

The Alan Turing Institute

RESEARCH ASSISTANT – ELSA, SAFE AND ETHICAL AI PROGRAMME

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 purpose is to make great leaps in data science and AI research to change the world for the better. Its goals are to advance world-class research and apply it to national and global challenges, build skills for the future by contributing to training people across sectors and career stages, and drive an informed public conversation by providing balanced and evidence-based views on data science and AI.

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

Turing Lead Adrian Weller received a grant as Co-Investigator (Co-I) to be part of the EU Horizon European Lighthouse on Secure and Safe AI project ('ELSA'). The goal of ELSA is to create a European virtual centre of excellence for safe and secure AI that will address major challenges hampering the deployment of AI technologies. Research will focus on "technical robustness and safety", "privacy preserving techniques and infrastructures" and "human agency and oversight". Efforts to detect, prevent, and mitigate threats and enable recovery from harm focus on 3 grand challenges:

1. Robustness guarantees and certification
2. Private and robust collaborative learning at scale
3. Human-in-the-loop decision making: Integrated governance to ensure meaningful oversight

which cut across 6 use cases: health, autonomous driving, robotics, cybersecurity, multimedia, and document intelligence.

ELSA will focus on machine-learning-based methods and especially deep learning, which form the foundation for many modern AI applications. Throughout the project, robust technical approaches will be integrated with legal and ethical principles supported by meaningful and effective governance architectures to nurture and sustain the development and deployment of AI technology that serves and promotes foundational European values.

This project builds on and expands the internationally recognized network of excellence ELLIS, the European Laboratory for Learning and Intelligent Systems. ELLIS is made up of a network of over 100 organizations and more than 300 ELLIS Fellows and Scholars (113 ERC grants) committed to shared standards of excellence.

For more details about the project, please see [elsa: European Lighthouse on Secure and Safe AI \(elsa-ai.eu\)](https://elsa-ai.eu).

ROLE PURPOSE

The Turing are hiring a Research Assistant to support the delivery of parts of the ELSA project for which Turing will be funded. The successful candidate will work with Turing Lead, Adrian Weller at the Turing and the other ELSA project partners.

The Alan Turing Institute

The ideal candidate will contribute in one or more of the following ways: 1. Technical; 2. Law and Governance; 3. Ethics and Policy.

The project will develop new theory, methods and tools to help enable the secure and safe development and deployment of AI systems. An area of focus will be reliable tools for AI assurance to help monitor and enforce desired properties of systems in the context of ELSA project use cases, coordinating with legal expert Co-Investigators on the project, and partners to help develop technical solutions to ethical governance challenges raised. Work we will advance includes developing robust measures of transparency for fairness, incorporating methods to ensure appropriate privacy is preserved, and exploring scalable approaches to incorporating human knowledge and priors in the loop for effective human-machine team decisions.

For the reason above, we welcome applications from candidates with expertise in any of the following: Mathematics, Computer Science, as well as Law, Ethics or Policy.

Within transparency, the primary focus is interpretability of the outputs of machine learning systems which will be directly beneficial for practitioners, stakeholders or regulators; though we shall also consider broader issues, such as the process of how data and humans fit into decision pipelines.

The candidate will join a vibrant team of researchers and will have opportunities to engage with cutting-edge projects and experts at leading universities. These include Karen Yeung (Birmingham), and project teams funded by other Turing strategic partnerships.

DUTIES AND AREAS OF RESPONSIBILITY

General duties for Research Assistant posts:

- Contribute to report writing, academic and other publications, and presentations and support in writing other outputs such as blogs, explainers, academic outputs
- Assist in communicating research outputs to diverse stakeholders, through conferences, events, meetings, and press opportunities as appropriate.
- Liaise with a network of internal and external contacts and key stakeholders to assist with the work of the research project as appropriate
- Work collaboratively with academic experts and broader research partners from across the Turing and the wider Turing community.
- Contribute to the life of the Institute and support a diverse and inclusive community through embracing the Turing values.
- Undertake required administration tasks
- Adhere to and promote principles of reproducible and ethical data science and ensure secure handling of data and health and safety in all aspects of work.

Specific to the project/programme/role

- Undertake high-quality research, under supervision, contributing to the broader research aims of the ELSA project
- Support the development and application of state-of-the-art approaches for fairness, transparency, robustness and/or privacy to tackle research challenges of the ELSA project. These might include:
 - Robust methods for practically useful transparency, examining the extent of leakage of private information
 - Appropriate methods to monitor and mitigate bias and unfairness
 - Ways to improve practical human-machine team performance
 - Formal verification and validation

The Alan Turing Institute

OTHER DUTIES

- Travel may be necessary to meet requirements of the role

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:		
Education/Qualification		
Masters degree or equivalent in Computer Science, Mathematics, Policy, Law, Ethics, or a closely related discipline	E	A
In the process of pursuing a PhD or interest in pursuing a PhD	D	A
Knowledge and Experience		
Experience in design, development and implementation of research software tools and libraries, such as Python, Java, GPU programming (Tensorflow, PyTorch etc), symbolic verification (SAT, SMT)	E	A I
A solid background in one or more of the following: technical approaches to fairness, transparency, robustness or privacy; probabilistic models, human-machine teams	E	A I
Hands-on experience with Machine Learning methods	E	A I
Experience of conducting academic research	E	A
A developing record of scientific publication, which may include journal articles, book chapters, and policy reports/white papers	D	A I
Experience or demonstrable interest in Policy, Law or Ethics	D	A I
Communication		
Excellent written and verbal communication skills with the ability to clearly communicate complex or specialist information/research findings clearly and concisely	E	A I
Ability to present and communicate complex information and conceptual ideas to diverse audiences	E	I
Liaison and Networking		
Proven ability to collaborate successfully with others and engage effectively with colleagues and research partners	E	I
Ability to work effectively as part of a team, providing support to others as required, with an approachable and flexible attitude towards work.	E	I
Participates in networks within the organisation or externally to share knowledge and information in order to develop practice or help others learn	E	I
Decision-Making Processes and Outcomes		
Ability to organise working time, take initiative to focus on the delivery of work.	E	A I
Ability to independently makes decisions which are low risk and that mainly affect themselves or a small number of people and are guided by regulation and practice	E	I
Analysis and Research		

The Alan Turing Institute

Ability to adopt appropriate data analysis methods for the purpose, and produce reports to present the findings	E	I
Ability to use own judgement to analyse and solve problems	E	I
Other Requirements		
Commitment to EDI principles and to the Organisation values (outlined below)	E	I

*May be required for some roles, recruiting manager will confirm at interview

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 3862 3536 or email recruitment@turing.ac.uk.

CLOSING DATE FOR APPLICATIONS: SUNDAY 08 SEPTEMBER 2024 AT 23:59 (LONDON UK BST)

We reserve the right to close this vacancy early or to interview suitable candidates before the closing date if enough applications are received.

TERMS AND CONDITIONS

This full-time post is offered on a fixed term basis until September 2025. The annual salary is £41,352 plus excellent benefits, including flexible working and family friendly policies, <https://www.turing.ac.uk/work-turing/why-work-turing/employee-benefits>

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 making sure our recruitment process is accessible and inclusive. This includes making reasonable adjustments for candidates who have a disability or long-term condition. Please contact us at adjustments@turing.ac.uk to find out how we can assist you.

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.