

The Alan Turing Institute

RESEARCH ASSOCIATE: Probabilistic Program Scaffolds for large language models - Symbolic / Neuro-symbolic AI

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.

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.

The Probabilistic Program Scaffolds for LLMs 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 position, we are seeking candidates with strong backgrounds in answer set programming (ASP), probabilistic programming languages, or symbolic artificial intelligence. Experience with these declarative and logical approaches to AI is crucial for this research project on scaffolding large language models (LLMs). Responsibilities will involve leveraging the candidate's ASP, probabilistic programming, or symbolic AI expertise to map between natural language and formal programming frameworks. They will design experiments and benchmarks that use logical constraints to analyse LLM reasoning capabilities. We seek researchers with a deep understanding of symbolic methods like ASP and probabilistic programming to incorporate rigorous inferencing into neural LLMs. The candidate's knowledge of logically constraining AI systems will be essential for developing novel techniques to control and scaffold the behaviour of large generative models. Their expertise in the theory and application of probabilistic programming and logic-based programming (like ASP) is key for this project's focus on improving LLMs with logic-based constraints.

ROLE PURPOSE

The successful candidate will focus on evaluating and enhancing novel methods using answer set programming (ASP) to improve the mathematical and logical reasoning capabilities of large language models (LLMs).

Leveraging their background in ASP, probabilistic programming, or symbolic AI, they will develop techniques to map between natural language and declarative frameworks to impose logical constraints on LLM reasoning. The candidate will design experiments using logic-based scaffolds to systematically test and enhance performance on 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 the University of Leeds) to advance the state-of-the-art in robust and consistent language generation.

The Alan Turing Institute

This is an exciting opportunity for someone with expertise in constraint-based methods like ASP and experience applying semantic restrictions to align LLM outputs with robust, human-like reasoning. Their knowledge of symbolic techniques will be key for scaffolding LLMs to generate consistent and interpretable results across a variety of logical tasks.

DUTIES AND RESPONSIBILITIES

- Develop novel methods that apply declarative frameworks using the candidate's expertise in ASP, probabilistic programming, and symbolic AI to LLMs.
- Lead projects on scaffolding LLMs with logic-based constraints using ASP and probabilistic programming.
- Design experiments and benchmarks to analyse LLM reasoning capabilities using logic-based scaffolds.
- Collaborate to define research directions and implement objectives focused on neuro-symbolic integration.
- Stay current on literature related to integrating deep learning and symbolic methods.
- Disseminate findings leveraging knowledge of ASP and constraints for reliable natural language generation.
- Act as a point of contact, drawing on background in symbolic techniques to inform stakeholders.
- Supervise junior researchers, providing perspective from logic-based AI.
- Work closely with the team to execute work packages using declarative programming knowledge.
- Contribute ASP and symbolic AI expertise to advance the Probabilistic Programming theme.

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

A PhD or equivalent qualification/experience in Mathematics, Computer Science, or a closely related discipline.

E

A

Knowledge and Experience

A solid background in probabilistic programming, logic programming or symbolic models for artificial intelligence. Strong mathematical and logical reasoning abilities.

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 at least one of the following topics: ASP, Neural logic programming, inductive logic programming.

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 C++ and python with some experience 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

Ability to organise working time, take the initiative, and carry out research independently

E

I

The Alan Turing Institute

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 GMT

TERMS AND CONDITIONS

This full-time post is offered on a 2-year fixed-term basis starting in March 2024. The annual salary is £42,893 - £48,510 (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 £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 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