



University of Oxford Department of Computer Science

Job description and selection criteria

Job title	Postdoctoral Researcher in Reinforcement Learning
Division	MPLS
Department	Computer Science
Location	Wolfson Building, Parks Road, Oxford.
Grade and salary	Grade 7: Salary £30,434 – £37,394 p.a.
Hours	Full Time
Contract type	Fixed term for 2 years from 1 October 2015, or as soon as possible thereafter
Reporting to	Professor Shimon Whiteson
Vacancy reference	119004

Introduction

The University

The University of Oxford is a complex and stimulating organisation, which enjoys an international reputation as a world-class centre of excellence in research and teaching. It employs over 11,000 staff and has a student population of over 22,000.

Our annual income in 2013/14 was £1,174.4m. Oxford is one of Europe's most innovative and entrepreneurial universities: income from external research contracts exceeds £478.3m p.a., and more than 80 spin-off companies have been created.

Oxford is a collegiate university, consisting of the central University and colleges. The central University is composed of academic departments and research centres, administrative departments, libraries and museums. There is a highly devolved operational structure, which is split across four academic divisions, Academic Services and University Collections and University Administrative Services. For further information, please see:

www.ox.ac.uk/staff/about_the_university/new_to_the_university/structure_of_university.

For more information please visit <http://www.ox.ac.uk/about>

The Mathematical, Physical, and Life Sciences Division (MPLS)

The Mathematical, Physical, and Life Sciences (MPLS) Division is one of the four academic divisions of the University. Oxford is widely recognised as one of the world's leading science universities. The disciplines within the MPLS Division regularly appear at the highest levels in world rankings. In the results of the six-yearly UK-wide assessment of university research, REF2014, the MPLS division received the highest overall grade point average (GPA) and the highest GPA for outputs. We received the highest proportion of 4* outputs, and the highest proportion of 4* activity overall. More than 50 per cent of MPLS activity was assessed as world leading.

The MPLS Division's 10 departments and 3 interdisciplinary units span the full spectrum of the mathematical, computational, physical, engineering and life sciences, and undertake both fundamental research and cutting-edge applied work. Our research addresses major societal and technological challenges and is increasingly focused on key interdisciplinary issues. We collaborate closely with colleagues in Oxford across the medical sciences, social sciences and humanities, and with other universities, research organisations and industrial partners across the globe in pursuit of innovative research geared to address critical and fundamental scientific questions.

MPLS is proud to be the home of some of the most creative and innovative scientific thinkers and leaders working in academe. Our senior researchers have been awarded some of the most significant scientific honours (including Nobel prizes and prestigious titles such as FRS and FR.Eng) and we have a strong tradition of attracting and nurturing the very best early career researchers who regularly secure prestigious fellowships. The Division is also the proud holder of eight Athena Swan Awards (4 Silver and 4 Bronze) illustrating our commitment to ensure good practice and to encourage women in science at all levels in the division.

We have around 6,000 students and play a major role in training the next generation of leading scientists. Oxford's international reputation for excellence in teaching is reflected in its position at the top of the major league tables and subject assessments. MPLS academics educate students of high academic merit and potential from all over the world. Through a mixture of lectures, practical work and the distinctive college tutorial system, students develop their ability to solve major mathematical, scientific and engineering problems.

MPLS is dedicated to bringing the wonder and potential of science to the attention of audiences far beyond the world of academia. We have a strong commitment to supporting public engagement in science through initiatives including the Oxford Sparks portal (<http://www.oxfordsparks.net/>) and a large variety of outreach activities; these are crucial activities given so many societal and technological issues demand an understanding of the science that underpins them. We also endeavour to bring the potential of our scientific efforts forward for practical and beneficial application to the real world and our desire is to link our best scientific minds with industry and public policy makers.

For more information about the MPLS division, please visit: <http://www.mpls.ox.ac.uk/>

Department of Computer Science

The Department of Computer Science was established in 1957, making it one of the longest-established Computer Science departments in the country. It is one of the UK's leading Computer Science Departments (ranked first in a number of international rankings). The Research Excellence Framework (REF) in December 2014 resulted in 74 members of the Department having 53% of their research activity ranked in the top category of 4* (world-leading). Overall, we received an average of 3.34 across the department (3* being internationally excellent). A significant majority of the Department are active in externally sponsored research, with both government and industrial funding. At present there are 69 members of academic staff and almost 100 research staff.

The Department has close links with government, industry, and other departments within the University. Among the latter are Mathematics, Engineering, Physics, Statistics and a number of life sciences departments. The Department is housed across multiple sites within the University's South Parks Road Science area, facilitating strong collaborative links with research groups and institutes in closely allied areas (including the Oxford Internet Institute and the Oxford e-Research Centre). This is an essentially inter-disciplinary activity which is at present attracting major funding from a number of sources. At present the Department holds over £50m in external research contracts.

Research in the Department is currently managed in seven themes:

- *Programming Languages and Software Engineering* (led by Professor Jeremy Gibbons, and including Professor Jim Davies) works on a wide variety of areas including model-driven development, functional programming, and static analysis;
- *Security* (led by Professor Bill Roscoe) specialises in cybersecurity (Professor Sadie Creese leads a new Cybersecurity Centre), protocol analysis, trusted computing, networking, and human-centred computing;
- *Automated Verification* (led by Professor Marta Kwiatkowska) covers probabilistic and software model checking (Professor Daniel Kroening), time and concurrency (Professor Joel Ouaknine, Professor James Worrell, and Professors Roscoe and Lowe), and hardware (Professor Tom Melham);
- *Computational Biology* (led by Professor David Gavaghan, and including Professors Kevin Burrage, Helen Byrne, and Blanca Rodriguez) is one of the world's leading groups building computational models of biological systems, and is particularly well-known for its work on the heart;
- *Foundations, Logic and Structures*, (leader Professor Samson Abramsky) which includes groups working on quantum information and computation (Professors Samson Abramsky and Bob Coecke), game semantics and verification (Professor Luke Ong), and constraints (Professor Peter Jeavons);
- *Information Systems* (led by Professor Ian Horrocks, and including Professors Michael Benedikt, Nando de Freitas, Boris Motik, Georg Gottlob, and Michael Wooldridge) has groups working on databases, knowledge representation and reasoning, multi-agent systems, and computational linguistics (Professor Stephen Pulman);
- *Algorithms* (led by Professor Leslie Ann Goldberg, and including Professors Paul Goldberg and Elias Koutsoupias) covering computational complexity, algorithmic game theory, and constraint satisfaction.

For more information please visit: <http://www.cs.ox.ac.uk/>

Summary of the University's Equal Opportunities Policy

The policy and practice of the University of Oxford require that all staff are afforded equal opportunities within employment. Entry into employment with the University and progression within employment will be determined only by personal merit and the application of criteria which are related to the duties of each particular post and the relevant salary structure. In all cases, ability to perform the job will be the primary consideration. Subject to statutory provisions, no applicant or member of staff will be treated less favourably than another because of age, disability, gender reassignment, marriage or civil partnership, pregnancy or maternity, race, religion or belief, sex, or sexual orientation.

Job description

Research topic	Coevolutionary Policy Search
Principal Investigator / supervisor	Professor Shimon Whiteson
Funding partner	ERC

Overview of the role

The aim of this project, which is funded by a Starting Grant from the European Research Council, is to develop a new class of reinforcement learning and sample-based decision-theoretic planning methods that overcome fundamental obstacles to the efficient optimisation of control policies for autonomous agents. Creating agents that are effective in diverse settings is a key goal of artificial intelligence with large potential implications in robotics, e-commerce, information retrieval, traffic control, etc.

A key limitation in existing methods is the prohibitive cost of evaluating candidate policies, i.e., estimating their quality by averaging performance over multiple trials. This cost grows quickly with increases in task complexity (making trials more expensive) or stochasticity (necessitating more trials). To address this difficulty, this project will develop a new approach that simultaneously optimises both policies and the distributions over stochastic events that are sampled from when evaluating those policies. The key insight is that, in many tasks, many trials are wasted because they do not elicit the "controllable rare events" critical for efficiently distinguishing between policies. Thus, the project will develop methods that automatically discover the best events, instead of sampling them randomly.

The focus of the project is on policy-search methods for reinforcement learning as well as closely related methods for sample-based decision-theoretic planning. Techniques such as stochastic optimisation, Bayesian optimisation, and Bayesian quadrature are expected to play a key role. The project will involve both theoretical work as well as an extensive empirical analysis on challenging tasks from robotics and information retrieval.

Responsibilities/duties

- Assist in the management and administration of the project. This involves small-scale project management in order to co-ordinate multiple aspects of work to meet deadlines.
- Collaborate in the preparation of research papers for publication in the scientific literature.

- Contribute to the development and analysis of new algorithms within the project scope
- Test hypotheses and analyse scientific data from a variety of sources, reviewing and refining working hypotheses as appropriate.
- Develop ideas for generating research income, presenting detailed research proposals to senior researchers.
- Present papers at conferences or public meetings.
- Act as a source of information and advice to other members of the group on scientific procedures and experimental techniques.
- Participation in regular meetings with colleagues in Oxford and elsewhere
- Help with the organisation of seminars and workshops
- Assistance in the supervision of post-graduate students working on related projects.
- The postholder may have the opportunity to teach. This may include lecturing, small-group teaching, and tutoring of undergraduates and graduate students.
- The postholder may carry out any other duties as are within the scope, spirit and purpose of the job as requested by their line manager or the Principal Investigators.

Selection Criteria

Essential:

- A PhD (or very close to completion) in Computer Science, or a related area
- A documented track record of the ability to conduct and complete research projects, as witnessed by published peer-reviewed work (according to the experience of the candidate) in the areas of machine learning, decision-theoretic planning, or a related area
- Ability to manage own academic research and associated activities with an organised approach
- Ability to work independently, as well as collaborate with others
- Strong mathematics skills in probability and statistics
- Strong scripting and object-oriented programming skills
- Ability to contribute ideas for new research projects and research income generation
- Excellent communication skills, including the ability to write for publication, present research proposals and results, and represent the research group at meetings.
- Good knowledge of the current state-of-the-art in machine learning

Desirable:

- Good knowledge of the current state-of-the-art in reinforcement learning and decision-theoretic planning
- Experience conducting large-scale experiments with large-scale datasets
- Experience supervising PhD students and managing projects
- Experience of actively collaborating in the development of research articles for publication

Pre-employment screening

Please note that the appointment of the successful candidate will be subject to standard pre-employment screening, as applicable to the post. This will include right-to-work, proof of identity and references. All applicants must read the candidate notes on the University's pre-employment screening procedures, found at:

<https://www.ox.ac.uk/about/jobs/preemploymentscreening/>.

All academic and related posts (any grade above grade 5) are subject to the University's retirement policy. The University operates an employer justified retirement age, for which the retirement date is the 30 September immediately preceding the 68th birthday. Applicants should be aware that any employment beyond the University's retirement age is subject to approval through the procedures outlined at:

www.admin.ox.ac.uk/personnel/end/retirement/acrelretire/ejra/.

Working at the University of Oxford

For further information about working at Oxford, please see:

http://www.ox.ac.uk/about_the_university/jobs/research/

Salary and Benefits

The post, which is a full time appointment, is funded by ERC and is available for up to 2 years from 1 October 2015, or as soon as possible thereafter (with the possibility of extension), has a salary on the University grade 07S scale (currently £30,434 - £37,394 p.a.). This includes membership of the Universities Superannuation Scheme (USS) and has an annual leave entitlement of 38 days per year (inclusive of all public holidays and university closed periods).

How to apply

If you consider that you meet the selection criteria, click on the **Apply Now** button on the 'Job Details' page and follow the on-screen instructions to register as a user. You will then be required to complete a number of screens with your application details, relating to your skills and experience. When prompted, please provide details of two referees and indicate whether we can contact them at this stage. You will also be required to upload a CV and supporting statement. The supporting statement should describe what you have been doing over at least the last 10 years. This may have been employment, education, or you may have taken time away from these activities in order to raise a family, care for a dependant, or travel for example. Your application will be judged solely on the basis of how you demonstrate that that you meet the selection criteria outlined above and we are happy to consider evidence of transferable skills or experience which you may have gained outside the context of paid employment or education.

Please save all uploaded documents to show your name and the document type.

All applications must be received by **midday** on the closing date stated in the online advertisement.

Should you experience any difficulties using the online application system, please email recruitment.support@admin.ox.ac.uk

To return to the online application at any stage, please click on the following link www.recruit.ox.ac.uk

Please note that you will be notified of the progress of your application by automatic e-mails from our e-recruitment system. **Please check your spam/junk mail** regularly to ensure that you receive all e-mails.