Identifying Latent Stochastic Differential Equations

Ali Hasan*, João M. Pereira*, Sina Farsiu, and Vahid Tarokh

Abstract—We present a method for learning latent stochastic differential equations (SDEs) from high dimensional time series data. Given a high-dimensional time series generated from a lower dimensional latent unknown Itô process, the proposed method learns the mapping from ambient to latent space, and the underlying SDE coefficients, through a self-supervised learning approach. Using the framework of variational autoencoders, we consider a conditional generative model for the data based on the Euler-Maruyama approximation of SDE solutions. Furthermore, we use recent results on identifiability of latent variable models to show that the proposed model can recover not only the underlying SDE coefficients, but also the original latent variables, up to an isometry, in the limit of infinite data. We validate the method through several simulated video processing tasks, where the underlying SDE is known, and through real world datasets.

Index Terms—Stochastic differential equations, autoencoder, latent space, identifiablity, data-driven discovery.

I. Introduction

VARIATIONAL auto-encoders (VAEs) are a widely used tool to learn lower-dimensional latent representations of high-dimensional data. However, the learned latent representations often lack interpretability, and it is challenging to extract relevant information from the representation of the dataset in the latent space. In particular, when the high-dimensional data is governed by unknown and lower-dimensional dynamics, arising, for instance, from unknown physical or biological interactions, the latent space representation often fails to bring insight on these dynamics.

To address this shortcoming, we propose a VAE-based framework for recovering latent dynamics governed by stochastic differential equations (SDEs). SDEs are a generalization of ordinary differential equations, that contain both a deterministic term, denoted by drift coefficient, and a stochastic term, denoted by diffusion coefficient. SDEs are often used to study stochastic processes, with applications ranging from modeling physical and biological phenomena to financial markets. Moreover, their properties have been extensively studied in the fields of probability and statistics, and a rich set of tools for analyzing these have been developed. However, most tools are limited to lower dimensional settings, which further motivates recovering lower dimensional latent representations of the data.

To define the problem, suppose we observe a high-dimensional time-series $\{X_t\}_{t \in \boldsymbol{T}}$, for which there exists a unknown latent representation $\{Z_t\}_{t \in \boldsymbol{T}}$ which is governed by an SDE, with drift and diffusion coefficients that are also unknown. More specifically, the latent representation is defined by an injective function f, which we denote by latent mapping, such that

$$X_t = f(Z_t) + \epsilon_t, \quad t \in \boldsymbol{T} \tag{1}$$

where the noise terms $\{\epsilon_t\}_{t \in \boldsymbol{T}}$ are i.i.d. and independent of $\{Z_t\}_{t \in \boldsymbol{T}}$. In this paper, we propose a VAE-based model for recovering both the latent mapping and the coefficients of the SDE that governs Z_t.

We are also concerned with identifiability. Since the latent representation is unknown, applying any one-to-one mapping to Z_t yields another latent representation $\tilde{Z}_t$ of X_t, with different latent mapping and latent SDE dynamics. To pick out one latent representation, up to equivalence by one-to-one transformations on the latent space, we propose the following two-fold approach.

(i) Under some conditions over the coefficients of the SDE that governs Z_t, we show that there exists another latent representation $\tilde{Z}_t$ of X_t, such that $\tilde{Z}_t$ is governed by an SDE with an *isotropic* diffusion coefficient (Theorem 2).

(ii) We prove that the method proposed in this paper, in the limit of infinite data, is able to recover, up to an isometry, the mapping from X_t to $\tilde{Z}_t$ and the drift coefficient of the SDE that governs $\tilde{Z}_t$ (Theorems 5, 6).

By assuming the diffusion coefficient is isotropic, our approach has an easier task of learning the latent dynamics, since the diffusion coefficient does not need to be estimated. An example of the proposed method is presented in Fig. 1.

Our paper is organized as follows. First, we present an overview of previous work. Then we review the notion of SDEs, develop a generative model to study latent SDEs, and present the VAE framework that enables learning of the proposed model. Followingly, we show that the VAE proposed recovers the true model parameters up to an isometry and give some practical considerations on the method presented. Finally, we test the proposed method in several synthetic and real world video datasets, governed by lower dimensional SDE dynamics, and present a brief discussion on the results.

II. Related Work

Previous work in learning SDEs has been mostly focused on lower-dimensional data. Classical approaches assume fixed drift and diffusion coefficients with parameters that need to be estimated [1]. In [2], a method is proposed where the terms of the Fokker-Planck equation are estimated using sparse regression with a predefined dictionary of functions, and in [3] a similar idea is applied to the Kramers-Moyal expansion of
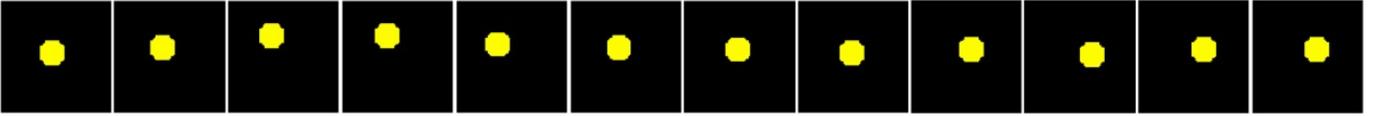
(a) Yellow ball moving according to a 2D Ornstein-Uhlenbeck process;
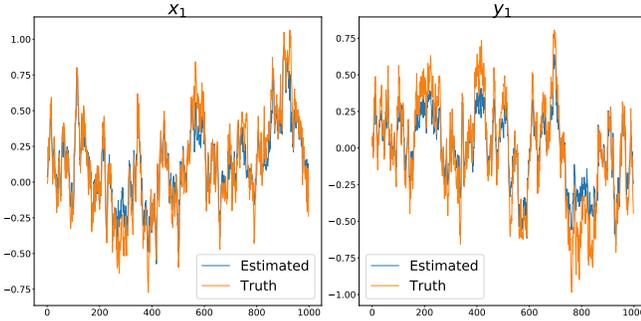
(b) Comparison between the true centers of the ball and the latent representation learned by the VAE at different frames of the video;
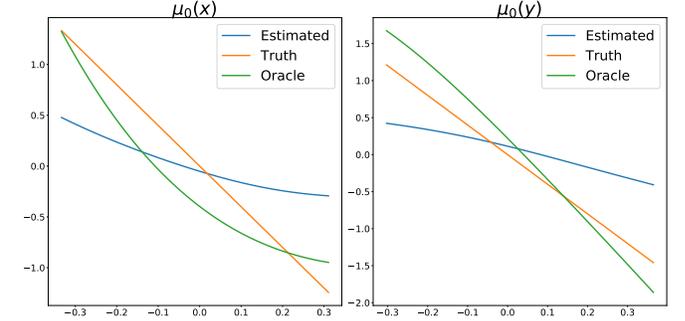
(c) Comparison between the true drift coefficient of the underlying SDE, the drift coefficient learnt by SDE–VAE, and the drift coefficient learnt by an oracle that knows the true latent variables. The difficulty in estimation is seen even in the oracle estimate;

Fig. 1: From a video of a yellow ball moving in the plane, according to a 2D Ornstein-Uhlenbeck process (Fig. 1a), the proposed model learns that the relevant latent representation of each frame are the x and y coordinates of the ball (Fig. 1b), and learns the drift coefficient of the SDE (Fig. 1c).

the SDE. In [4], the authors describe a method that makes use of a SDE driven by a counting process, while [5] describes a method for recovering an SDE using Gaussian processes. The statistical model we introduce for learning latent SDEs is similar to a Hidden Markov Model (HMM) with a complicated emission model. For this problem, spectral methods [6], and extensions with non-parametric emission models [7]–[9], have been proposed.

The work that resembles the most our contribution is [10], where a method is presented for uncovering the latent SDE for high-dimensional data using Gaussian processes. However, the method assumes that there is an intermediate feature space, such that the map from ambient space to feature space is known, and the map from feature space to latent space is linear and unknown. Its applicability is therefore limited when it is not clear what features of the data should be considered.

Regarding work that involves neural networks, [11] describes a variational inference scheme for SDEs using neural networks, and in [12], a method using variational auto-encoders is presented to recover latent second-order ordinary differential equations from data, but the dynamics are assumed to be governed by a deterministic ODE. Other related works that exploit knowledge the existence of SDEs are [13], where the adjoint sensitivity method generalized for backpropagating through an SDE solver to train neural networks and [14] which describes a method for using an autoencoder with path integrals in control scenarios. However, none are interested in recovering an underlying SDE or analyzing what SDE was recovered.

Finally, in the case of image/video data, recent works in stochastic video prediction [15], [16] describe methods for stochastic predictions of video. While these are favorable on reproducing the dynamics of the observed data, the latent variables lack interpretability. In [17], a method based on recurrent neural networks is presented that decomposes the latent space and promotes disentanglement, in an effort to provide more meaningful features in the latent space.

In all of the related work, none address the problem of recovering an underlying SDE given high dimensional measurements, balancing both interpretability and efficacy in modeling complex data sets. The proposed method aims to fill this gap.

III. STOCHASTIC DIFFERENTIAL EQUATIONS

Here we review the definition of SDE. For a time interval $\mathbb{T} = [0, T]$, let $\{W_t\}_{t \in \mathbb{T}}$ be a d-dimensional Wiener process. We say the stochastic process $\{Z_t\}_{t \in \mathbb{T}} \in \mathbb{R}^d$ is a solution to the Itô SDE

$$\mathrm{d}Z_t = \mu(Z_t, t)\,\mathrm{d}t + \sigma(Z_t, t)\,\mathrm{d}W_t, \tag{2}$$

if Z_0 is independent of the σ-algebra generated by W_t, and

$$Z_T = Z_0 + \int_0^T \mu(Z_t, t)\,\mathrm{d}t + \int_0^T \sigma(Z_t, t)\,\mathrm{d}W_t. \tag{3}$$

Here we denote the *drift coefficient* by $\mu : \mathbb{R}^d \times \mathbb{T} \to \mathbb{R}^d$, the *diffusion coefficient* by $\sigma : \mathbb{R}^d \times \mathbb{T} \to \mathbb{R}^{d \times d}$ and the second integral in (3) is the Itô stochastic integral [18]. When the coefficients are globally Lipschitz, that is,

$$\|\mu(x, t) - \mu(y, t)\| + \|\sigma(x, t) - \sigma(y, t)\| \leq D\|x - y\|$$
$$\forall x, y \in \mathbb{R}^d, t \in \mathbb{T}, \tag{4}$$

for some constant $D > 0$, there exists a unique t-continuous strong solution to (2) [18, Theorem 5.2.1]. Finally, throughout

the paper we can assume $\sigma(z, t)$ is a symmetric positive semi-definite matrix for all $z \in \mathbb{R}^d$ and $t \in \mathbb{T}$, which follows from [18, Theorem 7.3.3].

For ease of exposition, we present our main results for SDEs with time independent coefficients, and extend the results to time-dependent coefficients in Section VII.B.

IV. PROBLEM DEFINITION

In this paper, we consider a high-dimensional stochastic process $\{X_t\}_{t \in \mathbb{T}} \in \mathbb{R}^n$, which has a latent representation $\{Z_t\}_{t \in \mathbb{T}} \in \mathbb{R}^d$, with $n \geq d$, as defined in (1). Moreover, Z_t is governed by an SDE, with drift coefficient $\mu : \mathbb{R}^d \to \mathbb{R}^d$ and diffusion coefficient $\sigma : \mathbb{R}^d \to \mathbb{R}^{d \times d}$. The aim of this paper is to recover the latent mapping f and the coefficients of the SDE that governs Z_t (μ and σ) from X_t. We consider the following problem.

Problem 1. *Find (f, μ, σ) such that (1) holds, and $\{Z_t\}_{t \in \mathbb{T}}$ is a solution to the SDE with drift and diffusion coefficients μ and σ, respectively.*

By definition, the latent space is unknown, so any one-to-one transformation of the latent space cannot be recovered from the observed data. Therefore there is an inherent ambiguity of one-to-one functions for Problem 1, which we formalize as follows.

Proposition 1. *Consider the equivalence relation,*

$$(f, \mu, \sigma) \sim (\tilde{f}, \tilde{\mu}, \tilde{\sigma}), \qquad (5)$$

if there is an invertible function $g : \mathbb{R}^d \to \mathbb{R}^d$ such that

- *For any solution Y_t of (2), $g(Y_t)$ is a solution to (2) with drift and diffusion coefficients $\tilde{\mu}$ and $\tilde{\sigma}$, respectively[1];*
- *$\tilde{f}(z) = (f \circ g^{-1})(z) \quad \forall z \in \mathbb{R}^d$.*

Then if (f, μ, σ) is a solution to Problem 1 and $(\tilde{f}, \tilde{\mu}, \tilde{\sigma}) \sim (f, \mu, \sigma)$, then $(\tilde{f}, \tilde{\mu}, \tilde{\sigma})$ is also a solution to Problem 1. In particular, we can only recover (f, μ, σ) up to its equivalence class.

Proof. If $(f, \mu, \sigma) \sim (\tilde{f}, \tilde{\mu}, \tilde{\sigma})$, then there exists an invertible function g such that $\tilde{\mu}$ and $\tilde{\sigma}$ are the drift and diffusion coefficients of the SDE that governs $\tilde{Z}_t = g(Z_t)$. We have $\tilde{f} = f \circ g^{-1}$, which implies,

$$\tilde{f}(\tilde{Z}_t) = f(g^{-1}(g(Z_t))) = f(Z_t),$$

and (1) also holds with f and Z_t replaced by $\tilde{f}$ and $\tilde{Z}_t$, respectively, thus $(\tilde{f}, \tilde{\mu}, \tilde{\sigma})$ is also a solution to Problem 1. □

Since we can only recover (f, μ, σ) up to its equivalence class, we should focus on recovering an element of the equivalence class which is easier to describe. The following theorem achieves that: under some conditions on μ and σ, there is other element $(\tilde{f}, \tilde{\mu}, \tilde{\sigma})$ in the same equivalence class of (f, μ, σ) for which $\tilde{\sigma}$ is *isotropic*, that is, $\tilde{\sigma}(y) = I_d$ for all $y \in \mathbb{R}^d$, where I_d is the identity matrix of size d.

---

[1]Itô's Lemma [19] implies that if Y_t is the solution of an SDE, then $g(Y_t)$ is also the solution of another SDE, for which the drift and diffusion coefficients, $\tilde{\mu}$ and $\tilde{\sigma}$, can be explicitly written in terms of μ, σ and g.

Theorem 2. *Suppose that (f, μ, σ) is a solution to Problem 1, and that the following conditions are satisfied:*

- *(2.i) μ and σ are globally Lipschitz as in (4), and $\sigma(y)$ is symmetric positive definite for all $y \in \mathbb{R}^d$.*
- *(2.ii) σ is differentiable everywhere and for all $y \in \mathbb{R}^d$.*

$$\frac{\partial \sigma(y)}{\partial y_k} \sigma(y)^{-1} e_j = \frac{\partial \sigma(y)}{\partial y_j} \sigma(y)^{-1} e_k, \qquad (6)$$

where e_j is the j-th canonical basis vector of $\mathbb{R}^d$.

Then there exists a solution $(\tilde{f}, \tilde{\mu}, \tilde{\sigma})$ to Problem 1 such that $\tilde{\sigma}$ is isotropic.

Proof. Using Proposition 1, it suffices to find an invertible function g such that $g(Z_t)$ is governed by an SDE with an isotropic diffusion coefficient. An SDE for which such a function exists is called reducible and *2.i)* and *(2.ii)* are necessary and sufficient conditions for an SDE to have this property. See [20, Proposition 1] for a formal statement of that result and respective proof. The proof of Theorem 2 then follows by letting $\tilde{f} = f \circ g^{-1}$ and defining $\tilde{\mu}$ in terms of g, μ and σ, using Itô's Lemma [19]. □

Finally, we provide a lemma to further the understanding of Theorem 2, in particular when *(2.ii)* holds.

Lemma 3. *Suppose that σ satisfies (2.i), then any of the following conditions are sufficient for (2.ii) to hold.*

- *(3.i) The latent dimension is 1 ($d = 1$) and $\sigma(y)$ is positive.*
- *(3.ii) $\sigma(y)$ is a positive diagonal matrix and the i-th diagonal element depends only on coordinate i, that is, there exist functions $f_i : \mathbb{R} \to \mathbb{R}$ such that $\sigma(y)_{ii} = f_i(y_i)$, for all $y \in \mathbb{R}^d$.*
- *(3.iii) There exists a $d \times d$ invertible matrix M and a function $\Lambda : \mathbb{R}^d \to \mathbb{R}^{d \times d}$ such that $\sigma(y) = M\Lambda(M^{-1}y)M^T$, and $\Lambda(y)$ satisfies (3.ii), for all $y \in \mathbb{R}^d$.*
- *(3.iv) σ^{-1} is the Hessian of a convex function.*

Moreover, (3.iv) is also necessary. As an example, a Brownian motion is a reducible SDE. This condition also holds in many cases of practical interest, see examples in [21].

Proof. See Appendix B.1 □

V. ESTIMATING THE LATENT SDE USING A VAE

Motivated by the previous section, in this section we assume the latent space is governed by an SDE with an an isotropic diffusion coefficient, and describe a method for recovering (f, μ), using a VAE. While $\{X_t\}_{t \in \mathbb{T}}$ is a stochastic process defined for all $t \in \mathbb{T}$, in practice, we sample X_t at discrete times and, for ease of exposition, we assume unless stated otherwise that the sampling frequency is constant.

A. Generative Model

In order to learn the decoder and the drift coefficient, we consider pairwise consecutive time series observations $\boldsymbol{X} = (X_{t+\Delta t}, X_t)$, which correspond to the latent variables $\boldsymbol{Z} = (Z_{t+\Delta t}, Z_t)$. Accordingly, we consider the following
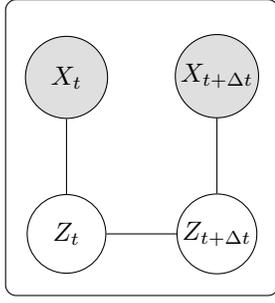
Fig. 2: Probabilistic graphical model for our generative model. The ambient variables X_t and $X_{t+\Delta t}$ only depend on each other through the latent variables Z_t and $Z_{t+\Delta t}$.

conditional generation model, with model parameters $\phi = (f, \mu, \gamma)$.

$$p_\phi(\boldsymbol{X}, \boldsymbol{Z}) = \\ p_f(X_{t+\Delta t}|Z_{t+\Delta t})p_\mu(Z_{t+\Delta t}|Z_t)p_f(X_t|Z_t)p_\gamma(Z_t). \quad (7)$$

where

- The terms $p_f(X_{t+\Delta t}|Z_{t+\Delta t})$ and $p_f(X_t|Z_t)$ are defined by (1), which implies that

$$p_f(X_t|Z_t) = p_\epsilon(X_t - f(Z_t)), \quad (8)$$

  where p_ϵ is the probability distribution function of ϵ_t.
- The prior distribution on the latent space is given by $p_\gamma(Z_t)$. This term is added to ease the training of the VAE.
- The term $p_\mu(Z_{t+\Delta t}|Z_t)$ is related to the SDE dynamics. In order to model this equation with a conditional generation model, we use the Euler-Maruyama method, which provides an approximation for the distribution of $Z_{t+\Delta t}$ that is valid if Δt is small enough. Recalling that Z_t is a solution to an SDE with an isotropic diffusion coefficient, and drift coefficient $\mu : \mathbb{R}^d \to \mathbb{R}^d$, we have

$$Z_{t+\Delta t} \approx Z_t + \mu(Z_t)\Delta t + W_{t+\Delta t} - W_t. \quad (9)$$

Since $W_{t+\Delta t} - W_t$ is distributed as a multivariate centered Gaussian variable with variance $\Delta t I_d$, we define

$$p_\mu(Z_{t+\Delta t}|Z_t) = \\ \frac{1}{\sqrt{2\pi\Delta t}^d} \exp\left(-\frac{\|Z_{t+\Delta t} - Z_t - \mu(Z_t)\Delta t\|^2}{2\Delta t}\right). \quad (10)$$

For the probabilistic generative model we consider, the ambient variables X_t and $X_{t+\Delta t}$ only depend on each other through the latent variables Z_t and $Z_{t+\Delta t}$. The corresponding Markov network model is drawn in Fig. 2.

B. VAE encoder and training loss

We describe an encoder $q_\psi(\boldsymbol{Z}|\boldsymbol{X})$, that approximates the true posterior $p_\phi(\boldsymbol{Z}|\boldsymbol{X})$, which is computationally intractable. It follows from (7) and Fig. 2 that Z_t is independent of $X_{t+\Delta t}$ conditioned on $Z_{t+\Delta t}$, and $Z_{t+\Delta t}$ is independent of X_t conditioned on Z_t, thus we can factorize

$$p_\phi(\boldsymbol{Z}|\boldsymbol{X}) = p_\phi(Z_{t+\Delta t}|X_{t+\Delta t}, Z_t)p_\phi(Z_t|X_t). \quad (11)$$

Accordingly, we can factorize our encoder as

$$q_\psi(\boldsymbol{Z}|\boldsymbol{X}) = q_{\psi_1}(Z_{t+\Delta t}|X_{t+\Delta t}, Z_t)q_{\psi_2}(Z_t|X_t). \quad (12)$$

Regarding the training loss, let $\mathcal{D} = \{x_{t+\Delta t}, x_t\}_{t \in \boldsymbol{T}}$ be the observed data, already paired into consecutive observations, and $q_\mathcal{D}$ the empirical distribution in $\mathcal{D}$. We then train the VAE by minimizing the loss

$$\mathcal{L}(\phi, \psi) \\ = D_{KL}\left(q_\psi(\boldsymbol{Z}|\boldsymbol{X})q_\mathcal{D}(\boldsymbol{X}) \,\|\, p_\phi(\boldsymbol{Z}|\boldsymbol{X})q_\mathcal{D}(\boldsymbol{X})\right) \\ \qquad\qquad\qquad\qquad - \mathbb{E}_{q_\mathcal{D}(\boldsymbol{X})}\left[p_\phi(\boldsymbol{X})\right], \quad (13) \\ = \mathbb{E}_{q_\mathcal{D}(\boldsymbol{X})}\left[\mathbb{E}_{q_\psi(\boldsymbol{Z}|\boldsymbol{X})}\left[\log q_\psi(\boldsymbol{Z}|\boldsymbol{X}) - \log p_\phi(\boldsymbol{X}, \boldsymbol{Z})\right]\right]. \\ (14)$$

Here (13) can be thought of as the negative evidence lower bound; minimizing it forces $q_\psi(\boldsymbol{Z}|\boldsymbol{X})$ to approximate $p_\phi(\boldsymbol{Z}|\boldsymbol{X})$ while maximizing the likelihood of $p_\phi(\boldsymbol{X})$ under the distribution $q_\mathcal{D}$.

To calculate (14), we use the reparametrization trick [22] to backpropagate through the SDE, and the exact form of the KL divergence between two Gaussians which we describe in Appendix A. The training algorithm then proceeds as a regular VAE.

C. An approximate encoder

For training the VAE, it is convenient to consider the simplified encoder:

$$\tilde{q}_\psi(\boldsymbol{Z}|\boldsymbol{X}) = \tilde{q}_\psi(Z_{t+\Delta t}|X_{t+\Delta t})\tilde{q}_\psi(Z_t|X_t). \quad (15)$$

This decomposition allows for using the same encoder twice, and therefore eases the training of the VAE. If $\epsilon = 0$ in (8), and since f injective, we would be able to determine $Z_{t+\Delta t}$ from $X_{t+\Delta t}$. In particular, that would imply $Z_{t+\Delta t}$ was conditionally independent of Z_t, given $X_{t+\Delta t}$, and that (15) was exact. Although the noise ϵ is not 0, we assume it is relatively small compared with the noise related to the SDE term $p_\mu(Z_{t+\Delta t}|Z_t)$. Intuitively, that implies $X_{t+\Delta t}$ gives much more information about $Z_{t+\Delta t}$ than Z_t, and we can consider the approximation

$$p_\phi(Z_{t+\Delta t}|X_{t+\Delta t}, Z_t) \approx p_\phi(Z_{t+\Delta t}|X_{t+\Delta t}), \quad (16)$$

without losing too much information. We formalize this argument in the following proposition, using mutual information. On one hand, the quantity $I(Z_{t+\Delta t}; X_{t+\Delta t})$, measures the information one gets of $Z_{t+\Delta t}$ by learning $X_{t+\Delta t}$, and $I(Z_{t+\Delta t}; Z_t|X_{t+\Delta t})$ measures the additional information one gets of $Z_{t+\Delta t}$ by further knowing Z_t. On other hand, the KL divergence term that appears in the definition of mutual information will be the same that appears in the training loss of the VAE (14). Our assumption that the noise ϵ is small compared to the SDE term can be formalized as $I(Z_{t+\Delta t}; X_{t+\Delta t}) \gg I(Z_{t+\Delta t}; Z_t)$, and this hypothesis can be used to justify our argument.

Proposition 4. *If $I(Z_{t+\Delta t}; Z_t) \ll I(Z_{t+\Delta t}; X_{t+\Delta t})$, then*

$$I(Z_{t+\Delta t}; Z_t|X_{t+\Delta t}) \ll I(Z_{t+\Delta t}; X_{t+\Delta t}), \quad (17)$$

and $I(Z_{t+\Delta t}; X_{t+\Delta t}, Z_t) \approx I(Z_{t+\Delta t}; X_{t+\Delta t})$.

Proof. The proof follows from applying the chain rule and non-negativity of the mutual information. We present the details and recall the definition of mutual information in Appendix B.2 ☐

VI. IDENTIFIABILITY

In this section, we return to the topic of identifiability. Previously, we showed that we can assume that the diffusion coefficient is isotropic, and introduced the prior parameter γ as a mechanism to ease the training of the VAE. Here we provide identifiability results for the remaining model parameters $\phi = (f, \mu, \gamma)$.

A crucial element of our analysis concerns the probability distribution of $\boldsymbol{X} = (\hat{X}_{t+\Delta t}, \hat{X}_t)$, that is, the distribution of pairwise consecutive data points. The probability distribution of $\boldsymbol{X}$ is defined by ϕ, through equation (7), by integrating over $\boldsymbol{Z}$. Suppose that $\phi^* = (f^*, \mu^*, \gamma^*)$ are the true model parameters of the data, and $\phi = (f, \mu, \gamma)$ are model parameters such that

$$p_\phi(x_{t+\Delta t}, x_t) = p_{\phi^*}(x_{t+\Delta t}, x_t) \quad \forall x_{t+\Delta t}, x_t \in \mathbb{R}^n, \quad (18)$$

Then, since these two generative models coincide, and the latent space is unknown, it is not possible to determine which of these two models provides a description, through equation (7), of the true latent space. In other words, both models provide plausible descriptions of the latent space.

It is therefore important, for identifiability purposes, to characterize all parameter configurations where the generative models coincide. In the following theorem, we show that if the generative models coincide, then the corresponding model parameters are equal up to an isometry. Recalling Proposition 1, it becomes clear that it is only possible to recover ϕ up to an isometry: if g is an isometry and Z_t is a solution to an SDE with an isotropic diffusion coefficient, then $g(Z_t)$ is a solution to another SDE also with an isotropic diffusion coefficient.

Theorem 5. *Suppose that the true generative model of $\boldsymbol{X}$ has parameters $\phi^* = (f^*, \mu^*, \gamma^*)$, and that the following technical conditions hold:*

1) *The set $\{x \in \mathcal{X} | \varphi_\epsilon(x) = 0\}$ has measure zero, where φ_ϵ is the characteristic function of the density p_ϵ defined in (8).*
2) *f^* is injective and differentiable.*
3) *μ^* is differentiable almost everywhere.*

Then, for almost all values of Δt,[2] if $\phi = (f, \mu, \gamma)$ are other model parameters such that (18) holds, then ϕ and ϕ^* are equal up to an isometry. That is, there exists an orthogonal matrix Q and a vector b, such that for all $z \in \mathbb{R}^d$:

$$f(z) = f^*(Qz + b), \quad (19)$$

$$\mu(z) = Q^T \mu^*(Qz + b), \quad (20)$$

and

$$p_\gamma(z) = p_{\gamma^*}(Qz + b). \quad (21)$$

[2]Specifically, there is a finite set S such that if $\Delta t \notin S$, the condition holds.

The proof of Theorem 5 is closely related with the theory developed in [23], and is available in Appendix C.Finally, we show the VAE framework presented in this paper can obtain the true model parameters in the limit of infinite data.

Theorem 6. *Let $\{q_\psi(\boldsymbol{Z}|\boldsymbol{X})\}_{\psi \in \Psi}$ be an encoder that can be factorized as in (12), where Ψ includes all parameter configurations of the encoder, and assume the following:*

- *The family $\{q_\psi(\boldsymbol{Z}|\boldsymbol{X})\}_{\psi \in \Psi}$ includes $p_{\phi^*}(\boldsymbol{Z}|\boldsymbol{X})$,*
- *$\mathcal{L}(\phi, \psi)$ is minimized with respect to both ϕ and ψ.*

Then, in the limit of infinite data, we obtain the true model parameters $\phi^* = (f^*, \mu^*, \gamma^*)$, up to an isometry.

Proof. See [23, Supplemental Material B.6]. ☐

We note that in this result we consider a general encoder that can be factorized as in (12), and not the simpler encoder that we introduce in (15). Empirically, we observe that using this simplification introduces a model generalization error that is small compared with the data generalization error.

VII. PRACTICAL CONSIDERATIONS

We made number of simplifying assumptions that may not hold in practical cases. Here we discuss some of their implications on the proposed method.

A. Variable sampling frequency

In order to simplify the exposition of the results, we have assumed that the sampling frequency is fixed. However the proposed framework can also accommodate variable sampling frequency with some modifications. Specifically, for two consecutive observations at times t_1 and t_2, (9) becomes

$$Z_{t_2} \approx Z_{t_1} + \mu(Z_{t_1}, t_1)(t_2 - t_1) + W_{t_2} - W_{t_1}, \quad (22)$$

and $p_\mu(Z_{t_2}|Z_{t_1})$ is defined analogously to (10). Furthermore, Theorems 5 and 6 also hold for this modification.

We note however that this approach depends on the validity of approximation (22). If $t_2 - t_1$ is too large, an adjustment of the underlying integrator may be necessary. One possible integrator is to split the interval in multiple sub-intervals, use Euler-Maruyama in each sub-interval, and use the parametrization trick for training. Other possible integrators use diffusion bridges or a multi-resolution MCMC approach inspired by the results in [24], [25].

B. SDEs with time dependence

While our primary focus is on time-independent SDEs due to their prevalence in the literature, we additionally describe how our method can also be used for time-dependent SDEs. Time-dependent SDEs have relevant applications in finance, see for example [26], [27]. We consider a similar conditional generation model as in (7), where (9) should be rewritten as (22), which implies that (10) becomes

$$p_\mu(Z_{t_2}|Z_{t_1}, t_1) = \frac{1}{\sqrt{2\pi(t_2 - t_1)}^d}$$
$$\exp\left(-\frac{\|Z_{t_2} - Z_{t_1} - \mu(Z_{t_1}, t_1)(t_2 - t_1)\|^2}{2(t_2 - t_1)}\right).$$

The encoder can also depend on time by appending the time value to the last linear layer of the encoder, if the approximation given by (12) is insufficient. Modifying Theorems 5 and 6 to accommodate time-dependent drift coefficients is straightforward, see Theorem D.10 for an example on how Theorem 5 is also valid for time-dependent SDEs. For Theorem 2, the crucial part is the following extension of [20, Proposition 1] to time-dependent SDEs, which we prove in Appendix DA.

Theorem 7 (Multivariate time-dependent Lamperti transform). *Suppose that $\{Y_t\}_{t \in \mathbb{T}} \in \mathbb{R}^d$ is a solution to the SDE:*

$$dY_t = \mu(Y_t, t)\, dt + \sigma(Y_t, t)\, dW_t, \qquad (23)$$

where $\mu : \mathbb{R}^d \times \mathbb{T} \to \mathbb{R}^d$ and $\sigma : \mathbb{R}^d \times \mathbb{T} \to \mathbb{R}^{d \times d}$. Moreover, suppose the following conditions are satisfied:

(i) *μ and σ are globally Lipschitz, that is, (4) holds, and $\sigma(y, t)$ is symmetric positive definite for all $y \in \mathbb{R}^d, t \in \mathbb{T}$.*

(ii) *σ is differentiable everywhere and for all $y \in \mathbb{R}^d$, $t \in \mathbb{T}$ and $j, k \in \{1, \ldots, d\}$*

$$\frac{\partial \sigma(y, t)}{\partial y_k} \sigma(y, t)^{-1} e_j = \frac{\partial \sigma(y, t)}{\partial y_j} \sigma(y, t)^{-1} e_k, \qquad (24)$$

where e_j is the j-th canonical basis vector of $\mathbb{R}^d$.

Then there exists a function $g : \mathbb{R}^d \times \mathbb{T} \to \mathbb{R}^d$ and $\tilde{\mu} : \mathbb{R}^d \times \mathbb{T} \to \mathbb{R}^d$ such that $Y_t = g(Z_t, t)$ and $\{Z_t\}_{t \in \mathbb{T}}$ is a solution to the SDE:

$$dZ_t = \tilde{\mu}(Z_t, t)\, dt + dW_t, \qquad (25)$$

Using the time-dependent Lamperti transform combined with the time dependence results in Theorem D.10 allows for the straightforward extension to time dependent SDEs.

C. Determining the latent dimension

In order to learn the latent dimension, we suggest using the following architecture search heuristic. Instead of considering an isotropic diffusion coefficient, we set $\sigma(y) = D$ for all $y \in \mathbb{R}^d$, where D is a diagonal matrix with learnable diagonal entries. Starting with a guess for the latent dimension, we increase it if the image reconstruction is unsatisfactory, and decrease it if some of the diagonal entries of D are close to 0 (adding an ℓ_1 regularization to the diagonal entries of D will promote sparsity and help drive some of its values to 0).

Using the likelihood in the linear case. As a first step in obtaining theoretical guarantees for determining the latent dimension, we provide a result on using the likelihood to recover the true latent dimension for the case where f, ψ are linear functions.

Theorem 8 (Latent size recovery with the likelihood). *Suppose the true generative model of $\mathbf{X}$ is generated according to a full rank linear transformation of a latent SDE Z_t with conditions on μ, σ as above*

$$X_t = AZ_t, \quad A \in \mathbb{R}^{n \times d}$$
$$dZ_t = \mu(Z_t, t)dt + \sigma(Z_t, t)dW_t, \quad Z_t \in \mathbb{R}^d.$$

Moreover suppose that $\mathrm{rank}(A) = d$ and let the estimate of A with dimension j be $\hat{A}_j$. Then in the limit of infinite data, the model with latent size j satisfying

$$\max_{j \in \mathbb{N}} \log \mathcal{L}(\hat{A}_j), \quad \hat{A}_j \in \mathbb{R}^{n \times j} \qquad (26)$$

will recover the proper latent dimension $j = d$.

Proof. The proof involves considering the likelihood of the transformed variables for different latent dimensions. See Appendix D.C for more details. $\square$

Estimating the diffusion coefficient. In Appendix DB, we present an interpretability result that considers learnable diffusion coefficients. Unfortunately, this result requires conditions that do not apply for simpler SDEs, such as Brownian random walks, therefore we decided to present Theorem 5 in the paper instead.

VIII. Experiments

We consider 4 synthetic and one real-world datasets to illustrate the efficacy of SDE-VAE.

A. Datasets

1) Moving Yellow Ball: For this dataset, we simulate the stochastic motion of a yellow ball moving according to a given SDE using the Euler-Maruyama method. That is to say, the x and y coordinates of the center of the ball are governed by an SDE with an isotropic diffusion coefficient, and the drift coefficients are defined for $(x, y) \in \mathbb{R}^2$ as follows.

$$\textbf{Constant}: \quad \mu(x, y) = (-1/4, 1/4);$$
$$\textbf{OU}: \quad \mu(x, y) = (-4x, -4y);$$
$$\textbf{Circle}: \quad \mu(x, y) = (-x - 3y, y - 3x);$$

where **OU** stands for the Ornstein-Uhlenbeck process. The latent space dimension is 2, corresponding to the x and y coordinates of the ball. Fig. 1a shows an example of the movement of the balls. We train the model on one realization of the SDE for 1000 time steps with $\Delta t = 0.01$. We rescale the realization of the SDE so that the x and y coordinates of the ball are always between 0 and 1; in practice, this only changes the map from latent to ambient space, and should not affect the ability of our method to recover the latent SDE realization. We add an extension to this dataset using *5 Moving Blue Balls* with a 10 dimensional latent space. We study an **OU** process where each of the balls reverts to a specific section within the image. This dataset is challenging because of the number of objects and due to the changes in illumination when balls overlap.

2) Moving Red Digits: To further investigate the generative properties of the proposed method, we consider images of 2 digits from the MNIST dataset moving according to an SDE in the image plane, similarly to the dataset above, and use the Euler-Maruyama method to simulate the spatial positions of the two digits. In this case, the latent space is 4-dimensional, corresponding to the x and y coordinates of the center of each of the digits. Letting $\boldsymbol{x} \in \mathbb{R}^2$ and $\boldsymbol{y} \in \mathbb{R}^2$ be the x and y

coordinates of each ball, then the drift coefficient is defined for $(\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{R}^4$ as follows.

$$\begin{aligned}
\textbf{Constant}: \quad & \mu(\boldsymbol{x}, \boldsymbol{y}) = (\boldsymbol{0}, \boldsymbol{0}); \\
\textbf{OU}: \quad & \mu(\boldsymbol{x}, \boldsymbol{y}) = (-\boldsymbol{x} - 1, -\boldsymbol{y} + 1); \\
\textbf{Circle}: \quad & \mu(\boldsymbol{x}, \boldsymbol{y}) = (-x_0 - 2x_1, -x_1 + 2x_0, \\
& \qquad\qquad - y_0 - 2y_1, -y_1 + 2y_0);
\end{aligned}$$

Here, $\mu(\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{R}^4$ and the product of vectors is defined entry-wise. For each digit pair, we generate 10 trajectories of the SDE with 100 frames in each trajectory and $\Delta t = 0.01$. In this dataset, we also rescale the realization of the SDE so that the x and y coordinates of the digits are always between 0 and 1. All diffusion matrices are the identity matrix.

3) Wasserstein Interpolation: We are interested in investigating the algorithm's efficacy on a non movement dataset by generating a series of interpolations between two images according to their Wasserstein barycenters [28]. This experiment aims to consider the method's performance on a more complicated dataset than the previous two. For instance, spatial position-based encodings of the objects in the images will not work in this case. We sample images from the COIL-20 dataset [29] to interpolate according to the realization of an SDE (an example is given in the third row of Fig. 3). To further illustrate the idea, Fig. 5 shows the output of the decoder when interpolating between two images using a Brownian bridge. We simulate a 1-D SDE which determines the relative weight of each image within the interpolation. The drift coefficient is defined as follows.

$$\begin{aligned}
\textbf{OU}: \quad & \mu(x) = -2x \quad \forall x \in \mathbb{R}; \\
\textbf{Double Well}: \quad & \mu(x) = 2x(1 - x^2) \quad \forall x \in \mathbb{R}; \\
\textbf{GBM}: \quad & \mu(x) = \frac{1}{2}x; \; \sigma(x) = x;
\end{aligned}$$

where **GBM** refers to geometric Brownian motion. We simulate 1000 images with $\Delta t = 0.01$. The realization of the SDE is rescaled so that the relative weight of each image is always between 0 and 1. For the **GBM** case, the Lamperti transform of the original SDE results in an Itô process with 0 drift and unit diffusion. We therefore compare the learned latent μ to the constant zero drift. All diffusion coefficients are constant unit apart from the GBM case.

4) Moving Ball with Wasserstein Background: As an additional challenge, we will consider a dataset where we have a 3 dimensional SDE where one component modulates the background while another moves a ball in the foreground. We use the Wasserstein interpolation as the background process and the yellow ball as the foreground. This is a complicated dataset that adds occlusion and multiple moving parts to the underlying SDE. We choose the following two drift coefficients, one OU process similar to previous experiments and another process similar to the Cox-Igersoll-Ross process with anisotropic diffusion. Letting x, y be the x, y coordinates of the ball, and z be the Wasserstein barycenter, we define the

drift coefficients, and diffusion coefficient for the anisotropic SDE, as follows.

$$\begin{aligned}
\textbf{OU}: \quad & \mu(x, y, z) = (-x, -2y, -3z); \\
\textbf{Cauchy}: \quad & \mu(x, y, z) = \left(\frac{-x}{1 + x^2}, \frac{-2y}{1 + y^2}, \frac{-3z}{1 + z^2} \right) \\
\textbf{Anisotropic}: \quad & \mu(x, y, z) = (-x, -2y, 0.6 - 0.3z); \\
& \sigma(x, y, z) = \begin{pmatrix} 1 & 2 & 0 \\ 2.5 & 3 & 0 \\ 0 & 0 & \sqrt{z} \end{pmatrix}
\end{aligned}$$

Taking the multivariate Lamperti transform of the anisotropic case, we obtain a new drift of the form

$$\mu(x, y, z) = \left(1.5x - 2y, -1.25x + y, \frac{1.4 - 1.2z}{4\sqrt{z}} \right)$$

with which we estimate the recovery efficacy.

5) Fluorescent DNA: The last dataset consists of videos of a DNA molecule floating in solution undergoing random thermodynamic fluctuations as described in the work of [30]. The videos undergo minimal pre-processing, through histogram equalization based on [31] and normalization of the pixel values. The "ground truth" latent variables are obtained by segmenting the molecule using a method similar to the one described in [30] and using the center of the segmented molecule. Using these latent variables as ground truth, we compute the best affine mapping between these and the ones estimated using our method and report the results in Table I. For this experiment we analyze three datasets, two with additional noise added in one direction of the molecule and the other with no noise added (denoted by $V = 1, 2$ and $V = 0$ respectively). The $V = 1$, $V = 2$ datasets have an anisotropic diffusion coefficient, with additional intensity given in the y variable whereas in the $V = 0$ case the noise is isotropic with the identity matrix as the diffusion following equations (3) and (5) from [30]. In both cases, the ground truth should be a random walk corresponding with no drift, that is, $\mu = (0, 0)$. We use these as the ground truth drift values and compute the MSE between the estimated and the theoretical drift. Since for this dataset the diffusion coefficient is not known, we compute the latent variables using an affine transformation, rather than an orthogonal transformation. Since the datasets are fixed, we repeat the experiment 5 times with different initializations of the neural networks to report the values in the table.

B. Experiment Setup

For all experiments, we use a convolutional encoder-decoder architecture. The latent drift coefficient μ is represented as a multi layer perceptron (MLP). We consider the same network architectures between all experiments in order to maintain consistency. All architecture and hyper-parameter specifications are available in the supplemental material. All image sizes are $64 \times 64 \times 3$ making the ambient dimension of size 12,288.

First, we test the proposed method on learning the latent mapping in these datasets. In order to show evidence of the theoretical results presented in Section VI, we measure the mean square error (MSE) between the true latent representation and the true drift coefficient, with the ones estimated by

Dataset	SDE Type	$\mathcal{L}_{\text{latent}}$	$\mathcal{L}_\mu$	$\mathcal{L}_\mu$ CRLB	Reconstruction MSE
	Constant	$1.08(\pm0.47)\times10^{0}$	$5.07(\pm2.99)\times10^{-1}$		$1.71(\pm1.85)\times10^{-2}$
Balls	OU	$1.41(\pm1.20)\times10^{-2}$	$9.20(\pm2.23)\times10^{-1}$	2.00×10^{-1}	$3.23(\pm12.0)\times10^{-3}$
	Circle	$4.65(\pm2.56)\times10^{-2}$	$1.47(\pm0.62)\times10^{0}$		$5.10(\pm3.55)\times10^{-3}$
Multiple Balls	OU	$2.90(\pm0.79)\times10^{-1}$	$2.61(\pm0.51)\times10^{0}$	1.00×10^{0}	$1.69(\pm1.05)\times10^{-2}$
	Constant	$4.36(\pm1.11)\times10^{-1}$	$7.74(\pm4.65)\times10^{-1}$		$4.58(\pm1.13)\times10^{-3}$
Digits	OU	$1.99(\pm0.47)\times10^{-1}$	$7.69(\pm4.91)\times10^{-1}$	4.00×10^{-1}	$2.57(\pm0.96)\times10^{-3}$
	Circle	$2.38(\pm1.15)\times10^{-1}$	$2.15(\pm0.48)\times10^{0}$		$3.02(\pm0.88)\times10^{-3}$
	OU	$3.53(\pm4.31)\times10^{-2}$	$8.79(\pm1.21)\times10^{-1}$		$5.55(\pm3.03)\times10^{-4}$
Wasserstein	Double Well	$5.20(\pm6.40)\times10^{-2}$	$1.45(\pm0.76)\times10^{0}$	1.00×10^{-1}	$2.08(\pm3.33)\times10^{-3}$
	GBM	$3.28(\pm2.32)\times10^{-1}$	$1.23(\pm2.45)\times10^{-1}$		$2.23(\pm2.81)\times10^{-3}$
	OU	$4.23(\pm2.70)\times10^{-2}$	$7.08(\pm2.79)\times10^{-1}$		$1.76(\pm0.68)\times10^{-2}$
Ball + Wasserstein	Cauchy	$1.24(\pm0.90)\times10^{-1}$	$8.82(\pm7.36)\times10^{-1}$	3.00×10^{-1}	$1.91(\pm0.81)\times10^{-2}$
	Anisotropic	$1.03(\pm0.82)\times10^{0}$	$1.74(\pm0.50)\times10^{0}$		$1.90(\pm0.43)\times10^{-2}$
Ball + $\mathcal{N}(0,4)$		$1.16(\pm0.51)\times10^{-2}$	$1.22(\pm0.39)\times10^{0}$		$8.99(\pm5.28)\times10^{-3}$
Ball + t-noise		$1.14(\pm0.30)\times10^{-2}$	$1.29(\pm0.17)\times10^{0}$		$6.99(\pm1.99)\times10^{-3}$
Wasserstein + $\mathcal{N}(0,4)$	OU	$1.16(\pm0.51)\times10^{-2}$	$9.51(\pm5.63)\times10^{-1}$	N/A	$2.09(\pm0.43)\times10^{-3}$
Wasserstein + t-noise		$3.83(\pm0.55)\times10^{-2}$	$9.95(\pm1.12)\times10^{-1}$		$1.61(\pm0.43)\times10^{-3}$
Ball + Wasserstein + $\mathcal{N}(0,4)$		$8.57(\pm4.33)\times10^{-2}$	$1.34(\pm1.05)\times10^{0}$		$2.23(\pm0.24)\times10^{-2}$
Ball + Wasserstein + t-noise		$9.62(\pm5.22)\times10^{-2}$	$6.28(\pm1.39)\times10^{-1}$		$2.15(\pm0.39)\times10^{-2}$
	$V=0$ (Constant)	$7.93(\pm0.01)\times10^{-2}$	$1.49(\pm1.89)\times10^{-2}$		$1.05(\pm0.01)\times10^{0}$
Fluorescing DNA	$V=1$ (Constant)	$1.87(\pm0.05)\times10^{-1}$	$8.04(\pm4.03)\times10^{-2}$	N/A	$1.00(\pm0.00)\times10^{0}$
	$V=2$ (Constant)	$2.03(\pm0.15)\times10^{-1}$	$2.39(\pm2.09)\times10^{-1}$		$9.45(\pm0.02)\times10^{-1}$

TABLE I: Comparison of the MSE defined in (27), (28), for the proposed method across different datasets. The estimation error for the drift coefficient $\mathcal{L}_\mu$ is reported for learning μ with an MLP, and compared with the information theoretical CRLB.

the VAE. Since Theorems 5 and 6 imply that the true latent space and the one obtained by the VAE are equal up to an isometry, we measure the MSE using the following formulas.

$$\mathcal{L}_{\text{latent}} = \frac{1}{N} \min_{Q,b} \sum_{t=0}^{N} \|Q\tilde{f}^{-1}(X_t) + b - Z_t\|_2^2, \qquad (27)$$

$$\mathcal{L}_\mu = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} \|Q\hat{\mu}(Q^T(x - b)) - \mu(x)\|_2^2. \qquad (28)$$

Here, the minimum is over all orthogonal matrices $Q \in \mathbb{R}^{d \times d}$ and vectors $b \in \mathbb{R}^d$, and $\tilde{f}^{-1}$ is the function learned by the VAE encoder. The minimizers of (27) are calculated using a closed form solution which we describe in the Appendix F. To calculate (28), we use the minimizers of (27), and $\mathcal{X}$ is the set of sampled points in the run. All experiments are repeated over 5 independent runs, and the average and standard deviation are reported in Table I. We also compare $\mathcal{L}_\mu$ obtained by learning μ using an MLP and a Cramér Rao lower bound (CRLB) for $\mathcal{L}_\mu$. The CRLB is obtained using an information theoretic argument, and provides a lower bound for the MSE of any estimator of μ, therefore can considered as a baseline of what is theoretically achievable. The derivation of the CRLB for these experiments is described in Appendix F.

Finally, regarding the oracle in Fig. 1c, we use the same network architecture for the drift coefficient μ, which is trained by maximizing the log-likelihood of the Euler-Maruyama approximation of the latent SDE.

C. Interpretation of results

Comparing $\mathcal{L}_{\text{latent}}$ in Table I between the different experiments, we see that the proposed method is able to learn the latent representation better for the yellow ball, which was expected since this was the simplest dataset. Comparing $\mathcal{L}_\mu$ with the theoretical CRLB, we observe that our method is able to recover the drift coefficient within the same order of magnitude for the OU process in all datasets, the constant drift process (Brownian random walk) for the digits dataset, and the geometric Brownian motion case for the Wasserstein distance. On the other hand, we believe the performance for the yellow ball dataset, constant process, was hindered by the fact that any solution to that SDE is unbounded, and when we rescale the SDE, so that the coordinates are between 0 and 1, we lose information. Additionally, in the circle cases, the data is largely concentrated in the circular region, but certain jumps from the noise cause the extreme points to be poorly learned, leading to a higher MSE. This behavior also is exhibited in the double well case where the bulk of the data is within the wells but regions outside the potential have higher MSEs.

The final column of Table I describes the MSE between the ground truth image and the reconstruction from the decoder $\|X_t - \hat{f}(\hat{Z}_t)\|_2^2$. Examples of the reconstructions are given in Fig. 3, the reconstruction of the test set data qualitatively matches with respect to the locations of the original images. This suggests that the generative capabilities of the method are effective in generating new images conditioned on proper coordinates given by the latent SDE.

For the noisy synthetic datasets, we compare the MSE to the original, denoised image. In these cases, the MSE for the noisy experiments and the noiseless experiments are within one order of magnitude, suggesting the proposed method is effectively denoising the image and tracking only the relevant object governed by the SDE. Examples of the denoising are
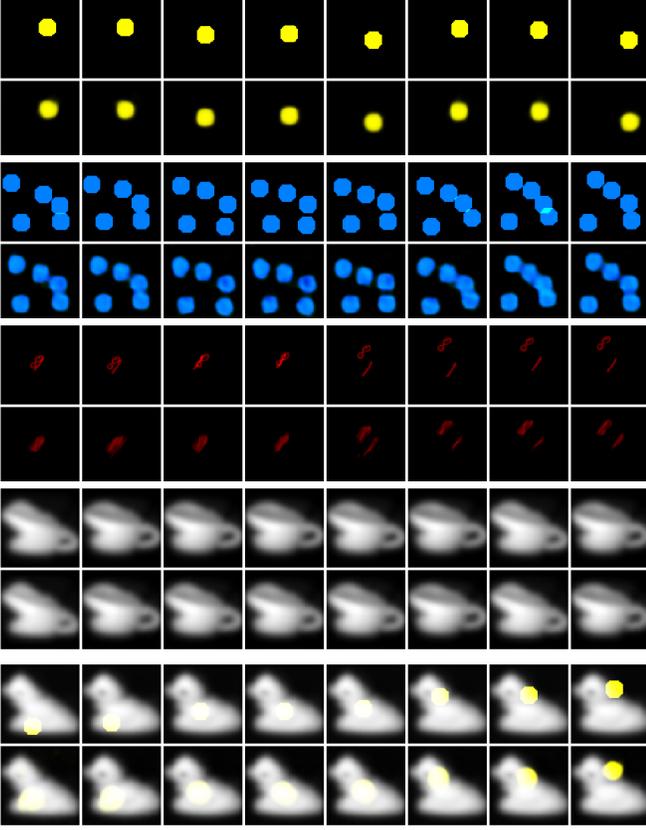
Fig. 3: Qualitative comparison of image reconstructions for the Ornstein-Uhlenbeck process on held-out test data. Top row is the original data and bottom row is the reconstruction.
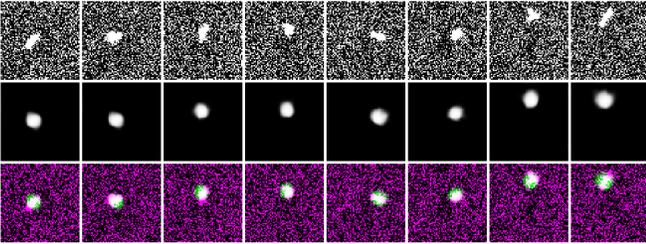
Fig. 4: Reconstruction of the DNA molecule on test data from a 2D latent space. The proposed method effectively tracks the location of the molecule while disregarding noise. Top row: histogram normalized data, middle row: reconstruction based on 2D latent space, bottom row: overlay of original input and reconstructed output.

given in Fig. 6. For the DNA datasets, we do not have a ground truth and report the MSE between the noisy original images. As expected, the MSE is high due to the autoencoder denoising the image.

D. Adding observation noise

In many real applications, such as in the DNA molecule datasets, the observation is corrupted by noise. Techniques for dealing with this type of data have been extensively

Fig. 5: Sampling a Brownian bridge in the Wasserstein OU dataset. Top two rows: ground truth generated bridge between endpoints. Bottom two rows: sampled bridge between endpoints from latent space using proposed method. First, the endpoint images are embedded into the latent space. Then, we sample a Brownian bridge between the two latent points. Finally, we decode the latent bridge using the decoder.

studied in the field of filtering [32]. We conduct additional experiments where we use the proposed method to uncover latent SDEs with observation noise. Although this violates our previous assumption that the observation noise is small, we wanted nevertheless to analyze the empirical performance of our method for more noisy datasets. For these experiments, we generate the movies according to the same procedures described in previous sections, but we add additional noise to the final output. That is, we observe $X_t + \epsilon_t$ where ϵ_t is sampled either from a Gaussian distribution or a Student's t-distribution where the degrees of freedom parameter is set to 3. We include the t-distributed noise experiments since the Student's t-distribution exhibits a fatter tail, which should make estimation harder. We illustrate examples of the image reconstruction capabilities in Fig. 6 for the moving ball and Wasserstein datasets with both Student-t and Gaussian noise.

The results suggest that, even though the encoder approximation is valid when the observation noise is low, as stated in Proposition 4, empirically our method performs well for datasets with considerable observation noise: all experiments are within 25% error of the noise-free experiments.

E. Learning the latent dimension

In order to validate the heuristic described in Section VII-C, we analyze three additional experiments where we attempt to learn the dimension of the latent space using that heuristic. We assume the ground truth is 2-dimensional in all cases, corresponding to the planar movement of the molecule. However, there exists additional movement in the orientation of the molecule, which may contribute to the latent dynamics and a latent space larger than 2. We illustrate the diagonal entries obtained by the heuristic, sorted in decreasing order, in Fig. 7. We use the same architecture and parameters as in the experiments described in Section VIII-A5, but for these experiments we include a learnable parameter for the diffusion term and set the latent space dimension to 16.

The estimated diagonal entries show a rapid decay towards zero, with only the first few having values greater than 0.1.
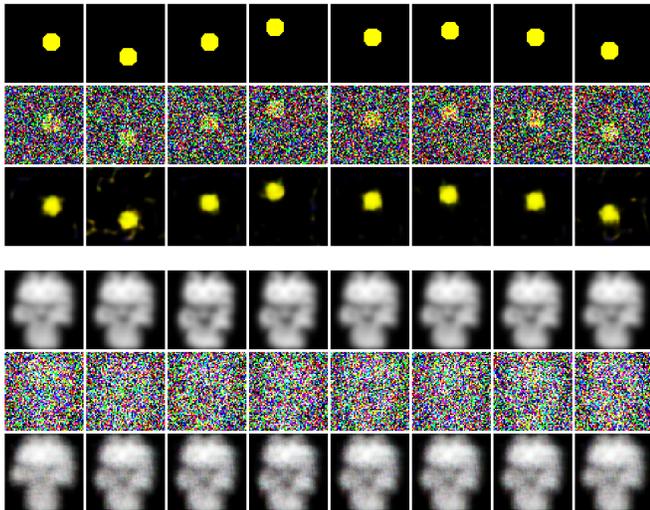
Fig. 6: Two experiments with different noise corruption. Top row: original data; middle row: observation; bottom row: reconstruction. Top experiment with t-distributed noise, bottom with Gaussian distributed noise.
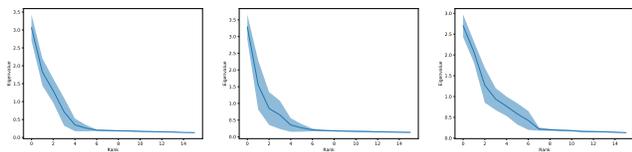
Fig. 7: Sorted diagonal entries learned by the heuristic described in Section VII-C when applied to the fluorescent DNA datasets [30] with $V = 0$ (left), $V = 1$ (middle) and $V = 2$ (right). The latent space has dimension 16.

Practically speaking, one would need to choose a threshold for which the diagonal values below this threshold are considered noise and not part of the true latent space.

Additional experiments are included in Appendices J and Appendix K.
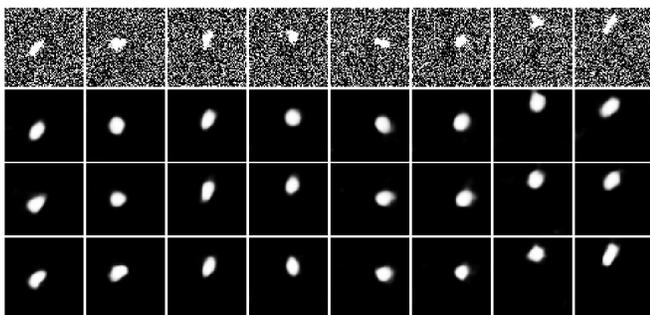
Fig. 8: Comparison of reconstructions on test set for different size latent space. Increasing latent size improves reconstruction of orientation. From top to bottom, first row: Original data; second row: 4-D latent space; third row: 6-D latent space; fourth row: 8-D latent space.

IX. DISCUSSION

In this paper we present a novel approach to learn latent SDEs using a VAE framework. We describe a method that applies to very high dimensional cases, including video data. We showed that a large class of latent SDEs can be reduced to latent SDEs with isotropic diffusion coefficients. We prove that the proposed method is able to recover a the latent SDE in this class, up to isometry, and validate our results with numerical experiments. In most cases, the experiments suggest the method is learning the appropriate SDE up to the order of the Cramér-Rao lower bounds we obtained, with a few cases being more difficult than others.

We anticipate the proposed theory and numerical results to lay the foundations for a multitude of downstream applications. As an example, the proposed method could be used to learn an SDE governing a time series of patient imaging data. The latent drift function could be used to determine whether or not a patient is at high risk for significant deterioration and can help with intervention planning.

There are a variety of additional avenues for expanding the method. Recently, the work by [33] describes an extension to the framework established by [23] wherein the authors propose a method for identifiability that does not require knowledge of the intrinsic dimension. A similar approach can be employed in this method where eigenvalues of the estimated diffusion coefficient diminish to form a low rank matrix, indicating unnecessary components. Another extension would be to extend our results to diffusion coefficients that do not satisfy (6). We can also consider alternative metrics rather than the KL divergence for regularizing the increments, such as the Fisher divergence or Wasserstein distance.

In other directions, we consider extending our analysis for problems where the samples are sparsely sampled in time. Specifically, how the changes in an underlying integrator affect the theoretical analysis and how existing integrators such as the one proposed in [24] can be merged into the proposed method. Furthermore, a thorough analysis of using a neural network to model the drift coefficient, versus parametric forms, or a dictionary of candidate functions, similar to [2], is warranted, in the interest of interpretability. Moreover, the recovery of the latent SDE provides a natural extension for stochastic control and reinforcement learning where guarantees can be achieved based on the recovered drift function.

A very promising extension of the work is to the case where there is significant observation noise or the observation satisfies an SDE driven by non-Gaussian noise, as is common in filtering problems. A straightforward way to extend our method for this purpose would be a practical implementation of an encoder that decomposes as in (12). This type of problem was addressed by [10] where they consider observations with point processes and with additional noise corruption. Considering such observations can increase the applicability of the proposed method to additional problems where this type of observation noise is prevalent.

Finally, considering other types of stochastic processes, such as Lévy flights or jump processes, could provide additional meaningful avenues for applications of the work.

APPENDIX A
CALCULATING THE VAE LOSS

To calculate the expectation in (14) we use three techniques: Exact formula (when possible), the reparametrization trick [22] and a first order Taylor approximation, which we explain as follows. Let Z be a random variable with expectation μ, and suppose we want to estimate $\mathbb{E}[f(Z)]$ for some function f. We can use the reparametrization trick to estimate this, but an even simpler way is the rough estimate $\mathbb{E}[f(Z)] \approx f(\mu)$. This corresponds to the exact value of a first order Taylor approximation of f:

$$
\begin{aligned}
\mathbb{E}[f(Z)] &= \mathbb{E}[f(\mu) + J_f(\mu)(Z - \mu) \\
&\quad + \{2^{\text{nd}} \text{ order and higher terms}\}], \\
&= f(\mu) + \{2^{\text{nd}} \text{ order and higher terms}\}.
\end{aligned}
$$

Going back to calculating the training loss, we expand (14), getting

$$
\begin{aligned}
\mathcal{L}(\phi, \psi) = \mathbb{E}_{q_{\mathcal{D}}(\boldsymbol{X})} \big[\mathbb{E}_{\tilde{q}_\psi(\boldsymbol{Z}|\boldsymbol{X})} \big[& \log \tilde{q}_\psi(Z_{t+\Delta t}|X_{t+\Delta t}) \\
& + \log \tilde{q}_\psi(Z_t|X_t) \quad \text{(A.1)} \\
& - \log p_\gamma(Z_t) \quad \text{(A.2)} \\
& - \log p_\mu(Z_{t+\Delta t}|Z_t) \quad \text{(A.3)} \\
- \log p_f(X_{t+\Delta t}|Z_{t+\Delta t}) & - \log p_f(X_t|Z_t) \big] \big] \quad \text{(A.4)}
\end{aligned}
$$

We now explain how we calculate/estimate each of the expectations (A.1), (A.2), (A.3), and (A.4). As in vanilla VAEs, our encoder $\tilde{q}_\psi$ is given by a Gaussian distribution conditioned on X_t. We define $\boldsymbol{\mu}_{\tilde{q}}(X_t)$ and $\boldsymbol{L}_{\tilde{q}}(X_t)$ as the neural networks that encode its mean and its Cholesky decomposition of the covariance matrix. That is, letting $\boldsymbol{\Sigma}_{\tilde{q}}(X_t)$ be the covariance matrix, we have $\boldsymbol{\Sigma}_{\tilde{q}}(X_t) = \boldsymbol{L}_{\tilde{q}}(X_t)\boldsymbol{L}_{\tilde{q}}(X_t)^\top$. The probability distribution function $\tilde{q}_\psi$ is defined by:

$$
\tilde{q}_\psi(Z_t|X_t) = \\
\frac{1}{\sqrt{2\pi}^d |\det \boldsymbol{L}_{\tilde{q}}(X_t)|} \exp\left(-\frac{\|\boldsymbol{L}_{\tilde{q}}(X_t)^{-1}(Z_t - \boldsymbol{\mu}_{\tilde{q}}(X_t))\|^2}{2} \right).
$$

This implies $\mathbb{E}[Z_t|X_t] = \boldsymbol{\mu}_{\tilde{q}}(X_t)$ and

$$
\begin{aligned}
\mathbb{E}[Z_t Z_t^\top|X_t] &= \boldsymbol{L}_{\tilde{q}}(X_t)\boldsymbol{L}_{\tilde{q}}(X_t)^\top + \boldsymbol{\mu}_{\tilde{q}}(X_t)\boldsymbol{\mu}_{\tilde{q}}(X_t)^\top, \\
&= \boldsymbol{\Sigma}_{\tilde{q}}(X_t) + \boldsymbol{\mu}_{\tilde{q}}(X_t)\boldsymbol{\mu}_{\tilde{q}}(X_t)^\top,
\end{aligned}
$$

thus (A.1) can be calculated exactly[3]:

$$
\begin{aligned}
& \mathbb{E}_{q_\psi(\boldsymbol{Z}|\boldsymbol{X})} \left[\log q_\psi(Z_{t+\Delta t}|X_{t+\Delta t})|\boldsymbol{X} \right] \\
&= -\log|\det \boldsymbol{L}_{\tilde{q}}(X_t)| \\
&\quad -\frac{1}{2}\mathbb{E}\left[\operatorname{tr}\left((Z_t - \boldsymbol{\mu}_{\tilde{q}}(X_t))(Z_t - \boldsymbol{\mu}_{\tilde{q}}(X_t))^\top \boldsymbol{L}_{\tilde{q}}(X_t)^{-\top}\boldsymbol{L}_{\tilde{q}}(X_t)^{-1} \right)|X_t \right], \\
&= -\log|\det \boldsymbol{L}_{\tilde{q}}(X_t)| - \frac{1}{2}\operatorname{tr}\left(\boldsymbol{\Sigma}_{\tilde{q}}(X_t)\boldsymbol{\Sigma}_{\tilde{q}}(X_t)^{-1} \right), \\
&= -\log|\det \boldsymbol{L}_{\tilde{q}}(X_t)| - \frac{d}{2}.
\end{aligned}
$$

Regarding (A.2), we observed that having the prior distribution p_γ to be Gaussian, with a fixed isotropic covariance controlled

---

[3]We do not include the term $\sqrt{2\pi}^d$ since it is a constant that does not influence the optimization.

---

by an hyper-parameter ν, worked best empirically. In this case, the expectation can then be calculated exactly:

$$
\begin{aligned}
& \mathbb{E}[-\log p_\gamma(Z_t)|X_t] \\
&= -\frac{d}{2}\log\nu + \frac{\nu}{2}\mathbb{E}\left[\|Z_t\|^2|X_t \right], \\
&= -\frac{d}{2}\log\nu + \frac{\nu}{2}\operatorname{tr}\left(\boldsymbol{\Sigma}_{\tilde{q}}(X_t) \right) + \frac{\nu}{2}\|\boldsymbol{\mu}_{\tilde{q}}(X_t)\|^2.
\end{aligned}
$$

For calculating the loss we can ignore the term $\frac{d}{2}\log\nu$, since it stays constant during training. The term (A.3) is equal to:

$$
\begin{aligned}
& \mathbb{E}_{q_\psi(\boldsymbol{Z}|\boldsymbol{X})}\left[-\log p_\mu(Z_{t+\Delta t}|Z_t)|\boldsymbol{X} \right] \\
&= \frac{d}{2}\log\Delta t + \frac{1}{2\Delta t}\mathbb{E}\Big[\operatorname{tr}\big((Z_{t+\Delta t} - Z_t - \mu(Z_t)\Delta t) \\
&\qquad\qquad (Z_{t+\Delta t} - Z_t - \mu(Z_t)\Delta t)^\top \big)|\boldsymbol{X} \Big].
\end{aligned}
$$

We note that the only expectations that cannot be calculated exactly are the ones involving $\mu(Z_t)$. Conditioned on X_t, the distribution of Z_t is given by the encoder q_ψ, thus we use the first order approximation $\mu(Z_t) \approx \mu(\boldsymbol{\mu}_{\tilde{q}}(X_t))$. Calculating the other expectations, and noting that the definition $q_\psi(\boldsymbol{Z}|\boldsymbol{X})$ implies that $\operatorname{Cov}(Z_t, Z_{t+\Delta t}) = 0$, we get the formula:

$$
\begin{aligned}
& \mathbb{E}_{q_\psi(\boldsymbol{Z}|\boldsymbol{X})}\left[-\log p_\mu(Z_{t+\Delta t}|Z_t) \right] \\
&= \frac{d}{2}\log\Delta t + \frac{1}{2\Delta t}\operatorname{tr}(\boldsymbol{L}_{\tilde{q}}(X_t)\boldsymbol{L}_{\tilde{q}}(X_t)^\top) \\
&\quad + \frac{1}{2\Delta t}\operatorname{tr}(\boldsymbol{L}_{\tilde{q}}(X_{t+\Delta t})\boldsymbol{L}_{\tilde{q}}(X_{t+\Delta t})^\top) \\
&\quad + \frac{1}{2\Delta t}\|\boldsymbol{\mu}_{\tilde{q}}(X_{t+\Delta t}) - \boldsymbol{\mu}_{\tilde{q}}(X_t) - \mu(\boldsymbol{\mu}_{\tilde{q}}(X_t))\|^2.
\end{aligned}
$$

Finally we model the noise ϵ in (A.4) as a centered Gaussian random variable with variance τI, where τ is a hyper-parameter. We have,

$$
\mathbb{E}_{q_\psi(\boldsymbol{Z}|\boldsymbol{X})}\left[-\log p_f(X_t|Z_t) \right] = \frac{d}{2}\log\tau + \frac{1}{2\tau}\mathbb{E}\left[\|f(Z_t) - X_t\|^2 \right],
$$

and use the reparametrization trick to calculate this.

APPENDIX B
PROOF OF LEMMAS

A. Proof of Lemma 3

Proof. The implications *(3.i)* $\Rightarrow$ *(3.ii)* $\Rightarrow$ *(3.iii)* are trivial and the Brownian motion is a particular case of *(3.iii)*, thus we show that *(3.iii)* $\Rightarrow$ *(3.iv)* and *(3.iv)* $\Leftrightarrow$ *(2.ii)*. To show *(3.iii)* $\Rightarrow$ *(3.iv)*, notice that for all $y \in \mathbb{R}^d$, $\sigma^{-1}(y) = M^{-\top}\Lambda(M^{-1}y)^{-1}M^{-1}$. Since Λ satisfies condition *(3.ii)*, we have $\Lambda(y)$ is a diagonal matrix and letting $\tilde{m}_i$ be the i-th row of M^{-1}, $\Lambda(M^{-1}y)_{ii} = f_i(\tilde{m}_i^\top y)$ for all $y \in \mathbb{R}^d$ and $i \in \{1,\ldots,d\}$. Since $\sigma^{-1}(y) = M^{-\top}\Lambda(M^{-1}y)^{-1}M^{-1}$, and and Λ is diagonal, we have,

$$
\begin{aligned}
\sigma^{-1}(y) &= \sum_{i=1}^d \Lambda(M^{-1}y)_{ii}\tilde{m}_i\tilde{m}_i^\top, \quad \forall y \in \mathbb{R}^d \\
&= \sum_{i=1}^d f_i(\tilde{m}_i^\top y)\tilde{m}_i\tilde{m}_i^\top, \quad \forall y \in \mathbb{R}^d
\end{aligned}
$$

Since f_i is positive, there is a convex function g_i such that $g_i'' = f_i$, which implies that $f_i(\tilde{m}_i^\top y)\tilde{m}_i\tilde{m}_i^\top$ is the Hessian

of $g_i(\tilde{m}_i^\top y)$. σ^{-1} is then the hessian of a positive linear combination of convex functions, which is thus a convex function. To show that *(2.ii) $\Rightarrow$ (3.iv)*, we note that (24) implies that for all j, k

$$\frac{\partial}{\partial y_k}\left(\sigma(y)^{-1}\right)e_j = -\sigma(y)^{-1}\frac{\partial\sigma(y)}{\partial y_k}\sigma(y)^{-1}e_j$$
$$= -\sigma(y)^{-1}\frac{\partial\sigma(y)}{\partial y_j}\sigma(y)^{-1}e_k,$$
$$= \frac{\partial}{\partial y_j}\left(\sigma(y)^{-1}\right)e_k,$$

thus [34, Theorem 11.49] implies that there exists a function $h : \mathbb{R}^d \to \mathbb{R}^d$ such that $\sigma(y)^{-1} = J_h(y)$ for all $y \in \mathbb{R}^d$, where J_h is the Jacobian of h. Moreover, since $\sigma^{-1}(y)$ is symmetric,

$$\frac{\partial}{\partial y_j}h_i(y) = (\sigma(y)^{-1})_{ij} = (\sigma(y)^{-1})_{ji} = \frac{\partial}{\partial y_i}h_j(y),$$

thus applying [34, Theorem 11.49] again, there is a function $g : \mathbb{R}^d \to \mathbb{R}$ such that $h(y) = \nabla g(y)$, which implies then that $\sigma(y)^{-1} = H_g(y)$, and since $\sigma(y)^{-1}$ is positive definite, g is convex. To show the other direction, suppose that $\sigma(y)^{-1} = H_g(y)$, then, for all $i, j, k \in \{1, \ldots, d\}$,

$$\frac{\partial}{\partial y_k}\left(\sigma(y)^{-1}\right)_{ij} = \frac{\partial^3 g(y)}{\partial y_i \partial y_j \partial y_k},$$
$$= \frac{\partial}{\partial y_j}\left(\sigma(y)^{-1}\right)_{ik},$$

which finally implies

$$\frac{\partial\sigma(y)}{\partial y_k}\sigma(y)^{-1}e_j = -\sigma(y)\frac{\partial}{\partial y_k}\left(\sigma(y)^{-1}\right)e_j$$
$$= -\sigma(y)\frac{\partial}{\partial y_j}\left(\sigma(y)^{-1}\right)e_k$$
$$= \frac{\partial\sigma(y)}{\partial y_j}\sigma(y)^{-1}e_k$$

$\square$

B. Proof of Proposition 4

First, we recall the definition of mutual information. The mutual information between the random variables X and Y, with joint distribution $p_{X,Y}$, marginal distributions p_X and p_Y, respectively, and conditional distribution $p_{Y|X}$, is defined as:

$$I(X; Y) := D_{KL}(p_{X,Y}(X, Y) \| p_X(X)p_Y(Y)),$$
$$= \mathbb{E}_{p_X(X)}\left[D_{KL}(p_{Y|X}(Y|X) \| p_Y(Y))\right]. \quad \text{(B.1)}$$

We now proceed with the proof of Proposition 4. By the chain rule of mutual information, the following identities hold:

$$I(Z_{t+\Delta t}; Z_t | X_{t+\Delta t})$$
$$= I(Z_{t+\Delta t}, X_{t+\Delta t}; Z_t) - I(X_{t+\Delta t}; Z_t),$$
$$\leq I(Z_{t+\Delta t}, X_{t+\Delta t}; Z_t),$$
$$= I(Z_{t+\Delta t}; Z_t) + I(X_{t+\Delta t}; Z_t | Z_{t+\Delta t}). \quad \text{(B.2)}$$

Note that (7) implies that $X_{t+\Delta t}$ and Z_t are conditionally independent given $Z_{t+\Delta t}$, which implies that $I(X_{t+\Delta t}; Z_t | Z_{t+\Delta t}) = 0$ and $I(Z_{t+\Delta t}; Z_t | X_{t+\Delta t}) \leq$

$I(Z_{t+\Delta t}; Z_t) \ll I(Z_{t+\Delta t}; X_{t+\Delta t})$. Finally, applying the chain rule,

$$I(Z_{t+\Delta t}; X_{t+\Delta t}) \leq I(Z_{t+\Delta t}, X_{t+\Delta t}; Z_t),$$
$$= I(Z_{t+\Delta t}; X_{t+\Delta t}) + I(Z_{t+\Delta t}; Z_t | X_{t+\Delta t}),$$
$$\leq I(Z_{t+\Delta t}; X_{t+\Delta t}) + I(Z_{t+\Delta t}; Z_t),$$
$$\approx I(Z_{t+\Delta t}; X_{t+\Delta t}).$$

APPENDIX C
PROOF OF THEOREM 5

Our proof is very similar to the proof of [23, Theorem 1]. However, we cannot use that result directly, since our generative model is slightly different. Therefore we replicate the proof here, adapted to our generative model, but suggest consulting [23] to understand more intricate details of the proof.

Our proof is split in 3 steps:

(I) First, we write our generation model as a convolution that depends on the noise ϵ, to reduce the equality $p_\theta(\mathbf{x}) = p_{\theta^*}(\mathbf{x})$ to the noiseless case, obtaining (C.1).
(II) Next, we use some algebraic tricks to get a linear dependence between f_*^{-1} and f^{-1}, obtaining (C.13).
(III) Finally, we keep using algebraic manipulations to get the rest of the statements in the theorem.

Step (I): Let $\mathcal{Z} = \mathbb{R}^d$ be the latent space, $\mathcal{X} = f(\mathcal{Z})$, that is, $x \in \mathcal{X}$ if there is $z \in \mathcal{Z}$ such that $f(z) = x$. First, we write $p_\theta(x_{t+\Delta t}, x_t)$ as a convolution.

$$p_\theta(x_{t+\Delta t}, x_t)$$
$$= \int_\mathcal{Z}\int_\mathcal{Z} p_\epsilon(x_{t+\Delta t} - f(z_{t+\Delta t}))p_\epsilon(x_t - f(z_t)) \quad \text{(i)}$$
$$\qquad q_{\mu,\gamma}(z_{t+\Delta t}, z_t)\,\mathrm{d}z_{t+\Delta t}\,\mathrm{d}z_t,$$
$$= \int_\mathcal{X}\int_\mathcal{X} p_\epsilon(x_{t+\Delta t} - \tilde{x}_{t+\Delta t})p_\epsilon(x_t - \tilde{x}_t) \quad \text{(ii)}$$
$$\qquad q_{\mu,\gamma,f^{-1}}(\tilde{x}_{t+\Delta t}, \tilde{x}_t)\,\mathrm{d}\tilde{x}_{t+\Delta t}\,\mathrm{d}\tilde{x}_t,$$
$$= \int_{\mathbb{R}^n}\int_{\mathbb{R}^n} p_\epsilon(x_{t+\Delta t} - \tilde{x}_{t+\Delta t})p_\epsilon(x_t - \tilde{x}_t) \quad \text{(iii)}$$
$$\qquad q_{\mu,\gamma,f^{-1},\mathcal{X}}(\tilde{x}_{t+\Delta t}, \tilde{x}_t)\,\mathrm{d}\tilde{x}_{t+\Delta t}\,\mathrm{d}\tilde{x}_t,$$
$$= \left[(p_\epsilon \times p_\epsilon) * q_{\mu,\gamma,f^{-1},\mathcal{X}}\right](\tilde{x}_{t+\Delta t}, \tilde{x}_t), \quad \text{(iv)}$$

wherein:

(i) We write $p_\theta(x_{t+\Delta t}, x_t)$ as the integration over $z_t, z_{t+\Delta t}$ of $p_\theta(x_{t+\Delta t}, x_t, z_{t+\Delta t}, z_t)$, expand p_θ as in (7), and let

$$q_{\mu,\gamma}(z_{t+\Delta t}, z_t) = p_\mu(z_{t+\Delta t}|z_t)p_\gamma(z_t).$$

(ii) We do the change of variable $\tilde{x}_t = f(z_t)$ and similarly for $\tilde{x}_{t+\Delta t}$. As in [23], the change of variable volume term for $\tilde{x}_t$ is

$$\mathrm{vol}\, J_{f^{-1}}(\tilde{x}_t) := \sqrt{\det\left(J_f(f^{-1}(\tilde{x}_t))^\top J_f(f^{-1}(\tilde{x}_t))\right)},$$

and $\tilde{x}_{t+\Delta t}$ is analogous. We then let

$$q_{\mu,\gamma,f^{-1}}(\tilde{x}_{t+\Delta t}, \tilde{x}_t)$$
$$= \mathrm{vol}\, J_{f^{-1}}(\tilde{x}_t)\,\mathrm{vol}\, J_{f^{-1}}(\tilde{x}_{t+\Delta t})q_{\mu,\gamma}(f^{-1}(\tilde{x}_{t+\Delta t}), f^{-1}(\tilde{x}_t)).$$

(iii) We let

$$q_{\mu,\gamma,f^{-1},\mathcal{X}}(\tilde{x}_{t+\Delta t}, \tilde{x}_t)$$
$$= \mathbb{1}_{\mathcal{X}}(\tilde{x}_{t+\Delta t})\mathbb{1}_{\mathcal{X}}(\tilde{x}_t)q_{\mu,\gamma,f^{-1}}(\tilde{x}_{t+\Delta t}, \tilde{x}_t),$$

so that the domain of the integral may be $\mathbb{R}^n$ instead of $\mathcal{X}$.

(iv) We finally notice the convolution formula, with $(p_\epsilon \times p_\epsilon)(\tilde{x}_{t+\Delta t}, \tilde{x}_t) = p_\epsilon(\tilde{x}_{t+\Delta t})p_\epsilon(\tilde{x}_t)$.

We now have that if $p_\theta(x_{t+\Delta t}, x_t) = p_{\theta^*}(x_{t+\Delta t}, x_t)$ for all $x_{t+\Delta t}, x_t \in \mathbb{R}^n$, then by taking the Fourier transform we get that for all $x_{t+\Delta t}, x_t \in \mathbb{R}^n$ (or $\omega_{t+\Delta t}, \omega_t \in \mathbb{R}^n$ when applicable)

$$\left[(p_\epsilon \times p_\epsilon) * q_{\mu,\gamma,f^{-1},\mathcal{X}}\right](x_{t+\Delta t}, x_t)$$
$$= \left[(p_\epsilon \times p_\epsilon) * q_{\mu_*,\gamma_*,f_*^{-1},\mathcal{X}_*}\right](x_{t+\Delta t}, x_t),$$
$$\mathcal{F}q_{\mu,\gamma,f^{-1},\mathcal{X}}(\omega_{t+\Delta t}, \omega_t)\varphi_\epsilon(\omega_{t+\Delta t})\varphi_\epsilon(\omega_t)$$
$$= \mathcal{F}q_{\mu_*,\gamma_*,f_*^{-1},\mathcal{X}_*}(\omega_{t+\Delta t}, \omega_t)\varphi_\epsilon(\omega_{t+\Delta t})\varphi_\epsilon(\omega_t),$$
$$\mathcal{F}q_{\mu,\gamma,f^{-1},\mathcal{X}}(\omega_{t+\Delta t}, \omega_t) = \mathcal{F}q_{\mu_*,\gamma_*,f_*^{-1},\mathcal{X}_*}(\omega_{t+\Delta t}, \omega_t),$$
$$q_{\mu,\gamma,f^{-1},\mathcal{X}}(x_{t+\Delta t}, x_t) = q_{\mu_*,\gamma_*,f_*^{-1},\mathcal{X}_*}(x_{t+\Delta t}, x_t),$$

where $\mathcal{F}$ denotes Fourier Transform and $\varphi_\epsilon = \mathcal{F}p_\epsilon$ is the characteristic function of ϵ. Here condition 1 guarantees that we can divide by $\varphi_\epsilon(\omega_{t+\Delta t})\varphi_\epsilon(\omega_t)$. This equality for all $x_t, x_{t+\Delta t}$ implies that $\mathcal{X} = f(\mathcal{Z}) = f^*(\mathcal{Z}) = \mathcal{X}^*$ and, for all $x_t, x_{t+\Delta t}$ in $\mathcal{X}$,

$$\log \mathrm{vol}\, J_{f_*^{-1}}(x_{t+\Delta t}) + \log \mathrm{vol}\, J_{f_*^{-1}}(x_t)$$
$$+ \log p_{\mu_*}(f_*^{-1}(x_{t+\Delta t})|f_*^{-1}(x_t)) + \log p_{\gamma_*}(f_*^{-1}(x_t))$$
$$= \log \mathrm{vol}\, J_{f^{-1}}(x_{t+\Delta t}) + \log \mathrm{vol}\, J_{f^{-1}}(x_t)$$
$$+ \log p_\mu(f^{-1}(x_{t+\Delta t})|f^{-1}(x_t)) + \log p_\gamma(f^{-1}(x_t)). \quad (C.1)$$

Step (II): We have

$$\log p_\mu(f^{-1}(x_{t+\Delta t})|f^{-1}(x_t))$$
$$= -\frac{\|f^{-1}(x_{t+\Delta t}) - f^{-1}(x_t) - \mu(f^{-1}(x_t))\Delta t\|^2}{2\Delta t}.$$

Let

$$\lambda(x_t) = f^{-1}(x_t) - \mu(f^{-1}(x_t))\Delta t, \quad (C.2)$$

$$\psi(x_t) = \mathrm{vol}\, J_{f^{-1}}(x_t) + \log p_\gamma(f^{-1}(x_t)) - \frac{\|\lambda(x_t)\|^2}{2\Delta t}, \quad (C.3)$$

and

$$\varphi(x_{t+\Delta t}) = \mathrm{vol}\, J_{f^{-1}}(x_{t+\Delta t}) - \frac{\|f^{-1}(x_{t+\Delta t})\|^2}{2\Delta t}, \quad (C.4)$$

and define analogously λ_*, ψ_* and φ_*. We have for all $x_t, x_{t+\Delta t}$ in $\mathcal{X}$,

$$\psi(x_t) + \frac{1}{\Delta t}\left\langle f^{-1}(x_{t+\Delta t}), \lambda(x_t)\right\rangle + \varphi(x_{t+\Delta t})$$
$$= \psi_*(x_t) + \frac{1}{\Delta t}\left\langle f_*^{-1}(x_{t+\Delta t}), \lambda_*(x_t)\right\rangle + \varphi_*(x_{t+\Delta t}). \quad (C.5)$$

Let w_0 be some element in $\mathcal{X}$, then subtracting two equations we obtain

$$\left\langle f^{-1}(x_{t+\Delta t}), \lambda(x_t) - \lambda(w_0)\right\rangle + \zeta(x_t) - \zeta(w_0)$$
$$= \left\langle f_*^{-1}(x_{t+\Delta t}), \lambda_*(x_t) - \lambda_*(w_0)\right\rangle, \quad (C.6)$$

where $\zeta(x) = (\psi(x_t) - \psi_*(x_t))\Delta t$. Let $z_0 = f_*^{-1}(w_0)$ and $z_1, \ldots, z_d$ a set of vectors in $\mathcal{Z}$ such that $z_1 - z_0, \ldots, z_d - z_0$ are linearly independent. Let $\mathbf{Z}$ the $d \times d$ matrix such that its i-th column is $z_i - z_0$ (by the hypothesis, $\mathbf{Z}$ is an invertible matrix). Moreover let $\mathbf{M}_*$ the $d \times d$ matrix with i-th column defined by $\mathbf{M}_i = \mu_*(z_i) - \mu_*(z_0)$, and finally let $\mathbf{\Lambda}_*$ be the matrix with i-th column defined by $\mathbf{\Lambda}_{*i} = \lambda_*(f_*(z_i)) - \lambda_*(w_0)$. By construction we have $\mathbf{\Lambda}_* = \mathbf{Z} - \Delta t\mathbf{M}_*$, or equivalently

$$\mathbf{\Lambda}_*\mathbf{Z}^{-1} = I - \Delta t\mathbf{M}_*\mathbf{Z}^{-1}. \quad (C.7)$$

Therefore $\mathbf{\Lambda}\mathbf{Z}^{-1}$ is singular only if $\frac{1}{\Delta t}$ is an eigenvalue of $\mathbf{M}_*\mathbf{Z}^{-1}$. Since that set of eigenvalues is finite, we set S in the theorem statement to be

$$S = \left\{\lambda : \frac{1}{\lambda} \text{ is a positive eigenvalue of } \mathbf{M}_*\mathbf{Z}^{-1}\right\} \quad (C.8)$$

If now $\Delta t \in \mathbb{R}^+ \backslash S$, the matrix $\mathbf{\Lambda}_*\mathbf{Z}^{-1}$ is non-singular, which implies $\mathbf{\Lambda}_*$ is non-singular. Let $\mathbf{\Lambda}_*$ defined as above, define $\mathbf{\Lambda}$ analogously for λ, and let $\psi \in \mathbb{R}^d$ such that

$$\psi_i = \zeta(f_*(z_i)) - \zeta(w_0).$$

We write (C.6) in matrix notation for $x_t = f(z_0), \ldots, f(z_d)$.

$$\mathbf{\Lambda}_*^\top f_*^{-1}(x_{t+\Delta t}) = \mathbf{\Lambda}^\top f^{-1}(x_{t+\Delta t}) + \psi, \quad (C.9)$$

which implies that

$$f_*^{-1}(x_{t+\Delta t}) = \mathbf{\Lambda}_*^{-\top}\mathbf{\Lambda}^\top f^{-1}(x_{t+\Delta t}) + \mathbf{\Lambda}_*^{-\top}\psi, \quad (C.10)$$

holds for all $x_{t+\Delta t} \in \mathcal{X}$. Equivalently, for all $z \in \mathcal{Z}$, $f_*(z) \in \mathcal{X}$ and

$$z = \mathbf{\Lambda}_*^{-\top}\mathbf{\Lambda}^\top f^{-1}(f_*(z)) + \mathbf{\Lambda}_*^{-\top}\psi, \quad (C.11)$$

If we take the derivatives on z on both sides, we get

$$I = \mathbf{\Lambda}_*^{-\top}\mathbf{\Lambda}^\top J_{f^{-1} \circ f_*}(z). \quad (C.12)$$

Since the matrix I is non-singular, this implies $\mathbf{\Lambda}_*^{-\top}\mathbf{\Lambda}^\top$ is a non-singular matrix. Therefore there exists an invertible matrix Q and a vector b such that

$$f_*^{-1}(x_{t+\Delta t}) = Qf^{-1}(x_{t+\Delta t}) + b. \quad (C.13)$$

Step (III): Equation (C.13) implies that for all $z \in \mathcal{Z}$

$$f(z) = f_*(f_*^{-1}(f(z))),$$
$$= f_*(Qf^{-1}(f(z)) + b),$$
$$= f_*(Qz + b).$$

Moreover $J_f = J_{f_*}Q$, and $\log \mathrm{vol}\, J_{f^{-1}} = \log \mathrm{vol}\, J_{f_*^{-1}} + \log|\det Q|$. Now replacing $x_{t+\Delta t}$ by $f(z)$ on (C.5) and using (C.13) and (C.3), we get for all $z \in \mathcal{Z}$ and $x_t \in \mathcal{X}$,

$$\psi(x_t) + \frac{1}{\Delta t}\langle z, \lambda(x_t)\rangle - \frac{\|z\|^2}{2\Delta t} + \log|\det Q|$$
$$= \psi_*(x_t) + \frac{1}{\Delta t}\langle Qz + b, \lambda_*(x_t)\rangle - \frac{\|Qz + b\|^2}{2\Delta t}. \quad (C.14)$$

Taking derivatives with respect to z, we get

$$\lambda(x_t) - z = Q^\top\lambda_*(x_t) - Q^\top Qz - Q^\top b. \quad (C.15)$$

Taking derivatives again we get $Q^\top Q = I$, thus Q is orthogonal. Replacing this back in (C.15), we get

$$\lambda_*(x_t) = Q\lambda(x_t) + b. \tag{C.16}$$

Replacing (C.2), letting $f(z) = x_t$ and using (C.13), we get for all $z \in \mathcal{Z}$,

$$Qz + b - \mu_*(Qz + b)\Delta t = Qz - Q\mu(z)\Delta t + b, \tag{C.17}$$

which implies $\mu(z) = Q^\top \mu_*(Qz + b)$ for all $z \in \mathcal{Z}$. Finally, replacing all equations obtained in (C.14), and using (C.3) and $|\det Q| = 1$ (since Q is orthogonal), we get

$$\log p_\gamma(f^{-1}(x_t)) = \log p_{\gamma_*}(f_*^{-1}(x_t)), \tag{C.18}$$

and (21) follows from taking $x_t = f(z)$ and using (C.13). $\qquad\square$

APPENDIX D
SUPPLEMENTARY RESULTS REGARDING PRACTICAL CONSIDERATIONS

Here we complement present the proof of several results presented in the main submission.

A. Proof of Theorem 7

Proof. We prove the lemma by constructing the functions g and $\tilde\mu$ such that the SDE that governs $Y_t = g(Z_t, t)$ is given by (23). By Ito's lemma [19], we have

$$dY_t = \left(\frac{\partial g}{\partial t}(Z_t, t) + \frac{\partial g}{\partial y}(Z_t, t)\tilde\mu(Z_t, t) + \frac{1}{2}\Delta g(Z_t, t) \right) dt + \frac{\partial g}{\partial y}(Z_t, t) dW_t, \tag{D.1}$$

where $\frac{\partial g}{\partial y}(z, t)$ is the Jacobian of g, only in terms of y, and Δg is the Laplacian, defined by

$$(\Delta g(z, t))_i = \sum_{k=1}^{d} \frac{\partial^2 g(z, t)_i}{\partial y_k^2}.$$

We now choose h and g as in Lemma D.9, noting that h is the inverse of g, therefore $Z_t = h(Y_t, t)$ and

$$\frac{\partial g}{\partial y}(Z_t, t) = \frac{\partial g}{\partial y}(h(Y_t, t), t),$$
$$= \sigma(Y_t, t).$$

We just obtained that the diffusion terms in (23) and (D.1) are equal, and the drift terms coincide if we define

$$\tilde\mu(z, t) := \frac{\partial g}{\partial y}(z, t)^{-1}\left(\mu(g(z, t), t) - \frac{\partial g}{\partial t}(z, t) - \frac{1}{2}\Delta g(z, t) \right).$$

$\qquad\square$

Lemma D.9. *Suppose σ follows conditions (i) and (ii) in Theorem 2 and define*

$$h(y, t) := \int_0^1 \sigma(\tau y, t)^{-1} y\, d\tau. \tag{D.2}$$

Then the following conditions hold:

(i)
$$\frac{\partial h}{\partial y}(y, t) = \sigma(y, t)^{-1}. \tag{D.3}$$

(ii) There exists a function $g : \mathbb{R}^d \times \mathbb{T} \to \mathbb{R}^d$ such that for all $z \in \mathbb{R}^d, t \in \mathbb{T}$, $g(z, t) = y$ whenever $h(y, t) = z$.
(iii) g is differentiable everywhere and

$$\frac{\partial g}{\partial y}(h(y, t), t) = \sigma(y, t).$$

Proof. The proof is largely algebraic manipulations, so we leave the full details to the supplementary materials. $\qquad\square$

B. Interpretability for learnable diffusion coefficients

Here we present a result similar to Theorem 5, but considering a learnable diffusion coefficient. That is, we rewrite (10),

$$p_{\mu,\sigma}(Z_{t+\Delta t}|Z_t, t) = \frac{1}{(2\pi\Delta t)^{\frac{d}{2}} \det \sigma(Z_t, t)^{\frac{1}{2}}}$$
$$\exp\left(-\frac{1}{2\Delta t}(Z_{t+\Delta t} - \lambda(Z_t))^\top \sigma(Z_t, t)^{-1}(Z_{t+\Delta t} - \lambda(Z_t)) \right).$$

where $\sigma(Z_t, t)$ is the diffusion coefficient and $\lambda(Z_t) = Z_t + \Delta t\mu(Z_t, t)$. We first state the result, then explain some of our reservations against it, and finally present its proof.

A brief note on notation: for a symmetric $n \times n$ matrix M, we denote by $\operatorname{vsym} M$, a vector of dimension $\binom{n+1}{2}$, which consists of M flattened to a vector, such that the entries off-diagonal only appear once, and are multiplied by $\sqrt{2}$. This definition implies that, for 2 symmetric matrices A and B, we have $\operatorname{tr}(AB) = (\operatorname{vsym} A)^\top \operatorname{vsym} B$. Moreover, we denote by $u \oplus v$ the concatenation of vectors u and v.

Theorem D.10. *Suppose that the true generative model with arbitrary diffusion coefficients has parameters $\theta^* = (f^*, \mu^*, \sigma^*, \gamma^*)$, and that the following conditions hold:*

1) The set $\{x \in \mathcal{X} | \varphi_\epsilon(x) = 0\}$ has measure zero, where φ_ϵ is the characteristic function of the density p_ϵ defined in (8).
2) f^* is injective and differentiable.
3) Letting $N = \binom{n+1}{2} + n$, there exist $N + 1$ vectors $z_0, \ldots, z_N$ and scalars $t_0, \ldots, t_N$ such that the vectors $\mathbf{\Lambda}_1 - \mathbf{\Lambda}_0, \ldots, \mathbf{\Lambda}_N - \mathbf{\Lambda}_0$, with

$$\mathbf{\Lambda}_i := \operatorname{vsym}(\sigma(z_i, t_i)^{-1}) \oplus \left(-2\sigma(z_i, t_i)^{-1}(z_i - \mu(z_i, t_i)) \right), \tag{D.4}$$

are linearly independent.

Then if $\theta = (f, \mu, \sigma, \gamma)$ are other parameters that yield the same generative distribution, that is

$$p_\theta(x_{t+\Delta t}, x_t) = p_{\theta^*}(x_{t+\Delta t}, x_t) \quad \forall x_{t+\Delta t}, x_t \in \mathbb{R}^n, \tag{D.5}$$

then θ and θ^* are equal up to an affine transformation. That is, there exists an invertible matrix A and a vector b, such that for all $z \in \mathbb{R}^d$:

$$f(z) = f^*(Az + b), \tag{D.6}$$
$$\mu(z, t) = A^{-1}\mu^*(Az + b, t) \quad \forall t \in \mathbb{T}, \tag{D.7}$$
$$\sigma(z, t) = A^{-1}\sigma^*(Az + b, t)A^{-\top} \quad \forall t \in \mathbb{T}, \tag{D.8}$$

and

$$p_\gamma(z) = |\det A|^{-1}p_{\gamma^*}(Az + b). \tag{D.9}$$

Before we show the details of the proof, we explain why we have chosen not to include this result in the main submission. The main reason is that condition 3, in theorem statement, is not satisfied for simpler diffusion coefficients, such as a constant diffusion coefficient. We felt our theory was not satisfactory if it did not apply for a simple Brownian motion, which is one of the most simple SDEs that exist.

Other issue is related to identifiability: first we lose uniqueness up to an isometry, getting an affine transformation, and then we lose the connection between SDEs arising from the Itô's lemma. Itô's lemma implies that employing a change of variable effectively leads to another SDE, and therefore we can never estimate the true latent variable up to this change of variables. However Theorem D.10 fails to capture this. We finally present the proof of this Theorem in the supplementary materials due to its similarity to the proof of Theorem 5.

C. Proof of Theorem 8

Proof. For this proof, since it concerns recovering the latent dimension size, we assume that $\hat{A}, \mu, \sigma$ are well recovered in the sense that (14) is minimized given the conditions on the latent size in $\hat{A}$. Applying Itô's lemma, we obtain a new SDE for the transformed data. Since the transformation is linear, it is easy to characterize the distribution of the transformed space. Using the Euler-Maruyama discretization as above, we obtain a distribution on the increments $I_t = X_{t+\Delta t} - X_t$ of X_t

$$I_t \sim \mathcal{N}\left(A\mu(A^{-1}X_t, t)\Delta t, A\sigma(A^{-1}X_t, t)A^\top \Delta t\right) \quad \text{(D.10)}$$
$$Z_{t+\Delta t} - Z_t \sim \mathcal{N}\left(\mu(Z_t, t)\Delta t, \sigma(Z_t, t)\Delta t\right)$$

and recover the likelihood of **X**

$$\log \mathcal{L}(X_{t+\Delta t}|X_t, A) = -\frac{1}{2}\log \det A\sigma(A^{-1}X_t, t)A^\top \Delta t$$
$$-\frac{1}{2}\left(I_t - A\mu(A^{-1}X_t, t)\Delta t\right)\left(A\sigma(A^{-1}X_t, t)A^\top \Delta t\right)^{-1}$$
$$\left(I_t - A\mu(A^{-1}X_t, t)\Delta t\right)^\top - \frac{d}{2}\log 2\pi$$

where determinants and inverses are understood as pseudo-determinants and pseudoinverses. To recover the size of the latent dimension, we must estimate the map $\hat{A} \in \mathbb{R}^{n \times j}$ from a j dimensional latent space. We then consider the cases when $j > d$ and $j < d$ to show that at the minimum of (26), $j = d$.

$(j > d)$: Suppose $j > d$, then it is straightforward to show that $\hat{A}^\top \hat{A}$ is singular since at minimum $\text{rank}\,\hat{A} \leq d$. When $\hat{A}^\top \hat{A}$ is singular, $\mathcal{L}(\hat{A})$ is undefined since $\left(\hat{A}\sigma\hat{A}^\top\right)^{-1}$ and $\hat{A}^{-1}$ do not exist. Followingly, (26) will be undefined, therefore when $j > d$, the undefined likelihood makes j the incorrect choice.

$(j < d)$: Suppose now that $j < d$. Since $\text{rank}\,\sigma(x, t) = d$, then $\text{rank}\,A\sigma(Z_t, t)A^\top \Delta t = d$ but $\text{rank}\,\hat{A}\sigma(Z_t, t)\hat{A}^\top \Delta t = j < d$. Since $\text{rank}\,\text{Cov}(I_t) = d$ but the estimated rank of the covariance is j, $\hat{A}$ cannot achieve the maximum likelihood estimate. Then $\mathcal{L}(\hat{A}) < \mathcal{L}(A)$.

Finally, since $j \not< d$ and $j \not> d$, $j = d$. □

APPENDIX E
CALCULATING OPTIMAL ORTHOGONAL AND AFFINE TRANSFORMATIONS

Here we describe the minimizers of (27). Let A be a $d \times N$ matrix with columns given by $A_t = \tilde{f}^{-1}(X_t)$, and let B a $d \times N$ matrix with columns given by $B_t = Z_t$. Then we have (27) is equivalent to:

$$\mathcal{L}_{\text{latent}} = \frac{1}{N}\min_{Q,b}\|QA + b\mathbf{1}^\top - B\|_F^2, \quad \text{(E.1)}$$

where $\mathbf{1}$ is the all-ones vector and $\|\cdot\|_F$ denotes Frobenious norm, defined by $\|M\|_F = \sqrt{\text{tr}(M^\top M)}$. Followingly, we use the Frobenious dot-product of matrices, defined as $\langle M, O\rangle = \text{tr}(M^\top O)$. We present closed form solutions of

$$\min \|QA + b\mathbf{1}^\top - B\|_F^2$$
$$\text{s.t. } Q \in \mathcal{Q}, b \in \mathbb{R}^d \quad \text{(E.2)}$$

for both cases when $\mathcal{Q}$ is the space of square matrices and of orthogonal matrices. We first calculate b in terms of Q.

$$\|QA + b\mathbf{1}^\top - B\|_F^2$$
$$= \|QA - B\|_F^2 + \langle b\mathbf{1}^\top, b\mathbf{1}^\top - 2(QA - B)\rangle$$
$$= \|QA - B\|_F^2 + \text{tr}\left(\mathbf{1}b^\top(b\mathbf{1}^\top - 2(QA - B))\right)$$
$$= \|QA - B\|_F^2 + \text{tr}\left(Nb^\top b - 2b^\top(QA - B))\mathbf{1}\right)$$

Since this is a convex function of b, the minimum is obtained when we set the gradient to zero

$$2Nb - 2(QA - B)\mathbf{1} = \mathbf{0}$$
$$\Rightarrow b = \frac{1}{N}(QA - B)\mathbf{1} \quad \text{(E.3)}$$

Replacing b in (E.2), we get

$$\min_Q \|Q\tilde{A} - \tilde{B}\|_F^2, \quad \text{(E.4)}$$

where $\tilde{A} = A - \frac{1}{N}A\mathbf{1}\mathbf{1}^\top$ and $\tilde{B} = B - \frac{1}{N}B\mathbf{1}\mathbf{1}^\top$. We note that $\frac{1}{N}A\mathbf{1}$ and $\frac{1}{N}B\mathbf{1}$ are the average of the columns of A and B, respectively, thus $\tilde{A}$ and $\tilde{B}$ are centered versions of A and B. We first find the minimum of (E.4) for general square matrices. In this case the objective is again a convex function of the entries of Q, so we can find the minimizer by equating the gradient of Q.

$$2Q\tilde{A}\tilde{A}^\top - 2\tilde{B}\tilde{A}^\top = \mathbf{0}$$
$$\Rightarrow Q = \tilde{B}\tilde{A}^\top(\tilde{A}\tilde{A}^\top)^{-1}$$

On the other hand, if Q is orthogonal, then (E.4) is an instance of the orthogonal Procrustes problem [35], and letting $USV^\top = \tilde{B}\tilde{A}^\top$ be the singular value decomposition of $\tilde{B}\tilde{A}^\top$, the minimizer is $Q = UV^\top$. With Q calculated, we replace it in (E.3) to finally calculate b.

APPENDIX F
CRAMER-RAO BOUNDS FOR ESTIMATING THE DRIFT COEFFICIENT

Here we derive the Cramér-Rao lower bound (CRLB) we use for estimating SDEs. The CRLB gives an information theoretically lower bound on the MSE of any estimator, and in

particular it is also a lower bound for estimating μ using the proposed VAE. However, it is hard to calculate the CRLB for some of the SDEs considered, thus we focus on the simplified problem: determining μ up to a global shift. That is,

$$\mu(z) = \mu^*(z) + \eta, \quad \forall z \in \mathcal{Z},$$

where the function $\mu^*(z)$ is known but $\eta \in \mathbb{R}^d$ is unknown. We note that this is the exact CRLB for the constant drift SDE, and is still a lower bound on the estimation error for the OU process.

Using the Euler-Maruyama approximation, the increments of the SDE are distributed as $\mathcal{N}((\mu^*(z) + \eta)\Delta t, \sigma \Delta t)$. The CRLB for an estimator of the mean $\hat{\mu}$ of the Gaussian distribution $\mathcal{N}(\mu, \sigma)$ is

$$\mathrm{Cov}(\hat{\mu}) \succeq \frac{1}{N}\sigma,$$

which implies that $\mathbb{E}[\|\hat{\mu} - \mu\|^2] \geq \frac{1}{N}\mathrm{tr}\,\sigma$. Substituting the values from our increment distribution, and taking into account that $\sigma = I$, we get

$$\mathbb{E}[\|(\hat{\eta} - \eta)\Delta t\|^2] \geq \frac{\mathrm{tr}(I)\Delta t}{N}$$
$$\mathbb{E}[\|\hat{\mu}(z) - \mu(z)\|^2] \geq \frac{d}{\Delta t N}$$

References

[1] S. M. Iacus, *Simulation and inference for stochastic differential equations: with R examples*. Springer Science & Business Media, 2009.

[2] S. H. Rudy, S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Data-driven discovery of partial differential equations," *Science Advances*, vol. 3, no. 4, p. e1602614, 2017.

[3] L. Boninsegna, F. Nüske, and C. Clementi, "Sparse learning of stochastic dynamical equations," *The Journal of chemical physics*, vol. 148, no. 24, p. 241723, 2018.

[4] J. Jia and A. R. Benson, "Neural jump stochastic differential equations," *arXiv preprint arXiv:1905.10403*, 2019.

[5] C. Yildiz, M. Heinonen, J. Intosalmi, H. Mannerstrom, and H. Lahdesmaki, "Learning stochastic differential equations with Gaussian processes without gradient matching," in *2018 IEEE 28th International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, 2018, pp. 1–6.

[6] D. Hsu, S. M. Kakade, and T. Zhang, "A spectral algorithm for learning hidden Markov models," *Journal of Computer and System Sciences*, vol. 78, no. 5, pp. 1460–1480, 2012.

[7] K. Kandasamy, M. Al-Shedivat, and E. P. Xing, "Learning hmms with nonparametric emissions via spectral decompositions of continuous matrices," *arXiv preprint arXiv:1609.06390*, 2016.

[8] L. Song, B. Boots, S. Siddiqi, G. J. Gordon, and A. Smola, "Hilbert space embeddings of hidden Markov models," 2010.

[9] L. Song, A. Anandkumar, B. Dai, and B. Xie, "Nonparametric estimation of multi-view latent variable models," in *International Conference on Machine Learning*. PMLR, 2014, pp. 640–648.

[10] L. Duncker, G. Bohner, J. Boussard, and M. Sahani, "Learning interpretable continuous-time models of latent stochastic dynamical systems," in *International Conference on Machine Learning*, 2019, pp. 1726–1734.

[11] B. Tzen and M. Raginsky, "Neural stochastic differential equations: Deep latent Gaussian models in the diffusion limit," *arXiv preprint arXiv:1905.09883*, 2019.

[12] C. Yildiz, M. Heinonen, and H. Lahdesmaki, "ODE2VAE: Deep generative second order ODEs with Bayesian neural networks," in *Advances in Neural Information Processing Systems*, 2019, pp. 13 412–13 421.

[13] X. Li, T.-K. L. Wong, R. T. Chen, and D. Duvenaud, "Scalable gradients for stochastic differential equations," *arXiv preprint arXiv:2001.01328*, 2020.

[14] J.-S. Ha, Y.-J. Park, H.-J. Chae, S.-S. Park, and H.-L. Choi, "Adaptive path-integral autoencoders: Representation learning and planning for dynamical systems," *Advances in Neural Information Processing Systems*, vol. 31, pp. 8927–8938, 2018.

[15] M. Babaeizadeh, C. Finn, D. Erhan, R. H. Campbell, and S. Levine, "Stochastic variational video prediction," in *International Conference on Learning Representations*, 2018. [Online]. Available: https://openreview.net/forum?id=rk49Mg-CW

[16] M. Kumar, M. Babaeizadeh, D. Erhan, C. Finn, S. Levine, L. Dinh, and D. Kingma, "Videoflow: A conditional flow-based model for stochastic video generation," in *International Conference on Learning Representations*, 2020.

[17] J.-T. Hsieh, B. Liu, D.-A. Huang, L. F. Fei-Fei, and J. C. Niebles, "Learning to decompose and disentangle representations for video prediction," in *Advances in Neural Information Processing Systems*, 2018, pp. 517–526.

[18] B. Øksendal, "Stochastic differential equations," in *Stochastic differential equations*. Springer, 2003, pp. 65–84.

[19] H. Kunita and S. Watanabe, "On square integrable martingales," *Nagoya Mathematical Journal*, vol. 30, pp. 209–245, 1967.

[20] Y. Aït-Sahalia, "Closed-form likelihood expansions for multivariate diffusions," *The Annals of Statistics*, vol. 36, no. 2, pp. 906–937, 2008.

[21] ——, "Maximum likelihood estimation of discretely sampled diffusions: a closed-form approximation approach," *Econometrica*, vol. 70, no. 1, pp. 223–262, 2002.

[22] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," *arXiv preprint arXiv:1312.6114*, 2013.

[23] I. Khemakhem, D. P. Kingma, and A. Hyvärinen, "Variational autoencoders and nonlinear ICA: A unifying framework," *arXiv preprint arXiv:1907.04809*, 2019.

[24] G. O. Roberts and O. Stramer, "On inference for partially observed nonlinear diffusion models using the Metropolis–Hastings algorithm," *Biometrika*, vol. 88, no. 3, pp. 603–621, 2001.

[25] S. Kou, B. P. Olding, M. Lysy, and J. S. Liu, "A multiresolution method for parameter estimation of diffusion processes," *Journal of the American Statistical Association*, vol. 107, no. 500, pp. 1558–1574, 2012.

[26] J. Hull and A. White, "Pricing interest-rate-derivative securities," *The review of financial studies*, vol. 3, no. 4, pp. 573–592, 1990.

[27] F. Black, E. Derman, and W. Toy, "A one-factor model of interest rates and its application to treasury bond options," *Financial analysts journal*, vol. 46, no. 1, pp. 33–39, 1990.

[28] J. Solomon, F. De Goes, G. Peyré, M. Cuturi, A. Butscher, A. Nguyen, T. Du, and L. Guibas, "Convolutional Wasserstein distances: Efficient optimal transportation on geometric domains," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 4, pp. 1–11, 2015.

[29] S. A. Nene, S. K. Nayar, H. Murase *et al.*, "Columbia object image library (coil-20)," 1996.

[30] S. Lameh, L. Ding, and D. Stein, "Controlled amplification of dna brownian motion using electrokinetic noise," *Physical Review Applied*, vol. 14, no. 5, p. 054042, 2020.

[31] K. Zuiderveld, "Contrast limited adaptive histogram equalization," *Graphics gems*, pp. 474–485, 1994.

[32] M. H. Davis and S. I. Marcus, "An introduction to nonlinear filtering," in *Stochastic systems: The mathematics of filtering and identification and applications*. Springer, 1981, pp. 53–75.

[33] P. Sorrenson, C. Rother, and U. Köthe, "Disentanglement by nonlinear ICA with general incompressible-flow networks (GIN)," *arXiv preprint arXiv:2001.04872*, 2020.

[34] J. M. Lee, "Smooth manifolds," in *Introduction to Smooth Manifolds*. Springer, 2013.

[35] J. C. Gower, G. B. Dijksterhuis *et al.*, *Procrustes problems*. Oxford University Press on Demand, 2004, vol. 30.

[36] B. B. Avants, N. J. Tustison, M. Stauffer, G. Song, B. Wu, and J. C. Gee, "The insight toolkit image registration framework," *Frontiers in neuroinformatics*, vol. 8, p. 44, 2014.

APPENDIX G
SUPPLEMENTAL PROOFS

A. Proof of Lemma D.9

Proof. First, (24) implies that for all j, k

$$\frac{\partial}{\partial y_k}\left(\sigma(y,t)^{-1}\right)e_j = -\sigma(y,t)^{-1}\frac{\partial\sigma(y,t)}{\partial y_k}\sigma(y,t)^{-1}e_j$$
$$= -\sigma(y,t)^{-1}\frac{\partial\sigma(y,t)}{\partial y_j}\sigma(y,t)^{-1}e_k,$$
$$= \frac{\partial}{\partial y_j}\left(\sigma(y,t)^{-1}\right)e_k,$$

and (i) follows from [34, Theorem 11.49]. We now prove that $h(\cdot, t)$ is an one-to-one function, thus proving its inverse $g(\cdot, t)$ exists and is well-defined.

We start by showing that for all t the set $\mathcal{H}_M := \{y : \|h(y,t)\| \le M\}$ is compact. Since h is continuous, $\mathcal{H}_M$ is a closed set, thus it remains to prove it is bounded. Condition (i) in Theorem 2 implies not only that $\sigma(y,t)$ is positive definite but also that its norm is bounded by $\|y\|$.

$$\|\sigma(y,t)\| \le \|\sigma(y,t) - \sigma(0,t)\| + \|\sigma(0,t)\|, \qquad \text{(G.1)}$$
$$\le D\|y\| + \|\sigma(0,t)\|, \qquad \text{(G.2)}$$
$$\le \tilde{D}(1 + \|y\|), \qquad \text{(G.3)}$$

where $\tilde{D} = \max\{D, \max_{t \in \mathbb{T}} \|\sigma(0,t)\|\}$. This implies the eigenvalues of $\sigma(y,t)$ are upper-bounded by $(1 + \|y\|)\tilde{D}$, and since $\sigma(y,t)$ is a PSD matrix, the eigenvalues of its inverse are lower bounded by $1/((1 + \|y\|)\tilde{D})$, thus

$$w^T\sigma(y,t)^{-1}w \ge \frac{\|w\|^2}{(1 + \|y\|)\tilde{D}} \quad \text{for all} \quad y, w \in \mathbb{R}^d. \quad \text{(G.4)}$$

Applying this and the Cauchy-Schwarz inequality,

$$\|y\|\|h(y,t)\| \ge y^T h(y,t)$$
$$= \int_0^1 y^T \sigma(\tau y, t)^{-1} y \, d\tau,$$
$$\ge \int_0^1 \frac{\|y\|^2}{(1 + \tau\|y\|)\tilde{D}} \, d\tau,$$
$$= \frac{\|y\|}{\tilde{D}} \log(1 + \|y\|).$$

Therefore, if $M \ge \|h(y,t)\|$, then $M \ge \frac{1}{\tilde{D}} \log(1 + \|y\|)$, thus

$$\|y\| \le \exp(M\tilde{D}) - 1,$$

which proves $\mathcal{H}_M$ is bounded. We now prove h is surjective. Let $z \in \mathbb{R}^d$ and

$$y_* = \underset{y \in \mathcal{H}_{1+2\|z\|}}{\arg\min} \|h(y,t) - z\|^2. \qquad \text{(G.5)}$$

Since $\mathcal{H}_{1+2\|z\|}$ is compact, the minimum is achieved by a point in $\mathcal{H}_{1+2\|z\|}$, and this point is not in the boundary. Since by definition $h(0,t) = 0$, we have $0 \in \mathcal{H}_{1+2\|z\|}$ and for any element $w \in \partial\mathcal{H}_{1+2\|z\|}$

$$\|h(w,t) - z\| \ge \|h(w,t)\| - \|z\|,$$
$$= 2\|z\| + 1 - \|z\|,$$
$$> \|z\| = \|h(0,t) - z\|,$$

thus w does not achieve the minimum. Since the minimum is achieved by an interior point of $\mathcal{H}_{1+2\|z\|}$ and h is differentiable, y_* is a critical point of (G.5). We then have

$$0 = \nabla\|h(y_*,t) - z\|^2,$$
$$= 2\frac{\partial h}{\partial y}(y_*,t)(h(y_*,t) - z),$$
$$= 2\sigma(y_*,t)^{-1}(h(y_*,t) - z).$$

Since $\sigma(y_*,t)^{-1}$ is non-singular, we must have $h(y_*,t) = z$, thus h is surjective. We now prove that $h(y,t) \ne h(w,t)$ for all $y, w \in \mathbb{R}^d$. By the Fundamental Theorem of Calculus

$$(w - y)^T (h(w,t) - h(y,t))$$
$$= \int_0^1 (w - y)^T \frac{\partial h}{\partial y}(y + \tau(w - y), t)(w - y) \, d\tau,$$
$$= \int_0^1 (w - y)^T \sigma(y + \tau(w - y), t)^{-1}(w - y) \, d\tau$$
$$> 0,$$

where the last line follows from (G.4), thus $h(y,t) \ne h(w,t)$, h is bijective, and g in (ii) is well defined. Finally, (iii) follows from the Inverse Function Theorem

$$\frac{\partial g}{\partial y}(h(y,t),t) = \frac{\partial h}{\partial y}(y,t)^{-1},$$
$$= \sigma(y,t).$$

$\square$

B. Proof of Theorem D.10

Proof. The proof is again very similar to the proof of Theorem 5. The 3 steps are the same as in the previous proof, and step (I) follows exactly in the same way, so we start on step (II).

Step (II): From Step (I), we have that for all $x_t, x_{t+\Delta t}$ in $\mathcal{X}$,

$$\log \text{vol} \, J_{f_*^{-1}}(x_{t+\Delta t}) + \log \text{vol} \, J_{f_*^{-1}}(x_t)$$
$$+ \log p_{\mu_*, \sigma_*}(f_*^{-1}(x_{t+\Delta t})|f_*^{-1}(x_t), t) + \log p_{\gamma_*}(f_*^{-1}(x_t))$$
$$= \log \text{vol} \, J_{f^{-1}}(x_{t+\Delta t}) + \log \text{vol} \, J_{f^{-1}}(x_t)$$
$$+ \log p_{\mu, \sigma}(f^{-1}(x_{t+\Delta t})|f^{-1}(x_t), t) + \log p_\gamma(f^{-1}(x_t)).$$
$$\text{(G.6)}$$

Let

$$\lambda(x_t) = f^{-1}(x_t) - \mu(f^{-1}(x_t), t)\Delta t, \qquad \text{(G.7)}$$
$$\eta(x_t) = \sigma(f^{-1}(x_t), t)^{-1}\lambda(x_t), \qquad \text{(G.8)}$$

and define analogously $\lambda_*(x_t), \eta_*(x_t)$. Let $z_t = f^{-1}(x_t)$, and similarly for $z_{t+\Delta t}$. We then have

$$\log p_{\mu, \sigma}(z_{t+\Delta t}|z_t, t)$$
$$= -\frac{1}{2} \log \det \sigma(z_t, t)$$
$$- \frac{1}{2\Delta t}(z_{t+\Delta t} - \lambda(x_t))^T \sigma(z_t, t)^{-1}(z_{t+\Delta t} - \lambda(x_t)),$$
$$= -\frac{1}{2} \log \det \sigma(z_t, t)$$
$$- \frac{1}{2\Delta t}\left(z_{t+\Delta t}^T \sigma(z_t, t)^{-1} z_{t+\Delta t} - 2z_{t+\Delta t}^T \eta(x_t) + \lambda(x_t)^T \eta(x_t)\right).$$

We now write this as an augmented linear system. Let $\mathcal{A}(z) = \mathrm{vsym}(zz^T) \oplus z$,

$$\xi(x_t) = \mathrm{vsym}(\sigma(f^{-1}(x_t), t)^{-1}) \oplus (-2\eta(x_t)), \qquad \text{(G.9)}$$

and define $\xi_*(x_t)$ analogously. We have

$$z_{t+\Delta t}^T \sigma(z_t, t)^{-1} z_{t+\Delta t}$$
$$= \left\langle \mathrm{vsym}(z_{t+\Delta t} z_{t+\Delta t}^T), \mathrm{vsym}(\sigma(z_t, t)^{-1}) \right\rangle,$$

thus

$$\log p_{\mu,\sigma}(z_{t+\Delta t}|z_t, t)$$
$$= -\frac{1}{2} \log \det \sigma(z_t, t)$$
$$- \frac{1}{2\Delta t} \left(\langle \mathcal{A}(z_{t+\Delta t}), \xi(x_t) \rangle + \lambda(x_t)^T \eta(x_t) \right),$$
$$= -\frac{1}{2} \log \det \sigma(f^{-1}(x_t), t)$$
$$- \frac{1}{2\Delta t} \left(\langle \mathcal{A}(f^{-1}(x_{t+\Delta t})), \xi(x_t) \rangle + \lambda(x_t)^T \eta(x_t) \right).$$

Finally define

$$\psi(x_t) = \mathrm{vol}\, J_{f^{-1}}(x_t) + \log p_\gamma(f^{-1}(x_t))$$
$$- \frac{1}{2} \log \det \sigma(f^{-1}(x_t), t) - \frac{\langle \lambda(x_t), \eta(x_t) \rangle}{2\Delta t}, \quad \text{(G.10)}$$

$$\varphi(x_{t+\Delta t}) = \mathrm{vol}\, J_{f^{-1}}(x_{t+\Delta t}), \qquad \text{(G.11)}$$

and ψ_* and φ_* analogously. We finally have for all $x_t, x_{t+\Delta t}$ in $\mathcal{X}$, an equation similar to (C.5).

$$\psi(x_t) - \frac{1}{2\Delta t} \left\langle \mathcal{A}(f^{-1}(x_{t+\Delta t})), \xi(x_t) \right\rangle + \varphi(x_{t+\Delta t})$$
$$= \psi_*(x_t) - \frac{1}{2\Delta t} \left\langle \mathcal{A}(f_*^{-1}(x_{t+\Delta t})), \xi_*(x_t) \right\rangle + \varphi_*(x_{t+\Delta t}).$$
$$\text{(G.12)}$$

We now proceed in the same way as before to obtain a similar equality for all $x_{t+\Delta t} \in \mathcal{X}$.

$$\mathcal{A}(f_*^{-1}(x_{t+\Delta t})) = \mathbf{\Xi}_*^{-T} \mathbf{\Xi}^T \mathcal{A}(f^{-1}(x_{t+\Delta t})) + \mathbf{\Xi}_*^{-T} \psi,$$
$$\text{(G.13)}$$

with the exception that the invertibility of $\mathbf{\Xi}_*^T$ is now a consequence of condition 3 in the theorem statement. To prove $\mathbf{\Xi}_*^{-T} \mathbf{\Xi}$ is invertible, let $z_0, z_1, \ldots, z_N \in \mathbb{R}^d$, with N defined as in the theorem statement, such that the vectors $\mathcal{A}(z_i) - \mathcal{A}(z_0)$, $i = 1, \ldots, N$ are linearly independent. We note that the entries of $\mathcal{A}(z)$ are linear independent polynomials so finding such vectors is always possible. Let $x_i = f_*(z_i)$, $i = 0, \ldots, N$, define $\mathbf{A}_*$ as the $N \times N$ matrix with columns $\mathbf{A}_{*,i}$

$$\mathbf{A}_{*,i} = \mathcal{A}(f_*^{-1}(x_i)) - \mathcal{A}(f_*^{-1}(x_0)),$$

and define $\mathbf{A}$ analogously. Equation (G.13) implies

$$\mathbf{A}_* = \mathbf{\Xi}_*^{-T} \mathbf{\Xi}^T \mathbf{A}. \qquad \text{(G.14)}$$

Since $\mathbf{A}_*$ is invertible by construction, so is $\mathbf{\Xi}_*^{-T} \mathbf{\Xi}^T$. From (G.13), we can now get a linear dependence between f_*^{-1} and f^{-1}. Let $\nu(z) = f_*^{-1}(f(z))$, then (G.13) implies that there are symmetric matrices M_i, vectors v_i and scalars c_i such that

$$\nu(z)_i = z^T M_i z + v_i^T z + c_i, \qquad \text{(G.15)}$$

and that there are symmetric matrices $\tilde{M}_i$, vectors $\tilde{v}_i$ and scalars $\tilde{c}_i$ such that

$$\nu(z)_i^2 = z^T \tilde{M}_i z + \tilde{v}_i^T z + \tilde{c}_i. \qquad \text{(G.16)}$$

Subtracting the square of (G.15) with the (G.16) for all z implies that $M_i = 0$, $\tilde{M}_i = v_i v_i^T$, $\tilde{v}_i = 2c_i v_i$ and $\tilde{c}_i = c_i^2$. This and the invertibility of $\mathbf{\Xi}_*^{-T} \mathbf{\Xi}^T$ finally imply that there is an invertible matrix A and a vector b such that for all $x_{t+\Delta t}$

$$f_*^{-1}(x_{t+\Delta t}) = A f^{-1}(x_{t+\Delta t}) + b. \qquad \text{(G.17)}$$

Step (III): Equation (G.17) implies that for all $z \in \mathcal{Z}$

$$f(z) = f_*(f_*^{-1}(f(z))),$$
$$= f_*(A f^{-1}(f(z)) + b),$$
$$= f_*(Az + b).$$

Moreover $J_f = J_{f_*} A$, and $\log \mathrm{vol}\, J_{f^{-1}} = \log \mathrm{vol}\, J_{f_*^{-1}} + \log |\det A|$. Now replacing $x_{t+\Delta t}$ by $f(z)$ on (G.12) and using (G.17) and (G.10), we get for all $z \in \mathcal{Z}$ and $x_t \in \mathcal{X}$

$$\psi(x_t) - \frac{1}{\Delta t} \langle z, \eta(x_t) \rangle + \frac{z^T \sigma(f^{-1}(x_t), t)^{-1} z}{2\Delta t} + \log |\det A|$$
$$= \psi_*(x_t) - \frac{1}{\Delta t} \langle Az + b, \eta_*(x_t) \rangle$$
$$+ \frac{(Az + b)^T \sigma_*(f_*^{-1}(x_t), t)^{-1} (Az + b)}{2\Delta t} \qquad \text{(G.18)}$$

Taking derivatives with respect to z, we get

$$- \eta(x_t) + \sigma(f^{-1}(x_t), t)^{-1} z$$
$$= -A^T \eta_*(x_t) + A^T \sigma_*(f_*^{-1}(x_t), t)^{-1} (Az + b) \qquad \text{(G.19)}$$

Taking derivatives again we get

$$\sigma(f^{-1}(x_t), t)^{-1} = A^T \sigma_*(f_*^{-1}(x_t), t)^{-1} A \qquad \text{(G.20)}$$

Letting $f(z) = x_t$ and using (G.17), we get for all $z \in \mathcal{Z}$,

$$\sigma(z, t) = A^{-1} \sigma_*(Az + b, t) A^{-T} \qquad \text{(G.21)}$$

Replacing this back in (G.19), using (G.8) and multiplying by the inverse of (G.20), we get

$$\lambda(x_t) = A^{-1} \lambda_*(x_t) - A^{-1} b, \qquad \text{(G.22)}$$

or $\lambda_*(x_t) = A\lambda(x_t) + b$. Using (G.20), (G.7), letting $f(z) = x_t$ and using (C.13), we get for all $z \in \mathcal{Z}$,

$$Az + b - \mu_*(Az + b, t)\Delta t = Az - A\mu(z, t)\Delta t + b \qquad \text{(G.23)}$$

which implies $\mu(z, t) = A^{-1}\mu_*(Az + b, t)$ for all $z \in \mathcal{Z}$. Finally, replacing all equations obtained in (G.18), and using (G.10), we get

$$\log p_\gamma(f^{-1}(x_t)) = \log p_{\gamma_*}(f_*^{-1}(x_t)) - \log |\det A|, \quad \text{(G.24)}$$

and (D.9) follows from taking $x_t = f(z)$ and using (G.17). $\qquad \square$

Dataset	Balls	Digits	Wasserstein	Balls + Wasserstein	DNA
Batch Size			Full Trajectory		
	(800)	(100)	(800)	(800)	(100)
Validation Size			100		
AE LR			0.001		0.0001
$\hat{\mu}$ LR			0.001		0.0001
$\hat{\mu}$ Width			16		
$\hat{\mu}$ Depth			4		
$\hat{\mu}$ Activation			Softplus		
τ	0.01	0.005	0.05	0.01	0.01
ν			0		
Optimizer			Adam		
LR Decay			Exponential		
	0.997	0.997	0.999	0.999	0.9998
Epochs		1500		2000	150

TABLE II: Hyperparameters for all experiments.

Encoder

Input size: (3, 64, 64)
Conv(3, 8, 5, 1, 2), LeakyReLU, MaxPool(2, 2, 0)
Conv(8, 16, 5, 1, 2), BN, LeakyReLU, MaxPool(2, 2, 0)
Conv(16, 32, 5, 1, 2), BN, LeakyReLU, MaxPool(2, 2, 0)
Conv(32, 64, 5, 1, 2), BN, LeakyReLU
Flatten
$\text{FC}_\mu(4096, d)$, $\text{FC}_\sigma(4096, d)$
Output size : $(d,)$

TABLE III: Encoder architectures for all experiments. BN refers to batch normalization.

Decoder

Input size: $(d,)$
FC(d, 4096)
Unflatten
U, Conv(64, 64, 5, 1, 2), BN, LeakyReLU
U, Conv(64, 32, 5, 1, 2), BN, LeakyReLU
U, Conv(32, 16, 5, 1, 2), BN, LeakyReLU
U, Conv(16, 8, 5, 1, 2), BN, LeakyReLU
U, Conv(8, 3, 5, 1, 2), Sigmoid
Output size : (3, 64, 64)

TABLE IV: Decoder architecture for all experiments. U refers to an upsampling layer.

APPENDIX H
HYPERPARAMETERS

The list of hyperparameters used for the experiments are available in Table II. Architecture details for all experiments are given in Tables III and IV. For each run, we run the algorithm 3 times on the same run and choose the one with the lowest validation loss. We performed limited hyperparameter tuning, instead leaving most parameters the same across all datasets. For the noise datasets, we use the same hyperparameters as without the noise.

APPENDIX I
FLUORESCENT DNA DETAILS

The original data are given as spatially varying raw counts over period of 100 frames per video. Let $V \in \mathbb{N}^{100 \times 512 \times 512}$

represent a video with 100 frames with 512 pixels as the width and height. We define

$$\bar{V} = V - \frac{1}{100} \sum_i V_{i,j,k}$$

and

$$\tilde{V} = \frac{\bar{V} - \min \bar{V}}{\max \bar{V} - \min \bar{V}}.$$

After normalization, we pass the frames through a 4×4 Gaussian filter with $\sigma = 3$ and compute a new image using a maximum filter over 3×3 blocks of the image. From this, we compute the maximum value corresponding to the center of the molecule. We then take the average of all the centers as a guideline for a refinement iteration where we again compute the centers conditioned on the mean center.

For the input to the neural network, we first compute

$$\hat{V} = \frac{V - \min V}{\max V - \min V}.$$

and then apply the adaptive histogram equalization [31] algorithm to each frame in the video.

APPENDIX J
COMPARISON WITH [10]

We add a comparison where we apply the method proposed in [10] to the true latent SDE. Since [10] has the advantage of observing the true latent SDE realization, rather than the high dimensional ambient space observations, the method should act as a lower bound to the proposed method. The results are reported in Table V along with the Cramér-Rao lower bounds previously obtained and p-values on a two sample t-test between the statistics of the two methods. In this case, we see that the method in [10] performs very well in low dimensional cases with easier SDEs (namely the 1D and 2D OU and the constant) but does worse in the other cases. Moreover, when considering the statistical significance of the differences between the methods, the method by [10] exhibits statistical significance in the low dimensional OU processes (for $p < 0.01$) and in the 4D constant case (for $p < 0.05$). The proposed method exhibits statistical significance for the Circle and Double Well experiments (for $p < 0.01$). This provides greater evidence of the efficacy of the proposed method since, with the much more difficult estimation task

Dataset	SDE Type	Oracle + [10]	Proposed	CRLB	p-value
	Constant	$2.37(\pm0.10)\times10^{-1}$	$5.07(\pm2.99)\times10^{-1}$		0.0783
Balls	OU	$2.17(\pm0.10)\times10^{-1}$	$9.20(\pm2.23)\times10^{-1}$	2.00×10^{-1}	0.0001
	Circle	$4.43(\pm0.74)\times10^{0}$	$1.47(\pm0.62)\times10^{0}$		0.0001
Multiple Balls	OU	$2.74(\pm0.51)\times10^{0}$	$2.61(\pm0.51)\times10^{0}$	1.00×10^{0}	0.6975
	Constant	$1.92(\pm0.49)\times10^{-1}$	$7.74(\pm4.65)\times10^{-1}$		0.0238
Digits	OU	$5.66(\pm0.42)\times10^{-1}$	$7.69(\pm4.91)\times10^{-1}$	4.00×10^{-1}	0.3839
	Circle	$2.96(\pm0.22)\times10^{0}$	$2.15(\pm0.48)\times10^{0}$		0.0089
	OU	$1.54(\pm0.48)\times10^{-1}$	$8.79(\pm1.21)\times10^{-1}$		< 0.0001
Wasserstein	Double Well	$5.37(\pm2.45)\times10^{0}$	$1.45(\pm0.76)\times10^{0}$	1.00×10^{-1}	0.0091
	GBM	–	$1.23(\pm2.45)\times10^{-1}$		–
	OU	$5.97(\pm3.07)\times10^{-1}$	$7.08(\pm2.79)\times10^{-1}$		0.5661
Ball + Wasserstein	Cauchy	$5.09(\pm0.04)\times10^{-1}$	$8.82(\pm7.36)\times10^{-1}$	3.00×10^{-1}	0.2899
	Anisotropic	-	$1.74(\pm0.50)\times10^{0}$		-

TABLE V: Comparison of the MSE in estimating the drift coefficient, $\mathcal{L}_\mu$, defined in (28), between the proposed method with the benchmark of [10], across different datasets.

DNA Dataset	$V = 0$	$V = 1$	$V = 2$
ITK	2.58×10^{0}	4.27×10^{0}	8.28×10^{0}
Proposed	7.93×10^{-2}	1.87×10^{-1}	2.03×10^{-1}

TABLE VI: MSE between estimated displacements and approximate coordinates estimated by intensity segmentation.

from the ambient space, the proposed method performs within one order of magnitude in error in all cases and better with higher dimensional latent space.

APPENDIX K
COMPARISON WITH IMAGE REGISTRATION TECHNIQUES

We add a comparison between our method and image registration techniques in estimating the molecules displacement in the Fluorescent DNA datasets. We used the `TranslationTransform` method in the Insight ToolKit (ITK) [36] to estimate the displacement between frames. The results of the MSE between the estimated displacement from the registration algorithm and the coordinates approximated from the intensity based particle tracking are presented in Table VI. We observe that the proposed method outperforms the image registration method in the DNA datasets. We believe this may be due to noise, which makes tracking the position of the molecule difficult. Additional preprocessing may be necessary to achieve better results with image registration techniques.