EQDerivatives

For this maiden edition of the Journal of Machine Learning in Finance, we selected five papers according to the latest themes in the application of machine learning such as reinforcement learning, natural language processing based signals combined with ESG ratings and options pricing using deep learning.

We open this first edition with a thought-provoking position paper that will set the stage and summarizes the debate regarding limitations and opportunities in machine learning for finance.

In this position paper, the authors discuss some of the most persisting myths related to applications of machine learning. They selected, through consultation with numerous experts, current discussions ranging from supervised learning to reinforcement learning for trading and investment. This position article identifies the major red flags and sets out guidelines and solutions to avoid them, illustrating best practices with simple detailed case studies.

This first volume also reflects current focus (and hopes) on reinforcement learning from the machine learning community. Deemed as the closest form of “real” intelligence, reinforcement learning is a very promising technique for finance across investment management and trading.

In this paper, the authors give an overview and outlook in the field of reinforcement learning as it applies to solving financial applications of intertemporal choice. Through numerical examples in trading and hedging the authors demonstrate the added value of reinforcement learning, illustrating that reinforcement learning can solve dynamic optimization problems in an almost model-free way, relaxing the assumptions often needed for classical approaches. One main contribution of this article is the elucidation of the link between these dynamic optimization problems and reinforcement learning, concretely addressing how to formulate expected intertemporal utility maximization problems using modern machine learning techniques.

The second paper on reinforcement learning focuses on execution, specifically deep execution. The paper compares two popular approaches, value-based and policy-based agent, so to tackle the optimal execution problem by trading against a popular market benchmark, the TWAP. They showed that deep reinforcement learning, even at its infant stage, can reach the theoretically derived optimum through acting on the environment directly. They also found supporting evidence that agents can learn to capitalize on long term price trends without directly observing the price and exploit available information to create dynamic strategies as an informed trader and thus outperform static strategies such as TWAP.

EQDerivatives

In the following paper, Ravenpack's research team propose a ruled-based approach for selecting stocks, blending ESG ratings provided by a major index provider and sentiment scores based on natural language processing techniques. They explore if such a combination could lead to socially responsible portfolios that outperform the market while minimizing the risk. Following a statistical approach, they found supporting evidence in the interest of using ESG rating scores as downside protection and that combining ESG rating with sentiment score could provide additional value for longer term holding period.

For the closing paper, we selected a derivatives pricing methodology using Deep Learning and synthetically created data. Indeed, there have been surprisingly few applied papers in machine learning for options' pricing literature. Machine learning tools, being non-parametric by nature, seemed like a natural approach in pricing problems where empirical reality does not equate with the proposed previous theoretical framework. In other words, when conventional theory fails to explain behavior, machine learning might give an exciting model substitute.

Derivatives pricing is a promising research area for applied machine learning in finance. In this paper, the authors present an entirely data-driven approach using deep learning with synthetically generated data, including term structure of yield and volatility. Their results show a promising reduction in computing time and the added value of using synthetic methods when data are not widely available.

Tony Guida

Editor-in-Chief

Journal of Machine Learning in Finance