Abstract

This thesis explores the predictability of financial returns across hourly, daily, weekly, and monthly horizons using Long Short-Term Memory (LSTM) networks. Despite advancements in machine learning, its application in finance faces unique challenges, such as small datasets and low signal-to-noise ratios. Our research aims to address the limitations of existing studies, which predominantly focus on the daily horizon and although some studies analyze different horizons, direct comparisons are challenging due to the varied methodologies and datasets employed. By utilizing a consistent dataset and methodology, we enable a direct comparison of models' performance across various horizons. We enhance predictive models by incorporating fractionally differentiated series to retain memory in financial data and realized volatility from high-frequency data to capture market fluctuations. Our study also extends beyond equities to include futures markets. The key takeaway of our research is that LSTM networks are particularly effective for short-term financial return predictions at hourly and daily horizons. Their performance decreases for longer horizons, such as weekly and monthly, possibly due to fewer market inefficiencies to exploit. Furthermore, the inclusion of futures data does not enhance model performance but reveals interesting trends in feature selection.