

Gregory Farquhar

10 Rectory Road
Oxford, Oxfordshire, OX4 1BW
gregory.farquhar@gmail.com
<https://greg-farquhar.github.io>
Mobile: +447900218375

Education

2015-present **Oxford University** **DPhil Candidate (expected graduation Jan 2020)**

- Enrolled in the Center for Doctoral Training in Autonomous Intelligent Machines and Systems
- Research focus on deep reinforcement learning under the supervision of Prof. Shimon Whiteson

2011-2015 **Oxford University** **MSci in Physics (1st class honours)**

- Received 1st class marks in all examinations, ranking in the top 10 students in my cohort's Preliminary Examinations in Physics

2006-2011 **Grammar School at Leeds**

- International Baccalaureate: 45/45 (Maths HL, Physics HL, Chemistry HL, History HL, English SL, French SL)
- GCSE: 12 A*

Experience

Autumn 2018 **Facebook AI Research – Research Internship**

- Worked on action-space curricula for deep reinforcement learning, applied to tactics in the real-time strategy game StarCraft

Summer 2014 **Rutherford Appleton Laboratory – Summer Research Fellowship**

- Developed an automated system to test and characterise the properties of ABC130 silicon detector chips for the planned upgrade to the LHC's ATLAS experiment
- Designed and carried out roboticised tests to verify successful powering and configuration, as well as to analyse the chips' analogue and digital behaviours

Summer 2013 **NASA/JPL – Summer Research Fellowship**

- Developed a data-reduction software package to extract spatially-resolved spectral data from infrared images

Summer 2012 **NASA/JPL – Summer Research Fellowship**

- Developed tools and techniques for data reduction, archiving, and visualisation to improve the efficiency of dozens of research fellows' analytical and archival data-flow

Selected Publications and Preprints

Recent highlights

- **Farquhar, G.**, Whiteson, S., Foerster, J. (2019). Loaded DiCE: Trading off Bias and Variance in Any-Order Score Function Gradient Estimators for Reinforcement Learning. NeurIPS 2019.
 - Low-variance estimators of any-order derivatives for RL, with advantage estimation and more
- **Farquhar, G.**, Rocktäschel, T., Igl, M., & Whiteson, S. (2018). TreeQN and ATreeC: Differentiable Tree Planning for Deep Reinforcement Learning. ICLR 2018.
 - Building intuitions from model-based planning into architectures for deep RL
- Foerster, J.*, **Farquhar, G.***, Afouras, T., Nardelli, N., & Whiteson, S. (2018). Counterfactual multi-agent policy gradients. AAAI 2018.
 - Centralised critic with decentralised actors; per-agent advantage estimate to reduce variance
 - AAAI 2018 Outstanding Student Paper
- **Farquhar, G.**, Gustafson, L., Lin, Z., Whiteson, S., Usunier, N., Synnaeve, G. (2019). Growing Action Spaces. Under review, preprint available.
 - Curriculum learning with progressively larger action spaces to shape exploration

Other publications, chronological

- Foerster, J. N., de Witt, C. A. S., **Farquhar, G.**, Torr, P. H., Boehmer, W., & Whiteson, S. (2019). Multi-Agent Common Knowledge Reinforcement Learning. NeurIPS 2019.
 - Joint-action-space policies that are decentralisable by conditioning on common knowledge
- Luketina, J., Nardelli, N., **Farquhar, G.**, Foerster, J., Andreas, J., Grefenstette, E., ... & Rocktäschel, T. (2019). A Survey of Reinforcement Learning Informed by Natural Language. IJCAI 2019.
 - Arguing for new research directions and benchmarks to integrate modern NLP with RL
- Samvelyan, M., Rashid, T., Schroeder de Witt, C., **Farquhar, G.**, Nardelli, N., Rudner, T. G., ... & Whiteson, S. (2019). The starcraft multi-agent challenge. Proc AAMAS 2019.
 - Set of benchmark tasks for partially observable decentralised control
- Rashid, T., Samvelyan, M., Witt, C.S., **Farquhar, G.**, Foerster, J.N., & Whiteson, S. (2018). QMIX: Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. ICML 2018.
 - Decentralisable joint-action value functions using a monotonically combined factorisation
- Foerster, J.N., **Farquhar, G.**, Al-Shedivat, M., Rocktäschel, T., Xing, E.P., & Whiteson, S. (2018). DiCE: The Infinitely Differentiable Monte-Carlo Estimator. ICML 2018.
 - Correct and easy higher-order score-function gradient estimation
- Foerster, J., Nardelli, N., **Farquhar, G.**, Afouras, T., Torr, P. H., Kohli, P., & Whiteson, S. (2017). Stabilising Experience Replay for Deep Multi-Agent Reinforcement Learning. ICML 2017.
 - Approaches to dealing with nonstationarity in off-policy multi-agent learning

Awards

- AAAI 2018 Outstanding Student Paper Award for “Counterfactual multi-agent policy gradients”
- Oxford MPhys Commendation for outstanding practical work
- Top 50 students nationally in British Physics Olympiad and UKMT Mathematics Challenge
- Distinction in the British Informatics Olympiad

Teaching, Outreach, and Service

2016-present

- Co-supervisor of two masters students, and five 1st year PhD students on short projects
- Program Committee member for NeurIPS (top 400 reviewers), ICML, ICLR, AAAI
- Organise Oxford’s Deep/Reinforcement Learning reading group, regularly present academic work
- Designed and led practical sessions for graduate-level short course in reinforcement learning
- Invited speaker at GTC San Diego 2018
- Outreach: research demos at The AI Summit 2017, Women in CS Oxford open days, speaking at local groups

Programming Experience

- Experienced with Python
- Deep learning frameworks (PyTorch, Tensorflow)
- C++, Lua, Matlab (some experience)
- Core contributor to <https://github.com/oxwhirl/pymarl> and <https://github.com/oxwhirl/smac> (framework and benchmarks for decentralised multi-agent RL)

Other

2009-2016

Competitive Debating

- As Chair of Oxford Debate Selections Committee, managed and balanced a £50k budget to fund debating events and activities at the University of Oxford
- Organised national events, selected and oversaw convening teams for international competitions
- Organised debate workshops and coaching for hundreds of students

Competitive debating involves giving speeches on often specialist topics with 15 minutes preparation time and no access to research resources.

- Competed for England at the World Schools Debate Championships in 2010 and 2011, both times ranking in the top 10 speakers individually
- Represented Oxford University at 6 national and 2 international tournaments – was ranked in the top 50 speakers at the 2012 World Universities Debate Championships in Manila
- Conversational German, some French
- Passionately play and write music – piano and guitar