



University of Oxford Department of Computer Science

Job description and selection criteria

Job title	Senior Research Associate on AI Threat Detection
Division	Mathematical, Physical and Life Sciences (MPLS)
Department	Computer Science
Location	Robert Hooke Building, Parks Road, Oxford
Grade and salary	Grade 8: Salary range between £48,235- £57,255 per annum (with the potential to under-fill at Grade 7 with salaries in the range of £38,674 - £46,913 p.a.)
Hours	Full time (part time can be considered)
Contract type	Fixed term contract until 31 August 2026
Reporting to	Professor Sadie Creese
Vacancy Reference	176971
Additional information	Whilst the role is a Grade 8 position, we would be willing to consider candidates with potential but less experience who are seeking a development opportunity, for which an initial appointment would be at Grade 7 (Grade 7: £38,674 - £46,913 p.a.) with the responsibilities adjusted accordingly. This would be discussed with applicants at interview/appointment where appropriate.

Research topic	Threat Detection for Cybersecurity of Organisations Using AI
Principal Investigator / supervisor	Prof Sadie Creese (Computer Science) / Prof Michael Goldsmith (Computer Science) / Prof Stephen Roberts (Engineering Science)
Funding partner	The funds supporting this research project are provided by the Oxford Martin School (https://www.oxfordmartin.ox.ac.uk/)
Project web site	www.gcsc.ox.ac.uk/



Overview of the role

The Department of Computer Science is looking to employ a Postdoctoral Researcher to work within an interdisciplinary team on AI Threat Detection. The successful candidate will work as part of a project team, consisting of researchers in the departments of Engineering and Computer Science, supporting the Oxford Martin Programme on AI Threat Detection, as well as engaging across the wide local network of experts in AI, cybersecurity, AI safety & governance.

The Oxford Martin Programme on AI Threat Detection aims to fill a critical gap in AI security by developing advanced methods to detect attacks on AI systems. While much research focuses on preventing AI vulnerabilities, there is currently no way to effectively detect when systems have been compromised. The programme will create a continuous monitoring tool to identify threatening activity across AI data inputs, learning models, and algorithms. By researching optimal detection methods and building a database of attack-response information, the programme seeks to improve the accuracy and resilience of AI threat detection in various organisational contexts.

The successful candidate will be required to work primarily under the direction of Professor Sadie Creese and the supervision of Professor Michael Goldsmith, with guidance from other senior academics.

The proposed research considers:

- What data-points might present useful features for indicating that an AI system or component has been compromised.
- Whether we can determine the likely normal range of outputs from the AI system, and how we might be able to characterise threatening outputs.
- To what extent is that context specific, and how much might be more generally applicable; specifically, in the case of the AI model, whether we can utilise the body of work in understanding model-drift, anomaly-detection and adversarial-responses as mechanisms for detecting model evolutions that are threatening.
- Whether it might be possible to construct a feature vector that combines communication output channels and model evolution assessments, to improve accuracy of threat detection.
- Other data-points, such as human-system interactions, that might feed into the detection strategy.

Flexible working

This position is on-site. Working hours can be flexible subject to the line manager's approval.

Responsibilities/duties

Specific Tasks

- Lead the development of a threat detection architecture for detecting threats to AI systems
- Lead on research methods for identifying and modelling AI system behaviour baselines
- Collaborate with EngSci team on researching the spectrum of threat and vulnerability models for the AI systems
- Collaborate with EngSci team to investigate, through experimentation, the most efficient feature sets for determining AI system behaviours for the models in the library

- Collaborate with EngSci team to design and run experiments exploring the utility of the prototype threat detection systems for protecting the various AI models in the library, investigating the resilience of AI models to relevant security threats and interpreting results
- Agree clear task objectives, organise, and delegate work to other members of the team and coach other members of the group on specialist methodologies or procedures
- Raise research funds through grant applications, manage own area of a larger research budget
- Share responsibility for shaping the research group's plans and the writing of group-funding applications for new research projects
- Liaise with funding bodies and provide information to project stakeholders and represent the research group at external meetings/seminars, either with other members of the team or alone
- Write and publish academic papers detailing findings
- Present research findings at international conferences [and contribute to grant proposals]
- Mentor graduate students and organise seminars
- Develop software tools and relevant databases

Additional Tasks

- Managing own academic research and administrative activities to meet team milestones and deadlines. This involves small-scale project management to co-ordinate multiple aspects.
- Adapt existing and develop new research methodologies and materials as required
- Prepare working theories and analyse qualitative and/or quantitative data from a variety of sources, reviewing and refining theories as appropriate.
- Contribute ideas for new research tasks, and presentation of concepts to other senior researchers.
- Develop ideas for generating research income, and present detailed research proposals to senior researchers.
- Collaborate in the preparation of scientific reports and journal articles, and occasionally present papers.
- Act as a source of information and advice to other members of the group on methodologies or procedures.
- Represent the research group at external meetings and seminars, either with other members of the group or alone.
- Carry out collaborative projects with colleagues in partner institutions and research groups.
- The researcher may have the opportunity to undertake ad-hoc paid teaching (this includes lecturing, demonstrating, small-group teaching, tutoring of undergraduates and graduate students and supervision of Masters projects in collaboration with principal investigators). Permission must be sought in advance for each opportunity.
- Any other duties appropriate with the role.

Selection criteria

Essential

- Hold a relevant PhD/DPhil and post-qualification research experience in a cybersecurity-relevant field (e.g. computer science, international relations, economics, political science, social or physical sciences, philosophy, cognitive or social psychology, law, anthropology, development or sociology).
- Previous experience of contributing to high-quality and relevant publications/presentations.

- Ability to work independently and as part of a team.
- Excellent communication skills, including the ability to write for publication, present research proposals and results, and represent the research group at meetings.
- Strong publication record and familiarity with the existing literature and research in the field.
- Possess sufficient specialist knowledge in the discipline to further develop the research project and methodologies.
- Ability to independently plan and manage a research project, including a research budget.
- Ability to raise funds through making grant applications.

Desirable

- Experience with research in the adversarial resilience of AI models, or related approaches.
- Experience of supervising staff.
- Experience of managing a research budget.
- Experience of making grant applications.

Pre-employment screening

Standard checks

If you are offered the post, the offer will be subject to standard pre-employment checks. You will be asked to provide: proof of your right-to-work in the UK; proof of your identity; and (if we haven't done so already) we will contact the referees you have nominated. You will also be asked to complete a health declaration so that you can tell us about any health conditions or disabilities for which you may need us to make appropriate adjustments.

Please read the candidate notes on the University's pre-employment screening procedures at: <https://www.jobs.ox.ac.uk/pre-employment-checks>

About the University of Oxford

Welcome to the University of Oxford. We aim to lead the world in research and education for the benefit of society both in the UK and globally. Oxford's researchers engage with academic, commercial and cultural partners across the world to stimulate high-quality research and enable innovation through a broad range of social, policy and economic impacts.

We believe our strengths lie both in empowering individuals and teams to address fundamental questions of global significance, while providing all our staff with a welcoming and inclusive workplace that enables everyone to develop and do their best work. Recognising that diversity is our strength, vital for innovation and creativity, we aspire to build a truly diverse community which values and respects every individual's unique contribution.

While we have long traditions of scholarship, we are also forward-looking, creative and cutting-edge. Oxford is one of Europe's most entrepreneurial universities and we rank first in the UK for university

spin-outs, and in recent years we have spun out 15-20 new companies every year. We are also recognised as leaders in support for social enterprise.

Join us and you will find a unique, democratic and international community, a great range of staff benefits and access to a vibrant array of cultural activities in the beautiful city of Oxford.

For more information, please visit www.ox.ac.uk/about/organisation.

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). Our Computer Science and Informatics submission to the UK Research Excellence Framework (REF) in December 2021 resulted in 81% of research activity ranked as 4* (world-leading) and the rest ranked as 3* (internationally excellent). A significant majority of the Department are active in externally sponsored research, with both government and industrial funding. At present, there are 74 members of academic staff and 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 life sciences. 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). At present, the Department holds over £50m in external research contracts.

Research in the Department is currently managed in ten themes:

- *Algorithms & Complexity Theory*, led by Professor Leslie Ann Goldberg, focusses on determining the inherent difficulty of computational problems, classifying problems according to this inherent difficulty, and designing and analysing algorithms that use computational resources as efficiently as possible;
- *Artificial Intelligence & Machine Learning*, led by Professor Michael Wooldridge, focuses on theoretical foundations of AI, multiagent systems, deep learning, reinforcement learning, and computational linguistics;
- *Automated Verification*, led by Professor Marta Kwiatkowska, investigates theory and practice of formal verification and correct-by-construction synthesis for software and hardware systems;
- *Computational Biology & Health Informatics*, led by Professor Blanca Rodriguez, is concerned with computational approaches for biomedical research and healthcare innovation;
- *Human-Centred Computing*, led by Professor Nigel Shadbolt, includes human-computer interaction, social computing, and the worldwide web;
- *Data and Knowledge & Action*, led by Professor Ian Horrocks, includes databases, knowledge representation and reasoning;
- *Programming Languages*, led by Professor Sam Staton, includes functional programming, program analysis, and programming language foundations;
- *Quantum*, led by Professor Jonathan Barrett, focusses on quantum computing including quantum software, causality in quantum theory, quantum cryptography and foundations of quantum computing;
- *Security*, led by Professor Ivan Martinovic, specialises in cybersecurity, protocol analysis, systems security, trusted computing, and networking.
- *Systems*, led by Professor Niki Trigoni, focusses especially on cyber physical systems. We plan to substantially broaden our research in systems to complement our existing research areas.

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

The Department of Computer Science holds a bronze Athena Swan award to recognise advancement of gender equality: representation, progression and success for all.

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 and the MPLS Division is home to our non-medical sciences, with 9 academic departments that 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 tackles major societal and technological challenges – whether developing new energy solutions or improved cancer treatments, understanding climate change processes, or helping to preserve biodiversity, 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 and we have a strong tradition of attracting and nurturing the very best early career researchers who regularly secure prestigious fellowships and faculty positions. MPLS continues in its work to support diversity in its staffing, seeing that it will bring benefits to all, and we are pleased to note that all academic departments in the Division hold Athena Swan Awards.

We have around 7,000 full and part-time students (including approximately 3,500 graduate 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 diverse 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 (www.oxfordsparks.ox.ac.uk) 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 bring the potential of our scientific efforts forward for practical and beneficial application to the real world and our desire, aided by the work of Oxford University Innovation and Oxford Sciences Innovation, is to link our best scientific minds with industry and public policy makers.

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

How to apply

Applications are made through our online recruitment portal. Information about how to apply is available on our Jobs website <https://www.jobs.ox.ac.uk/how-to-apply>.

Your application will be judged solely on the basis of how you demonstrate that you meet the selection criteria stated in the job description.

As part of your application you will be asked to provide details of two referees and indicate whether we can contact them now.

You will be asked to upload a CV and a supporting statement. The supporting statement must explain how you meet each of the selection criteria for the post using examples of your skills and experience. This may include experience gained in employment, education, or during career breaks (such as time out to care for dependants)

Please upload all documents **as PDF files** with your name and the document type in the filename.

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

Information for priority candidates

A priority candidate is a University employee who is seeking redeployment because they have been advised that they are at risk of redundancy, or on grounds of ill-health/disability. Priority candidates are issued with a redeployment letter by their employing department(s).

If you are a priority candidate, please ensure that you attach your redeployment letter to your application (or email it to the contact address on the advert if the application form used for the vacancy does not allow attachments).

If you need help

Application FAQs, including technical troubleshooting advice is available at:

<https://staff.web.ox.ac.uk/recruitment-support-faqs>

Non-technical questions about this job should be addressed to the recruiting department directly (hr@cs.ox.ac.uk)

To return to the online application at any stage, please go to: www.recruit.ox.ac.uk.

Please note that you will receive an automated email from our online recruitment portal to confirm receipt of your application. **Please check your spam/junk mail** if you do not receive this email.

Important information for candidates

Data Privacy

Please note that any personal data submitted to the University as part of the job application process will be processed in accordance with the GDPR and related UK data protection legislation. For further information, please see the University's Privacy Notice for Job Applicants at: <https://compliance.admin.ox.ac.uk/job-applicant-privacy-policy>. The University's Policy on Data Protection is available at: <https://compliance.admin.ox.ac.uk/data-protection-policy>.

The University's policy on retirement

The University operates an Employer Justified Retirement Age (EJRA) for very senior research posts at **grade RSIV/D35 and clinical equivalents E62 and E82** of 30 September before the 70th birthday. The justification for this is explained at: <https://hr.admin.ox.ac.uk/the-ejra>.

For **existing** employees on these grades, any employment beyond the retirement age is subject to approval through the procedures: <https://hr.admin.ox.ac.uk/the-ejra>.

There is no normal or fixed age at which staff in posts at other grades have to retire. Staff at these grades may elect to retire in accordance with the rules of the applicable pension scheme, as may be amended from time to time.

Equality of opportunity

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. No applicant or member of staff shall be discriminated against because of age, disability, gender reassignment, marriage or civil partnership, pregnancy or maternity, race, religion or belief, sex, or sexual orientation.

Benefits of working at the University

Employee benefits

University employees enjoy 38 days' paid holiday, generous pension schemes, travel discounts, and a variety of professional development opportunities. Our range of other employee benefits and discounts also includes free entry to the Botanic Gardens and University colleges, and discounts at University museums. See www.admin.ox.ac.uk/personnel/staffinfo/benefits.

University Club and sports facilities

Membership of the University Club is free for all University staff. The University Club offers social, sporting, and hospitality facilities. Staff can also use the University Sports Centre on Iffley Road at discounted rates, including a fitness centre, powerlifting room, and swimming pool. See www.club.ox.ac.uk and www.sport.ox.ac.uk/oxford-university-sports-facilities.

Information for staff new to Oxford

If you are relocating to Oxfordshire from overseas or elsewhere in the UK, the University's Welcome Service website includes practical information about settling in the area, including advice on relocation, accommodation, and local schools. See www.welcome.ox.ac.uk.

There is also a visa loan scheme to cover the costs of UK visa applications for staff and their dependents. See www.admin.ox.ac.uk/personnel/permits/reimburse&loanscheme/.

Family-friendly benefits

With one of the most generous family leave schemes in the Higher Education sector, and a range of flexible working options, Oxford aims to be a family-friendly employer. We also subscribe to My Family Care, a service that provides practical advice and support for employees who have caring responsibilities. The service offers a free telephone advice line, and the ability to book emergency back-up care for children, adult dependents and elderly relatives. See www.admin.ox.ac.uk/personnel/staffinfo/benefits/family/mfc/.

Childcare

The University has excellent childcare services, including five University nurseries as well as University-supported places at many other private nurseries.

For full details, including how to apply and the costs, see www.admin.ox.ac.uk/childcare/.

Disabled staff

We are committed to supporting members of staff with disabilities or long-term health conditions. For further details, including information about how to make contact, in confidence, with the University's Staff Disability Advisor, see www.admin.ox.ac.uk/eop/disab/staff.

Staff networks

The University has a number of staff networks including the Oxford Research Staff Society, BME staff network, LGBT+ staff network and a disabled staff network. You can find more information at www.admin.ox.ac.uk/eop/inpractice/networks/.

The University of Oxford Newcomers' Club

The University of Oxford Newcomers' Club is an organisation run by volunteers that aims to assist the partners of new staff settle into Oxford, and provides them with an opportunity to meet people and make connections in the local area. See www.newcomers.ox.ac.uk.