TS-CorrNet: Explicit Linear and Non-linear Decomposition for Time-Series Imputation

Anonymous authors
Paper under double-blind review

Abstract

Existing time series imputation techniques generally adhere to a single modeling paradigm. Still, real-time series require three attributes simultaneously: (i) dynamics mixed with linear and non-linearity, (ii) heteroskedastic reconstruction difficulty, and (iii) non-randomized missing data (MNAR), which many benchmarks do not even evaluate. This paper proposes **TS-CorrNet** (Time-Series Correction Network). It is a four-stage hierarchical model that explicitly separates a linear component (BRITS $\rightarrow$ ARIMA) and two non-linear correction branches (Mask-conditioned BiLSTM and Mask-attention Transformer) before recombining them via a learnable gate. *Volatility-Weighted Loss* is used for training; this is not an arbitrary heuristic, but rather a risk-weighted M-estimator that accounts for heteroskedasticity. Across four standard public benchmarks and a financial application domain—and spanning MCAR, MAR, and MNAR conditions—TS-CorrNet outperformed five baselines (ranging from classical to diffusion-based methods) in all 15 dataset$\times$ mechanism cells. Furthermore, it demonstrated the most significant improvement (a 70% reduction in MAE) in the domain with the greatest data scarcity. Ablation results for the six components identified explicit linear/nonlinear decomposition and learnable gates as the two most critical design elements ($p < 10^{-3}$, $\sim$80% degradation when removed). The code is included in the supplementary material.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.