This is the author's accepted version of the article. The final version published by IEEE is Pablo Fernández Pérez, Claudio Fiandrino and Joerg Widmer, "Ph.D. Forum: Explainable AI for Time Series Analysis in 5G/6G Operations" *26th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, Forth Worth, Texas, USA, 2025, pp. TBD, doi: TBD.

Ph.D. Forum: Explainable AI for Time Series Analysis in 5G/6G Operations

Pablo Fernández Pérez
Supervised by: Claudio Fiandrino, and Joerg Widmer
IMDEA Networks Institute, Madrid, Spain
Email: pablo.fernandez@imdea.org

The emergence of 5G networks has transformed the mobile ecosystem. Higher access speeds have driven explosive growth in mobile traffic demand. 5G is projected to handle 75 % of all mobile traffic by 2029, compared to just 26 % at the end of 2024. Traffic forecasting at both base station and city levels has become crucial for network operators. It enables optimization across multiple domains including network deployment planning [1], and enhanced mobility management [2].

AI-based time series forecasting has advanced significantly. Transformer-based models have surpassed traditional statistical approaches across varied datasets.

These advanced forecasting models operate as black boxes, with decision-making processes that human observers cannot easily interpret. This fundamental lack of explainability creates significant barriers to their adoption in production networks. Network managers understandably hesitate to trust outputs from systems whose reasoning remains opaque. The concern extend beyond mere skepticism. Network engineers face substantial challenges when troubleshooting these models, and the opacity creates vulnerabilities to adversarial attacks. The research has demonstrated how carefully crafted perturbations—such as artificially increased load or targeted jamming—can remain undetectable to human monitors while significantly degrading predictor accuracy [3]. This situation highlights the critical need for developing robust explainability tools that help users understand and trust AI forecasting models. By making the decision-making processes transparent and interpretable, these tools bridge the gap between sophisticated algorithmic performance and practical implementation in production networks.

Legacy explainable AI techniques such as SHAP [4], and LIME [5] have been adapted or directly applied to time series analysis. However these approaches fundamentally fall short in providing truly useful explanations beyond basic input relevance assessments [6]. Their critical limitation lies in their inability to reveal the hidden causes of model errors that simultaneously address both the model's internal reasoning processes and the specific characteristics of current input data. These approaches also generate explanation object that can become impractical due to their excessive size, further limiting their real-world utility. For instance, in multi-horizon multivariate forecasting scenarios, the explanation object size becomes $L \in R^{N \times L \times C \times H \times C}$. Where N is the batch dimension, L is the length of the input sequence, C is the number of variates or channels and H is the prediction length.

To address these challenges, we introduced two novel explainability techniques: AICHRONOLENS [6] and CHRONOPROF [7]. ChronoProf specifically targets the explanation of AI/ML models operating on time series data, accommodating both classification and regression tasks. This technique explains how feature importance evolves dynamically over time, providing crucial insights into the model's decision-making process across varying observations. The tool enhances the interpretability of deep neural networks (DNNs) by leveraging linear model concepts through SHAP values. It computes virtual weights for each feature, analogous to linear model coefficients, enabling clear model explanations for specific inputs. Unlike LIME, these virtual weights are unique for each time step, ensuring temporal coherence across time steps.

On the other hand, AICHRONOLENS is specifically designed for univariate time series. Its architecture correlates SHAP values with Gramian Angular Fields to generate a vector of correlations that is used to identify points of interest in the time series. These points are crucial for detecting potential model errors and anomalies within the data.

Visual examples. Figure 1 illustrates a use case example of CHRONOPROF. The figure displays the distribution of virtual weights for two features, showing how they cluster around specific values that are well separated from each other. This clustering enhances confidence in feature importance analysis as a reliable information source. This phenomenon appears consistently across all tested models and contrasts with SHAP values, which consistently overlap. This separation makes it possible to explore the different decision boundaries established by each model and understand their decision-making processes while maintaining a global explanation of the task.

Fig. 2 shows qualitatively that there exist consecutive errors with high magnitude with a change of sign (*e.g.*, first positive than negative or vice versa). AICHRONOLENS identifies this behavior with triangles of positive or negative correlations over time that are interrupted by a full column with a weak correlation.

Limitations. AICHRONOLENS method is limited to univariate time series, which restricts its application to multivariate scenarios. Furthermore, when dealing with multi-horizon time series, explanations must be generated individually for

©2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

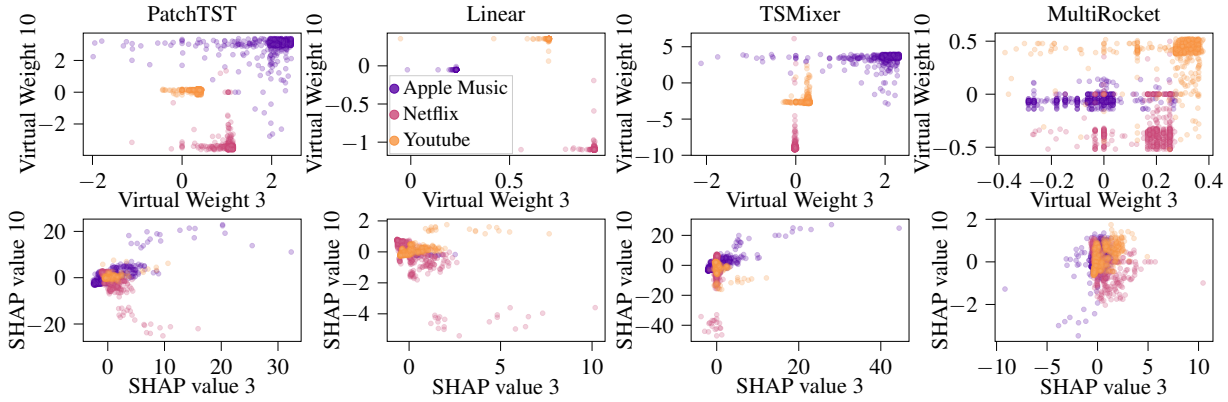


Fig. 1. Demonstrating the advantages of CHRONOPROF in network traffic classification applications. The inputs consist of one day long traffic time series with hourly granularity labeled by application. The plots emphasizes the contribution of two specific hours (features) from the lookback period, highlighting the differences between SHAP and virtual weights.

each prediction point, making the approach impractical for long-horizon tasks. This limitation further emphasizes the dimensionality challenges that explainers typically face.

CHRONOPROF can effectively address applications where dimensionality is reduced, such as computing the variance of virtual weights to observe dynamic changes over time. However, certain applications remain subject to dimensionality challenges, such as observing actual changes in virtual weights between two consecutive time steps. Additionally, explanations remain difficult for non-expert users to interpret, which hinders widespread adoption.

Future work. Exploring task-specific explanation techniques rather than developing universal explainers could prove valuable. Current approaches often attempt to create universal explanation methods, but tailoring explanations to particular domains or tasks might yield more insightful results. Additionally, investigating XAI applications in other domains presents interesting opportunities, such as physics-informed neural networks (PINNs) in network contexts and symbolic regression for generating analytical functions from data. PINNs, which incorporate physical laws into neural network architectures, could provide explanations that are inherently consistent with domain knowledge. Meanwhile, symbolic regression techniques could transform black-box models into analytical mathematical expressions, bridging the gap between performance and explainability in complex systems.

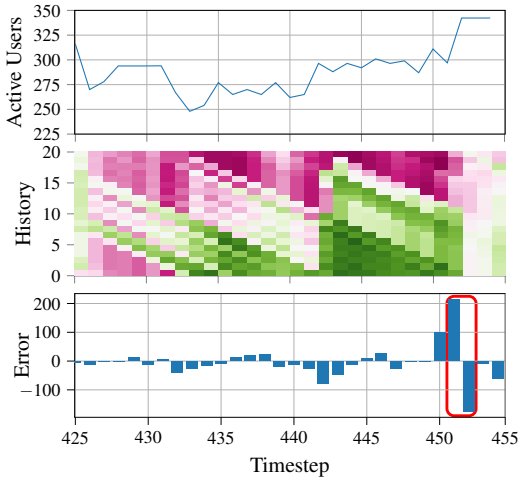


Fig. 2. Consecutive errors with high magnitude and sign change

ACKNOWLEDGMENTS

"This work is partially supported by bRAIN project PID2021-128250NB-I00 funded by MCIN/ AEI /10.13039/501100011033/ and the European Union ERDF A way of making Europe" and by the TUCAN6-CM regional project (TEC-2024/COM-460); P. Fernandéz received funding from Programa Investigo" (grant 2022-C23.I01.P03.S0020-0000038) funded by EU-NextGenerationEU and SEPE/PRTR."

REFERENCES

[1] P. D. Francesco *et al.*, "Assembling and using a cellular dataset for mobile network analysis and planning," *IEEE Trans. on Big Data*, vol. 4, no. 4, pp. 614–620, 2018.

[2] S. Zhao *et al.*, "Cellular network traffic prediction incorporating handover: A graph convolutional approach," in *Proc. of IEEE SECON*, 2020, pp. 1–9.

[3] D. Adesina *et al.*, "Adversarial machine learning in wireless communications using RF data: A review," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 1, pp. 77–100, 2023.

[4] S. M. Lundberg *et al.*, "A unified approach to interpreting model predictions," in *Proc. of NIPS*, 2017, pp. 4768–4777.

[5] M. T. Ribeiro *et al.*, ""Why Should I Trust You?": Explaining the predictions of any classifier," in *Proc. of ACM SIGKDD*, 2016, p. 1135–1144.

[6] C. Fiandrino *et al.*, "AIChronoLens: advancing explainability for time series AI forecasting in mobile networks," in *Proc. of IEEE INFOCOM*, 2024.

[7] P. Fernández Pérez *et al.*, "ChronoProf: profiling time series forecasters and classifiers in mobile networks with explainable AI," in *2025 IEEE 26th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, 2025.