Optimal execution using Reinforcement learning

In the domain of quantitative finance, there are many common problems in which there are precisely specified objectives and constraints. An important example is the problem of optimized trade execution. In this problem, the goal is to buy (or sell) a quantity Q within a fixed time horizon H, in a manner so that the execution cost is minimized.

Reinforcement learning (RL) is the subfield of machine learning concerned with decision-making and control. It studies how an agent can learn how to achieve goals in a complex, uncertain environment. At each step, the agent takes an *action*, whereupon it receives an *observation* and a *reward* from the environment. An RL algorithm seeks to maximize the agent's total reward, through a learning process that usually involves lots of trial and error. RL has a long history, and has recently gained new attention through the successful applications within games by Google DeepMind, [1] and [2]. Other classic well-known examples are robot walking, helicopter flying, and two-player board games like backgammon.

This thesis aims at applying RL methods to the optimal execution problem using microstructure data. One major part of the thesis will be to explore and verify the findings in [3]. In this paper, the authors use state variables such as remaining time, remaining quantity to execute and market price spread combined with simple actions to find an optimal policy. The result is an optimal trading scheme for execution, which outperforms a set of naïve policies out-of-sample. A possible continuation in the thesis is then to make use of other RL techniques such as model-free control or value function approximation. See [4] for a course on those techniques and Reinforcement learning in general.

We seek applicants that have good understanding of machine learning methods, and preferably experience of RL or dynamic programming. The optimal execution problem as presented here is a challenging programming task, so applicants should have a keen interest in programming, especially functional and/or distributed programming.

To apply, please send your CV, grades and a short letter, in Swedish or English, to examensarbete@lynxhedge.se before 2017-10-15.

[1] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, M. Riedmiller. Playing Atari with Deep Reinforcement Learning. Tech Report, 19 Dec. 2013, http://arxiv.org/abs/1312.5602

[2] AlphaGo. https://en.wikipedia.org/wiki/AlphaGo

[3] M. J. Kearns, Y. Nevmyvaka, Y. Feng. 2006. Reinforcement learning for optimized trade execution. In Proceedings of the International Conference on Machine Learning.

[4] UCL Course on RL, Advanced Topics, 2015.

LYNX ASSET MANAGEMENT AB

Box 7060, SE-103 86 Stockholm
Visit Norrmalmstorg 12

Tel +46 8 663 33 60
Fax +46 8 663 33 28

www.lynxhedge.se
info@lynxhedge.se

Reg. no 556573-1782
Registered office Stockholm