Post-Doc Position in Explainable Artificial Intelligence over evolving IoT data streams

Institution: Cergy Paris University and ENSEA
Research Laboratory: ETIS lab, DATA&AI team
Location: Cergy-Pontoise, France
Duration: 18 months
Start Date: before October 2025

Position Overview:
We are excited to announce an opening for a Post-Doctoral researcher to join our dynamic research team DATA&AI (ex MIDI team) at the ETIS laboratory of the CY Cergy Paris University and the ENSEA school of engineering, member of the CNRS national research center. The successful candidate will contribute in developing **explainable methods** for **AI models trained** over **evolving IoT data streams (AIoT) for industrial applications such as predictive maintenance**. They will also participate in the project management and dissemination of research outcomes.

Subject:
Predictive maintenance (PDM) in industrial settings spans from identifying anomalies and categorizing failures in already observed data, to prognostically predicting the Remaining Useful Life (RUL) and the Failure Time (FT) of machines, appliances, etc., in the future[1]. Typically such predictive tasks are implemented using Deep Learning and/or statistical analysis techniques, which may be complex to interpret, while their performance is challenged by multiple sources of errors (sensors' tuning, aging, failures, etc.) and the non-stationary nature of IoT data streams. In particular, concept drifts can originate from the changes in the underlying data generation mechanism that reflects different states of the monitored system. In this context, the predictive performance of already trained models f(X,Y) (e.g., classification, regression) may start degrading after a certain point in time and hence models need to be adapted at the right frequency. However, not all types of changes in the joint probability distribution P(X,Y) have the same impact on model performance[2]. In this respect, we need to distinguish between changes in the posterior probability distribution P(Y|X) (i.e., Model drifts) between input features X and the target variable Y, from class-conditioned data distribution changes P(X|Y) (i.e., likelihood drifts) and changes in the distribution of input features P(X) (i.e., covariate drifts). Clearly, not all types of changes of the joint probability distribution $P(X, y)$ influence predictive models in the same way and hence, they require different mitigation actions[3]. In this project we are particularly interested in how data quality and concept drifts affect the performance of PDM tasks such RUL and FT.
Moreover, besides developing performant, robust, and stable FT and RUL prediction algorithms through time, we are also interested in enhancing interpretability of their results. On the one hand, engineers need to know the root causes for a predicted machine failure at a time t in the

future, so that they may take the best possible action towards preventing the failure to happen, or replace a machine in time before having to take the system down for replacement of the compromised machine. Such explanations should cover both the time parameter (why a failure will happen after a time interval) and the type of failure (why a specific type of error will happen). On the other hand, fine-grained explanations of the different types of concept drifts can guide data analysts to take *timely, and informed* actions for adapting the prediction algorithm to the observed concept drift. While explainability has been a major research interest in recent years[4], explanation methods for concept drift are still in their infancy. Some of the approaches aim for the detection and quantification of drift[5] [6], its localization in space[7] [8] or its visualization[9], while others focus on feature-wise representations of drift [10]. In this project, we aim to investigate actionable concept drift explanations, adding in the equation *weak and strong signals* for failure events. We believe that *concept drift explanations* constitute a form of *actionable* explanations responding to both aforementioned expert needs, and thus can be more valuable than standard feature importance explanations.

Context:

The position is funded by the prestigious EU Horizon PANDORA project, *A Comprehensive Framework enabling the Delivery of Trustworthy Datasets for Efficient AIoT Operation*. The main goal of the project is to contribute towards the creation of a dynamic AI pipeline in the context of IoT applications, focusing on i) the creation of synthetic but trustworthy data, and ii) on the development of AI algorithms for various tasks (from classic classification, regression and anomaly detection, to forecasting and predictive maintenance) taking into account the specific data characteristics of continuous data generation settings and leveraging the domain knowledge when available. Explainability is a fundamental property for rendering the systems reliable, as it can help i) enhance the acceptability of the system by the system experts and further aid in decision making, ii) optimize model performance (time and accuracy) by revealing the actual causes behind predictions, and iii) repair data acquisition processes or model training/updating by exposing errors and/or drifts that may arise through time. The project gathers over 20 academic and industrial partners, with real and challenging use-cases and thus provides a unique opportunity to contribute to cutting-edge research with significant real-world impact.

Responsibilities/Opportunities:

- Conduct high quality research in **Explainable AI in non-stationary settings**.
- Develop novel algorithms and methodologies for **Predictive Maintenance using IoT streams**.
- Publish in **top-tier conferences and journals**.
- Collaborate with an interdisciplinary team of researchers and industry partners.
- Participate in project meetings and contribute to the project's management.

Requirements:

- PhD in **Machine Learning, AI, Data Science, Statistics, or a related field**.
- Strong background in at least one of the following: **deep learning, continual learning, time series analysis, or predictive maintenance**.
- Experience with **explainable AI (XAI)** methods.
- Proficiency in **Python and relevant ML frameworks** (PyTorch, TensorFlow, Scikit-learn).
- Excellent publication record in **Data management, Artificial Intelligence, Machine Learning, or IoT applications**.

Contact Information:
For inquiries regarding the position, please contact:

Aikaterini Tzompanaki, Associate Professor, CY Cergy Paris University
Email: aikaterini.tzompanaki@cyu.fr
Website: https://perso.etis-lab.fr/tzompanaki/
Vassilis Christophides, Full Professor, ENSEA
Email: Vassilis.Christophides@ensea.fr
Website: https://www.etis-lab.fr/2022/01/13/vassilis-christophides/

How to Apply:
Interested candidates should send **in one pdf** their
- CV,
- motivation letter and/or research statement,
- publication list,
- at least one relevant publication, and
- two contacts for recommendations,

to aikaterini.tzompanaki@cyu.fr and Vassilis.Christophides@ensea.fr with the subject line "**Post-Doc Position in Explainable AIoT**."

Visit us online:
- PANDORA project: https://pandora-heu.eu/
- DATA&IA team in ETIS Lab: https://www.etis-lab.fr/midi/
- CY Cergy Paris University: https://www.cyu.fr/
- ENSEA: https://www.ensea.fr/en

References

[1] Pashami, S., Nowaczyk, S., Fan, Y., Jakubowski, J., Paiva, N., Davari, N., Bobek, S., Jamshidi, S., Sarmadi, H., Alabdallah, A. and Ribeiro, R.P., 2023. Explainable predictive maintenance. arXiv:2306.05120.
[2] Sijie Dong, Qitong Wang, Soror Sahri, Themis Palpanas, and Divesh Srivastava. 2024. Efficiently Mitigating the Impact of Data Drift on Machine Learning Pipelines. Proc. VLDB Endow. 17, 11 (July 2024), 3072–3081. https://doi.org/10.14778/3681954.3681984
[3] Lu, J., Liu, A., Dong, F., Gu, F., Gama, J., Zhang, G.: Learning under concept drift: A review. IEEE transactions on knowledge and data engineering 31(12), 2346–2363 (2018)
[4] C. Molnar, Interpretable Machine Learning, 2020, https://christophm.github.io/interpretable-ml-book/
[5] J. Lu, A. Liu, F. Dong, F. Gu, J. Gama, G. Zhang, Learning under concept drift: A review, IEEE TKDE (2018) http://dx.doi.org/10.1109/tkde.2018.2876857.
[6] G. Webb, L. Lee, B. Goethals, F. Petitjean, Analyzing concept drift and shift from sample data, Data Min. Knowl. Discov. 32 (2018) http://dx.doi.org/10.1007/s10618-018-0554-1.
[7] A. Liu, Y. Song, G. Zhang, J. Lu, Regional concept drift detection and density synchronized drift adaptation, in: IJCAI, 2017, http://dx.doi.org/10.24963/ijcai.2017/317.
[8] F. Hinder, V. Vaquet, J. Brinkrolf, A. Artelt, B. Hammer, Localization of concept drift: Identifying the drifting datapoints, in: 2022 International Joint Conference on Neural Networks, IJCNN, 2022, pp. 1–9, http://dx.doi.org/10.1109/IJCNN55064.2022.9892374.
[9] X. Wang, W. Chen, J. Xia, Z. Chen, D. Xu, X. Wu, M. Xu, T. Schreck, ConceptExplorer: Visual analysis of concept drifts in multi-source time-series data, in: 2020 IEEE Conference on Visual Analytics Science and Technology, VAST, 2020.
[10] G. Webb, L. Lee, B. Goethals, F. Petitjean, Analyzing concept drift and shift from sample data, Data Min. Knowl. Discov. 32 (2018) http://dx.doi.org/10.1007/s10618-018-0554-1.