

The Alan Turing Institute

RESEARCH ASSOCIATE – PROJECT 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 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 several 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

At The Alan Turing Institute, the AI programme's goal is to advance world-class research into artificial intelligence, its applications and its implications for society, building on the wealth of expertise and knowledge across our academic network. Our programme expertise draws on the cross-disciplinarity at the heart of the Turing; bringing together technical experts to forge ahead with breakthroughs in fundamental AI research and its applications, with social scientists, ethicists, legal experts, industry, and policy-makers who are considering the implications of AI in real-world applications. The programme will contribute to leadership around this area, drawing on our links with industry and government to help ensure that the UK remains at the forefront of scientific innovation while building an ethical and regulatory framework for the use of AI that prevents misuse and inappropriate discrimination.

Our aim is to build a leading centre of excellence for technical work on safe and ethical AI, in line with the government's Industrial Strategy and in step with global demand for research and guidance in this domain.

Programme Director Adrian Weller has recently 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.

The Alan Turing Institute

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 Associate to support and enable the delivery of parts of the ELSA project for which Turing will be funded. The successful candidate will work with the Programme Director, Adrian Weller at the Turing and the other ELSA project partners.

The role 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.

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

- Develop and apply state-of-the-art approaches for fairness, transparency, robustness and/or privacy to tackle real-world-inspired 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
- Scope, pilot and deliver high quality research activity under the direction of the Programme Director
- Drive collaboration with academic experts and broader research partners from across the Turing and the ELSA community
- Publish and disseminate high-quality research papers and publications detailing research output and project case-studies.
- To participate in and to develop internal and external partnerships, for example, to identify sources of funding, generate income, obtain projects, or build relationships for future activities.

Other duties:

- Present, disseminate and explain our work at conferences and workshops, and internal and external events hosted by Turing and/or the partner.
- Contribute to the life of the Institute and support and help develop its community.

The Alan Turing Institute

Please note that job descriptions cannot be exhaustive, and the post-holder 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.

PERSON SPECIFICATION

Skills and Requirements Post holders will be expected to demonstrate the following	Essential (E) Desirable (D)	Tested at application(A) Tested at interview(I)
Education		
Research Associate level: PhD in Mathematics, Computer Science or a closely related discipline.	E	A
Research Assistant level: Near completion of a PhD or equivalent level of professional qualification in Mathematics, Computer Science or closely related discipline.	E	A
Knowledge and Experience		
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
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
Track record of the ability to initiate, develop and deliver high quality research aligned with the research strategy indicated by the PI and any industrial stakeholders and to publish in leading journals and conferences.	E	A&I
Hands-on experience with Machine Learning methods	E	A&I
Track record of outstanding research and in delivering impact appropriate to career stage	E	A
Experience in publishing research papers, code libraries or technical reports and giving presentations or classes on technical subjects.	E	A&I
Ability to rapidly assimilate new computational and mathematical ideas and techniques on the job and apply them successfully.	D	A/I
Ability to create and promote a collegial and collaborative approach to interdisciplinary research activities.	D	A/I
Communication		
Excellent writing skills and proven ability to communicate complex, specialist or conceptual information/research findings clearly and persuasively to diverse audiences, including the ability to explain technical concepts to technical and non-technical audiences.	E	I
Ability to write research reports and papers in styles accessible to both academic and lay audiences.	D	I
Analysis and Research		

The Alan Turing Institute

Ability to organise working time, take the initiative, and carry out research independently, under the guidance of the PI	E	I
Ability to use own judgement to analyse and solve problems	E	I
Liasion and Networking		
Participates in networks within the organisation or externally to share knowledge and information in order develop practice or help others learn	E	A/I
Decision Making Processes and Outcomes		
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	A/I
Other Requirements		
Commitment to meeting deadlines	E	I
Commitment to EDI principles and to the Organisation values	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 at recruitment@turing.ac.uk.

TERMS AND CONDITIONS

This full-time post is offered on a 3-year fixed-term basis. The annual salary is £38,850 - £46,200 (depending on experience) 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 £36,236 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 continuous eligibility 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.