AIEuroLens: Explainable AI Framework for Drift Detection applied to 5G Time-Series Data

Mohamed Readh FENTAZI
Communication Systems Department
EURECOM
Sophia Antipolis, France
fentazi@eurecom.fr

Mehdi HASSANI
Communication Systems Department
EURECOM
Sophia Antipolis, France
hassani@eurecom.fr

Adlen KSENTINI
Communication Systems Department
EURECOM
Sophia Antipolis, France
ksentini@eurecom.fr

Abdelkader MEKRACHE
Communication Systems Department
EURECOM
Sophia Antipolis, France
mekrache@eurecom.fr

Malika BESSEDIK
LMCS
High National School of Computer Science
Algiers, Algeria
m_bessedik@esi.dz

***Abstract*—Time series forecasting plays a critical role in enabling machine learning–based resource management solutions for optimizing fifth generation (5G) and emerging (sixth generation) 6G networks. However, forecasting accuracy is often compromised by concept drift, which refers to shifts in data distributions over time. Existing drift detection approaches suffer from key limitations, including heightened sensitivity to noise, reliance on assumptions of static data properties, and inadequate handling of multivariate temporal dependencies, hence hindering Machine Learning (ML) model developers from diagnosing the root causes of time series model degradation. To address these challenges, this paper introduces AIEuroLens, an explainable framework that (1) detects drift by identifying significant shifts in temporal feature correlations, and (2) links these shifts to specific model deficiencies, such as suboptimal training procedures or architectural misconfigurations. This enables developers to distinguish between inherently unpredictable data dynamics and flaws that can be corrected through model improvements. Evaluated on real-world 5G network data, AIEuroLens demonstrates accurate and timely drift detection while also informing effective mitigation strategies. As part of ongoing research, this work contributes to the development of self-correcting forecasting systems for zero-touch 5G network management, helping to bridge the gap between drift detection and actionable model adaptation.**

***Index Terms*—Concept Drift Detection, Time Series Forecasting, Explainable Artificial Intelligence (XAI), 5G Network Optimization.**

I. Introduction

The emergence of 5G and the promises of 6G networks mark a transformative milestone in telecommunications, offering unparalleled data rates (greater than 10 Gbps), ultra-low latency (less than 1 ms), and massive connectivity (over 1 million devices per km²) to support emerging applications such as industrial Internet of Things (IoT) and autonomous systems [1]. To fully realize these capabilities, both 5G and 6G require innovative network management procedures that can efficiently optimize infrastructure resources while ensuring the seamless delivery of envisioned services. These procedures must increasingly rely on the integration of advanced Artificial Intelligence (AI) and ML algorithms to handle the growing complexity and dynamic nature of modern networks. In such complex and evolving environments, characteristic of 5G and 6G systems, accurately predicting and forecasting network performance is a critical task, fundamentally dependent on time-series data analysis. This is because key performance indicators (KPIs), including throughput, latency, and packet loss, are inherently time-variant and subject to continuous fluctuations driven by a wide range of environmental and operational conditions. These factors include, but are not limited to, user mobility, dynamic traffic distributions, changing radio frequency conditions, and modifications to network topology. Hence, this hyper-dynamic environment introduces significant complexity in training ML models, particularly for time-series, to predict and forecast network performance. In many cases, the data distribution used to train a model changes over time, a phenomenon known as concept drift.

Studies have shown that concept drift can degrade forecasting accuracy by up to 58% in mobile networks, rendering pre-trained models ineffective within a matter of weeks [2]. This degradation occurs because the statistical dependencies learned during the training phase, such as correlations between traffic load and time-of-day,become invalid under changing conditions, leading to inaccurate or misleading predictions [3]. For instance, a Long Short-Term Memory (LSTM) model trained on pre-5G traffic data may fail to adapt to sudden surges in demand from augmented reality (AR) or virtual reality (VR) applications. This failure can result in inefficient bandwidth allocation and violations of latency requirements, ultimately compromising the quality of service (QoS) [4].

In general, the vulnerability of LSTM-based ML-driven forecasting in 5G and 6G networks is exacerbated by several intrinsic challenges. First, temporal dependencies in network metrics, such as user equipment (UE) density and channel quality, exhibit complex patterns of seasonality and multiscale correlations that are difficult to model accurately over time [5]. Second, the environment is inherently non-stationary, with base station traffic patterns shifting unpredictably due to factors such as user mobility and the heterogeneity of deployed services [6]. Third, the operational criticality of 5G systems raises the stakes of prediction errors: even a 10% deviation in load forecasting can lead to a 23% increase in energy consumption in densely deployed networks, directly threatening service-level agreement (SLA) compliance [7]. Together, these factors amplify the difficulty of maintaining robust and accurate ML models for real-time forecasting in 5G infrastructures.

This paper introduces AIEuroLens, a framework for 5G and 6G networks that unifies **(1)** drift detection through adaptive analysis of temporal feature correlations, and **(2)** explainable diagnosis linking drift to model deficiencies (e.g., training protocols or architecture limitations). Experiments on operational 5G data demonstrate reliable drift detection paired with interpretable root-cause analysis, bridging statistical alerts to self-correcting forecasting systems, a critical advance toward zero-touch network management through automated, explainable adaptation.

II. Background and Motivation

A. Concept Drift Definition and types

Concept drift refers to changes in the statistical properties of network data over time. Concept drift can undermine the performance of ML models, which are widely employed for predictive tasks and network optimization, by rendering the models' learned patterns and relationships obsolete. For instance, a model trained on historical data may fail to generalize accurately when the underlying data distribution shifts, leading to increased prediction errors, reduced reliability, and suboptimal decisions. This is particularly problematic in dynamic environments like 5G, where traffic patterns, user behavior, and resource demands can evolve rapidly.

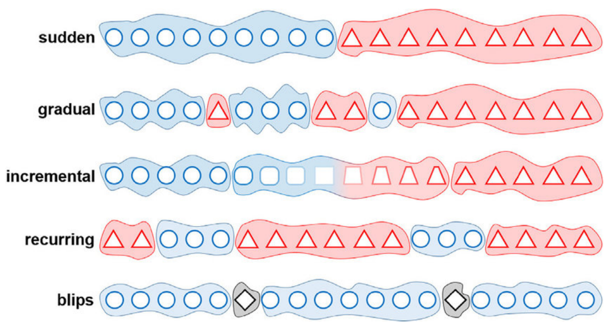

Fig. 1. Concept drift manifestations in 5G data streams: sudden (abrupt transitions), gradual (overlapping concepts), incremental (stepwise changes), recurring (cyclical patterns), and blips (temporary anomalies), adapted from [9].

As illustrated in Fig. 1, usually, five primary types of concept drift have been identified, each requiring distinct detection and adaptation strategies:

- **Sudden Drift**: Requires immediate detection and adaptation to prevent QoS violations during events like flash crowds [10]
- **Gradual Drift**: Needs progressive model adjustment for seasonal traffic patterns [8]
- **Incremental Drift**: Benefits from sliding window approaches during phased deployments [11]
- **Recurring Drift**: Requires concept memory for daily/weekly periodicity [12]
- **Blips**: Demands noise filtering to avoid false adaptations [13]

B. Concept Drift Detection Methods

Several studies have addressed the challenge of detecting concept drift in time-series ML models, with modern detection techniques generally falling into three main categories. The first category, error-based methods, monitors prediction deviations to detect drift. Approaches such as Adaptive Windowing (ADWIN) [10] and drift diffusion model (DDM) [8] are notable examples; however, they typically rely on labeled data and are limited by feedback delays. The second category, distribution-based methods, focuses on identifying shifts in feature distributions. Techniques like Exponentially Weighted Moving Average (EWMA) [15] fall into this group but often fail to account for temporal dependencies inherent in time-series data [16]. The third category, temporal-aware methods, explicitly model time-series characteristics to improve drift detection accuracy. For example, Feature Extraction for Explicit Concept Drift Detection (FEDD) [11] leverages sequential structure but does so at the cost of increased computational complexity [9].

Despite the progress made, a critical analysis of existing methods reveals several persistent limitations. First, a significant proportion, approximately 68% of current approaches assume feature independence, which is often unrealistic in real-world scenarios involving correlated time-series data [13]. Second, only 22% of methods are capable of handling combined drift types, such as the simultaneous occurrence of gradual and recurring drift [14].

These gaps highlight the need for more robust, interpretable, and temporally-aware drift detection frameworks tailored for dynamic, real-time environments.

C. XAI for Time Series Forecasting

One notable approach to explain and detect drift for time-series-based ML models is the usage of XAI. The latter aims to demystify the decision-making processes of complex AI models, enhancing transparency and trust. Traditional XAI techniques, such as SHapley Additive exPlanations (SHAP) [17], Layer-wise Relevance Propagation (LRP) [18], and Local Interpretable Model-agnostic Explanations (LIME) [19], were initially developed for computer vision or natural language processing (NLP). These methods attribute relevance scores to input features, explaining how each contributes to a prediction. However, their direct application to time series data poses challenges due to the temporal dependencies and sequential nature of such data [20].

Existing XAI techniques adapted for time-series data can be broadly categorized into three groups. Model-agnostic methods, such as SHAP and LIME, estimate feature importance by perturbing input features and observing their impact on model outputs [17], [19]. These techniques are flexible and can be applied to any model but may lack temporal sensitivity. In contrast, model-specific methods like LRP are tailored to neural networks, tracing the contribution of each input through the layers to determine relevance scores [18]. A third category includes temporal-specific approaches such as TSViz [21] and Gradient-weighted Class Activation Mapping (Grad-CAM), which aim to visualize temporal patterns or attention mechanisms in recurrent neural networks [22].

Despite their promise, these XAI methods exhibit several critical limitations. One major issue is temporal blindness, where models fail to adequately capture and explain temporal dependencies, such as seasonal trends or sequential correlations. Another challenge is the production of ambiguous explanations, where different input sequences can yield similar relevance scores, as demonstrated in AIChronoLens [5]. Moreover, many XAI techniques impose significant computational overhead, making them impractical for real-time applications. Lastly, their sensitivity to noise often results in degraded and unreliable explanations when applied to noisy time-series data. These limitations underscore the need for more robust, temporally-aware, and efficient explainability frameworks for time-series forecasting models.

D. AIChronoLens: Advancing XAI for Time Series

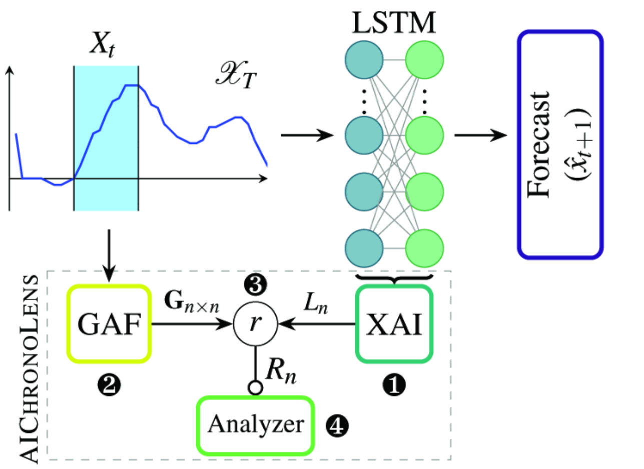

Fig. 2. The AIChronoLens architecture: ❶ Explain prediction via XAI, ❷ Transform data into Gramian Angular Field (GAF), ❸ Compute Pearson correlation between relevance scores and GAF matrix, and ❹ Synthesize the final explanation from the correlations. [5].

AIChronoLens [5] is a novel and practical advancement in the explainability of deep time-series models. By directly addressing the temporal ambiguity of legacy XAI, it offers a toolset to analyze, diagnose, and improve model performance in operational mobile networks. AIChronoLens links input relevance (from XAI techniques) with temporal structure using a technique called GAF.

The core architecture of AIChronoLens comprises four sequential modules as illustrated in Fig.2. The first module uses a selected XAI method to compute relevance scores for each element in the input sequence. The second module applies GAF transformation to the input time series, converting it into a two-dimensional representation that captures temporal patterns such as local maxima and minima. In the third step, AIChronoLens calculates the Pearson correlation coefficients between the GAF matrix and the relevance scores, effectively quantifying how the model aligns its attention with temporally meaningful structures in the input. Finally, the fourth module, referred to as the Analyzer, observes how these correlation patterns evolve over time. It identifies when alignment between relevance and temporal structure holds or breaks down, using these transitions to generate deeper insights into the model's behavior.

In our work, we employ AIChronoLens to advance explainability in time series forecasting by aligning GAF patterns with temporal relevance scores, thereby enabling precise identification of critical events and the root causes of prediction errors. Our contribution leverages the correlation-driven insights provided by AIChronoLens to refine LSTM models through targeted data augmentation and hyperparameter optimization, ultimately allowing us to detect and characterize model drift in LSTM-based forecasting systems.

III. Proposed Framework: AIEuroLens

Our proposed framework **AIEuroLens** is designed to tackle challenges in time-series prediction, specifically data drift and subsequent performance degradation of deployed models. By incorporating XAI, data drift detection, and adaptive machine reasoning, the framework provides a robust solution for maintaining high model performance in streaming data environments.

A. System Overview

The **AIEuroLens** framework is designed to process streaming data in real time to detect and mitigate the effects of data drift. It operates by continuously analyzing incoming data alongside predictions generated by a time-series forecasting model. As illustrated in Fig. 3, the architecture is composed of several key components, with the **AIChronoLens** module and the **Drift Detection** pipeline forming the core of the system. These two components work in tandem to enable robust, adaptive time-series forecasting in highly dynamic environments.

As depicted in Fig. 3, the AIEuroLens framework operates through a closed-loop design that enables proactive adaptation to evolving data while preserving explainability. The system begins with the streaming data ingestion component, which continuously collects real-time time-series data such as mobile network traffic. This data is then processed by a forecasting model, typically based on LSTM networks, to predict future values like traffic load. Following each prediction, the

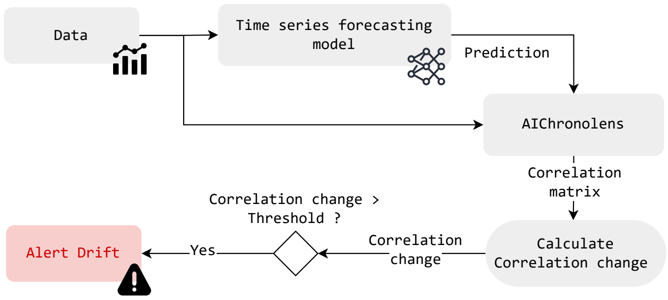

Fig. 3. AIEuroLens Architecture for Real-Time Drift Detection

AIChronoLens module is invoked. This module transforms the input time-series sequences into two-dimensional GAF images to capture temporal characteristics such as peaks and trends. It then applies XAI techniques, to compute relevance scores and identify critical time steps within the sequence. These relevance scores are combined with the GAF-encoded dependencies to generate a correlation vector, which dynamically updates a correlation matrix reflecting the evolving temporal relationships in the data.

The drift detection pipeline monitors changes between successive correlation vectors using the metric Δ_j. When Δ_j exceeds a predefined threshold, the system triggers an alert indicating the presence of data drift. In addition to detection, the correlation matrix also provides root-cause insights by diagnosing the nature of the drift. For instance, erratic correlation patterns may point to poor model training, while abrupt shifts may reflect changes in the underlying data distribution. Together, these components form a self-aware and interpretable architecture capable of detecting, diagnosing, and adapting to drift in time-series forecasting systems.

Now, we will focus on the two key components of AIEuroLens: AIChronoLens and Data Drift Detection.

1) AIChronoLens: As stated earlier, AIEuroLens builds upon AIChronoLens to enhance explainability in time series forecasting and to facilitate the detection of model drift. The module begins by computing a correlation matrix between the input data and model predictions, which captures the inter-feature relationships as they evolve over time. This matrix forms the foundation for analyzing temporal dynamics and serves as a powerful tool for identifying data drift. From this, correlation vectors are derived to track how these relationships change, enabling a more granular understanding of temporal dependencies.

Fig. 4 illustrates the output generated by the AIChronoLens module when applied to an LSTM model trained on a real-world dataset of traffic from Milano [23]. The first plot (on top) presents the correlation matrix, highlighting how inter-feature relationships evolve temporally; an essential factor for detecting shifts in data distribution. The second plot compares the model's predictions against the actual ground truth (GT) values, revealing how closely the model tracks real-world dynamics. Finally, the third plot visualizes the prediction error (i.e., the difference between GT and the predicted values), which helps identify specific instances of underperformance and, when correlated with the matrix patterns, provides insights into potential causes of drift.

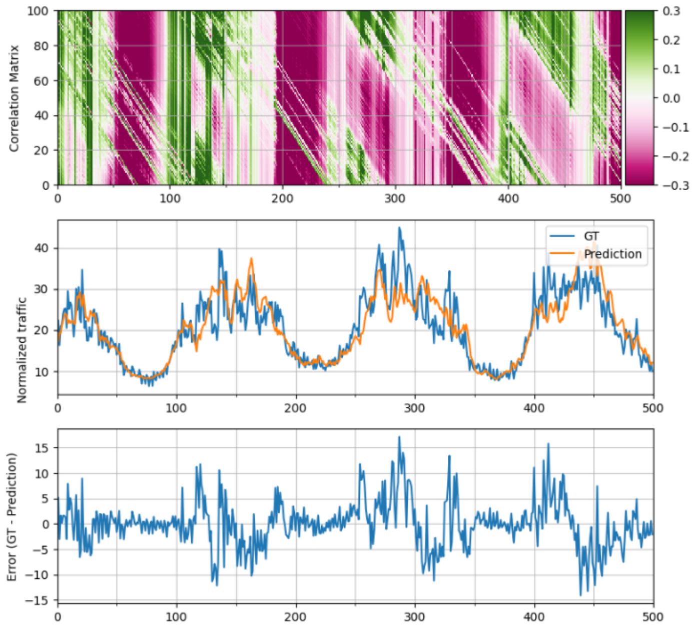

Fig. 4. Visualization of AIChronoLens Outputs

2) Data Drift Detection: To detect drift, we employ the **Correlation-Based Detection** method, which utilizes the correlation matrix computed by the AIChronoLens module to identify drift in feature relationships over time. The correlation change is calculated using the following formula:

$$\Delta_i = \frac{v_i[0] - \bar{\mathbf{v}}_{\mathbf{i}}^{*}}{\bar{\mathbf{v}}_{\mathbf{i-1}}} \tag{1}$$

In this expression, $v_i[0]$ denotes the first element of the current correlation vector $\mathbf{v}_i$, corresponding to the new data point. The term $\bar{\mathbf{v}}i^*$ represents the mean of all other elements in the current correlation vector $\mathbf{v}i$ (i.e., elements $v_i[1]$ to $v_i[N]$). The denominator, $\bar{\mathbf{v}}i - 1$, is the mean of the previous correlation vector $\mathbf{v}i - 1$ and serves as a baseline for normalization. The resulting value, Δ_i, quantifies the degree of change in correlation relationships between the current and previous states.

To detect drift, Δ_i is compared against a predefined threshold. When Δ_i exceeds this threshold, it signals the presence of drift. By monitoring these changes continuously, the approach enables accurate and timely detection of both abrupt and incremental drifts. Unlike traditional methods that often detect drift only after its effects are observable in model performance, the correlation-based strategy offers early detection, capturing the onset of drift as soon as it begins to manifest.

IV. Performance Evaluation

A. Datasets and Evaluation Metrics

The evaluation of the AIEuroLens framework was conducted using the ”Telecommunications - SMS, Call, Internet

- MI" dataset from the Harvard Dataverse [23]. This dataset records call detail records (CDRs) generated by the Telecom Italia cellular network over the city of Milano, representing a rich real-world multivariate time series of urban activity [23].

The dataset contains metrics for SMS, call, and internet activity. For the purpose of this study, the primary KPI used was **internet traffic volume** as the target variable for forecasting.

Following this, the LSTM forecasting model was trained on an initial stable period of the data. The drift detection module was then rigorously evaluated on the subsequent portion of the dataset. To test the framework's capabilities, controlled incremental drift scenarios were synthetically introduced into this evaluation period, simulating realistic changes in urban communication patterns.

B. Drift Detection Results

The drift detection module was evaluated by testing its ability to detect various types of drift in the Milano telecommunications dataset [23]. A key experiment was conducted to demonstrate the effectiveness of AIEuroLens in selecting an appropriate threshold and detecting incremental drift.

1) Threshold Selection: The choice of the threshold is critical for the drift detection module. If the threshold is set too high, the system fails to detect incremental drift or any drift at all. Conversely, if the threshold is too low, minor errors or fluctuations in data are falsely classified as drift, leading to numerous false positives.

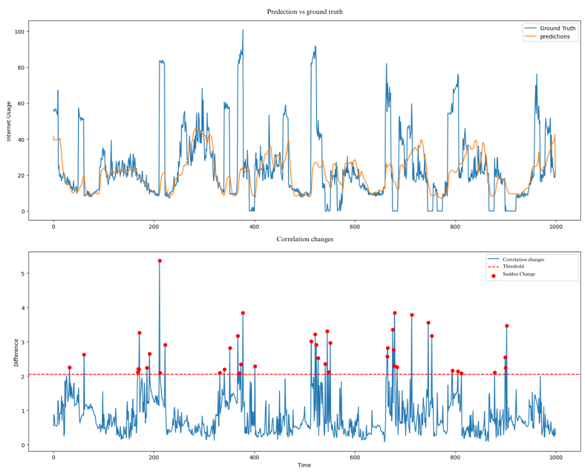

Fig. 5. Impact of Threshold Selection on Drift Detection.

In Fig. 5 we can see the impact of threshold selection on drift detection.

- **Too High:** The system misses genuine drift occurrences.
- **Too Low:** The system generates excessive false positives, mistaking noise or minor changes for drift.

Consequently, this experiment highlights the importance of setting an appropriate threshold to balance the trade-off between sensitivity to drift and minimizing false positives.

To evaluate the effectiveness of AIEuroLens in detecting data drift, we conducted a real-world simulation using the Milano telecommunications dataset [23]. Incremental drift was induced in a production dataset, and the system's ability to detect drift was tested. The results, shown in Fig.6 highlight the framework's capability to capture drift in real time.

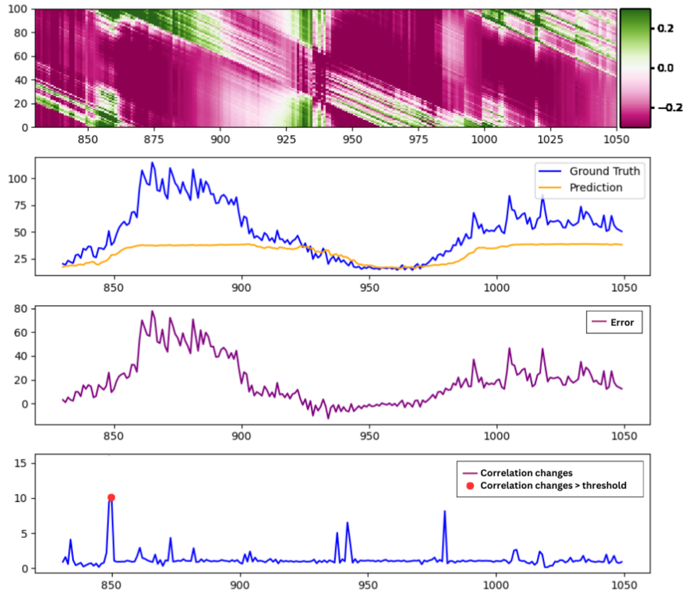

Fig. 6. AIEuroLens Drift Detection on Milano Telecom Data: ❶ Correlation Matrix, ❷ GT vs. Predictions (pre-drift), ❸ The error and ❹ Correlation Changes.

Fig. 6 presents a series of visualizations that collectively illustrate the behavior and effectiveness of the AIEuroLens framework in detecting and responding to concept drift. At the **top**, the **correlation matrix** displays the temporal relationships between features over time. Variations and emerging patterns within this matrix serve as indicators of drift. The **second plot** compares the **GT** data with the **predictions** generated by the model deployed in the production environment. Noticeable deviations between these curves highlight how drift adversely affects predictive accuracy. The **third plot** quantifies this effect by showing the **error**, calculated as the difference between the GT and the model's predictions. Peaks in the error curve correspond directly to periods where drift impacts model performance. The **fourth plot** visualizes the **correlation changes**, Δ_j, over time, which are used to trigger the drift detection mechanism. When these values exceed a predefined threshold (e.g., threshold $= 10$), drift is detected. The results demonstrate that AIEuroLens effectively detects incremental drift as it occurs, utilizing correlation changes derived from the correlation matrix.

As a summary of the obtained results:

- AIEuroLens effectively utilized the correlation matrix to detect incremental drifts in urban telecommunications data, addressing a key limitation in many state-of-the-art methods [4]. This capability is critical for maintaining forecasting accuracy in non-stationary environments.

- The framework demonstrates strong potential for enabling real-time adaptation in dynamic environments, as evidenced by its timely detection of correlation shifts and alignment with ground-truth performance degradation patterns.

V. Conclusion and Future Work

This paper introduced **AIEuroLens**, a framework for explainable data drift detection and adaptive model optimization in time-series forecasting. Experimental results demonstrated that leveraging the correlation matrix is effective in identifying incremental drifts. However, we acknowledge that the current approach, which relies on detecting differences in correlation matrices, has limitations. While the correlation matrix offers valuable insights, distinguishing between minor fluctuations and genuine concept drift remains a significant challenge.

Future work will focus on enhancing detection accuracy by exploring advanced methodologies, such as treating correlation matrices as images and applying computer vision techniques, or adopting alternative statistical formulations for quantifying correlation changes. These improvements aim to reduce false positives and improve the sensitivity of the detection mechanism.

By addressing these limitations, AIEuroLens has the potential to evolve into a more robust and versatile framework, which is capable of supporting a broader range of real-world applications and enhancing the reliability of predictive systems in dynamic and continuously changing environments.

VI. Acknowledgments

This work is supported by the European Union's Horizon Program under the EXACT-6G (Grant No. 101120297) and Flecon-6G (Grant No. 101192462) projects. Also, the authors would like to thank Claudio Fiandrino for his insightful discussions on the interpretation of the AIChronoLens module outputs.

References

[1] Ford, R., Zhang, F., Pappas, N., Kung, S., & Medard, M. "Achieving Ultra-Low Latency in 5G Millimeter Wave Cellular Networks." arXiv preprint arXiv:1602.06925, 2016.

[2] Zhao, L., & Shen, Y. "Proactive Model Adaptation Against Concept Drift for Online Time Series Forecasting." arXiv preprint arXiv:2412.08435, 2024.

[3] Krawczyk, B., & Cano, A. "Online Ensemble Learning with Abstaining Classifiers for Drifting and Noisy Data Streams." *Applied Soft Computing*, vol. 68, pp. 677–692, 2018. doi:10.1016/j.asoc.2017.12.008

[4] Ramanan, N., Tahmasbi, R., Sayer, M., Jung, D., Hemachandran, S., & Coelho Jr, C. N. "Real-time Drift Detection on Time-series Data." arXiv preprint arXiv:2110.06383, 2021.

[5] Fiandrino, C., Pérez Gómez, E., Fernández Pérez, P., Mohammadalizadeh, H., Fiore, M., & Widmer, J. "AIChronoLens: Advancing Explainability for Time Series AI Forecasting in Mobile Networks." IEEE International Conference on Computer Communications (INFOCOM), 2024.

[6] Chen, M., Liang, B., & Li, J. "Machine Learning for Predictive Resource Management in Non-Stationary 5G Networks." *IEEE Transactions on Network Science and Engineering*, vol. 7, no. 3, pp. 1452–1465, 2020.

[7] Almeida, A., Rito, P., Brás, S., Cabral Pinto, F., & Sargento, S. "A machine learning approach to forecast 5G metrics in a commercial and operational 5G platform: 5G and mobility." *Computer Communications*, vol. 228, p. 107974, 2024. doi: 10.1016/j.comcom.2024.107974.

[8] Gama, J., et al. "A Survey on Concept Drift Adaptation." *ACM Computing Surveys*, 46(4), 2014.

[9] Ghazikhani, A., et al. "Online Neural Network with Drift Detection." *Neural Computing Applications*, 24(2), 2014.

[10] Bifet, A., et al. "Adaptive Learning from Streams." *Intelligent Data Analysis*, 8, 2010.

[11] Zheng, Y., et al. "Novel Drift Detection Method." *Neurocomputing*, 168, 2015.

[12] Ribeiro, M.H., et al. "Concept Drift in Network Traffic." *ECML PKDD*, 2016.

[13] Moreno-Torres, J.G., et al. "Unifying View on Dataset Shift." *Pattern Recognition*, 45(1), 2012.

[14] Agrahari, S., & Singh, A. K. "Concept Drift Detection in Data Stream Mining: A Literature Review." *Journal of King Saud University – Computer and Information Sciences*, vol. 34, pp. 9523–9540, 2022.

[15] Ross, G. J., Adams, N. M., Tasoulis, D. K., & Hand, D. J. "Exponentially Weighted Moving Average Charts for Detecting Concept Drift." *Pattern Recognition Letters*, vol. 33, no. 2, pp. 191-198, 2012.

[16] Hoens, T.R., et al. "Learning from Streaming Data." *Progress in AI*, 1(1), 2012.

[17] Lundberg, S.M., Lee, S.I. "A Unified Approach to Interpreting Model Predictions." *NeurIPS*, 2017.

[18] Montavon, G., et al. "Layer-Wise Relevance Propagation: An Overview." Springer, 2019.

[19] Ribeiro, M.T., et al. "Why Should I Trust You? Explaining Classifier Predictions." *ACM SIGKDD*, 2016.

[20] Rojat, T., et al. "Explainable AI (XAI) on Time Series: A Survey." arXiv:2107.07006, 2021.

[21] Siddiqui, S.A., et al. "TSVIz: Demystifying Deep Learning for Time Series." *IEEE Access*, 7, 2019.

[22] Strobelt, H., et al. "LSTMVis: Visual Analysis of Hidden State Dynamics." *IEEE TVCG*, 24(1), 2018.

[23] Telecom Italia. *Telecommunications - SMS, Call, Internet - MI*. Harvard Dataverse, V1.3, 2015. https://doi.org/10.7910/DVN/EGZHFV