
Learning Deep Hidden Nonlinear Dynamics from Aggregate Data

Yisen Wang^{1,2,3} Bo Dai^{1,4} Lingkai Kong¹ Sarah Monazam Erfani⁵ James Bailey⁵ Hongyuan Zha¹
¹Georgia Tech ²Tsinghua University ³Ant Financial Services Group ⁴Google Brain ⁵University of Melbourne
wangys14@mails.tsinghua.edu.cn, bohr.dai@gmail.com, zha@gatech.edu

Abstract

Learning nonlinear dynamics from diffusion data is a challenging problem since the individuals observed may be different at different time points, generally following an aggregate behaviour. Existing work cannot handle the tasks well since they model such dynamics either directly on observations or enforce the availability of complete longitudinal individual-level trajectories. However, in most of the practical applications, these requirements are unrealistic: the evolving dynamics may be too complex to be modeled directly on observations, and individual-level trajectories may not be available due to technical limitations, experimental costs and/or privacy issues. To address these challenges, we formulate a model of diffusion dynamics as the *hidden stochastic process* via the introduction of hidden variables for flexibility, and learn the hidden dynamics directly on *aggregate observations* without any requirement for individual-level trajectories. We propose a dynamic generative model with Wasserstein distance for LEarninG dEep hidden Nonlinear Dynamics (LEGEND) and prove its theoretical guarantees as well. Experiments on a range of synthetic and real-world datasets illustrate that LEGEND has very strong performance compared to state-of-the-art baselines.

1 INTRODUCTION

Diffusion data is a widespread form of data that involves spatial or status transitions over time, *e.g.*, Brownian movement in physics, cell differentiation or gene expression in biology, molecular motion in chemistry, bird migration in ecology, traffic flows in transportation, population trends in social sciences and so on. Learning

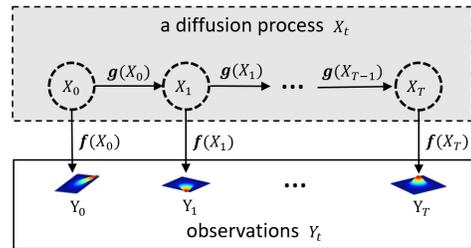


Figure 1: An illustration of the framework which builds dynamics on an auxiliary hidden variable X_t with a observation function. Observation Y_t is the aggregate formatted data.

the underlying dynamics which governs the evolution of such data is a fundamental problem. It reveals the nature of the dynamical phenomenon, based on which we can make better future predictions. However, in these areas, complete longitudinal individual-level trajectories (*i.e.*, the tracking of one individual over the entire diffusion process) may often not be available due to technical limitations, experimental costs and/or privacy issues. Rather, one often instead observes a random group of independently sampled individuals from the population, and the observations can contain different individuals at different time points. This is common for catch and release experiments in ecology (*e.g.*, bird migration) where it is difficult to observe a single bird twice (Bartholomew & Bohnsack, 2005), and in biological research where a cell may need to be sacrificed in order for an observation on it to be made (Banks & Potter, 2004). We refer to observations made in these scenarios as *aggregate observations* to differentiate them from the case of individual-level trajectories which provide full information.

Modeling the dynamics on aggregate observations have been investigated recently in (Hashimoto et al., 2016), where a stochastic differential equation (SDE) has been used to capture the transition directly on observations Y_t . However, its performance degrades when the dynamics

become complex due to their limited expressive ability, as illustrated later in our experiments. Instead of modeling dynamics directly on observations, a hidden variable X_t can be introduced for modeling complicated dynamics, which can be decomposed into a relatively simple hidden dynamic on X_t with a complicated observation function. As illustrated in Figure 1, $Y_t(t \in [0, T])$ is a series of aggregated observations of a diffusion process. We formulate that Y_t is determined by the hidden dynamic on X_t and the observation function $\mathbf{f}(X_t)$. Existing models such as Hidden Markov Model (HMM) (Eddy, 1996), Kalman Filter (KF) (Harvey, 1990) and Particle Filter (PF) (Djuric et al., 2003) are popular methods with hidden variables. However, these models and their variants (Langford et al., 2009; Hefny et al., 2015) require individual-level trajectories, which may not be available, as was mentioned earlier. It consequently still remains an open issue as to how one can learn the underlying dynamics directly from aggregate observations with a hidden stochastic process, for complicated real-world scenarios.

To address these challenges, we propose a novel framework to incorporate the use of hidden variables into the modeling of diffusion dynamics from evolving distributions (as those approximated from aggregate observations). We bypass the need for likelihood-based estimation of model parameters and posterior estimation of hidden variables by using Wasserstein distance learning. The model we propose is named LEGEND (LEarninG dEep hidden Nonlinear Dynamics) and the main contributions are:

- We propose a framework for learning complicated nonlinear dynamics from aggregate data via a hidden continuous stochastic process.
- We extend Wasserstein learning to likelihood-free and posterior-free estimations of dynamic parameter learning and hidden state inference.
- We theoretically provide a generalization bound and convergence analysis of our framework, which is the first theoretical result as far as we know.
- We empirically demonstrate the effectiveness of our framework for learning nonlinear dynamics from aggregate observations on both synthetic and real-world datasets.

2 PROBLEM DEFINITION

We first introduce a continuous model of diffusion dynamics using a stochastic differential equation (SDE), then formally define the tasks of filtering and smoothing based inference, then review the Wasserstein distance objective as one of the distribution metric.

Hidden Continuous Nonlinear Dynamics. To characterize the dynamics of observations, we introduce a hidden continuous nonlinear dynamical system as shown in Figure 1, together with a measurement of hidden states. In particular, the hidden state $X_t \in \mathbb{R}^n$ is the underlying auxiliary variable that cannot be accessed directly, and it follows a SDE:

$$dX_t = \mathbf{g}(X_t)dt + \Sigma^{1/2}d\omega_t, \quad (1)$$

where $\mathbf{g} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a nonlinear deterministic drift function, and $\omega_t \in \mathbb{R}^n$ is a Brownian motion process with noise covariance $\Sigma \in \mathbb{R}^n \times \mathbb{R}^n$. At each time point, the observation Y_t is written as a measurement of the hidden state X_t :

$$Y_t = \mathbf{f}(X_t), \quad (2)$$

where $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a nonlinear observation function. Together, Eqs. (1) and (2) define the nonlinear dynamics with continuous hidden states.

Aggregate Observations. We obtain a collection of independent and identically distributed (i.i.d) samples $\{y_t^i\}_{i=1}^N$ of Y_t at some time point t , that we term aggregate or distributional observations. The observed individuals in previous time observations $\{y_{t-1}^i\}_{i=1}^N$ are often not identical to those for the current time observations $\{y_t^i\}_{i=1}^N$, implying it is not possible to construct the full trajectory of a single individual. However, we can approximate the probability distribution in terms of a finite number of samples as

$$p(Y_t) \approx \frac{1}{N} \sum_{i=1}^N \delta(Y_t - y_t^i). \quad (3)$$

Thus, we can treat these aggregate samples from the same time point together as a distribution which evolves in the dynamic system.

Problems. Under this aggregate setting for observations, it is difficult to obtain individual-level trajectories due to technical limitations, experimental costs and/or privacy issues. Therefore, we propose a new framework to learn the nonlinear dynamics for the distributions (as approximated from aggregate observations) without the need for individual-level trajectories. That is, we treat $\{y_t^i\}_{i=1}^N$ as an empirical approximation to the distribution at time t and its dynamics is learned via an auxiliary hidden variable X_t . Once the dynamics are learned, we are faced with two inference tasks:

- 1) *Filtering based inference*: the task is to infer the next future observation Y_{T+1} , given observations $\{Y_0, Y_1, \dots, Y_T\}$.
- 2) *Smoothing based inference*: the task is to infer a missing intermediate observation $Y_k(0 < k < T)$, given $\{Y_0, \dots, Y_{k-1}, Y_{k+1}, \dots, Y_T\}$.

Wasserstein Distance Objective. Following our distribution-based problem definition, the metric on distributions is a key criterion for our objection function, just like the mean squared error (MSE) criterion for regression problems. Among the well-known distribution-based measures, such as Total Variation (TV) distance, Kullback-Leibler (KL) divergence, Jensen-Shannon (JS) divergence, Wasserstein distance has recently been shown to possess more appealing properties for distance measurement of distributions (Arjovsky et al., 2017). We therefore choose to adopt the Wasserstein distance for measuring the discrepancy between the learned distributions and their ground truth.

The definition of Wasserstein-1 distance (also named the Earth-Mover distance (EM)) is:

$$W(\mathbb{P}_r, \mathbb{P}_g) = \inf_{\gamma \in \Pi(\mathbb{P}_r, \mathbb{P}_g)} \mathbb{E}_{(x,y) \sim \gamma} [\|x - y\|], \quad (4)$$

where $\Pi(\mathbb{P}_r, \mathbb{P}_g)$ denotes the set of all joint distributions $\gamma(x, y)$ whose marginals are \mathbb{P}_r and \mathbb{P}_g respectively. The infimum in (4) is highly intractable. However, Kantorovich-Rubinstein duality (Villani, 2008) shows that

$$W(\mathbb{P}_r, \mathbb{P}_g) = \sup_{\|D\|_{L \leq 1}} \mathbb{E}_{x \sim \mathbb{P}_r} [D(x)] - \mathbb{E}_{x \sim \mathbb{P}_g} [D(x)], \quad (5)$$

where the supremum is over all 1-Lipschitz functions D . We can assume a parameterized family of functions $\{D_w\}_{w \in \mathcal{W}}$ lying in the 1-Lipschitz function space. Therefore, Eq. (5) could be solved by

$$\max_{w \in \mathcal{W}} \mathbb{E}_{x \sim \mathbb{P}_r} [D_w(x)] - \mathbb{E}_{x \sim \mathbb{P}_g} [D_w(x)]. \quad (6)$$

With Wasserstein distance, our objectives for dynamic learning and inference tasks can be unified to minimize the Wasserstein distance between the generated distribution and the target distribution, which will be instantiated in the following Section 3.

3 PROPOSED FRAMEWORK

In this section, we first discuss our methodology for parameter learning of dynamics within the LEGEND framework, and then introduce in detail how the framework addresses the filtering and smoothing based inference problems.

3.1 Parameter Learning of Dynamics

In order to efficiently solve the SDE of hidden state X_t , we adopt an approximate numerical solution called the Euler-Maruyama method (Talay, 1994). Suppose the SDE is defined on $[0, T]$, then the Euler-Maruyama approximation to the true solution of SDE is a Markov chain defined

as follows:

$$X_{t+\Delta t} = X_t + \mathbf{g}(X_t)\Delta t + \Sigma^{1/2}\Delta\omega_t, \quad (7)$$

where the interval $[0, T]$ is partitioned into M equal sub-intervals of width $\Delta t = T/M > 0$ and $\Delta\omega_t$ are independent and identically distributed normal random variables with expected value zero and variance Δt . Correspondingly, observations Y_t are functions of X_t :

$$Y_t = \mathbf{f}(X_t). \quad (8)$$

Given a sequence of distributional observations, we need to minimize the Wasserstein distance between the generated distribution and the observed distribution at each time point to learn functions \mathbf{f} and \mathbf{g} . The objective function for parameter learning of dynamics is

$$\min_{\mathbf{f}, \mathbf{g}} \sum_t W(\mathbb{P}(Y_t), \mathbb{P}(\mathbf{f}(X_t))), \quad (9)$$

where $X_t \sim \mathbb{P}(X_t|X_{t-1})$. The evolving process¹ from X_{t-1} to X_t is controlled by function \mathbf{g} following the SDE in Eq. (7). One common approach for learning \mathbf{f} and \mathbf{g} is to calculate the likelihood of Y_t under distributions, however in many cases, this is intractable. Here, we propose to use generative models to directly generate samples which satisfy the target distribution Y_t . Following the minimization of the Wasserstein distance between generations and observations, the generative model can eventually learn the dynamics of Y_t . The specific form of parameterization will be described in Section 4.

3.2 Filtering based Inference

The inference of Y_{T+1} given observations $\mathcal{Y}_T = \{Y_0, Y_1, \dots, Y_T\}$ can be solved by:

$$Y_{T+1} = (\mathbf{f} \circ \mathbf{g})(X_T). \quad (10)$$

To achieve this, we need to obtain the hidden state X_T first, that is, find the posterior probability $p(X_T|\mathcal{Y}_T)$ of the hidden state conditioned on the entire sequence of observations \mathcal{Y}_T , which is a filtering problem.

To obtain the posterior of the hidden state, the classical forward algorithm needs to solve one dynamic programming problem per sample, which requires individual-level trajectory for posterior inference. However, for our aggregate observation setting, we alternatively treat the Bayesian inference problem from an optimization perspective following (Dai et al., 2016).

We first briefly introduce the idea of the optimization method, then generalize it to solve our problem. Dai

¹There may be several Δt intervals between time $t - 1$ and t

et al. (2016) use a probability $q(U)$ to approximate the posterior probability $p(U|V)$ by minimizing

$$\min_{q(U) \in \mathcal{P}} -\langle q(U), \log p(V|U) \rangle + KL(q(U) \parallel p(U)) \quad (11)$$

over the space of all valid densities \mathcal{P} . $\langle \cdot \rangle$ is the inner product, KL is the Kullback-Leibler divergence, U is the hidden variable and V is the observation variable. Thus, $p(V|U)$ is the likelihood of observation and $p(U)$ is the prior of the hidden variable. Assuming we have the trajectory for a single individual ($x_t \sim X_t, y_t \sim Y_t$), then the posterior probability of filtering is

$$p(x_t|y_{1:t}) = \frac{p(y_t|x_t)p(x_t|y_{1:t-1})}{\int p(y_t|x_t)p(x_t|y_{1:t-1})dx_t}, \quad (12)$$

where $p(x_t|y_{1:t-1})$ is the propagation probability and $p(x_t|y_{1:t})$ is the updated probability. Generally, $p(x_t|y_{1:t-1})$ could be regarded as the prior of x_t for the updated probability $p(x_t|y_{1:t})$. Following the idea of Eq. (11), we can use a probability $\pi(x_t)$ to approximate the posterior probability $p(x_t|y_{1:t})$ by recursively optimizing

$$\min_{\pi(x_t) \in \mathcal{P}} -\langle \pi(x_t), \log p(y_t|x_t) \rangle + KL(\pi(x_t) \parallel p(x_t|y_{1:t-1})), \quad (13)$$

where

$$\begin{aligned} p(x_t|y_{1:t-1}) &= \int p(x_t, x_{t-1}|y_{1:t-1})dx_{t-1} \\ &= \int p(x_t|x_{t-1})p(x_{t-1}|y_{1:t-1})dx_{t-1} \\ &= \int p(x_t|x_{t-1})\pi(x_{t-1})dx_{t-1}. \end{aligned} \quad (14)$$

In the following, we generalize Eqs. (13) and (14) which were defined on an individual trajectory, to the case for aggregate/distributional data. In Eq. (13), to obtain the optimal solution, we need to maximize the inner product (the first term) and minimize the KL divergence (the second term). We redefine the two terms using Wasserstein distance. For the first term, since maximizing the inner product is equivalent to minimizing the distance between distributions, we replace the inner product with the Wasserstein distance between the distributions of $\mathbf{f}(\pi_t)$ (generated) and Y_t (ground truth). For the second term, we replace KL divergence with Wasserstein distance which is a better measurement for distributions and is thus more suitable for aggregate data (Arjovsky et al., 2017). In Eq. (14), the relationship between two consecutive time of hidden variables is replaced by function \mathbf{g} . We then can generalize Eqs. (13) and (14) to our filtering objective function under aggregate observations:

$$\min \sum_t W(\mathbb{P}(\mathbf{f}(\pi_t)), \mathbb{P}(Y_t)) + W(\mathbb{P}(\pi_t), \mathbb{P}(\mathbf{g}(\pi_{t-1}))), \quad (15)$$

where $\pi_t \sim \mathbb{P}(X_t|Y_t)$ is our target filtering distribution.

3.3 Smoothing based Inference

Smoothing based inference is for predicting the missing intermediate observation $Y_k (0 < k < T)$ given observations $\mathcal{Y}_{T \setminus k} = \{Y_0, \dots, Y_{k-1}, Y_{k+1}, \dots, Y_T\}$. One method (Desbouvier et al., 2011) to solve this is

$$Y_k = (\mathbf{f} \circ \mathbf{g})(X_{k-1}). \quad (16)$$

To achieve this, we need to obtain the hidden state X_{k-1} first. This is a smoothing problem which focuses on a hidden state somewhere in the middle of a sequence conditioned on the whole sequence of observations.

Different from the filtering task where the current state is estimated recursively from all past observations, smoothing computes the best state estimates given all available observations from both the past and the future. One well-known and simple approach for smoothing is the forward-backward smoother. During a forward pass the standard filtering algorithm is applied to the observations. Afterwards, on the backward pass, inverse filtering is applied to the same time series of observations. Finally the filtering estimates of both the forward and backward pass are combined into the smoothed estimates (Briers et al., 2010). Since the information from the observation should be incorporated only once into the smoothed estimate, we need to combine the posterior estimate of the forward pass with the prior estimate of the backward pass and vice versa.

Thus, following the idea in the above filtering problem (treating the posterior estimation from a optimization perspective), we first learn a forward estimate of the hidden state π_t^f and also a backward estimate π_t^b using Eq. (15). These then form a weighted Wasserstein barycenter problem (Agueh & Carlier, 2011) whose solution is the posterior of smoothing (Kitagawa, 1994). That is, we can obtain the optimal smoothing result π_t^s by optimizing the Wasserstein distance to the observations and the weighted Wasserstein barycenter:

$$\begin{aligned} \min \sum_t W(\mathbb{P}(\mathbf{f}(\pi_t^s)), \mathbb{P}(Y_t)) \\ + \lambda_1 W(\mathbb{P}(\pi_t^s), \mathbb{P}(\pi_t^f)) \\ + \lambda_2 W(\mathbb{P}(\pi_t^s), \mathbb{P}(\pi_t^b)), \end{aligned} \quad (17)$$

where $\pi_t^s \sim \mathbb{P}(X_t|Y_T)$ is our target smoothing distribution. And the weights λ_1 and λ_2 are hyperparameters, which are given intuitively with $\lambda_1 = t/T$ and $\lambda_2 = 1 - \lambda_1$ such that smoothing problem becomes filtering problem when $t = T \rightarrow \lambda_1 = 1$. Actually, there are other alternatives one could use for the weights, but these basic settings already work well in our experiments.

4 MODEL PARAMETERIZATION

As stated above, we adopt Wasserstein distance to measure the difference between distributions, and have defined our objectives accordingly. In this section, we establish a dynamic generative model via neural network parameterization based on Wasserstein distance.

According to the dual formulation of Wasserstein distance Eq. (6), our distribution-based objective of parameter learning of dynamics in Eq. (9) becomes

$$\min_{\mathbf{f}, \mathbf{g}} \sum_t \left(\max_{D_t^1} (\mathbb{E}_{y_t \sim \mathbb{P}(Y_t)} [D_t^1(y_t)]) - \mathbb{E}_{x_t \sim \mathbb{P}(x_t | \mathbf{g}(x_{t-1}))} [D_t^1(\mathbf{f}(x_t))] \right). \quad (18)$$

For the filtering and smoothing tasks, we introduce a new function \mathbf{h} to characterize the target filtering or smoothing distributions. The filtering objective in Eq. (15) becomes

$$\begin{aligned} \min_{\mathbf{h}} \sum_t \left(\max_{D_t^1} (\mathbb{E}_{y_t \sim \mathbb{P}(Y_t)} [D_t^1(y_t)]) \right. \\ \left. - \mathbb{E}_{\pi_t \sim \mathbb{P}(\mathbf{h}(Y_t))} [D_t^1(\mathbf{f}(\pi_t))] \right) \\ + \max_{D_t^2} (\mathbb{E}_{\pi_t \sim \mathbb{P}(\mathbf{h}(Y_t))} [D_t^2(\pi_t)]) \\ \left. - \mathbb{E}_{\pi_{t-1} \sim \mathbb{P}(\mathbf{h}(Y_{t-1}))} [D_t^2(\mathbf{g}(\pi_{t-1}))] \right). \quad (19) \end{aligned}$$

And the smoothing objective in Eq. (17) becomes

$$\begin{aligned} \min_{\mathbf{h}^s} \sum_t \left(\max_{D_t^1} (\mathbb{E}_{y_t \sim \mathbb{P}(Y_t)} [D_t^1(y_t)]) \right. \\ \left. - \mathbb{E}_{\pi_t^s \sim \mathbb{P}(\mathbf{h}^s(Y_t))} [D_t^1(\mathbf{f}(\pi_t^s))] \right) \\ + \lambda_1 \max_{D_t^2} (\mathbb{E}_{\pi_t^s \sim \mathbb{P}(\mathbf{h}^s(Y_t))} [D_t^2(\pi_t^s)]) \\ \left. - \mathbb{E}_{\pi_t^f \sim \mathbb{P}(\mathbf{h}^f(Y_t))} [D_t^2(\pi_t^f)] \right) \\ + \lambda_2 \max_{D_t^3} (\mathbb{E}_{\pi_t^s \sim \mathbb{P}(\mathbf{h}^s(Y_t))} [D_t^3(\pi_t^s)]) \\ \left. - \mathbb{E}_{\pi_t^b \sim \mathbb{P}(\mathbf{h}^b(Y_t))} [D_t^3(\pi_t^b)] \right). \quad (20) \end{aligned}$$

In traditional implicit generative models, given a random variable z with a fixed distribution $p(z)$, we can pass it through a parametric generator G_θ (typically a neural network) which directly generates samples following a certain distribution \mathbb{P}_θ . Such design is of high flexibility, as by varying the parameters θ of the neural networks, we can change this distribution to any distribution of interest. While in our framework, we need a dynamic generative model to match distributions at each time step which can be regarded as a combination of several Generative Adversarial Networks (GANs) (Goodfellow et al., 2014). Specifically, functions \mathbf{f} , \mathbf{g} , \mathbf{h} are all generators (sharing parameters over time) and D_t is a discriminator at time

t . We formulate functions \mathbf{f} , \mathbf{g} and D_t as normal feed-forward neural networks²:

$$f^l = \sum_k \sigma(w_k^f f^{l-1} + b_k^f), \quad (21)$$

$$g^l = \sum_k \sigma(w_k^g g^{l-1} + b_k^g), \quad (22)$$

$$D_t^l = \sum_k \sigma(w_k^{D_t} D_t^{l-1} + b_k^{D_t}), \quad (23)$$

where f^l , g^l , D_t^l are the l -th layers of the neural networks, σ is the activation function, and w_k , b_k are the neural network parameters. Note the output layer of the discriminator D_t only has one single neuron to output a scalar value.

As for the function \mathbf{h} (in both filtering and smoothing), we use a recurrent neural network (RNN) to model it, similar to (Mogren, 2016). For the purposes of simplicity and clarity of exposition, we illustrate the computational process here using a vanilla RNN, whereas the actual recursive unit used in our experiments is LSTM unit (Hochreiter & Schmidhuber, 1997). Given inputs as sequences of observations $\{Y_0, Y_1, \dots, Y_t, \dots, Y_T\}$ and outputs as filtering or smoothing hidden states $\{\pi_0, \pi_1, \dots, \pi_t, \dots, \pi_T\}$, the parameterization function \mathbf{h} works as follows:

$$s_t = \sigma(AY_t + Bs_{t-1} + b), \quad (24)$$

$$\pi_t = \sigma(Cs_t + b), \quad (25)$$

where s_t is the memorized history information and A, B, C are parameter matrices of RNN.

Note that we need to enforce the Lipschitz constraints when solving Wasserstein distance from duality in Eq. (5). To achieve this, we adopt the strategy of gradient penalty in (Gulrajani et al., 2017) to regularize the Wasserstein distance³.

For parameter learning of dynamics in Eq. (18), we can obtain the optimal discriminator and gradients by

$$D_t^* = \arg \max_{D_t} (E_{y_t^i} [D_t(y_t^i)] - E_{x_{t-1}^i} [D_t(\mathbf{f} \circ \mathbf{g})(x_{t-1}^i)]) \quad (26)$$

$$g_{\mathbf{f}, \mathbf{g}} = - \sum_t \nabla_{\mathbf{f}, \mathbf{g}} E_{x_{t-1}^i} [D_t^*((\mathbf{f} \circ \mathbf{g})(x_{t-1}^i))], \quad (27)$$

where the gradient of \mathbf{g} needs to back propagate through the entire chain. In practice, we use gradient decent to update the discriminator. The parameter learning procedure of our model is presented in Algorithm 1. Similarly, we can derive the results for filtering and smoothing from Eq. (19) and (20), respectively.

² \mathbf{g} could be several nested \mathbf{g} due to the Δt in SDE.

³We omit this term in our equations for simplicity.

Algorithm 1 Parameter learning of dynamics

```

for # training iterations do
  for  $k$  steps do
    Sample  $\{\varepsilon^i\}_{i=1}^N \sim \mathbb{P}(\varepsilon)$ 
    for time  $t$  in  $[0:T]$  do
      Sample  $\{y_t^i\}_{i=1}^N \sim \mathbb{P}(Y_t)$ 
      for  $i=1$  to  $N$  do
         $x_0^i = \varepsilon^i$ ,
         $x_{t+\Delta t}^i = x_t^i + g(x_t^i)\Delta t + \Sigma^{1/2}\mathcal{N}(0, \Delta t)$ 
      end for
    end for
    Update the discriminator  $D_t$  by
     $\nabla_{D_t} \frac{1}{N} \sum_{i=1}^N D_t(y_t^i) - \nabla_{D_t} \frac{1}{N} \sum_{i=1}^N D_t((\mathbf{f} \circ \mathbf{g})(x_{t-1}^i))$ 
    end for
    Update  $\mathbf{f}, \mathbf{g}$  by ascending its stochastic gradient
     $-\sum_t \nabla_{\mathbf{f}, \mathbf{g}} \frac{1}{N} \sum_{i=1}^N D_t((\mathbf{f} \circ \mathbf{g})(x_{t-1}^i))$ 
  end for

```

5 THEORETICAL ANALYSIS

In this section, we provide a generalization error analysis and a convergence guarantee for our learning framework. Our analysis mainly focuses on the parameter learning component of our method, however, similar results can also be derived for filtering and smoothing based inference. For the purpose of simplicity, we briefly present our main results here and leave detailed theorems and proofs to Appendix A.

Generalization Error. We denote \mathcal{F} and \mathcal{G} as the function spaces of \mathbf{f} and \mathbf{g} , respectively, and the \mathcal{D} as the function space of the $\{D_t\}_{t=0}^T$, and $\mathbf{g}^{ot}(x, \xi_t) = \underbrace{((I + \mathbf{g}) \circ (I + \mathbf{g}) \circ \dots \circ (I + \mathbf{g}))}_t(x) + \xi_t$ with $\xi_t \sim \mathcal{N}(0, \Delta t)$. We define

$$\ell(\mathbf{f}, \mathbf{g}) = \mathbb{E}_{y_{0:T}, x_0 \sim p(x), \xi_{0:T}} \left[\sum_{t=0}^T \max_{D_t \in \mathcal{D}} \left[D_t(y_t) - D_t((\mathbf{f} \circ \mathbf{g}^{ot}(x_0, \xi_t))) \right] \right]. \quad (28)$$

Theorem 1. *Without loss of generality, we assume in each timestamp the number of the observations is N . Given the samples $\mathcal{Y} = \{(y_t^i)_{i=1}^N\}_{t=0}^T$ ($|\mathcal{Y}|_\infty = C_Y$) where $y_{0:T} = (y_t^i)_{i=1}^N$ are sampled i.i.d. from the underline stochastic processes, and $\mathcal{X} = \{x_0^i\}_{i=1}^N$, $\Xi = \{\xi_{0:T}^i\}_{i=1}^N$ are also i.i.d. sampled. Assume \mathcal{D} is a subset of k -Lipschitz functions and denote the $\mathfrak{R}(\mathcal{F} \circ \mathcal{G}^{ot})$ as the Rademacher complexity of the function space $\mathcal{F} \circ \mathcal{G}^{ot}$. We*

have

$$\frac{1}{T} \ell(\mathbf{f}, \mathbf{g}) \leq \frac{1}{T} \hat{\ell}(\mathbf{f}, \mathbf{g}) + \frac{4kC}{\sqrt{N}} + 4k \frac{\sum_{i=1}^T \mathfrak{R}(\mathcal{F} \circ \mathcal{G}^{ot})}{T}. \quad (29)$$

For the different parametrizations, *i.e.*, different function spaces \mathcal{F} and \mathcal{G} , the Rademacher complexity of $\mathfrak{R}(\mathcal{F} \circ \mathcal{G}^{ot})$ will be different. For example, if we parametrize the $\mathbf{f}(z) = \sigma(W_f z)$ and $\mathbf{g}(z, \xi) = I^\top[z, \xi] + \sigma(W_g z)$ as single layer neural networks, where σ satisfies some mild condition (Bartlett et al., 2017), following (Bartlett et al., 2017), we have the $\mathfrak{R}(\mathcal{F} \circ \mathcal{G}^{ot}) = \tilde{O}\left(\frac{\sqrt{C_1(W_f)C_2^t(I, W_g)(\sum_{i=1}^t C_3(I, W_g))}}{N}\right)$, where C_1 , C_2 and C_3 are some constants related to the parameters. For the details of the conditions on σ and the exact formulation of the constants, please refer to (Bartlett et al., 2017).

Convergence Analysis. Convexity-concavity no longer holds for the objectives in the learning and inference parts, therefore, convergence analysis for convex-concave saddle point problem in (Nemirovski et al., 2009) cannot be directly applied. Inspired by (Dai et al., 2017), we can see that once we obtain D_t^* , Algorithm 1 can be understood as a special case of stochastic gradient descent for a non-convex problem. Thus, we have the following finite-step convergence guarantee for our framework.

Theorem 2. *Assume that the parametrized empirical loss function $\hat{\ell}(\mathbf{f}, \mathbf{g})$ is K -Lipschitz and variance of its stochastic gradient is bounded by ς^2 . Let the algorithm run M iterations with stepsize $\zeta = \min\{\frac{1}{K}, \frac{C'}{\varsigma\sqrt{M}}\}$ for some $C' > 0$ and output $(w_f^1, w_g^1), \dots, (w_f^M, w_g^M)$. Setting the candidate solution to be $w = (\hat{w}_f^M, \hat{w}_g^M)$ randomly chosen from $(w_f^1, w_g^1), \dots, (w_f^M, w_g^M)$ such that $P(w = w^j) = \frac{2\zeta - K\zeta^2}{\sum_{j=1}^M (2\zeta - K\zeta^2)}$, then it holds that*

$$\mathbb{E} \left[\left\| \nabla \hat{\ell}(\hat{\mathbf{f}}_w^M, \hat{\mathbf{g}}_w^M) \right\| \right] \leq \frac{KC^2}{M} + (C' + \frac{C}{C'}) \frac{\varsigma}{\sqrt{M}}, \quad (30)$$

where $C := \sqrt{2(\hat{\ell}(w_f^1, w_g^1) - \min \hat{\ell}(w_f, w_g))}/K$ represents the distance of the initial solution to the optimal solution.

6 EXPERIMENTS

In this section, we evaluate our LEGEND framework on various types of synthetic and real-world datasets.

Baselines: We compare our model with two recently proposed methods that learn dynamics directly from aggregate observations — modeling directly on Y_t using a SDE. The baselines differ from one other in their characterization of the drift term $\mathbf{g}(X_t)$ of the SDE. The two

baselines considered in our experiments are two representatives from parametric and non-parametric categories: 1) OU (Orstein-Uhlenbeck (Huang et al., 2016)): modeling the drift term using an Orstein-Uhlenbeck process (Gillespie, 1996), which is a stationary Gauss-Markov process with the drift term $\theta(\mu - x_t)$ (θ, μ are parameters); and 2) NN (Hashimoto et al., 2016): modeling the drift term using a neural network (NN) which is a sum of ramps.

6.1 Synthetic Data

We first assess our model on three synthetic datasets generated using the following three diffusion dynamics:

Synthetic-1:

$$\begin{aligned} x_0 &\sim \mathcal{N}(0, \Sigma_1), \\ x_{t+\Delta t} &= x_t + \frac{1}{4}x_t\Delta t + \mathcal{N}(0, \Sigma_0), \\ y_t &= 2x_t. \end{aligned} \quad (31)$$

Synthetic-2:

$$\begin{aligned} x_0 &\sim \mathcal{N}(0, \Sigma_2), \\ x_{t+\Delta t} &= x_t + (0.1x_t^2 + 0.5x_t)\Delta t + \mathcal{N}(0, \Sigma_0), \\ y_t &= \exp(x_t). \end{aligned} \quad (32)$$

Synthetic-3:

$$\begin{aligned} x_0 &\sim \mathcal{U}([-2, 2]), \\ x_{t+\Delta t} &= x_t + (0.5x_t + |x_t|)\Delta t + \mathcal{N}(0, \Sigma_0), \\ y_t &= \log |x_t|. \end{aligned} \quad (33)$$

Synthetic-1 is a linear dynamic on x_t with linear measurement y_t where x_0 are sampled from multivariate normal distributions with covariance matrix Σ_1 (diagonal elements are 0.04 and others are 0.032). A nonlinear dependency between x_t and y_t is formulated in *Synthetic-2*: a quadratic dynamic on x_t and an exponential dependency of y_t on x_t where x_0 are sampled from multivariate normal distributions with covariance matrix Σ_2 (diagonal elements are 0.01 and others are 0.008). In *Synthetic-3*, we test a more complex scenario: highly nonlinear dynamics on x_t with highly nonlinear measurement y_t where x_0 are sampled from a uniform distribution $\mathcal{U}([-2, 2])$. For each synthesized dynamic, we obtain x_t like $\{x_0, x_1, x_2, x_3\}$ every $5\Delta t$ time following a discretized SDE in Eq. (7), and generate 1000 samples at each time step from x_t out of which only 500 samples are chosen as observations y_t like $\{y_0, y_1, y_2, y_3\}$. We consider population evolution \mathbb{R}^d with three different dimensions: $d = 2$, $d = 5$ and $d = 10$. Note that Δt is set to 0.2 for all datasets. The stochastic terms are all sampled from multivariate

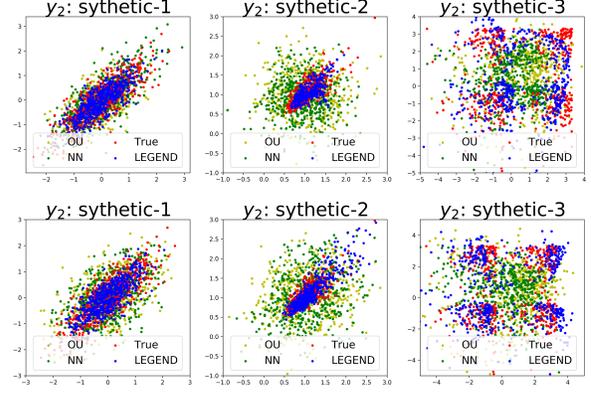


Figure 2: The true and predicted distributions for different models in filtering (top row) and smoothing (bottom row) based inference tasks on 2-dimensional synthetic-1 (left column), synthetic-2 (middle column) and synthetic-3 (right column) datasets.

normal distributions with covariance matrix Σ_0 (diagonal elements are 0.0025 and others are 0.002).

Our proposed model along with the baselines are evaluated on two tasks: 1) filtering based inference: given observations y_0 and y_1 , the task is to predict y_2 ; and 2) smoothing based inference: given observations y_0, y_1 and y_3 , the task is to predict y_2 .

Experimental Setup: For our LEGEND model, we set $D, \mathbf{f}, \mathbf{g}$ as a four-layer, two-layer and four-layer feed-forward neural network respectively with ReLU (Glorot et al., 2011) activation function, and set \mathbf{h} as a one-layer RNN with LSTM unit. In terms of training, we use the Adam optimizer (Kingma & Ba, 2014) with learning rate 10^{-4} , $\beta_1 = 0.5$ and $\beta_2 = 0.9$. The baselines OU and NN are configured with respect to their settings in the original papers without using pre-training (Hashimoto et al., 2016).

Results: We first show the capability of our model for learning low-dimensional ($d = 2$) diffusion dynamics. As visualized in Figure 2, given $\{y_0\}$ and $\{y_1\}$, our model can precisely learn the dynamics and correctly predict y_2 (top row) where a better match was observed between the predictions (blue points) and the ground truth (red points). Note that the dynamics on y_t become more and more complicated from synthetic-1 to synthetic-3. It can be seen from Figure 2 that our model works well on both simple and complex dynamics, while baselines OU and NN only work well on simple dynamics. Similar results are also observed in the smoothing based inference task, as shown in the bottom row of Figure 2.

We then evaluate our model using Wasserstein error for both low-dimensional ($d = 2$) and high-dimensional ($d =$

Table 1: The Wasserstein error of different models on synthetic-1/2/3 (Syn-1/2/3), RNA-seq (RNA) and bird migration (Bird) datasets. The best results are highlighted in **bold**.

Data	Target	Task	NN	OU	LEG- END
Syn-1	filtering y_2	$d = 2$	0.30	0.29	0.06
		$d = 5$	3.09	2.52	0.06
		$d = 10$	11.19	9.61	0.18
	smoothing y_2	$d = 2$	0.70	0.80	0.04
		$d = 5$	3.40	2.92	0.08
		$d = 10$	9.58	8.96	0.12
Syn-2	filtering y_2	$d = 2$	0.87	1.36	0.17
		$d = 5$	3.49	4.38	0.47
		$d = 10$	8.55	10.42	1.37
	smoothing y_2	$d = 2$	1.62	1.75	0.22
		$d = 5$	5.28	4.17	0.57
		$d = 10$	11.14	9.91	2.84
Syn-3	filtering y_2	$d = 2$	8.55	10.79	3.79
		$d = 5$	31.95	35.17	13.21
		$d = 10$	113.21	116.42	42.52
	smoothing y_2	$d = 2$	8.43	9.08	2.22
		$d = 5$	28.37	31.26	11.26
		$d = 10$	102.65	109.80	38.73
RNA	Krt8	D7	6.16	9.82	2.31
		D4	27.98	24.54	4.89
	Krt18	D7	6.86	9.80	3.16
		D4	24.75	25.88	4.21
Bird	GrayJay	June	1.9e3	2.5e3	1.2e3
		April	1.5e3	1.1e3	0.3e3

5, 10) diffusion dynamics. Wasserstein error measures the difference between predicted distribution and the true distribution. As reported in Table 1, our model achieves much lower Wasserstein error than the two baselines on all the 3 datasets for 2/5/10-dimensional dynamics. The poor performance of OU and NN may due to the fact that y_t becomes more and more complicated as dimension increases on all three datasets. The superior performance of our model verifies the importance of hidden variables — they are necessary for the modeling of complex nonlinear dynamics and complex measurements of hidden states.

6.2 Real Data: Single-cell RNA-seq

In this section, we evaluate our model on a typical application of distribution based continuous diffusion dynamics in biology: learning the diffusion process where embryonic stem cells differentiate into mature cells (Klein et al., 2015). The cell population begins to differentiate from embryonic stem cells after the removal of LIF (leukemia inhibitory factor) at day 0 (D0). Single-cell RNA-seq measurements (or observations) are sampled at day 0 (D0), day 2 (D2), day 4 (D4), and day 7 (D7). At each time

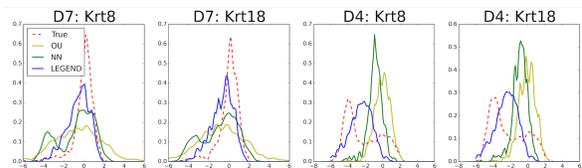


Figure 3: The true and predicted marginal distributions of the differentiating genes at D7 (filtering based inference task) and D4 (smoothing based inference task) for different models.

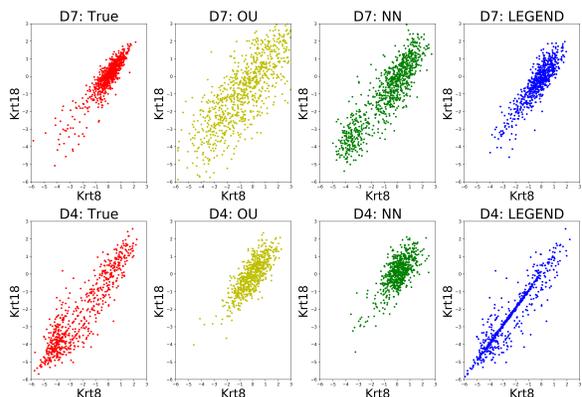


Figure 4: The true (left column) and predicted (right 3 columns) correlations between Krt8 and Krt18 at D7 (top row) and D4 (bottom row). The closer to the true correlation the better the performance.

point, the expression of 24,175 genes for several hundreds cells are measured (933, 303, 683 and 798 cells at D0, D2, D4, and D7 respectively). We focus on the dynamics of cell differentiation for the two main epithelial makers studied in (Klein et al., 2015), *i.e.*, Keratin 8 (Krt8) and Keratin 18 (Krt18). We evaluate two tasks on this data: 1) filtering based inference: predicting the gene expression level at D7 given only the observations at D0 and D4; and 2) smoothing based inference: predicting the gene expression level at D4 given D0, D2 and D7.

Experimental Setup: We set f as a one-layer feed-forward neural network, g as a three-layer feed-forward neural network and h as a one-layer RNN with LSTM neurons. For preprocessing, we apply standard normalization procedures (Hicks et al., 2015) to correct for batch effects, and impute missing expression levels using non-negative matrix factorization, similarly as it did in (Hashimoto et al., 2016). The stochastic term Σ in SDE are sampled from multivariate normal distributions with diagonal covariance matrix (diagonal elements are 1). Other configurations and baselines are the same as those in Section 6.1.

Results: We first show in Table 1 that compared to other baselines our model achieves the lowest Wasserstein error in both filtering (D7) and smoothing (D4) tasks on both Krt8 and Krt18. This proves that our model is capable of learning the precise differentiation dynamics and the distributions of the two studied gene expressions. We further provide a closer look into the learned distributions of the two genes in Figure 3. As can be seen, the distributions of Krt8 and Krt18 predicted by our model (curves in blue) are much closer to their true distributions (curves in red) at both D4 and D7, as compared to the baseline models. Moreover, our model can effectively identify the correlations between Krt8 and Krt18, as shown in Figure 4. This implies that our model can accurately learn the dynamics even considering the correlational structure of the true dynamics.

6.3 Real Data: Bird Migration

We also evaluate our model on another typical application of distribution based diffusion dynamics: bird migration research in ecology. We use the eBird basic dataset (EBD), which gathers large volumes of information on where and when birds occur in the world (Sullivan et al., 2009). We down-sampled EBD to only include the tracking records for the species GrayJay between January, 2017 and June, 2017 (monthly data) at United States where 400 samples are randomly selected as observations for each month. There are again two tasks evaluated here: 1) filtering based inference: we apply our model on the months February and April so as to predict the population at June; 2) smoothing based inference: we apply our model on months February, March and June so as to predict the population at April. The experimental setups are the same as those in Section 6.2.

Results: We plot the true and predicted locations (longitude and latitude) of the species GrayJay in Figure 5, and report the Wasserstein error⁴ in Table 1. Again, our model achieves the lowest Wasserstein error in both filtering (June) and smoothing (April) based inference tasks. The evolving dynamics of bird migration can be very complicated and extremely difficult to learn, mostly because bird migration could be affected by many irregular factors related to the specific time. Even so, with the introduction of the hidden state variable, our model can predict locations which are close to the ground truth, and with better performance than OU and NN which directly build models on observations. This result demonstrates the advantages of our framework in solving real-world problems involving complicated diffusion dynamics.

⁴Our model could be further improved if considering more complex hidden diffusion process, *e.g.*, jump diffusion process, but the framework is the same to this paper.

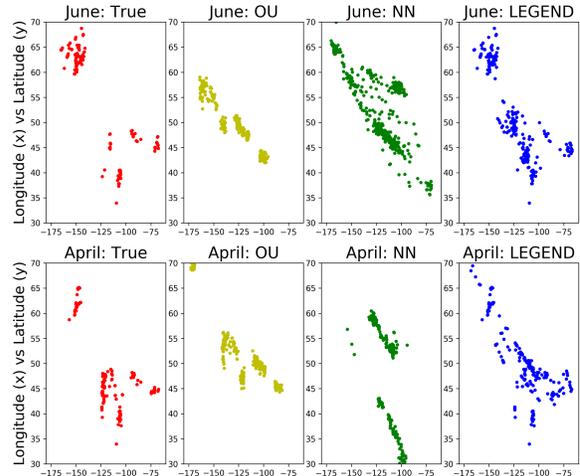


Figure 5: The true (left column) and predicted (right 3 columns) distributions of GrayJay species at month June (filtering based inference task) and April (smoothing based inference task) for different models.

7 CONCLUSIONS

In this paper, we formulated a novel technique to learn nonlinear continuous diffusion dynamics from *aggregate observations*. In particular, we showed how one can model dynamics as a *hidden continuous stochastic process*, and proposed a framework that employs a dynamic generative model with Wasserstein distance to learn the evolving dynamics. In addition to deriving solutions for both filtering and smoothing based inference tasks, we also established theoretical guarantees on the generalization and convergence properties of our framework. Through comprehensive experimental evaluation on synthetic and real-world datasets, we demonstrated that our approach has very strong performance compared to state-of-the-art techniques on both filtering and smoothing based inference tasks.

Acknowledgements

This work is partially supported by NSF IIS-1717916, CIMM-1745382 and China Scholarship Council (CSC).

References

- Agueh, Martial and Carlier, Guillaume. Barycenters in the wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.
- Arjovsky, Martin, Chintala, Soumith, and Bottou, Léon. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.

- Banks, HT and Potter, Laura K. Probabilistic methods for addressing uncertainty and variability in biological models: Application to a toxicokinetic model. *Mathematical Biosciences*, 192(2):193–225, 2004.
- Bartholomew, Aaron and Bohnsack, James A. A review of catch-and-release angling mortality with implications for no-take reserves. *Reviews in Fish Biology and Fisheries*, 15(1-2):129–154, 2005.
- Bartlett, Peter L, Foster, Dylan J, and Telgarsky, Matus J. Spectrally-normalized margin bounds for neural networks. In *NIPS*, 2017.
- Briers, Mark, Doucet, Arnaud, and Maskell, Simon. Smoothing algorithms for state–space models. *Annals of the Institute of Statistical Mathematics*, 62(1):61, 2010.
- Dai, Bo, He, Niao, Dai, Hanjun, and Song, Le. Provable bayesian inference via particle mirror descent. In *AISTATS*, 2016.
- Dai, Bo, Shaw, Albert, Li, Lihong, Xiao, Lin, He, Niao, Chen, Jianshu, and Song, Le. Smoothed dual embedding control. *arXiv preprint arXiv:1712.10285*, 2017.
- Desbouvries, François, Petetin, Yohan, and Ait-El-Fquih, Boujemaa. Direct, prediction- and smoothing-based kalman and particle filter algorithms. *Signal Processing*, 91(8):2064–2077, 2011.
- Djuric, Petar M, Kotecha, Jayesh H, Zhang, Jianqui, Huang, Yufei, Ghirmai, Tadesse, Bugallo, Mónica F, and Miguez, Joaquin. Particle filtering. *IEEE Signal Processing Magazine*, 20(5):19–38, 2003.
- Eddy, Sean R. Hidden markov models. *Current Opinion in Structural Biology*, 6(3):361–365, 1996.
- Ghadimi, Saeed and Lan, Guanghui. Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- Gillespie, Daniel T. Exact numerical simulation of the ornstein-uhlenbeck process and its integral. *Physical Review E*, 54(2):2084, 1996.
- Glorot, Xavier, Bordes, Antoine, and Bengio, Yoshua. Deep sparse rectifier neural networks. In *AISTATS*, 2011.
- Goodfellow, Ian, Pouget-Abadie, Jean, Mirza, Mehdi, Xu, Bing, Warde-Farley, David, Ozair, Sherjil, Courville, Aaron, and Bengio, Yoshua. Generative adversarial nets. In *NIPS*, 2014.
- Gulrajani, Ishaan, Ahmed, Faruk, Arjovsky, Martin, Dumoulin, Vincent, and Courville, Aaron C. Improved training of wasserstein gans. In *NIPS*, 2017.
- Harvey, Andrew C. *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press, 1990.
- Hashimoto, Tatsunori, Gifford, David, and Jaakkola, Tommi. Learning population-level diffusions with generative rnns. In *ICML*, 2016.
- Hefny, Ahmed, Downey, Carlton, and Gordon, Geoffrey J. Supervised learning for dynamical system learning. In *NIPS*, 2015.
- Hicks, Stephanie C, Teng, Mingxiang, and Irizarry, Rafael A. On the widespread and critical impact of systematic bias and batch effects in single-cell rna-seq data. *BioRxiv*, 2015.
- Hochreiter, Sepp and Schmidhuber, Jürgen. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- Huang, Gang, Jansen, HM, Mandjes, Michel, Spreij, Peter, and De Turck, Koen. Markov-modulated ornstein-uhlenbeck processes. *Advances in Applied Probability*, 48(1):235–254, 2016.
- Kingma, Diederik P and Ba, Jimmy. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kitagawa, Genshiro. The two-filter formula for smoothing and an implementation of the gaussian-sum smoother. *Annals of the Institute of Statistical Mathematics*, 46(4):605–623, 1994.
- Klein, Allon M, Mazutis, Linas, Akartuna, Ilke, Tallapragada, Naren, Veres, Adrian, Li, Victor, Peshkin, Leonid, Weitz, David A, and Kirschner, Marc W. Droplet barcoding for single-cell transcriptomics applied to embryonic stem cells. *Cell*, 161(5):1187–1201, 2015.
- Langford, John, Salakhutdinov, Ruslan, and Zhang, Tong. Learning nonlinear dynamic models. In *ICML*, 2009.
- Mogren, Olof. C-rnn-gan: Continuous recurrent neural networks with adversarial training. *arXiv preprint arXiv:1611.09904*, 2016.
- Nemirovski, Arkadi, Juditsky, Anatoli, Lan, Guanghui, and Shapiro, Alexander. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- Sullivan, Brian L, Wood, Christopher L, Iliff, Marshall J, Bonney, Rick E, Fink, Daniel, and Kelling, Steve. ebird: A citizen-based bird observation network in the biological sciences. *Biological Conservation*, 142(10):2282–2292, 2009.
- Talay, Denis. Numerical solution of stochastic differential equations. 1994.
- Villani, Cédric. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.