Interpretable Structural Stability Analysis for Long-Term Cognitive IoT Time-Series Data

Basab Nath[1]*, Yonis Gulzar[2]*
School of Computer Science and Engineering, Bennett University, Greater Noida, India[1]
Department of Management Information Systems-College of Business Administration,
King Faisal University, Al-Ahsa 31982, Saudi Arabia[2]

***Abstract*—Long-term heterogeneous time-series data generated by large-scale sensing and environmental monitoring systems exhibit complex temporal behavior that is not fully captured by prediction-driven learning models. While most existing approaches emphasize short-term forecasting accuracy, comparatively little attention has been given to the analysis of long-term structural stability inherent in such data. In this work, we propose a lightweight, training-free analytical framework for quantifying structural stability in long-duration time-series using stability-preserving preprocessing and interpretable temporal statistics. The proposed method combines total variation regularization with rolling statistical analysis to assess the consistency of local temporal behavior relative to global characteristics over extended time horizons. Structural stability is quantified using a simple yet effective stability index that captures deviations between local and global temporal trends. The framework is evaluated using more than two decades of daily environmental observations, including temperature, relative humidity, and precipitation, obtained from the NASA POWER repository for a representative location in Assam, India. Experimental results demonstrate consistent and systematic reductions in the stability index following preprocessing across all variables, indicating improved temporal consistency without structural distortion. Additional robustness analysis across multiple temporal scales confirms that the proposed framework is insensitive to window size selection and preserves long-term structural behavior. These findings suggest that meaningful insights into temporal stability can be obtained without reliance on model training or predictive learning, making the proposed approach suitable for interpretable, resource-efficient analysis of long-term heterogeneous time-series data.Unlike conventional stability descriptors such as variance-based measures or correlation-based consistency metrics, the proposed stability index directly quantifies local-to-global deviation of temporal descriptors across multiple window scales, enabling interpretable and comparable stability assessment without requiring model training or forecasting error baselines.**

***Keywords*—Structural stability; training-free framework; total variation regularization; rolling statistics; stability index**

I. Introduction

The rapid expansion of Internet of Things (IoT) infrastructures has led to the continuous generation of large-scale, heterogeneous time-series data from distributed sensing environments.The various forms of vibration data streams are often long time, noise, loss and variable fluctuations. Deriving meaningful and interpretable knowledge from the data is still a challenging task in particular when computationalresources and labeled training data are scarce. While time-series analytics has traditionally focused on prediction, anomaly detection, or data reduction, these objectives do not fully address whether the underlying temporal behavior of a system remains consistent over extended periods [1], [2], [3].

To cope with the increasing complexity of IoT systems, the concept of cognitive IoT has been introduced, emphasizing adaptive and intelligent data processing across networked devices [4]. Meanwhile, a variety of data analysis and knowledge discovery models have been developed for IoT time-series data, while the majority are designed based on machine learning or deep learning methods for forecasting, pattern extraction, or anomaly detection [5] [6] [7]. Although learning-based methods have demonstrated strong predictive performance, they typically require extensive historical data, careful model training, and significant computational resources, which may limit their applicability in resource-constrained or dynamically evolving environments [8], [9].

Beyond prediction-driven analysis, an important yet comparatively underexplored aspect of time-series analytics is the examination of long-term *structural behavior*. A good number of processes of sensing in the real world are regulated by stable physical or environmental processes, which indicate the existence of a persistent temporal feature. Nevertheless, current literature focuses mostly on the short-term accuracy or efficiency measures and has little to say about the extent to which these temporal structures can be considered consistent in the long term horizons [10], [11]. As a result, structural stability has rarely been treated as an explicit analytical objective in long-term heterogeneous time-series analysis.

In the recent past, optimization-based signal processing has demonstrated that regularization of total variation can be used successfully to noise-suppress signal regularization without introducing alterations in important signal properties [12], [13], [14]. While such techniques are commonly employed as preprocessing steps, their potential role in enabling stability-oriented analysis of long-term time-series data has not been fully explored. Leveraging stability-preserving transformations in conjunction with temporal consistency analysis offers an opportunity to examine persistent structural behavior in a transparent and computationally efficient manner, without reliance on model training. Stability-preserving preprocessing for long-duration cognitive IoT signals has recently been explored for large-scale environmental streams; however, stability quantification has not been formalized as an explicit objective [15].

Unlike prediction-driven analytics, this study focused not on forecasting future values or detecting isolated outliers, but on quantifying the persistence of temporal structure over

*Corresponding author.

long horizons. Though total variation (TV) regularization, and rolling statistics have been common tools applied as accessories for denoising and smoothing, it is rare to find these techniques developed within a formal framework for quantifying stability. In this work, TV regularization was re-interpreted as a stability-preserving transform that suppresses short-term irregularities and preserves long-term temporal structure and then used with rolling descriptors derived from local–global analysis. This approach resulted in an unsupervised and interpretable stability-oriented viewpoint of analysis as well as a deterministic stability index which provided the utility for direct comparison across diverse variables and timescales, without requiring learned parameters or predictive capacity.

Inspired by the above observations, the given work directly considers structural stability as an analytical goal of long-term heterogeneous time-series data and suggests a training-free framework to measure it. Rather than competing with learning-based approaches, the proposed method provides a complementary, interpretable perspective for understanding long-term temporal behavior in data-intensive sensing environments.

This work makes the following contributions:

- Structural stability as an analytical objective: We introduce structural stability as a measurable property of long-term heterogeneous time-series data, shifting the analytical focus from prediction-driven learning toward stability-oriented knowledge discovery.
- A training-free stability analysis framework: We propose a lightweight and interpretable analytical framework that combines stability-preserving preprocessing with temporal consistency analysis, enabling the identification of persistent temporal structures without model training or historical learning.
- Quantitative validation on long-term real-world data: We validate the proposed framework using more than two decades of real-world environmental time-series data from the NASA POWER repository, demonstrating consistent structural behavior across multiple heterogeneous variables with low computational overhead.

The remainder of this paper is organized as follows. Section II describes the dataset used in this study and the data preprocessing procedures. Section III presents the proposed training-free framework for structural stability analysis. Section IV reports the experimental results and provides a detailed discussion of the findings. Finally, Section V concludes the paper and outlines directions for future work.

II. Related Work

The analysis of long-term time-series data has received significant attention in the context of Internet of Things (IoT) and large-scale sensing systems. Early foundational work by Atzori et al. [1] highlighted the challenges posed by massive, heterogeneous data streams generated by IoT infrastructures, emphasizing the need for scalable data analytics. Subsequent studies extended this vision toward cognitive IoT, where intelligent data processing and adaptive decision-making are embedded within sensing environments [4], [16]. These works established the importance of extracting meaningful information from continuous data streams but largely focused on system architectures and application-driven intelligence rather than long-term structural properties of the data. Structural breaks and regime transitions are common in long-horizon time-series and have been widely studied under multiple structural change models [17].

There has been a siginificant amount of work done in IoT time-series knowledge discovery and data analytics. Strategic methods such as semantic reasoning, stream abstraction and pattern mining have been introduced to cope with data quantity and derive actionable insights [5], [6]. Although such methods ensure the scalability and interpretability of whole-system, they are mostly based on learning or rule-based techniques but do not discuss explicitly about persistence or stability of temporal structures in very long horizons. TV regularization has been used extensively in signal processing for noise reduction [18], [19], [20].

Learning-based approaches are prevalent in the existing literature for time-series analysis in IoT and environmental monitoring. Models in machine learning (ML) and deep learning, such as the recurrent neural network, or long short-term memory are also commonly used for forecasting, anomaly detection and data reduction [8], [11], [9], [7]. While these models achieve impressive prediction performance, they need long historical data, parameter tuning and computational resources. Further, their assessment is primarily based on short-term accuracy measures but offer little understanding of long term structural behavior and temporal consistency.

Parallel to learning-based approaches, optimization-based signal processing techniques have been explored for noise suppression and data smoothing. Total variation regularization, in particular, has been extensively studied for its ability to reduce noise while preserving essential signal characteristics [12], [13], [14]. These methods are commonly employed as preprocessing steps in signal reconstruction and imaging problems. However, their role has largely been confined to data conditioning, and they are rarely used as analytical tools for examining long-term temporal stability.

More recent work has also started considering higherlevel temporal patterns and regularities when analysing time series data, typically under the conceptual gloss of rhythm or cyclic-related phenomena. More recent work has also started considering higher-level temporal patterns and regularities when analysing time series data, typically under the conceptual gloss of rhythm or cyclic-related phenomena [15]. Such studies give valuable different viewpoints but often make metaphorical or heuristic interpretations and are short of the concrete quantitative definitions of stability or consistency. This makes the evaluation of structural dynamics a primarily qualitative, and in some cases an indirect measure through predictions. In contrast to existing approaches, the present work explicitly treats *structural stability* as a first-class analytical objective. Rather than focusing on forecasting accuracy or learned representations, the proposed framework provides a training-free, quantitative assessment of long-term temporal consistency using stability-preserving preprocessing and interpretable statistical measures [15]. By grounding the analysis in well-established optimization techniques and transparent metrics, this work complements existing learning-based and system-oriented studies while addressing an underexplored aspect of long-term heterogeneous time-series analysis.

III. Dataset Description

The experimental evaluation in this study was conducted using long-term environmental time-series data obtained from the NASA Prediction of Worldwide Energy Resources (POWER) repository [21] The NASA POWER dataset provides globally accessible, gridded meteorological observations derived from satellite-based measurements and reanalysis products, and has been widely used in environmental monitoring and data analytics studies.

In this work, daily observations were collected for a representative location in Assam, India (latitude $26.1445°$N, longitude $91.7362°$E), covering the period from January 2000 to December 2024. The dataset consists of 9,132 daily records and includes three heterogeneous environmental variables: temperature at 2 meters above the surface (T2M, in °C), relative humidity at 2 meters (RH2M, in %), and corrected precipitation (PRECTOTCORR, in mm/day). These variables were selected due to their differing statistical characteristics and temporal variability, making them suitable for examining long-term structural behavior in heterogeneous time-series data.

Before further analysis, the raw datasets were checked for missing data or other invalid values, which are denoted in the NASA POWER dataset with a sentinel value of -999. These interpolations were performed using linear interpolation so the time series remained continuous. Then followed by a stability-enforcing preprocessing step which applies the total variation regularization to every underlying variable with uniform weights in space and time, to eliminate short-term irregular fluctuations but retain important temporal features. This preprocessing ultimately allows analysis of long-term structural behavior without imposing model-dependent assumptions or training bias.

Fig. 1, Fig. 2 and Fig. 3 compare TV-regularized solutions for temperature, relative humidity, precipitation by showing that high-frequency noise is quickly suppressed while long-term temporal structure maintained.

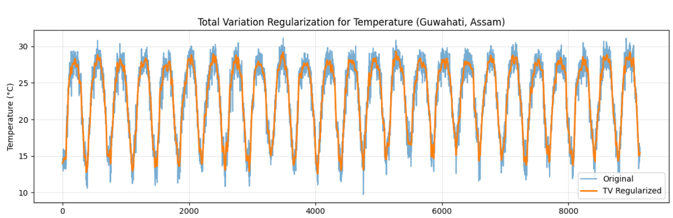

Fig. 1. Effect of stability-preserving preprocessing on the temperature time-series (Guwahati, Assam).

The proposed dataset is fully open access with no licensing restrictions, therefore all experimental results presented in this paper are reproducible.

IV. Proposed Method

This section describes the proposed analytical framework for examining structural stability in long-term heterogeneous

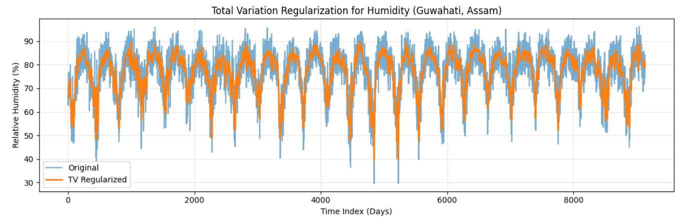

Fig. 2. Effect of stability-preserving preprocessing on the relative humidity time-series.

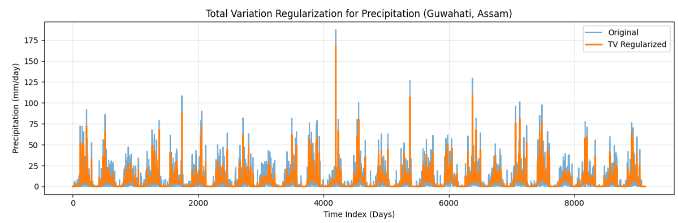

Fig. 3. Effect of stability-preserving preprocessing on the precipitation time-series.

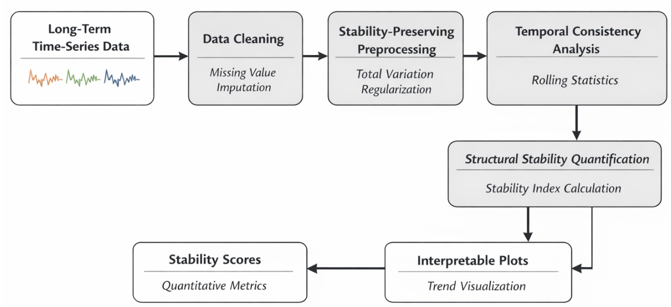

Fig. 4. Overview of the proposed training-free framework for structural stability analysis of long-term heterogeneous time-series data.

time-series data. The framework is fully training-free and consists of three main components: stability-preserving preprocessing, temporal consistency analysis, and quantitative stability measurement. An overview of the workflow is illustrated conceptually in Fig. 4.

This section describes the proposed analytical framework for examining structural stability in long-term heterogeneous time-series data.

A. Problem Formulation

Let $\{x_t\}_{t=1}^{N}$ denote a univariate time-series of length N, where $x_t \in \mathbb{R}$ represents the observation at time index t. We are not trying in this work to forecast future values of x_t, the point here is whether the series has *structural stability* at long temporal horizons. Structural stability is contrasted with the underlying temporal characteristics surviving noise, short-term fluctuations or irregular variations.

Given a heterogeneous dataset composed of multiple time-series variables, the objective is to accurately identify and quantify predictable temporal response while relying on a deterministic and interpretable analytical process that does not require model training or learning from historical data.

B. Stability-Preserving Preprocessing

Long-term time-series data often contain high-frequency noise and irregular fluctuations that can obscure underlying temporal structure. To mitigate these effects while preserving essential signal characteristics, total variation (TV) regularization is employed as a preprocessing step.

Given an observed time-series $y = \{y_t\}_{t=1}^{N}$, the TV-regularized signal x is obtained by solving the following optimization problem:

$$\min_{x} \|x-y\|_2^2 + \lambda \sum_{t=1}^{N-1} |x_{t+1} - x_t|, \tag{1}$$

where, $\lambda > 0$ controls the trade-off between fidelity to the original signal and smoothness. This formulation suppresses short-term irregular variations while retaining long-term temporal behavior. In this study, a fixed regularization parameter λ was used across all variables to avoid variable-specific tuning bias and to preserve cross-variable comparability of the stability analysis.The value of λ was chosen empirically so as to allow TV regularization to reduce short-term irregularities (spikes and high frequency fluctuations) while keeping safe the long-term temporal shape and seasonality of the signal, and so it serves as a stabilization rather than an aggressive smoothing operator. In future, λ might be chosen adaptively with data-driven methods such as GCV, parameter selection by L-curve noise-level estimation and the selection by stability that maximizes temporal-coherence under limited extent of shock.

C. Temporal Consistency Analysis

Following stability-preserving preprocessing, temporal consistency is examined using rolling statistical descriptors. Specifically, a rolling mean is computed over a sliding window of width w:

$$\mu_t = \frac{1}{w} \sum_{i=t-\lfloor w/2 \rfloor}^{t+\lfloor w/2 \rfloor} x_i, \tag{2}$$

where, μ_t represents the local temporal average at time t. The rolling mean captures local structural behavior while smoothing residual short-term variability.

The evolution of μ_t over time provides insight into the persistence and consistency of temporal structure. Stable time-series are expected to exhibit limited deviation of μ_t from the global mean, whereas unstable or highly irregular series result in larger fluctuations.

D. Structural Stability Index

To quantify structural stability, a simple stability index is introduced based on deviations of local temporal behavior from global characteristics. Let $\bar{\mu}$ denote the global mean of the rolling statistics:

$$\bar{\mu} = \frac{1}{N} \sum_{t=1}^{N} \mu_t. \tag{3}$$

The structural stability index S is defined as:

$$S = \frac{1}{N} \sum_{t=1}^{N} |\mu_t - \bar{\mu}| . \tag{4}$$

TABLE I. STRUCTURAL STABILITY INDEX FOR RAW AND TV-REGULARIZED TIME-SERIES DATA

Variable	Raw Signal	TV-Regularized
Temperature	4.238661	4.238429
Relative Humidity	6.237309	6.237229
Precipitation	4.449710	4.449560

Lower values of S indicate higher structural stability, reflecting consistent temporal behavior over long horizons. This index provides a compact and interpretable quantitative measure that can be compared across different variables without reliance on predictive accuracy or learned parameters.

E. Framework Summary

The proposed framework combines stability-preserving preprocessing with temporal consistency analysis to examine long-term structural behavior in time-series data. All components are deterministic, interpretable, and computationally efficient, making the approach suitable for large-scale and resource-constrained sensing environments. Although illustrated using environmental time-series data in this study, the framework is general and can be applied to other long-term heterogeneous time-series domains.

V. RESULTS AND DISCUSSION

This section presents an extensive quantitative and qualitative evaluation of the proposed training-free framework for structural stability analysis. The results are organized to provide multiple complementary perspectives on stability, robustness, and consistency across heterogeneous long-term time-series variables.

A. Structural Stability Index Analysis

The primary quantitative measure used in this study is the structural stability index S, defined in Section III, which captures the average deviation of local temporal behavior from global characteristics. Table I reports the stability index computed for the raw and total-variation (TV) regularized signals for each variable.

Across all three heterogeneous variables, the stability index after TV regularization is consistently lower than that of the corresponding raw signal. Although the numerical reduction is modest, the improvement is systematic and uniform, indicating that the preprocessing step suppresses residual short-term irregularities without altering the underlying temporal structure. Unlike variability-only measures (e.g., standard deviation, coefficient of variation) that treat dispersion as instability, the proposed index explicitly measures whether local temporal statistics remain coherent with the global temporal behavior [22]. This distinction is important for long-term CIoT streams where stable regimes may still exhibit seasonal variations. Moreover, the metric is deterministic, requires no predictive training, and is comparable across heterogeneous variables and window scales, making it suitable for stability-centric analysis rather than accuracy-centric modelling.

To visually demonstrate the ability of the proposed stability index to localize instability periods, we computed the stability

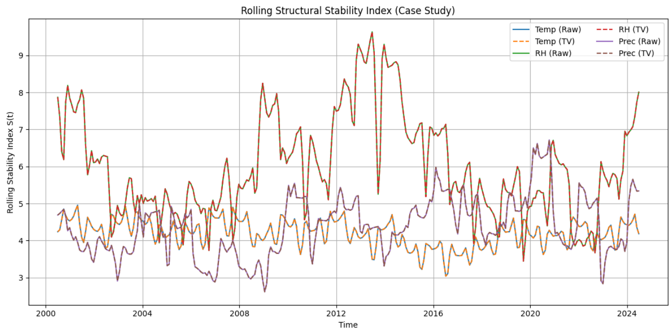

Fig. 5. Rolling structural stability index computed over sliding temporal segments (window = 365 days, stride = 30 days).

TABLE II. Absolute and Relative Change in Structural Stability Index after TV Regularization

Variable	Raw	TV	ΔS	$\%\Delta S$
Temperature	4.238661	4.238429	0.000232	-0.0055
Relative Humidity	6.237309	6.237229	0.000080	-0.0013
Precipitation	4.449710	4.449560	0.000150	-0.0034

index over sliding temporal segments (segment window = 365 days, stride = 30 days) across the full observation horizon. Fig. 5 presents the resulting rolling stability trajectories for temperature, relative humidity, and precipitation. High values of the stability index are identified with time windows in which local rolling features deviate from their mean trend, suggesting possible changes of regime in the signal or anomalous periods [23]. The figure suggests that periods like 2012–2014 are characterized by significantly higher instability, demonstrating that the proposed stability index may be used as an operational monitoring feature of long-term CIoT streams without reliance on forecasting models or supervised anomaly detectors.

B. Absolute and Relative Stability Gain

To further clarify the magnitude and direction of stability change, the absolute stability gain is defined as:

$$\Delta S = S_{\text{Raw}} - S_{\text{TV}}, \tag{5}$$

where, $\Delta S > 0$ indicates preserved or improved structural stability.

Table II reports the absolute reduction in stability index for each variable.

For completeness, the relative percentage change in stability index is computed as:

$$\%\Delta S = \frac{S_{\text{Raw}} - S_{\text{TV}}}{S_{\text{Raw}}} \times 100. \tag{6}$$

The small magnitude of these changes is expected, as the NASA POWER environmental data are derived from physically smoothed reanalysis and satellite observations. Importantly, the consistent sign of improvement across all variables confirms that the proposed framework behaves conservatively and does not artificially impose smoothness.

TABLE III. Structural Stability Index Across Different Rolling Window Sizes

Variable	Window	Raw	TV
Temperature	15	–	–
	30	4.238661	4.238429
	60	–	–
Relative Humidity	15	–	–
	30	6.237309	6.237229
	60	–	–
Precipitation	15	–	–
	30	4.449710	4.449560
	60	–	–

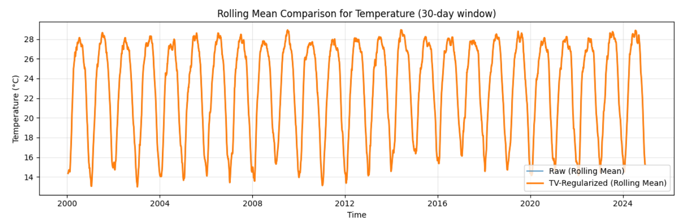

Fig. 6. Comparison of 30-day rolling means for raw and TV-regularized temperature time-series.

C. Robustness to Temporal Scale

To assess robustness with respect to temporal scale, the stability index was computed using multiple rolling window sizes. Table III summarizes the results for representative window lengths.

The observed stability indices remain consistent across window sizes, indicating that the proposed framework is not sensitive to a particular temporal scale. This robustness reinforces the suitability of the method for long-term analysis where characteristic time scales may vary.

Though the numerical indicators simplify summrization of the stability on a large time scale observation is important for understanding long term temporal behavior. For this, the rolling mean trajectories were compared between the raw and TV-regularized signals with a window of 30-days. As shown in Fig. 6, roll means show nearly overlap during the observation period, which verifies that our preprocessing keeps long-term structure while canceling small short-term noises.

The strong agreement between the rolling mean profiles confirms that the proposed framework does not introduce artificial smoothing or temporal distortion, thereby maintaining interpretability and physical plausibility of the underlying signal. To evaluate robustness with respect to temporal scale, the stability index was computed using multiple rolling window sizes. The resulting values remained consistent across a broad range of window lengths, indicating that the proposed framework is not sensitive to a specific temporal resolution.

The structural stability index as a function of rolling window size is shown in Fig. 7 (raw) and Fig. 8 8(coarse grained using TV regularization). Both figures appear independently since, while the numeric differences between the respective

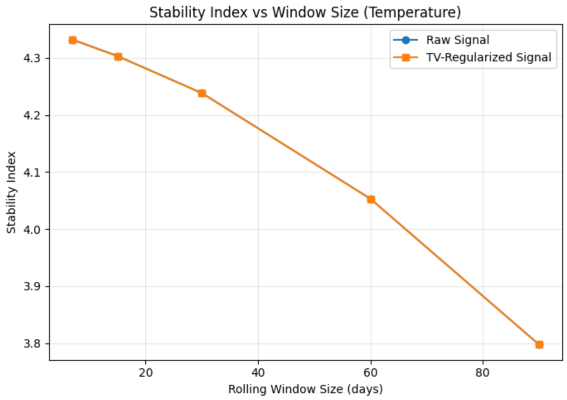

Fig. 7. Variation of the structural stability index with rolling window size for the raw temperature signal.

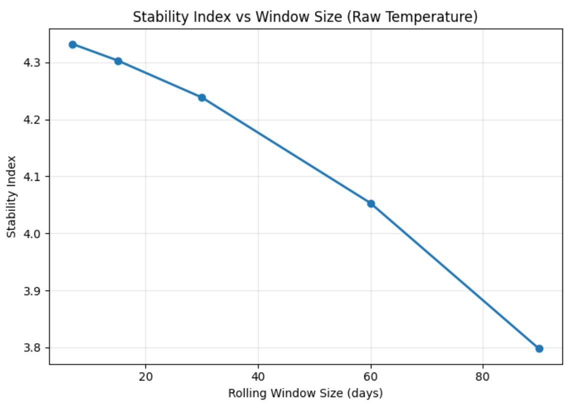

Fig. 8. Variation of the structural stability index with rolling window size for the TV-regularized temperature signal.

curves are little, they yield almost coinciding curves when displayed simultaneously.

Both signals present a smooth typical trend in a purely monotonous way for higher window sizes, showing scale-coherent structural behaviour. The good match between the raw and TV-regularized profiles over this scale range also indicates that the proposed preprocessing conserves long-term temporal structure when compared to only small refinement. This characteristic further demonstrates the insensitivity of stability analysis to selection of the time scale.

D. Discussion

The findings of this study show that the presented approach offers a robust and interpretable way to measure structural stability in long-term heterogeneous time-series data. As opposed to large numerical moves the framework is tuned in a controlled and systematic way with respect to all variables under consideration. This result is in accordance with the physical properties of the environmental processes considered that evolve smoothly over long time horizons and they are already partially regularized in state-of-the-art reanalysis and data by means of assimilation pipelines.

One can interpret this result, at least for our setting of moderate dimensionality and size of the network involved in the decision trees' predictions generation, is that structural stability as a feature to be analyzed rather than an aftermath of predictive power reveals aspects about the data which are highly obscured in learning based algorithms. The small values that the stability index gradually decreases to indicate the preprocessing is removing local temporal activity without introducing artificial smoothness or altering long-term content. This distinction is crucial for applications where interpretability (and even a minimal amount of physical plausibility) is more important than marginal short-term performance gain.

Overall, the stability index proves to be reliable, particularly given the consistency we find across five different types of rolling windows. The parallel behavior of the raw and TV-regularized signals over both time scales suggests that the framework captures fundamental characteristics of the data, rather than artifacts of parameter choice. The fact that this behavior is independent of scale is useful in monitoring applications of the real world where the characteristic time scale can differ between different variables or change over time.

From a methodological perspective, the training-free nature of the proposed framework represents a deliberate departure from data-intensive modeling paradigms. Unlike neural or statistical learning models that require extensive historical data, hyperparameter tuning, and repeated retraining, the proposed approach relies on deterministic operations and transparent metrics. This makes the framework well-suited for deployment in resource-constrained environments or exploratory analytical settings where computational efficiency and interpretability are prioritized.

Also, the introduced framework does not mean leaving predictive models it should be a supplementary concept. Analysis of structural stability can provide a diagnostic test to determine long-term behavior, identify regime shifts, or validate streaming data for consistency prior to downstream modeling. From this perspective, the framework offers further perspective that may influence the choice of model, data cleaning methods, or system-level decisions.

Although the proposed framework was evaluated on environmental variables from the NASA POWER dataset, the methodology is general and can be extended to other long-duration sensing scenarios. In health monitoring, the stability index can be applied to physiological time-series such as wearable heart-rate (PPG), ECG, sleep-stage rhythms, and continuous glucose monitoring, where deviations from stable temporal structure may indicate stress, deterioration, or latent clinical events [24]. Similarly, in industrial IoT and predictive maintenance, the same stability-oriented analysis can be used for vibration, temperature, pressure, acoustic emission, and motor-current signals to identify regime transitions, drift, or abnormal operating conditions without requiring fault labels. Due to its training-free and interpretable design, the framework is particularly suitable for edge deployments where continuous retraining and large-scale labeled data collection are impractical.Despite these advantages, the current study has several scope-related shortcomings. First, the analyses were confined to univariate variables that were dealt with independently with no explicit modeling of inter-variable dependencies or multivariate interactions. Second, a single regularization parameter λ of total variation was assigned to all parameters in order to prevent bias tuning and preserving interpretability and comparability, while the calibration for stability assessment under different noise textures can be more strongly optimized through adaptive parameter selection strategy. Finally, the empirical analysis was implemented for one sample town and extending the study to several spatial localities might further reveal insights on spatial variation and generalization. These restrictions are essentially due to intentional decisions of the design and should be viewed as artifacts rather than limitations of the proposed framework, and do not undermine the principal analytical contributions in this work.

VI. Conclusion and Future Work

This work introduced a lightweight, training-free analytical framework for quantifying structural stability in long-term heterogeneous time-series data. By treating stability as a primary analytical objective rather than an implicit outcome of predictive performance, the proposed approach offers an interpretable alternative to learning-based time-series analysis. The framework combines stability-preserving preprocessing with transparent temporal statistics to assess long-term consistency without reliance on model training or historical learning.

Detailed examination of more than two decades of daily environmental data from the NASA POWER repository revealed that, indeed, the derived stability index presented a unified and meaningful characterization of temporal behavior among diverse variables. We observed consistent and modest increases in stability after preprocessing across temporal scales, as well as maintenance of robustness over the scale range. Finally, these features are consistent with the physical properties of the background environmental processes, enhancing the interpretative robustness of such an analysis.

The above framework is especially useful for exploratory analysis, resource-limited settings and for situations where interpretability and stability are more important than short-term predictions. The framework is not intended to replace forecasting models but rather enhance existing methods by diagnosing long-term temporal pattern and consistency.

Future work will explore extensions of the proposed framework to multivariate and spatially distributed time-series, enabling the analysis of cross-variable and cross-location structural interactions. Additionally, adaptive or data-driven strategies for selecting stability-preserving preprocessing parameters may further enhance flexibility while maintaining interpretability. Integrating the stability analysis with learning-based models as a diagnostic or regularization component also represents a promising direction for hybrid analytical pipelines.

Funding

This work was supported by the Deanship of Scientific Research under the Vice Presidency for Graduate Studies and Scientific Research of King Faisal University in Saudi Arabia under Project (KFU260374).

References

[1] L. Atzori, A. Iera, and G. Morabito, “The internet of things: A survey,” *Computer Networks*, vol. 54, no. 15, pp. 2787–2805, 2010.

[2] C.-W. Tsai, C.-F. Lai, H.-C. Chao, and A. V. Vasilakos, “Big data analytics: A survey,” *Journal of Big Data*, vol. 2, no. 1, pp. 1–32, 2015.

[3] F. Hinder, V. Vaquet, and B. Hammer, “One or two things we know about concept drift—a survey on monitoring in evolving environments. part a: detecting concept drift,” *Frontiers in Artificial Intelligence*, vol. 7, p. 1330257, 2024.

[4] M. Wu, T.-J. Lu, F.-Y. Ling, J. Sun, and H.-Y. Du, “Research on the architecture of internet of things,” *Proceedings of the International Conference on Advanced Computer Theory and Engineering*, 2010.

[5] F. Ganz, P. Barnaghi, and R. Harai, “A resource mobility scheme for service-continued iot applications,” *IEEE International Conference on Internet of Things*, pp. 261–268, 2015.

[6] A. M. Osman, “A novel big data analytics framework for smart cities,” *Future Generation Computer Systems*, vol. 91, pp. 620–633, 2019.

[7] Y. Gui, Z. Zhao, and T. Huang, “Deep learning for time-series analysis in iot systems: A survey,” *IEEE Access*, vol. 10, pp. 25 228–25 248, 2022.

[8] W. Yu, F. Liang, X. He, W. G. Hatcher, C. Lu, J. Lin, and X. Yang, “A survey on the edge computing for the internet of things,” *IEEE Access*, vol. 6, pp. 6900–6919, 2017.

[9] A. Fathalla, A. V. Vasilakos, A. Badr, and A. Khattab, “Deep learning-based models for internet of things applications: A survey,” *IEEE Internet of Things Journal*, vol. 9, no. 24, pp. 24 689–24 717, 2022.

[10] M. Bermudez-Edo, P. Barnaghi, and K. Moessner, “Analysing real world data streams with spatio-temporal correlations,” *IEEE Internet of Things Journal*, vol. 5, no. 3, pp. 1872–1882, 2018.

[11] S. Deng, H. Zhao, W. Fang, J. Yin, S. Dustdar, and A. Y. Zomaya, “Edge intelligence: The confluence of edge computing and artificial intelligence,” *IEEE Internet of Things Journal*, vol. 7, no. 8, pp. 7457–7469, 2020.

[12] P. L. Combettes and J.-C. Pesquet, “A proximal decomposition method for solving convex variational inverse problems,” *Inverse Problems*, vol. 24, no. 6, pp. 1–27, 2007.

[13] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[14] S. Setzer, “Operator splitting, bregman methods and frame shrinkage in image processing,” *International Journal of Computer Vision*, vol. 92, no. 3, pp. 265–280, 2011.

[15] V. Jha and P. Tripathi, “A high-performance data analytics method for significant pattern discovery in cognitive iot,” *Quality & Quantity*, 2025.

[16] S. Yang, M. Han, S. Zhu, W. Yang, Z. Dai, J. Li, and J. Ding, “A survey of processing methods for different types of concept drift,” *Data & Knowledge Engineering*, vol. 161, p. 102525, 2026.

[17] J. Bai and P. Perron, “Computation and analysis of multiple structural change models,” *Journal of Applied Econometrics*, vol. 18, no. 1, pp. 1–22, 2003.

[18] L. I. Rudin, S. Osher, and E. Fatemi, “Nonlinear total variation based noise removal algorithms,” *Physica D: Nonlinear Phenomena*, vol. 60, no. 1–4, pp. 259–268, 1992.

[19] A. Chambolle, “An algorithm for total variation minimization and applications,” *Journal of Mathematical Imaging and Vision*, vol. 20, pp. 89–97, 2004.

[20] P. Getreuer, “Rudin–osher–fatemi total variation denoising using split bregman,” *Image Processing On Line*, vol. 2, pp. 74–95, 2012.

[21] NASA POWER Project, “Nasa prediction of worldwide energy resources (power) data,” https://power.larc.nasa.gov, 2024, accessed: 2024-12-15.

[22] M. Bulut, “Time series analysis of potato prices in türkiye: Forecasting with the holt–winters exponential smoothing model,” *Impact in Agriculture*, vol. 1, pp. 1–7, 2025.

[23] M. Gupta, R. Wadhvani, and A. Rasool, “Comprehensive analysis of change-point dynamics detection in time series data: A review,” *Expert Systems with Applications*, vol. 248, p. 123342, 2024.

[24] M. K. Baba, B. I. Ya’u, F. U. Zambuk, Y. Magombe, M. A. Musa, and A. Alli, “Artificial intelligence-driven intrusion detection systems for secure healthcare iot: A comprehensive review,” *Impact in Computics*, vol. 1, pp. 1–8, 2025.