Time Series Forecasting (TSF) Using Various Deep Learning Models

Authors : Jimeng Shi, Mahek Jain, Giri Narasimhan

Abstract : Time Series Forecasting (TSF) is used to predict the target variables at a future time point based on the learning from previous time points. To keep the problem tractable, learning methods use data from a fixed-length window in the past as an explicit input. In this paper, we study how the performance of predictive models changes as a function of different look-back window sizes and different amounts of time to predict the future. We also consider the performance of the recent attention-based Transformer models, which have had good success in the image processing and natural language processing domains. In all, we compare four different deep learning methods (RNN, LSTM, GRU, and Transformer) along with a baseline method. The dataset (hourly) we used is the Beijing Air Quality Dataset from the UCI website, which includes a multivariate time series of many factors measured on an hourly basis for a period of 5 years (2010-14). For each model, we also report on the relationship between the performance and the look-back window sizes and the number of predicted time points into the future. Our experiments suggest that Transformer models have the best performance with the lowest Mean Average Errors (MAE = 14.599, 23.273) and Root Mean Square Errors (RSME = 23.573, 38.131) for most of our single-step and multi-steps predictions. The best size for the look-back window to predict 1 hour into the future appears to be one day, while 2 or 4 days perform the best to predict 3 hours into the future.

Keywords : air quality prediction, deep learning algorithms, time series forecasting, look-back window
Conference Title : ICMLA 2022 : International Conference on Machine Learning and Applications
Conference Location : Boston, United States
Conference Dates : April 21-22, 2022