Deep Reinforcement Learning for Quantitative Trading

Yifei Huang

HKUST

April 12, 2019

DRL for Quantitative Trading

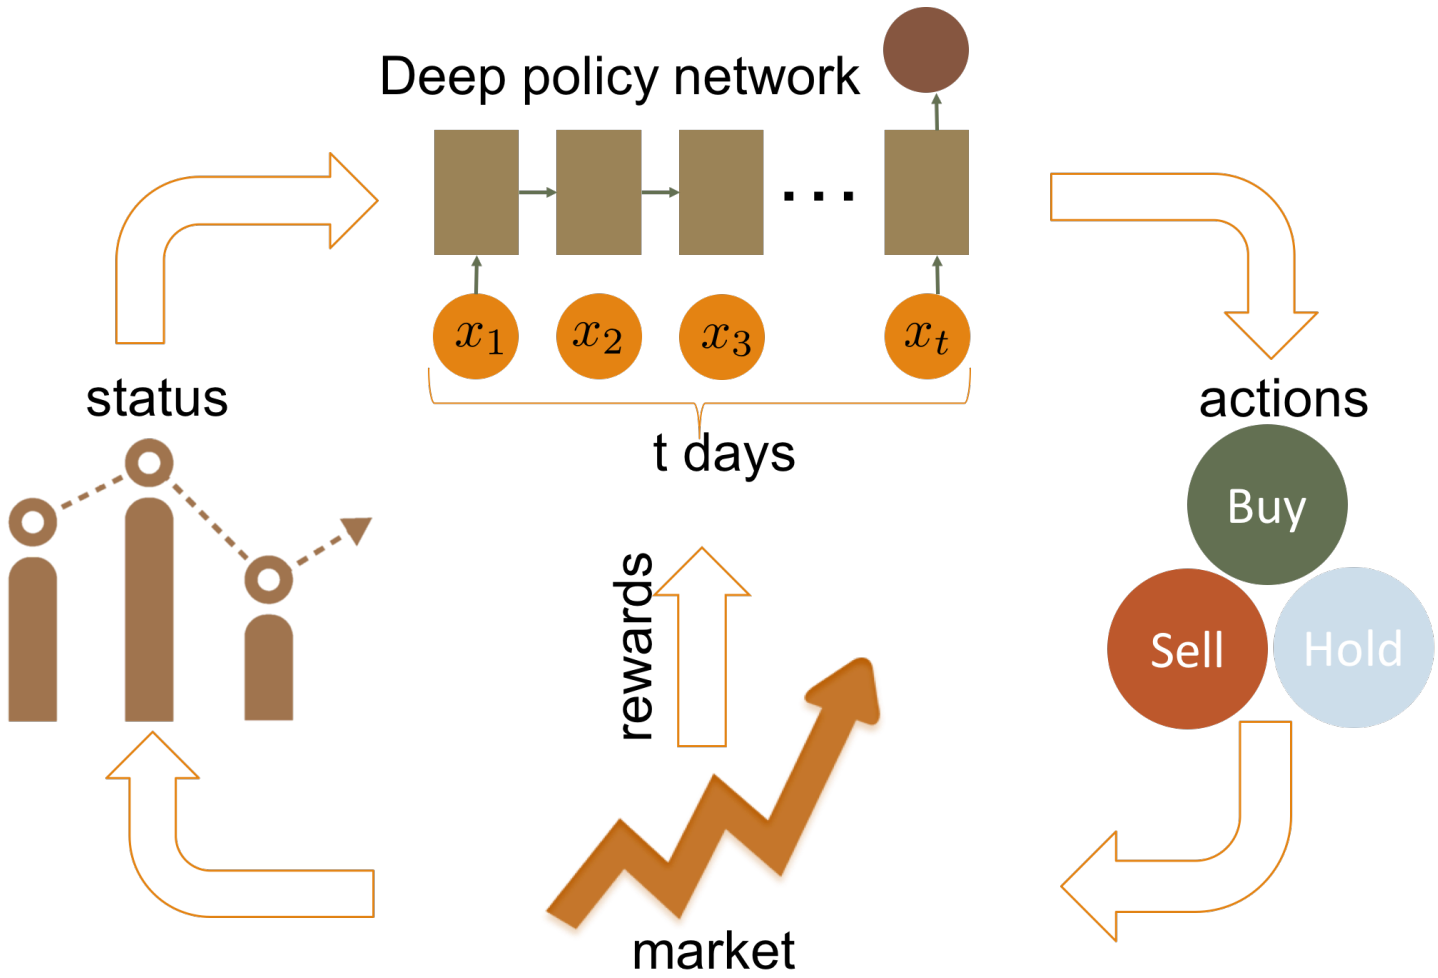

Figure: The RL training process for trading

DRL for Quantitative Trading

State: Agent will receive the current state from environment

- Historical data until now which usually contains OHLCV and some financial indicators such as EMA, MACD, RSI.....

Action: Agent will return the corresponding actions to environment based current state

- *Buy Or Sell:* You buy or sell some shares at current state
- *Hold:* You just hold and do nothing to avoid risk

Reward: Environment will receive the action from agent, then return a reasonable reward for this action and the next state

- *Sharpe Ratio:* $\frac{E[R_a - R_b]}{\sqrt{var[R_a - R_b]}}$ where R_a is the asset return, R_b is the risk free rate
- *Return Ratio:* $R = \frac{ClosePrice_{next}}{ClosePrice_{current}}$
- *Custom Reward:* Any other reasonable reward

DRL Algorithms

Discrete Action:

- Policy Gradient [Sutton et al.(2000)]
- Double DQN [Hasselt et al.(2016)]
- Dueling-DQN [Wang et al.(2015)]
- A3C [Mnih et al.(2016)]
- ...

Continuous Action:

- Deep Deterministic Policy Gradient (DDPG) [Lillicrap et al.(2015)]
- A3C [Mnih et al.(2016)]
- ...

References

https://github.com/ceruleanacg/Personae

Continuous control with deep reinforcement learning
Lillicrap et al. 2015 *preprint arXiv:1509.02971*

Deep reinforcement learning with double q-learning
Van Hasselt et al. 2016 *Thirtieth AAAI Conference on Artificial Intelligence*

Dueling network architectures for deep reinforcement learning
Wang et al. 2015 *preprint arXiv:1511.06581*

Policy gradient methods for reinforcement learning with function approximation
Sutton et al. 2000 *Advances in neural information processing systems*

Asynchronous methods for deep reinforcement learning
Mnih et al. 2016 *International conference on machine learning*