Self-Guidance in Task-Agnostic Diffusion Models for Probabilistic Time Series Forecasting

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 11 peer-reviewed papers addressing the following research question: To what extent does self-guidance in task-agnostic diffusion models improve probabilistic forecast calibration metrics compared to conditional diffusion approaches on standard time series benchmarks. 13 claims were extracted from source literature; 7 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.9/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Predict, Refine, Synthesize: Self-Guiding Diffusion Models for Probabilistic Time Series Forecasting. Research question: To what extent does self-guidance in task-agnostic diffusion models improve probabilistic forecast calibration metrics compared to conditional diffusion approaches on standard time series benchmarks?.

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 6.9/10.

3 Results

11 papers retrieved. 13 claims extracted; 7 independently verified. Quality review score: 6.9/10.

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).	✓	0.16
CRPS was approximated by the normalized average quantile loss using 100 sample paths.	×	0.03
Results report means and standard deviations over three independent runs.	×	0.01
TSDiff is an unconditional diffusion model for time series.	✓	0.21
TSDiff proposes a self-guidance mechanism that enables conditioning the model during inference without requiring auxilia	✓	0.20
The self-guidance mechanism allows the unconditional model to be used for arbitrary forecasting and imputation tasks tha	✓	0.15
TSDiff can generate probabilistic forecasts in the presence of missing values.	×	0.06
The implicit probability density learned by TSDiff can be leveraged to refine the predictions of base forecasters.	✓	0.20
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 without requiring conditional	✓	0.16
The code for the study is available at github.com/amazon-science/unconditional-time-series-diffusion.	×	0.08

References

- `http://arxiv.org/abs/2110.13179v8`
- `http://arxiv.org/abs/2307.11494v3`
- `http://arxiv.org/abs/2503.20240v3`