Προτεινόμενα Θέματα Διπλωματικής Εργασίας

Από

Ηλία Αλεβίζο

1) **Neuro-symbolic complex event forecasting**

Abstract:

In traditional machine learning, the goal is to "predict" the output of a function on previously unseen input data. The input data need not necessarily have a temporal dimension and the term "prediction" refers to the output of the learned function on a new data point. On the other hand, the task of forecasting is to predict the temporally future output of some function or the occurrence of an event. Time is thus a crucial component for forecasting. What makes forecasting particularly challenging is the fact that, from the (current) timepoint where a forecast is produced until the (future) timepoint for which we try to make a forecast, no data is available.

For this thesis, the candidate will have the opportunity to work with and/or develop a wide range of techniques (symbolic, statistical, machine learning) that can perform efficient event forecasting in real-time. Particular focus will be placed on combining deep learning and symbolic methods in order to develop a neuro-symbolic system for forecasting with the capacity to provide as its output forecasts that are not only accurate but also explainable and trustworthy. The developed system may be tested on a variety of application domains, from the analysis of biomedical data (epidemiological data, digital patients, evolution of patient trajectories) to the processing of moving objects trajectories and to the guidance of robotic swarms.

Bibliography:

Alevizos, E., Artikis, A. and Paliouras, G., 2022. Complex event forecasting with prediction suffix trees. The VLDB Journal, 31(1), pp.157-180.

Marra, G., Dumančić, S., Manhaeve, R. and De Raedt, L., 2024. From statistical relational to neurosymbolic artificial intelligence: A survey. Artificial Intelligence, p.104062.

Garza, A. and Mergenthaler-Canseco, M., 2023. TimeGPT-1. arXiv preprint arXiv:2310.03589.

Supervisors: Elias Alevizos, Nikos Katzouris

Deliverable (besides the thesis itself): Publication quality report, 12 double-column pages

2) **Unsupervised complex event forecasting**

Abstract:

In traditional machine learning, the goal is to “predict” the output of a function on previously unseen input data. The input data need not necessarily have a temporal dimension and the term "prediction" refers to the output of the learned function on a new data point. On the other hand, the task of forecasting is to predict the temporally future output of some function or the occurrence of an event. Time is thus a crucial component for forecasting. What makes forecasting particularly challenging is the fact that, from the (current) timepoint where a forecast is produced until the (future) timepoint for which we try to make a forecast, no data is available.

In particular, complex event forecasting assumes that we already have a pattern that we need to detect on event streams and that we need to then build a probabilistic model in order to forecasts occurrences of the pattern as well. For this thesis, the candidate will address the issue of unsupervised complex event forecasting, i.e., the situation where the pattern is unknown and the only available information concerns the labels of pattern occurrences. Another additional challenge is that current forecasting models are constructed based on “manually” provided features. The goal is to automate this process in order to discover the most important features through dimensionality reduction, feature extraction and feature selection techniques.

Bibliography:

Alevizos, E., Artikis, A. and Paliouras, G., 2022. Complex event forecasting with prediction suffix trees. The VLDB Journal, 31(1), pp.157-180.

Ghojogh, B., Samad, M.N., Mashhadi, S.A., Kapoor, T., Ali, W., Karray, F. and Crowley, M., 2019. Feature selection and feature extraction in pattern analysis: A literature review. *arXiv preprint arXiv:1905.02845*.

Supervisors: Elias Alevizos

Deliverable (besides the thesis itself): Publication quality report, 12 double-column pages

3) **Time series forecasting with statistical and deep learning methods**

Abstract:

Time series forecasting has been a standard task of time series analysis. Until recently, purely neural methods for time series forecasting were considered sub-optimal in comparison to well known statistical methods (like ARIMA or Vector AutoRegression) or hybrid ones. However, recently an avalanche of deep learning methods for time series forecasting has appeared, claiming to outperform previous state-of-the-art methods.

For this thesis, the candidate will develop a benchmark suite for time series forecasting with various methods, such as rolling mean, ARIMA, LSTM, Temporal Fusion Transformer, DeepAR, NBeats, Graph Neural Networks, Moirai/TimeGPT. The methods will be compared for their accuracy and performance in the task of uni- and multi-variate forecasting for multiple horizons.

Bibliography:

Rob J Hyndman and George Athanasopoulos. Forecasting: Principles and Practice, 3rd ed.

Miller, John A., Mohammed Aldosari, Farah Saeed, Nasid Habib Barna, Subas Rana, I. Budak Arpinar, and Ninghao Liu. "A survey of deep learning and foundation models for time series forecasting." *arXiv preprint arXiv:2401.13912* (2024).

Benidis, Konstantinos, Syama Sundar Rangapuram, Valentin Flunkert, Yuyang Wang, Danielle Maddix, Caner Turkmen, Jan Gasthaus et al. "Deep learning for time series forecasting: Tutorial and literature survey." *ACM Computing Surveys* 55, no. 6 (2022): 1-36.

Challu, Cristian, Kin G. Olivares, Boris N. Oreshkin, Federico Garza Ramirez, Max Mergenthaler Canseco, and Artur Dubrawski. "Nhits: Neural hierarchical interpolation for time series forecasting." In *Proceedings of the AAAI conference on artificial intelligence*, vol. 37, no. 6, pp. 6989-6997. 2023.

Ekambaram, Vijay, Arindam Jati, Nam Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. "Tsmixer: Lightweight mlp-mixer model for multivariate time series forecasting." In *Proceedings of the 29th ACM SIGKDD conference on knowledge discovery and data mining*, pp. 459-469. 2023.

Supervisors: Elias Alevizos

Deliverable (besides the thesis itself): A Python library (possibly using other libraries) running and comparing multiple time series forecasting methods. The process should be highly automated.

4) **High performance real-time AI on streaming data**

Abstract:

A Complex Event Processing (CEP) system takes as input a stream of events, along with a set of patterns, defining relations among the input events, and detects instances of pattern satisfaction, thus producing an output stream of complex events. Typically, a CEP system needs to handle high-volume and high-velocity data streams and must be able to do so with strict latency requirements. It is therefore crucial for such systems to be extremely efficient.

For this thesis, the candidate will develop an expressive and highly efficient CEP engine in Python which will have the capacity to work upon real-time streams of events. The already existing CEP engine, Wayeb, will be used as a reference and baseline. Optimization techniques, like query re-ordering, will be explored, as well as the potential for applying distribution/parallelization methods.

Bibliography:

Alevizos, E., Artikis, A. and Paliouras, G., 2022. Complex event forecasting with prediction suffix trees. The VLDB Journal, 31(1), pp.157-180.

Giatrakos, Nikos, Elias Alevizos, Alexander Artikis, Antonios Deligiannakis, and Minos Garofalakis. "Complex event recognition in the big data era: a survey." *The VLDB Journal* 29 (2020): 313-352.

https://github.com/ElAlev/Wayeb

Supervisors: Elias Alevizos

Deliverable (besides the thesis itself): A Python CEP engine, semantically equivalent to Wayeb.