Generative Modeling of Irregular Time Series via SDE-Induced Continuous-Discrete Variational Inference

Zexin Yuan[1], Qinliang Su[1,2*], Junxi Xiao[1]

[1] School of Computer Science and Engineering, Sun Yat-sen University
[2] Guangdong Key Laboratory of Big Data Analysis and Processing

Background - Modeling Irregular Time Series

- **Irregular:** Observations occur at **unequal intervals**
- **Downstream Applications:** Finance, IoT, Healthcare.
- **Challenges in Traditional Methods:**
  - Fail to model continuous-time dynamics.
  - Assume dynamic determinism and uniqueness.
- This motivates the use of **Stochastic Differential Equations(SDEs)** to capture complex dynamics and uncertainty.

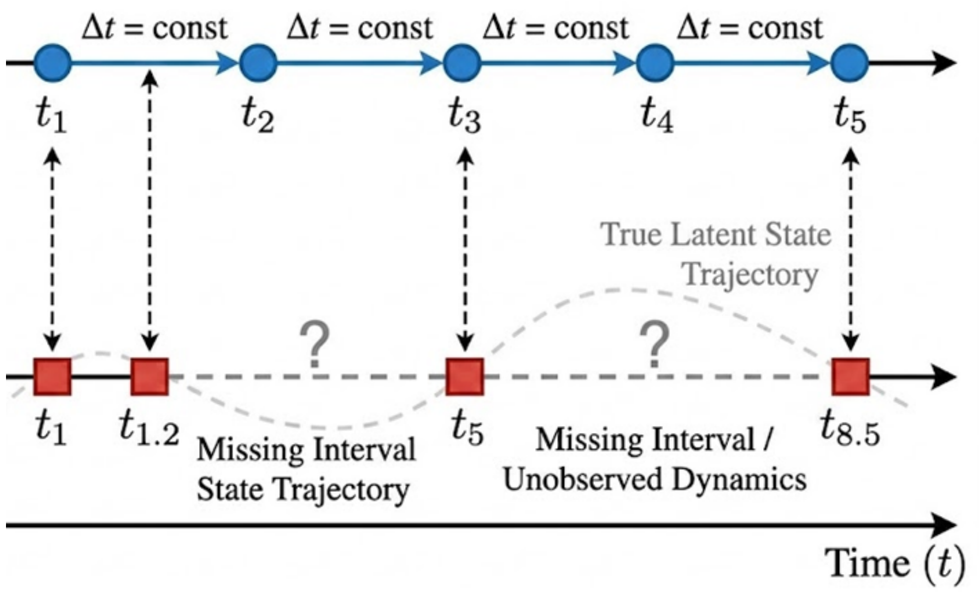

Path-based Variational Inference

- **Objective(Path-Based ELBO): Infer posterior paths over the entire continuous interval $[0,T]$.**

$$\mathscr{L}_{path}(\theta, \phi) = \mathbb{E}_{\mathbb{Q}_\phi}\left[\sum_{i=0}^{N} \log p_\theta(x_{t_i} | z_{t_i})\right] - D_{KL}(\mathbb{Q}_\phi || \mathbb{P}_\theta)$$

- **Critical Limitations:**
  - Requires **expensive numerical simulations** over continuous stochastic paths, also **unstable.**
  - Linear restrictions to enable analytic solving cause **severe expressivity bottlenecks.**
  - **Non-Markovian posterior process.**

Marginal-based Variational Inference

- **The Paradigm Shift**: From learning full trajectories to learning **discrete distributions.**
- **Marginal-Based ELBO:**

$$\mathscr{L}_{marg}(\theta, \phi) = \mathbb{E}_{q_\phi}\left[\sum \log p_\theta(x_{t_i} | z_{t_i})\right] - D_{KL}(q_\phi || p_\theta)$$

- **Theorem 4.1:**

$$\mathscr{L}_{marg}(q) \geq \mathscr{L}_{path}(\mathbb{Q})$$

Marginal-based ELBO provides a strictly tighter lower bound

SDE-VI

- The challenge lies in ensuring the joint distribution over discrete timestamps is governed by an underlying SDE and maintaining its **expressivity and efficiency**.
- Parameterize discrete Gauss-Markov distribution $\{\mu_{t_i}, \sigma_{t_i}, \gamma_{t_i,t_{i+1}}\}_{i=0}^{N}$
  - **Theorem 4.2**: $\gamma_{t_i,t_{i+1}} = \exp(-\bar{s} \cdot \Delta t_{i,i+1})$ guarantees induction by a valid LTV SDE.
- Time-dependent Conditional Normalizing Flows progressively map the unimodal base to a complex, time-evolving multimodal density.
  - **Proposition 4.3**: The refined process is characterized as **a non-linear Itô SDE.**
- Generalizes to the complex domain to model oscillatory dynamics.
  - **Theorem 4.4**: Predicts only one additional parameter (average angular frequency $\bar{\omega}_i$)

Results

- Datasets: Human Activity, Pendulum, Physionet, USHCN
- Metrics: Accuracy, MSE, NLL

Table 2. Test MSE ($\times 10^{-2}$) for inter/extra-polation on Physionet and USHCN. Baselines are taken from (Park et al., 2025).

	Interpolation		Extrapolation	
Model	Physionet	USHCN	Physionet	USHCN
mTAND	0.208 ± 0.025	1.766 ± 0.009	0.340 ± 0.020	2.360 ± 0.038
RKN-Δ_t	0.186 ± 0.030	0.009 ± 0.002	0.703 ± 0.050	1.491 ± 0.272
GRU-Δ_t	0.271 ± 0.057	0.090 ± 0.059	0.870 ± 0.077	2.081 ± 0.054
GRU-D	0.338 ± 0.027	0.944 ± 0.011	0.873 ± 0.071	1.718 ± 0.015
Latent-ODE	0.212 ± 0.027	1.798 ± 0.009	0.725 ± 0.072	2.034 ± 0.005
ODE-RNN	0.236 ± 0.009	0.831 ± 0.008	0.467 ± 0.006	1.955 ± 0.466
GRU-ODE-B	0.521 ± 0.038	0.841 ± 0.142	0.798 ± 0.071	5.437 ± 1.020
CRU	0.182 ± 0.091	0.016 ± 0.006	0.629 ± 0.093	1.273 ± 0.066
ACSSM	0.116 ± 0.011	$\mathbf{0.006 \pm 0.001}$	0.627 ± 0.019	0.941 ± 0.014
SDE-VI (Ours)	$\mathbf{0.052 \pm 0.002}$	0.044 ± 0.012	$\mathbf{0.335 \pm 0.006}$	$\mathbf{0.455 \pm 0.003}$

Results

- Datasets: Human Activity, Pendulum, Physionet, USHCN
- Metrics: Accuracy, MSE, NLL
- By formulating the recursive sampling as an associative scan, we achieve $O(\log N)$ parallel complexity.

Table 1. Accuracy (%) for Human Activity Classification and Test MSE($\times 10^{-3}$) for Pendulum Regression. †Our implementation. Baseline results are taken from (Park et al., 2025).

Model	Activity Acc (↑)	Pendulum MSE (↓)
Latent-ODE	87.0 ± 2.8	15.70 ± 0.29
CRU	$86.0 \pm 1.6^{\dagger}$	4.63 ± 1.07
Latent-SDE$_{\mathrm{H}}$	90.6 ± 0.4	3.84 ± 0.35
mTAND	91.1 ± 0.2	3.20 ± 0.60
ACSSM	91.4 ± 0.4	2.98 ± 0.30
SDE-VI (Ours)	$\mathbf{91.9 \pm 0.7}$	$\mathbf{2.68 \pm 0.17}$

Table 3. Training runtime (s/epoch) and speedup relative to CRU.

Model	USHCN		Physionet	
	Time (↓)	Speedup (↑)	Time (↓)	Speedup (↑)
RKN-Δ_t	32.1	1.3x	114.9	2.6x
GRU-D	99.6	0.4x	2457.0	0.1x
Latent-ODE	37.5	1.1x	560.0	0.5x
ODE-RNN	27.6	1.5x	295.4	1.0x
GRU-ODE-B	132.6	0.3x	527.7	0.6x
CRU	41.6	1.0x	302.7	1.0x
ACSSM	21.6	1.9x	78.6	3.9x
SDE-VI (Ours)	9.3	4.5x	40.2	7.5x
w/o Flow	9.6	4.3x	**37.2**	**8.1x**
w/o Complex	**9.0**	**4.6x**	38.8	7.8x

Takeaways

- **Methodological Paradigm Shift**
  - Shift from path-based integration to marginal-based inference
- **Theoretical Guarantees**
  - Discrete distributions strictly induced by valid LTV and nonlinear SDEs.
- **Expressivity & Scalability**
  - Captures nonlinear, multimodal, and oscillatory dynamics.
  - $O(\log N)$ parallel scan and sliding-window compatibility for ultra-long sequences.