Does the task-agnostic self-guidance approach in diffusion models demonstrate superior scaling properties with

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: Does the task-agnostic self-guidance approach in diffusion models demonstrate superior scaling properties with increased model size compared to task-specific conditional models on probabilistic forecasting metrics?.

2 Methodology

Systematic literature search across multiple databases yielded 9 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

9 papers retrieved. 14 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
TSDiff can generate probabilistic forecasts, even in the presence of missing values.	×	0.05
The implicit probability density learned by TS-Diff can be leveraged to refine the predictions of base forecasters.	✓	0.17
Synthetic samples generated by TSDiff are adequate for training downstream forecasters.	×	0.14
Experiments were conducted on eight univariate time series datasets from different domains, available in GluonTS.	×	0.08
The quality of probabilistic forecasts was evaluated using the continuous ranked probability score (CRPS).	×	0.15
The CRPS was approximated by the normalized average quantile loss using 100 sample paths.	×	0.03
TSDiff's use cases include Predict, Refine, and Synthesize.	×	0.12
TSDiff can be conditioned during inference to perform predictive tasks such as forecasting.	×	0.09
Predictions of base forecasters can be improved by leveraging the implicit probability density of TSDiff.	×	0.15
Realistic samples generated by TSDiff can be used to train downstream forecasters achieving good performance on real tes	×	0.10
TSDiff is an unconditional diffusion model for time series.	✓	0.23
TSDiff proposes two inference schemes to utilize the model for forecasting.	×	0.05
TSDiff proposes a self-guidance mechanism that enables conditioning the model during inference, without requiring auxili	✓	0.20
The self-guidance approach of TSDiff is competitive against task-specific models on several datasets and across multiple	✓	0.16

References

- http://arxiv.org/abs/2110.13179v8

- http://arxiv.org/abs/2312.02246v4
- http://arxiv.org/abs/2307.11494v3