Deep Learning for Forecasting

Tim Januschowski

Lorenzo Stella

Kashif Rasul

Ali Caner Türkmen

tim.januschowski@zalando.de stellalo@amazon.com

rasul.kashif@morganstanley.com atturkmen@amazon.com

In this in-person workshop, we aim at covering neural forecasting methods from 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.

The workshop 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]). Furthermore, we will dive into recent work that combines neural networks with probabilistic models such as deep state space (Rangapuram et al., 2018) and deep factor (Wang et al., 2019) models. Finally, we will introduce GluonTS (Alexandrov et al., 2020), a time series modeling toolkit primarily aimed at forecasting which is available in open-source.

To complement the lectures, we will offer practical sessions for the workshop participants where we will rely on GluonTS (Alexandrov et al., 2020).

Target Audience and Requirements

This workshop is appropriate for anyone with a solid programming background and a general interest in neural networks. Prior knowledge in neural networks is recommended but not necessary. Knowledge of forecasting, 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.

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.