
Continuous Latent Process Flows

Ruizhi Deng^{1,2*} **Marcus A. Brubaker**^{1,3,4} **Greg Mori**^{1,2} **Andreas M. Lehrmann**¹
¹Borealis AI ²Simon Fraser University ³York University ⁴Vector Institute

Abstract

Partial observations of continuous time-series dynamics at arbitrary time stamps exist in many disciplines. Fitting this type of data using statistical models with continuous dynamics is not only promising at an intuitive level but also has practical benefits, including the ability to generate continuous trajectories and to perform inference on previously unseen time stamps. Despite exciting progress in this area, the existing models still face challenges in terms of their representation power and the quality of their variational approximations. We tackle these challenges with continuous latent process flows (CLPF), a principled architecture decoding continuous latent processes into continuous observable processes using a time-dependent normalizing flow driven by a stochastic differential equation. To optimize our model using maximum likelihood, we propose a novel piecewise construction of a variational posterior process and derive the corresponding variational lower bound using importance weighting of trajectories. An ablation study demonstrates the effectiveness of our contributions and comparisons to state-of-the-art baselines show our model’s favourable performance on both synthetic and real-world data.

1 Introduction

Sparse and irregular observations of continuous dynamics are common in many areas of science, including finance [15, 36], healthcare [16], and physics [30]. Time-series models driven by stochastic differential equations (SDEs) provide an elegant framework for this challenging scenario and have recently gained popularity in the machine learning community [11, 18, 24]. The SDEs are typically implemented by neural networks with trainable parameters and the latent processes defined by the SDEs are then decoded into an observable space with complex structure. Due to the lack of closed-form transition densities for most SDEs, dedicated variational approximations have been developed to maximize the observational log-likelihoods [2, 18, 24].

Despite the progress of existing works, challenges still remain for SDE-based models to be applied to irregular time-series data. One major challenge is the model’s representation power. The continuous-time flow process (CTFP; [11]) uses a series of invertible mappings continuously indexed by time to transform a simple Wiener process to a more complex stochastic process. The use of a simple latent process and invertible transformations permits CTFP models to evaluate the exact likelihood of observations on any time grid efficiently, but they also limit the set of stochastic processes that CTFP can express to some specific form which can be obtained using Ito’s Lemma and excludes many commonly seen stochastic processes. Another constraint of representation power in practice is the Lipschitz property of transformations in the models. The latent SDE model proposed by Hasan et al. [18] and CTFP both transform a latent stochastic process with constant variance to an observable one using injective mappings. Due to the Lipschitz property existing in many invertible neural network architectures, some processes that can be written as a non-Lipschitz transformation of a simple process, like geometric Brownian motion, cannot be expressed by these models unless specific choices of non-Lipschitz decoders are used. Apart from the model’s representation power, variational inference is another challenge in training SDE-based models. The latent SDE model in the work of Li et al. [24]

*This work was done during an internship at Borealis AI. Correspondance to wsdmdeng@gmail.com.




Figure 1: **Overview.** Our architecture uses a stochastic differential equation (SDE; left) to drive a time-dependent normalizing flow (NF; right). At time t_1, t_2 (grey bars), the values of the SDE trajectories (colored trajectories on the left) serve as conditioning information for the decoding of a simple base process (grey trajectories on the right) into a complex observable process (colored trajectories on the right). The color gradient on the right shows the individual trajectories of this transformation, which is driven by an augmented neural ODE. Since all stochastic processes and mappings are time-continuous, we can model observed data as partial realizations of a continuous process, enabling modelling of continuous dynamics and inference on irregular time grids.

uses a principled method of variational approximation based on importance weighting of trajectories between a variational posterior and a prior process. The variational posterior process is constructed using a single SDE conditioned on all observations and is therefore restricted to be a Markov process. This approach may lack the flexibility to approximate the true posterior process well enough in complex inference tasks, e.g., in an online setting or with variable-length observation sequences.

In this work we propose Continuous Latent Process Flows (CLPFs)², a model that is governed by latent dynamics defined by an expressive generic stochastic differential equation. Inspired by [11], we then use dynamic normalizing flows to decode each latent trajectory into a continuous observable process. Driven by different trajectories of the latent stochastic process continuously evolving with time, the dynamic normalizing flow can map a simple base process to a diverse class of observable processes. We illustrate this process in Fig. 1. This decoding is critical for the model to generate continuous trajectories and be trained to fit observations on irregular time grids using a variational approximation. Good variational approximations and proper handling of complex inference tasks like online inference depend on a flexible variational posterior process. Therefore, we also propose a principled method of defining and sampling from a non-Markov variational posterior process that is based on a piecewise evaluation of SDEs and can adapt to new observations. The proposed model excels at fitting observations on irregular time grids, generalizing to observations on more dense time grids, and generating trajectories continuous in time.

Contributions. In summary, we make the following contributions: (1) We propose a flow-based decoding of a generic SDE as a principled framework for continuous dynamics modeling of irregular time-series data. (2) We improve the variational approximation of the observational likelihood through a flexible non-Markov posterior process based on a piecewise evaluation of the underlying SDE; (3) We validate the effectiveness of our contributions in a series of ablation studies and comparisons to state-of-the-art time-series models, both on synthetic and real-world datasets.

2 Preliminaries

2.1 Stochastic Differential Equations

SDEs can be viewed as a stochastic analogue of ordinary differential equations (ODEs) in the sense that $\frac{d\mathbf{Z}_t}{dt} = \boldsymbol{\mu}(\mathbf{Z}_t, t) + \text{random noise} \cdot \boldsymbol{\sigma}(\mathbf{Z}_t, t)$. Let \mathbf{Z}_t be a variable which continuously evolves with time. An m -dimensional SDE describing the stochastic dynamics of \mathbf{Z}_t usually takes the form

$$d\mathbf{Z}_t = \boldsymbol{\mu}(\mathbf{Z}_t, t) dt + \boldsymbol{\sigma}(\mathbf{Z}_t, t) d\mathbf{W}_t, \quad (1)$$

where $\boldsymbol{\mu}$ maps to an m -dimensional vector, $\boldsymbol{\sigma}$ is an $m \times k$ matrix, and \mathbf{W}_t is a k -dimensional Wiener process. The solution of an SDE is a continuous-time stochastic process \mathbf{Z}_t that satisfies the integral equation $\mathbf{Z}_t = \mathbf{Z}_0 + \int_0^t \boldsymbol{\mu}(\mathbf{Z}_s, s) ds + \int_0^t \boldsymbol{\sigma}(\mathbf{Z}_s, s) d\mathbf{W}_s$ with initial condition \mathbf{Z}_0 , where

²Code available at <https://github.com/BorealisAI/continuous-latent-process-flows>

the stochastic integral should be interpreted as a traditional Itô integral [27, Chapter 3.1]. For each sample trajectory $\omega \sim \mathbf{W}_t$, the stochastic process \mathbf{Z}_t maps ω to a different trajectory $\mathbf{Z}_t(\omega)$.

Latent Dynamics and Variational Bound. SDEs have been used as models of latent dynamics in a variety of contexts [2, 18, 24]. As closed-form finite-dimensional solutions to SDEs are rare, variational approximations are often used in practice. Li et al. [24] propose a principled way of re-weighting latent SDE trajectories for variational approximations using Girsanov’s theorem [27, Chapter 8.6]. Specifically, consider a prior process and a variational posterior process in the interval $[0, T]$ defined by two stochastic differential equations $d\mathbf{Z}_t = \boldsymbol{\mu}_1(\mathbf{Z}_t, t) dt + \boldsymbol{\sigma}(\mathbf{Z}_t, t) d\mathbf{W}_t$ and $d\hat{\mathbf{Z}}_t = \boldsymbol{\mu}_2(\hat{\mathbf{Z}}_t, t) dt + \boldsymbol{\sigma}(\hat{\mathbf{Z}}_t, t) d\mathbf{W}_t$, respectively. Furthermore, let $p(\mathbf{x}|\mathbf{Z}_t)$ denote the probability of observing \mathbf{x} conditioned on the trajectory of the latent process \mathbf{Z}_t in the interval $[0, T]$. If there exists a mapping $\mathbf{u} : \mathbb{R}^m \times [0, T] \rightarrow \mathbb{R}^k$ such that $\boldsymbol{\sigma}(\mathbf{z}, t)\mathbf{u}(\mathbf{z}, t) = \boldsymbol{\mu}_2(\mathbf{z}, t) - \boldsymbol{\mu}_1(\mathbf{z}, t)$ and \mathbf{u} satisfies Novikov’s condition [27, Chapter 8.6], we obtain the variational lower bound

$$\log p(\mathbf{x}) = \log \mathbb{E}[p(\mathbf{x}|\mathbf{Z}_t)] = \log \mathbb{E}[p(\mathbf{x}|\hat{\mathbf{Z}}_t)\mathbf{M}_T] \geq \mathbb{E}[\log p(\mathbf{x}|\hat{\mathbf{Z}}_t) + \log \mathbf{M}_T], \quad (2)$$

where $\mathbf{M}_T = \exp(-\int_0^T \frac{1}{2} \|\mathbf{u}(\hat{\mathbf{Z}}_t, t)\|^2 dt - \int_0^T \mathbf{u}(\hat{\mathbf{Z}}_t, t)^T d\mathbf{W}_t)$. See [24] for a formal proof.

2.2 Normalizing Flows

Normalizing flows [3, 8, 12, 13, 21, 22, 23, 28, 29, 31] employ a bijective mapping $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ to transform a random variable \mathbf{Y} with a simple base distribution $p_{\mathbf{Y}}$ to a random variable \mathbf{X} with a complex target distribution $p_{\mathbf{X}}$. We can sample from a normalizing flow by first sampling $\mathbf{y} \sim p_{\mathbf{Y}}$ and then transforming it to $\mathbf{x} = f(\mathbf{y})$. As a result of invertibility, normalizing flows can also be used for density estimation. Using the change-of-variables formula, we have $\log p_{\mathbf{X}}(\mathbf{x}) = \log p_{\mathbf{Y}}(g(\mathbf{x})) + \log \left| \det \left(\frac{\partial g}{\partial \mathbf{x}} \right) \right|$, where g is the inverse of f .

Continuous Indexing. More recently, normalizing flows have been augmented with a continuous index [6, 10, 11]. For instance, the continuous-time flow process (CTFP; [11]) models irregular observations of a continuous-time stochastic process. Specifically, CTFP transforms a simple d -dimensional Wiener process \mathbf{W}_t to another continuous stochastic process \mathbf{X}_t using the transformation $\mathbf{X}_t = f(\mathbf{W}_t, t)$, where $f(\mathbf{w}, t)$ is an invertible mapping for each t . Despite its benefits of exact log-likelihood computation of arbitrary finite-dimensional distributions, the expressive power of CTFP to model stochastic processes is limited in the following two aspects: (1) An application of Itô’s lemma [27, Chapter 4.2] shows that CTFP can only represent stochastic processes of the form

$$d\mathbf{f}(\mathbf{W}_t, t) = \left\{ \frac{d\mathbf{f}}{dt}(\mathbf{W}_t, t) + \frac{1}{2} \text{Tr}(\mathbf{H}_{\mathbf{w}}\mathbf{f}(\mathbf{W}_t, t)) \right\} dt + (\nabla_{\mathbf{w}}\mathbf{f}^T(\mathbf{W}_t, t))^T d\mathbf{W}_t, \quad (3)$$

where $\mathbf{H}_{\mathbf{w}}\mathbf{f}$ is the Hessian matrix of \mathbf{f} with respect to \mathbf{w} and $\nabla_{\mathbf{w}}\mathbf{f}$ is the derivative. A variety of stochastic processes, from simple processes like the Ornstein-Uhlenbeck (OU) process to more complex non-Markov processes, fall outside of this limited class and cannot be learned using CTFP (see Appendix A for formal proofs); (2) Many normalizing flow architectures are compositions of Lipschitz-continuous transformations [7, 8, 17]. It is therefore challenging to model certain stochastic processes that are non-Lipschitz transformations of simple processes using CTFP without prior knowledge about the functional form of the observable processes and custom-tailored normalizing flows with non-Lipschitz transformations (see Appendix B for an example).

A latent variant of CTFP is further augmented with a static latent variable to introduce non-Markovian behavior. It models continuous stochastic processes as $\mathbf{X}_t = f(\mathbf{W}_t, t; \mathbf{Z})$, where \mathbf{Z} is a latent variable with standard Gaussian distribution and $f(\cdot, \cdot; \mathbf{z})$ is a CTFP model that decodes each sample \mathbf{z} of \mathbf{Z} into a stochastic process with continuous trajectories. Latent CTFP can be used to estimate finite-dimensional distributions using a variational approximation. However, it is unclear how much a latent variable \mathbf{Z} with finite dimensions can improve CTFPs representation power.

3 Model

Equipped with these tools, we can now describe our problem setting and model architecture. Let $\{(\mathbf{x}_{t_i}, t_i)\}_{i=1}^n$ denote a sequence of d -dimensional observations sampled at arbitrary points in time,

where \mathbf{x}_{t_i} and t_i denote the value and time stamp of the observation, respectively. The observations are assumed to be partial realizations of a continuous-time stochastic process \mathbf{X}_t . Our training objective is the maximization of the observational log-likelihood induced by \mathbf{X}_t on a given time grid,

$$\mathcal{L} = \log p_{\mathbf{X}_{t_1}, \dots, \mathbf{X}_{t_n}}(\mathbf{x}_{t_1}, \dots, \mathbf{x}_{t_n}), \quad (4)$$

for an inhomogeneous collection of sequences with varying lengths and time stamps. At test-time, in addition to the maximization of log-likelihoods, we are also interested in sampling sparse, dense, or irregular trajectories with finite-dimensional distributions that conform with these log-likelihoods. We model this challenging scenario with Continuous Latent Process Flows (CLPF). In Section 3.1, we present our model in more detail. Training and inference methods will be discussed in Section 3.2.

3.1 Continuous Latent Process Flows

A Continuous Latent Process Flow consists of two major components: an SDE describing the continuous latent dynamics of an observable stochastic process and a continuously indexed normalizing flow serving as a time-dependent decoder. The architecture of the normalizing flow itself can be specified in multiple ways, e.g., as an augmented neural ODE [14] or as a series of affine transformations [10]. The following paragraphs discuss the relationship between these components in more detail.

Continuous Latent Dynamics. Analogous to our overview in Section 2.1, we model the evolution of an m -dimensional time-continuous latent state \mathbf{Z}_t in the time interval $[0, T]$ using a flexible stochastic differential equation driven by an m -dimensional Wiener Process \mathbf{W}_t ,

$$d\mathbf{Z}_t = \boldsymbol{\mu}_\gamma(\mathbf{Z}_t, t) dt + \boldsymbol{\sigma}_\gamma(\mathbf{Z}_t, t) d\mathbf{W}_t, \quad (5)$$

where γ denotes the (shared) learnable parameters of the drift function $\boldsymbol{\mu}$ and variance function $\boldsymbol{\sigma}$. In our experiments, we implement $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ using deep neural networks (see Appendix E for details). Importantly, the latent state \mathbf{Z}_t exists for each $t \in [0, T]$ and can be sampled on any given time grid, which can be irregular and different for each sequence.

Time-Dependent Decoding. Latent variable models decode a latent state into an observable variable with complex distribution. As an observed sequence $\{(\mathbf{x}_{t_i}, t_i)\}_{i=1}^n$ is assumed to be a partial realization of a continuous-time stochastic process, continuous trajectories of the latent process \mathbf{Z}_t should be decoded into continuous trajectories of the observable process \mathbf{X}_t , and not discrete distributions. Following recent advances in dynamic normalizing flows [6, 10, 11], we model \mathbf{X}_t as

$$\mathbf{X}_t = F_\theta(\mathbf{O}_t; \mathbf{Z}_t, t), \quad (6)$$

where \mathbf{O}_t is a d -dimensional stochastic process with closed-form transition density³ and $F_\theta(\cdot; \mathbf{z}_t, t)$ is a normalizing flow parameterized by θ for any \mathbf{z}_t, t . The transformation F_θ decodes each sample path of \mathbf{Z}_t into a complex distribution over continuous trajectories \mathbf{X}_t if F_θ is a continuous mapping and the sampled trajectories of the base process \mathbf{O}_t are continuous with respect to time t . Unlike [11], who use a Wiener process as base process, we use the Ornstein–Uhlenbeck (OU) process, which has a stationary marginal distribution and bounded variance. As a result, the variance of the observation process does not increase due to the increase of variance in the base process and is primarily determined by the latent process \mathbf{Z}_t and flow transformation F_θ .

Flow Architecture. The continuously indexed normalizing flow $F_\theta(\cdot; \mathbf{z}_t, t)$ can be implemented in multiple ways. Deng et al. [11] use ANODE [14], defined as the solution to the initial value problem

$$\frac{d}{d\tau} \begin{pmatrix} \mathbf{h}(\tau) \\ \mathbf{a}(\tau) \end{pmatrix} = \begin{pmatrix} f_\theta(\mathbf{h}(\tau), \mathbf{a}(\tau), \tau) \\ g_\theta(\mathbf{a}(\tau), \tau) \end{pmatrix}, \quad \begin{pmatrix} \mathbf{h}(\tau_0) \\ \mathbf{a}(\tau_0) \end{pmatrix} = \begin{pmatrix} \mathbf{o}_t \\ (\mathbf{z}_t, t)^T \end{pmatrix}, \quad (7)$$

where $\tau \in [\tau_0, \tau_1]$, $\mathbf{h}(\tau) \in \mathbb{R}^d$, $\mathbf{a}(\tau) \in \mathbb{R}^{m+1}$, $f_\theta : \mathbb{R}^d \times \mathbb{R}^{m+1} \times [\tau_0, \tau_1] \rightarrow \mathbb{R}^d$, $g_\theta : \mathbb{R}^{m+1} \times [t_0, t_1] \rightarrow \mathbb{R}$, and F_θ is defined as the solution of $\mathbf{h}(\tau)$ at $\tau = \tau_1$. Note the difference between t and τ : while $t \in [0, T]$ describes the continuous process dynamics, $\tau \in [\tau_0, \tau_1]$ describes the continuous time-dependent decoding at each time step t . Alternatively, Cornish et al. [10] propose a variant of continuously indexed normalizing flows based on a series of N affine transformations f_i ,

$$\mathbf{h}_t^{(i+1)} = f_i(\mathbf{h}_t^{(i)}; \mathbf{z}_t, t) = k(\mathbf{h}_t^{(i)}) \cdot \exp(-u^{(i)}(\mathbf{z}_t, t)) - v^{(i)}(\mathbf{z}_t, t), \quad (8)$$

³More precisely, the conditional distribution $p_{\mathbf{O}_{t_i} | \mathbf{O}_{t_j}}(\mathbf{o}_{t_i} | \mathbf{o}_{t_j})$ must exist in closed form for any $t_j < t_i$.




Figure 2: **Graphical Model of Generation and Inference.** Red circles represent latent variables Z_{t_i} . Unfilled blue circles represent samples from the OU process at discrete time points O_{t_i} . Blue diamonds in Fig. 2a indicate each X_{t_i} is the result of a deterministic mapping of Z_{t_i} and O_{t_i} . Filled blue circles in Fig. 2b represent observed variables X_{t_i} .

where $\mathbf{h}_t^{(0)} = \mathbf{o}_t$, $\mathbf{h}_t^{(N)} = \mathbf{x}_t$, $u^{(i)}$ and $v^{(i)}$ are unconstrained transformations, and k is an invertible mapping like a residual flow [8]. The temporal relationships among Z_t , O_t , and X_t from a graphical model point of view are shown in Fig. 2a.

3.2 Training and Inference

With the model fully specified, we can now focus our attention on training and inference. Computing the observational log-likelihood (Eq.(4)) induced by a time-dependent decoding of an SDE (Eq.(6)) on an arbitrary time grid is challenging, because only few SDEs have closed-form transition densities. Consequently, variational approximations are needed for flexible SDEs such as Eq.(5). We propose a principled way of approximating the observational log-likelihood with a variational lower bound based on a novel piecewise construction of the posterior distribution of the latent process. In summary, we first express the observational log-likelihood as a marginalization over piecewise factors conditioned on a latent trajectory, then approximate this intractable computation with a piecewise variational posterior process, and finally derive a variational lower bound for it.

Observational Log-Likelihood. The observational log-likelihood can be written as an expectation over latent trajectories of the conditional likelihood, which can be evaluated in closed form,

$$\begin{aligned}
\mathcal{L} &= \log p_{\mathbf{X}_{t_1}, \dots, \mathbf{X}_{t_n}}(\mathbf{x}_{t_1}, \dots, \mathbf{x}_{t_n}) \\
&= \log \mathbb{E}_{\omega \sim \mathbf{W}_t} \left[p_{\mathbf{X}_{t_1}, \dots, \mathbf{X}_{t_n} | \mathbf{Z}_t}(\mathbf{x}_{t_1}, \dots, \mathbf{x}_{t_n} | \mathbf{Z}_t(\omega)) \right] \\
&= \log \mathbb{E}_{\omega \sim \mathbf{W}_t} \left[\prod_{i=1}^n p_{\mathbf{X}_{t_i} | \mathbf{X}_{t_{i-1}}, \mathbf{Z}_{t_i}, \mathbf{Z}_{t_{i-1}}}(\mathbf{x}_{t_i} | \mathbf{x}_{t_{i-1}}, \mathbf{Z}_{t_i}(\omega), \mathbf{Z}_{t_{i-1}}(\omega)) \right],
\end{aligned} \tag{9}$$

where $\mathbf{Z}_t(\omega)$ denotes a sample trajectory of \mathbf{Z}_t driven by $\omega \sim \mathbf{W}_t$. For simplicity, we assume w.l.o.g. and in this section only $\mathbf{Z}_0, \mathbf{X}_0$ to be given. As a result of invertibility, the conditional likelihood terms $p_{\mathbf{X}_{t_i} | \mathbf{X}_{t_{i-1}}, \mathbf{Z}_{t_i}, \mathbf{Z}_{t_{i-1}}}$ can be computed using the change-of-variables formula,

$$\begin{aligned}
&\log p_{\mathbf{X}_{t_i} | \mathbf{X}_{t_{i-1}}, \mathbf{Z}_{t_i}, \mathbf{Z}_{t_{i-1}}}(\mathbf{x}_{t_i} | \mathbf{x}_{t_{i-1}}, \mathbf{Z}_{t_i}(\omega), \mathbf{Z}_{t_{i-1}}(\omega)) \\
&= \log p_{O_{t_i} | O_{t_{i-1}}}(\mathbf{o}_{t_i} | \mathbf{o}_{t_{i-1}}) - \log \left| \det \frac{\partial F_\theta(\mathbf{o}_{t_i}; \mathbf{Z}_{t_i}(\omega), t_i)}{\partial \mathbf{o}_{t_i}} \right|,
\end{aligned} \tag{10}$$

where $\mathbf{o}_{t_i} = F_\theta^{-1}(\mathbf{x}_{t_i}; \mathbf{Z}_{t_i}(\omega), t_i)$.

Piecewise Construction of Variational Posterior. Directly computing the observational log-likelihood is intractable and we use a variational approximation during both training and inference. Good variational approximations rely on variational posteriors that are close enough to the true posterior of the latent trajectory conditioned on observations. Previous methods [24] use a single SDE to propose the variational posterior conditioned on all observations. Instead, we develop a more flexible method that can update the posterior process parameters when a new observation is seen and naturally adapts to different time grids. Our posterior process is not constrained by the Markov property of SDE solutions. Moreover, the proposed method serves as the basis for a principled approach to online inference tasks using variational posterior processes.

Our construction makes use of a further decomposition of the observational log-likelihood based on the following facts: $\{\mathbf{W}_{s+t} - \mathbf{W}_s\}_{t \geq 0}$ is also a Wiener process $\forall s \geq 0$ and the solution to Eq.(5) is a Markov process. Specifically, let $\{(\Omega^{(i)}, \mathcal{F}_{t_i - t_{i-1}}^{(i)}, P^{(i)})\}_{i=1}^n$ be a series of probability spaces on which n independent m -dimensional Wiener processes $\mathbf{W}_t^{(i)}$ are defined. We can sample a complete trajectory of the Wiener process \mathbf{W}_t in the interval $[0, T]$ by sampling independent trajectories of length $t_i - t_{i-1}$ from $\Omega^{(i)}$ and adding them, i.e., $\omega_t = \sum_{\{i: t_i < t\}} \omega_{t_i - t_{i-1}}^{(i)} + \omega_{t - t_{i^*} - 1}^{(i^*)}$, where $i^* = \max\{i : t_i < t\} + 1$. As a result, we can solve the latent stochastic differential equations in a piecewise manner. Samples of \mathbf{Z}_{t_i} can be obtained by solving the stochastic differential equation

$$d\tilde{\mathbf{Z}}_t = \boldsymbol{\mu}_\gamma(\tilde{\mathbf{Z}}_t, t + t_{i-1}) dt + \boldsymbol{\sigma}_\gamma(\tilde{\mathbf{Z}}_t, t + t_{i-1}) d\mathbf{W}_t^{(i)} \quad (11)$$

in the interval $[0, t_i - t_{i-1}]$, with $\mathbf{Z}_{t_{i-1}}$ as the initial condition. The log-likelihood of the observations can thus be rewritten as

$$\begin{aligned} \mathcal{L} &= \log \mathbb{E}_{\omega^{(1)}, \dots, \omega^{(n)} \sim \mathbf{W}_t^{(1)} \times \dots \times \mathbf{W}_t^{(n)}} \left[\prod_{i=1}^n p(\mathbf{x}_{t_i} | \mathbf{x}_{t_{i-1}}, \mathbf{z}_{t_i}, \mathbf{z}_{t_{i-1}}, \omega^{(i)}) \right] \\ &= \log \mathbb{E}_{\omega^{(1)} \sim \mathbf{W}_t^{(1)}} \left[p(\mathbf{x}_{t_1} | \mathbf{x}_{t_0}, \mathbf{z}_{t_1}, \mathbf{z}_{t_0}, \omega^{(1)}) \dots \right. \\ &\quad \left. \mathbb{E}_{\omega^{(i)} \sim \mathbf{W}_t^{(i)}} \left[p(\mathbf{x}_{t_i} | \mathbf{x}_{t_{i-1}}, \mathbf{z}_{t_i}, \mathbf{z}_{t_{i-1}}, \omega^{(i)}) \mathbb{E}_{\omega^{(i+1)} \sim \mathbf{W}_t^{(i+1)}} [\dots] \dots \right] \right]. \end{aligned} \quad (12)$$

It is worth noting that the value of \mathbf{x}_{t_i} is determined by \mathbf{z}_{t_i} and \mathbf{o}_{t_i} . With the distribution of \mathbf{O}_{t_i} depending on the value of $\mathbf{o}_{t_{i-1}}$, and thus on $\mathbf{z}_{t_{i-1}}$ and $\mathbf{x}_{t_{i-1}}$, after marginalizing over \mathbf{o}_{t_i} each conditional likelihood term in Eq.(12) is conditioned on $\mathbf{x}_{t_{i-1}}$, \mathbf{z}_{t_i} , and $\mathbf{z}_{t_{i-1}}$. For simplicity we also omit conditioning of each term on Wiener process trajectories up to t_{i-1} , i.e., $\omega^{(1)}, \omega^{(2)}, \dots, \omega^{(i)}$. In preparation of our variational approximation we can now introduce one posterior SDE

$$d\tilde{\mathbf{Z}}_t = \boldsymbol{\mu}_{\phi_i}(\tilde{\mathbf{Z}}_t, t + t_{i-1}) dt + \boldsymbol{\sigma}_\gamma(\tilde{\mathbf{Z}}_t, t + t_{i-1}) d\mathbf{W}_t^{(i)} \quad (13)$$

for each expectation $\mathbb{E}_{\omega^{(i)} \sim \mathbf{W}_t^{(i)}} \left[p(\mathbf{x}_{t_i} | \mathbf{x}_{t_{i-1}}, \mathbf{z}_{t_i}, \mathbf{z}_{t_{i-1}}, \omega^{(i)}) \mathbb{E}_{\omega^{(i+1)} \sim \mathbf{W}_t^{(i+1)}} [\dots] \dots \right]$ in Eq.(12).

Variational Lower Bound with Piecewise Importance Weighting. The posterior SDEs in Eq.(13) form the basis for a variational approximation of the expectations in Eq.(12). Specifically, sampling $\tilde{\mathbf{z}}$ from the posterior SDE, the expectation can be rewritten as

$$\mathbb{E}_{\omega^{(i)} \sim \mathbf{W}_t^{(i)}} \left[p(\mathbf{x}_{t_i} | \mathbf{x}_{t_{i-1}}, \tilde{\mathbf{z}}_{t_i}, \mathbf{z}_{t_{i-1}}, \omega^{(i)}) \mathbf{M}^{(i)}(\omega^{(i)}) \mathbb{E}_{\omega^{(i+1)} \sim \mathbf{W}_t^{(i+1)}} [\dots] \dots \right], \quad (14)$$

where $\mathbf{M}^{(i)} = \exp(-\int_0^{t_i - t_{i-1}} \frac{1}{2} |\mathbf{u}(\tilde{\mathbf{Z}}_s, s)|^2 ds - \int_0^{t_i - t_{i-1}} \mathbf{u}(\tilde{\mathbf{Z}}_s, s)^T d\mathbf{W}_s^{(i)})$ serves as importance weight for the sampled trajectory between the prior latent SDE (Eq.(11)) and posterior latent SDE (Eq.(13)), with \mathbf{u} satisfying $\boldsymbol{\sigma}_\gamma(\mathbf{z}, s + t_{i-1}) \mathbf{u}(\mathbf{z}, s) = \boldsymbol{\mu}_{\phi_i}(\mathbf{z}, s + t_{i-1}) - \boldsymbol{\mu}_\gamma(\mathbf{z}, s + t_{i-1})$. We can define such a posterior latent SDE for each interval. As a result, a sample $\tilde{\mathbf{z}}_{t_i}$ can be obtained by solving the posterior SDEs defined on the intervals up to t_i given the initial value $\tilde{\mathbf{z}}_{t_0}$ and sample paths of Wiener processes up to t_i , i.e., $\omega^{(1)}, \omega^{(2)}, \dots, \omega^{(i)}$. After applying the importance weight to each interval and Jensen's inequality, we can derive the following evidence lower bound (ELBO) of the log-likelihood:

$$\begin{aligned} \mathcal{L} &= \log \mathbb{E}_{\omega^{(1)} \sim \mathbf{W}_t^{(1)}} \left[p(\mathbf{x}_{t_1} | \mathbf{x}_{t_0}, \tilde{\mathbf{z}}_{t_1}, \tilde{\mathbf{z}}_{t_0}, \omega^{(1)}) \mathbf{M}^{(1)} \dots \mathbb{E}_{\omega^{(i)} \sim \mathbf{W}_t^{(i)}} \left[p(\mathbf{x}_{t_i} | \mathbf{x}_{t_{i-1}}, \tilde{\mathbf{z}}_{t_i}, \tilde{\mathbf{z}}_{t_{i-1}}, \omega^{(i)}) \mathbf{M}^{(i)} \dots \right] \right] \\ &= \log \mathbb{E}_{\omega^{(1)}, \dots, \omega^{(n)} \sim \mathbf{W}_t^{(1)} \times \dots \times \mathbf{W}_t^{(n)}} \left[\prod_{i=1}^n p(\mathbf{x}_{t_i} | \mathbf{x}_{t_{i-1}}, \tilde{\mathbf{z}}_{t_i}, \tilde{\mathbf{z}}_{t_{i-1}}, \omega^{(i)}) \mathbf{M}^{(i)}(\omega^{(i)}) \right] \\ &\geq \mathbb{E}_{\omega^{(1)}, \dots, \omega^{(n)} \sim \mathbf{W}_t^{(1)} \times \dots \times \mathbf{W}_t^{(n)}} \left[\sum_{i=1}^n \log p(\mathbf{x}_{t_i} | \mathbf{x}_{t_{i-1}}, \tilde{\mathbf{z}}_{t_i}, \tilde{\mathbf{z}}_{t_{i-1}}, \omega^{(i)}) + \sum_{i=1}^n \log \mathbf{M}^{(i)}(\omega^{(i)}) \right]. \end{aligned} \quad (15)$$

The bound above can be further extended into a tighter bound in IWAE [5] form by drawing multiple independent samples from each $\mathbf{W}_t^{(i)}$. The variational parameter ϕ_i is the output of an encoder RNN that takes as inputs the sequence of observations up to t_i , $\{\mathbf{x}_{t_1}, \dots, \mathbf{x}_{t_i}\}$, and the sequence

of previously sampled latent states, $\{z_{t_1}, \dots, z_{t_{i-1}}\}$. The parameter ϕ_i is thus updated based on both the latest observation and past history. As a result, the variational posterior distributions of the latent states Z_{t_i} are no longer constrained to be Markov and the parameterization of the variational posterior can flexibly adapt to different time grids. The dependency structure between latent and observed variables during inference is depicted in Fig. 2b.

4 Related Work

The introduction of neural ODEs [7] unified modeling approaches based on differential equations and modern machine learning. Explorations of time-series modeling leveraging this class of approach have been conducted, which we review here.

The latent ODE model [7, 32] propagates a latent state across time using ordinary differential equations. As a result, the entire latent trajectory is solely determined by its initial value. Even though latent ODE models have continuous latent trajectories, the latent state is decoded into observations at each time step independently. Neural controlled differential equations (CDEs) [20] and rough differential equations (RDEs) [25] propagate a hidden state across time continuously using controlled differential equations driven by functions of time interpolated from observations on irregular time grids. Neural ODE Processes (NDPs) [26] construct a distribution over neural ODEs in order to effectively manage uncertainty over the underlying process that generates the data.

Among existing time-series works with continuous dynamics, the latent SDE model [24] is most similar to ours. The SDE model includes an adjoint sensitivity method for training SDEs; the derivation of the variational lower bound in our proposed model is based on the same principle of trajectory importance weighting between two stochastic differential equations. The posterior process there is defined as a global stochastic differential equation. Our model further exploits the given observation time grid of each sequence to induce a piecewise posterior process with richer structure.

The continuous-time flow process [11] (CTFP) models irregular time-series data as an incomplete realization of continuous-time stochastic processes obtained by applying normalizing flows to the Wiener process. We have discussed the limits of CTFP in Sec. 2.2. Because CTFP is a generative model that is guaranteed to generate continuous trajectories, we use it as the decoder of a latent process for better inductive bias in modeling continuous dynamics. Apart from CTFP, there are also works outside the deep learning literature that apply invertible transformations to stochastic processes [33, 34]. Warped Gaussian Processes [33] transform Gaussian processes to non-Gaussian processes in observation space using monotonic functions. Copula Processes [34] extend the concept of copulas from multivariate random variables to stochastic processes. They transform a stochastic process via a series of marginal cumulative distribution functions to obtain another process while preserving the underlying dependency structure.

Alternative training frameworks have also been explored for neural SDEs [18, 19]. Stochastic differential equations can be learned as latent dynamics with a variational approximation [18]. Connections between generative adversarial network (GAN) objectives and neural SDEs have been drawn [19]. Brownian motion inputs can be mapped to time-series outputs. Alongside a CDE-based discriminator GAN-based training can be conducted to obtain continuous-time generative models.

5 Experiments

We compare our proposed architecture against several baseline models with continuous dynamics that can be used to fit irregular time-series data, including CTFP, latent CTFP, latent SDE, and latent ODE. We also run experiments with Variational RNN (VRNN) [9], a model that can be used to fit sequential data. However, VRNN does not define a continuous dynamical model and cannot generate trajectories with finite-dimensional distributions that are consistent with its log-likelihood estimation. All implementation and training details can be found in the supplementary material.

5.1 Synthetic Data

We evaluate our model on synthetic data sampled from known stochastic processes to verify its ability to capture a variety of continuous dynamics. We use the following processes:

Table 1: **Quantitative Evaluation (Synthetic Data)**. We show test negative log-likelihoods (NLLs) of four synthetic stochastic processes across different models. Below each process, we indicate the intensity of the Poisson process from which the time stamps for the test sequences were sampled for testing. [GBM: geometric Brownian motion (ground truth NLLs: $[\lambda = 2, \lambda = 20] = [0.388, -0.788]$); LSDE: linear SDE; CAR: continuous auto-regressive process; SLC: stochastic Lorenz curve]

Model	GBM		LSDE		CAR		SLC	
	$\lambda = 2$	$\lambda = 20$	$\lambda = 2$	$\lambda = 20$	$\lambda = 2$	$\lambda = 20$	$\lambda = 20$	$\lambda = 40$
VRNN [9]	0.425	-0.650	-0.634	-1.665	1.832	2.675	2.237	1.753
Latent ODE [32]	1.916	1.796	0.900	0.847	4.872	4.765	9.117	9.115
CTFP [11]	2.940	0.678	-0.471	-1.778	383.593	51.950	0.489	-0.586
Latent CTFP [11]	1.472	-0.158	-0.468	-1.784	249.839	43.007	1.419	-0.077
Latent SDE [24]	1.243	1.778	0.082	0.217	3.594	3.603	7.740	8.256
CLPF (ours)	0.444	-0.698	-0.831	-1.939	1.322	-0.077	-2.620	-3.963

Geometric Brownian Motion $d\mathbf{X}_t = \mu\mathbf{X}_t dt + \sigma\mathbf{X}_t d\mathbf{W}_t$. Even though geometric Brownian motion can theoretically be captured by the CTFP model, it would require the normalizing flow to be non-Lipschitz; there is no such constraint for the proposed CLPF model.

Linear SDE $d\mathbf{X}_t = (a(t)\mathbf{X}_t + b(t)) dt + \sigma(t) d\mathbf{W}_t$. The drift term of the SDE is a linear transformation of \mathbf{X}_t and the variance term is a deterministic function of time t . An application of Itô’s lemma shows that the solution is a stochastic process that cannot be captured by CTFP.

Continuous AR(4) Process. This process tests our model’s ability to capture non-Markov processes (see Appendix C for implementation details):

$$\begin{aligned} \mathbf{X}_t &= [d, 0, 0, 0]\mathbf{Y}_t, \\ d\mathbf{Y}_t &= A\mathbf{Y}_t dt + e d\mathbf{W}_t, \end{aligned} \quad \text{where } A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ a_1 & a_2 & a_3 & a_4 \end{bmatrix}. \quad (16)$$

Stochastic Lorenz Curve. We evaluate our model’s performance on multi-dimensional (3D) data:

$$\begin{aligned} d\mathbf{X}_t &= \sigma(\mathbf{Y}_t - \mathbf{X}_t) dt + \alpha_x d\mathbf{W}_t, \\ d\mathbf{Y}_t &= (\mathbf{X}_t(\rho - \mathbf{Z}_t) - \mathbf{Y}_t) dt + \alpha_y d\mathbf{W}_t, \\ d\mathbf{Z}_t &= (\mathbf{X}_t\mathbf{Y}_t - \beta\mathbf{Z}_t) dt + \alpha_z d\mathbf{W}_t. \end{aligned} \quad (17)$$

In all cases, we sample the observation time stamps from a homogeneous Poisson process with intensity λ . To demonstrate our model’s ability to generalize to different time grids, we evaluate it using different intensities λ . An approximate numerical solution to the SDEs is obtained using the Euler-Maruyama scheme for the Itô integral.

Our results on synthetic data are displayed in Table 1. We train and evaluate all models on observations of sample trajectories in the interval $[0, 30]$ with observation time stamps sampled from a Poisson process with intensity $\lambda = 2$, except for the stochastic Lorenz curve which is sampled in the interval $[0, 2]$. The results demonstrate the proposed CLPF model’s favourable performance across the board. In particular, our model outperforms both CTFP variants on all four synthetic datasets. We attribute this competitive edge to the expressive power of a generic SDE over a static latent variable. The CTFP models also perform relatively poorly on the Continuous AR(4) process, which is non-Markov. As the Continuous AR(4) process has an underlying 4-dim. stochastic process \mathbf{Y}_t and is generated by projecting \mathbf{Y}_t to a 1-dim. observation space \mathbf{X}_t , models driven by a high-dimensional latent process like latent SDE and CLPF show better performance in this case.

To evaluate our model’s performance in capturing continuous dynamics, we increase the (average) density of observations and generate observation time stamps from a Poisson process with intensity $\lambda = 20$ ($\lambda = 40$ for SLC). The results show that models that can generate continuous trajectories, including CTFP, latent CTFP and CLPF, generalize better to dense observations than the other models. In the first row of Fig. 3, we visualize sample sequences from CLPF and VRNN models trained on LSDE data and compare them with samples from an LSDE ground truth process. We run both CLPF and VRNN on a time grid between 0 and 30, with a gap of 0.01 between steps. Samples from the CLPF models share important visual properties with samples from the ground truth process. On the other hand, VRNN fails to generate sequences visually similar to the ground truth, despite its competitive density estimation results on sparse time grids.

Table 2: **Likelihood Estimation (Real-World Data)**. We show test negative log-likelihoods (NLLs; lower is better). For CTFP, the reported values are exact; for the other models, we report IWAE bounds using $K = 125$ samples. CLPF-ANODE stands for a CLPF model implemented with augmented neural ODEs. CLPF-iRes stands for a CLPF model implemented with indexed residual flows.

Model	Mujoco [32]	BAQD [35]	PTBDB [4]
VRNN [9]	-15.876	-1.204	-2.035
Latent ODE [32]	23.551	2.540	-0.533
Latent SDE [24]	3.071	1.512	-1.358
CTFP [11]	-7.598	-0.170	-1.281
Latent CTFP [11]	-12.693	-0.480	-1.659
CLPF-ANODE (ours)	-14.694	-0.619	-1.575
CLPF-iRes (ours)	-10.873	-0.486	-1.519

Table 3: **Ablation Study (Synthetic Data)**. We show test negative log-likelihoods across different variants of the proposed model. We report IWAE bounds using $K = 125$ samples with observation time stamps sampled from a Poisson point process with $\lambda = 2$. [CLPF-Global: single global posterior SDE in latent SDE style [24]; CLPF-Independent: independent decoder instead of CTFP-decoder; CLPF-Wiener: Wiener base process instead of OU-process]

Model	GBM	LSDE	CAR	SLC
CLPF-Global	0.447	-0.821	1.552	-3.304
CLPF-Independent	0.800	-0.326	4.970	7.924
CLPF-Wiener	0.390	-0.790	1.041	-1.885
Latent SDE	1.243	0.082	3.594	7.740
CLPF	0.444	-0.831	1.322	-2.620

Table 4: **Sequential Prediction (Real-World Data)**. We report the average L2 distance between prediction results and ground truth observations in a sequential prediction setting. The prediction is based on the average of 125 samples. Results are reported in the format *mean*, [25th percentile, 75th percentile].

Model	Mujoco [32]	BAQD [35]	PTBDB [4]
VRNN [9]	1.599, [0.196, 1.221]	0.519, [0.168, 0.681]	0.037, [0.005, 0.032]
Latent ODE [32]	13.959, [9.857, 15.673]	1.416, [0.936, 1.731]	0.224, [0.114, 0.322]
Latent SDE [24]	7.627, [2.384, 8.381]	0.848, [0.454, 1.042]	0.092, [0.032, 0.111]
CTFP [11]	1.969, [0.173, 1.826]	0.694, [0.202, 0.966]	0.055, [0.006, 0.046]
Latent CTFP [11]	1.983, [0.167, 1.744]	0.680, [0.189, 0.943]	0.065, [0.007, 0.059]
CLPF-ANODE (ours)	1.629, [0.149, 1.575]	0.542, [0.150, 0.726]	0.048, [0.005, 0.041]
CLPF-iRes (ours)	1.846, [0.177, 1.685]	0.582, [0.183, 0.805]	0.055, [0.006, 0.049]

5.2 Real-world Data

Many real-world systems, despite having continuous dynamics, are recorded via observations at a fixed sampling rate. Mujoco-Hopper [32], Beijing Air-Quality Dataset (BAQD) [35] and PTB Diagnostic Database (PTBDB) [4, 16] are examples of such datasets. We create training/testing data at irregular times by drawing time stamps from a Poisson process and mapping them to the nearest observed sample points (see Appendix D for details).

We demonstrate our model’s ability to fit this type of observation as a surrogate for its ability to capture real-world continuous dynamics. Additionally, we also evaluate CLPF on a sequential prediction task to illustrate its benefits in an online inference setting. The results for these experiments are reported in Table 2 and Table 4. In the likelihood estimation task our model achieves competitive performance with respect to other methods that generate continuous dynamics. We do observe low negative log-likelihoods (NLLs) by VRNN on these data, though note that VRNN does not generate true continuous dynamics. In the sequential prediction task CLPF outperforms the other baselines with continuous dynamics; the performance gap between CLPF and VRNN is significantly reduced compared to likelihood estimation.

In the second row of Fig. 3, we also visualize one dimension of sample sequences generated from CLPF and VRNN models trained on a dense time grid of BAQD data (gap between steps: 0.01) and compare with ground truth samples from the dataset. We can see that VRNN fails to generate trajectories visually similar to samples from the ground truth data on a dense time grid, despite its favourable NLL on a sparse and irregular time grid. The sampled trajectories tend to converge and show little variance after running VRNN for a large number of iterations (1500 for $t = 1.5$). We also observe that VRNN fails to generate trajectories with consistent visual properties when we sample observations on different time grids (see supplementary material for more qualitative samples).

5.3 Ablation Study

We compare the proposed CLPF models against three variants that are obtained by making changes to key components of CLPF. The results with these models are shown in Table 3. **CLPF-Global** is a variant created by replacing our piecewise construction with a global posterior process similar to latent SDE [24]. **CLPF-Independent** is a variant where the continuously indexed normalizing flows

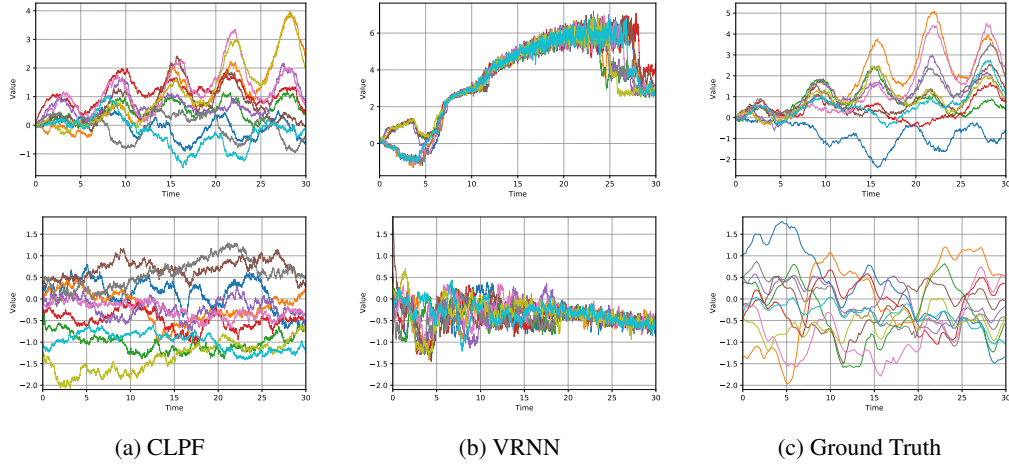


Figure 3: **Qualitative Samples on Dense Time Grid of CLPF and VRNN Models.** We generate the sample sequences from CLPF and VRNN trained with LSDE (first row) and BAQD (second row) data and compare the samples generated from the model to samples from the ground truth process (LSDE) or training set (BAQD).

and CTFP-decoding are replaced with static distributions and independent decoders. **CLPF-Wiener** is a variant that uses a Wiener process as the base process for the CTFP decoder.

CLPF-Global differs from CLPF in a manner similar to the difference between latent SDE and CLPF-Independent. In both comparisons we can see modest improvements of performance when we replace the posterior process defined by a single SDE with a piecewise constructed posterior. Both the comparisons between CLPF and CLPF-Independent and between CLPF-Global and latent SDE show that a CTFP-style decoding into continuous trajectories is indeed a better inductive bias when modeling data with continuous dynamics. Although the comparison with CLPF-Wiener does not show a clear trend, CLPF still delivers an overall better performance on the synthetic datasets. Therefore we use the OU process as our base process in all CLPF experiments.

6 Discussion and Conclusion

Limits. While being more expressive than the CTFP model, the proposed CLPF model is not capable of exact likelihood evaluation and relies on variational approximations for estimation and optimization. The CLPF model can also comprise up to two differential equations that need to be numerically solved, potentially incurring higher computational cost. Fortunately, this computational cost can be controlled via additional parameters: (1) The approximation error of numerical ODE/SDE solutions can be controlled through the tolerance parameters of numerical solvers and the inversion error of residual flows can be controlled through the number of fixed-point iterations. [3]; (2) The tightness of the IWAE bound for likelihood estimation can be controlled through the number of latent samples.

Broader Impact. As an expressive generative model for continuous time-series, we believe CLPF can be used to model continuous dynamics from partial observations in a wide range of areas, including physics, healthcare, and finance. Furthermore, it can potentially be leveraged in a variety of downstream tasks, including data generation, forecasting, and interpolation. However, care must be taken that CLPF models are not used for adverse inference on restricted data. Protection of critical unobserved information is an active area of research [1] that can be explored as an orthogonal component in future versions of our model.

Conclusion. We have presented Continuous Latent Process Flows (CLPF), a generative model of continuous dynamics that enables inference on arbitrary real time grids, a complex operation for which we have also introduced a powerful piecewise variational approximation. Our architecture is built around the representation power of a flexible stochastic differential equation driving a continuously indexed normalizing flow. An ablation study as well as a comparison to state-of-the-art baselines for continuous dynamics have demonstrated the effectiveness of our contributions on both synthetic and real-world datasets. A set of qualitative results support our findings. In the future, we plan to explore the theoretical properties of our model in more detail, including an analysis of its universal approximation properties and its ability to capture non-stationary dynamics.

References

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep Learning with Differential Privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016.
- [2] Cedric Archambeau, Dan Cornford, Manfred Opper, and John Shawe-Taylor. Gaussian Process Approximations of Stochastic Differential Equations. *Proceedings of Machine Learning Research (PMLR)*, 2007.
- [3] Jens Behrmann, Will Grathwohl, Ricky T. Q. Chen, David Duvenaud, and Joern-Henrik Jacobsen. Invertible Residual Networks. *International Conference on Machine Learning (ICML)*, 2019.
- [4] R Bousseljot, D Kreiseler, and A Schnabel. Nutzung der EKG-Signaldatenbank CARDIODAT der PTB über das Internet. *Biomedizinische Technik/Biomedical Engineering*, 1995.
- [5] Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. Importance Weighted Autoencoders. *International Conference on Learning Representations (ICLR)*, 2016.
- [6] Anthony Caterini, Rob Cornish, Dino Sejdinovic, and Arnaud Doucet. Variational Inference with Continuously-Indexed Normalizing Flows. *arXiv:2007.05426*, 2020.
- [7] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural Ordinary Differential Equations. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [8] Ricky T. Q. Chen, Jens Behrmann, David Duvenaud, and Jörn-Henrik Jacobsen. Residual Flows for Invertible Generative Modeling. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [9] Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio. A Recurrent Latent Variable Model for Sequential Data. *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.
- [10] Rob Cornish, Anthony Caterini, George Deligiannidis, and Arnaud Doucet. Relaxing Bijectivity Constraints with Continuously Indexed Normalising Flows. *International Conference on Machine Learning (ICML)*, 2020.
- [11] Ruizhi Deng, Bo Chang, Marcus A Brubaker, Greg Mori, and Andreas Lehrmann. Modeling Continuous Stochastic Processes with Dynamic Normalizing Flows. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [12] Laurent Dinh, David Krueger, and Yoshua Bengio. NICE: Non-linear Independent Components Estimation. *International Conference on Learning Representations (ICLR)*, 2015.
- [13] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density Estimation using Real NVP. *International Conference on Learning Representations (ICLR)*, 2017.
- [14] Emilien Dupont, Arnaud Doucet, and Yee Whye Teh. Augmented Neural ODEs. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [15] Ramazan Gençay, Michel Dacorogna, Ulrich A Muller, Olivier Pictet, and Richard Olsen. *An Introduction to High-frequency Finance*. Elsevier, 2001.
- [16] Ary L Goldberger, Luis AN Amaral, Leon Glass, Jeffrey M Hausdorff, Plamen Ch Ivanov, Roger G Mark, Joseph E Mietus, George B Moody, Chung-Kang Peng, and H Eugene Stanley. PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals. *Circulation*, 2000.
- [17] Will Grathwohl, Ricky T. Q. Chen, Jesse Bettencourt, and David Duvenaud. Scalable Reversible Generative Models with Free-form Continuous Dynamics. *International Conference on Learning Representations (ICLR)*, 2019.
- [18] Ali Hasan, João M Pereira, Sina Farsiu, and Vahid Tarokh. Identifying Latent Stochastic Differential Equations with Variational Auto-Encoders. *arXiv:2007.06075*, 2020.
- [19] Patrick Kidger, James Foster, Xuechen Li, Harald Oberhauser, and Terry Lyons. Neural SDEs as Infinite-dimensional GANs. *International Conference on Machine Learning (ICML)*, 2021.
- [20] Patrick Kidger, James Morrill, James Foster, and Terry Lyons. Neural Controlled Differential Equations for Irregular Time Series. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [21] Durk P Kingma and Prafulla Dhariwal. Glow: Generative Flow with Invertible 1x1 convolutions. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [22] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved Variational Inference with Inverse Autoregressive Flow. *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [23] Ivan Kobyzev, Simon Prince, and Marcus A Brubaker. Normalizing Flows: Introduction and Ideas. *arXiv:1908.09257*, 2019.

- [24] Xuechen Li, Ting-Kam Leonard Wong, Ricky T. Q. Chen, and David Duvenaud. Scalable Gradients for Stochastic Differential Equations. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- [25] James Morrill, Cristopher Salvi, Patrick Kidger, James Foster, and Terry Lyons. Neural Rough Differential Equations for Long Time Series. *International Conference on Machine Learning (ICML)*, 2021.
- [26] Alexander Norcliffe, Cristian Bodnar, Ben Day, Jacob Moss, and Pietro Lio. Neural ODE Processes. *International Conference on Learning Representations (ICLR)*, 2021.
- [27] Bernt Oksendal. *Stochastic Differential Equations: An Introduction with Applications*. Springer Science & Business Media, 2013.
- [28] George Papamakarios, Theo Pavlakou, and Iain Murray. Masked Autoregressive Flow for Density Estimation. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [29] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing Flows for Probabilistic Modeling and Inference. *Journal of Machine Learning Research (JMLR)*, 2021.
- [30] Kira Rehfeld, Norbert Marwan, Jobst Heitzig, and Jürgen Kurths. Comparison of Correlation Analysis Techniques for Irregularly Sampled Time Series. *Nonlinear Processes in Geophysics*, 2011.
- [31] Danilo Rezende and Shakir Mohamed. Variational Inference with Normalizing Flows. *International Conference on Machine Learning (ICML)*, 2015.
- [32] Yulia Rubanova, Ricky T. Q. Chen, and David Duvenaud. Latent Ordinary Differential Equations for Irregularly-Sampled Time Series. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [33] Edward Snelson, Carl Edward Rasmussen, and Zoubin Ghahramani. Warped Gaussian Processes. *Advances in Neural Information Processing Systems (NeurIPS)*, 2004.
- [34] Andrew G Wilson and Zoubin Ghahramani. Copula Processes. *Advances in Neural Information Processing Systems (NeurIPS)*, 2010.
- [35] Shuyi Zhang, Bin Guo, Anlan Dong, Jing He, Ziping Xu, and Song Xi Chen. Cautionary Tales on Air Quality Improvement in Beijing. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 2017.
- [36] Gilles Zumbach and Ulrich Müller. Operators on Inhomogeneous Time Series. *International Journal of Theoretical and Applied Finance*, 2001.