WaveStitch: Flexible and Fast Conditional Time Series Generation with Diffusion Models

Aditya Shankar [1], Lydia Chen [1,2], Arie van Deursen [1], and Rihan Hai [1]

[1] Delft University of Technology, The Netherlands
{a.shankar,arie.vandeursen,r.hai}@tudelft.nl
[2] Université de Neuchâtel, Switzerland
lydiaychen@ieee.org

Synthesising time series under conditions is crucial for forecasting, imputation, and generative tasks. Time series often have metadata and partially observed signals that jointly influence the generated values. However, existing methods face three key limitations: (1) they condition only on the metadata or on the observed values, but not both together; (2) they adopt either training-time approaches that fail to generalise to unseen scenarios, or inference-time approaches that ignore metadata; and (3) they suffer from trade-offs between generation speed and temporal coherence across time windows, choosing either slow and coherent autoregressive methods or fast and incoherent parallel ones. We propose WaveStitch, a novel diffusion-based method to overcome these hurdles through: (1) dual-sourced conditioning on both metadata and partially observed signals; (2) a hybrid training-inference architecture, incorporating metadata during training and observations at inference via gradient-based guidance; and (3) a novel pipeline-style paradigm that generates time windows in parallel while preserving coherence through an inference-time conditional loss and a stitching mechanism. Across diverse datasets, WaveStitch demonstrates adaptability to arbitrary patterns of observed signals, achieving 1.81x lower mean-squared-error compared to the state-of-the-art, and generates data up to 166.48x faster than autoregressive methods while maintaining coherence. Our code is available at: `https://github.com/adis98/WaveStitch`.