M2 Research Internship: Discovering dynamic arbitrage with deep reinforcement learning

September 17, 2020

Laboratory name: CFM Chair of Econophysics & Complex Systems (`www.econophysix.com`), LadHyX
CNRS identification code: UMR CNRS 7646
Internship location: Ecole polytechnique, Palaiseau, and Capital Fund Management, Paris.
Thesis possibility after internship: YES
Funding: YES

Supervision: Stephen Hardiman (Capital Fund Management)
Michael Benzaquen (Ecole Polytechnique) & Jean-Philippe Bouchaud (Capital Fund Management)

Email: michael.benzaquen@polytechnique.edu

Subject

Trading a large quantity on a financial markets has an effect on the market price. On average, other market participants and liquidity providers will respond to your trading activity by revising their quotes in a way which pushes the price (statistically speaking) in the direction of the trade. This phenomenon is known as market impact and like spread costs and exchange fees, it contributes to the effective cost of trading for large institutional investors.

To accurately estimate expected performance of an investment strategy which takes into account all expected trading costs, practitioners must formulate a theory for how their trading activity impacts prices. Although a proposed impact theory may accurately model certain observed empirical responses to real trading activity, it may still permit the possibility of arbitrage (or price manipulation) by means of a "pump and dump" scheme or a non-intuitive sequence of trades that can be repeated to provide consistent gains. Such a non-physical weakness in an impact theory would disqualify it for use as a realistic model for estimating expected investment performance.

The aim of this project is to test the potential of deep reinforcement learning to discover trading policies that can deliver infinite investment returns in such non-realistic environments. An algorithm such as DDPG (a state of the art reinforcement learning algorithm for control problems with continuous action spaces) would learn to gaming the weakness of the impact theory through exploration and feedback from simulation. The objectives are:

- Build a simulation framework with synthetic data for simulating investment performance with a chosen impact theory.
- Adapt DDPG (or equivalent state of the art method) to optimally trade the investment problem.
- Discover dynamic arbitrage strategies when varying the form of impact and decay kernels.
- Tune hyper-parameters and architecture to optimize performance metrics: average reward, speed and stability of convergence.
- Extend to noisier environments, improve stability and convergence with rewards drawn from noisy price returns.

The internship will be held within the CFM Chair of Econophysics and Complex Systems in close collaboration with members of the CFM Trading Research team. Applicants need not have any prior knowledge of financial markets or their microstructure but the ideal candidate would have some experience or training in the practical application of deep reinforcement learning methods. The data is accessible with the aid of Python, R and C++ so familiarity with some of these languages would be beneficial. The applicant would have access to the TensorFlow machine learning framework and a powerful high performance computing cluster.

References

[1] Chaouki, Hardiman, Schmidt, Sérié & De Lataillade (2020). *Deep deterministic portfolio optimization.* The Journal of Finance and Data Science, 6, 16-30.

[2] Gatheral (2010). *No-dynamic-arbitrage and market impact.* Quantitative finance, 10(7), 749-759.

[3] Huberman & Stanzl (2004). *Price manipulation and quasi-arbitrage.* Econometrica, 72(4), 1247-1275.

[4] Lillicrap *et al.* (2015). *Continuous control with deep reinforcement learning.* arXiv preprint 1509.02971.