Workshop 4: Deep Learning & Foundation Models for Forecasting

Tim Januschowski, Databricks, tim.januschowski@databricks.com;

Kashif Rasul, Morgen Stanley Research, rasul.kashif@morganstanley.com

In this in-person workshop, we aim to cover neural forecasting methods from the ground up, starting from the very basics of deep learning to well-established deep forecasting model such as DeepAR (Salinas et al., 2019) to more recent forecasting models, including foundation models for which we'll review selected models.

The workshop will be in-person, with a mix of theoretical lectures and practical sessions. In the lectures, we will focus on the fundamentals of deep learning such as the various architecture types (e.g. feed-forward, convolutional, recurrent neural networks and transformers) and the most important breakthroughs that established the strength of neural networks. Then, we will see how deep learning can be applied to forecasting by reviewing several state-of-the-art neural forecasting models (e.g., WaveNet (Van Den Oord et al., 2016), DeepAR (Salinas et al., 2020), NBEATS (Oreshkin et al., 2019) and the sequence-to-sequence model family [7, 10]). We will introduce AutoGluon-Timeseries, (Shchur et al., 2023) an open-source toolkit for easily training highly accurate probabilistic forecasting models, which automatically combines and tunes both statistical forecasting and deep learning models from frameworks such as GluonTS (Alexandrov et al., 2020) and in particular foundation models which we will also review in some depth.

To complement the lectures, we will offer practical sessions for the workshop participants where we will rely on AutoGluon.

Target Audience and Requirements

This workshop is appropriate for anyone with a solid programming background and a general interest in neural networks. Prior knowledge of neural networks is recommended but optional. Knowledge of forecasting, and basic statistical and machine learning knowledge are a prerequisite. For the practical material, Python programming knowledge is essential.

We will inform the participants of more detailed set-ups closer to the workshop.

References

Alexander Alexandrov, Konstantinos Benidis, Michael Bohlke-Schneider, Valentin Flunkert, Jan Gasthaus, Tim Januschowski, Danielle C Maddix, Syama Sundar Rangapuram, David Salinas, Jasper Schulz, et al. Gluonts: Probabilistic and neural time series modeling in Python. Journal of Machine Learning Research, 21(116):1–6, 2020.

Boris N Oreshkin, Dmitri Carpov, Nicolas Chapados, and Yoshua Bengio. N-beats: Neural basis expansion analysis for interpretable time series forecasting. arXiv preprint arXiv:1905.10437, 2019

Syama Sundar Rangapuram, Matthias W Seeger, Jan Gasthaus, Lorenzo Stella, Yuyang Wang, and Tim Januschowski. Deep state space models for time series forecasting. In Advances in Neural Information Processing Systems, pp. 7785–7794, 2018.

David Salinas, Michael Bohlke-Schneider, Laurent Callot, Roberto Medico, and Jan Gasthaus. High- dimensional multivariate forecasting with low-rank gaussian copula processes. In Advances in Neural Information Processing Systems 32, 2019.

David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. International Journal of Forecasting, 36(3):1181–1191, 2020.

Shchur, O., Turkmen, A.C., Erickson, N., Shen, H., Shirkov, A., Hu, T. and Wang, B., 2023, August. AutoGluon–TimeSeries: AutoML for Probabilistic Time Series Forecasting. In AutoML Conference 2023 (ABCD Track).

Aäron Van Den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. SSW, 125, 2016.

Yuyang Wang, Alex Smola, Danielle Maddix, Jan Gasthaus, Dean Foster, and Tim Januschowski. Deep factors for forecasting. In International Conference on Machine Learning, pp. 6607–6617, 2019.

Kashif Rasul, Arjun Ashok, Andrew Robert Williams, Arian Khorasani, George Adamopoulos, Rishika Bhagwatkar, Marin Biloš, Hena Ghonia, Nadhir Hassen, Anderson Schneider, Sahil Garg, Alexandre Drouin, Nicolas Chapados, Yuriy Nevmyvaka, Irina Rish.

Lag-llama: Towards foundation models for time series forecasting. R0-FoMo: Robustness of Few-shot and Zero-shot Learning in Large Foundation Models. 2023.

Abdul Fatir Ansari, Lorenzo Stella, Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin Shen, Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor, Jasper Zschiegner, Danielle C Maddix, Michael W Mahoney, Kari Torkkola, Andrew Gordon Wilson, Michael Bohlke-Schneider, Yuyang Wang.

Chronos: Learning the language of time series. In TMLR 2024