אוניברסיטת בן-גוריון בנגב
Ben-Gurion University of the Negev
DSRC@BGU
ATA SCIENCE RESEARCH CENTER
BEN-GURION UNIVERSITY OF THE NEGEV

A Diffusion Model for Regular Time Series Generation from Irregular Data with Completion and Masking

Gal Fadlon*, Idan Arbiv*, Nimrod Berman, Omri Azencot

Curious? Scan to dive in

NEURAL INFORMATION PROCESSING SYSTEMS

Contributions

- We introduce a **novel model,** a vision-inspired diffusion framework that efficiently handles irregularly sampled time series from short to very long sequences.
- Using a **masking strategy**, we optimize directly on observed signals instead of assuming full data availability.
- Our approach delivers **state-of-the-art results,** improving benchmarks by up to 70% while cutting computational cost by 85%.

Background

Regular Time Series Generation from Irregular Data

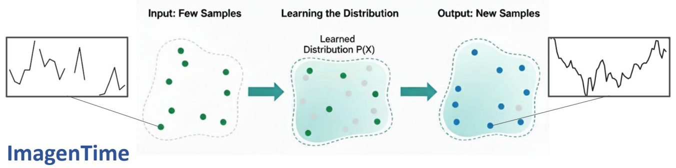

ImagenTime

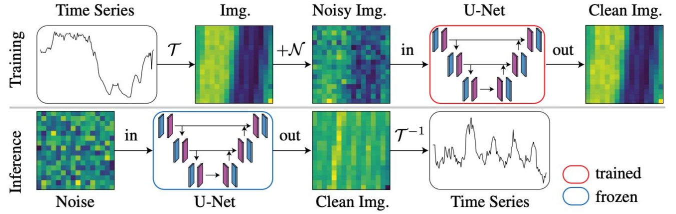

ImagenTime Introduces a **novel connection** between time series modeling and image generation and has achieved state-of-the-art performance on regular generative tasks. By transforming time series into images via **invertible transforms**, it leverages vision-based **diffusion models** to handle sequences from very short to ultra-long within a **single framework**.

Model Overview

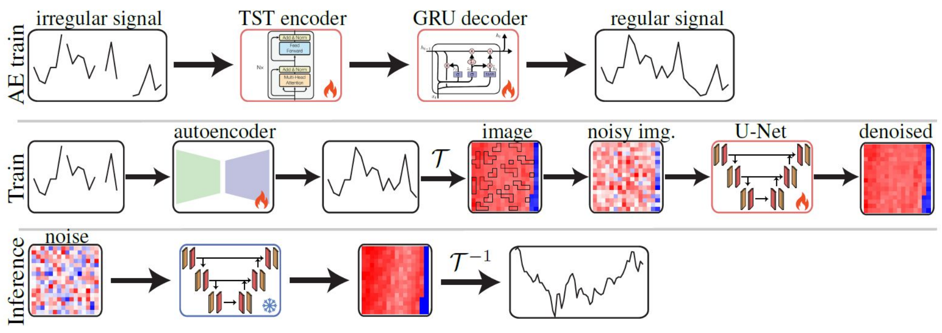

Method

Step 1: Completion. In the first step of our **Two-Step Method for Natural Neighborhoods**, a scalable and efficient **Time Series Transformer** (TST)–based autoencoder is trained to fill missing values in the irregular sequence. TST is chosen to **avoid the computational overhead** associated with NCDE-based preprocessing. This produces a regularly sampled sequence, naturally forming pixel neighborhoods. Training uses a **masked reconstruction loss.**

Step 2: Vision Diffusion with Masking. **Vision-based diffusion** model is trained on the completed sequence, first transforming it to an image representation using the **delay embedding.** Imputed values are used only as a **weak conditioning signal**, preventing over-reliance on potentially inaccurate data. The diffusion loss is computed only on observed pixels, effectively masking out imputed values.

Results

Table 1: Averaged results over 30%, 50%, 70% missing rates for length 24.

	Model	ETTh1	ETTh2	ETTm1	ETTm2	Weather	Electricity	Energy	Sine	Stock
Disc.	TimeGAN-Δt	0.499	0.499	0.499	0.499	0.497	0.499	0.474	0.497	0.479
	GT-GAN	0.471	0.369	0.412	0.366	0.481	0.427	0.325	0.338	0.249
	KoVAE	0.197	0.081	0.050	0.067	0.332	0.498	0.323	0.043	0.118
	Ours	**0.037**	**0.009**	**0.012**	**0.011**	**0.057**	**0.384**	**0.080**	**0.010**	**0.008**
Pred.	TimeGAN-Δt	0.267	0.336	0.235	0.314	0.394	0.262	0.457	0.334	0.072
	GT-GAN	0.186	0.092	0.125	0.094	0.145	0.148	0.069	0.096	0.020
	KoVAE	0.057	0.054	0.045	0.050	0.057	**0.047**	0.050	0.074	0.017
	Ours	**0.053**	**0.046**	**0.044**	**0.044**	**0.022**	0.049	**0.047**	**0.069**	**0.012**
FID	TimeGAN-Δt	3.140	3.199	3.419	3.218	2.378	23.39	6.507	2.780	2.668
	GT-GAN	2.212	8.635	14.29	6.385	2.758	9.993	1.531	1.698	2.181
	KoVAE	1.518	0.248	0.180	0.280	3.699	6.163	0.629	0.037	0.369
	Ours	**0.124**	**0.035**	**0.047**	**0.024**	**0.170**	**3.580**	**0.132**	**0.015**	**0.036**
Corr.	TimeGAN-Δt	3.743	1.051	2.350	0.579	1.200	13.24	3.765	2.424	1.399
	GT-GAN	7.148	0.916	2.467	0.356	0.791	14.92	3.889	3.282	0.261
	KoVAE	0.183	0.177	0.130	0.262	2.899	4.283	2.630	0.041	0.064
	Ours	**0.084**	**0.054**	**0.065**	**0.039**	**0.396**	**2.031**	**0.922**	**0.015**	**0.019**

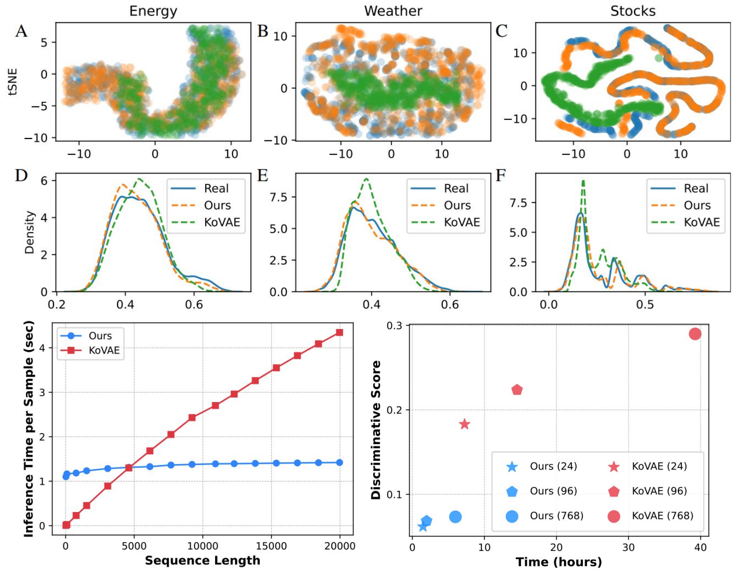

Table 2: Ablation Studies, disc. scores with 30%, 50% and 70% for sequence length of 24

Model	30%		50%		70%	
	Energy	Stock	Energy	Stock	Energy	Stock
KoVAE + TST	0.399	0.109	0.407	0.064	0.408	0.037
TimeAutoDiff + TST	0.293	0.100	0.329	0.101	0.468	0.375
TransFusion + TST	0.201	0.050	0.279	0.058	0.423	0.065
Ours (Mask Only)	0.157	0.087	0.269	0.168	0.372	0.237
Ours (Without Mask)	0.158	0.025	0.307	0.045	0.444	0.013
Ours	**0.048**	**0.007**	**0.065**	**0.007**	**0.128**	**0.007**