

THE UNIVERSITY OF MANCHESTER
PARTICULARS OF APPOINTMENT
FACULTY OF SCIENCE & ENGINEERING
SCHOOL OF ENGINEERING
DIVISION OF COMPUTER SCIENCE

RESEARCH ASSOCIATE IN SECURE AND VERIFIABLE AI MODELS

VACANCY REF: SAE-018999

Salary:	£33,309 - £40,927 per annum dependent on experience
Hours:	Full Time
Duration:	Fixed Term From 1 July 2022 until 30 June 2023
Location:	Oxford Road, Manchester

Enquiries about vacancy shortlisting and interviews:

Name: Mustafa Mustafa

Email: mustafa.mustafa@manchester.ac.uk

BACKGROUND

The EnnCore project aims to address a fundamental security problem in neural-based (NB) architectures, allowing system designers to specify and verify a conceptual/behavioural hardcore to the system, which can be used to safeguard NB systems against unexpected behaviour and attacks.

It pioneers the dialogue between contemporary explainable neural models and full-stack neural software verification. We, therefore, develop methods, algorithms and tools to achieve fully-verifiable intelligent systems, which are explainable, whose correct behaviour is guaranteed, and that are privacy-preserving and robust towards attacks.

OVERALL PURPOSE OF THE JOB

You will enjoy designing, developing, and evaluating novel AI models (deep neural networks) that are secure and robust against attacks. The project will involve continuous interaction with experts in explainable AI and formal software verification.

You will also have the opportunity to build use cases and collaborate with domain experts in cancer research and energy trading. You will design, develop and evaluate new models in the context of their accuracy, privacy protection, and robustness. This position may include research on various techniques such as federated learning, differential privacy, homomorphic encryption, formal verification, and adversarial methods. The post is initially for one year, with the possibility for extensions.

KEY RESPONSIBILITIES, ACCOUNTABILITIES OR DUTIES

The range of duties will include:

- Design, develop and evaluate new AI models and supporting methodologies
- Independent prototyping and empirical analysis
- Be involved in the supervision of projects
- Assist in the development of student research skills
- Conduct individual and collaborative research projects
- Write up research work for publication
- Continually update knowledge and understanding in field or specialism
- Translate knowledge of advances in the subject area into research activity
- Communicate material of a specialist or highly technical nature
- Liaise with colleagues and students
- Build internal contacts and participate in internal networks for the exchange of information and to form relationships for future collaboration
- Work with colleagues on joint projects, as required
- Collaborate with academic colleagues on areas of shared research interest
- Attend and contribute to relevant meetings
- Use new research techniques and methods
- Use initiative and creativity to identify areas for research, develop new research methods and extend the research portfolio
- Contribute to collaborative decision making with colleagues in areas of research

PERSON SPECIFICATION

Essential:

- Relevant PhD (or about to finish a PhD) in machine learning, privacy-enhancing techniques or formal verification
- Specialist knowledge in machine learning (either foundational or applied), preferably federated learning
- Specialist knowledge in privacy-enhancing techniques or formal verification
- Confident in prototyping complex systems independently
- Academic writing skills evidenced by publications at competitive conferences and journals
- Experience in designing or verifying privacy-preserving protocols, systems or AI models
- Excellent communication and interpersonal skills
- Excellent time management and organisational skills

- Ability to work independently and as part of a team
- Ability to liaise confidently and effectively with a range of individuals
- Flexible approach to dealing with research problems as they arise
- Willingness to learn and develop
- Ability to present in both written and oral publications
- Ability to meet deadlines
- Ability to assess and organise resources

Desirable:

- Previous academic work on privacy-enhancing techniques, privacy-preserving AI models, safety in AI systems, or adversarial methods
- Previous implementation experience with a formal analysis tool
- Knowledge and experience in reliability assessment of neural networks and concurrent systems.
- Publications in top conferences and journals in AI/ML (e.g., AAAI, IJCAI, ICML), systems security (e.g., IEEE S&P, EuroS&P, NDSS, CCS, USENIX security, PETS) or formal verification (e.g., TACAS, ACE, CAV, ICSE)
- Solid mathematical understanding of AI, formal verification or privacy-enhancing techniques