Deep Learning for Forecasting

June 11, 2020

In this virtual workshop, we aim at covering neural forecasting methods from ground up, starting from the very basics of deep learning up to recent forecasting model improvements such as [5].

Given that the workshop is fully virtual, we will primarily rely on lectures. In these, we will focus on the fundamentals of deep learning such as the various architecture types (e.g. feed-forward, convolutional and recurrent neural networks) 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 [8], DeepAR [6], NBEATS [3] 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 [4] and deep factor [9] models, and bayesian techniques such as [2]. Finally, we will introduce GluonTS [1], a time series modelling toolkit primarily aimed at forecasting which is available in open-source.

To complement the lectures, we will provide notebooks for the workshop participants to interactively explore some of these ideas themselves. Given that the workshop is workshop, this will be self-study largely. Notebooks will rely on GluonTS [1].

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 pre-requistite.

For the practical material, python programming knowledge is essential.

References

[1] Alexander Alexandrov, Konstantinos Benidis, Michael Bohlke-Schneider, Valentin Flunkert, Jan Gasthaus, Tim Januschowski, Danielle C Maddix, Syama Rangapuram, David Salinas, and Jasper Schulz. Gluonts: Probabilistic time series models in python. *arXiv preprint arXiv:1906.05264*, 2019.

[2] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. *arXiv preprint arXiv:1505.05424*, 2015.

[3] 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.

[4] 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*, pages 7785–7794, 2018.

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

[6] David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. Deepar: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 2019.

[7] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to sequence learning with neural networks. In *Proc. NIPS*, Montreal, CA, 2014.

[8] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alexander Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. In *Arxiv*, 2016.

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

[10] Ruofeng Wen Wen, Kari Torkkola, and Balakrishnan Narayanaswamy. A multi-horizon quantile recurrent forecaster. In *NIPS Time Series Workshop*. 2017.