How robust is TSDiff's performance to variations in hyperparameters (e.g., diffusion steps, noise scheduling)

Assignee Research

June 10, 2026

Abstract

Advances in diffusion models for generative artificial intelligence have recently propagated to the time series (TS) domain, demonstrating state-of-the-art performance on various tasks. However, prior works on TS diffusion models often borrow the framework of existing works proposed in other domains without considering the characteristics of TS data, leading to suboptimal performance. In this work, we propose Adaptive Noise schedule for Time series diffusion models (ANT), which automatically predetermines proper noise schedules for given TS datasets based on their statistics representing non-s

1 Introduction

This paper examines: ANT: Adaptive Noise Schedule for Time Series Diffusion Models. Research question: How robust is TSDiff's performance to variations in hyperparameters (e.g., diffusion steps, noise scheduling) across different domains in the Monash benchmark, and what is the trade-off between sample quality and computational efficiency?.

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 3.8/10.

3 Results

11 papers retrieved. 7 claims extracted; 0 independently verified. Quality review score: 3.8/10.

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
ANT yields a 27.8% gain compared to a linear schedule of TSDiff on the M4 dataset.	×	0.03
ANT is adaptive, flexible, model-agnostic, and efficient.	×	0.06
ANT proposes a noise schedule that minimizes the discrepancy between the ideal linear line and the non-stationarity curv	×	0.12
ANT decreases the non-stationarity of TS on a linear scale.	×	0.13
Diffusion step embedding (DE) is unnecessary for TS diffusion models when a linear schedule is employed.	×	0.13
Non-linear schedules are more robust to the change of the number of diffusion steps than linear schedules in terms of pe	×	0.11
ANT outperforms the baseline methods in extensive experimental results across various datasets.	×	0.06

References

- http://arxiv.org/abs/2307.11494v3
- http://arxiv.org/abs/1811.00620v1
- http://arxiv.org/abs/2410.14488v1