Novelty Assessment Report

Paper: Latent Stochastic Interpolants
PDF URL: https://openreview.net/pdf?id=txiGUfI4yF
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2025-12-27

Abstract

Stochastic Interpolants (SI) are a powerful framework for generative modeling, capable of flexibly transforming between two probability distributions. However, their use in jointly optimized latent variable models remains unexplored as they require direct access to the samples from the two distributions. This work presents Latent Stochastic Interpolants (LSI) enabling joint learning in a latent space with end-to-end optimized encoder, decoder and latent SI models. We achieve this by developing a principled Evidence Lower Bound (ELBO) objective derived directly in continuous time. The joint optimization allows LSI to learn effective latent representations along with a generative process that transforms an arbitrary prior distribution into the encoder-defined aggregated posterior. LSI sidesteps the simple priors of the normal diffusion models and mitigates the computational demands of applying SI directly in high-dimensional observation spaces, while preserving the generative flexibility of the SI framework. We demonstrate the efficacy of LSI through comprehensive experiments on the standard large scale ImageNet generation benchmark.

Core Task Landscape

This paper addresses: **Generative Modeling in Latent Space with Continuous Time Dynamics**

A total of **50 papers** were analyzed and organized into a taxonomy with **20 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Continuous-Time Latent Dynamical Systems**
- **Diffusion Models in Latent Space**
- **Latent Stochastic Interpolants and Generative Bridges**
- **Spatial and Graph Dynamics**
- **Domain-Specific Applications**
- **Discrete-Time and Hybrid Approaches**
- **Inference and Optimization Methods**

Complete Taxonomy Tree

- Generative Modeling in Latent Space with Continuous Time Dynamics Survey Taxonomy
- Continuous-Time Latent Dynamical Systems
  - Neural ODE-Based Generative Models
  - Latent ODE for Time Series (5 papers)
    - [1] Generative learning for nonlinear dynamics (Gilpin, 2024) View paper
    - [6] Path-minimizing latent ODEs for improved extrapolation and inference (Matt L Sampson, 2025) View paper
    - [8] Continuous latent process flows (Deng RuiZhi, 2021) View paper
    - [31] Time-aware neural ordinary differential equations for incomplete time series modeling (Zhuoqing Chang, 2023) View paper
    - [38] Deep latent state space models for time-series generation (Zhou, 2023) View paper
  - Neural ODE for Image and Visual Synthesis (3 papers)
    - [13] Modeling GAN Latent Dynamics using Neural ODEs (W Xia, 2023) View paper
    - [28] ODE-based generative modeling: Learning from a single natural image (Jian Yue, 2025) View paper
    - [40] Learning a generative motion model from image sequences based on a latent motion matrix (Julian Krebs, 2021) View paper
  - Stochastic Latent Dynamics
  - SDE-Based Generative Frameworks (4 papers)
    - [2] Generative modeling of neural dynamics via latent stochastic differential equations (El-Gazzar Ahmed, 2024) View paper
    - [23] Multivariate stochastic modeling for transcriptional dynamics with cell-specific latent time using SDEvelo (Xu Liao, 2024) View paper
    - [29] Generative time series models with interpretable latent processes for complex disease trajectories (C Trottet, 2023) View paper
    - [37] Latent sdes on homogeneous spaces (Zeng, 2023) View paper
  - Stochastic Process Diffusion (1 papers)
    - [32] Modeling temporal data as continuous functions with stochastic process diffusion (BiloÅ¡, 2023) View paper
  - Flow-Based Generative Models (2 papers)
  - [22] Monge-AmpÃ¨re Flow for Generative Modeling (Zhang, 2018) View paper
  - [30] Manifold Interpolating Optimal-Transport Flows for Trajectory Inference (Guillaume Huguet, 2022) View paper
- Diffusion Models in Latent Space
  - Latent Diffusion for Static Data (3 papers)
  - [4] Devil is in the Details: Density Guidance for Detail-Aware Generation with Flow Models (Karczewski, 2025) View paper

    - [20] Exploit Your Latents: Coarse-Grained Protein Backmapping with Latent Diffusion Models (Rongchao Zhang, 2025) View paper
    - [49] Controllable and Compositional Generation with Latent-Space Energy-Based Models (Nie, 2021) View paper
    - Latent Diffusion for Temporal and Sequential Data (4 papers)
    - [3] Latent diffusion-based data augmentation for continuous-time dynamic graph model (Yuxing Tian, 2024) View paper
    - [7] Latent Conditional Diffusion-based Data Augmentation for Continuous-Time Dynamic Graph Model (Tian, 2024) View paper
    - [9] Extraction and recovery of spatio-temporal structure in latent dynamics alignment with diffusion models (Wang Yule, 2023) View paper
    - [41] Generative Pre-trained Autoregressive Diffusion Transformer (Zhang Yuan, 2025) View paper
    - Consistency and Few-Step Latent Diffusion (2 papers)
    - [27] AudioLCM: Text-to-Audio Generation with Latent Consistency Models (Liu, 2024) View paper
    - [50] AudioLCM: Efficient and High-Quality Text-to-Audio Generation with Minimal Inference Steps (Huadai Liu, 2024) View paper
    - Unified and Theoretical Diffusion Frameworks (3 papers)
    - [21] A survey on diffusion models for time series and spatio-temporal data (Yiyuan Yang, 2024) View paper
    - [35] The principles of diffusion models (Lai, 2025) View paper
    - [43] Unified Continuous Generative Models (Sun Peng, 2025) View paper
- Latent Stochastic Interpolants and Generative Bridges ★ (2 papers)
    - [0] Latent Stochastic Interpolants (Anon et al., 2026) View paper
    - [11] ARTEMIS integrates autoencoders and SchrÃ¶dinger Bridges to predict continuous dynamics of gene expression, cell population, and perturbation from time-series single-cell data (Sayali Anil Alatkar, 2025) View paper
- Spatial and Graph Dynamics
    - Continuous-Time Dynamic Graphs (2 papers)
    - [5] TG-GAN: Continuous-time temporal graph generation with deep generative models (Zhang Liming, 2020) View paper
    - [33] Continuous-time generative graph neural network for attributed dynamic graphs: student research abstract (Moallemy-Oureh, 2022) View paper
    - Spatial Dynamics and Relational Event Models (2 papers)
    - [15] Dynamic latent space relational event model (I Artico, 2023) View paper
    - [19] Spate-gan: Improved generative modeling of dynamic spatio-temporal patterns with an autoregressive embedding loss (Konstantin Klemmer, 2022) View paper
- Domain-Specific Applications
    - Biological and Biomedical Systems (2 papers)
    - [16] T-CACE: A Time-Conditioned Autoregressive Contrast Enhancement Multi-Task Framework for Contrast-Free Liver MRI Synthesis, Segmentation, and Diagnosis (Xiao, 2025) View paper
    - [39] STODE: A Deep Generative Framework for Continuous Spatiotemporal Dynamics in Spatial Transcriptomics (Koichiro Majima, 2025) View paper
    - Molecular and Material Science (2 papers)
    - [36] Molecular latent space simulators (Sidky, 2020) View paper
    - [48] Generative Latent Space Dynamics of Electron Density (Chiang Yuan, 2025) View paper
    - Physical and Environmental Systems (4 papers)
    - [10] Generating Unseen Nonlinear Evolution in Sea Surface Temperature Using a Deep Learning-Based Latent Space Data Assimilation Framework (Zheng Qing-yu, 2024) View paper
    - [24] Optimization of geological carbon storage operations with multimodal latent dynamic model and deep reinforcement learning (Zhongzheng Wang, 2024) View paper
    - [25] Evolve Smoothly, Fit Consistently: Learning Smooth Latent Dynamics For Advection-Dominated Systems (Wan, 2023) View paper
    - [42] Enabling dynamic 3D coherent diffraction imaging via adaptive latent space tuning of generative autoencoders (Alexander Scheinker, 2024) View paper
    - Specialized Generative Tasks (3 papers)
    - [14] Generative Plant Growth Simulation from Sequence-Informed Environmental Conditions (Mohamed Debbagh, 2024) View paper
    - [45] Gendef: Learning generative deformation field for video generation (WANG Wen, 2023) View paper
    - [47] Towards Scene Graph Anticipation (Rohith Peddi, 2024) View paper
- Discrete-Time and Hybrid Approaches
    - Hierarchical and Structured Latent Models (2 papers)
    - [17] Deep generative model with hierarchical latent factors for time series anomaly detection (Challu, 2022) View paper
    - [34] Latent space energy-based neural odes (Cheng Sheng, 2024) View paper
    - Discrete-Time Generative Models with Latent Dynamics (2 papers)
    - [12] Exploring Variational Autoencoders and Generative Latent Time-Series Models for Synthetic Data Generation and Forecasting (Dodda, 2024) View paper
    - [26] Anode-gan: incomplete time series imputation by augmented neural ode-based generative adversarial networks (Zhuoqing Chang, 2023) View paper
    - Representation Learning and Latent Space Analysis (1 papers)
    - [44] From latent dynamics to meaningful representations (Dedi Wang, 2024) View paper
- Inference and Optimization Methods (2 papers)
    - [18] GOKU-UI: Ubiquitous Inference through Attention and Multiple Shooting for Continuous-time Generative Models (German Abrevaya, 2023) View paper
    - [46] Contraction and entropy production in continuous-time Sinkhorn dynamics (Srinivasan Anand, 2025) View paper

Narrative

Core task: generative modeling in latent space with continuous time dynamics. This field addresses how to learn and sample from complex data distributions by evolving latent representations through continuous-time processes. The taxonomy reveals several major branches: Continuous-Time Latent Dynamical Systems focus on neural ODEs and SDEs that model temporal evolution in learned embeddings, often applied to irregular time series or physical systems. Diffusion Models in Latent Space adapt score-based generative methods to compressed representations, balancing sample quality with computational efficiency. Latent Stochastic Interpolants and Generative Bridges construct explicit transport maps or stochastic paths between distributions, offering principled frameworks for generation and interpolation. Spatial and Graph Dynamics extend these ideas to structured data such as dynamic graphs or spatiotemporal fields, while Domain-Specific Applications demonstrate successes in areas ranging from molecular simulation to audio synthesis. Discrete-Time and Hybrid Approaches blend continuous formulations with practical discretization strategies, and Inference and Optimization Methods tackle the computational challenges of training and sampling in these models.

A particularly active line of work explores how to design efficient interpolation schemes that minimize path complexity or enforce physical constraints, as seen in Path Minimizing ODEs[6] and Density Guidance Flow[4]. Another contrasting direction emphasizes flexible stochastic bridges that can handle irregular or multimodal data, exemplified by Neural Dynamics Latent SDEs[2] and ARTEMIS[11]. Latent Stochastic Interpolants[0] sits naturally within the Generative Bridges branch, proposing a framework for constructing stochastic paths in latent space that interpolate between data and noise distributions. Compared to ARTEMIS[11], which focuses on adaptive temporal modeling for irregular observations, Latent Stochastic Interpolants[0] emphasizes the design of principled interpolation dynamics that can be efficiently sampled. This work contributes to ongoing efforts to unify optimal transport perspectives with generative modeling, addressing trade-offs between theoretical guarantees, sampling speed, and the expressiveness of learned latent dynamics.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. ARTEMIS integrates autoencoders and SchrÃ¶dinger Bridges to predict continuous dynamics of gene expression, cell population, and perturbation from time-series single-cell data

Authors: Sayali Anil Alatkar, Daifeng Wang | **Year/Venue**: 2025 • Bioinformatics | **URL**: View paper

Abstract

Abstract Summary Cellular processes like development, differentiation, and disease progression are highly complex and dynamic (e.g. gene expression). These processes often undergo cell population changes driven by cell birth, proliferation, and death. Single-cell sequencing enables gene expression measurement at the cellular resolution, allowing us to decipher cellular and molecular dynamics underlying these processes. However, the high costs and destructive nature of sequencing restrict observa...

Relationship Analysis

Both papers belong to the Latent Stochastic Interpolants and Generative Bridges category, using stochastic processes to transform distributions in latent space. While the original paper (LSI) develops a general framework for generative modeling with joint encoder-decoder-latent model training using continuous-time SDEs and ELBO objectives for tasks like ImageNet generation, ARTEMIS applies Schrödinger bridges specifically to biological time-series single-cell data, incorporating unbalanced dynamics to model cell population changes and gene expression trajectories. The key difference is that LSI focuses on general-purpose image generation with flexible priors, whereas ARTEMIS specializes in modeling biological processes with population dynamics using domain-specific constraints.

Contributions Analysis

Overall novelty summary. The paper introduces Latent Stochastic Interpolants (LSI), a framework for jointly learning encoder, decoder, and latent generative dynamics via stochastic interpolation. It resides in the 'Latent Stochastic Interpolants and Generative Bridges' leaf, which contains only two papers total. This is a notably sparse research direction within the broader taxonomy of 50 papers across 20 leaf nodes, suggesting the specific combination of stochastic interpolants with end-to-end latent variable learning remains relatively unexplored compared to more crowded areas like latent diffusion for static data.

The taxonomy reveals several neighboring branches: 'Diffusion Models in Latent Space' (11 papers across four sub-leaves) focuses on score-based methods in learned representations, while 'Stochastic Latent Dynamics' (5 papers) emphasizes SDE-based frameworks with variational inference. 'Flow-Based Generative Models' (2 papers) explores optimal transport flows. LSI bridges these directions by combining stochastic interpolation—a transport-inspired approach—with joint latent space optimization, distinguishing itself from standard latent diffusion's fixed encoder-decoder paradigms and from pure SDE models that lack the interpolant formulation's explicit distribution-matching guarantees.

Among 24 candidates examined, the continuous-time ELBO contribution shows overlap: 2 of 4 examined papers provide refutable prior work, indicating this theoretical component has substantial precedent within the limited search scope. The LSI framework itself (10 candidates, 0 refutations) and the unifying perspective (10 candidates, 0 refutations) appear more novel relative to the examined literature. The analysis explicitly covers top-K semantic matches plus citation expansion, not an exhaustive survey, so these statistics reflect novelty within a focused but incomplete sample of related work.

Given the sparse taxonomy leaf and limited refutations for two of three contributions, the work appears to occupy a relatively underexplored niche. However, the ELBO derivation's overlap with prior continuous-time variational methods suggests incremental theoretical refinement rather than foundational innovation in that component. The assessment is constrained by the 24-paper search scope and may not capture all relevant precedents in adjacent communities.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Latent Stochastic Interpolants (LSI) framework

Description: The authors introduce LSI, a framework that enables end-to-end joint learning of an encoder, decoder, and generative model in an unobserved latent space with continuous-time dynamics. This extends Stochastic Interpolants to support jointly optimized latent variable models, which was previously not possible since SI requires direct access to samples from both distributions.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Social lode: human trajectory prediction with latent odes

URL: View paper

Brief Assessment

Social LODE[67] focuses on human trajectory prediction using latent ODEs within a VAE framework for sequential data. The original paper addresses generative modeling with stochastic interpolants in latent spaces for image generation, which is a fundamentally different application domain and technical approach.

---

2. Temporal latent auto-encoder: A method for probabilistic multivariate time series forecasting

URL: View paper

Brief Assessment

Temporal Latent Autoencoder[70] focuses on multivariate time series forecasting using encoder-decoder architectures with temporal regularization in latent space, not on extending Stochastic Interpolants to jointly optimized latent variable models with continuous-time dynamics for generative modeling.

---

3. Stochastic Latent Talking Face Generation Toward Emotional Expressions and Head Poses

URL: View paper

Brief Assessment

Stochastic Talking Face[69] focuses on modeling time-varying facial motion distributions for emotional expressions using deep state space models and variational autoencoders, not on general continuous-time latent generative frameworks with stochastic interpolants.

---

4. A conditional latent autoregressive recurrent model for generation and forecasting of beam dynamics in particle accelerators

URL: View paper

Brief Assessment

Beam Dynamics Autoregressive[71] focuses on charged particle beam dynamics in accelerators using CV AE and LSTM for spatiotemporal forecasting, not on joint encoder-decoder-generative model training with continuous-time latent dynamics as in LSI.

---

5. Conditional Image-to-Video Generation with Latent Flow Diffusion Models

URL: View paper

Brief Assessment

Latent Flow Diffusion[66] focuses on conditional image-to-video generation using latent flow sequences for warping, not on joint end-to-end learning of encoder-decoder-generative models with continuous-time dynamics in latent space as LSI does.

---

6. Simultaneous Modeling of Protein Conformation and Dynamics via Autoregression

URL: View paper

Brief Assessment

Protein Conformation Autoregression[65] focuses on protein conformational dynamics using autoregressive modeling with SE(3) diffusion, not on joint encoder-decoder-generative model training in continuous-time latent spaces for general data distributions.

---

7. Learning the intrinsic dynamics of spatio-temporal processes through Latent Dynamics Networks

URL: View paper

Brief Assessment

Latent Dynamics Networks[68] focuses on learning spatio-temporal dynamics of PDEs through continuous-time latent variables with a reconstruction network, but does not address the specific challenge of extending Stochastic Interpolants to jointly optimized latent variable models or the ELBO derivation in continuous time that LSI introduces.

---

8. A personalized time-resolved 3D mesh generative model for unveiling normal heart dynamics

URL: View paper

Brief Assessment

Heart Dynamics Mesh[64] focuses on cardiac mesh generation using geometric encoders and temporal transformers for medical imaging, not on joint training of encoder-decoder-generative models with continuous-time latent dynamics as in LSI.

---

9. Latent Conditional Diffusion-based Data Augmentation for Continuous-Time Dynamic Graph Model

URL: View paper

Brief Assessment

Dynamic Graph Conditional Diffusion[7] focuses on data augmentation for continuous-time dynamic graphs using a conditional diffusion model in latent space, not on joint encoder-decoder-generative model training with continuous-time latent dynamics as in LSI.

---

10. Generative learning for nonlinear dynamics

URL: View paper

Brief Assessment

Generative Nonlinear Dynamics[1] focuses on learning representations and generators for chaotic dynamical systems from time series data, not on joint encoder-decoder-generative model training in latent spaces with continuous-time stochastic interpolants.

Contribution 2: Principled continuous-time ELBO objective

Description: The authors derive a novel Evidence Lower Bound (ELBO) training objective formulated directly in continuous time. This objective enables simulation-free scalable training while preserving the flexibility of arbitrary prior distributions and providing data log-likelihood control, addressing the computational challenges of applying SI in high-dimensional spaces.

This contribution was assessed against **4 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Denoising Diffusion Variational Inference: Diffusion Models as Expressive Variational Posteriors

URL: View paper

Prior Art Analysis

Diffusion Variational Inference[62] demonstrates prior work on continuous-time ELBO objectives for diffusion models that enable simulation-free training with flexible priors. The candidate paper derives a continuous-time ELBO for variational inference using diffusion models in latent space, presenting a 'Markovian ELBO' and 'regularized ELBO' formulations. Both papers derive ELBO objectives in continuous time for diffusion-based models, use simulation-free training approaches, and support flexible prior distributions. The candidate's work on continuous-time ELBO formulations for diffusion models predates or parallels the original paper's claims of novelty in this area.

Evidence

Evidence 1 - **Rationale**: Both papers claim to derive novel ELBO objectives in continuous time for diffusion-based models. The candidate explicitly states they train with 'a novel regularized evidence lower bound (elbo)' for diffusion models, directly challenging the original's claim of being first to derive such an objective. - **Original**: we achieve this by developing a principled evidence lower bound (elbo) objective derived directly in continuous time. the joint optimization allows lsi to learn effective latent representations along with a generative process that transforms an arbitrary prior distribution into the encoder-defined a... - **Candidate**: we propose denoising diffusion variational inference (ddvi), a black-box variational inference algorithm for latent variable models which relies on diffusion models as flexible approximate posteriors. specifically, our method introduces an expressive class of diffusion-based variational posteriors t...

Evidence 2 - **Rationale**: The candidate derives a continuous-time ELBO formulation (the 'Markovian ELBO') for diffusion models. This demonstrates prior work on deriving ELBO objectives in continuous time for diffusion-based variational inference. - **Original**: our key innovation lies in deriving a principled, flexible and scalable training objective as an evidence lower bound (elbo) directly in continuous time. this objective, like si, provides data log-likelihood control, while enabling scalable end-to-end training - **Candidate**: the standard approach to fit auxiliary-variable generative models (maaløe et al. 2016) is to apply the elbo twice: log pθ(x) ≥ log pθ(x) - kl(qϕ(z|x)||pθ(z|x)) (3) ≥ log pθ(x) - kl(qϕ(z|x)||pθ(z|x)) (4) - eqϕ(z|x)[kl(qϕ(y|x, z)||r(y|x, z))] = eqϕ(y,z|x)[log pθ(x|z)] (5) - kl(qϕ(y, z|x)||r(y|x, z)p(z...

Evidence 3 - **Rationale**: The candidate demonstrates simulation-free training by leveraging the Markov structure of diffusion processes, similar to the original's claim of simulation-free scalable training. - **Original**: our formulation admits simulation-free training analogous to observation-space diffusion and si models, while preserving the flexibility of si framework. - **Candidate**: optimizing equation (5) requires tractably dealing with the prior regularization term lreg(x, θ, ϕ) := -kl(qϕ(y, z|x)||r(y|x, z)p(z)), which we equivalently rewrite as: lreg = eqϕ(y,z|x)[log(r(y|x, z)p(z))] + h(q). (6) we can expand the first term by leveraging the markov structure of r, qto rewrite...

---

2. Simulator-free stochastic variational inference for neural SDEs

URL: View paper

Prior Art Analysis

Simulator-free Neural SDEs[63] demonstrates that a continuous-time ELBO objective for neural SDEs with simulation-free training was developed prior to the original paper. The candidate explicitly presents a 'reparametrized ELBO' that enables 'simulator-free svi for nsdes' and states this approach allows training 'without a simulator/sde solver in the training loop.' The candidate's abstract and theoretical sections establish that their method optimizes 'a lower bound on the evidence' in continuous time while avoiding simulation, directly addressing the same computational challenges the original paper claims to solve.

Evidence

Evidence 1 - **Rationale**: Both papers claim to develop methods for optimizing an evidence lower bound in continuous time without requiring simulation during training, addressing the same core computational challenge. - **Original**: we achieve this by developing a principled evidence lower bound (elbo) objective derived directly in continuous time. the joint optimization allows lsi to learn effective latent representations along with a generative process that transforms an arbitrary prior distribution into the encoder-defined a... - **Candidate**: this thesis focuses on the development of methods for training neural stochastic differential equations (nsdes) by optimizing a lower bound on the evidence without a simulator/sde solver in the training loop.

Evidence 2 - **Rationale**: Both papers describe using a continuous-time ELBO framework with flexible priors to address computational challenges in high-dimensional spaces, though applied to different problem domains. - **Original**: our approach allows transforming arbitrary prior distributions into the encoder-defined aggregated posterior, simultaneously aligning data representations with a high-fidelity generative process using that representation. lsi's single elbo objective provides a unified, scalable framework that avoids... - **Candidate**: the abstract formalism of stochastic variational inference (svi) allows us to evaluate the performance of our approach across a number of problems in wide-reaching domains including state estimation, governing equations discovery , and, of course, time-series forecasting.

Evidence 3 - **Rationale**: Both papers emphasize simulation-free training as a key advantage, with the candidate demonstrating significant computational improvements over baseline methods. - **Original**: our formulation admits simulation-free training analogous to observation-space diffusion and si models, while preserving the flexibility of si framework. - **Candidate**: in addition to the asymptotic time cost improvements offered by the methods we develop here, we show that our approach can reduce the number of model evaluations required in training nsdes by more than an order of magnitude in practice compared even to neural ordinary differential equations (nodes) wh...

Evidence 4 - **Rationale**: The candidate's table of contents explicitly lists sections on 'simulator-free svi' and 'reparametrized elbo' for neural SDEs, indicating prior work on the same principled continuous-time ELBO objective claimed by the original paper. - **Original**: 3)principled elbo objective:a new elbo as a principled training objective that retains strengths of si - simple simulation free training and flexible prior choice - while enabling the benefits of joint training in a latent space. - **Candidate**: theorem 1 for simulator-free svi for nsdes ... the reparametrized elbo ... amortized svi for latent nsdes

---

3. ITF-VAE: Variational Auto-Encoder Using Interpretable Continuous Time Series Features

URL: View paper

Brief Assessment

ITF-VAE[61] focuses on time series generation using interpretable continuous function features with a standard VAE ELBO, not on continuous-time stochastic processes or simulation-free training for flexible priors in latent variable models.

---

4. Amortized Control of Continuous State Space Feynman-Kac Model for Irregular Time Series

URL: View paper

Brief Assessment

Feynman-Kac Irregular Series[60] focuses on irregular time series with multi-marginal Doob's h-transform for discrete observations, while the original paper addresses latent variable models for generative modeling with continuous-time dynamics and arbitrary priors.

Contribution 3: Unifying perspective on continuous-time latent variable models

Description: The authors provide a theoretical perspective that unifies Stochastic Interpolants with latent variable models through continuous-time stochastic processes. This perspective connects diffusion bridges, variational posteriors, and stochastic interpolants to enable joint optimization in latent spaces.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. ARTEMIS integrates autoencoders and SchrÃ¶dinger Bridges to predict continuous dynamics of gene expression, cell population, and perturbation from time-series single-cell data

URL: View paper

Brief Assessment

ARTEMIS[11] focuses on biological applications (cellular trajectories, gene expression dynamics) using VAE with Schrödinger bridges for single-cell data, not on providing a theoretical unifying perspective for continuous-time latent variable models in generative modeling.

---

2. Denoising diffusion bridge models

URL: View paper

Brief Assessment

Denoising Diffusion Bridges[52] focuses on diffusion bridge processes for image-to-image translation in observation space, not on unifying continuous-time stochastic processes with latent variable models for joint optimization in latent spaces.

---

3. Consistency diffusion bridge models

URL: View paper

Brief Assessment

Consistency Diffusion Bridges[54] focuses on diffusion bridge models for paired data translation tasks, not on unifying continuous-time latent variable models with stochastic interpolants in latent spaces as the original paper does.

4. Latent schr {\" o} dinger bridge diffusion model for generative learning

URL: View paper

Brief Assessment

Latent Schrodinger Bridge[59] focuses on pre-training encoder-decoder structures and establishing end-to-end convergence rates for latent diffusion models, rather than providing a unifying theoretical perspective connecting diffusion bridges, variational posteriors, and stochastic interpolants for joint optimization.

5. Diffusion bridges vector quantized variational autoencoders

URL: View paper

Brief Assessment

Diffusion Bridges VQ-VAE[55] focuses on discrete latent representations using diffusion bridges between continuous coded vectors and priors for VQ-VAE models. The original paper's unifying perspective connects diffusion bridges, variational posteriors, and stochastic interpolants for joint optimization in continuous latent spaces, which is a different theoretical framework and application domain.

6. Diffusion bridge implicit models

URL: View paper

Brief Assessment

Diffusion Bridge Implicit[56] focuses on accelerating sampling for diffusion bridge models through non-Markovian processes on discretized timesteps, not on unifying continuous-time latent variable models with stochastic interpolants as the original paper does.

7. Controllable fashion rendering via brownian bridge diffusion model with latent sketch encoding

URL: View paper

Brief Assessment

Fashion Brownian Bridge[58] focuses on fashion image rendering using Brownian bridge diffusion models with sketch encoding. This is a domain-specific application paper, not a theoretical framework for unifying continuous-time latent variable models with stochastic interpolants.

8. Modeling temporal data as continuous functions with process diffusion

URL: View paper

Brief Assessment

Process Diffusion Continuous[57] focuses on modeling temporal data as continuous functions using stochastic processes for noise, not on unifying diffusion bridges with latent variable models through continuous-time stochastic processes as the original paper does.

9. Learning diffusion bridges on constrained domains

URL: View paper

Brief Assessment

Diffusion Bridges Constrained[53] focuses on learning diffusion processes on constrained domains (discrete sets, bounded spaces) using Doob's h-transform, not on unifying latent variable models with stochastic interpolants in continuous latent spaces.

10. Stochastic Interpolants: A Unifying Framework for Flows and Diffusions

URL: View paper

Brief Assessment

Stochastic Interpolants Framework[51] focuses on bridging arbitrary probability distributions in observation space using continuous-time stochastic processes, not on latent variable models with joint optimization of encoder-decoder architectures. The framework does not address the specific challenge of learning in unobserved latent spaces with end-to-end training that the original paper tackles.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Latent Stochastic Interpolants View paper
- [1] Generative learning for nonlinear dynamics View paper
- [2] Generative modeling of neural dynamics via latent stochastic differential equations View paper
- [3] Latent diffusion-based data augmentation for continuous-time dynamic graph model View paper
- [4] Devil is in the Details: Density Guidance for Detail-Aware Generation with Flow Models View paper
- [5] TG-GAN: Continuous-time temporal graph generation with deep generative models View paper
- [6] Path-minimizing latent ODEs for improved extrapolation and inference View paper
- [7] Latent Conditional Diffusion-based Data Augmentation for Continuous-Time Dynamic Graph Model View paper
- [8] Continuous latent process flows View paper
- [9] Extraction and recovery of spatio-temporal structure in latent dynamics alignment with diffusion models View paper
- [10] Generating Unseen Nonlinear Evolution in Sea Surface Temperature Using a Deep Learning-Based Latent Space Data Assimilation Framework View paper
- [11] ARTEMIS integrates autoencoders and SchrÃ¶dinger Bridges to predict continuous dynamics of gene expression, cell population, and perturbation from time-series single-cell data View paper
- [12] Exploring Variational Autoencoders and Generative Latent Time-Series Models for Synthetic Data Generation and Forecasting View paper
- [13] Modeling GAN Latent Dynamics using Neural ODEs View paper
- [14] Generative Plant Growth Simulation from Sequence-Informed Environmental Conditions View paper
- [15] Dynamic latent space relational event model View paper
- [16] T-CACE: A Time-Conditioned Autoregressive Contrast Enhancement Multi-Task Framework for Contrast-Free Liver MRI Synthesis, Segmentation, and Diagnosis View paper
- [17] Deep generative model with hierarchical latent factors for time series anomaly detection View paper

- [18] GOKU-UI: Ubiquitous Inference through Attention and Multiple Shooting for Continuous-time Generative Models View paper
- [19] Spate-gan: Improved generative modeling of dynamic spatio-temporal patterns with an autoregressive embedding loss View paper
- [20] Exploit Your Latents: Coarse-Grained Protein Backmapping with Latent Diffusion Models View paper
- [21] A survey on diffusion models for time series and spatio-temporal data View paper
- [22] Monge-AmpÃ¨re Flow for Generative Modeling View paper
- [23] Multivariate stochastic modeling for transcriptional dynamics with cell-specific latent time using SDEvelo View paper
- [24] Optimization of geological carbon storage operations with multimodal latent dynamic model and deep reinforcement learning View paper
- [25] Evolve Smoothly, Fit Consistently: Learning Smooth Latent Dynamics For Advection-Dominated Systems View paper
- [26] Anode-gan: incomplete time series imputation by augmented neural ode-based generative adversarial networks View paper
- [27] AudioLCM: Text-to-Audio Generation with Latent Consistency Models View paper
- [28] ODE-based generative modeling: Learning from a single natural image View paper
- [29] Generative time series models with interpretable latent processes for complex disease trajectories View paper
- [30] Manifold Interpolating Optimal-Transport Flows for Trajectory Inference View paper
- [31] Time-aware neural ordinary differential equations for incomplete time series modeling View paper
- [32] Modeling temporal data as continuous functions with stochastic process diffusion View paper
- [33] Continuous-time generative graph neural network for attributed dynamic graphs: student research abstract View paper
- [34] Latent space energy-based neural odes View paper
- [35] The principles of diffusion models View paper
- [36] Molecular latent space simulators View paper
- [37] Latent sdes on homogeneous spaces View paper
- [38] Deep latent state space models for time-series generation View paper
- [39] STODE: A Deep Generative Framework for Continuous Spatiotemporal Dynamics in Spatial Transcriptomics View paper
- [40] Learning a generative motion model from image sequences based on a latent motion matrix View paper
- [41] Generative Pre-trained Autoregressive Diffusion Transformer View paper
- [42] Enabling dynamic 3D coherent diffraction imaging via adaptive latent space tuning of generative autoencoders View paper
- [43] Unified Continuous Generative Models View paper
- [44] From latent dynamics to meaningful representations View paper
- [45] Gendef: Learning generative deformation field for video generation View paper
- [46] Contraction and entropy production in continuous-time Sinkhorn dynamics View paper
- [47] Towards Scene Graph Anticipation View paper
- [48] Generative Latent Space Dynamics of Electron Density View paper
- [49] Controllable and Compositional Generation with Latent-Space Energy-Based Models View paper
- [50] AudioLCM: Efficient and High-Quality Text-to-Audio Generation with Minimal Inference Steps View paper
- [51] Stochastic Interpolants: A Unifying Framework for Flows and Diffusions View paper
- [52] Denoising diffusion bridge models View paper
- [53] Learning diffusion bridges on constrained domains View paper
- [54] Consistency diffusion bridge models View paper
- [55] Diffusion bridges vector quantized variational autoencoders View paper
- [56] Diffusion bridge implicit models View paper
- [57] Modeling temporal data as continuous functions with process diffusion View paper
- [58] Controllable fashion rendering via brownian bridge diffusion model with latent sketch encoding View paper
- [59] Latent schr {\" o} dinger bridge diffusion model for generative learning View paper
- [60] Amortized Control of Continuous State Space Feynman-Kac Model for Irregular Time Series View paper
- [61] ITF-VAE: Variational Auto-Encoder Using Interpretable Continuous Time Series Features View paper
- [62] Denoising Diffusion Variational Inference: Diffusion Models as Expressive Variational Posteriors View paper
- [63] Simulator-free stochastic variational inference for neural SDEs View paper
- [64] A personalized time-resolved 3D mesh generative model for unveiling normal heart dynamics View paper
- [65] Simultaneous Modeling of Protein Conformation and Dynamics via Autoregression View paper
- [66] Conditional Image-to-Video Generation with Latent Flow Diffusion Models View paper
- [67] Social lode: human trajectory prediction with latent odes View paper
- [68] Learning the intrinsic dynamics of spatio-temporal processes through Latent Dynamics Networks View paper
- [69] Stochastic Latent Talking Face Generation Toward Emotional Expressions and Head Poses View paper
- [70] Temporal latent auto-encoder: A method for probabilistic multivariate time series forecasting View paper
- [71] A conditional latent autoregressive recurrent model for generation and forecasting of beam dynamics in particle accelerators View paper