

The Alan Turing Institute

Research Associates; Probabilistic Programming Language

BACKGROUND ON 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.

NEW PROBABILISTIC PROGRAMMING PILLAR

In 2022, the Alan Turing Institute signalled its intention to establish a portfolio of foundational AI research, which would complement the strengths of the institute around applications of AI and AI policy. An initial portfolio of research across three pillars, foundation models, game theory, and probabilistic programming, was launched in early 2023.

We are looking for two highly creative and motivated **Research Associates** to join the team. They will work on the broad "**Composable, Parallelisable and User-Friendly Inference**" theme within the Probabilistic Programming Pillar. All positions will involve direct collaboration with professors Hong Ge, Robert Goudie (University of Cambridge), George Nicholson (University of Oxford), William Moses (UIUC, US), and Paul Birrell (UKHSA). These Research associates roles will have the opportunity to interact with world-class research groups, including Cambridge Machine Learning Group (<https://mlg.eng.cam.ac.uk/>), Oxford CSML group (<https://csml.stats.ox.ac.uk/>), MRC Biostatistics Unit (<https://www.mrc-bsu.cam.ac.uk>) and other project partners.

DEVELOPMENT OF COMPOSABLE, PARALLELISABLE AND USER-FRIENDLY INFERENCE AND GROWING THE COMMUNITY OF THE TURING PROBABILISTIC PROGRAMMING LANGUAGE

Turing.jl (<https://turinglang.org>) is a general-purpose probabilistic programming language (PPL), written fully in Julia. Turing.jl has been used extensively for applied data science applications. Turing.jl is rapidly growing: it has been downloaded more than 16,000 times in 2022, has about 900 Slack members, 1,900 GitHub stars, and over 300 citations.

This project aims to extend Turing.jl to facilitate collaborative, multi-dataset analyses through compositionality and modularity; to enable Bayesian PPL-enabled analysis of huge data through parallelisation; and to improve its usability through community consultation to inform improved application programming interface (API) design.

Research Area 1: Compositional and modular modelling and inference

A key strength of PPLs is that they enable users to design customised probabilistic models that are tailored to their specific settings. Nevertheless, the same modelling patterns and components repeatedly arise in PPL models, but in existing PPLs such components are usually re-implemented from scratch by the PPL user in each model. This is cumbersome and wasteful of the PPL user's time and effort.

The Alan Turing Institute

The research associates will work on probabilistic programming techniques that better facilitate integrative and collaborative modelling and develop methods for composing modular Turing.jl models such as Markov melding (Goudie et al. 2019), cut and semi-modular inference (Nott et al. 2024). The candidate should have extensive experience in probabilistic programming and Monte Carlo/MCMC methodologies.

Research Area 2: Scalable approximate inference and MCMC

Performing probabilistic programming inference at scale is critical for real-world use cases, for example, pandemic forecasting (Nicholson et al. 2022), economic modelling (e.g. DSGE models used by central banks), and high throughput 'omics' (Berger, Yu, 2023). Inference using PPLs in such large-scale applications is often slow, and obtaining fast inference currently requires Turing.jl users to undertake onerous profile-optimize-parallelise steps for their specific PPL model.

The research associates will develop and implement scalability and speed improvements in Turing.jl, such as optimising automatic differentiation, and development and implementation of parallel algorithms. The candidate should have extensive experience in probabilistic machine learning, Gaussian processes, and approximate inference methods like sequential Monte Carlo and MCMC.

Research Software Engineering (RSE) support

Both associates will have access to significant dedicated Research Software Engineering support from the Alan Turing Institute throughout the project. The RSEs will support the research programme's emphasis on the dissemination of research results in the form of open-source software, together with publications at the highest possible standard. The Research Associates researchers will also have opportunities to work on community projects, including improving PPL usability and developing educational tools and materials to support the dissemination and adoption of probabilistic programming technologies.

DUTIES AND AREAS OF RESPONSIBILITY

The duties of the research associates will include:

- Work closely with other project team members.
- Keep on top of the state of the art in the relevant literature.
- Develop Turing.jl, as described in research areas 1 and 2 above and update its documentation accordingly.
- Help mentor the Research Software Engineering team members.
- Help to write academic research papers.
- Disseminate the research output to the research community, e.g., by giving talks at international conferences.
- Present, disseminate and explain our work at internal and external events hosted by the ATI.
- Contribute to the life of the ATI and support its community.

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 ATI and the role develop.

The Alan Turing Institute

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, Physics, Engineering, Statistics, Machine Learning, or a closely related discipline.	E	A
Research Assistant level: Near completion (thesis submitted) of a PhD or equivalent level of professional qualification in Mathematics, Computer Science, Physics, Engineering, Statistics, Machine Learning, or a closely related discipline.	E	A
Knowledge and Experience		
A solid background in one or more of the following: Bayesian inference, MCMC, graphical models	E	A&I
One or more topics in probabilistic machine learning, e.g. variational inference, state space models, sequential Monte Carlo, Gaussian processes, causality	D	A&I
Experience in design, development and implementation of research software tools and libraries	D	A&I
Deep and demonstrable experience of software development in at least one programming language, e.g. Julia, Python	D	A&I
Experience of GitHub or similar for distributed version control	D	A&I
Record of accomplishment of the ability to initiate, develop and deliver high-quality research aligned with the research strategy indicated by the PI and to publish in peer-reviewed journals and conferences; 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.	E	A/I
Ability to create and promote a collegial and collaborative approach to work, particularly in interdisciplinary research.	D	A/I
Communication		
Excellent writing skills and a 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	A/I
Analysis and Research		
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

The Alan Turing Institute

Liaison and Networking		
Evidence of participation within an organisation or discipline-related network to share knowledge and information, develop best practices, and facilitate learning among peers.	D	A/I
Other Requirements		
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.

CLOSING DATE FOR APPLICATIONS: Sunday 25 February 2024 (London, UK GMT)

TERMS AND CONDITIONS

This full-time post is offered on a 1-year fixed-term basis with the option to extend it for a further 2 years. 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.

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.