

This is a published notice on the Find a Tender service: <https://www.find-tender.service.gov.uk/Notice/020406-2024>

Planning

ARIA Safeguarded AI: Technical Area 3 Applications

ADVANCED RESEARCH AND INVENTION AGENCY

F01: Prior information notice

Prior information only

Notice identifier: 2024/S 000-020406

Procurement identifier (OCID): ocds-h6vhtk-047a89

Published 4 July 2024, 11:15am

Section I: Contracting authority

I.1) Name and addresses

ADVANCED RESEARCH AND INVENTION AGENCY

96 EUSTON ROAD,

LONDON

NW12DB

Email

tenders@aria.org.uk

Country

United Kingdom

Region code

UKI31 - Camden and City of London

Justification for not providing organisation identifier

Not on any register

Internet address(es)

Main address

<https://www.aria.org.uk>

Buyer's address

<https://www.aria.org.uk/programme-safeguarded-ai/>

I.3) Communication

Additional information can be obtained from the above-mentioned address

I.4) Type of the contracting authority

Other type

Non departmental body

I.5) Main activity

General public services

Section II: Object

II.1) Scope of the procurement

II.1.1) Title

ARIA Safeguarded AI: Technical Area 3 Applications

II.1.2) Main CPV code

- 73000000 - Research and development services and related consultancy services

II.1.3) Type of contract

Services

II.1.4) Short description

About the programme

Backed by £59M, ARIA's programme 'Safeguarded AI' explores if it's possible to formally verify the safety of AI systems through quantitative methods-specifically, exploring a possible pathway for developing a "gatekeeper" AI to understand the real-world interactions and consequences of an autonomous AI agent, and ensure it only operates within agreed-upon guardrails for a given application.

The programme is split into three Technical Areas (TAs) which will:

- TA1 (Scaffolding) build an extendable, interoperable language and platform to maintain real-world models/specifications and check proof certificates.
- TA2 (Machine learning) use frontier AI to help domain experts build best-in-class mathematical models for real-world complex dynamics and leverage frontier AI to train autonomous systems
- TA3 (Applications) unlock significant economic value with quantitative safety guarantees by deploying a gatekeeper-safeguarded autonomous AI system in a critical cyber-physical operating context.

II.1.6) Information about lots

This contract is divided into lots: No

II.2) Description

II.2.2) Additional CPV code(s)

- 38970000 - Research, testing and scientific technical simulator
- 48000000 - Software package and information systems

II.2.3) Place of performance

NUTS codes

- UK - United Kingdom

II.2.4) Description of the procurement

About TA3

Through TA3, we will look to demonstrate a 'gatekeeper' workflow that could be used to create and maintain decision-support tools and/or safeguard autonomous AI systems to

solve specific, economically valuable challenges.

As part of this work, teams will elicit functional and non-functional requirements from customers for a proposed product, then outline simplified versions for very concrete use cases. These specifications will serve as the benchmark evaluations for our entire research programme, as we'll look to develop AI systems specifically tailored to solve them.

A full set of applications will be determined after the TA3 solicitation, but some applications we're currently considering include: energy system optimisation; infectious disease epidemiology, climate and weather prediction; aircraft and spaceflight dynamics; and control systems for robots in human environments.

II.2.14) Additional information

Detail on TA3 solicitation:

In early summer, we will launch a solicitation for TA3, in which we're looking for individuals or existing entities (including start-ups, SMEs, non-profits, large companies) interested in using our gatekeeper AI workflow to build safeguarded products for specific tasks.

The solicitation will run in two phases:

Phase one: In the initial phase, we'll distribute funding to deeply understand customer needs and elicit requirements, this may include: sourcing datasets, designing evaluation suites to validate the performance of predictive models and autonomous or semi-autonomous controllers.

Phase two: Successful teams (estimated 2-4 groups total) will then advance to the second phase, where they will receive significant funding to apply the tools developed in the other Technical Areas to their application domain. Our thesis Safeguarded AI: constructing guaranteed safety (<https://www.aria.org.uk/wp-content/uploads/2024/01/ARIA-Safeguarded-AI-Programme-Thesis-V1.pdf>) gives a detailed overview of how this workstream compliments the whole programme.

ARIA Creators (recipients of ARIA's funding), gain access to the programme's research community, ARIA's programme leadership and wider community.

II.3) Estimated date of publication of contract notice

31 July 2024

Section IV. Procedure

IV.1) Description

IV.1.8) Information about the Government Procurement Agreement (GPA)

The procurement is covered by the Government Procurement Agreement: No