

# The Alan Turing Institute

**Research Associate – Reinforcement Learning, Machine Learning, Security and Privacy (NeurIPS, AAI, ICML, ICLR, etc)**

## THE ALAN TURING INSTITUTE

There has never been a more significant time to work in data science and AI. There is recognition of the importance of these technologies to our economic and social future: the so-called fourth industrial revolution. The technical challenge of keeping our data secure and private has grown in its urgency and importance. At the same time, voices from academia, industry, and government are coming together to debate how these technologies should be governed and managed.

The Alan Turing Institute, as the UK's national institute for data science and artificial intelligence, plays an important part in driving forward advances in these technologies in order to change the world for the better.

The Institute is named in honour of Alan Turing, whose pioneering work in theoretical and applied mathematics, engineering and computing is considered to have laid the foundations for modern-day data science and artificial intelligence. The Institute's goals are to undertake world-class research, apply its research to real-world problems, driving economic impact and societal good, lead the training of a new generation of scientists, and shape the public conversation around data and algorithms.

After launching in 2015 with government funding from EPSRC and five founding universities, the Institute has grown an extensive network of university partners from across the UK and launched a number of major partnerships with industry, public and third sector. Today it is home to more than 500 researchers, a rapidly growing team of in house research software engineers and data scientists and a business team.

## BACKGROUND

The Defence & Security programme at the Turing is looking to expand its AI for Cyber Defence (AICD) Research Centre, a national hub for AIxCyber related research in the UK.

As a team, we aim to autonomously defend computer networks and systems through our ground-breaking research at the intersection between traditional computer security and machine learning. Day to day, we collaborate with technical and subject matter experts from our partner organisations as well as academics, software engineers, and data scientists from across the Turing's research community. We present our work to a range of audiences including research colleagues, senior decision makers and non-technical stakeholders.

We are a cross-disciplinary team and encourage applications from both generalists and specialists including computer scientists, software engineers, machine learning practitioners, mathematicians, statisticians and data engineers. Applicants with PhDs in either systems/information security or deep reinforcement learning are most welcome.

The team practices an agile, experiment-driven approach and values a positive, supportive and collaborative environment in which 'radical candour' and 'lifelong learning' are encouraged. We take ownership of our work and operate with a high level of autonomy in our roles, to deliver measurable impact (e.g., publications in top-tier venues) to our partners. Thus, you will have the space to initiate and lead your own projects as well as join ongoing projects/papers led by other researchers in the team. Collaborating with external researchers, supervising MSc thesis projects and proposing projects for the Turing enrichment (PhD) students etc are encouraged.

## ROLE PURPOSE

AICD is aiming to become a world-class research centre, at the Alan Turing Institute, focussed on delivering the science needed for developing autonomous and resilient cyber defence using deep reinforcement learning (DRL) and other autonomous decision making techniques. This involves both the application of existing AI algorithms and techniques as well as fundamentally advancing AI where necessary.

The technical scope of this role includes:

- The application of modern DRL techniques including (but not limited to):
  - Transformers and attention techniques for both episodic process memory and reduced action and observation spaces.

# The Alan Turing Institute

- Multi-agent approaches such as swarms of specialised agents
- Curiosity and related techniques for internally-generated reward signals.
- Meta-learning and generalisability to novel environments.
- Genetic techniques e.g., for improving RL algorithm performance and generalisability.
- Adversarial approaches to RL policies as well as other AI systems.
- Foundational models for decision making.
- Explainable RL.
- Relevant foundational research on AI including making improvements to existing techniques and proposing alternatives that advance the state of the art.
- Advancing open security and privacy problems that might be solved with modern AI techniques. The areas we are working on include (but are not limited to):
  - Fully autonomous cyber operations and network defence.
  - Active attacks on anonymity networks.
  - Validating ciphers, protocols and their implementations.
  - Systems security attacks and defences. Strengthening defences by discovering new adversarial techniques and models (e.g., vulnerability discovery, automated red teaming).
  - Adaptive fuzzing e.g., of web protocols, binary executables and hardware implementations.
- Writing papers for submission to high quality peer review venues (e.g., AAI, ICML, ICLR, NeurIPS, USENIX, ACM CCS, IEEE S&P).

## DUTIES AND AREAS OF RESPONSIBILITY

The research associate will work closely with the Centre Leads based at the Turing Institute to:

- Pursue collaborative research of high quality, consistent with making a full active research contribution in line with the research strategy outlined by the Centre Leads.
- Write or contribute to publications or disseminate research findings using other appropriate media.
- Attend and present research findings and papers at academic and professional conferences, and to contribute to the external visibility of the Institute.
- Ensure compliance with secure handling of data and health and safety in all aspects of work.
- Participate in and develop internal and external partnerships, for example to identify sources of funding, generate income, obtain projects, or build relationships for future activities.
- Contribute to the running of workshops to showcase early-stage research in the area.
- Participate in international research challenges and competitions.
- Contribute to the running of interactive events such as data study groups and capture the flag competitions.
- Establish regular ways to convene researchers such as with a seminar series, special interest and/or reading group(s).

Please note that job descriptions cannot be exhaustive, and the postholder may be required to undertake other duties, which are broadly in line with the above key responsibilities. This job description is written at a specific time and is subject to changes as the demands of the Institute and the role develop.

# The Alan Turing Institute

PERSON SPECIFICATION		
Skills and Requirements	Essential (E) Desirable (D)	Tested at Application (A) Tested at Interview (I)
Post holders will be expected to demonstrate the following:		
<b>Education/Qualification</b>		
PhD or an equivalent qualification in machine learning, computer security or a closely related discipline.	E	A
<b>Knowledge and Experience</b>		
General machine learning knowledge.	E	A/I
Familiarity with reinforcement learning, DRL or related control techniques.	D	A/I
Prior experience developing software in a scientific computing context, ideally in Python. Experience in frameworks such as NumPy, Tensorflow, PyTorch, Ray/RLlib, Stable Baselines. Experience in development suites, systems and versioning products (e.g., Git, IDEs, Linux).	D	A/I
The ability to initiate, plan, organise, implement and deliver programmes of work to tight deadlines.	E	A/I
The ability to initiate, develop and deliver high quality research aligned with the research strategy indicated by the Centre Leads and any industrial stakeholders and to publish in peer reviewed conferences and journals.	E	A/I
The ability to work in a team and interact professionally within a team of researchers and students.	E	A/I
A desire to collaborate with experts across domains.	E	A/I
Ability to work independently and proactively to lead one's own work, including planning and execution, and to prioritise work to meet deadlines.	E	A/I
Ability to organise working time, take the initiative, and carry out research independently, under the guidance of the PI.	E	A/I
Publication record in top-tier peer reviewed international conferences or journals.	E	A/I
<b>Communication</b>		
Good effective communication (oral and written) skills, presentation and training skills.	E	I
Adapts the style of communication to the audience and ensures understanding.	E	I
<b>Project Management &amp; Project Delivery</b>		
Proactive approach to managing stakeholders and their requirements and identifying opportunities for collaboration.	E	A&I
Adapts services and systems to meet stakeholders' needs and identifies ways of improving standards. Learns from issues and takes action to resolve them.	E	A&I
<b>Decision Making</b>		
Ability to lead own work Independently, and make independent decisions which are low risk and that mainly affect themselves or a small number of people and are guided by regulation and practice	E	A

# The Alan Turing Institute

Work with others to make collaborative decisions that may be operational or strategic and impact immediate team or work area.	E	A&I
Recommend and advise on available options for decisions that affect operational processes, taking into account any risks.	E	A&I
<b>Initiative and Problem Solving</b>		
Uses judgement to analyse and solve problems, and take action to prevent recurrence of problems.	E	I
Consider possible solutions to identify those which offer wider benefits, and obtain evidence to support thinking.	E	I
<b>Analysis and Research</b>		
Ability to plan and implement rigorous analysis plans.	E	I
Identify and use a range of standard sources to gather and analyse routine data and produce reports that can be interpreted by others.	E	I
Understand when additional data is required and identifies appropriate sources. Produces reports that identify key issues and findings.	E	I
<b>Other Requirements</b>		
Commitment to EDI principles and to the Organisation values	E	I
Eligible for UK/NATO security clearance in principle, e.g., by nationality or 5+ years NATO residency etc	E	I

# The Alan Turing Institute

## OUR VALUES

The Alan Turing Institute is committed to equality diversity and inclusion and to eliminating discrimination. All employees are expected to embrace, follow and promote our [EDI Principles](#) and Our Values.

### Our values

- Trust**  
We create an environment where we have trust and can be trusted
- Inclusivity**  
We expect our Turing community to contribute to a culture that is inclusive and free of barriers
- Respect**  
We all have different roles, priorities and challenges but our shared purpose is the same
- Leadership**  
Leadership is everyone's business; Turing leaders set the right tone and lead by example
- Transparency**  
Everyone should understand the how and the why of our decisions and actions
- Integrity**  
We are all ambassadors for the Turing's mission of changing the world for the better

## APPLICATION PROCEDURE

If you are interested in this opportunity, please click the apply button below. You will need to register on the applicant portal and complete the application form including your CV and covering letter. If you have questions about the role or would like to apply using a different format, please contact us on 020 3970 2148 or 0203 862 3340, or email [recruitment@turing.ac.uk](mailto:recruitment@turing.ac.uk).

**CLOSING DATE FOR APPLICATIONS: 4 June 2023 at 23:59**

## TERMS AND CONDITIONS

This full-time is offered on a fixed-term basis until March 2025. The annual salary is £42,893 - £48,510 plus excellent benefits, including flexible working and family friendly policies, <https://www.turing.ac.uk/work-turing/why-work-turing/employee-benefits>

Candidates who have not yet been officially awarded their PhD will be appointed as Research Assistant at a salary of £40,148 per annum.

## EQUALITY, DIVERSITY AND INCLUSION

The Alan Turing Institute is committed to creating an environment where diversity is valued and everyone is treated fairly. In accordance with the Equality Act, we welcome applications from anyone who meets the specific criteria of the post regardless of age, disability, ethnicity, gender reassignment, marital or civil partnership status, pregnancy and maternity, religion or belief, sex and sexual orientation.

We are committed to building a diverse community and would like our leadership team to reflect this. We therefore welcome applications from the broadest spectrum of backgrounds.

Reasonable adjustments to the interview process will be made for any candidates with a disability.

***Please note all offers of employment are subject to obtaining and retaining the right to work in the UK and satisfactory pre-employment security screening which includes a DBS Check.***

***Full details on the pre-employment screening process can be requested from [HR@turing.ac.uk](mailto:HR@turing.ac.uk).***