

The Alan Turing Institute

SENIOR RESEARCH ASSOCIATE – 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 Senior 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

The Alan Turing Institute

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.

Research work will focus on understanding the potential and limits of third-party certification mechanisms in establishing AI assurance as well as on investigating the integration of technical methods with legal and ethical regimes to ensure that the design and deployment of AI systems will serve legal and ethical value.

In line with the focus and components of the project, the candidate is expected to have professional experience that provides suitable foundations for understanding technical as well as governance related aspects concerning the design, development and use of AI technologies, including ethical and regulatory considerations. They are expected to have a strong interest in existing and emerging legal and regulatory frameworks and requirements with relevance to AI technologies at the national and international level.

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 Senior Research Associate posts:

- Provide technical project management and leadership for research projects, ensuring successful outcomes
- Define the research direction in collaboration with Principal Investigators (PIs).
- Be a point of contact, supporting PIs in engaging with stakeholders regarding projects and deputising in meetings where necessary.
- Take the lead on writing up findings as they emerge, producing and developing reports, and developing publications in peer reviewed journals, in collaboration with the research team.
- Present, disseminate and explain our work at meetings/events and contribute to both the internal and external visibility of the Institute.
- Take responsibility for driving collaboration with academic experts and broader research partners from across the Turing, and the wider Turing / project community.
- Supervise the work of junior members of the team and provide guidance as required
- Line manage direct reports if required
- Contribute to the life of the Institute and support a diverse and inclusive community through embracing the Turing values.
- 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

- Play a leading role in undertaking high-quality independent research, actively contributing, and steering to the broader research aims of the Safe and Ethical AI Programme

The Alan Turing Institute

- 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

- Teaching may be required as part of collaboration work.
- Travel may be necessary to meet the 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		
A PhD (or equivalent experience and/or qualifications) in any of the following relevant areas: Policy, Law, Ethics, Mathematics, Computer Science or a 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
Substantial research experience in their field of expertise with commensurate output	E	A I
Evidence of high-quality publication(s) in a relevant field commensurate with your career stage	E	A
Experience in scoping and investigating independent research questions within a given overarching theme	D	A I
Previous experience overseeing and supervising more junior colleagues' work	D	A I
Communication		
Excellent written and verbal communication skills including the ability to present complex or technical information, and to communicate effectively with diverse audiences	E	A I
Ability to adapt the style of communication to the audience and ensure understanding	E	A I
Ability to negotiate and influence others	E	I
Analysis and Research		
Ability to carry out research independently and take the lead on writing up findings as they emerge	E	A I
Take the lead on research direction in collaboration with the PIs	E	I

The Alan Turing Institute

Project Management & Project Delivery		
Ability to initiate, plan and prioritise research projects and ensure effective use of research resources	E	I
Proactive approach to managing stakeholders and their requirements and identifying opportunities for collaboration	E	I
Ability to keep accurate and up to date knowledge of services available in own and related areas of work	E	A I
Willing to seek opportunities to make improvements and be proactive in identifying sources of funding/income	D	I
Teamwork and Motivation		
Ability to work effectively both as part of a team and in cross collaboration with other teams as required by the role	E	I
Ability to develop contacts and research collaborations within the Turing and wider community	E	I
Experience in managing and supervising resources	E	A I
Decision- Making Processes and Outcomes		
Ability to advise on available options for decisions that affect operational processes, considering any risks	E	I
Ability to work with others to make collaborative decisions	E	I
Teaching and Learning		
Teaching may be required as part of the role	E	A I
Other Requirements		
Commitment to EDI principles and to the Organisation values (outlined below)	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 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 £55,184 - £56,779 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.