
MINGFU WANG, York University

Risk-sensitive Policies for Portfolio Management Problems

Portfolio management is the decision-making process of allocating investment into different financial investment products. However, seeking the optimal trading strategy in a complex and dynamic stock market is challenging due to high uncertainty and massive noise in the financial market. Nowadays, Artificial Intelligence (AI) is well-developed in the financial markets. Trained an AI model as an automated agent can beat experienced human traders. Deep Deterministic Policy Gradient (DDPG), proposed by Google DeepMind, is an actor-critic Deep Reinforcement Learning (RL) algorithm. Although the DDPG achieves a remarkable performance in financial market trading, most of the literature surprisingly ignores the possible risk of rare occurrences of catastrophic events and the effects of the worst-case scenarios on trading decisions. It consists of two-level policies: the lower-level policy and the higher-level policy. The lower-level policy aims to maximize the expected future discounted rewards when the portfolio risk under a certain level of risk, and when the portfolio risk exceeds a certain level of risk, the higher-level policy is to adjust the action from the lower-level policy to reduce the portfolio risk and provides a very conservative trading strategy. In addition, we are the first to apply the distributional DDPG to the portfolio management problems for maximizing the protection of investors. It refers to seek a risk-sensitive policy for portfolio management problems by modeling the distribution of future returns and maximizing the worst-case performance under this distribution. The learned risk-sensitive policy can map the same state to different actions depending on the propensity for risk.