

The Alan Turing Institute

SENIOR RESEARCH ASSOCIATE: Probabilistic program scaffolds for large language models – NLP & LLM

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

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. The initial portfolio of research contains projects in two themes around foundation models and game theory. Five projects within a new probabilistic programming theme are being launched in early 2024. This role is being recruited to be part of one of the new projects in the probabilistic programming theme.

PROBABILISTIC PROGRAM SCAFFOLDS FOR LARGE LANGUAGE MODELS (Focus NLP/LLM) Project

This project aims to improve the reliability and consistency of large language models (LLMs) by integrating the LLMs with probabilistic answer set programming (P-ASP). This will allow for enforcement of logical, semantic, and syntactic constraints on LLM outputs. This role will involve developing methods to map between natural language and P-ASP frameworks, designing experiments to analyse LLM reasoning capabilities, developing methods for implementing P-ASP solvers to reformulate prompts, and develop strategies for constrained inference in LLMs.

The Turing Institute is hiring Research Associates to support this research project. This involves scaffolding large language models (LLMs) with declarative constraints to improve their reliability. There are two positions available. For this senior position, we are seeking candidates with strong backgrounds in natural language processing and experience working with LLMs. Responsibilities will include developing techniques to map between natural language and probabilistic programming frameworks, designing experiments to analyse reasoning capabilities, and incorporating multimodal data to further constrain LLM sampling. The ideal candidate

The Alan Turing Institute

will have expertise in analysing the inner workings of large language models as well as experience applying logical and semantic constraints to natural language generation systems.

ROLE PURPOSE

The successful candidate will focus on evaluating and enhancing novel methods with Answer Set Programming based approaches to improve the mathematical and logical reasoning abilities of LLMs. This work will involve developing techniques to map between natural language and declarative programming frameworks to impose logical constraints on LLM reasoning. The candidate will design experiments and benchmarks to systematically test and improve performance on various mathematical, logical, and commonsense reasoning benchmarks.

The candidate will collaborate with a vibrant team of researchers at leading universities (City, University of London, Imperial College London, and University of Leeds) to advance the state-of-the-art in robust natural language generation. This is an exciting opportunity for someone with expertise in analysing LLMs and experience applying semantic and syntactic constraints for consistent and robust reasoning.

DUTIES AND RESPONSIBILITIES

- Develop novel methods to map between natural language and probabilistic programming frameworks, building on the candidate's expertise with NLP and LLMs.
- Lead research projects on scaffolding LLMs with declarative constraints
- Design computational experiments and benchmarks to analyse commonsense reasoning capabilities of LLMs using P-ASP scaffolds.
- Collaborate with the project investigators to define the research direction and implement the objectives described in the proposal.
- Stay up to date on relevant literature at the intersection of foundation models, neuro-symbolic AI, and commonsense reasoning.
- Take the lead on disseminating findings through publications, conferences, and Turing Institute events.
- Act as a point of contact with stakeholders and supervise junior researchers on the project team.
- Work closely with cross-disciplinary team members to execute the work packages focused on robust mapping, LLM-ASP inference, and multimodal extensions.
- Contribute domain expertise in LLMs to advance the Probabilistic Programming research theme.

Please note that the 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.

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 / Qualification		
A PhD or equivalent qualification/experience in Mathematics, Computer Science, or a closely related discipline.	E	A
Knowledge and Experience		
A solid background in two or more of the following: use of Large Language Models such as GPT and open LLMs (such as Llama). Hands on experience with controlling inference in LLMs and transformer-based sequence to sequence models.	E	A
Track record of the ability to initiate, develop and deliver high quality research aligned with the research strategy and any industrial stakeholders and to publish in peer reviewed journals and conferences.	E	A/I
Hands-on experience with Machine Learning methods	E	A/I
Track record of outstanding research and 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
Experience in design, development and implementation of research software tools and libraries, such as Python, with specific knowledge of PyTorch and JAX).	D	A/I
Ability to rapidly assimilate new computational and mathematical ideas and techniques on the job and apply them successfully.	D	I
Ability to create and promote a collegial and collaborative approach to interdisciplinary research activities.	D	I
Communication		
Ability to communicate complex, specialist, or conceptual information clearly and persuasively to diverse audiences.	E	A
Ability to write research reports and papers in styles accessible to both academic and lay audiences.	E	I
Analysis and Research		

The Alan Turing Institute

Ability to organise working time, take the initiative, and carry out research independently	E	I
Ability to use own judgement to analyse and solve problems	E	I
Expert in data gathering and analysis, able to develop hypotheses to explain results and confidently present findings.	E	I
Liaison 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
A proven ability to collaborate successfully in a multidisciplinary environment to meet project objectives.	E	A/I
Teamwork and Motivation		
Ability to work effectively across disciplinary boundaries, both as part of an interdisciplinary team and in close collaboration with external partners in different disciplines.	E	I
Ability to develop contacts and research collaborations within the Institute and the wider community.	E	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 7 January 2024 at 23:59 (London, UK BST)

TERMS AND CONDITIONS

This full-time post is offered on a 2-year fixed-term basis. The annual salary is £53,576 - £55,125 (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>

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