

Smart Liquidity: Reinforcement Learning Approaches for AMM Optimization in Decentr. Exch.

Graduate



Ahmet Cemil Karaosman

Initial Situation:

Decentralized financial markets (DeFi) are developing dynamically and provide novel trading mechanisms through the use of Automated Market Makers (AMMs). These enable continuous liquidity without central order books, but are associated with risks for liquidity providers. With the introduction of concentrated liquidity in Uniswap V3, capital is deployed more efficiently, but it requires complex and active management.

At the same time, Uniswap V4 and its programmable hooks open up new possibilities for dynamically adjusting trading fees and automatically rebalancing liquidity ranges. Against this backdrop, this thesis investigates how Reinforcement Learning (RL) can be used to automate and optimize liquidity strategies.

Approach:

For the study, the RL algorithms Dueling Double Deep Q-Network (DDQN) and Proximal Policy Optimization (PPO) were implemented in a custom-developed simulation environment. The models were trained and evaluated based on historical ETH/USDT data for Uniswap V3. Additionally, an LSTM-based price model was used to support the RL agents with market forecasts.

Since Uniswap V4 was not yet released at the time of this work, its new mechanisms, especially the programmable hooks, were simulated based on the official technical documentation.

The performance of all variants was compared with a static baseline using key metrics such as reward, slippage, divergence loss, and capital utilization.

Result:

The results show that both RL frameworks significantly outperform the static baseline. DDQN achieves the highest cumulative reward but exhibits increased slippage and lower stability. PPO proves to be a more robust solution with higher learning stability and improved capital utilization. Particularly in combination with the V4 hooks, the efficiency metrics are significantly increased. Divergence loss is nearly eliminated, capital utilization rises to over 90%, and the reward development stabilizes sustainably.

Overall, the work confirms that RL, in conjunction with the programmable extensions of Uniswap V4, represents a promising tool for the intelligent and adaptive management of liquidity in DeFi markets.

Advisor

Dr. Shao Jü Woo


Co-Examiner

Prof. Dr. Daniel Lenz


Subject Area

Computational Engineering

Liquidity Utilization [%]
Own presentment



Average Divergence Loss
Own presentment



Average Slippage Loss
Own presentment

