed in the metric approximation, and for different levels of dynamical noise $\sigma \in \{0.1, 0.25, 0.50\}$.

More precisely, among all sufficiently smooth curves $\mathbf{x} : [0, \tau] \to \mathbb{R}^d$ with boundary conditions

$$\mathbf{x}(0) = \mathbf{x}_a, \qquad \mathbf{x}(\tau) = \mathbf{x}_b, \tag{112}$$

the OM-minimiser is the path that minimises the **OM-action**

$$\mathcal{S}_{\mathrm{OM}}[\mathbf{x}] = \int_0^\tau \left[\frac{1}{4\,D} \|\dot{\mathbf{x}}(t) - \mathbf{f}(\mathbf{x}(t))\|_2^2 + \frac{1}{2} \nabla \cdot \mathbf{f}(\mathbf{x}(t)) \right] \mathrm{d}t. \tag{113}$$

Thereby, the OM-minimiser represents the most probable continuous-time path connecting two consecutive observations under the ground-truth stochastic dynamics for fixed inter-observation interval τ.

In the framework presented in the main text, we employ geodesic curves on the empirical manifold as guidance for path augmentation. Since these curves are not, in general, identical to OM-minimisers, it is natural to quantify how well they approximate the corresponding most probable paths, and contrast its performance to a simple interpolant (straight curve) that joins the considered observations. To that end, for each pair of observations, we compare the identified geodesic curve $\boldsymbol{\gamma}_{\mathrm{geo}}$ with the OM-minimiser $\mathbf{x}^\star_{\mathrm{OM}}$ by measuring both a path-space distance and an action-based discrepancy, and compare these two discrepancies with the discrepancies of a straight interpolant to the OM-minimiser.

We characterise the path-space discrepancy between the two curves using their time-averaged L^2 path distance

$$d_{\mathrm{path}} = \left(\frac{1}{\tau} \int_0^\tau \|\boldsymbol{\gamma}_{\mathrm{geo}}(t) - \mathbf{x}^\star_{\mathrm{OM}}(t)\|_2^2 \, \mathrm{d}t \right)^{1/2}. \tag{114}$$

This quantifies how close the two curves remain in state space over the inter-observation interval.

Second, we compare the curves in terms of their OM action. Since the OM minimiser is defined as the minimiser of Eq.113, the **action gap**, i.e. the difference between the action evaluated for the most probable path and for the geodesic curve

$$\Delta\mathcal{S}_{\mathrm{geo}} = \mathcal{S}_{\mathrm{OM}}[\boldsymbol{\gamma}_{\mathrm{geo}}] - \mathcal{S}_{\mathrm{OM}}[\mathbf{x}^\star_{\mathrm{OM}}]. \tag{115}$$

We expect $\Delta\mathcal{S} \geq 0$, and thus values close to zero indicate that the geodesic curve is nearly optimal in the OM sense.

Since the absolute value of the action can vary across settings, to compare how the discrepancy between the geodesic and the OM-minimiser changes for different values of $\sigma_{\mathcal{M}}$, total trajectory length T, and inter-observation interval τ, we consider the **relative action gap**

$$\Delta\mathcal{S}^{\mathrm{rel}}_{\mathrm{geo}} = \frac{\mathcal{S}_{\mathrm{OM}}[\boldsymbol{\gamma}_{\mathrm{geo}}] - \mathcal{S}_{\mathrm{OM}}[\mathbf{x}^\star_{\mathrm{OM}}]}{\left|\mathcal{S}_{\mathrm{OM}}[\mathbf{x}^\star_{\mathrm{OM}}]\right| + \varepsilon}, \tag{116}$$

where $\varepsilon = 10^{-12}$. This normalized quantity facilitates comparison across different parameter regimes.

We quantified the aforementioned discrepancies across different values of the metric-approximation parameter $\sigma_{\mathcal{M}}$, different total trajectory lengths, and different inter-observation intervals τ. For each parameter setting we used 3 independent runs, and within each run we evaluate the above quantities on 20 different inter-observation intervals. Thus, for each configuration we obtain empirical summaries of the path distance Eq. 114, the action gap Eq. 115, and the relative action gap Eq. 116 over a collection of independently generated bridge problems.

Discrete OM action. The OM minimiser is computed on a discretized time grid, that we denote here with

$$0 = t_0 < t_1 < \cdots < t_M = \tau. \tag{117}$$

$X_k \approx x(t_k)$ identifies a discrete path with fixed endpoints

$$\mathbf{X}_0 = \mathbf{x}_a, \qquad \mathbf{X}_M = \mathbf{x}_b. \tag{118}$$

Using midpoint discretisation, we approximate the continuous OM functional by

$$\mathcal{S}^{(M)}_{\mathrm{OM}}(\mathbf{X}_{0:M}) = \sum_{k=0}^{M-1} \left[\frac{1}{4D} \left\| \frac{\mathbf{X}_{k+1} - \mathbf{X}_k}{t_{k+1} - t_k} - f\left(\frac{\mathbf{X}_{k+1} + \mathbf{X}_k}{2} \right) \right\|^2 + \frac{1}{2} \nabla \cdot f\left(\frac{\mathbf{X}_{k+1} + \mathbf{X}_k}{2} \right) \right] (t_{k+1} - t_k). \tag{119}$$

We define then discrete OM-minimiser as

$$X^\star_{0:M} = \arg \min_{\mathbf{X}_1, \ldots, \mathbf{X}_{M-1}} \mathcal{S}^{(M)}_{\mathrm{OM}}(\mathbf{X}_{0:M}), \tag{120}$$

and optimise only the interior points, while keeping the endpoints fixed.

Numerical computation of the OM-minimiser. The continuous OM problem is a two-point boundary-value variational problem: we want to identify the path connecting the prescribed endpoints that minimises the OM-action equation 113. Rather than solving the Euler–Lagrange boundary-value equations directly, we solve the equivalent discrete optimisation problem equation 120.

To improve robustness, especially for longer intervals or smaller noise levels, where the action landscape can become stiff, we resorted to a multi-scale optimisation. More precisely, we start the optimisation on a coarse temporal grid obtained by subsampling the finest target grid, solve the discrete OM minimisation problem on that coarse grid, and then use the resulting path as the initialisation for the next finer discretisation. This refinement procedure is repeated until the target finest grid is reached. Thus, the optimisation proceeds through a sequence of nested discretisations (here we employed 3 nested discretisation).

Computation of action gap and relative action gap. We denote the geodesic curve on the discretised grid by $X^{\mathrm{geo}}_{0:M}$. We evaluate both the geodesic path and the identified OM-minimising path using the *same* discrete OM functional equation 119. The discrete action gap is therefore

$$\Delta\mathcal{S}^{(M)} = \mathcal{S}^{(M)}_{\mathrm{OM}}(\mathbf{X}^{\mathrm{geo}}_{0:M}) - \mathcal{S}^{(M)}_{\mathrm{OM}}(\mathbf{X}^{\star}_{0:M}), \tag{121}$$

and the corresponding discrete relative action gap is

$$\Delta\mathcal{S}^{(M)}_{\mathrm{rel}} = \frac{\mathcal{S}^{(M)}_{\mathrm{OM}}(\mathbf{X}^{\mathrm{geo}}_{0:M}) - \mathcal{S}^{(M)}_{\mathrm{OM}}(\mathbf{X}^{\star}_{0:M})}{\left|\mathcal{S}^{(M)}_{\mathrm{OM}}(\mathbf{X}^{\star}_{0:M})\right| + \varepsilon}. \tag{122}$$

Likewise, we compute the discrete path distance on the same grid as

$$d^{(M)}_{\mathrm{path}} = \left(\frac{1}{M+1}\sum_{k=0}^{M}\|\mathbf{X}^{\mathrm{geo}}_k - X^{\star}_k\|^2\right)^{1/2}. \tag{123}$$

Together, these quantities allow us to assess not only whether the geodesic curves remain geometrically close to associated OM-minimisers, but also whether they are nearly optimal in the sense of minimising the OM-action. The former addresses similarity in state space, while the latter directly measures how well the geodesic approximates the most probable path under the assumed stochastic dynamics.

Example OM-minimisers, geodesics curves and linear interpolants are illustrated in Fig. 14.

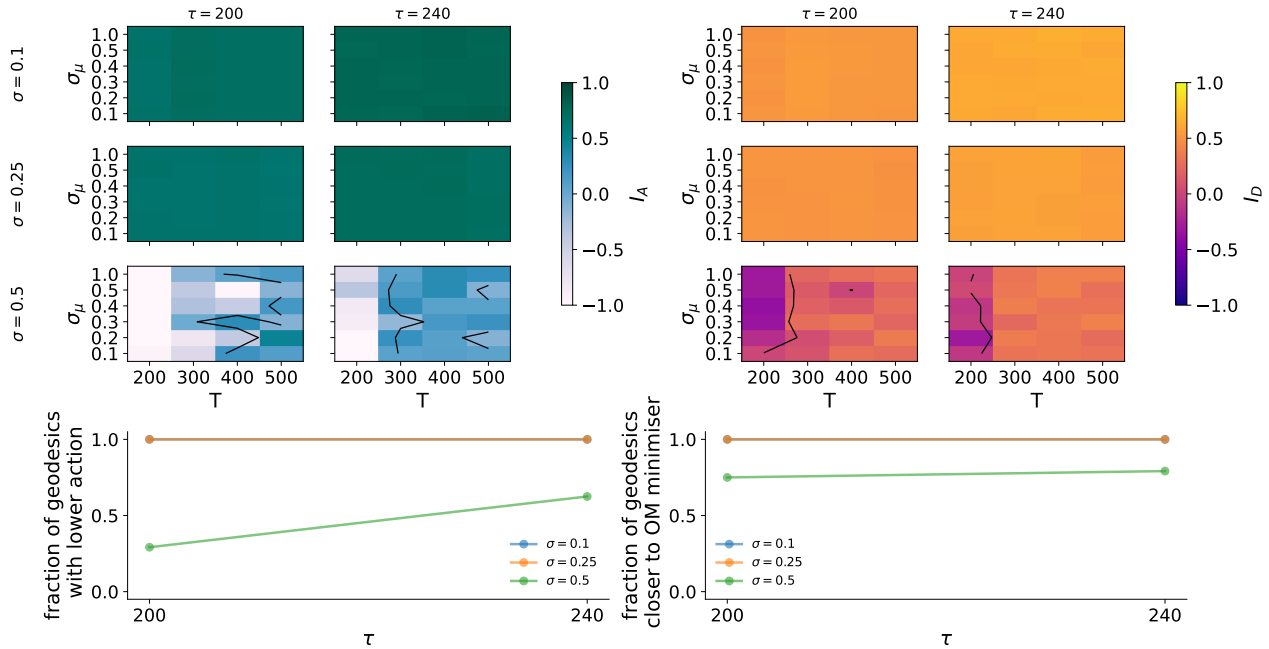

Figure 13

Relative performance of geodesic curves in approximating the OM-minimisers against linear interpolants. ($\mathrm{d}t = 0.01$) **a.)** improvement in OM action gap, quantified by $I_A = 1 - \frac{\Delta A_{\text{geo}}}{\Delta A_{\text{lin}}}$, $\Delta A_{\text{geo}} = A[\gamma_{\text{geo}}] - A[\gamma_{\text{OM}}]$, $\Delta A_{\text{lin}} = A[\gamma_{\text{lin}}] - A[\gamma_{\text{OM}}]$. **b.)** Improvement in path distance to the OM-minimiser, quantified by $I_D = 1 - \frac{d(\gamma_{\text{geo}},\gamma_{\text{OM}})}{d(\gamma_{\text{lin}},\gamma_{\text{OM}})}$. Positive values indicate parameter settings in which the geodesic curve provides an improved approximation to the OM-minimiser over the linear interpolant, $I = 0$ indicates equal performance, and negative values indicate settings in which the linear interpolant is closer to the OM minimiser. Results computed across trajectory length T, metric bandwidth σ_μ, dynamical noise σ, and inter-observation interval τ, whether the geodesic reference curve is closer to the OM minimiser than a straight-line interpolant between the same endpoints. The black contour marks the transition $I = 0$. **c.), d.)** Summary plots indicating for each τ and σ, the fraction of tested (T, σ_μ) settings in which the geodesic outperforms the linear interpolant. **c.)** Fraction of geodesics with smaller OM action gap, $\Delta A_{\text{geo}} < \Delta A_{\text{lin}}$. **d.)** Fraction of geodesics with smaller path distance to the OM minimiser, $d(\gamma_{\text{geo}}, \gamma_{\text{OM}}) < d(\gamma_{\text{lin}}, \gamma_{\text{OM}})$. Values closer to one indicate that the geodesic curves consistently provide a better approximation to the OM minimiser across metric bandwidths and trajectory lengths, while values closer to zero indicate that the straight-line interpolant provides a better approximation.

J. Details on numerical experiments

We simulated a two dimensional Van der Pol oscillator with drift function

$$f_1(x, y) = \mu(x - \frac{1}{3}x^3 - y) \tag{124}$$

$$f_2(x, y) = \frac{1}{\mu}x, \tag{125}$$

starting from initial condition $x0 = [1.81, -1.41]$ and under noise amplitudes $\sigma = \{0.25, 0.50, 0.75, 1.00\}$ for total duration of $T = \{500, 1000\}$ time units. The employed inter-observation intervals $\tau = \{80, 120, 160, 200, 240, 280, 320\} \times \mathrm{d}t$. The last inter-observation interval exceeds the half period of the oscillator and thus samples only a single state per period. This resulted in erroneous estimates. In this setting this indicates the upper limit of τ for which we can provide estimates. However for any inference method, if the observation process samples only one observation per period, identifying the underlying force field without additional assumptions is not possible with temporal methods. The discretisation time-step used for simulation of the ground truth dynamics, and path augmentation $\delta t = 0.01$. For sampling the controlled bridges we employed $N = 100$ particles evolving the associated ordinary differential equation as described in (Maoutsa and Opper, 2022). The logarithmic gradient estimator used $M = 40$ inducing points. The sparse Gaussian process for estimating the drift was based on a sparse kernel approximation of $S = 300$ points. In the presented simulation we have employed a weighting parameter $\beta = 0.5$ (Eq. 35). This provides a moderate pull towards the invariant density. The example in

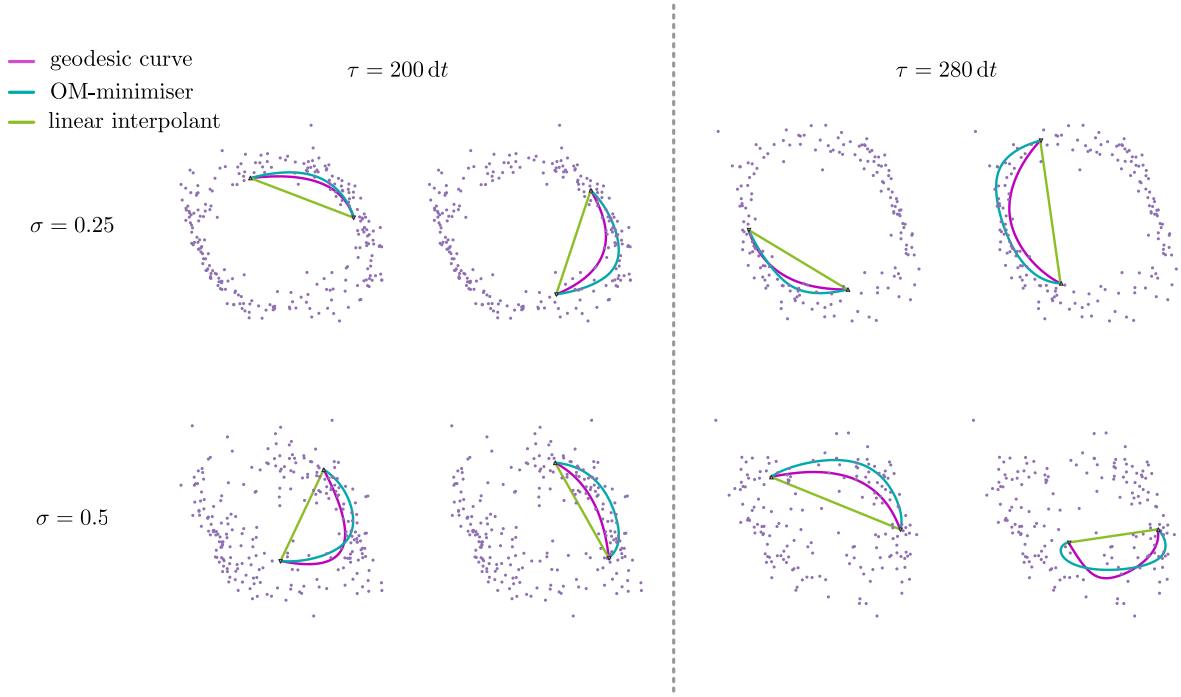

Figure 14

Example Onsager-Machlup most probable paths, identified geodesics, and linear interpolants for the Van der Pol system. Identified geodesics, OM-minimisers under the ground truth dynamics, and linear interpolants for two example inter-observation intervals for $\sigma = \{0.25, 0.50\}$ and inter-observation interval $\tau = 200\,\mathrm{d}t$ and $\tau = 280\,\mathrm{d}t$. ($\mathrm{d}t = 0.01$)

Figure 2 was constructed with $\beta = 1$ and provides a better approximation of the transition density, than $\beta = 0.5$.

For the **out-of-equilibrium process** with harmonic trapping and circulation and a Gaussian repulsive obstacle in the centre we followed the description presented in (Frishman and Ronceray, 2020) following the drift

$$f_\mu(\mathbf{x}) = -\Omega_{\mu\nu} x_\nu + \alpha e^{-x^2/2\sigma^2} x_\mu \quad \text{with} \quad \Omega = \begin{pmatrix} 2 & 2 \\ -2 & 2 \end{pmatrix}, \tag{126}$$

for $\alpha = 10$ and simulated the stochastic system with noise amplitude $\sigma = 0.5$ on a time grid of $\mathrm{d}t = 0.01$ steps, observed at inter-observation intervals $\tau = \{150, 200, 250\} \times \mathrm{d}t$ and for total duration $T = 1000$ time units.

For the **Hopf system** we used the drift

$$f_1(x_1, x_2) = z_2, \tag{127}$$
$$f_2(x_1, x_2) = -z_1 + \left(\mu - z_1^2\right) z_2, \tag{128}$$

with $\mu = 0.35$ and integrated the system with noise amplitude $\sigma = 0.15$ on a timegrid with $\mathrm{d}t = 0.01$ resolution, observed at $\tau = \{200, 300, 400\} \times \mathrm{d}t$ time intervals. This is the normal form of the Hopf bifurcation.

For the **Selkov glycolysis model** (Selkov, 1968) we employed the drift

$$f_1(x_1, x_2) = -x_1 + \alpha x_2 + x_1^2 x_2, \tag{129}$$
$$f_2(x_1, x_2) = 0.6 - \alpha x_2 - x_1^2 x_2, \tag{130}$$

with $a = 0.06$ and noise amplitude $\sigma = 0.05$ for inter-observation intervals $\tau = \{100, 200\} \times \mathrm{d}t$ and simulation time grid of $dt = 0.01$ spacing and for total duration $T = 1000$ time units.

This model is a minimal two-variable model of glycolytic oscillations, first introduced in (Selkov, 1968). It describes the autocatalytic feedback processes in the glycolysis pathway, focusing on how simple nonlinear interactions can give

rise to oscillatory dynamics in concentrations of intermediates. The first state variable x_1 represents the concentration of adenosine diphosphate, while x_2 represents the concentration of a glycolytic intermediate.

For the **Repressilator dynamics**, we considered the system introduced by Elowitz and Leibler. The repressilator consists of three genes arranged in a cyclic inhibitory loop, where each gene represses the expression of the next one in the cycle.

The deterministic dynamics of the system are

$$\dot{x}_1 = \frac{\beta}{1 + \left(\frac{x_3}{k}\right)^n} - \gamma x_1, \tag{131}$$

$$\dot{x}_2 = \frac{\beta}{1 + \left(\frac{x_1}{k}\right)^n} - \gamma x_2, \tag{132}$$

$$\dot{x}_3 = \frac{\beta}{1 + \left(\frac{x_2}{k}\right)^n} - \gamma x_3, \tag{133}$$

where β controls the maximal production rate, γ is the degradation rate, k is the repression threshold, and n is the Hill coefficient controlling the nonlinearity of repression.

In our simulations we used $\beta = 3.0$, $\gamma = 1.0$, $k = 1.0$, and $n = 4.0$. We simulated the system with $\mathrm{d}t = 0.01$ for $T = 1000$ time units, and observed the trajecotry every $\tau = 170 \times \mathrm{d}t$ time units. We used $M = 120$ inducing points for the logarithmic gradient estimation, $N = 400$ particles for sampling the flows, and employed noise amplitude $\sigma = 0.2$

J.1. On computation of geodesic curves

For the computation of geodesic curves we followed the framework introduced in (Arvanitidis et al., 2019). The geodesic equation relies on a non-parametric estimation of the Riemannian metric, which is constructed using kernel-weighted local diagonal covariances, and has computational complexity $\mathcal{O}(ND)$, where D is the dimensionality of the problem and N denotes the number of samples. The computational cost of solving the geodesic equation scales sublinearly with increasing dimensionality.

J.2. Details on baseline methods

We compared the performance of our method against a series of competing methods for inference of stochastic dynamics. In particular, we compared our method against methods specifically designed for inference of stochastic systems from single trajectories, and against systems that infer population dynamics.

We employed the following methods that assume single trajectories for drift inference:

1. Gaussian process regression without state estimation (**GP**)
2. path augmentation with Ornstein-Uhlenbeck dynamics with Gaussian process inference (**OU**) (Batz et al., 2018)
3. sparse variational inference with state estimation (**SVISE**) (Course and Nair, 2023a)
4. basis function approximation of Kramers-Moyal coefficients, i.e. the drift function (**KM-basis**) (Nabeel et al., 2025)
5. latent SDE inference with amortized reparameterization with (**LatentSDE+GP-pre**) and without pre-training (**LatentSDE**) (Course and Nair, 2023b).

We further compared our method with recent Schroedinger bridge generating frameworks that primary aim to infer population dynamics from snapshot data. In particular we considered the following frameworks:

I. Metric Flow Matching **(MFM)** (Kapusniak et al., 2024)

II. Generalized Schrödinger Bridge Matching **(GSBM)** (Liu et al., 2024)

III. Simulation-free Schrödinger bridges via score and flow matching (**[SF]^2 M**) (Tong et al., 2024)

For these methods, we clustered the observations of each system into *disjoint* subsets of adjacent points. We employed the k-Nearest neighbours algorithm (Fix, 1985; Cover and Hart, 1967) to construct the clusters as local neighbourhoods on the state space, comprising each at most 64 and minimum 20 observations. We paired each cluster $\mathcal{J}_b$ with the set $\mathcal{J}_b^+ \doteq \{\, \boldsymbol{\mathcal{O}}_{k+1} \,:\, \boldsymbol{\mathcal{O}}_k \in \mathcal{J}_b \,\}$ of the next observation of each cluster member. We then considered each cluster pair $(\mathcal{J}_b, \mathcal{J}_b^+)$ as the initial and terminal condition for a Schrödinger bridge problem, i.e.

$$\pi_0^b \doteq \{\, \boldsymbol{\mathcal{O}}_k : \boldsymbol{\mathcal{O}}_k \in \mathcal{J}_b \,\} \tag{134}$$

$$\pi_1^b \doteq \{\, \boldsymbol{\mathcal{O}}_\ell : \boldsymbol{\mathcal{O}}_\ell \in \mathcal{J}_b^+ \,\}. \tag{135}$$

These serve as samples of the densities required as boundaries conditions for the Schrödinger bridges.

For the multi-marginal setting, starting from the cluster that contained the observation $\boldsymbol{\mathcal{O}}_1$ and subsequently created a sequence of cluster following the time ordering of the observations, i.e.

$$\pi_i^0 = \{\, \boldsymbol{\mathcal{O}}_{k+i} : k \in \mathcal{J}_0 \,\}. \tag{136}$$

We then employed a sequence of 50 marginal densities $\{\pi_i^0\}_{i=0}^{49}$ as snapshot observations required by the framework.

Metric Flow Matching. For the Metric Flow Matching framework, we trained on observations resulting from total simulation length $T_{\text{MFM}} = 3\,T = 1500$ (time units) to ensure sufficient data for each bridge. For each constructed bridge indexed by b, the flow network trained with the flow matching objective represents the velocity of the samples $\mathbf{u}_b(\mathbf{x}, t)$ transferred within the normalised time $t \in [0, 1]$ from the initial boundary condition to the terminal one. We approximate a time-independent local drift $\hat{\mathbf{f}}_b(\mathbf{x})$ by rescaling the velocity field $\mathbf{u}_b(\mathbf{x}, t)$ with the inter-observation interval τ, i.e.,

$$\hat{\mathbf{f}}_b(\mathbf{x}) \;=\; \frac{1}{\tau}\,\mathbf{u}_b(\mathbf{x}, t). \tag{137}$$

To obtain a global drift estimate from the individual local estimates, we compute "responsibilities" or weights of each individual drift for each point $\mathbf{x}_m$ of a pre-defined two-dimensional evaluation grid that covers the state space region occupied by the observations. These weights indicate how relevant each bridge b was for estimating the drift at each grid point $\mathbf{x}_m$. For each bridge, we compute support weights $\omega_b(\mathbf{x})$ on the grid employing kernel density estimation (KDE) over the bridge boundary condition samples. Then, for each grid point $\mathbf{x}_m$, we compute bridge responsibilities as

$$\rho_b(\mathbf{x}_m) \;=\; \frac{\omega_b(\mathbf{x}_m)}{\sum_{j=1}^{B} \omega_j(\mathbf{x}_m)}\,, \qquad \sum_{b=1}^{B} \rho_i(\mathbf{x}_m) = 1. \tag{138}$$

We estimate the global drift at each grid point by weighting the local estimated drifts with the corresponding bridge responsibility, i.e.,

$$\hat{\mathbf{f}}(\mathbf{x}_m) \;=\; \sum_{b=1}^{B} \rho_b(\mathbf{x}_m)\,\hat{\mathbf{f}}_b(\mathbf{x}_m). \tag{139}$$

K. Computational complexity of the proposed method

We provide here a breakdown of the computational cost of the main components of the proposed framework. Let K denote the number of sparse observations, d the state-space dimension, I_{EM} the number of E.M. iterations, L the number of temporal discretisation points used within each inter-observation interval, N_p the number of particles used for bridge augmentation, and M the number of inducing points used in the sparse Gaussian-process drift estimator. The total number of augmented state samples generated in one E.M. iteration is

$$N_{\text{aug}} \approx (K-1)LN_p.$$

The computational cost separates into four main components: metric estimation, geodesic construction, geometry-guided bridge augmentation, and drift estimation.

Metric estimation. The Riemannian metric is estimated once from the observed states. In the diagonal local-covariance approximation used in the experiments, evaluating the metric at N_q query locations by direct summation over all observations costs

$$\mathcal{O}(N_q K d),$$

with memory cost $\mathcal{O}(N_q d)$ for storing the diagonal metric values. This term is not repeated at every E.M. iteration unless the metric is re-estimated. In the present implementation, the metric is computed from the observed empirical density and then kept fixed. This step is therefore not the dominant cost in the reported experiments.

Geodesic construction. For each pair of consecutive observations (O_k, O_{k+1}), we compute a geodesic with respect to the learned metric. If each geodesic is represented by L_γ discretisation points and the boundary-value solver requires N_γ optimisation or solver iterations, the cost is approximately

$$\mathcal{O}((K-1) N_\gamma L_\gamma C_H(d)),$$

where $C_H(d)$ denotes the cost of evaluating the metric and its derivatives along the curve. For the diagonal metric used here, this cost is substantially cheaper than for a full dense metric, for which metric inversion and Christoffel-symbol computations may introduce additional polynomial dependence on d. Geodesic construction is also performed independently for each inter-observation interval and is therefore parallelisable across the $K-1$ intervals. Since the geodesics are computed before the E.M. iterations and then reused, this component is typically smaller than the repeated bridge-sampling cost.

Geometry-guided bridge augmentation. The main computational bottleneck is the geometry-guided bridge augmentation step. At each E.M. iteration, and for each inter-observation interval, we solve a controlled bridge problem using a particle representation. The leading cost scales with the number of intervals, time-discretisation points, and particles:

$$\mathcal{O}(I_{\mathrm{EM}}(K-1) L N_p C_{\mathrm{step}}),$$

where C_{step} is the cost of one particle propagation, weight update, and control evaluation. In the simplest implementation, C_{step} includes evaluating the current drift estimate, applying the observation and geodesic potentials, and performing particle-filtering operations such as weighting and resampling. If the control or density-gradient terms are estimated using interactions among particles, the worst-case per-time-step cost may increase from $\mathcal{O}(N_p)$ to $\mathcal{O}(N_p^2)$. Thus, depending on the implementation, the bridge step scales between

$$\mathcal{O}(I_{\mathrm{EM}}(K-1) L N_p C_f(d)) \quad \text{and} \quad \mathcal{O}\big(I_{\mathrm{EM}}(K-1) L N_p^2\big),$$

where $C_f(d)$ is the cost of evaluating the current drift model at one state. This is the dominant computational component of the current implementation, especially as d, L, or N_p increase. The memory requirement is

$$\mathcal{O}((K-1) L N_p d)$$

if all augmented trajectories are stored, but can be reduced by streaming Monte Carlo estimates of the sufficient statistics needed for drift learning.

Sparse Gaussian-process drift estimation. Given the augmented trajectories, the drift is updated using a sparse Gaussian-process approximation with M inducing points. For N_{aug} augmented samples, the leading cost of sparse GP inference scales as

$$\mathcal{O}\big(d\, N_{\mathrm{aug}} M^2 + d\, M^3\big),$$

when the d drift components are inferred independently. The corresponding memory cost is approximately

$$\mathcal{O}\big(N_{\mathrm{aug}} M + M^2\big),$$

up to factors depending on whether all augmented samples are stored explicitly or used only through Monte Carlo estimates of the required integrals. Since $M \ll N_{\mathrm{aug}}$, this sparse approximation avoids the cubic cost $\mathcal{O}(N_{\mathrm{aug}}^3)$ of full Gaussian-process regression.

Overall cost. Combining the above terms, the total runtime is approximately

$$\mathcal{O}(N_q K d) + \mathcal{O}((K-1) N_\gamma L_\gamma C_H(d)) + \mathcal{O}(I_{\mathrm{EM}}(K-1) L N_p C_{\mathrm{step}}) + \mathcal{O}\big(I_{\mathrm{EM}} d\, N_{\mathrm{aug}} M^2 + I_{\mathrm{EM}} d\, M^3\big)\,.$$

In practice, the bridge augmentation term dominates because it is repeated at every E.M. iteration and scales with the number of particles, the temporal resolution of the bridge, and the number of inter-observation intervals. By contrast, metric estimation and geodesic construction are performed once and are parallelisable across query locations or observation pairs. The $K-1$ controlled bridge problems are also independent conditional on the current drift estimate, so the augmentation step is naturally parallelisable across inter-observation intervals.

Component	Leading cost	Main scaling variables
Metric estimation	$\mathcal{O}(N_q K d)$	observations, query points, dimension
Geodesic construction	$\mathcal{O}((K-1) N_\gamma L_\gamma C_H(d))$	intervals, geodesic resolution, metric cost
Bridge augmentation	$\mathcal{O}(I_{\mathrm{EM}}(K-1) L N_p C_{\mathrm{step}})$	E.M. iterations, particles, bridge resolution
Sparse GP drift update	$\mathcal{O}(I_{\mathrm{EM}} d N_{\mathrm{aug}} M^2 + I_{\mathrm{EM}} d M^3)$	augmented samples, inducing points, dimension

Table 4

Runtime scaling of the main components of the geometry-guided augmentation framework. The dominant term in the current implementation is the particle-based bridge augmentation step.

In the present work we prioritised demonstrating the inference method in the sparse-observation setting and establishing the usefulness of geometry-guided augmentation for drift inference. The current implementation is therefore not intended to be a fully optimised scalable solver for high-dimensional systems. More scalable implementations could use amortised bridge samplers, adaptive particle allocation, reduced-rank or local approximations of the control, mini-batched drift updates, or GPU-exploiting implementation. Developing such improvements is an important direction for future work.

L. Algorithmic details

Here we provide the outline algorithm for each constituent component of our work. Algorithm A1 provides the main skeleton of the framework. For the geometric approximation and the construction of the geodesics we defer the readers to (Arvanitidis et al., 2019). Algorithm A2 outlines the solution of the control problem that implements the path augmentation. This part is an adapted version of the main algorithm proposed by (Maoutsa and Opper, 2021). Finally, Algorithm A3 describes the solution of the Gaussian process inference given the path augmentations (bridges) created for each augmentation pair. For the simulation of Fokker-Planck equation solutions we used the deterministic particle framework of (Maoutsa et al., 2020).

Algorithm A1: Skeleton of the proposed framework.

Input: $\boldsymbol{\mathcal{O}} = \{(\mathbf{x}_k, t_k)\}_{k=1}^{K}$: observed states at timepoint t_k

Output: $\hat{\mathbf{f}}$: posterior estimate of the drift function

$\boldsymbol{B}^{(j)}$: augmented paths of latent states (optional)

```
// initialise f̂ with a coarse drift estimate
```

Initialise drift estimate $\hat{\mathbf{f}}^{(0)}$ according to Eq. 18

```
// Approximate Riemannian metric from observations (Eq. 32)
```

$H_{dd} = \text{ApproximateMetric}\big(\{\boldsymbol{\mathcal{O}}_{\boldsymbol{k}}\}_{k=1}^{K}\big)$

```
//
// Construct geodesics between O_k and O_{k+1} under the estimated metric as shortest paths
```

$\boldsymbol{\Gamma}^{(\ell)} = \text{ConstructGeodesics}(\boldsymbol{\mathcal{O}}_k, \boldsymbol{\mathcal{O}}_{k+1}, H_{dd})$

// $\boldsymbol{\Gamma}^{k} = \{\boldsymbol{\gamma}_{t'}^{k}\}_{k=1}^{K}$ geodesic curves between selected observation pairs

for *each iteration* j **do**

```
// augment paths along geodesics using particle flow
```

$\boldsymbol{B}^{(j)} = \text{AugmentPaths}(\{\boldsymbol{\mathcal{O}}\}_{k=1}^{K}, \boldsymbol{\Gamma}^{(j)}, \hat{\mathbf{f}}^{(j-1)})$

// uses the deterministic particle flow / bridge construction (Alg. A2) to sample augmented trajectories with $\hat{\mathbf{f}}^{(j-1)}$

```
// Gaussian process inference of the drift function
```

$\hat{\mathbf{f}}^{(\ell)} = \text{GPDriftInference}(\{\boldsymbol{\mathcal{O}}\}_{k=1}^{K}, \boldsymbol{B}^{(j)})$

// update GP posterior over $\mathbf{f}$ using original and augmented data

end

Algorithm A2: Path augmentation algorithm employing Deterministic Particle Flow control

Input: N, M: scalars, number of particles and number of inducing points
t_k, t_{k+1}, dt: scalars, initial and final timepoints, and discretisation step
$\boldsymbol{\mathcal{O}}_k, \boldsymbol{\mathcal{O}}_{k+1}$: $1 \times d$, $1 \times d$ initial and target state
$\hat{\mathbf{f}}$: current drift estimate
σ: noise amplitude
$\boldsymbol{\gamma}_t$: geodesic curve in functional representation

Output: F: $d \times N \times (t_{k+1} - t_k)/dt$, sample representation of $q_t(x)$

$\ell = \frac{(t_{k+1} - t_k)}{\mathrm{d}t}$ // number of timesteps

// Forward filtering $\rho_t(x)$ (Eq. 38)

$\epsilon = 10^{-3}$

$\mathbf{Z}_{ti=0} = \mathcal{N}(\boldsymbol{\mathcal{O}}_k, \epsilon\, \mathbf{I}_d)$ // initialise particles' positions

$\mathbf{Z}_{ti=1} = \mathbf{Z}_0 + \mathrm{d}t \left(\hat{\mathbf{f}}(\mathbf{Z}_0, t_0) - \frac{1}{2}\sigma^2 \frac{\mathbf{Z}_0 - \boldsymbol{\mathcal{O}}_k}{\epsilon}\right)$ // 1st step is with analytic score

For $ti = 2 : \ell$ // deterministic propagation

$\mathbf{Z}_{ti+1} = \mathbf{Z}_{ti} + dt \left(\hat{\mathbf{f}}(\mathbf{Z}_{ti}, t) - \frac{1}{2}\sigma^2 \nabla \log \rho(\mathbf{Z}_{ti}; \mathbf{Z}_{ti})\right)$

$W = \exp\left(-U(\mathbf{Z}_{ti+1}, t)\, dt\right)$

$T^* = \text{EnsembleTransformParticleFilter}(\mathbf{Z}_{ti+1}, W)$

$\mathbf{Z}_{ti+1} = \mathbf{Z}_{ti+1} \cdot T^*$

// Time-reversed propagation of flow $q_t(\mathbf{x})$

$\mathbf{B}_{ti=\ell} = \mathcal{N}(\boldsymbol{\mathcal{O}}_{k+1}, \epsilon\, \mathbf{I}_d)$ // initialise particles' positions

// 1st step is stochastic

$\mathbf{B}_{ti=\ell-1} = \mathbf{B}_\ell - \mathrm{d}t \left(\hat{\mathbf{f}}(\mathbf{B}_\ell, t_1) + \frac{1}{2}\sigma^2 \nabla \log \rho(\mathbf{B}_\ell; \mathbf{Z}_\ell) - \frac{1}{2}\sigma^2 \frac{\mathbf{B}_\ell - \boldsymbol{\mathcal{O}}_{k+1}}{\epsilon}\right)$

For $ti = \ell - 2 : 0$ // deterministic propagation

$\mathbf{B}_{ti-1} = \mathbf{B}_{ti} - \mathrm{d}t \left(\hat{\mathbf{f}}(\mathbf{B}_{ti}, t) - \frac{1}{2}\sigma^2 \nabla \log \rho(\mathbf{B}_{ti}, \mathbf{Z}_{ti}) + \frac{1}{2}\sigma^2 \nabla \log q(\mathbf{B}_{ti}, \mathbf{B}_{ti})\right)$

// Compute control $u(x\,t)$ and controlled paths $\mathbf{F}_{0:T}$

For $ti = 1 : \ell$

$\mathbf{u}(\mathbf{x}, ti) = \sigma^2 \nabla \log q(\mathbf{x}; \mathbf{B}_{ti}) - \sigma^2 \nabla \log \rho(\mathbf{x}; \mathbf{Z}_{ti})$

$\mathbf{F}_{ti+1} = \mathbf{F}_{ti} + \mathrm{d}t \left(\hat{\mathbf{f}}(\mathbf{F}_{ti}, t) + \mathbf{u}(\mathbf{F}_{ti}, ti) - \frac{1}{2}\sigma^2 \frac{\mathbf{F}_0 - \boldsymbol{\mathcal{O}}_k}{\epsilon}\right)$

With the notation $\nabla \log q(\mathbf{x}; \mathbf{B}_{ti})$ we indicate the score function estimation in a functional form $(\mathbf{x})$ based on the density represented by the particles $\mathbf{B}_{ti}$, while $\nabla \log q(\mathbf{F}_{ti}; \mathbf{B}_{ti})$ indicates the same score function evaluated at locations $\mathbf{F}_{ti}$.

Algorithm A3: Gaussian process drift inference from an augmented path measure (part I)

Input: $\mathcal{Z} = \{\mathbf{z}_i\}_{i=1}^{S}$: inducing points for the sparse GP (Sp)
$\{\mathbf{X}_j(t_\ell)\}_{j=1,\dots,N}^{\ell=1,\dots,T'}$: particle positions from the path measure Q (BALL2)
$\{\mathbf{g}(\mathbf{X}_j(t_\ell), t_\ell)\}$: effective drift evaluated along particles (gbALL2)
$k^{\mathbf{f}}$: kernel with lengthscales ℓ_1, ℓ_2, ℓ_3 (shared across dimensions)
g: diffusion amplitude, $\sigma^2 = g^2$
Δt: time step of the particle simulation
d: state dimension, N: number of particles, T': number of time steps

Output: Approximations $I_1^{(i)}, I_2^{(i)}$ of the integrals over $A(\mathbf{x})$ and $B(\mathbf{x})$

// 0. shorthand and initialisation

Set $S \leftarrow |\mathcal{Z}|$ (number of inducing points)
Initialise $I_1 \in \mathbb{R}^{S\times S\times d}$ and $I_2 \in \mathbb{R}^{S\times d}$ to zero
Initialise $\Lambda \in \mathbb{R}^{S\times S\times d}$ and $\mathbf{d} \in \mathbb{R}^{S\times d}$ to zero

// 1. compute kernel matrices on the inducing points

Construct the inducing–inducing kernel matrix

$$\mathcal{K}_S = k^{\mathbf{f}}(\mathcal{Z}, \mathcal{Z}) \in \mathbb{R}^{S\times S}$$

and compute a regularised inverse

$$\mathcal{K}_S^{-1} = \big(\mathcal{K}_S + \varepsilon I\big)^{-1}, \quad \varepsilon \approx 10^{-3}.$$

Define the kernel map to inducing points

$$k^{\mathbf{f}}(\mathcal{Z}, \mathbf{x}) = \big(k^{\mathbf{f}}(\mathbf{z}_i, \mathbf{x})\big)_{i=1}^{S} \in \mathbb{R}^{S}.$$

// 2. sample-based approximation of $A(\mathbf{x})$ and $B(\mathbf{x})$

for $i = 1, \dots, d$ **do**

// loop over state dimensions

  for $\ell = 1, \dots, T'$ **do**

// loop over time

    Let $\mathbf{X}(t_\ell) \in \mathbb{R}^{d\times N}$ be the particle positions at time t_ℓ
    For each particle position $\mathbf{X}_j(t_\ell)$, compute

$$\mathbf{k}_j = k^{\mathbf{f}}(\mathcal{Z}, \mathbf{X}_j(t_\ell)) \in \mathbb{R}^{S}.$$

    Stack them column-wise to obtain

$$K_\ell = \big[\mathbf{k}_1, \dots, \mathbf{k}_N\big] \in \mathbb{R}^{S\times N}.$$

    Let $g_i(\mathbf{X}_j(t_\ell), t_\ell)$ denote the i-th component of the effective drift at particle j and time t_ℓ

// accumulate Monte Carlo estimates of the integrals

    Update

$$I_1^{(i)} \leftarrow I_1^{(i)} + K_\ell K_\ell^\top, \qquad I_2^{(i)} \leftarrow I_2^{(i)} + K_\ell \mathbf{g}_i(t_\ell),$$

    where $\mathbf{g}_i(t_\ell) = \big(g_i(\mathbf{X}_1(t_\ell), t_\ell), \dots, g_i(\mathbf{X}_N(t_\ell), t_\ell)\big)^\top$.

  end

// normalise by time and number of particles

$$I_1^{(i)} \leftarrow \frac{\Delta t}{N} I_1^{(i)}, \qquad I_2^{(i)} \leftarrow \frac{\Delta t}{N} I_2^{(i)}.$$

end

In this algorithm Here $I_1^{(i)}$ approximates $\int k^{\mathbf{f}}(\mathcal{Z}, \mathbf{x}) A(\mathbf{x}) k^{\mathbf{f}}(\mathbf{x}, \mathcal{Z})\, \mathrm{d}\mathbf{x}$, and $I_2^{(i)}$ approximates $\int k^{\mathbf{f}}(\mathcal{Z}, \mathbf{x}) B_i(\mathbf{x})\, \mathrm{d}\mathbf{x}$.

Algorithm A4: Gaussian process drift inference from an augmented path measure (part II)

Input: Same inputs as Alg. 3
$I_1^{(i)}, I_2^{(i)}$: Monte Carlo approximations from Alg. 3
$\mathcal{K}_S, \mathcal{K}_S^{-1}$: inducing–inducing kernel matrix and its regularised inverse

Output: Component-wise drift estimators $\hat{f}_i(\mathbf{x})$, $i = 1, \dots, d$
Expected negative log data likelihood $\mathcal{L}_{\text{path}}$ under Q_f

// 3. compute Λ and $\mathbf{d}$ for each component

for $i = 1, \dots, d$ **do**

// match Eq. equation 40 with sparse GP parametrisation

$$\Lambda^{(i)} \leftarrow \frac{1}{\sigma^2} \mathcal{K}_S^{-1} I_1^{(i)} \mathcal{K}_S^{-1}, \qquad \mathbf{d}^{(i)} \leftarrow \frac{1}{\sigma^2} \mathcal{K}_S^{-1} I_2^{(i)}.$$

end

This matches the definitions

$$\Lambda = \frac{1}{\sigma^2} \mathcal{K}_S^{-1} \left(\int k^{\mathbf{f}}(\mathcal{Z}, \mathbf{x}) A(\mathbf{x}) k^{\mathbf{f}}(\mathbf{x}, \mathcal{Z}) \mathrm{d}\mathbf{x} \right) \mathcal{K}_S^{-1}, \quad \mathbf{d} = \frac{1}{\sigma^2} \mathcal{K}_S^{-1} \left(\int k^{\mathbf{f}}(\mathcal{Z}, \mathbf{x}) B(\mathbf{x}) \mathrm{d}\mathbf{x} \right).$$

// 4. define the component-wise drift estimators

For each component $i = 1, \dots, d$, define

$$\hat{f}_i(\mathbf{x}) = k^{\mathbf{f}}(\mathbf{x}, \mathcal{Z}) \Big(I + \Lambda^{(i)} \mathcal{K}_S \Big)^{-1} \mathbf{d}^{(i)},$$

so that the full drift estimate is

$$\hat{\mathbf{f}}_S(\mathbf{x}) = \big(\hat{f}_1(\mathbf{x}), \dots, \hat{f}_d(\mathbf{x}) \big)^\top.$$

// 5. compute expected negative log data likelihood under Q_f

Initialise accumulators $S_{\|f\|} \leftarrow 0$, $S_{\nabla \cdot f} \leftarrow 0$, $S_{f \cdot g} \leftarrow 0$

for $\ell = 1, \dots, T'$ **do**

For all particle positions $\mathbf{X}_j(t_\ell)$, evaluate $\hat{\mathbf{f}}_S(\mathbf{X}_j(t_\ell))$

Accumulate

$$S_{\|f\|} \leftarrow S_{\|f\|} + \sum_{j=1}^{N} \|\hat{\mathbf{f}}_S(\mathbf{X}_j(t_\ell))\|^2,$$

$$S_{f \cdot g} \leftarrow S_{f \cdot g} + \sum_{j=1}^{N} \hat{\mathbf{f}}_S(\mathbf{X}_j(t_\ell))^\top \mathbf{g}(\mathbf{X}_j(t_\ell), t_\ell),$$

and compute the trace of the Jacobian $\nabla \cdot \hat{\mathbf{f}}_S(\mathbf{X}_j(t_\ell))$ via automatic differentiation, accumulating it into $S_{\nabla \cdot f}$

end

Approximate the expected negative log data likelihood (up to constants) as

$$\mathcal{L}_{\text{path}} = \frac{\Delta t}{N} \left(\tfrac{1}{2} S_{\|f\|} + S_{\nabla \cdot f} + S_{f \cdot g} \right),$$

which corresponds to evaluating the quadratic form in Eq. equation 40 under the approximate posterior Q_f.