Preprints.org

Review

Not peer-reviewed version

A Survey on Data Generation for Time Series: Taxonomy, Review and Prospects

Xu Zhang , Chang Xu * , Hao Li , Yuhao Huang , Qiushui Xu , Yuxuan Liang , Chenghao Liu , Ming Jin , Qingsong Wen , Peng Wang , Wei Wang , Jiang Bian

Posted Date: 11 May 2026

doi: 10.20944/preprints202605.0725.v1

Keywords: time series generation; synthetic data; data augmentation; temporal data mining; deep learning

Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026

doi:10.20944/preprints202605.0725.v1

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Review

A Survey on Data Generation for Time Series: Taxonomy, Review and Prospects

Xu Zhang [1], Chang Xu [2,*], Hao Li [3], Yuhao Huang [4], Qiushui Xu [5], Yuxuan Liang [6,†], Chenghao Liu [7], Ming Jin [8], Qingsong Wen [9], Peng Wang [1,†], Wei Wang [1] and Jiang Bian [2]

[1] College of Computer Science and Artificial Intelligence, Fudan University, Shanghai, China
[2] Microsoft Research, Beijing, China
[3] Department of Bioengineering, Imperial College London, London, United Kingdom
[4] School of Computer Science, Nanjing University, Nanjing, China
[5] Department of Industrial and Manufacturing Engineering, Pennsylvania State University, State College, Pennsylvania State, United States
[6] The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China
[7] Data Dog AI Research, New York, United States. E-mail: twinsken@gmail.com.
[8] School of ICT, Griffith University, Queensland, Australia
[9] Squirrel Ai Learning, Seattle, United States. E-mail: qingsongedu@gmail.com.
[*] Correspondence: chanx@microsoft.com
[†] Member, IEEE.

Abstract

Time series generation (TSG) plays a fundamental role in data engineering and knowledge discovery, serving as a key enabler for data augmentation, representation learning, privacy preservation, and scenario simulation analysis in temporal data mining. By synthesizing realistic and controllable time series, TSG directly benefits downstream tasks such as forecasting, anomaly detection, and classification by data augmentation and improving model generalization. This survey presents a comprehensive and systematic review of TSG methodologies, spanning traditional non-deep learning approaches, such as rule-based, statistical, and simulation-based methods, and modern deep generative models, including variational autoencoders, generative adversarial networks, diffusion models, normalizing flows, and large language models. To organize the rapidly growing literature, we introduce a unified multi-level taxonomy to organize the technical landscape of TSG and clarify the relationships among diverse approaches. The taxonomy emphasizes underlying modeling principles rather than individual representative models. Specifically, it categorizes TSG methods according to modeling backbones, generation settings (conditional versus unconditional), sampling regularity, and modality characteristics. Building on this taxonomy, we systematically review advanced methods within each category and discuss how they address key challenges such as irregular sampling, long-sequence modeling, multivariate dependencies, and controllable generation, and how they are applied in real-world scenarios. Furthermore, we present a holistic evaluation framework encompassing five core dimensions, including fidelity, diversity, controllability, downstream performance, and privacy protection. Finally, we identify critical open challenges and outline promising research directions from data characteristics, application scenarios, and modeling paradigm perspectives. This survey aims to serve as a structured reference and roadmap for researchers and practitioners in this rapidly evolving field.

Keywords: time series generation; synthetic data; data augmentation; temporal data mining; deep learning

---

1. Introduction

Time series generation (TSG) concerns the synthesis of realistic and diverse time series data that preserve essential temporal dynamics. As a fundamental component of modern data management and knowledge discovery systems, TSG serves as a core data engineering primitive for constructing,

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

enriching, and maintaining temporal datasets. Time series data are ubiquitous in real-world applications, ranging from industrial monitoring [1], healthcare [2], finance [3,4], and web services [5,6]. In the data mining communities, TSG is increasingly recognized as a critical enabler for downstream analytics, particularly in three key scenarios: (i) data augmentation [7], where synthetic sequences alleviate data scarcity and class imbalance; (ii) conditional generation [8], which enables controllable synthesis for scenario analysis and what-if simulations; and (iii) privacy-preserving data sharing [9], where generated time series reduce the risk of sensitive information leakage while retaining utility. Reflecting its growing importance, academic interest in this field has steadily increased, as shown in Figure 1.

Although both time series generation and forecasting are concerned with producing temporal sequences, they are fundamentally different. Forecasting aims to estimate future values conditioned on historical observations. In contrast, TSG seeks to model the full joint distribution of temporal processes. **By learning the underlying generative distribution of observed data, TSG enables the synthesis of realistic and diverse sequences that reflect complex temporal dynamics, including uncertainty, variability, and rare events**. This distinction is crucial: while forecasting is often concerned with producing one or a limited number of future trajectories, TSG provides a principled way to explore the space of plausible temporal behaviors, thereby supporting applications such as data augmentation, stress testing, and scenario simulation.

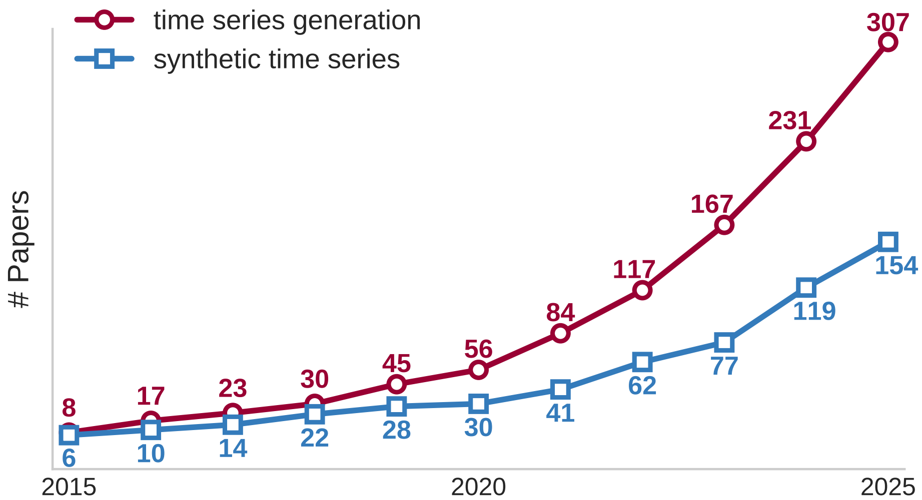

Figure 1. Trends in the cumulative number of studies on TSG in recent years. The numbers are retrieved from DBLP respectively using "time series generation" and "synthetic time series" as keywords.

Over the past few decades, TSG methods have evolved along several methodological paradigms, including rule-based approaches [10–19], simulation–based methods [20,21], and deep learning–based techniques. Among these paradigms, deep learning approaches have become increasingly dominant, while rule-based, statistical, and simulation-based methods are comparatively less prevalent, as illustrated in Figure 2. These TSG methods have been widely applied to tasks such as time series data augmentation, representation learning, privacy-preserving, and scenario simulation analysis. Rule-based methods [22] synthesize new time series data by recombining predefined patterns [13]. Statistical approaches typically rely on stochastic processes, such as Markov models, to generate time series for scenario simulation. Software simulation based methods generate time series data by constructing virtual environments and specifying interaction rules [20]. In recent years, deep generative models have achieved breakthrough progress in fields such as image, speech, and text processing, and have gradually extended to time series modeling [23–30]. Unlike traditional statistical methods that rely on explicit distributional assumptions, deep generative models can directly learn complex temporal dynamics and underlying generative mechanisms through high-dimensional nonlinear network architectures.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

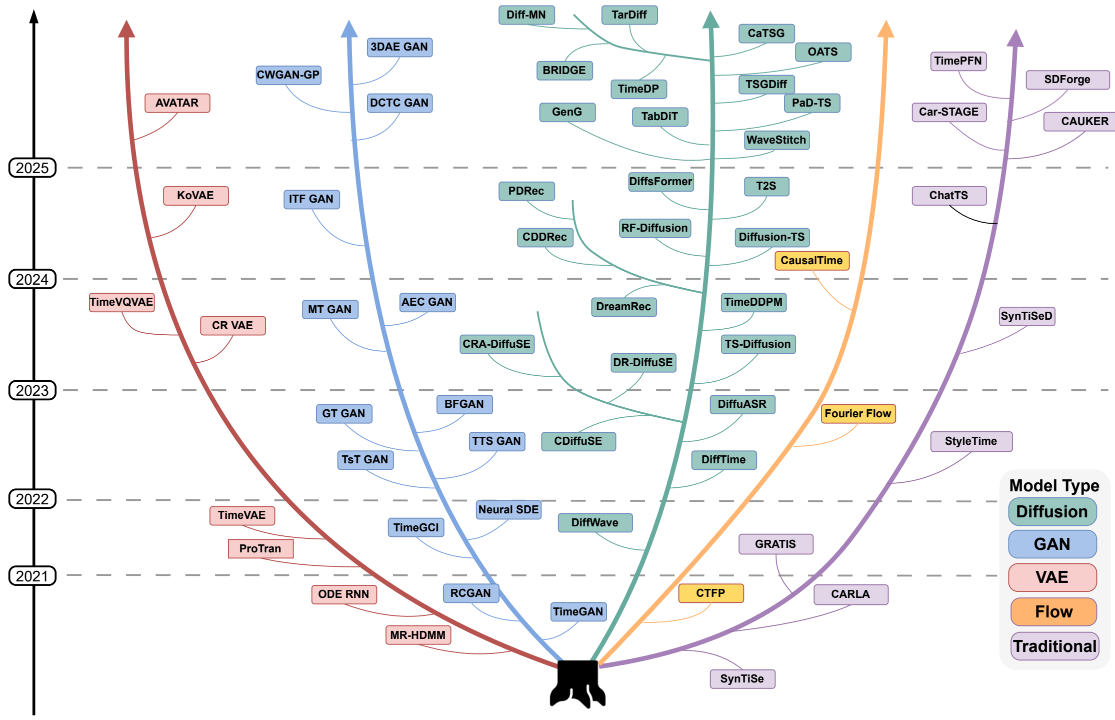

Figure 2. The evolutionary tree of time series generative models traces the development of representative methods until 2025 and highlights influential models across the field. Models on the same branch exhibit closer methodological relationships. Different branches correspond to major model families and are indicated by distinct colors. The vertical position of each model on the timeline denotes its year of introduction, illustrating the temporal progression and diversification of generative paradigms for time series data.

As summarized in Table 1, existing surveys on TSG are largely organized around specific model families (e.g., GANs or diffusion models), particular application domains, or isolated evaluation aspects. In contrast, our survey offers a comprehensive, generation-centric synthesis of the TSG field, systematically summarizing its **distinctive properties** from a generative perspective, such as irregular sampling, long-sequence generation, and condition controllability, rather than being limited to model architecture alone. Moreover, we cover a broad spectrum of methodologies and provide structured taxonomies of domain-specific applications and evaluation metrics, together with a systematic discussion of open challenges and future research directions.

Table 1. Comparison between existing surveys and this work.

Survey	Generation Centric	Method Coverage	Distinct TSG Properties Taxonomy	Domain Application Taxonomy	Evaluation Metric Taxonomy
Sommers et al. [31]	✗	GAN and Diffusion	✗	✗	✗
Yang et al. [32]	✗	Diffusion	✗	✓	✗
Hu et al. [33]	✓	Traditional methods, VAE, and GAN	✗	✗	✗
Brophy et al. [34]	✗	GAN	✗	Partial	Downstream performance, Fidelity, and privacy
Ang et al. [35]	✓	GAN, VAE and Flow	✗	✗	Fidelity
Ours	✓	Traditional methods, GAN, VAE, Flow, Diffusion, LLM	✓	✓	Downstream performance, Fidelity, Privacy, Diversity, Controllability

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

Based on a systematic review of more than 150 representative papers selected from leading venues in artificial intelligence and data science, we weave together the fragmented strands of rule-based, statistical, simulation-based, and deep generative research. Beyond cataloging existing methods, our synthesis highlights emerging trends toward more general, multimodal, and controllable generation paradigms. The remainder of this survey is structured as follows. We first establish the theoretical foundations and a hierarchical taxonomy of TSG methods in Sections 2 and 3. We then examine practical applications across key industries in Section 4, followed by a systematic review of evaluation metrics in Section 5. Finally, Section 6 explores open challenges and future research directions

Overall, this survey provides a unified TSG perspective that integrates traditional methods with modern deep learning under a consistent taxonomy, serving as a structured reference and roadmap for future research, applications, and evaluation of multimodal, controllable, and diverse TSG systems.

2. Background

This section provides the theoretical background necessary to contextualize the diverse approaches covered in this survey. We first formalize the TSG problem and then review the core generative paradigms. Specifically, we focus on the mathematical formulation of TSG and the adaptation of representative frameworks, including Variational Autoencoder, Generative Adversarial Network, Diffusion Model, and Normalizing Flow to the temporal domain. Figure 3 illustrates these paradigms.

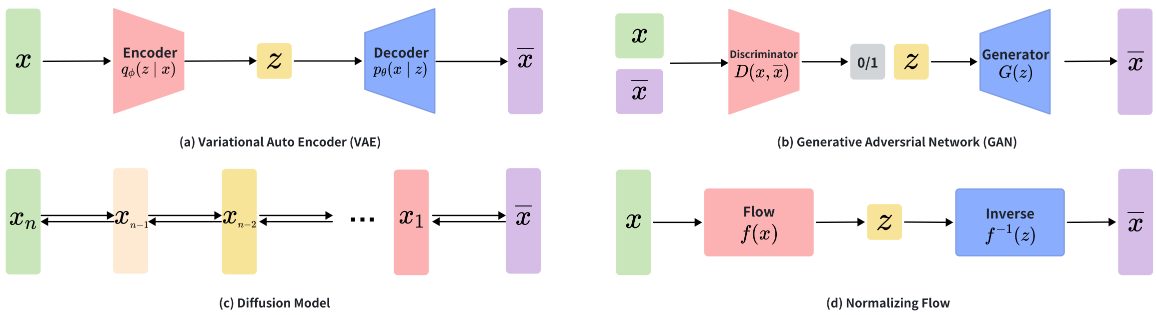

Figure 3. Overview of Representative Deep Generative Modeling Paradigms. (a) Variational Autoencoder (VAE): VAEs encode data into a latent distribution and generate samples by decoding latent variables drawn from a predefined prior. (b) Generative Adversarial Network (GAN): GANs learn data distributions implicitly through adversarial training between a generator and a discriminator. (c) Diffusion Models: Diffusion models transform data into noise via a forward process and generate samples by reversing this process through iterative denoising.

2.1. Time Series Generation

A TS is a sequence of observations $\mathbf{x}_{1:T} = (x_1, \ldots, x_T)$, where $x_t \in \mathbb{R}^d$ is indexed by time t. The goal of TSG is to learn a generative model $p_\theta(\mathbf{x}_{1:T})$ that approximates the true underlying data distribution $p_{\text{data}}(\mathbf{x}_{1:T})$. This enables sampling synthetic trajectories $\tilde{\mathbf{x}}_{1:T} \sim p_\theta(\mathbf{x}_{1:T})$ that are statistically indistinguishable from real data, as illustrated in Figure 3.

Compare with TS forecasting, which typically focuses on minimizing the prediction error of future values given a fixed history (i.e., estimating the conditional expectation $\mathbb{E}[\mathbf{x}_{t+1:T} \mid \mathbf{x}_{1:t}]$), TSG targets the *complete joint distribution* rather than deterministic extrapolation. This distinction brings two unique characteristics to TSG: **(1) Global Distribution Modeling:** Unlike forecasting, which prioritizes local trend accuracy, TSG must capture global semantic coherence and long-term dependencies (e.g., seasonality, non-stationarity) to ensure the generated samples possess realistic statistical properties across the entire horizon. **(2) Controllability and Diversity:** TSG extends beyond fitting historical patterns by enabling *conditional generation*, which allows control over specific attributes (e.g., class labels or textual descriptions) and supports the generation of diverse counterfactual scenarios rather than a single deterministic future.

Next, we categorize TSG tasks across three key dimensions: conditioning, domain, and modality.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

2.1.1. Unconditional and Conditional TSG

Unconditional models approximate the marginal distribution $p_\theta(\mathbf{x}_{1:T})$, synthesizing sequences based solely on intrinsic dynamics. In contrast, conditional generation extends this to: $p_\theta(\mathbf{x}_{1:T} \mid \mathbf{c})$, where $\mathbf{c}$ represents auxiliary controls (e.g., class labels, text descriptions). This formulation enables targeted simulation, which is essential for applications requiring semantically grounded generation.

2.1.2. Single- and Multi-Domain TSG

A domain defines a distributional context with consistent temporal patterns. In single-domain settings, training and generation occur within a homogeneous distribution. Multi-domain TSG involves heterogeneous datasets $\{\mathcal{D}_i\}$ where dynamics vary across sources. This setting requires models to disentangle shared temporal structures from domain-specific variations, enabling generalization across distinct data sources.

2.1.3. Single- and Multi-Modality TSG

Modality M refers to distinct signal types (e.g., ECG, text). Single-modality TSG focuses on synthesizing one signal source. Multi-modality TSG models the joint distribution across aligned signals:

$$p_\theta\left(x_{1:T}^{(1)}, \ldots, x_{1:T}^{(M)}\right), \tag{1}$$

or generates one modality conditioned on another, capturing inter-modality correlations and synchronization essential for complex systems like physiological monitoring.

2.2. Classic Generative Models

As illustrated in Figure 3, we review representative paradigms adapted for TSG, ranging from explicit density estimators to implicit adversarial approaches.

2.2.1. Variational Autoencoders (VAEs)

It maps sequences to a probabilistic latent space to capture temporal dynamics[36]. Given a time series $\mathbf{x}_{1:T}$, VAEs approximate the intractable posterior with an encoder $q_\phi(\mathbf{z} \mid \mathbf{x}_{1:T})$ and generate synthetic trajectories via a decoder $p_\theta(\mathbf{x}_{1:T} \mid \mathbf{z})$. The model is optimized by maximizing the Evidence Lower Bound (ELBO):

$$\mathcal{L}_{\text{ELBO}} = \mathbb{E}_{q_\phi}\left[\log p_\theta(\mathbf{x}_{1:T} \mid \mathbf{z})\right] - \text{KL}\big(q_\phi(\mathbf{z} \mid \mathbf{x}_{1:T}) \,\|\, p(\mathbf{z})\big). \tag{2}$$

In TSG, the latent variable $\mathbf{z}$ typically encodes global temporal characteristics, ensuring that sampled trajectories $\tilde{\mathbf{x}}_{1:T} \sim p_\theta(\cdot \mid \mathbf{z})$ maintain consistent statistical properties across time steps.

2.2.2. Generative Adversarial Networks (GANs)

It bypasses explicit likelihood estimation by establishing a two-player game between a generator G_θ and a discriminator D_ψ [37]. The generator maps a latent noise vector $\mathbf{z}$ to a synthetic sequence $\tilde{\mathbf{x}}_{1:T}$, while the discriminator attempts to distinguish real trajectories from generated ones. The training objective follows a minimax formulation:

$$\min_\theta \max_\psi \; \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}}[\log D_\psi(\mathbf{x}_{1:T})] + \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})}[\log(1 - D_\psi(G_\theta(\mathbf{z})))]. \tag{3}$$

For time series, D_ψ often uses recurrent or convolutional structures to penalize temporal discontinuities, encouraging G_θ to learn the joint distribution of sequence steps.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

2.2.3. Diffusion Models

It generates data via an iterative denoising process [38]. A forward process gradually corrupts clean data into Gaussian noise over diffusion timesteps, while a parameterized reverse process learns to remove the noise.

In TSG, let $\mathbf{x}_0 = \mathbf{x}_{0,1:T} \in \mathbb{R}^{T\times d}$ denote a clean trajectory of length T. The forward process perturbs the entire trajectory across diffusion steps $k \in \{1, \ldots, K\}$, producing a noisy sample $\mathbf{x}_k$. Here, the temporal dimension $1{:}T$ is preserved within $\mathbf{x}_k$. Training minimizes the noise prediction objective:

$$\mathcal{L}_{\text{simple}} = \mathbb{E}_{\mathbf{x}_0,k,\boldsymbol{\epsilon}}\left[\|\boldsymbol{\epsilon} - \epsilon_\theta(\mathbf{x}_k, k)\|^2\right], \tag{4}$$

where ϵ_θ, typically implemented with a U-Net or Transformer, learns to denoise the full trajectory at each diffusion step, enabling the modeling of long-range temporal dependencies.

2.2.4. Normalizing Flows

It constructs complex temporal distributions by transforming a simple base distribution (e.g., Gaussian) through a sequence of invertible and differentiable mappings f_θ [39]. A synthetic trajectory is generated as $\tilde{\mathbf{x}}_{1:T} = f_\theta(\mathbf{z})$. Unlike GANs, flows allow for exact log-likelihood computation via the change-of-variables formula:

$$\log p_\theta(\mathbf{x}_{1:T}) = \log p_0(\mathbf{z}) + \sum_{k=1}^{K} \log\left|\det \frac{\partial f_k}{\partial \mathbf{h}_{k-1}}\right|. \tag{5}$$

where $\mathbf{h}_{k-1}$ denotes the intermediate state input to the k-th layer. This invertibility makes flows particularly suitable for applications requiring both generation and precise density estimation of temporal patterns [40].

3. Taxonomy and Roadmap

3.1. Overview

To bring conceptual clarity to the diverse and rapidly evolving landscape of TSG, this chapter introduces a multi-level taxonomy aligned with the hierarchical framework shown in Figure 4. Rather than aiming for exhaustive coverage or enumerating individual methods, our taxonomy seeks to abstract the *personalized or shared modeling assumptions and control mechanisms* that underlie different TSG approaches. Specifically, we organize existing techniques by examining how they conceptualize the data generation process, what assumptions they impose on temporal dynamics and stochasticity, and how control or conditioning signals are incorporated into generation.

This taxonomy is structured to emphasize the *logical relationships* across hierarchical levels, ranging from high-level generative paradigms to key sub-categories and their defining technical properties. In doing so, it highlights how different methodological branches diverge and converge along critical dimensions such as generative assumptions, controllability, and conditioning flexibility.

By focusing on these fundamental design choices rather than individual implementations, we aim to provide a principled roadmap for understanding the technical space of TSG and its research directions. All related works are systematically summarized in Table 2.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

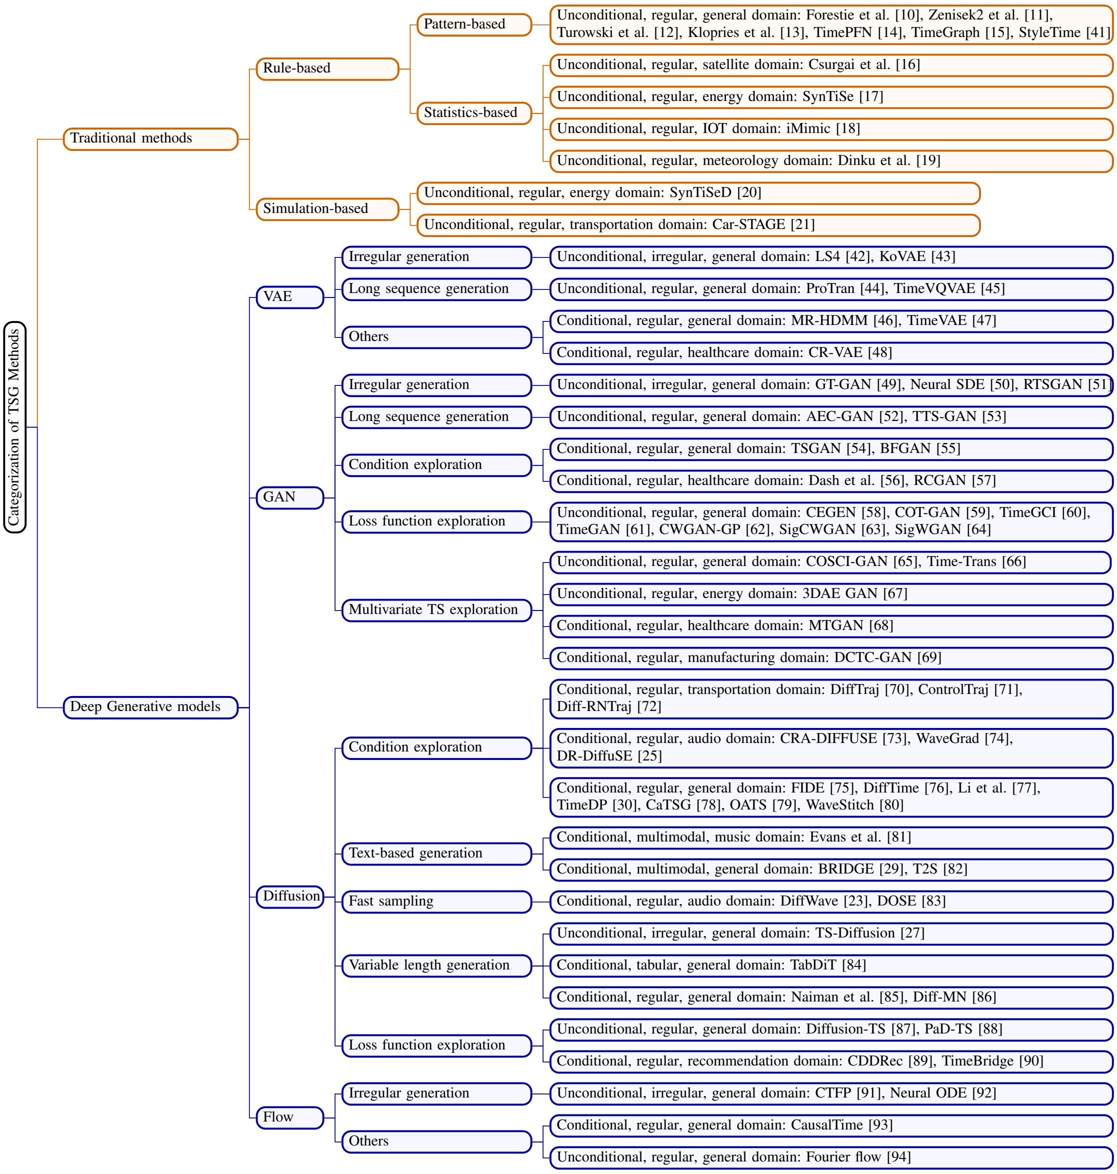

Figure 4. Research directions and core focuses of contemporary TSG methods.

3.2. Rule-Based Methods

3.2.1. Pattern-Based

To address data scarcity in TS applications, a line of *pure rule-based, pattern-driven* methods synthesizes data via explicit, human-designed rules rather than data-driven models. The synthetic data is used for pretraining TS models, aiming to augment training corpora with high-quality, controllable, and diverse synthetic data [22]. Existing synthesis methods used in this context can be broadly categorized into two classes.

The first class follows seasonal/trend/noise based generation, where synthetic series are constructed by composing interpretable components (e.g., trend, seasonality, noise, regime shifts) [41,42]. StyleTime [43] follows this paradigm by explicitly separating *trend* and *distributional style*, and transferring statistics such as autocorrelation, volatility, and power spectral density through analytical rules and gradient-based optimization, without deep learning. Forestier et al. [10] extend Dynamic Time Warping (DTW) barycenter averaging into a rule-based framework, where predefined weighting rules (Dirichlet averaging, nearest-neighbor selection, and distance-decayed weighting) control sequence fusion, alleviating overfitting in sparse datasets. Zenisek et al. [11] propose a configuration-driven generator that encodes explicit rules for stationarity (Gaussian processes), periodicity and

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

wear (sine/linear functions), historical replay, and noise injection, enabling controllable sensor-stream synthesis. Turowski et al. [12] formalize common smart-meter anomaly patterns into injection rules (e.g., zero windows or amplitude scaling), producing realistic anomalies that improve classical detectors. Klopries et al. [13] synthesize data via compositional rules over trend, seasonality, events, and noise, combined additively or multiplicatively for autoencoder pretraining. Similarly, Kegel et al. [44] generate scenario-specific time series by decomposing data into trend, seasonal, and residual components and modifying them according to predefined feature constraints. ChatTS [45] bridges time series and LLMs using attribute-level generation rules (e.g., trends, frequencies, noise, spikes) and rule-based instruction evolution.

The second class is based on kernel-driven synthesis, where time series are sampled from stochastic processes defined by kernel functions. Chronos [46] introduces KernelSynth, where kernel-composition rules (randomly combining basis kernels via addition or multiplication) generate synthetic temporal functions, and these functions are sampled from Gaussian processes to produce univariate time series. The synthetic data is used to pretrain the base model, enabling enhanced generalization for zero-shot forecasting. Further, TimePFN [14] introduces LMC-Synth, where kernel-composition rules generate latent temporal functions and linear coregionalization rules mix them into multivariate series, enabling strong zero- and few-shot generalization. CAUKER [47] presents a time series classification-focused synthetic framework, fusing Gaussian Process kernel composition with Structural Causal Models. It uses kernel-composition rules to form composite kernels and causal propagation rules to generate coherent, clustered series, enabling strong zero-shot generalization for the pretrained TS foundation model. Finally, TimeGraph [15] constructs rule-synthetic datasets for causal discovery by explicitly modeling dependencies, irregular sampling, missingness, and latent confounders, supporting systematic evaluation under controlled temporal complexity.

3.2.2. Statistics-Based TSG

Csurgai et al. [16] propose a Markov chain-based model for land mobile satellite links, combining a two-state fading process with a partitioned Fritchman Markov model to capture threshold-dependent fade durations, validated via Monte Carlo simulation. For wind power applications, Denaxas et al. [17] introduce SynTiSe, a multi-regime MCMC framework that preserves diurnal and seasonal patterns by fitting regime-specific models and matching key statistics such as PDF (Probability Density Function), ACF (Autocorrelation Function), and ramp characteristics. To address IoT data scarcity, Gkoulis et al. [18] develop iMimic, a modular statistics-driven library that combines trend, seasonality, and stochastic noise components, supporting both regular and irregular sampling with strong reproducibility guarantees. In climate modeling, Dinku et al. [19] generate historical rainfall time series by statistically fusing satellite estimates and rain gauge data, ensuring consistency in rainfall statistics across data-sparse regions. Focusing on **controllable generation**, GRATIS [48] synthesizes diverse time series using Gaussian mixture autoregressive models, with parameters tuned via genetic algorithms to meet user-specified targets.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

Table 2. Summary and main papers about TSG. Blue indicates diffusion models, red is Generative Adversarial Network (GAN), green is Variational Autoencoder (VAE), yellow is normalizing flows and grey is others.

Source	Data	Backbone	Task	Scenarios	Modality	Venue	Year
DiffWave [23]	Univariate	Diffusion	Conditional	Audio	Single	ICLR	2021
WaveGrad [49]	Univariate	Diffusion	Conditional	Audio	Single	ICLR	2021
SGMSE [24]	Univariate	Diffusion	Conditional	Audio	Single	Interspeech	2022
CDiffuSE [50]	Univariate	Diffusion	Conditional	General	Single	ICASSP	2022
DOSE [51]	Univariate	Diffusion	Conditional	Audio	Single	NeurIPS	2023
DR-DiffuSE [25]	Univariate	Diffusion	Unconditional	Audio	Single	AAAI	2023
CRA-DiffuSE [52]	Univariate	Diffusion	Conditional	Audio	Single	ICME	2023
DiffuASR [53]	Univariate	Diffusion	Conditional	Recommendation	Single	CIKM	2023
DreamRec [26]	Univariate	Diffusion	Conditional	Recommendation	Single	NeurIPS	2023
TimeDDPM [54]	Multivariate	Diffusion	Conditional	Industry	Single	IEEE Sensors Journal	2023
TS-Diffusion [27]	Multivariate	Diffusion	Conditional	General	Single	ArXiv	2023
DiffTraj [55]	Multivariate	Diffusion	Conditional	Transportation	Single	NeurIPS	2023
DiffCharge [56]	Multivariate	Diffusion	Conditional	Energy	Single	TSG	2023
CLDM [57]	Multivariate	Diffusion	Conditional	Energy	Single	TSTE	2023
Fu et al [58]	Multivariate	Diffusion	Conditional	Energy	Single	Energy and Buildings	2024
RF-Diffusion [28]	Multivariate	Diffusion	Conditional	Network	Single	MobiCom	2024
Klein et al. [59]	Multivariate	Diffusion	Conditional	Healthcare	Single	ArXiv	2024
FTS-Diffusion [60]	Multivariate	Diffusion	Conditional	General	Single	ICLR	2024
TIME WEAVER [8]	Multivariate	Diffusion	Conditional	General	Single	ICML	2024
FIDE [61]	Multivariate	Diffusion	Conditional	General	Single	NeurIPS	2024
Diffusion-TS [62]	Multivariate	Diffusion	Conditional	General	Single	ICLR	2024
Diff-RNTraj [63]	Multivariate	Diffusion	Conditional	Transportation	Single	TKDE	2024
ControlTraj [64]	Multivariate	Diffusion	Conditional	Transportation	Single	KDD	2024
PDRec. [65]	Univariate	Diffusion	Conditional	Recommendation	Single	AAAI	2024
TSDM [66]	Univariate	Diffusion	Unconditional	Industry	Single	MSSP	2024
DiffsFormer [67]	Univariate	Diffusion	Conditional	Finance	Single	ArXiv	2024
CDDRec [68]	Univariate	Diffusion	Conditional	Recommendation	Single	PAKDD	2024
Wang et al. [69]	Univariate	Diffusion	Conditional	Energy	Single	TSG	2024
TarDiff [70]	Multivariate	Diffusion	Conditional	Healthcare	Single	KDD	2025
Li et al. [71]	Multivariate	Diffusion	Conditional	General	Single	ICLR	2025
PaD-TS [72]	Multivariate	Diffusion	Conditional	General	Single	ArXiv	2025
WaveStitch [73]	Multivariate	Diffusion	Conditional	General	Single	ArXiv	2025
BRIDGE [29]	Univariate	Diffusion	Conditional	General	Multimodal	ICML	2025
TimeDP [30]	Univariate	Diffusion	Conditional	General	Single	AAAI	2025
T2S [74]	Univariate	Diffusion	Conditional	General	Multimodal	IJCAI	2025
TSGDiff [74]	Multivariate	Diffusion	Unconditional	General	Single	AAAI	2026
CaTSG [75]	Multivariate	Diffusion	Conditional	General	Single	ArXiv	2026
OATS [76]	Multivariate	Diffusion	Conditional	General	Single	ArXiv	2026
Diff-MN [77]	Multivariate	Diffusion	Unconditional	General	Single	ArXiv	2026
C-RNN-GAN [78]	Multivariate	GAN	Unconditional	Music	Single	ArXiv	2016
RCGAN [79]	Multivariate	GAN	Conditional	Healthcare	Single	ArXiv	2017
Zhang et al. [7]	Multivariate	GAN	Conditional	Energy	Single	SmartGridComm	2018
CGAN [80]	Multivariate	GAN	Conditional	Finance	Single	ArXiv	2019
Dash et al. [81]	Multivariate	GAN	Conditional	Healthcare	Single	AIME	2020
COT-GAN [82]	Multivariate	GAN	Conditional	General	Single	NeurIPS	2020
Lin et al. [83]	Multivariate	GAN	Conditional	Network	Single	IMC	2020
TSGAN [84]	Multivariate	GAN	Conditional	General	Single	ArXiv	2020
CR-GAN [85]	Multivariate	GAN	Conditional	Manufacturing	Single	TII	2020
Pinceti et al. [86]	Multivariate	GAN	Conditional	Energy	Single	PESGM	2021
RTSGAN [87]	Multivariate	GAN	Unconditional	General	Single	ICDM	2021
Neural SDE [88]	Multivariate	GAN	Unconditional	General	Single	ICML	2021
Leznik et al. [89]	Multivariate	GAN	Unconditional	Network	Single	ICPE	2021
Sun et al. [90]	Univariate	GAN	Unconditional	Industry	Single	TII	2022
TsT-GAN [91]	Multivariate	GAN	Unconditional	General	Single	ArXiv	2022
PSA-GAN [92]	Multivariate	GAN	Unconditional	General	Single	ICLR	2022
COSCI-GAN [93]	Multivariate	GAN	Unconditional	General	Single	NeurIPS	2022
GT-GAN [94]	Multivariate	GAN	Unconditional	General	Single	NeurIPS	2022
CEGEN [95]	Multivariate	GAN	Unconditional	General	Single	AAAI	2022
BFGAN [96]	Multivariate	GAN	Conditional	General	Single	TKDE	2022
MTGAN [97]	Multivariate	GAN	Conditional	Healthcare	Single	TKDE	2023
AEC-GAN [98]	Multivariate	GAN	Unconditional	General	Single	AAAI	2023
ITF-GAN [99]	Multivariate	GAN	Unconditional	General	Single	KBS	2024
CWGAN-GP [100]	Multivariate	GAN	Conditional	Transportation	Single	ESWA	2025
3DAE GAN [101]	Multivariate	GAN	Conditional	Energy	Single	TII	2025
DCTC-GAN [102]	Multivariate	GAN	Conditional	Manufacturing	Single	KBS	2025
MR-HDMM [103]	Multivariate	VAE	Conditional	General	Single	ICML	2018
ProTran [104]	Multivariate	VAE	Unconditional	General	Single	NeurIPS	2021
TimeVAE [105]	Multivariate	VAE	Conditional	General	Single	ArXiv	2021
LS4 [106]	Multivariate	VAE	Unconditional	General	Single	ICML	2023
TimeVQVAE [107]	Multivariate	VAE	Unconditional	General	Single	ICML	2023
CR-VAE [108]	Multivariate	VAE	Conditional	Healthcare	Single	AAAI	2023
KoVAE [109]	Multivariate	VAE	Conditional	General	Single	ICLR	2024
AVATAR [110]	Multivariate	VAE	Unconditional	General	Single	SDM	2025
NM-VQTSG [111]	Univariate	VAE	Unconditional	General	Single	ArXiv	2025
Neural ODE [112]	Multivariate	VAE	Unconditional	General	Single	NeurIPS	2018
CTFP [113]	Multivariate	Normalizing flow	Unconditional	General	Single	NeurIPS	2020
Fourier flow [114]	Multivariate	Normalizing flow	Unconditional	General	Single	ICLR	2021
CausalTime [115]	Multivariate	Normalizing flow	Unconditional	General	Single	ICLR	2023
Csurgai et al. [16]	Univariate	Markov	Unconditional	Satellite	Single	MWCS	2007
Dinku et al. [19]	Univariate	Statistics	Unconditional	Meteorology	Single	IJC	2016
SynTiSe [17]	Multivariate	Markov	Unconditional	Energy	Single	SysCon	2016
Forestie et al. [10]	Univariate	Rule	Unconditional	General	Single	ICDM	2017
Zenisek et al. [11]	Multivariate	Rule	Unconditional	Industry	Single	IFAC-PapersOnLine	2018
Turowski et al. [12]	Univariate	Rule	Unconditional	Energy	Single	ACM e-Energy	2022
Klopries et al. [13]	Univariate	Rule	Unconditional	General	Single	ETFA	2022
SynTiSeD [20]	Multivariate	Simulation	Unconditional	Energy	Single	MSCPES	2023
TimeGraph [15]	Univariate	Rule	Unconditional	General	Single	ArXiv	2025
iMimic [18]	Univariate	Statistics	Unconditional	General	Single	IOT	2025
Aloni et al. [116]	Univariate	Fourier transform	Unconditional	Environment	Single	EMS	2025
SBTS [117]	Multivariate	SDE	Unconditional	General	Single	ArXiv	2025
TimePFN [14]	Multivariate	Rule	Unconditional	General	Single	AAAI	2025
Car-STAGE [21]	Multivariate	Simulation	Unconditional	Transportation	Single	ArXiv	2025
CENTS [118]	Multivariate	General framework	Conditional	General	Single	ArXiv	2025

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

3.3. Simulation-Based TSG

It produces data by explicitly simulating real-world system dynamics, rather than relying on purely statistical or neural models. SynTiSeD [20] follows this paradigm for smart-living energy data by extracting appliance power consumption patterns from real measurements and simulating resident behaviors through predefined action rules (e.g., appliance on/off events), enabling multi-household scalability, modeling of rare scenarios, and ground-truth generation via interactions among households, residents, and appliances. In autonomous driving, Car-STAGE [21] builds on the CARLA [119] simulator to generate high-dimensional sensor time series by simulating complete driving environments, producing RGB, LiDAR, and Inertial Measurement Unit (IMU) sequences together with actor-level ground truth through physics-based interaction simulation.

3.4. Deep Generative Models

3.4.1. Variational Autoencoders (VAEs)

Next, we review representative VAE-based approaches for TSG across several key research directions.

Irregular generation. Chen et al. [112] introduce Neural Ordinary Differential Equations (Neural ODEs) into VAEs, proposing Latent ODEs to model irregularly sampled time series in continuous time. By parameterizing latent dynamics with ODEs solved by black-box solvers, the model handles arbitrary time intervals with constant memory cost and supports generation via continuous normalizing flows. Building on this, Zhou et al. [106] propose LS4, which combines ODE-based latent dynamics with the convolutional efficiency of deep state space models (S4). LS4 injects stochasticity across the entire latent trajectory and replaces explicit hidden-state updates with Fast Fourier Transform (FFT) based convolutions, enabling efficient modeling of sharp transitions and achieving significantly faster training on long, irregular sequences. Further, Naiman et al. [109] propose KoVAE by incorporating Koopman theory into the VAE prior, modeling latent dynamics with linear Koopman operators. Coupled with an NCDE embedding to handle irregular sampling, this design improves stability and interpretability while delivering robust generation quality across regular and irregular time series.

Long sequence generation. For long-horizon synthesis, VAEs become particularly effective when combined with Transformers. ProTran [104] introduces a probabilistic Transformer that integrates state-space modeling with self-attention in the latent space, enabling non-Markovian, long-range dependency modeling. By using hierarchical stochastic latent layers, ProTran supports non-autoregressive long-sequence generation with uncertainty estimation, outperforming RNN- and SSM (State-Space Models)-based baselines. Leveraging vector quantization, TimeVQVAE [107] proposes a two-stage framework for long TSG. It first decomposes sequences in the time-frequency domain into low- and high-frequency components, encodes and quantizes them into discrete tokens, and then employs bidirectional Transformers to model global temporal structure. This design stabilizes training and enables both unconditional and conditional long-sequence generation. Further, SDformer [120] further leverages the VQ-VAE with similarity-driven vector quantization to learn high-quality discrete token representations of time series, and then employs discrete Transformer models to model token-level data distribution for generation. The discrete-token mechanism reduces the modeling complexity of long sequences, and when combined with the long-range modeling capability of Transformers, it enables SDFormer to perform effectively on long-sequence generation tasks.

Others. We briefly review VAE-based approaches that target multi-rate modeling, interpretability, and causal generation. Che et al. [103] propose the Multi-Rate Hierarchical Deep Markov Model (MR-HDMM) for multi-rate multivariate time series. It employs a hierarchical latent structure with learnable switches and auxiliary connections to directly model dependencies across different sampling rates, avoiding explicit up-/down-sampling, and supports generation, forecasting, and interpolation via variational inference. Desai et al. [105] introduce TimeVAE, which balances realism, interpretability, and efficiency. It includes a convolutional VAE for end-to-end generation and an interpretable variant whose decoder explicitly models level, trend, and seasonality, enabling structured generation while

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

benefiting from the denoising effect of the VAE bottleneck. Li et al. [108] propose the Causal Recurrent VAE, which embeds Granger causality into a recurrent VAE. By using a sparsity-regularized multi-head decoder to learn a causal adjacency matrix and an error-compensation module for noise, it enables causality-aware, high-fidelity TSG with causal graph discovery.

3.4.2. Generative Adversarial Networks (GANs)

Below, we review representative GAN-based methods, highlighting their innovations in irregular sampling, conditional control, loss design, and multivariate modeling.

Irregular generation. Bridging GANs and stochastic processes, Neural SDEs [88] formulate TSG in continuous time by using neural stochastic differential equations as generators within a Wasserstein GAN. Noise is injected via Brownian motion to produce stochastic trajectories, while a neural-controlled differential equation (NCDE) discriminator handles irregular sampling. This adversarial setup learns arbitrary drift and diffusion functions directly from data, outperforming latent ODEs and continuous normalizing flows in capturing irregularity and stochasticity. GT-GAN [94] unifies regular and irregular generation by combining GANs, autoencoders, and neural differential equations. An autoencoder maps sequences to fixed-dimensional latent vectors, which are adversarially generated and decoded into continuous-time paths, enabling flexible resampling under missing or sparse observations. RTSGAN [87] addresses variable-length and missing-value time series via an autoencoder–Wasserstein GAN framework with a length-invariant latent space. Its extended version explicitly models missingness patterns by first predicting observation masks and timestamps, then generating values, yielding superior realism on clinical datasets with severe irregularity. Finally, SparseGAN employs sparse self-attention with the 1.5-entmax transformation to model cross-time feature dependencies, selectively ignoring irrelevant or missing time steps and avoiding fixed-interval assumptions in irregular TSG.

Long sequence generation. AEC-GAN mitigates distribution shift and bias accumulation in autoregressive generation by introducing an error correction module that reconstructs original data from adversarially perturbed sequences, enabling dynamic correction of historical biases and stable generation of arbitrarily long sequences. In contrast, TTS-GAN avoids the limitations of RNN-based GANs by adopting a pure transformer encoder for both generator and discriminator. By representing time series as $1 \times W$ images, partitioning them into patches with positional encoding, and leveraging self-attention to capture global temporal dependencies, it efficiently generates realistic long multi-dimensional sequences without recurrence.

Condition exploration. TSGAN [84] adopts a two-stage WGAN [121] architecture for few-shot settings, where one WGAN generates spectrograms from noise and a conditional WGAN synthesizes time series conditioned on these frequency-domain representations, effectively exploiting spectral priors to alleviate data scarcity. To address class imbalance, BFGAN [96] introduces an *importance label* condition derived from kNN and local outlier probability scores, guiding generation toward class-boundary and sparse regions via a CatGAN-style [122] auxiliary classifier. For multi-label clinical EHR data, MTGAN [97] uses a smooth conditional matrix that propagates disease labels across visits, mitigating rare-disease imbalance. RCGAN [79] conditions recurrent GANs on auxiliary labels (e.g., patient states) to generate realistic multivariate sequences, while HealthGAN [81] conditions generation on both static covariates and time series summary statistics, enabling joint synthesis of longitudinal signals and covariate-dependent dynamics.

Loss function exploration. TimeGAN [123] augments adversarial training with a supervised autoregressive loss that aligns stepwise conditional distributions between real and generated data, improving temporal consistency and training stability. CEGEN [95] derives its loss from Euler-discretized stochastic differential equations, introducing a conditional Wasserstein loss that explicitly matches transition distributions to accurately estimate drift and volatility. COT-GAN [82] further refines optimal-transport-based losses by incorporating causal constraints and entropic regularization, ensuring non-anticipative temporal dependencies while enabling efficient Sinkhorn optimization [124]. Beyond adversarial losses, TimeGCI [125] combines moment matching with a contrastive energy-

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

based objective to stabilize long-horizon generation, and TsT-GAN [91] enhances Transformer-based GANs with masked modeling and moment-matching losses alongside LS-GAN loss [126]. In contrast, SigCWGAN [127] replaces adversarial training with a supervised Sig-W_1 [128] loss based on path signatures, replacing the trainable parametric discriminator with an analytic discriminant criterion and substantially improving training stability. Finally, CWGAN-GP [100] adapts conditional Wasserstein GANs with gradient penalty to data-scarce settings, where its balanced loss design helps better generate synthetic electric vehicle demand time series and directly improves the downstream forecasting task.

Multivariate TS exploration. COSCI-GAN [93] employs channel-wise generators driven by a shared latent noise and a central discriminator, preserving both per-channel fidelity and cross-channel correlations. 3DAE-GAN [101] combines autoencoders with 3D convolutions to reshape 1D MTS into 3D representations, enabling joint modeling of temporal and inter-feature dependencies without external conditions. Conditional MTS generation further incorporates side information to guide synthesis. MTGAN [97] conditions on disease labels using a GRU-based generator and a smooth conditional matrix to improve rare-disease generation under data imbalance. DCTC-GAN [102] leverages industrial labels and a dual-channel Transformer (temporal and spatial encoders) to jointly capture long-range temporal dynamics and sensor correlations. Time-Trans [129] integrates autoencoders and GANs, combining TCNs for local patterns and Transformers for global dependencies via parallel layers and bidirectional cross-attention. Finally, RCGAN [79] conditions on auxiliary labels and uses recurrent networks in both generator and discriminator, enabling controllable generation of real-valued multivariate time series.

3.4.3. Diffusion Models

Next, we begin by introducing diffusion models for TSG with condition exploration.

Condition exploration. DiffTraj [55] introduces a conditional diffusion framework for trajectory generation, using a Traj-UNet with residual blocks and multi-scale fusion to model spatio-temporal dynamics, while embedding attributes such as velocity and departure time via Wide & Deep networks to guide denoising. ControlTraj [64] extends this idea with topology-aware conditions, employing a Masked Road Autoencoder (RoadMAE) to encode road structure and a GeoUNet with geographic attention to fuse structural and trajectory attributes for structurally consistent generation. FIDE [61] focuses on extreme-value preservation, using block maxima (peak values within predefined time windows) as explicit conditions and incorporating the Generalized Extreme Value (GEV) distribution into the loss to align both tail behavior and overall distributions. DiffTime [130] enables flexible conditional control without retraining by imposing differentiable constraints (e.g., fixed values, trends, extrema, and multivariate relations) at inference time to steer a pretrained diffusion model. Li et al. [71] propose risk-sensitive diffusion, which conditions generation on auxiliary risk vectors encoding sample quality (e.g., imputation uncertainty or measurement error). These vectors parameterize a risk-sensitive SDE (stochastic differential equation) during denoising, enabling robust conditional generation under noisy and non-Gaussian time series. For multi-domain generation, TimeDP [30] introduces label-free domain prompts derived from prototype importance weights encoding basic temporal patterns, supporting adaptation to unseen domains. Similarly, OATS [76] leverages prototypes of high-value samples as conditioning signals to dynamically generate high-quality data during downstream model training, thereby enabling effective data augmentation. CaTSG [75] extends conditional TSG by incorporating Pearl's causal ladder to enable observational, interventional, and counterfactual generation, thereby mitigating spurious correlations induced by unobserved confounders. It leverages a diffusion framework with backdoor-adjusted guidance and a learnable environment bank. WaveStitch [73] conditions diffusion on temporal and geographic features, handles unseen conditions via constraint–signal interaction, and accelerates generation through parallel segment synthesis with stitching. DOSE [51] mitigates condition collapse via diffusion dropout, randomly discarding intermediate samples during training to strengthen reliance on conditioning signals.

Text-based generation. Evans et al. [131] compress raw audio time series into compact latent codes via an autoencoder, align text and audio semantics using a CLAP-based embedding model [132],

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

and employ a diffusion transformer (DiT) with block-wise attention and rotary positional encoding to support long-sequence generation. Similarly, Woo et al. [133] use a pre-trained BERT language model to parse unstructured text descriptions, and cross-attention mechanisms fuse text embeddings from BERT with feature maps of the temporal U-Net, enabling the diffusion model to align generation with text semantics. Recently, BRIDGE [29] addresses cross-domain and constrained generation by treating text as the control signal: it uses an LLM-based multi-agent framework to synthesize text-time series pairs and integrates semantic prototypes with textual descriptions in a hybrid diffusion model to enhance controllability. Further, VerbalTS [134] proposes a multi-focal text processor to extract structured multi-semantic representations (covering diffusion-stage, temporal, and spatial perspectives) via learnable anchors. These text features are then dynamically integrated into the diffusion process through a semantic alignment adapter. During denoising, a multi-view noise estimator models time series from temporal, spatial, and diffusion dimensions, using text-guided parameters to refine generation. Emphasizing domain-agnostic text-to-series generation, T2S [74] proposes a diffusion framework that unifies natural language and time series across domains, combining a length-adaptive VAE for variable-length encoding with flow matching and a DiT-based cross-modal alignment scheme [135], enabling arbitrary-length sequence generation.

Fast sampling. DiffWave [23] is a diffusion-based audio waveform generator that adopts a non-autoregressive bidirectional dilated convolution architecture, and accelerates sampling by collapsing hundreds of training diffusion steps into a small number of inference steps through a carefully designed variance schedule. Building on fast sampling, DOSE [51] further reduces inference cost via a model-agnostic two-step scheme: a *coarse estimation* + *fine correction* process replaces iterative denoising, and conditioning signals (e.g., noisy speech) are embedded once into an adaptive prior rather than injected at every step. as a model-agnostic method, DOSE can reuse standard diffusion backbones such as DiffWave without additional architectural complexity, achieving substantial sampling speedups.

Variable length generation. T2S [74] introduces a length-adaptive VAE (LA-VAE) that encodes variable-length time series into unified latent spaces via dynamic up/downsampling, combined with a latent consistency loss and interleaved multi-length training to avoid forgetting. A DiT denoiser with flow matching then enables arbitrary-length generation. To handle irregular and high-dimensional sequences, TS-Diffusion [27] implicitly supports variable lengths through a point-process formulation and neural ODE (Ordinary Differential Equation) encoder, learning continuous-time representations without fixed-length constraints, while its ODE decoder naturally regenerates sequences of varying durations. Extending diffusion to tabular time series, TabDiT [136] employs a non-holistic VAE to encode rows independently and a DiT to model temporal dependencies, explicitly predicting padding rows as end-of-sequence markers to allow flexible output lengths. Naiman et al. [137] map TS of arbitrary length into images via invertible transforms, apply visual diffusion models to variable-size images, and invert them back, unifying short to ultra-long sequence generation without modifying model architectures. Finally, Diff-MN [77] is a continuous TSG framework that enhances standard NCDE with a Mixture-of-Experts (MoE) dynamics function and a decoupled architecture for dynamics-focused training. It employs a diffusion model to parameterize NCDE temporal dynamics, jointly learning the distribution of TS data and MoE weights, enabling sample-specific MoE-NCDE for continuous generation of variable length and high-frequency TS with richer contextual information.

Loss function exploration. Diffusion-TS [62] improves interpretability by directly reconstructing clean time series at each diffusion step (instead of noise), using a Fourier-based loss that jointly preserves time- and frequency-domain characteristics, together with a Transformer encoder–decoder that disentangles trend, seasonality, and residual components. Addressing population-level distribution shifts overlooked by Diffusion-TS, PaD-TS [72] introduces a population-aware objective that augments reconstruction loss with Maximum Mean Discrepancy (MMD) regularization to better preserve value distributions and functional dependencies across datasets. For conditional sequential tasks, CDDRec [68] proposes a cross-divergence loss based on KL divergence to enforce ranking awareness, combined with contrastive (InfoNCE) objectives to mitigate collapse and over-smoothing

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

in diffusion-based recommendation. Finally, TimeBridge [138] replaces the fixed Gaussian prior with a diffusion bridge, enabling data- and time-dependent priors (e.g., Gaussian-process-based) and scale-preserving constraints for conditional generation. To support these flexible priors, it redesigns the loss with Fourier-term complements and adjusted weighting schedules for more effective optimization. TSGDiff [139] is a graph-based framework for multivariate TSG that models TS as dynamic graphs, leveraging a graph neural network encoder–decoder and a latent diffusion process. A hybrid loss combining reconstruction, KL divergence, denoising, and Fourier losses is employed to effectively capture structural distributions.

3.4.4. Normalizing Flows

Chen et al. [112] introduce Neural Ordinary Differential Equations (Neural ODEs) as continuous normalizing flows by replacing discrete network layers with ODE solvers, achieving constant-memory backpropagation via the adjoint method and tractable trace-based Jacobian computation. While this ensures continuity in latent dynamics, continuity across observed time points is not explicitly enforced. To address this, Deng et al. [113] propose the Continuous-Time Flow Process (CTFP), which deforms a Wiener process via a dynamic normalizing flow to model complex observable processes, also enabling efficient likelihood estimation, sampling, and interpolation for irregular time series. Beyond pure flow-based modeling, Cheng et al. [115] integrate normalizing flows into the CausalTime framework to generate realistic time series with explicit ground-truth causal graphs: a causally disentangled neural network models dynamics, flows capture noise distributions, and causal structures are extracted via DeepSHAP [140,141] or prior knowledge to decompose autoregressive dynamics into causal and residual components. Shifting to the frequency domain, Alaa et al. [114] propose Fourier Flows, which apply normalizing flows after transforming time series via the Discrete Fourier Transform. By learning spectral filters through affine coupling layers and bidirectional RNNs, this approach efficiently handles variable-length and irregularly sampled sequences, leveraging FFT for fast Jacobian computation while enabling expressive time-domain transformations.

4. Applications

TSG technology demonstrates significant practical value in alleviating data scarcity, enhancing model training, and supporting scenario simulation across diverse domains. This section reviews its typical applications in finance, healthcare, industry, and energy, including the core generative methods and their practical contributions in each field. The key information is summarized in Table 3.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

Table 3. Summary of Generative Time Series Models Across Application Fields.

Application Field	Source	Model	Core Summary	Application Task
Finance	Fu et al. [80]	GAN	Risk analysis (VaR/ES estimation, CCAR scenario generation)	Risk analysis, time series forecasting
	Gao et al. [67]	Transformer + Diffusion	Stock factor augmentation, mitigate data collision	Stock prediction
	Huang et al. [60]	Diffusion	Generate irregular/scale-invariant financial time series	Stock prediction (data augmentation)
	Ericson et al. [142]	GAN	Comparative analysis for VaR estimation, yield curve modeling	VaR estimation, yield curve modeling
	Hamdouche et al. [143]	Schrödinger Bridge (Generative)	Deep hedging of options to enhance strategy robustness	Financial risk management
	Lu and Sester [144]	MMD (Kernel-based)	Synthesize stylized financial time series	Reinforcement learning-based portfolio management
	Dogariu et al. [145]	GAN	Synthesize US stock OHLC prices/log returns	Forecasting model training (data augmentation)
	Wiese et al. [146]	GAN	Generate S&P 500 series with realistic market dynamics	Trading strategy design, financial valuation
Healthcare	Esteban et al. [79]	GAN	Generate realistic ICU medical time series (heart/respiratory rate)	Early warning system training, medical simulation
	Dash et al. [81]	GAN	Synthetic time series with static covariates (age/gender)	In-hospital mortality/decompensation prediction
	Lu et al. [147]	GAN	Address rare disease imbalance in EHRs	Heart failure/Parkinson's disease prediction
	Kuo et al. [148]	Diffusion	Generate mixed-type longitudinal EHRs	Clinical decision support (RL agent training)
	Schürch et al. [149]	GAN	Personalized insulin therapy, glucose trajectory prediction	Diabetes treatment optimization
	Zhu et al. [150]	GAN	Personalized glucose time series for Type 1 Diabetes	Glucose prediction, pre-clinical trials
	Klein et al. [59]	Diffusion	Generate subject-specific EEG signals (ERP paradigms)	BCI research (alleviate data scarcity)
	Deng et al. [70] (TarDiff)	Diffusion	Task-oriented synthetic EHR time series	Mortality/ICU stay prediction, disease diagnosis
Industry	Zhang et al. [85]	GAN	Remaining Useful Life (RUL) estimation	RUL estimation
	Zhang et al. [151]	Transformer + GAN	Multivariate TSG	RUL prediction
	Zenisek et al. [11]	Rule-based Simulation	Stream synthetic time series for condition monitoring	Predictive maintenance (MQTT protocol)
	Sun et al. [90]	GAN	Small-sample online reliability assessment	Online reliability evaluation
	Dai et al. [54] (TimeDDPM)	Diffusion	Diffusion-based generation with spatiotemporal extraction	Industrial soft sensing (quality prediction)
	Yi et al. [66] (TSDM)	Generative Model	Vibration signal generation for rotating machinery	Small-sample fault diagnosis
	Ma et al. [152]	GAN + CNN-BiLSTM	Address EHA degradation data scarcity	EHA degradation prediction
	Shangguan et al. [153]	GAN + Gamma Process	Train wheel wear data generation	High-speed railway safety (wear prediction)
Energy	Denaxas et al. [154] (SynTiSe)	MCMC (Probabilistic)	High-resolution wind power data generation	Power system planning
	Zhang et al. [7]	GAN	Alleviate smart grid data scarcity/privacy issues	Smart grid downstream tasks (data augmentation)
	Pinceti et al. [86]	GAN	Transmission-level load data (hourly week-long profiles)	Power system simulation
	Meiser et al. [20] (SynTiSeD)	Multi-agent Simulation	Synthetic energy data generation	Scenario simulation, NILM algorithm enhancement
	Dong et al. [57]	Diffusion	Short-term wind power scenario generation (NWP integration)	Stochastic unit commitment (reduce dispatch cost)
	Li et al. [56] (DiffCharge)	Diffusion	EV charging scenario generation (battery/station-level)	Grid operation, day-ahead market bidding
	Fu et al. [58]	Diffusion	Long-term annual building energy data generation	Building energy data analysis
	Wang and Zhang [69]	Diffusion	Customized electricity load profile synthesis	Load forecasting, anomaly detection
	Zhang and Sikdar [101]	GAN + Autoencoder	Multivariate smart grid data (household/EV)	Smart grid analysis (privacy/anti-adversarial)
	Ye et al. [155]	Transformer + GAN	Day-ahead wind power scenario generation (error/seasonal)	Power system dispatching (temporal correlation)
	Li et al. [156]	GMM-HMM (Probabilistic)	Correlated wind farm power output time series	Power system planning (spatiotemporal correlation)
	Talbot et al. [157]	FVARMA (Time Series)	Correlated synthetic time series (solar/electricity demand)	Hybrid energy system uncertainty analysis
	Brentan et al. [158]	Statistical Model	Water demand time series (trends/randomness/climate)	Water network modeling, demand forecasting
	Tabeshpour and Belvasi [159]	Fourier Series (Statistical)	Artificial ocean wave time series	WEC dynamic analysis, design optimization

4.1. Finance

Fu et al. [80] apply CGANs to financial risk analysis, enabling Value at Risk (VaR) and Expected Shortfall (ES) estimation, Comprehensive Capital Analysis and Review (CCAR) scenario generation, and economic time series forecasting. Gao et al. [67] propose DiffsFormer for stock factor augmentation, improving forecasting performance on China Securities Index 300 (CSI300) and China Securities Index 800 (CSI800) while alleviating data collision issues. Huang et al. [60] develop FTS-Diffusion to generate irregular and scale-invariant financial time series, effectively reducing stock prediction errors via data augmentation. Ericson et al. [142] compare historical simulation, parametric models, and deep generative approaches (e.g., CGAN, CWGAN) for VaR estimation and yield curve modeling. Hamdouche et al. [143] introduce a Schrödinger bridge (SB)-based generative model that produces realistic synthetic time series and demonstrates strong utility in financial risk management, particularly for deep hedging of options, improving strategy robustness, and reducing replication error.

Lu and Sester [144] employ Maximum Mean Discrepancy (MMD) with a signature kernel to generate stylized financial time series, enabling downstream applications such as reinforcement-learning-based portfolio management. In a similar application-oriented way, Dogariu et al. [145] use GAN variants to synthesize U.S. stock market data (e.g., OHLC prices and log returns) to enhance forecasting model training, while Wiese et al. [146] propose Quant GANs to generate realistic Standard & Poor's 500 (S&P 500) financial series, supporting trading strategy design and financial valuation tasks.

4.2. Healthcare

Esteban et al. [79] apply Recurrent Conditional GANs (RCGAN) to synthesize real-valued ICU time series (e.g., heart rate, respiratory rate) for training early warning systems and medical simulators. Dash et al. [81] propose HealthGAN to jointly generate temporal signals and static covariates (e.g., age, gender), validating its utility on MIMIC-III clinical tasks such as mortality and decompensation prediction. To address disease imbalance in EHRs, Lu et al. [147] introduce MTGAN to generate multi-label patient diagnosis sequences from MIMIC-III/IV, improving predictions for diseases like heart failure and Parkinson's. Kuo et al. [148] leverage diffusion probabilistic models to generate mixed-type longitudinal EHRs (numeric, binary, categorical), supporting reinforcement-learning-based clinical decision making for hypotension, HIV therapy, and sepsis with reduced privacy risk.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

In personalized treatment modeling, Schürch et al. [149] generate insulin strategies and glucose trajectories for hospitalized diabetes patients, while Zhu et al. [150] propose GluGAN to synthesize personalized glucose series for Type 1 Diabetes, aiding prediction models and pre-clinical evaluation. Beyond EHRs, Klein et al. [59] develop a conditional diffusion model to generate subject- and task-specific EEG signals for Brain-Computer Interfaces, and TarDiff [70] produces task-oriented synthetic healthcare time series for mortality, ICU stay prediction, and physiological signal-based diagnosis across multiple datasets, alleviating data scarcity and class imbalance.

4.3. Industry

Zhang et al. [85] apply a convolutional recurrent GAN to generate degradation time series for remaining useful life (RUL) estimation, validated on NASA aero-engine and lithium-ion battery datasets, while Zhang et al. [151] further propose a dual-channel Transformer conditional GAN for multivariate generation, improving RUL prediction on turbofan engines and FEMTO bearings. For industrial monitoring, Zenisek et al. [11] developed a configuration file-driven software tool to stream synthetic time series data, simulating condition monitoring of industrial production plants via the Message Queue Telemetry Transfer protocol for predictive maintenance. Sun et al. [90] presented a worm Wasserstein generative adversarial network for small-sample online reliability assessment, demonstrated through the online reliability evaluation of Lithium-ion battery cells with extremely limited time series data. TimeDDPM [54] leverages diffusion-based generation for industrial soft sensing, augmenting limited samples to improve quality-variable prediction in dynamic processes, while TSDM [66] targets vibration signal generation for rotating machinery, enhancing small-sample fault diagnosis across bearing datasets. Finally, TimeGAN-based solutions are applied to degradation prediction in practice: the CNN-BiLSTM-Attention model combined with TimeGAN [152] addresses data scarcity in electro-hydrostatic actuator degradation, and Shangguan et al. [153] propose an improved TimeGAN with a stationary gamma process to generate train wheel wear data for reliable degradation prediction and railway safety assurance.

4.4. Energy

SynTiSe [154] employs a modified multi-regime Markov Chain Monte Carlo (MCMC) method with percentile-based discretization to generate high-resolution wind power time series for power system planning. To address data scarcity and privacy in smart grids, Zhang et al. [7] and Pincet et al. [86] propose conditional GANs for distribution- and transmission-level load data, respectively, generating realistic customized load profiles for system simulation, while SynTiSeD [20] extends this line with a multi-agent framework for large-scale synthetic energy data generation, improving Nonintrusive Load Monitoring (NILM) performance. Diffusion models are increasingly adopted in energy applications: the Conditional Latent Diffusion Model (CLDM) [57] leverages Numerical Weather Prediction (NWP) data for short-term wind power scenario generation, reducing dispatch costs in stochastic unit commitment; DiffCharge [56] generates electric vehicle (EV) charging scenarios at battery and station levels for grid operation and market bidding; and Fu et al. [58] design a metadata-driven conditional diffusion model for long-term building energy data synthesis. Wang and Zhang [69] further propose an attention-based conditional diffusion approach for customized electricity load profile generation, supporting load forecasting and anomaly detection. Beyond diffusion, Zhang and Sikdar [101] introduce a 3D Autoencoder GAN (3DAE-GAN) to synthesize multivariate smart-grid data (household consumption and EV driving) with privacy protection, while Ye et al. [155] combine Informer [160] and TimeGAN [123] for day-ahead wind power scenarios with improved temporal correlation. Classical statistical models remain relevant: Li et al. [156] use a Gaussian Mixture Model–Hidden Markov Model (GMM-HMM) for correlated multi-wind-farm generation, Talbot et al. [157] propose Fourier vector autoregressive moving average models for energy system uncertainty analysis, Brentan et al. [158] generate urban water demand time series, and Tabeshpour and Belvasi [159] synthesize ocean wave time series for wave energy converter design and analysis.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

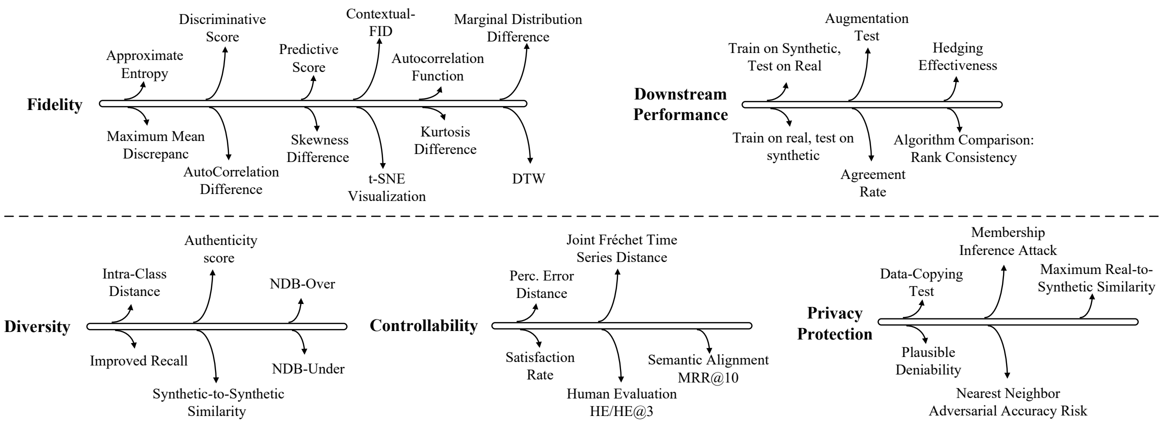

Figure 5. An overview of metrics for evaluating TSG quality across the dimensions of fidelity, diversity, controllability, downstream-task performance, and privacy protection.

5. Evaluation

Existing evaluations mainly focus on data quality and utility, with fidelity and downstream task performance as the primary dimensions [35,161,162]. By contrast, controllability, the ability to satisfy explicit generation conditions, is rarely considered as an independent evaluation dimension and is often absorbed into conditional generation or fidelity-based assessments.

In this review, we organize evaluation criteria into five core dimensions: *Fidelity*, *Diversity*, *Controllability*, *Downstream performance*, and *Privacy protection*. This categorization explicitly elevates controllability while maintaining a balanced focus on intrinsic quality, practical utility, and ethical considerations, providing a concise and application-oriented framework for comparing state-of-the-art TSG methods. A comprehensive overview of all metrics is shown in Figure 5.

5.1. Fidelity

Fidelity (also referred to as realism or authenticity) measures how closely synthetic time series resemble real data in terms of statistical properties, temporal dynamics, correlation structures, and overall distributional characteristics [163]. High fidelity is fundamental, as synthetic data that fails to capture real data characteristics is unlikely to be useful in downstream applications.

Approximate Entropy (ApEn) [164,165] measures the regularity and complexity of a time series by quantifying how often similar patterns persist over time. It is computed per channel for real and synthetic data, with squared differences used for comparison. Lower values indicate better preservation of temporal complexity. Discriminative Score (DS) [94,109] evaluates realism by training a GRU/LSTM-based classifier to distinguish real from synthetic series. Lower classification accuracy (or higher error) is better, as it implies synthetic data is indistinguishable from real data. Predictive Score (PS) [94,109,125] trains a forecasting model (e.g., GRU/LSTM) on synthetic data and evaluates it on real data using MAE or RMSE. Lower prediction error indicates better preservation of forecasting-relevant temporal dependencies. Contextual Fréchet Inception Distance (C-FID) [166] extends FID to time series by computing Fréchet distance between embeddings (e.g., from ts2vec) of real and synthetic data. Lower values reflect closer alignment of global and contextual features. Maximum Mean Discrepancy (MMD) [30,79,167,168] is a kernel-based distance in a reproducing kernel Hilbert space, measuring global distribution mismatch. lower values indicate better distributional similarity. Marginal Distribution Difference (MDD) [35,128] computes average absolute differences between marginal histograms across dimensions and time steps. Lower values indicate better preservation of basic statistics such as scale and variance.

Auto-Correlation Difference (ACD) [35,169,170] compares autocorrelation coefficients of real and synthetic series to assess temporal dependency preservation; lower values are better. Skewness Difference (SD) and Kurtosis Difference (KD) [35,161] measure discrepancies in distribution asymmetry and tail behavior. Lower values indicate better preservation of higher-order statistics. Dynamic

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

Time Warping (DTW) and Multivariate DTW (DTW_D) [35,171,172] quantify similarity under flexible temporal alignment for univariate and multivariate series, respectively; lower values indicate closer temporal pattern alignment. Autocorrelation Function (ACF) [146,161] compares lag-dependent correlations (e.g., of raw or transformed signals) between real and synthetic data to capture dependence properties such as volatility clustering. Smaller discrepancies imply better fidelity. Finally, distribution visualization using t-SNE and distribution plots [94,109,173] provides qualitative assessment by visualizing embedding overlap and comparing statistical distributions, where greater overlap and similar shapes indicate higher fidelity.

5.2. Downstream Performance

Downstream performance evaluates whether synthetic data is useful for real tasks such as classification, forecasting, and decision-making.

Train on Synthetic, Test on Real (TSTR) and Train on Real, Test on Synthetic (TRTS) [79,93] are widely used utility metrics. TSTR evaluates whether synthetic data contains sufficient information for real-world tasks, while TRTS assesses whether it preserves key statistical structures of real data. Higher scores are preferred for classification, whereas lower errors indicate better forecasting performance. Augmentation Test [174] measures performance gains when synthetic data augments real data, with larger improvements indicating stronger complementarity. Agreement Rate [175] quantifies prediction consistency between models trained on real and synthetic data; higher values imply better preservation of decision boundaries. Hedging Effectiveness [176] is a finance-specific metric evaluating the usefulness of synthetic data for risk modeling and hedging, where higher values are better. Algorithm Comparison via Rank Consistency [177] examines whether algorithm performance rankings on synthetic data match those on real data, with higher Spearman correlation indicating stronger consistency.

5.3. Diversity

Diversity measures whether synthetic time series cover the full variability of real data, rather than collapsing to a small subset of modes [87]. High diversity ensures representativeness and robustness, especially for data augmentation and rare-pattern modeling.

Intra-Class Distance (ICD) [164] measures the average distance among synthetic samples within the same class. Higher values indicate greater intra-class variability and reduced mode collapse. Synthetic-to-Synthetic Similarity (STS) [178] computes intra-class cosine similarity among synthetic samples, where values not close to 1 are preferred to avoid conditional mode collapse. Improved Recall [179] evaluates how well synthetic data covers the full spectrum of real data patterns, with higher values indicating better diversity and reduced mode dropping. NDB-Over and NDB-Under [180] partition real data into k-means bins and identify over- or under-represented regions by the generator. Values closer to $(0, 0)$ indicate balanced generation. Finally, the Authenticity score [163] measures the proportion of novel (non-duplicated) synthetic samples, where values closer to 1 indicate stronger diversity and originality.

5.4. Controllability

Controllability refers to a generator's ability to produce time series that satisfy predefined conditions (e.g., text descriptions), which is essential for targeted data synthesis. Despite its importance, controllable generation remains underexplored in existing evaluation surveys. Below, we summarize representative controllability metrics.

T2S [74] evaluates controllability by measuring semantic and numerical alignment between generated time series and text captions at point-, fragment-, and instance-level granularity. It adopts Mean Squared Error (MSE) and Weighted Absolute Percentage Error (WAPE) to quantify numerical deviations (lower is better), reflecting how accurately text-described trends are translated into values. Mean Reciprocal Rank at 10 (MRR@10) evaluates semantic alignment via cosine similarity (higher is better), indicating how closely generated sequences match the textual intent.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

BRIDGE [29] assesses controllability using both quantitative and human evaluation. Quantitatively, it employs Joint Fréchet Time Series Distance (J-FTSD), Mean Squared Error (MSE), and Mean Absolute Error (MAE), all lower-better metrics. J-FTSD jointly captures local shape and global distribution alignment, providing a more faithful measure of adherence to text-specified patterns. Human Evaluation (HE/HE@3) ranks outputs by text relevance and plausibility, where lower ranks indicate better perceptual alignment.

Coletta et al. [130] evaluate controllability by verifying compliance with predefined constraints (soft/hard, local/global). Percentage Error Distance (L2 distance to the target trend) measures adherence to soft trend constraints (lower is better), while Satisfaction Rate quantifies the proportion of samples meeting hard constraints such as fixed points or extrema (higher is better). Auxiliary metrics, including discriminative score, predictive score, and inference time, ensure that improved controllability does not compromise realism or efficiency.

5.5. Privacy Protection

Privacy protection evaluates whether synthetic data leaks sensitive information from training data. Membership Inference Attack (MIA) [181] quantifies membership inference risk using the F1 score. Lower values are better, indicating weaker traceability of synthetic samples to the training data. Nearest Neighbor Adversarial Accuracy Risk (NNAA) [9] measures re-identification risk by assessing how easily synthetic EHR time series can be distinguished from real data. Lower values are better, indicating reduced overfitting and privacy leakage. Plausible Deniability [175] ensures each synthetic sample could originate from multiple real samples. Higher values are better, reducing traceability. Data-Copying Test [182] detects overfitting on training samples by distance comparison. Values closer to 0 are better, indicating no memorization and preventing leakage of private information from the training set. Maximum Real-to-Synthetic Similarity (Max-RTS) [178] measures the highest similarity between real and synthetic samples. Lower is better, reducing re-identification risk.

6. Challenges and Opportunities

The field of TSG has made rapid progress, yet several open challenges continue to limit its applicability and reliability. These challenges arise from both the intrinsic complexity of time series data and the stringent demands of real-world applications, offering important directions for future research.

6.1. Perspective of Data Characteristics

First, **irregularly sampled and asynchronous multivariate data** pose persistent challenges. In domains such as healthcare and IoT, observations are event-driven and recorded at heterogeneous frequencies. Continuous-time models, including Neural ODEs [112], ODE-RNNs [183], and Neural Controlled Differential Equation (NCDE)-based methods [88,95,109], partially alleviate this issue but often incur high computational costs or struggle to capture complex cross-channel dependencies. Future work can explore more efficient continuous-time generative frameworks that natively model asynchronous dynamics without relying on lossy interpolation or imputation.

Second, **preserving long-range dependencies and global structure** remains non-trivial. Many models generate locally plausible sequences that violate long-term constraints such as physical laws, economic cycles, or seasonal patterns. While Transformers [91] and state-space models [104,106] improve long-context modeling, scaling them to ultra-long horizons with both fidelity and diversity remains open. **Future work can explore** hybrid architectures that combine global modeling (e.g., attention or state-space mechanisms) with local refinement via diffusion or autoregressive processes.

Third, **modeling extreme events and tail behavior is under-explored**. Since most objectives emphasize average-case fidelity, rare but critical events, such as financial crashes, equipment failures, or medical emergencies, are often underrepresented [147]. Methods that explicitly condition on extremes, such as block-maxima conditioning in FIDE [61], are promising but still limited in generality. **Future**

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

work can explore integrating extreme value theory, tail-aware loss functions, or importance-weighted objectives to better capture rare-event dynamics.

Fourth, **cross-domain generalization and multimodal generation** are emerging frontiers. Real-world data span multiple domains (e.g., hospitals, factories, climate regions) with distinct distributions, yet most models are trained within a single domain. Early cross-domain approaches such as TimeDP [30] remain preliminary. In parallel, multimodal generation, such as text to time series [29,74,131], raises challenges in aligning heterogeneous modalities (text, images, sensor signals). Future work can explore domain-agnostic and adaptive generative models via meta-learning, invariant representation learning, or causal modeling, alongside robust alignment mechanisms and evaluation protocols for cross-modal controllability.

Finally, **LLM-based TSG** is still nascent. SDForger [184] compresses periodic series into low-dimensional embeddings using Functional Principal Components (FPC) [185] or Fast Independent Component Analysis (FastICA) [186], converts them into text tokens for lightweight large language model (LLM) fine-tuning, and decodes them back into variable-length sequences for few-shot, text-conditioned generation. Meanwhile, emerging agentic systems point to new directions: GenG [187] separates high-level intent from low-level signal synthesis, retrieval-augmented multi-agent frameworks [188] adapt to domain drift, and LLM–diffusion hybrids [29] enable cross-domain, text-guided control. Future work can investigate temporal tokenization, scalability to irregular time series, and agent-based coordination that jointly optimizes context and stochastic generation without compounding errors.

6.2. Perspective of Application Scenarios

A key challenge is **fine-grained controllability and interpretability**. In safety-critical domains such as healthcare or industrial control, users often require precise, user-defined constraints (e.g., specific peak times or failure conditions). While conditional Generative Adversarial Networks (GANs) [79,81] and diffusion models [55,64,130] provide basic control, they typically lack the precision and flexibility needed for complex or counterfactual conditions. Moreover, their black-box nature limits interpretability, hindering trust and debugging. Future work can focus on unifying strong generative capacity with interpretable and human-understandable control mechanisms.

Another issue is the **absence of standardized, application-aligned evaluation protocols**. Current evaluations mainly rely on generic metrics such as fidelity and diversity, while the real value of synthetic data lies in downstream utility (e.g., improving rare-disease classification or reinforcement learning performance). There is a pressing need for task-specific evaluation frameworks that directly measure real-world impact rather than proxy scores [161].

Privacy preservation also remains insufficiently addressed. Although TSG is often positioned as a privacy-preserving solution, deep generative models may memorize and leak sensitive information. Formal guarantees such as Differential Privacy (DP) are rarely integrated. Future work can focus on developing TSG methods with quantifiable privacy guarantees while maintaining acceptable data utility, especially in regulated domains like healthcare and finance.

Finally, **scalability and computational efficiency** pose practical barriers. Many state-of-the-art models, including diffusion models [62] and large-scale VAE/Transformer-based approaches [104,107], are costly to train and sample. This limits deployment in resource-constrained or large-scale simulation settings. Future research into efficient architectures, faster sampling strategies (e.g., [23,51]), and model compression is crucial to bridge the gap between academic advances and industrial adoption.

6.3. Perspective on Modeling Paradigm

Beyond conventional deep generative models, **LLM-based agentic frameworks** have gained attention in areas such as data science automation, scientific discovery, and decision support, where Large Language Models (LLMs) act as reasoning-centric controllers that plan, invoke tools, and iteratively refine outputs. This paradigm introduces new opportunities for TSG. In particular, agentic LLMs can (i) reason over temporal patterns, trends, and event semantics to imagine time series; (ii)

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

generate sequences via explicit simulation of temporal or event-driven processes (e.g., regime shifts or causal chains); and (iii) orchestrate external time series models (e.g., diffusion or state-space models) through dynamic selection and conditioning.

Compared with traditional end-to-end TSG methods centered on distribution matching, agentic LLM-based systems offer notable advantages, including interpretable temporal reasoning, iterative self-refinement via feedback loops, and adaptive model integration through flexible tool orchestration. Together, these properties extend TSG toward more interpretable, controllable, and task-oriented time series synthesis.

7. Conclusions

In this survey, we organize existing TSG methods into a unified and comprehensive taxonomy spanning multiple paradigms, model architectures, research directions, and data characteristics. We review approaches from traditional non-deep learning techniques to modern deep generative frameworks. Our analysis reveals key trends: Deep learning now dominates TSG, diffusion models enable high-fidelity and controllable generation, and LLM-based approaches show promise for multimodal tasks such as text-to-time-series generation. Applications in finance, healthcare, industry, and energy demonstrate TSG's ability to mitigate data scarcity and enhance downstream tasks, including forecasting, anomaly detection, and decision-making. However, evaluation remains fragmented, highlighting the need for standardized, application-oriented protocols beyond generic fidelity metrics.

Despite rapid advances, challenges persist in handling irregular sampling, long-range dependencies, extreme events, cross-domain generalization, and privacy utility trade-offs. As temporal data grows in scale and complexity, TSG will play an increasingly important role in modeling dynamic systems, and this survey aims to guide future research and practical adoption across disciplines.

References

1. Zhang, X.; Wang, P.; Wang, C.; Xu, Z.; Nie, X.; Wang, W. Global Feature Enhancing and Fusion Framework for Strain Gauge Status Recognition. In Proceedings of the Companion Proceedings of the ACM on Web Conference 2025. ACM, 2025, WWW '25, p. 611–620. https://doi.org/10.1145/3701716.3715229.
2. Tian, M.; Chen, B.; Guo, A.; Jiang, S.; Zhang, A.R. Reliable generation of privacy-preserving synthetic electronic health record time series via diffusion models. *Journal of the American Medical Informatics Association* **2024**, *31*, 2529–2539. https://doi.org/10.1093/jamia/ocae229.
3. Zhang, X.; Huang, Z.; Wu, Y.; Lu, X.; Qi, E.; Chen, Y.; Xue, Z.; Wang, Q.; Wang, P.; Wang, W. Multi-period learning for financial time series forecasting. In Proceedings of the Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1, 2025, pp. 2848–2859.
4. Zhang, X.; Huang, Z.; Wu, Y.; Lu, X.; Qi, E.; Chen, Y.; Xue, Z.; Wang, P.; Wang, W. Self-Adaptive Scale Handling for Forecasting Time Series with Scale Heterogeneity. In Proceedings of the ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2024, pp. 7485–7489.
5. Zhang, X.; Wang, P.; Li, Y.; Wang, W. Amortized Predictability-aware Training Framework for Time Series Forecasting and Classification. *arXiv preprint arXiv:2602.16224* **2026**.
6. Zhang, X.; Wang, Q.; Wang, P.; Wang, W. SEMixer: Semantics Enhanced MLP-Mixer for Multiscale Mixing and Long-term Time Series Forecasting. *arXiv preprint arXiv:2602.16220* **2026**.
7. Zhang, C.; Kuppannagari, S.R.; Kannan, R.; Prasanna, V.K. Generative adversarial network for synthetic time series data generation in smart grids. In Proceedings of the 2018 IEEE international conference on communications, control, and computing technologies for smart grids (SmartGridComm). IEEE, 2018, pp. 1–6.
8. Narasimhan, S.S.; Agarwal, S.; Akcin, O.; Sanghavi, S.; Chinchali, S.P. Time Weaver: A Conditional Time Series Generation Model. *International Conference on Machine Learning* **2024**. https://doi.org/10.48550/arxiv.2403.02682.
9. Yale, A.; Dash, S.; Dutta, R.; Guyon, I.; Pavao, A.; Bennett, K.P. Generation and evaluation of privacy preserving synthetic health data. *Neurocomputing* **2020**, *416*, 244–255.
10. Forestier, G.; Forestier, G.; Petitjean, F.; Petitjean, F.; Dau, S.H.; Dau, H.; Webb, G.I.; Webb, G.I.; Keogh, E.; Keogh, E. Generating Synthetic Time Series to Augment Sparse Datasets. *Industrial Conference on Data Mining* **2017**. https://doi.org/10.1109/icdm.2017.106.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

11. Zenisek, J.; Wolfartsberger, J.; Sievi, C.; Affenzeller, M. Streaming synthetic time series for simulated condition monitoring. *IFAC-PapersOnLine* **2018**, *51*, 643–648.
12. Turowski, M.; Weber, M.; Neumann, O.; Heidrich, B.; Phipps, K.; Çakmak, H.K.; Mikut, R.; Hagenmeyer, V. Modeling and generating synthetic anomalies for energy and power time series. In Proceedings of the Proceedings of the Thirteenth ACM International Conference on Future Energy Systems, 2022, pp. 471–484.
13. Klopries, H.; Torres, D.O.S.; Schwung, A. Synthetic time series dataset generation for unsupervised autoencoders. *IEEE International Conference on Emerging Technologies and Factory Automation* **2022**. https://doi.org/10.1109/etfa52439.2022.9921598.
14. Taga, E.O.; Ildiz, M.E.; Oymak, S. TimePFN: Effective multivariate time series forecasting with synthetic data. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2025, Vol. 39, pp. 20761–20769.
15. Ferdous, M.H.; Hossain, E.; Gani, M.O. TimeGraph: Synthetic Benchmark Datasets for Robust Time-Series Causal Discovery. *arXiv preprint arXiv:2506.01361* **2025**.
16. Csurgai-Horvath, L.; Bito, J. Attenuation Time Series Synthesis for Land Mobile Satellite Links. In Proceedings of the 2007 16th IST Mobile and Wireless Communications Summit. IEEE, 2007, pp. 1–5.
17. Denaxas, E.A.; Bandyopadhyay, R.; Patiño-Echeverri, D.; Pitsianis, N. SynTiSe: A modified multi-regime MCMC approach for generation of wind power synthetic time series. In Proceedings of the 2015 Annual IEEE Systems Conference (SysCon) Proceedings. IEEE, 2015, pp. 668–674.
18. Gkoulis, D. Creating interpretable synthetic time series for enhancing the design and implementation of Internet of Things (IoT) solutions. *Internet of Things* **2025**, p. 101500.
19. Dinku, T.; Hailemariam, K.; Maidment, R.; Tarnavsky, E.; Connor, S. Combined use of satellite estimates and rain gauge observations to generate high-quality historical rainfall time series over Ethiopia. *International Journal of Climatology* **2014**, *34*, 2489–2504.
20. Meiser, M.; Duppe, B.; Zinnikus, I. Syntised–synthetic time series data generator. In Proceedings of the 2023 11th Workshop on Modelling and Simulation of Cyber-Physical Energy Systems (MSCPES). IEEE, 2023, pp. 1–6.
21. Almutairi, A.A.; LeBlanc, D.J.; Kusari, A. Car-STAGE: Automated framework for large-scale high-dimensional simulated time-series data generation based on user-defined criteria. *arXiv preprint arXiv:2503.03100* **2025**.
22. Liu, X.; Aksu, T.; Liu, J.; Wen, Q.; Liang, Y.; Xiong, C.; Savarese, S.; Sahoo, D.; Li, J.; Liu, C. Empowering Time Series Analysis with Synthetic Data: A Survey and Outlook in the Era of Foundation Models, 2025, [arXiv:cs.LG/2503.11411].
23. Kong, Z.; Ping, W.; Huang, J.; Zhao, K.; Catanzaro, B. DiffWave: A Versatile Diffusion Model for Audio Synthesis. In Proceedings of the International Conference on Learning Representations, 2021.
24. Welker, S.; Richter, J.; Gerkmann, T. Speech Enhancement with Score-Based Generative Models in the Complex STFT Domain. *Interspeech* **2022**. https://doi.org/10.48550/arxiv.2203.17004.
25. Tai, W.; Zhang, F.; Trajcevski, G.; Zhong, T. Revisiting Denoising Diffusion Probabilistic Models for Speech Enhancement: Condition Collapse, Efficiency and Refinement. *AAAI Conference on Artificial Intelligence* **2023**. https://doi.org/10.1609/aaai.v37i11.26597.
26. Yang, Z.; Wu, J.; Wang, Z.; Wang, S.; Yuan, Y.; He, X. Generate What You Prefer: Reshaping Sequential Recommendation via Guided Diffusion. *Neural Information Processing Systems* **2023**. https://doi.org/10.48550/arxiv.2310.20453.
27. Li, Y. Ts-diffusion: Generating highly complex time series with diffusion models. *arXiv preprint arXiv:2311.03303* **2023**.
28. Chi, G.; Yang, Z.; Wu, C.; Xu, J.; Gao, Y.; Liu, Y.; Han, T.X. RF-Diffusion: Radio Signal Generation via Time-Frequency Diffusion. *ACM/IEEE International Conference on Mobile Computing and Networking* **2024**. https://doi.org/10.1145/3636534.3649348.
29. Li, H.; Huang, Y.; Xu, C.; Schlegel, V.; Jiang, R.; Batista-Navarro, R.; Nenadic, G.; Bian, J. BRIDGE: Bootstrapping text to control time-series generation via multi-agent iterative optimization and diffusion modelling. In Proceedings of the Forty-second International Conference on Machine Learning, 2025.
30. Huang, Y.H.; Xu, C.; Wu, Y.; Li, W.J.; Bian, J. Timedp: Learning to generate multi-domain time series with domain prompts. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2025, Vol. 39, pp. 17520–17527.
31. Sommers, A.; Cummins, L.; Mittal, S.; Rahimi, S.; Seale, M.; Jaboure, J.; Arnold, T. A Survey of Transformer Enabled Time Series Synthesis . In Proceedings of the 2024 IEEE 10th International Con-

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

ference on Collaboration and Internet Computing (CIC), Los Alamitos, CA, USA, 2024; pp. 60–69. https://doi.org/10.1109/CIC62241.2024.00018.

32. Yang, Y.; Jin, M.; Wen, H.; Zhang, C.; Liang, Y.; Ma, L.; Wang, Y.; Liu, C.; Yang, B.; Xu, Z.; et al. A Survey on Diffusion Models for Time Series and Spatio-Temporal Data. *ACM Comput. Surv.* **2025**. Just Accepted, https://doi.org/10.1145/3783986.
33. Hu, C.; Sun, Z.; Li, C.; Zhang, Y.; Xing, C. Survey of Time Series Data Generation in IoT. *Sensors* **2023**, *23*. https://doi.org/10.3390/s23156976.
34. Brophy, E.; Wang, Z.; She, Q.; Ward, T. Generative adversarial networks in time series: A survey and taxonomy. *arXiv preprint arXiv:2107.11098* **2021**.
35. Ang, Y.; Huang, Q.; Bao, Y.; Tung, A.K.H.; Huang, Z. TSGBench: Time Series Generation Benchmark. *Proceedings of the VLDB Endowment* **2023**. https://doi.org/10.14778/3632093.3632097.
36. Kingma, D.P.; Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* **2013**.
37. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative adversarial networks. *Communications of the ACM* **2020**, *63*, 139–144.
38. Ho, J.; Jain, A.; Abbeel, P. Denoising diffusion probabilistic models. *Advances in neural information processing systems* **2020**, *33*, 6840–6851.
39. Papamakarios, G.; Nalisnick, E.; Rezende, D.J.; Mohamed, S.; Lakshminarayanan, B. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research* **2021**, *22*, 1–64.
40. Chen, R.T.; Rubanova, Y.; Bettencourt, J.; Duvenaud, D.K. Neural ordinary differential equations. *Advances in neural information processing systems* **2018**, *31*.
41. Dooley, S.; Khurana, G.S.; Mohapatra, C.; Naidu, S.V.; White, C. Forecastpfn: Synthetically-trained zero-shot forecasting. *Advances in Neural Information Processing Systems* **2023**, *36*, 2403–2426.
42. Lan, T.; Le, H.D.; Li, J.; He, W.; Wang, M.; Liu, C.; Zhang, C. Towards Foundation Models for Zero-Shot Time Series Anomaly Detection: Leveraging Synthetic Data and Relative Context Discrepancy, 2025, [arXiv:cs.LG/2509.21190].
43. El-Laham, Y.; Vyetrenko, S. Styletime: Style transfer for synthetic time series generation. In Proceedings of the Proceedings of the Third ACM International Conference on AI in Finance, 2022, pp. 489–496.
44. Kegel, L.; Hahmann, M.; Lehner, W. Feature-driven Time Series Generation. In Proceedings of the Grundlagen von Datenbanken, 2017, pp. 54–59.
45. Xie, Z.; Li, Z.; He, X.; Xu, L.; Wen, X.; Zhang, T.; Chen, J.; Shi, R.; Pei, D. ChatTS: Aligning Time Series with LLMs via Synthetic Data for Enhanced Understanding and Reasoning. *arXiv preprint arXiv:2412.03104* **2024**.
46. Ansari, A.F.; Stella, L.; Turkmen, C.; Zhang, X.; Mercado, P.; Shen, H.; Shchur, O.; Rangapuram, S.S.; Arango, S.P.; Kapoor, S.; et al. Chronos: Learning the language of time series. *arXiv preprint arXiv:2403.07815* **2024**.
47. Xie, S.; Feofanov, V.; Alonso, M.; Odonnat, A.; Zhang, J.; Palpanas, T.; Redko, I. Cauker: Classification time series foundation models can be pretrained on synthetic data only. *arXiv preprint arXiv:2508.02879* **2025**.
48. Kang, Y.; Hyndman, R.J.; Li, F. GRATIS: GeneRAting TIme Series with diverse and controllable characteristics. *Statistical Analysis and Data Mining: The ASA Data Science Journal* **2020**, *13*, 354–376.
49. Chen, N.; Zhang, Y.; Zen, H.; Weiss, R.J.; Norouzi, M.; Chan, W. WaveGrad: Estimating Gradients for Waveform Generation. In Proceedings of the International Conference on Learning Representations, 2021.
50. Lu, Y.J.; Wang, Z.; Watanabe, S.; Richard, A.; Yu, C.; Tsao, Y. Conditional Diffusion Probabilistic Model for Speech Enhancement. *IEEE International Conference on Acoustics, Speech, and Signal Processing* **2022**. https://doi.org/10.1109/icassp43922.2022.9746901.
51. Tai, W.; Lei, Y.; Zhou, F.; Trajcevski, G.; Zhong, T. DOSE: Diffusion Dropout with Adaptive Prior for Speech Enhancement. *Neural Information Processing Systems* **2023**. https://doi.org/null.
52. Qiu, Z.; Guo, Y.; Fu, M.; Huang, H.; Hu, Y.; He, L.; Sun, F. Cra-diffuse: Improved cross-domain speech enhancement based on diffusion model with tf domain pre-denoising. In Proceedings of the 2023 IEEE International Conference on Multimedia and Expo (ICME). IEEE, 2023, pp. 1709–1714.
53. Liu, Q.; Fan, Y.B.; Zhao, X.; Du, Z.Z.; Guo, H.; Tang, R.; Tian, F. Diffusion Augmentation for Sequential Recommendation. *International Conference on Information and Knowledge Management* **2023**. https://doi.org/10.1145/3583780.3615134.
54. Dai, Y.; Yang, C.; Liu, K.; Liu, A.; Liu, Y. TimeDDPM: Time series augmentation strategy for industrial soft sensing. *IEEE Sensors Journal* **2023**, *24*, 2145–2153.
55. Zhu, Y.; Ye, Y.; Zhang, S.; Zhao, X.; Yu, J.J.Q. DiffTraj: Generating GPS Trajectory with Diffusion Probabilistic Model. *Neural Information Processing Systems* **2023**. https://doi.org/null.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

56. Li, S.; Xiong, H.; Chen, Y. DiffCharge: Generating EV Charging Scenarios via a Denoising Diffusion Model. *IEEE Transactions on Smart Grid* **2023**. https://doi.org/10.1109/tsg.2024.3360874.
57. Dong, X.; Mao, Z.; Sun, Y.; Xu, X. Short-Term Wind Power Scenario Generation Based on Conditional Latent Diffusion Models. *IEEE Transactions on Sustainable Energy* **2023**. https://doi.org/10.1109/tste.2023.3327497.
58. Fu, C.; Kazmi, H.; Quintana, M.; Miller, C. Creating synthetic energy meter data using conditional diffusion and building metadata. *Energy and Buildings* **2024**. https://doi.org/10.1016/j.enbuild.2024.114216.
59. Klein, G.; Guetschel, P.; Silvestri, G.; Tangermann, M. Synthesizing EEG Signals from Event-Related Potential Paradigms with Conditional Diffusion Models. *arXiv.org* **2024**. https://doi.org/10.3217/978-3-99161-014-4-077.
60. Huang, H.; Chen, M.; Qiao, X. Generative Learning for Financial Time Series with Irregular and Scale-Invariant Patterns. *International Conference on Learning Representations* **2024**. https://doi.org/null.
61. Galib, A.; Tan, P.N.; Luo, L. FIDE: Frequency-Inflated Conditional Diffusion Model for Extreme-Aware Time Series Generation. *Neural Information Processing Systems* **2024**. https://doi.org/null.
62. Yuan, X.; Qiao, Y. Diffusion-TS: Interpretable Diffusion for General Time Series Generation. *International Conference on Learning Representations* **2024**. https://doi.org/10.48550/arxiv.2403.01742.
63. Wei, T.; Lin, Y.; Guo, S.; Lin, Y.; Huang, Y.; Xiang, C.; Bai, Y.; Ya, M.; Wan, H. Diff-RNTraj: A Structure-Aware Diffusion Model for Road Network-Constrained Trajectory Generation. *IEEE Transactions on Knowledge and Data Engineering* **2024**. https://doi.org/10.1109/tkde.2024.3460051.
64. Zhu, Y.; Yu, J.J.Q.; Zhao, X.; Liu, Q.; Ye, Y.; Chen, W.; Zhang, Z.; Wei, X.; Liang, Y. ControlTraj: Controllable Trajectory Generation with Topology-Constrained Diffusion Model. *Knowledge Discovery and Data Mining* **2024**. https://doi.org/10.1145/3637528.3671866.
65. Ma, H.; Xie, R.; Meng, L.; Chen, X.; Zhang, X.; Lin, L.; Kang, Z. Plug-in Diffusion Model for Sequential Recommendation. *AAAI Conference on Artificial Intelligence* **2024**. https://doi.org/10.48550/arxiv.2401.02913.
66. Yi, H.; Hou, L.; Jin, Y.; Saeed, N.A.; Kandil, A.; Duan, H. Time series diffusion method: A denoising diffusion probabilistic model for vibration signal generation. *Mechanical Systems and Signal Processing* **2024**, *216*, 111481.
67. Gao, Y.; Chen, H.; Wang, X.; Wang, Z.; Wang, X.; Gao, J.; Ding, B. DiffsFormer: A Diffusion Transformer on Stock Factor Augmentation. *arXiv.org* **2024**. https://doi.org/10.48550/arxiv.2402.06656.
68. Wang, Y.; Liu, Z.; Yang, L.; Yu, P.S. Conditional denoising diffusion for sequential recommendation. In Proceedings of the Pacific-Asia Conference on Knowledge Discovery and Data Mining. Springer, 2024, pp. 156–169.
69. Wang, Z.; Zhang, H. Customized load profiles synthesis for electricity customers based on conditional diffusion models. *IEEE Transactions on Smart Grid* **2024**, *15*, 4259–4270.
70. Deng, B.; Xu, C.; Li, H.; Huang, Y.; Hou, M.; Bian, J. TarDiff: Target-Oriented Diffusion Guidance for Synthetic Electronic Health Record Time Series Generation. *arXiv preprint arXiv:2504.17613* **2025**.
71. Li, Y.; Luyten, M.R.; Schaar, M. Risk-Sensitive Diffusion: Robustly Optimizing Diffusion Models with Noisy Samples. *International Conference on Learning Representations* **2025**. https://doi.org/null.
72. Li, Y.; Meng, H.; Bi, Z.; Urnes, I.T.; Chen, H. Population Aware Diffusion for Time Series Generation. *arXiv preprint arXiv:2501.00910* **2025**.
73. Shankar, A.; Chen, L.Y.; van Deursen, A.; Hai, R. WaveStitch: Flexible and Fast Conditional Time Series Generation with Diffusion Models. *arXiv preprint arXiv:2503.06231* **2025**.
74. Ge, Y.; Li, J.; Zhao, Y.; Wen, H.; Li, Z.; Qiu, M.; Li, H.; Jin, M.; Pan, S. T2s: High-resolution time series generation with text-to-series diffusion models. *arXiv preprint arXiv:2505.02417* **2025**.
75. Xia, Y.; Xu, C.; Liang, Y.; Wen, Q.; Zimmermann, R.; Bian, J. Causal Time Series Generation via Diffusion Models. *arXiv preprint arXiv:2509.20846* **2025**.
76. Deng, J.; Xu, C.; Ma, J.W.; Jin, M.; Liu, C.; Bian, J. OATS: Online Data Augmentation for Time Series Foundation Models. *arXiv preprint arXiv:2601.19040* **2026**.
77. Zhang, X.; Deng, J.; Xu, C.; Li, H.; Bian, J. Diff-MN: Diffusion Parameterized MoE-NCDE for Continuous Time Series Generation with Irregular Observations, 2026, [arXiv:cs.LG/2601.13534].
78. Mogren, O.; Mogren, O.; Mogren, O.; Mogren, O. C-RNN-GAN: Continuous recurrent neural networks with adversarial training. *arXiv: Artificial Intelligence* **2016**. https://doi.org/null.
79. Esteban, C.; Hyland, S.L.; Rätsch, G. Real-valued (medical) time series generation with recurrent conditional gans. *arXiv preprint arXiv:1706.02633* **2017**.
80. Fu, R.; Chen, J.; Chen, J.; Chen, J.; Zeng, S.; Zhuang, Y.; Sudjianto, A. Time Series Simulation by Conditional Generative Adversarial Net. *arXiv: Machine Learning* **2019**. https://doi.org/10.2139/ssrn.3373730.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

81. Dash, S.; Dash, S.; Yale, A.; Yale, A.; Guyon, I.; Guyon, I.; Guyon, I.; Bennett, K.P.; Bennett, K.P.; Bennett, K.P.; et al. Medical Time-Series Data Generation Using Generative Adversarial Networks. *Conference on Artificial Intelligence in Medicine in Europe* **2020**. https://doi.org/10.1007/978-3-030-59137-3_34.
82. Xu, T.; Xu, T.; Wenliang, L.K.; Wenliang, L.K.; Munn, M.; Munn, M.; Munn, M.; Acciaio, B.; Acciaio, B.; Acciaio, B. COT-GAN: Generating Sequential Data via Causal Optimal Transport. *Neural Information Processing Systems* **2020**. https://doi.org/null.
83. Lin, Z.; Lin, Z.; Jain, A.; Jain, A.; Wang, C.; Wang, C.; Kairouz, P.; Fanti, G.; Fanti, G.; Sekar, V.; et al. Using GANs for Sharing Networked Time Series Data: Challenges, Initial Promise, and Open Questions. *arXiv: Learning* **2020**. https://doi.org/10.1145/3419394.3423643.
84. Smith, K.E.; Smith, K.E.; Smith, A.O.; Smith, A.O.; Smith, A.O.; Smith, A.O. Conditional GAN for timeseries generation. *arXiv.org* **2020**. https://doi.org/null.
85. Zhang, X.; Qin, Y.; Yuen, C.; Jayasinghe, L.; Liu, X. Time-series regeneration with convolutional recurrent generative adversarial network for remaining useful life estimation. *IEEE Transactions on Industrial Informatics* **2020**, *17*, 6820–6831.
86. Pinceti, A.; Sankar, L.; Kosut, O. Synthetic Time-Series Load Data via Conditional Generative Adversarial Networks. *IEEE Power & Energy Society General Meeting* **2021**. https://doi.org/10.1109/pesgm46819.2021.9637821.
87. Pei, H.; Pei, H.; Ren, K.; Ren, K.; Yang, Y.; Yang, Y.; Liu, C.; Chang, L.; Qin, T.; Qin, T.; et al. Towards Generating Real-World Time Series Data. *Industrial Conference on Data Mining* **2021**. https://doi.org/10.1109/icdm51629.2021.00058.
88. Kidger, P.; Foster, J.; Li, X.; Oberhauser, H.; Lyons, T. Neural SDEs as Infinite-Dimensional GANs. *International Conference on Machine Learning* **2021**. https://doi.org/null.
89. Leznik, M.; Leznik, M.; Michalsky, P.; Michalsky, P.; Willis, P.; Willis, P.; Schanzel, B.; Schanzel, B.; Östberg, P.O.; Östberg, P.O.; et al. Multivariate Time Series Synthesis Using Generative Adversarial Networks. *International Conference on Performance Engineering* **2021**. https://doi.org/10.1145/3427921.3450257.
90. Sun, B.; Wu, Z.; Feng, Q.; Wang, Z.; Ren, Y.; Yang, D.; Xia, Q. Small sample reliability assessment with online time-series data based on a worm Wasserstein generative adversarial network learning method. *IEEE Transactions on Industrial Informatics* **2022**, *19*, 1207–1216.
91. Srinivasan, P.; Knottenbelt, W. Time-series Transformer Generative Adversarial Networks. *arXiv.org* **2022**. https://doi.org/10.48550/arxiv.2205.11164.
92. Jeha, P.; Bohlke-Schneider, M.; Mercado, P.; Nirwan, R.S.; Kapoor, S.; Flunkert, V.; Gasthaus, J.; Januschowski, T. PSA-GAN: Progressive Self Attention GANs for Synthetic Time Series. *International Conference on Learning Representations* **2022**. https://doi.org/null.
93. Seyfi, A. Generating multivariate time series with COmmon Source CoordInated GAN (COSCI-GAN). *Neural Information Processing Systems* **2022**. https://doi.org/null.
94. Jeon, J.; Jeon, J.; Kim, J.; Kim, J.; Song, H.; Song, H.W.; Cho, S.; Cho, S.; Park, N.; Park, N. GT-GAN: General Purpose Time Series Synthesis with Generative Adversarial Networks. *Neural Information Processing Systems* **2022**. https://doi.org/10.48550/arxiv.2210.02040.
95. Remlinger, C.; Mikael, J.; Élie, R. Conditional Loss and Deep Euler Scheme for Time Series Generation. *AAAI Conference on Artificial Intelligence* **2022**. https://doi.org/10.1609/aaai.v36i7.20782.
96. Lee, H.K.; Lee, H.K.; Lee, J.; Lee, J.Y.; Kim, S.B.; Kim, S.B. Boundary-Focused Generative Adversarial Networks for Imbalanced and Multimodal Time Series. *IEEE Transactions on Knowledge and Data Engineering* **2022**. https://doi.org/10.1109/tkde.2022.3182327.
97. Lu, C.; Reddy, C.K.; Wang, P.; Nie, D.; Ning, Y. Multi-Label Clinical Time-Series Generation via Conditional GAN. *IEEE Transactions on Knowledge and Data Engineering* **2023**. https://doi.org/10.1109/tkde.2023.3310909.
98. Wang, L.; Zeng, L.; Li, J. AEC-GAN: Adversarial Error Correction GANs for Auto-Regressive Long Time-Series Generation. *AAAI Conference on Artificial Intelligence* **2023**. https://doi.org/10.1609/aaai.v37i8.26208.
99. Klopries, H.; Schwung, A. ITF-GAN: Synthetic time series dataset generation and manipulation by interpretable features. *Knowledge-Based Systems* **2024**. https://doi.org/10.1016/j.knosys.2023.111131.
100. Chatterjee, S.; Hazra, D.; Byun, Y.C. GAN-based synthetic time-series data generation for improving prediction of demand for electric vehicles. *Expert Systems with Applications* **2025**, *264*, 125838.
101. Zhang, G.; Sikdar, B. Synthetic Time-Series Data Generation for Smart Grids Using 3D Autoencoder GAN. *IEEE Transactions on Industrial Informatics* **2025**.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

102. Zhang, Z.; Gao, H.; Sun, W.; Song, W.; Li, Q. Multivariate time series generation based on dual-channel Transformer conditional GAN for industrial remaining useful life prediction. *Knowledge-Based Systems* **2025**. https://doi.org/10.1016/j.knosys.2024.112749.
103. Che, Z.; Che, Z.; Purushotham, S.; Purushotham, S.; Li, G.; Li, G.; Li, G.; Jiang, B.; Jiang, B.; Liu, Y.; et al. Hierarchical Deep Generative Models for Multi-Rate Multivariate Time Series. *International Conference on Machine Learning* **2018**. https://doi.org/null.
104. Tang, B.; Matteson, D.S. Probabilistic Transformer For Time Series Analysis. *Neural Information Processing Systems* **2021**. https://doi.org/null.
105. Desai, A.; Freeman, C.; Wang, Z.; Beaver, I. Timevae: A variational auto-encoder for multivariate time series generation. *arXiv preprint arXiv:2111.08095* **2021**.
106. Zhou, L.; Poli, M.; Xu, W.; Massaroli, S.; Ermon, S. Deep Latent State Space Models for Time-Series Generation. *International Conference on Machine Learning* **2022**. https://doi.org/10.48550/arxiv.2212.12749.
107. Lee, D.; Malacarne, S.; Aune, E. Vector Quantized Time Series Generation with a Bidirectional Prior Model. In Proceedings of the International Conference on Artificial Intelligence and Statistics. PMLR, 2023, pp. 7665–7693.
108. Li, H.; Yu, S.; Príncipe, J.C. Causal Recurrent Variational Autoencoder for Medical Time Series Generation. *AAAI Conference on Artificial Intelligence* **2023**. https://doi.org/10.48550/arxiv.2301.06574.
109. Naiman, I.; Erichson, N.B.; Ren, P.; Mahoney, M.W.; Azencot, O. Generative Modeling of Regular and Irregular Time Series Data via Koopman VAEs. *International Conference on Learning Representations* **2023**. https://doi.org/10.48550/arxiv.2310.02619.
110. EskandariNasab, M.; Muhammad Hamdi, S.; Filali Boubrahimi, S. AVATAR: Adversarial Autoencoders with Autoregressive Refinement for Time Series Generation. In Proceedings of the Proceedings of the 2025 SIAM International Conference on Data Mining (SDM). SIAM, 2025, pp. 548–557.
111. Lee, D.; Malacarne, S.; Aune, E. Closing the Gap Between Synthetic and Ground Truth Time Series Distributions via Neural Mapping. *arXiv preprint arXiv:2501.17553* **2025**.
112. Chen, R.T.Q.; Chen, R.T.Q.; Chen, T.Q.; Rubanova, Y.; Rubanova, Y.; Bettencourt, J.; Bettencourt, J.; Duvenaud, D.; Duvenaud, D. Neural ordinary differential equations. *Neural Information Processing Systems* **2018**. https://doi.org/null.
113. Deng, R.; Chang, B.; Brubaker, M.A.; Mori, G.; Lehrmann, A.M. Modeling Continuous Stochastic Processes with Dynamic Normalizing Flows. *Neural Information Processing Systems* **2020**. https://doi.org/null.
114. Alaa, A.; Chan, A.J.; Schaar, M. Generative Time-series Modeling with Fourier Flows. *International Conference on Learning Representations* **2021**. https://doi.org/null.
115. Cheng, Y.; Wang, Z.; Xiao, T.; Zhong, Q.; Suo, J.; He, K. CausalTime: Realistically Generated Time-series for Benchmarking of Causal Discovery. *International Conference on Learning Representations* **2023**. https://doi.org/10.48550/arxiv.2310.01753.
116. Aloni, O.; Perelman, G.; Fishbain, B. Synthetic random environmental time series generation with similarity control, preserving original signal's statistical characteristics. *Environmental Modelling & Software* **2025**, *185*, 106283.
117. Alouadi, A.; Barreau, B.; Carlier, L.; Pham, H. Robust time series generation via Schr\" odinger Bridge: A comprehensive evaluation. *arXiv preprint arXiv:2503.02943* **2025**.
118. Fuest, M.; Cuesta, A.; Veeramachaneni, K. CENTS: Generating synthetic electricity consumption time series for rare and unseen scenarios. *arXiv preprint arXiv:2501.14426* **2025**.
119. Dosovitskiy, A.; Ros, G.; Codevilla, F.; Lopez, A.; Koltun, V. CARLA: An open urban driving simulator. In Proceedings of the Conference on robot learning. PMLR, 2017, pp. 1–16.
120. Chen, Z.; Feng, S.; Zhang, Z.; Xiao, X.; Gao, X.; Zhao, P. SDformer: Similarity-driven discrete transformer for time series generation. In Proceedings of the Proceedings of the 38th International Conference on Neural Information Processing Systems, Red Hook, NY, USA, 2024; NIPS '24.
121. Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; Courville, A.C. Improved training of wasserstein gans. *Advances in neural information processing systems* **2017**, *30*.
122. Springenberg, J.T. Unsupervised and semi-supervised learning with categorical generative adversarial networks. *arXiv preprint arXiv:1511.06390* **2015**.
123. Yoon, J.; Yoon, J.; Jarrett, D.; Jarrett, D.; Jarrett, D.; van der Schaar, M.; van der Schaar, M. Time-series Generative Adversarial Networks. *Neural Information Processing Systems* **2019**. https://doi.org/null.
124. Cuturi, M. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems* **2013**, *26*.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

125. Jarrett, D.; Bica, I.; van der Schaar, M. Time-series Generation by Contrastive Imitation. *Neural Information Processing Systems* **2023**. https://doi.org/10.48550/arxiv.2311.01388.
126. Mao, X.; Li, Q.; Xie, H.; Lau, R.Y.; Wang, Z.; Paul Smolley, S. Least squares generative adversarial networks. In Proceedings of the Proceedings of the IEEE international conference on computer vision, 2017, pp. 2794–2802.
127. Liao, S.; Ni, H.; Sabate-Vidales, M.; Szpruch, L.; Wiese, M.; Xiao, B. Sig-Wasserstein GANs for conditional time series generation. *Mathematical Finance* **2024**, *34*, 622–670.
128. Ni, H.; Szpruch, L.; Sabate-Vidales, M.; Xiao, B.; Wiese, M.; Liao, S. Sig-Wasserstein GANs for time series generation. In Proceedings of the Proceedings of the Second ACM International Conference on AI in Finance, 2021, pp. 1–8.
129. Liu, Y.; Wijewickrema, S.; Li, A.; Bester, C.; O'Leary, S.; Bailey, J. Time-transformer: Integrating local and global features for better time series generation. In Proceedings of the Proceedings of the 2024 SIAM International Conference on Data Mining (SDM). SIAM, 2024, pp. 325–333.
130. Coletta, A.; Gopalakrishan, S.; Borrajo, D.; Vyetrenko, S. On the Constrained Time-Series Generation Problem. *Neural Information Processing Systems* **2023**. https://doi.org/10.48550/arxiv.2307.01717.
131. Evans, Z.; Parker, J.; Carr, C.; Zukowski, Z.; Taylor, J.; Pons, J. Long-form music generation with latent diffusion. *International Society for Music Information Retrieval Conference* **2024**. https://doi.org/10.48550/arxiv.2404.10301.
132. Elizalde, B.; Deshmukh, S.; Al Ismail, M.; Wang, H. Clap learning audio concepts from natural language supervision. In Proceedings of the ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2023, pp. 1–5.
133. Woo, J.; Lee, J.; Iwana, B.K. Towards Time Series Generation Conditioned on Unstructured Natural Language. *arXiv preprint arXiv:2506.22927* **2025**.
134. Gu, S.; Li, C.; Jing, B.; Ren, K. VerbalTS: Generating Time Series from Texts. In Proceedings of the Forty-second International Conference on Machine Learning.
135. Peebles, W.; Xie, S. Scalable diffusion models with transformers. In Proceedings of the Proceedings of the IEEE/CVF international conference on computer vision, 2023, pp. 4195–4205.
136. Garuti, F.; Sangineto, E.; Luetto, S.; Forni, L.; Cucchiara, R. Diffusion Transformers for Tabular Data Time Series Generation. *International Conference on Learning Representations* **2025**. https://doi.org/10.48550/arxiv.2504.07566.
137. Naiman, I.; Berman, N.; Pemper, I.; Arbiv, I.; Fadlon, G.; Azencot, O. Utilizing Image Transforms and Diffusion Models for Generative Modeling of Short and Long Time Series. *Neural Information Processing Systems* **2024**. https://doi.org/10.48550/arxiv.2410.19538.
138. Park, J.; Lee, S.; Jeong, W.; Choi, Y.; Lee, J. Leveraging priors via diffusion bridge for time series generation. *arXiv preprint arXiv:2408.06672* **2024**.
139. Shen, L.; Li, X.; Long, L. TSGDiff: Rethinking Synthetic Time Series Generation from a Pure Graph Perspective. *arXiv preprint arXiv:2511.12174* **2025**.
140. Chen, H.; Lundberg, S.M.; Lee, S.I. Explaining a series of models by propagating Shapley values. *Nature communications* **2022**, *13*, 4512.
141. Lundberg, S.M.; Lee, S.I. A unified approach to interpreting model predictions. *Advances in neural information processing systems* **2017**, *30*.
142. Ericson, L.; Zhu, X.; Han, X.; Fu, R.; Li, S.; Guo, S.; Hu, P. Deep generative modeling for financial time series with application in VaR: A comparative review. *arXiv preprint arXiv:2401.10370* **2024**.
143. Hamdouche, M.; Henry-Labordere, P.; Pham, H. Generative Modeling for Time Series Via Schrödinger Bridge. *Available at SSRN 4412434* **2023**.
144. Lu, C.I.; Sester, J. Generative model for financial time series trained with MMD using a signature kernel. *arXiv preprint arXiv:2407.19848* **2024**.
145. Dogariu, M.; Ştefan, L.D.; Boteanu, B.A.; Lamba, C.; Ionescu, B. Towards realistic financial time series generation via generative adversarial learning. In Proceedings of the 2021 29th European Signal Processing Conference (EUSIPCO). IEEE, 2021, pp. 1341–1345.
146. Wiese, M.; Knobloch, R.; Korn, R.; Kretschmer, P. Quant GANs: Deep generation of financial time series. *Quantitative Finance* **2020**, *20*, 1419–1440.
147. Lu, C.; Reddy, C.K.; Wang, P.; Nie, D.; Ning, Y. Multi-label clinical time-series generation via conditional GAN. *IEEE Transactions on Knowledge and Data Engineering* **2023**, *36*, 1728–1740.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

148. Kuo, N.I.; Jorm, L.; Barbieri, S.; et al. Synthetic health-related longitudinal data with mixed-type variables generated using diffusion models. *arXiv preprint arXiv:2303.12281* **2023**.
149. Schürch, M.; Li, X.; Allam, A.; Hofer, G.; Mollaysa, A.; Cavelti-Weder, C.; Krauthammer, M. Generating personalized insulin treatments strategies with conditional generative time series models. In Proceedings of the Deep Generative Models for Health Workshop NeurIPS 2023, 2023.
150. Zhu, T.; Li, K.; Herrero, P.; Georgiou, P. Glugan: Generating personalized glucose time series using generative adversarial networks. *IEEE Journal of Biomedical and Health Informatics* **2023**, *27*, 5122–5133.
151. Zhang, Z.; Gao, H.; Sun, W.; Song, W.; Li, Q. Multivariate time series generation based on dual-channel Transformer conditional GAN for industrial remaining useful life prediction. *Knowledge-Based Systems* **2025**, *308*, 112749.
152. Ma, Z.; Sun, Y.; Ji, H.; Li, S.; Nie, S.; Yin, F. A CNN-BiLSTM-Attention approach for EHA degradation prediction based on time-series generative adversarial network. *Mechanical Systems and Signal Processing* **2024**, *215*, 111443.
153. Shangguan, A.; Xie, G.; Fei, R.; Mu, L.; Hei, X. Train wheel degradation generation and prediction based on the time series generation adversarial network. *Reliability Engineering & System Safety* **2023**, *229*, 108816.
154. Denaxas, E.A.; Denaxas, E.A.; Bandyopadhyay, R.; Bandyopadhyay, R.; Patiño-Echeverri, D.; Patiño-Echeverri, D.; Pitsianis, N.P.; Pitsianis, N.P. SynTiSe: A modified multi-regime MCMC approach for generation of wind power synthetic time series. *2015 Annual IEEE Systems Conference (SysCon) Proceedings* **2015**. https://doi.org/10.1109/syscon.2015.7116827.
155. Ye, L.; Peng, Y.; Li, Y.; Li, Z. A novel informer-time-series generative adversarial networks for day-ahead scenario generation of wind power. *Applied Energy* **2024**, *364*, 123182.
156. Li, Y.; Hu, B.; Niu, T.; Gao, S.; Yan, J.; Xie, K.; Ren, Z. GMM-HMM-based medium-and long-term multi-wind farm correlated power output time series generation method. *IEEE Access* **2021**, *9*, 90255–90267.
157. Talbot, P.W.; Rabiti, C.; Alfonsi, A.; Krome, C.; Kunz, M.R.; Epiney, A.; Wang, C.; Mandelli, D. Correlated synthetic time series generation for energy system simulations using Fourier and ARMA signal processing. *International Journal of Energy Research* **2020**, *44*, 8144–8155.
158. Brentan, B.M.; Meirelles, G.L.; Manzi, D.; Luvizotto, E. Water demand time series generation for distribution network modeling and water demand forecasting. *Urban Water Journal* **2018**, *15*, 150–158.
159. Tabeshpour, M.R.; Belvasi, N. Ocean waves time-series generation: Minimum required artificial wave time-series for wave energy converter analysis. *Journal of Marine Engineering & Technology* **2023**, *22*, 273–283.
160. Zhou, H.; Zhang, S.; Peng, J.; Zhang, S.; Li, J.; Xiong, H.; Zhang, W. Informer: Beyond efficient transformer for long sequence time-series forecasting. In Proceedings of the Proceedings of the AAAI conference on artificial intelligence, 2021, Vol. 35, pp. 11106–11115.
161. Stenger, M.; Leppich, R.; Foster, I.; Kounev, S.; Bauer, A. Evaluation is key: A survey on evaluation measures for synthetic time series. *Journal of Big Data* **2024**, *11*, 66.
162. Stenger, M.; Leppich, R.; Bauer, A.; Kounev, S. STEB: In Search of the Best Evaluation Approach for Synthetic Time Series. *arXiv preprint arXiv:2505.21160* **2025**.
163. Alaa, A.; Van Breugel, B.; Saveliev, E.S.; Van Der Schaar, M. How faithful is your synthetic data? sample-level metrics for evaluating and auditing generative models. In Proceedings of the International conference on machine learning. PMLR, 2022, pp. 290–306.
164. Leznik, M.; Lochner, A.; Wesner, S.; Domaschka, J. [SoK] The great GAN bake Off, an extensive systematic evaluation of generative adversarial network architectures for time series synthesis. *Journal of Systems Research* **2022**, *2*.
165. Leznik, M.; Michalsky, P.; Willis, P.; Schanzel, B.; Östberg, P.O.; Domaschka, J. Multivariate time series synthesis using generative adversarial networks. In Proceedings of the Proceedings of the ACM/SPEC International Conference on Performance Engineering, 2021, pp. 43–50.
166. Jeha, P.; Bohlke-Schneider, M.; Mercado, P.; Kapoor, S.; Nirwan, R.S.; Flunkert, V.; Gasthaus, J.; Januschowski, T. PSA-GAN: Progressive self attention GANs for synthetic time series. In Proceedings of the The tenth international conference on learning representations, 2022.
167. Gretton, A.; Borgwardt, K.M.; Rasch, M.J.; Schölkopf, B.; Smola, A. A kernel two-sample test. *The journal of machine learning research* **2012**, *13*, 723–773.
168. Kulkarni, V.; Tagasovska, N.; Vatter, T.; Garbinato, B. Generative models for simulating mobility trajectories. *arXiv preprint arXiv:1811.12801* **2018**.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026 doi:10.20944/preprints202605.0725.v1

169. Lai, G.; Chang, W.C.; Yang, Y.; Liu, H. Modeling long-and short-term temporal patterns with deep neural networks. In Proceedings of the The 41st international ACM SIGIR conference on research & development in information retrieval, 2018, pp. 95–104.
170. Parzen, E. On spectral analysis with missing observations and amplitude modulation. *Sankhyā: The Indian Journal of Statistics, Series A* **1963**, pp. 383–392.
171. Berndt, D.J.; Clifford, J. Using dynamic time warping to find patterns in time series. In Proceedings of the Proceedings of the 3rd international conference on knowledge discovery and data mining, 1994, pp. 359–370.
172. Shokoohi-Yekta, M.; Hu, B.; Jin, H.; Wang, J.; Keogh, E. Generalizing DTW to the multi-dimensional case requires an adaptive approach. *Data mining and knowledge discovery* **2017**, *31*, 1–31.
173. Van der Maaten, L.; Hinton, G. Visualizing data using t-SNE. *Journal of machine learning research* **2008**, *9*.
174. Wang, L.; Zhang, W.; He, X. Continuous patient-centric sequence generation via sequentially coupled adversarial learning. In Proceedings of the International Conference on Database Systems for Advanced Applications. Springer, 2019, pp. 36–52.
175. Bindschaedler, V.; Shokri, R.; Gunter, C.A. Plausible Deniability for Privacy-Preserving Data Synthesis. *Proceedings of the VLDB Endowment* **2017**, *10*.
176. Boursin, N.; Remlinger, C.; Mikael, J. Deep generators on commodity markets application to deep hedging. *Risks* **2022**, *11*, 7.
177. Lin, Z.; Jain, A.; Wang, C.; Fanti, G.; Sekar, V. Generating high-fidelity, synthetic time series datasets with doppelganger. *arXiv preprint arXiv:1909.13403* **2019**.
178. Norgaard, S.; Saeedi, R.; Sasani, K.; Gebremedhin, A.H. Synthetic sensor data generation for health applications: A supervised deep learning approach. In Proceedings of the 2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). IEEE, 2018, pp. 1164–1167.
179. Kynkäänniemi, T.; Karras, T.; Laine, S.; Lehtinen, J.; Aila, T. Improved precision and recall metric for assessing generative models. *Advances in neural information processing systems* **2019**, *32*.
180. Richardson, E.; Weiss, Y. On gans and gmms. *Advances in neural information processing systems* **2018**, *31*.
181. Shokri, R.; Stronati, M.; Song, C.; Shmatikov, V. Membership inference attacks against machine learning models. In Proceedings of the 2017 IEEE symposium on security and privacy (SP). IEEE, 2017, pp. 3–18.
182. Meehan, C.; Chaudhuri, K.; Dasgupta, S. A non-parametric test to detect data-copying in generative models. In Proceedings of the International conference on artificial intelligence and statistics, 2020.
183. Rubanova, Y.; Rubanova, Y.; Chen, R.T.Q.; Chen, R.T.Q.; Chen, T.Q.; Duvenaud, D.; Duvenaud, D. Latent Ordinary Differential Equations for Irregularly-Sampled Time Series. *Neural Information Processing Systems* **2019**. https://doi.org/null.
184. Rousseau, C.; Boschi, T.; Cornacchia, G.; Salwala, D.; Pascale, A.; Moreno, J.B. Forging Time Series with Language: A Large Language Model Approach to Synthetic Data Generation. In Proceedings of the The Thirty-ninth Annual Conference on Neural Information Processing Systems, 2025.
185. Ramsay, J.O.; Silverman, B.W. *Functional data analysis*; Springer, 2005.
186. Hyvarinen, A. Fast ICA for noisy data using Gaussian moments. In Proceedings of the 1999 IEEE international symposium on circuits and systems (ISCAS). IEEE, 1999, Vol. 5, pp. 57–61.
187. Zhou, X.; Jia, Q.; Hu, Y.; Xie, R.; Huang, T.; Yu, F.R. Geng: An llm-based generic time series data generation approach for edge intelligence via cross-domain collaboration. In Proceedings of the IEEE INFOCOM 2024-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), 2024, pp. 1–6.
188. Ravuru, C.; Sakhinana, S.S.; Runkana, V. Agentic retrieval-augmented generation for time series analysis. *arXiv preprint arXiv:2408.14484* **2024**.

Short Biography of Authors

Xu Zhang received the M.S. degree in Mechanical Engineering from the Dalian University of Technology, Liaoning, China, in 2022. He is currently pursuing the Ph.D. degree at the Fudan University, Shanghai, China. His research interests include data mining, deep learning, time-series analysis, foundation models, and post-training of GPT, as well as the development and application of intelligent agents. His academic contributions include multiple first-author publications at top-tier conferences such as ACM MM, AAAI, SIGKDD, WWW, ICLR, and ICML.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026
doi:10.20944/preprints202605.0725.v1

Chang Xu received the B.S. and Ph.D. degrees in computer science from Nankai University in 2014 and 2019, respectively. She is currently a Senior Researcher at Microsoft Research Asia, which she joined after her Ph.D. Her research interests include fundamental machine learning research, time series analysis, and generative modeling, as well as their applications in industry, including healthcare and finance. She has published dozens of papers with over a thousand citations in leading journals and conferences, and has served as a program committee member and peer reviewer for top-tier venues.

Hao Li received the Ph.D. degree from the Department of Computer Science, The University of Manchester, UK, in 2025. He is currently a Postdoctoral Researcher with Imperial College London, U.K. His research interests include Foundation Models, Reinforcement Learning and AI for Healthcare. His academic contributions include publications in top-tier conferences and journals such as NeurIPS, ICML, KDD, ACL, EMNLP, and NAACL.

Yuhao Huang received the B.S. degree from Nanjing University in 2021. He is currently a Ph.D candidate with the School of Computer Science, Nanjing University. His research interests include generative time series modeling, reinforcement learning, foundation models, and their applications in finance. He has published papers in top-tier international conferences such as NeurIPS, ICLR, ICML, AAAI and KDD.

Qiushui Xu received the B.S. degree from Lanzhou University in 2020. She is currently a Ph.D candidate at the Department of Industrial and Manufacturing Engineering, Pennsylvania State University. Her research interests include agentic systems, large language models, reinforcement learning, and convex optimization.

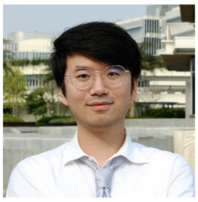

Yuxuan Liang is an Assistant Professor at The Hong Kong University of Science and Technology (Guangzhou), working on the research, development, and innovation of spatio-temporal data mining and urban computing. He obtained his PhD degree from School of Computing, National University of Singapore. He has published over 100 papers in refereed journals and conferences with 11,000 citations, and was recognized as Stanford/Elsevier Top 2% Scientists in 2024/25. He serves as Associate Editor of Neurocomputing and Area Chair for prestigious conferences such as KDD, ICML, NeurIPS, and ICLR. He serves as Tutorial Co-Chair for SSTD'25 and Workshop Co-Chair for PAKDD'26. He also regularly serves as Program Chair of the Workshop on Urban Computing (Urbcomp), Workshop on AI for Time Series (AI4TS), and Workshop on Mining and Learning from Time Series (MiLeTS) at KDD, WWW, IJCAI, etc. He received ICDM Best Paper Runner-up Award in 2025, ACM SIGSPATIAL China Chapter Rising Star Award in 2024, the 23rd China Patent Excellence Award in 2021, and SDSC Dissertation Research Fellowship in 2020.

Chenghao Liu is a Staff Research Scientist at Datadog AI Research, where he focuses on time series research and AI for observability. Previously, he was a Lead Applied Scientist at Salesforce AI Research, working on time series foundation models. He received B.S. and PH.D. degrees in Computer Science and Technology from Zhejiang University. His research interests include machine learning, data mining and their applications. He has published over 90 top-ranked AI conference and journal papers, had multiple Oral/Spotlight Papers at NeurIPS, ICML, KDD, ICLR, AAAI and IJCAI, etc..

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 11 May 2026
doi:10.20944/preprints202605.0725.v1

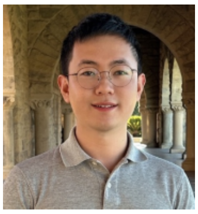

Dr Ming Jin is an assistant professor at the School of ICT, Griffith University. He obtained his Ph.D. degree from Monash University, Australia, in 2024. His research focuses on time series analysis and spatio-temporal data mining. He has published over 60 papers in prestigious venues (e.g., ICLR, ICML, NeurIPS, KDD, TPAMI, and TKDE) and gathered over 6,000 citations with an h-index of 25. He has been named among IEEE Computing's Top 30 Early Career Professionals 2025. He currently serves as an associate editor for Neurocomputing and regularly contributes as a senior program committee member for leading AI and data mining conferences.

Qingsong Wen is currently the Head of AI and Chief Scientist at Squirrel Ai Learning. He received his M.S. and Ph.D. degrees in Electrical and Computer Engineering from Georgia Institute of Technology, USA. His research interests include machine learning, data mining, and signal processing. He has published over 150 top-ranked AI conference and journal papers, had multiple Oral/Spotlight Papers at NeurIPS, ICML, and ICLR, had multiple Most Influential Papers at IJCAI, and received multiple IAAI Deployed Application Awards at AAAI. Currently, he serves as Chair of IEEE Task Force on AI for Time Series and Spatio-Temporal Data, and Vice Chair of INNS AI for Education Section. He also regularly serves as Area Chair of NeurIPS, ICML, ICLR, KDD, IJCAI, ICASSP, etc, and Associate Editor for IEEE Transactions on Pattern Analysis and Machine Intelligence.

Peng Wang received the PhD degree in computer science from Fudan University, China in 2007. He is a Professor at the School of Computer Science, Fudan University. His research includes time series data management and analytics. He has published over 100 papers in leading international journals and conferences.

Wei Wang received the PhD degree in computer science from Fudan University, China in 1998. He is a Professor at the School of Computer Science, Fudan University. His research includes database and data mining. He has published over 200 papers in leading international journals and conferences.

Jiang Bian received the bachelor's degree from Peking University, China, and the PhD degree in computer science fromt the Georgia Institute of Technology, USA. He is a senior principal research manager with Microsoft Research. He is leading the machine learning solutions and services group, with the main focus on designing cutting-edge machine learning algorithms into real-world application scenarios, including finance, healthcare, supply-chain and sustainability. Prior to this, he was a scientist with Yahoo! Labs in the United States, responsible for the content optimization and personalization and Web search modules in Yahoo! Homepage. After that, he jointed a leading content distribution platform in China, i.e., Yidian Inc., and became one of the core members of this startup company, with the major responsibility of developing advanced recommendation models. He has authored tens of research papers in many well-recognized international conferences and has submitted a couple of US patents. He has also served as PC member for several international conferences and Peer Reviewer for a few well-known journals.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.