To what extent can self-guidance mechanisms in unconditional diffusion models like TSDiff improve forecasting

Assignee Research

June 10, 2026

Abstract

Diffusion models have achieved state-of-the-art performance in generative modeling tasks across various domains. Prior works on time series diffusion models have primarily focused on developing conditional models tailored to specific forecasting or imputation tasks. In this work, we explore the potential of task-agnostic, unconditional diffusion models for several time series applications. We propose TSDiff, an unconditionally-trained diffusion model for time series. Our proposed self-guidance mechanism enables conditioning TSDiff for downstream tasks during inference, without requiring auxili

1 Introduction

This paper examines: Predict, Refine, Synthesize: Self-Guiding Diffusion Models for Probabilistic Time Series Forecasting. Research question: To what extent can self-guidance mechanisms in unconditional diffusion models like TSDiff improve forecasting accuracy on long-horizon prediction tasks in the Monash benchmark compared to autoregressive or transformer-based approaches?.

2 Methodology

Systematic literature search across multiple databases yielded 11 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 5.5/10.

3 Results

11 papers retrieved. 13 claims extracted; 4 independently verified. Quality review score: 5.5/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
Experiments were conducted on eight univariate time series datasets: Solar, Electricity, Traffic, Exchange, M4, UberTLC,	×	0.09
The datasets used in the experiments are available in GluonTS.	×	0.03
The quality of probabilistic forecasts was evaluated using the continuous ranked probability score (CRPS).	×	0.04
CRPS was approximated by the normalized average quantile loss using 100 sample paths.	×	0.02
Results report means and standard deviations over three independent runs.	×	0.01
TSDiff is an unconditional diffusion model for time series.	✓	0.22
TSDiff utilizes a self-guidance mechanism that enables conditioning the model during inference without requiring auxilia	✓	0.21
The self-guidance mechanism allows the unconditional model to perform arbitrary forecasting and imputation tasks that ar	×	0.14
TSDiff can generate probabilistic forecasts in the presence of missing values.	×	0.05
The implicit probability density learned by TSDiff can be leveraged to refine the predictions of base forecasters.	✓	0.21
Synthetic samples generated by TSDiff are adequate for training downstream forecasters.	✓	0.15
The self-guidance approach is competitive against task-specific models on several datasets and across multiple forecasti	×	0.14
The code for the study is available at github.com/amazon-science/unconditional-time-series-diffusion.	×	0.09

References

- `http://arxiv.org/abs/2410.11674v2`
- `http://arxiv.org/abs/2307.11494v3`
- `http://arxiv.org/abs/2106.09305v3`