

# Opportunity Seeds: Mathematics for Safe AI

## Call for proposals

**Round 1**

**Date: 15.01.25**

v1.0

## Table of Contents

Summary	3
Scope	4
The application	5
Criteria	6
Timeline	7
Guidance	8
Part 1: Core application	8
Part 2: Timeline, budget and administrative questions	10
How to apply	12

## SUMMARY

### Intro to opportunity seeds

Breakthroughs require finding people who think differently about what's possible – and empowering them to follow their vision.

To support scientific and technological breakthroughs, outside of programmes, ARIA Programme Directors can award opportunity seeds to support ambitious research aligned to the opportunity spaces they've identified. Find out more about how opportunity seeds work within our model [here](#).

### 'Mathematics for Safe AI' funding call

Programme Director David 'davidad' Dalrymple's opportunity space, '[Mathematics for Safe AI](#)' asks "how can we leverage mathematics – from scientific world-models to mathematical proofs – to ensure that powerful AI systems interact safely and as intended with real-world systems and populations". We're initially looking to fund around 10 high potential proposals. Projects can be up to £500k in funding each<sup>1</sup>, though we expect that most will be smaller than that.

We're looking for bold ideas that could challenge assumptions, open up new research paths, or provide steps towards new capabilities. If you have an important idea you're obsessed with, but you don't currently have the resources or support to take it forward, we'd like to hear from you.

Ideas could come from anywhere, so we welcome applications from individuals and teams who are early in their career or who have atypical backgrounds. We care more about your idea and your intrinsic motivation than we do about your CV.

## SCOPE

Beyond the scope described here and our [evaluation criteria](#), we don't have preconceptions about what ideas you might send our way. You don't need to contort your project to fit our goals: we want to hear what you really want to do.

---

<sup>1</sup> inclusive of VAT (where applicable) and all other costs such as overhead.

**In scope:**

- + Ideas that sit within the '**Mathematics for Safe AI**' opportunity space. By this, we mean your proposal should show how your idea **either** aligns with **or** challenges the assumptions of the Summary, Beliefs, or Observations in the [opportunity space document](#).
- + Ideas that could change the conversation about what is possible or valuable.
- + Ideas that range from early stage curiosity-driven research through to pre-commercial science and technology.

**Examples of projects we'd be excited to fund include (but are not limited to):**

- + Creating (largely) self-contained **prototypes/minimal-viable-products** of a Safeguarded AI workflow, similar to [this example](#) but pushing for incrementally more advanced environments (e.g. Atari games).
- + Analysing, developing, prototyping and/or applying AI and formal methods, to:
  - + Scale up the production of **verified software systems**
  - + **Verifiably secure hardware**, such as interval-bound solvers for Maxwell's equations in semiconductor devices.
- + Developing **benchmarks and evaluations** of AI capabilities that accelerate or automate formal methods, e.g. in areas such as theorem proving, autoformalization, and verified software pipelines.
- + Mathematical (formal) theories for **understanding AI training** or **internal representations**.
- + Conceptual and mathematical work aiming to characterise "**boundaries**" of systems we want to protect from harm (e.g. Boundaries part 3a), or notions of "**harm**" (e.g. [Quantifying Harm](#), or [MWER](#) – which are not yet compatible).
- + Creating **pedagogical materials** (e.g. educational videos, interactive tutorials, high-production-value expository blog posts, etc.) relevant to research produced in the context of Safeguarded AI's Technical Area 1 (or other parts of the programme), potentially in collaboration with programme Creators. The goal is to help upskill a broader population of engineers and scientists to be able to derive value from using the programme's modelling framework.
- + **Field-building efforts** aimed at building a flourishing, creative, diverse and inclusive community around ideas related to Mathematics for Safe AI, be that in the form of events, upskilling opportunities, etc.

- + Good-faith and **constructive criticisms** or [adversarial collaborations](#) regarding key assumptions underlying the thesis of Mathematics for Safe AI, and/or Safeguarded AI.

### **Out of scope:**

- + Ideas that live outside of the scope of the '**Mathematics for Safe AI**' opportunity space.
- + Ideas that are within scope of the '**Safeguarded AI**' programme, as these should be submitted to programme-specific funding calls.
- + Ideas that are undifferentiated or that are likely to happen without ARIA support.
- + Commercial products, or projects for which you wish to retain exclusive IP rights.

**Budget and length of projects:** We provide funding from between £10k and £500k, inclusive of VAT (where applicable) and all associated costs (both direct and indirect). There is no minimum length for a proposed project; the maximum length for this call is two years.

**Intellectual property:** By default, opportunity seeds in the 'Mathematics for Safe AI' opportunity space will come with a contractual obligation to release all intellectual property open-source, open-access (CC-BY), and under permissive software licences (MIT + Apache 2), as applicable.

**Who can apply:** We welcome applications from across the R&D ecosystem, including individuals (those not affiliated with an organisation), universities (including proposals from students, postdocs and staff), research institutions, small, medium and large companies, charities and public sector research organisations.

If you are an overseas applicant you should note that our primary focus will be on funding those who are based in the UK or those willing to conduct all or part of the project from the UK. However, funding will be available to applicants outside the UK if we believe the proposed project can significantly benefit the UK ([See below for more info](#)).

## THE APPLICATION

We've designed a short and straightforward application. We hope you can spend no more than two-three hours writing it, though you may wish to spend more time thinking through your idea before applying.

## Part 1: Core application

This is the main part of the application, which consists of five questions about you and your idea. We'll base our initial shortlist and selection on this part, so try to make it clear why your idea is important.

Your proposal can be **up to three pages**. Optionally, you may include diagrams or graphics as part of your proposal.

The questions are:

*(click on the link for guidance on answering each question):*

1. [What do you want to do?](#)
2. [Why is it of scientific or technological importance?](#)
3. [Why hasn't it been done yet?](#)
4. [How do you plan to make it happen?](#)
5. [Why are you the right person or team to work on it?](#)

## Part 2: Timeline, budget and administrative questions

Here you can provide a rough timeline and budget. If you're not sure, you can put a ballpark estimate and we can work out details later in the process. We also need you to answer a few more questions, such as which country you plan to operate from. These can be found under [additional questions](#). All Part 2 content is filled and/or uploaded directly in the application portal and does not count towards the 3 page limit of the core application.

See the section on "[Guidance](#)" for information about *how* to submit your application.

## CRITERIA

Once you have submitted your application, it will pass through an initial review to ensure it is within the opportunity space identified above. At this stage we will carry out some checks to verify your identity and review any national security risks.

**If a submission is not considered to be within the opportunity space or is otherwise outside of the [scope](#), it will be rejected prior to a full review, and you will be notified accordingly.**

We will evaluate the proposal against the criteria outlined below, rather than scoring each answer separately. We expect proposals to spike against our criteria, and to have different strengths and weaknesses. Final selection will be based on an assessment of the proposal as a whole.

In evaluating your proposal, we'll consider the following criteria:

- a) **Importance and vision of the proposal.** We're looking for ambitious and/or consensus-shifting ideas that have the potential to change the conversation about what's possible or valuable. So we want to understand: if your idea worked out (and we know it might not), would you learn or demonstrate something important? Might it provide a step towards unlocking an important scientific or technical capability? If initial results are promising, would we be compelled to double down on the bet with additional funding or support?
- b) **Our confidence that the individual or team can deliver on the proposal.** This doesn't mean your CV or credentials. Instead, we'll be looking for whether you can demonstrate deep problem knowledge, a well-considered plan, and intrinsic motivation to work on the idea.
- c) **Relevance to ARIA.** We're looking to support ideas that would be unlikely to happen without us; for example, because they are particularly speculative or are underserved by other funders. We also want to enrich the opportunity space, so we will consider the proposal's potential to diversify skill sets and backgrounds that can contribute to the space.

As part of our evaluation, we will also consider:

- Whether the proposed plan or experimental design is likely to deliver on the goals of the proposal.
- The reasonableness and realism of the proposed budget and timeline, to the extent appropriate.
- If you're not currently operating in the UK we'll also assess the extent to which the proposed work will benefit the UK ([See below for more info](#)).

For more information on our opportunity seed review and selection process, see [here](#).

## TIMELINES

This call for opportunity seed proposals will initially be open for 4 weeks. We may update timelines based on the volume of responses we receive or other learnings.

Application opens	Clarifying questions deadline <sup>2</sup>	Application deadline	Shortlisting <sup>3</sup>	Final award decision <sup>4</sup>
14 January 2025	04 February 2025	11 February 2025 (13:00 GMT)	27 February 2025	07 March 2025

## GUIDANCE

### Part 1: Core application

#### How to format your proposal

- Content layout: *A short title, followed by each application question and your answer*
- Page count: *No longer than 3 pages (3 sides) of A4, including diagrams, excluding references*
- Font: *Garamond, Computer Modern, or Arial, black, min 11pt*
- Margins: *At least 0.5 inches on all sides*
- File type: *PDF*

#### Guidance on answering the questions

- Where possible, be concrete and provide justifications for any claims you make.
- Write your application with simple and clear language that would be understandable to an expert from an adjacent field.

<sup>2</sup> Clarifying questions received after this date will not be reviewed.

<sup>3</sup> Shortlisted / unsuccessful applicants notified. Shortlisted applicants invited to meet with Programme Director.

<sup>4</sup> Successful / unsuccessful applicants notified



- Remember to review the evaluation criteria shared here above.

In addition, we've provided a few prompts below that might help spark your answers.

Application questions	Guidance
What is the title of your proposal?	<ul style="list-style-type: none"> <li>• <i>Keep your title short, ideally less than one line.</i></li> </ul>
What do you want to do?	<ul style="list-style-type: none"> <li>• <i>What's the hunch or big idea you can't get off your mind? Is there a problem space you want to explore? A hypothesis you want to test? A new method, tool, or technology you want to create?</i></li> <li>• <i>How does your proposal align with or challenge the assumptions of the <a href="#">Mathematics for Safe AI</a>?</i></li> <li>• <i>What will be the output of your effort?</i></li> </ul>
Why is it of scientific or technological importance?	<ul style="list-style-type: none"> <li>• <i>If it works out, how would it disrupt the current frontier of knowledge or supersede the state-of-the-art?</i></li> <li>• <i>What capabilities or future work could it unlock?</i></li> <li>• <i>Why do you want to spend your time on this?</i></li> </ul>
Why hasn't it been done yet?	<ul style="list-style-type: none"> <li>• <i>Have recent developments only now made it possible? Has the research community overlooked it? Perhaps no one has put the pieces together like this before?</i></li> <li>• <i>How is your approach different or unusual compared to what's being funded or developed elsewhere?</i></li> </ul>
How do you plan to make it happen?	<ul style="list-style-type: none"> <li>• <i>Where will you start and what steps will you take?</i></li> <li>• <i>What are the technical challenges?</i></li> <li>• <i>What does success look like to you? If there are ways to measure it, do you have early thoughts on how you'll do that?</i></li> </ul>
Why are you the right person (or team) to work on it?	<ul style="list-style-type: none"> <li>• <i>Where does your insight come from?</i></li> <li>• <i>Do you have any examples that showcase your motivation, niche expertise, or ability to commit?</i></li> <li>• <i>If you're looking to extend your team, what key expertise should the next member bring?</i></li> </ul>

<p>[Optional] If you would like, you can include diagrams or graphics as part of your application.</p>	<ul style="list-style-type: none"> <li>• Consider including one or more visual aids that help us understand your innovation. Examples could include: a diagram, schema, graph, graphical abstract, etc.</li> </ul>
--	--

Please note: Supporting data, preliminary data, journal articles, blogs, code or other materials are not required, but may be referenced or linked to in the submission if they directly support your application. However, ensure that your proposal stands on its own.

## Part 2: Timeline, budget and administrative questions

### How to format your answers

You can fill out these questions directly in the application portal, so there's no need to format them in a particular way.

### Guidance on answering the questions

<b>Application questions</b>	<b>Guidance</b>
<p>Why is ARIA funding essential for the success of the project?</p>	<p><i>Our goal is to empower and champion R&amp;D Creators in pushing the boundaries of what's possible. To better understand your project, we'd like to know why it is unlikely to happen without ARIA's support, and why our funding is essential compared to other potential sources, such as government research grants or venture capital. For example, the project might be underserved by existing funding avenues, or it could be too early-stage for venture capital backing.</i></p>
<p>Can you tell us more about you and your team?</p>	<p><i>We want to know who will be doing the work (not just the principal investigator or project lead) and what portion of their time will be dedicated to this project. We usually prefer project leads or other key researchers to be spending at least 50%, ideally 80%, of their time on the project. The project lead is not required to have any particular degree of seniority, though we do want to see that they have the requisite skills, expertise and motivation. When forming your teams we encourage you to consider</i></p>

	<p><i>the diversity of your team and who is best placed to lead the project.</i></p> <p><i>We understand that early-stage ideas may not yet have a full team in place. So if you don't know the individuals, just list the required roles. An example of how to answer this question is below:</i></p> <table border="1" data-bbox="743 604 1308 865"> <thead> <tr> <th>Individual</th> <th>Role / expertise</th> <th>Already in place? If not, how long after project kickoff are they likely to start?</th> <th>FTE</th> <th>Total time on project (months, rounded)</th> </tr> </thead> <tbody> <tr> <td>Sophia Fleissig</td> <td>Synthetic biologist, project lead (TA1.2)</td> <td>Currently assigned to a different project but could transfer to this project with 6 weeks notice</td> <td>80%</td> <td>28</td> </tr> <tr> <td>Unknown</td> <td>Expert in plant tissue culture and transformation (TA1.3)</td> <td>To be recruited, aiming to start within 3 months</td> <td>100%</td> <td>33</td> </tr> </tbody> </table>	Individual	Role / expertise	Already in place? If not, how long after project kickoff are they likely to start?	FTE	Total time on project (months, rounded)	Sophia Fleissig	Synthetic biologist, project lead (TA1.2)	Currently assigned to a different project but could transfer to this project with 6 weeks notice	80%	28	Unknown	Expert in plant tissue culture and transformation (TA1.3)	To be recruited, aiming to start within 3 months	100%	33
Individual	Role / expertise	Already in place? If not, how long after project kickoff are they likely to start?	FTE	Total time on project (months, rounded)												
Sophia Fleissig	Synthetic biologist, project lead (TA1.2)	Currently assigned to a different project but could transfer to this project with 6 weeks notice	80%	28												
Unknown	Expert in plant tissue culture and transformation (TA1.3)	To be recruited, aiming to start within 3 months	100%	33												
<p><b>What timeframe do you think you'll need for your proposed work in months?</b></p>	<p><i>There is no minimum length for a proposed project. The maximum length is 24 months.</i></p> <p><i>We value speed, so if you think you could do the work faster with a greater budget, we'd prefer you to bias your proposal in that direction.</i></p> <p><i>It's ok if you're not sure – you can input a rough estimate or range.</i></p>															
<p><b>What key milestones will mark your proposed projects progress?</b></p>	<p><i>Our aim is to help your research get started quickly, so we're committed to a fast negotiation and award process. To support this, we'd like you to propose two to three milestones that will help track the progress of your project.</i></p>															
<p><b>How much funding do you need to bring your idea to life?</b></p>	<p><i>Please provide an estimate in GBP (inclusive of VAT where applicable and all other costs) of what you consider a reasonable funding amount for your goals.</i></p> <p><i>The minimum funding amount is £10,000. The maximum amount is £500,000.</i></p>															

	<p><i>Ask for the amount of funding needed to get your idea somewhere exciting. Don't be afraid to be ambitious. It's ok if you're not sure – give your best estimate.</i></p>
<p><b>What is the high level breakdown of the funding you need?</b></p>	<p><i>We understand that you may not know the exact breakdown. We're asking for an estimate now, so you're confident that your project budget will be sufficient to deliver your goals and minimise your need to rescope.</i></p> <p><i>To answer this question, provide a high-level breakdown of your proposed budget using the table provided in the application portal.</i></p>
<p><b>[Optional] If you would like, you can provide:</b></p> <ul style="list-style-type: none"> <li>● Alternative budget and timeline scenarios</li> <li>● And/or a breakdown of costs</li> </ul>	<p><i>This is not required. However, we recognise that you might want to provide two different scenarios such as a Proof of Concept vs a more expansive plan.</i></p> <p><i>If so, let us know what the two scenarios would look like, and which one you think makes the most sense.</i></p>
<p><b>Does your proposal depend on background (pre-existing) IP?</b></p>	<p><i>If yes, give us an indication of:</i></p> <ul style="list-style-type: none"> <li>● <i>What background IP is required</i></li> <li>● <i>Whether you currently have rights to that IP</i></li> </ul>
<p><b>Have you already secured grant funding for a similar project or are you currently in the process of seeking support from other funding sources for the same idea?</b></p>	<p><i>If yes, tell us more about the funding you already have or are applying for.</i></p>
<p><b>Are there any other factors or restrictions that might impact your freedom to operate and deliver the idea?</b></p>	<p><i>These might include, for example: potential conflicts of interest; import/export restrictions; legal and regulatory restrictions that you are aware of.</i></p>

<p><b>Are you planning to give a portion of the work to external subcontractors?</b></p>	<p><i>If yes, let us know what work you plan to give to a subcontractor.</i></p> <p><i>Subcontractors are any proposed third parties that you plan to enter into a contract or agreement with for services necessary for the delivery or management of the project.</i></p>
<p><b>Do you intend to carry out the project as part of an organisation or as an individual?</b></p>	<p><i>In the application portal, please select the option that best describes your situation and provide details where required.</i></p>
<p><b>Where are you currently operating from and are you proposing to perform the majority of the proposed project outside of the UK?</b></p>	<p><i>Our primary focus is on funding those who are based in the UK. For the vast majority of applicants, we therefore require the majority of the project work to be conducted in the UK (i.e. &gt;50% of project costs and personnel time).</i></p> <p><i>If you are currently not operating in the UK, but the majority of the work will be conducted in the UK, we would like to know the proposed plan for achieving this. For example, are you intending to establish operations inside the UK, or partner with a UK based organisation.</i></p> <p><i>We can also award funding to applicants whose projects will primarily take place outside of the UK, if we believe it is unique/compelling and/or the UK would miss out if not funded.</i></p> <p><i>If you're planning to carry out the majority of the project outside of the UK, we'd like to know your location and your perspective on how the funded project would benefit the UK either immediately or in the long term.</i></p> <p><i>Benefits may include:</i></p> <ul style="list-style-type: none"> <li><i>● enhancing economic growth in the UK</i></li> <li><i>● fostering scientific innovation in the UK</i></li> <li><i>● enhancing the quality of life in the UK (or both in</i></li> </ul>

	<i>the UK and internationally)</i>
Have you read and understood our funding terms?	<p><i>Our goal is to ensure your research can get going quickly, so we want to ensure a fast negotiation and award process. We aim to have agreements signed within 6 weeks, which we recognise can be much faster than standard at some organisations. Before proceeding, please confirm that you have read and understand our <a href="#">funding terms</a>. If you are unsure which terms apply to you, you can find more guidance <a href="#">here</a>.</i></p> <p><i>You can also find our funding principles <a href="#">here</a>.</i></p>

## HOW TO APPLY

**Before submitting an application, please read this call for proposals in full and the [general ARIA funding FAQs](#), where you can find further guidance on how ARIA funds research, including more information on intellectual property, budgeting, and our standard grant and contracting agreements.**

If you have any further questions relating to the opportunity seed call for proposals, please email [clarifications@aria.org.uk](mailto:clarifications@aria.org.uk). Any questions or responses containing information relevant to all applicants will be provided to everyone that has started a submission within the application portal and published on our website [here](#).

Please read the portal instructions below and create your account before the application deadline. In case of any technical issues with the portal please contact [clarifications@aria.org.uk](mailto:clarifications@aria.org.uk).

Application [portal instructions](#).

## [APPLY HERE](#)