UNIVERSITY OF BIRMINGHAM

SIEMENS

Stationarity-Aware Retrieval-Augmented Time Series Forecasting

Shiqiao Zhou[1], Holger Schöner[2], Zipeng Wu[1], Edouard Fouché[2], IAG Wilson[1], Shuo Wang[1]*
[1]University of Birmingham, UK & [2]Siemens AG, Germany

KDD 2026 August 9–13, 2026 Jeju, Korea

Paper Code My Web

Motivation

Traditional retrieval-based forecaster assumes: **similar pasts imply similar futures.** However, real-world time series are often non-stationary: **similar histories may lead to different futures**. Thus, similarity-only retrieval becomes brittle and redundant.

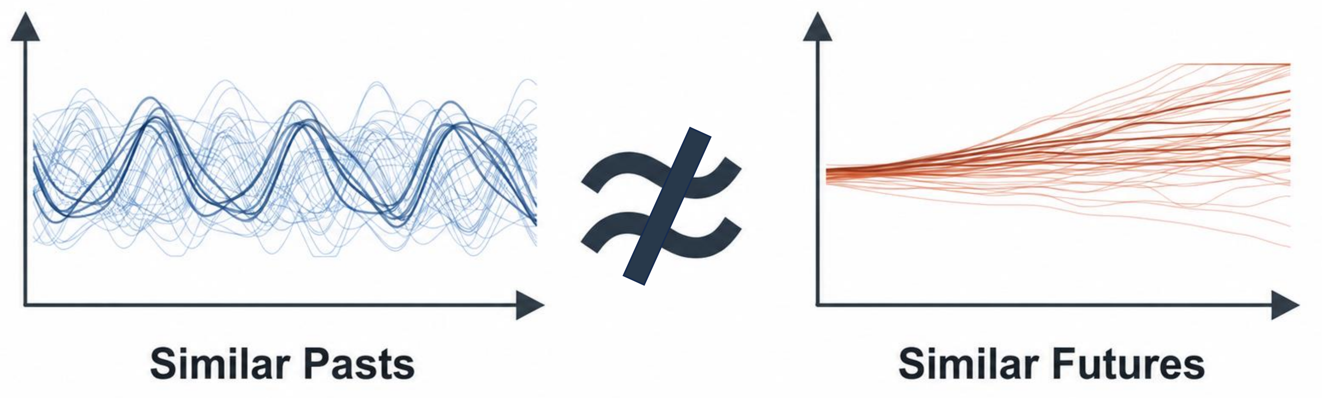

- **Examples**

Stationary Data

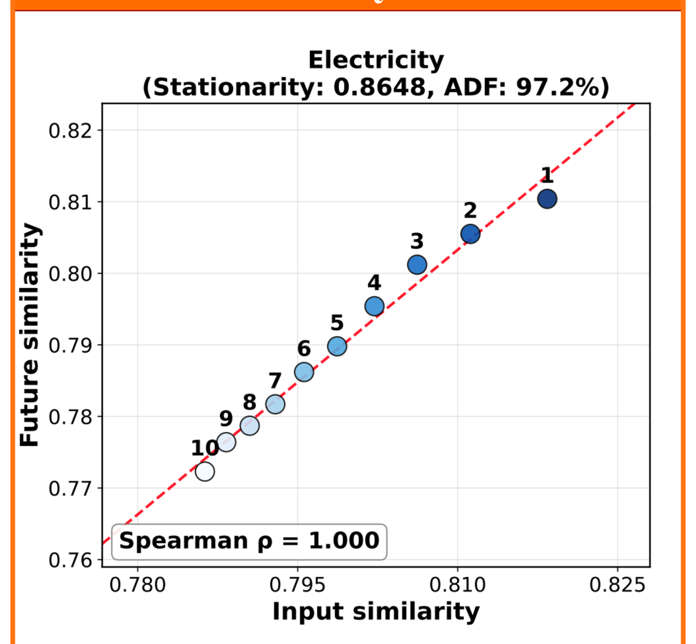

Assumption Holds: The most similar past perfectly corresponds to the most accurate future.

Non-Stationary Data

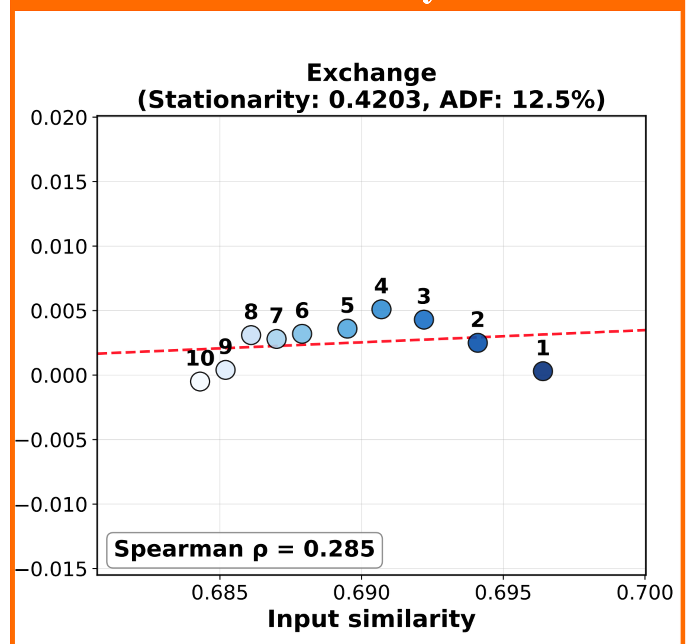

Assumption Broken: Superficially similar past segments have completely unrelated future.

Key Takeaways

1. **Time series similarity alone is not enough.** Under non-stationarity, similar historical patterns may lead to different futures,
2. **Retrieval should adapt to stationarity.** SARAF uses time-aligned matching, stationarity-guided diversity selection, and adaptive Gaussian fusion to balance relevance and diversity across different temporal regimes.
3. **Better retrieval improves forecasting.**

Method

Overview of SARAF:

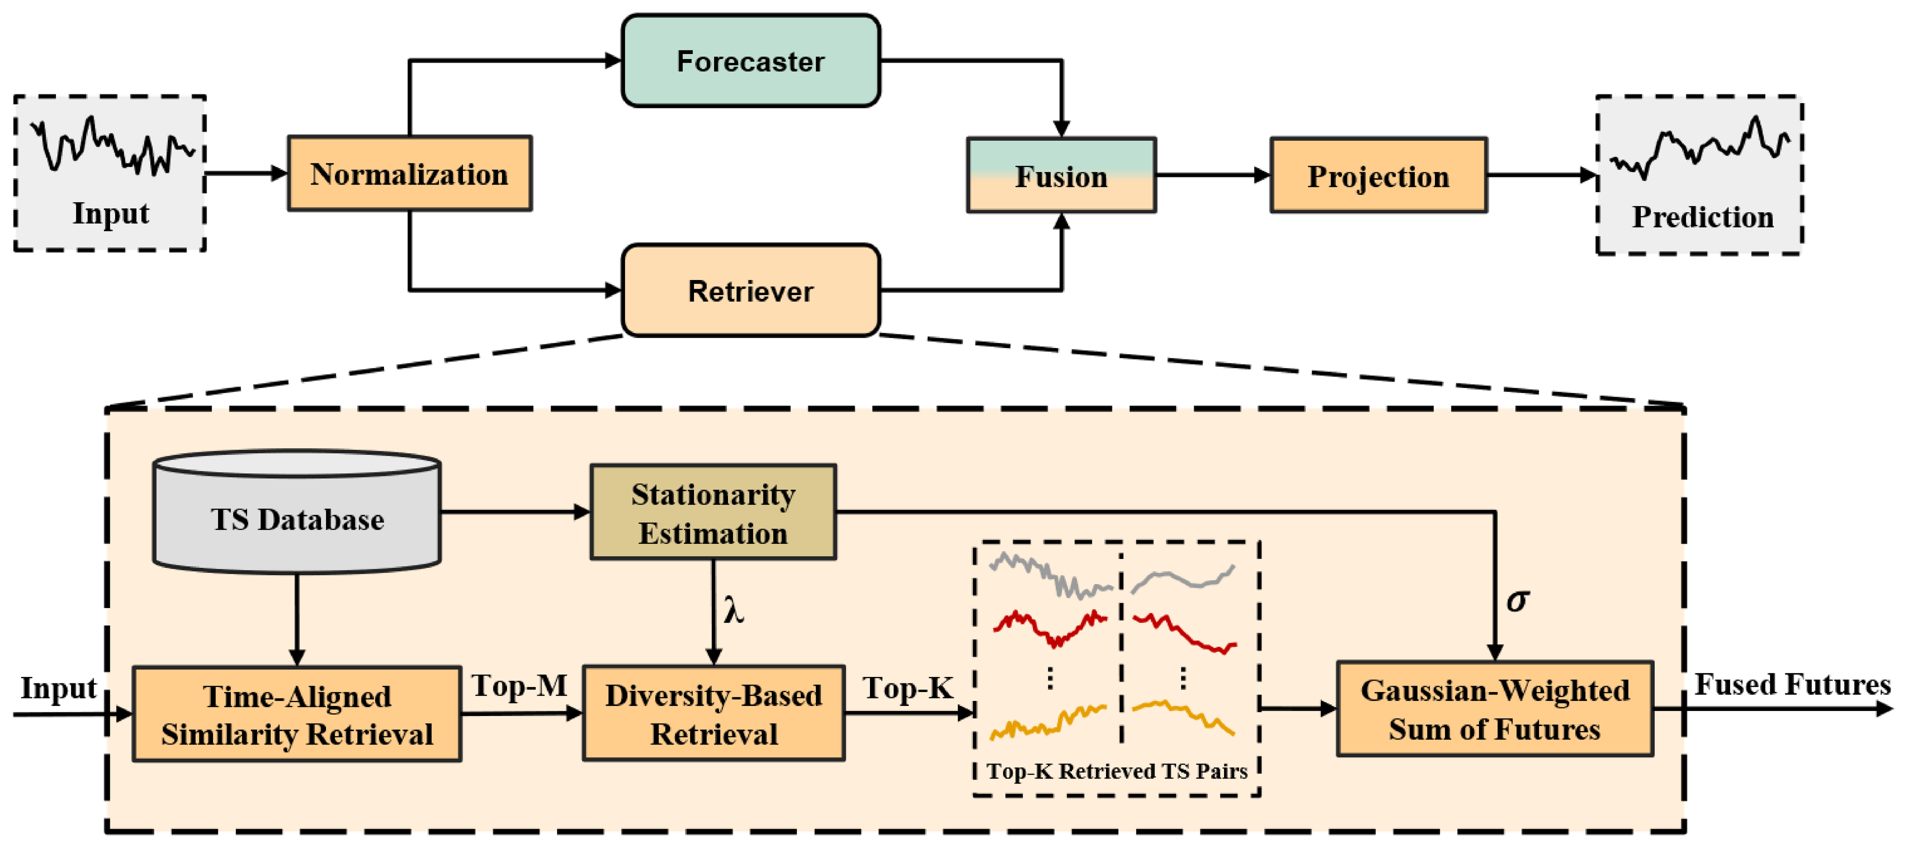

- **Stationarity Estimation.** To estimate dataset-level stationarity, We partition each sample $\mathbf{x}_i \in \mathbb{R}^{L\times C}$ into W sub-windows and compute the variation of local means and standard deviations across them. Let v_μ and v_σ denote these variations. The stationarity score $\bar{s}$ is defined by normalizing v_μ and v_σ with the global scale $\bar{\sigma}$ to achieve scale invariance.

$$\bar{s} = \frac{1}{N}\sum_{i=1}^{N} \tilde{s}_i\,, \quad \tilde{s}_i = \frac{1}{2}\left(\left[1 - \min\left(1, \frac{v_\mu}{\bar{\sigma}}\right)\right] + \left[1 - \min\left(1, \frac{v_\sigma}{\bar{\sigma}}\right)\right]\right) \in [0,1].$$

- **Time-Aligned Retrieval.** Pattern similarity alone may retrieve temporally misaligned cycles. We combine time series similarity with calendar-aware temporal cues to encourage retrieval from corresponding daily, weekly, and seasonal periods.

- **Stationarity-Guided Diversity Retrieval.** Highly similar neighbors are often redundant, especially for non-stationary series. We apply MMR-based selection to the Top-M candidates from time-aligned retrieval, with the relevance--diversity trade-off controlled by stationarity:

$$\lambda(\bar{s}) = \lambda_{\min} + \bar{s}(\lambda_{\max} - \lambda_{\min}).$$

- **Stationarity-Guided Gaussian Fusion**. The Top-K retrieved futures are aggregated using Gaussian-kernel weights:

$$w_k = \frac{\exp\left(-d_k^2/2\sigma(\bar{s})^2\right)}{\sum_{j=1}^{K}\exp\left(-d_j^2/2\sigma(\bar{s})^2\right)}, \quad \sigma(\bar{s}) = \sigma_{\min} + (1-\bar{s})(\sigma_{\max} - \sigma_{\min}).$$

Experiment Results

- **Main Results**

Model	SARAF (ours)		RAFT		DUET	
Metric	MSE	MAE	MSE	MAE	MSE	MAE
ETTh1	**0.415**	0.435	0.418	**0.434**	0.430	0.447
ETTh2	**0.340**	**0.392**	0.358	0.408	0.358	0.405
ETTm1	**0.346**	0.378	**0.347**	0.377	0.349	0.382
ETTm2	**0.248**	**0.314**	0.257	0.321	0.254	0.314
Exchange	**0.394**	**0.415**	0.449	0.444	0.469	0.446
Solar	0.206	0.258	0.211	0.255	**0.198**	**0.245**
ECL	**0.152**	**0.251**	0.156	0.253	0.163	0.254
Traffic	0.395	0.280	0.401	0.281	**0.380**	**0.249**

*Average performance over four horizons

- **Ablation Results**

Dataset	SARAF	w/o Ret	w/o For	w/o time	w/o div	w/o sta	w/o div+sta
ETTh1	**0.415**	0.421	1.328	0.42	0.415	0.416	0.416
ETTm1	**0.346**	0.359	1.272	0.347	0.347	0.348	0.348
Exchange	0.394	0.411	**0.377**	0.398	0.400	0.398	0.399
ECL	0.152	0.163	1.613	0.156	**0.151**	0.152	0.151
Traffic	**0.395**	0.414	2.768	0.402	0.396	0.396	0.397

*Average MSE over four horizons

Additional Analysis

- **Diversity-Based Retrieval Analysis.**

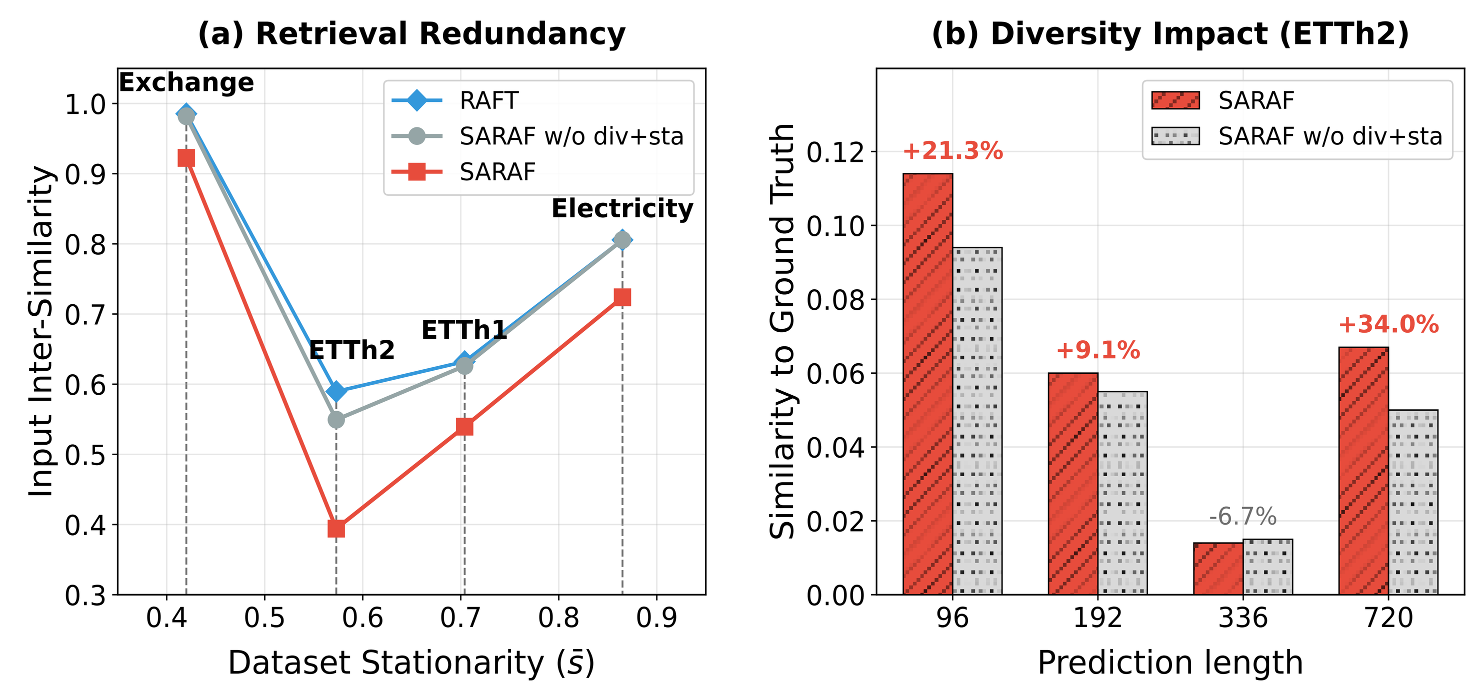

- **Efficiency**

Model	Param. (M)	Model Mem. (MiB)	All Mem. (MiB)	Speed (ms/iter)
SARAF (ours)	**0.088**	**0.335**	54.077	**0.334**
DUET	6.66	25.408	**36.03**	7.202
RAFT	0.104	0.397	122.996	0.59