

The Alan Turing Institute

Research Associate – Human Machine Teaming, Reinforcement Learning (CHI, AAI, NeurIPS, ICML, ICLR, SOUPS)

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 with an individual who is passionate about human-machine collaboration. As part of an exciting new collaboration with Army Research Labs (ARL) and the University of Birmingham, we are looking for someone to integrate autonomous intelligent agents and humans as part of an international and interdisciplinary team focussed on the allocation of function in human machine teams.

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

It is essential that autonomous agents can be cooperatively integrated with humans, enabling them to work together to enhance the security and privacy of digital systems. By leveraging the unique strengths of both humans and machines, we can develop more effective cyber defence strategies and stay ahead of evolving threats. Ultimately, the cooperative integration of autonomous agents and humans is key to building a more secure and resilient digital infrastructure that can withstand even the most sophisticated attacks.

The successful candidate will play a crucial role in advancing our understanding of how autonomous intelligent agents can be integrated with humans to optimise team performance and achieve better results. We aim to push the boundaries of

The Alan Turing Institute

human-machine teaming research and develop new technologies and strategies that can be applied in various team settings, including military and civilian applications. The technical scope of this role includes:

- Developing intelligent autonomous agents for human-machine collaboration (e.g., gym environments) including working with Deep Multi Agent Reinforcement Learning.
- Relevant foundational research on AI including making improvements to existing techniques and proposing alternatives that advance the state of the art.
- Advancing the state of the art in human machine teaming with a focus on the allocation of function in mixed teams.
- Authoring papers for submission to high quality peer review venues.

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		
<p>Skills and Requirements</p> <p>Post holders will be expected to demonstrate the following:</p>	<p>Essential (E)</p> <p>Desirable (D)</p>	<p>Tested at application (A)</p> <p>Tested at interview (I)</p>

The Alan Turing Institute

Education/Qualification		
PhD or an equivalent qualification in machine learning, computer security or a closely related discipline.	E	A
Knowledge and Experience		
Familiarity with human-computer interaction and human-machine teaming principles.	E	A/I
General machine learning knowledge.	E	A/I
Prior experience developing software in a scientific computing context, ideally in Python. Experience in frameworks such as NumPy, Tensorflow, PyTorch, Ray/RLlib, Gymnasium, PettingZoo, Stable Baselines. Experience in development suites, systems, and versioning products (e.g., Git, IDEs, Linux).	D	A/I
Familiarity with reinforcement learning, DRL or related control techniques.	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
Communication		
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
Project Management & Project Delivery		
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
Decision Making		
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
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
Initiative and Problem Solving		

The Alan Turing Institute

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
Analysis and Research		
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
Other Requirements		
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	A

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.

CLOSING DATE FOR APPLICATIONS: Sunday 17th September 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 for 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.

Successful candidates will need to undergo a security check by DSTL's security team.

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.