


The Alignment Project

Application Guidance



Contents

Purpose of this document	3
What is The Alignment Project?	3
Who is involved?	3
Why apply?	3
What we are looking for	3
Key dates	4
Eligibility	4
What to expect.....	5
Application Process Application Overview	5
Application Process Stage 1 - Eligibility Statement	6
Application Process Stage 1 - Expression of Interest.....	7
Application Process Stage 2 – Full Application	8
Selection Process Overview.....	8
Selection Process Stage 1 - Initial Screening and EOI review	8
Selection Process Stage 1 - Assessment Criteria.....	9
Selection Process Stage 2 - Full Application Assessment	10
What happens next?	10
Further support.....	11

Purpose of this document

This document serves as a resource for applicants to The Alignment Project. It outlines the objectives of the fund, the application process, and how applications will be assessed.

What is The Alignment Project?

An international, cross-sector coalition offering funding of up to £1 million to advance the field of alignment. Advanced Transformative AI has the potential to deliver unprecedented benefits to humanity. But this future depends on ensuring powerful AI systems reliably act as we intend them to, without unintended or harmful behaviours. Without advances in alignment research, future systems risk operating in ways we cannot fully understand or control, with profound implications for global safety and security. The Alignment Project is a £15 million global fund aiming to accelerate progress in AI alignment research and ensure advanced AI systems are developed safely, reliably, and for society's benefit. We provide funding of up to £1 million to researchers from across disciplines.

Who is involved?

The Alignment Project is supported by an international coalition of government, industry, and philanthropic funders – including the UK AI Security Institute, the Canadian AI Safety Institute, Schmidt Sciences, Amazon Web Services, Halcyon Futures, Safe AI Fund, UKRI, Anthropic and the Advanced Research and Invention Agency – and a world-leading team of expert advisors. By fostering interdisciplinary collaboration, providing financial support and dedicated compute resources, we are tackling one of AI's most urgent problems: developing AI systems that are beneficial, reliable and remain under human control at every step.

Why apply?

The Alignment Fund will provide funding, typically ranging from **£50,000 to £1,000,000** per project. The coalition partners may consider funding higher value projects.

As a recipient you will receive:

- **Funding:** Up to £1 million for researchers across disciplines from computer sciences to cognitive science
- **Compute access:** Dedicated compute resources from AWS, enabling technical experiments beyond typical academic reach,
- **Venture capital:** Investment from private funders to accelerate commercial alignment solutions.

What we are looking for

Alignment research is complex and multi-disciplinary; we want to expand the scale and diversity of ideas and approaches being brought to the challenge. We're focused on funding projects to develop mitigations to safety and security risks from misaligned systems. Priority research questions include:

- How can we prevent AI systems from carrying out actions that pose risks to our collective security, even when they are attempting to carry out such actions?

- How can we design AI systems which do not attempt to carry out such actions in the first place?

We welcome applications from everyone—whether you're deeply versed in these research areas or just beginning to explore them. Solving this problem will mean bringing expertise to bear from researchers across many fields who may have never considered this topic before.

Key dates

Fund Launch	▪ 30 th July 2025
Applications open for round 1	▪ 30 th July 2025
Information Webinar	▪ 11 th August 2025 – 4.00pm GMT
End of Clarification Questions	▪ 13 th August 2025
Application Deadline for round 1	▪ 10 th September
Stage 1: Assessment of Eligibility and Expression of Interest (EOI)	
Outcomes from the Stage 1 assessment	▪ 17 th September
Stage 2: Full Application Form Submission.	
Shortlisted applicants will work with a Research Sponsor for 4 weeks to iterate and complete their full application.	▪ 15 th October
Funding decision: AISI will aim to notify applicants of the outcome within 3-5 weeks of full application submission.	
	▪ 5 th - 19 th November
Project kick-off meeting	▪ Within 2 weeks of signing the Funding Agreement.
Payment	▪ Payments will be made monthly in arrears, subject to submission of milestone-based reporting.
	▪ In exceptional circumstances, depending on need, we may consider one payment in advance of activity. This will be decided on a case-by-case basis and reflected in the Funding Agreement.

Eligibility

Eligible applicants: Researchers based at academic institutions, non-profit organisations and for-profit entities, subject to due diligence requirements.

Eligible scope: The Alignment Project funds research which contributes to AI alignment and control. The full scope is set out in the Research Priorities, available on the website. Where uncertain, we encourage you to apply, and we will take a decision on whether the proposal falls within the established scope.

What to expect

If your application to The Alignment Project is successful, you will receive a customised support package designed to meet the needs of your project. Rather than a fixed funding structure, we take a flexible, case-by-case approach to funding.

Your funding package may include:

- Cash funding to support salaries, operations, or research expenses.
- Compute credits to access cloud-based resources for training, experimentation, and scaling.
- Support from expert researchers, drawn from our network of AI alignment experts and domain specialists.
- We will work with successful applicants to design a package that fits your project's goals and practical needs.

Funding decisions will take into account:

- The amount requested and the rationale behind it
- The duration of your project
- Your organisational status (e.g. academic team, start-up, independent researcher)
- Your technical infrastructure requirements, including compute and specialist input

Application Process | Application Overview

You can apply for the Alignment Project through the DSIT Funding Portal.

To ensure applicants meet the eligibility criteria and align with the objectives of the Alignment Project, we have implemented a **two-stage application process**:

Stage 1: Eligibility Statement and Expression of Interest (EOI) Submission	<ul style="list-style-type: none">▪ All prospective applicants are required to complete an Eligibility Statement and EOI form.▪ This allows us to assess preliminary eligibility and determine whether applicants' proposals are suitable to proceed to the next stage of development and assessment.
Stage 2: Shortlisting and Invitation to Submit a Full Application	<ul style="list-style-type: none">▪ Following the review of eligibility and EOIs, shortlisted applicants will be invited to submit a full application.▪ As part of this process, they will have the opportunity to work closely with an assigned Research Sponsor to refine their proposals and help strengthen their submissions before final assessment.▪ The full application will include:

-
- Eligibility Declaration – Confirmation that there have been no changes since Stage 1 (EOI) that would affect their eligibility.
 - Full Research Proposal – An opportunity for the applicant to further develop and refine initial ideas from the EOI stage.
 - Budget – This will include a breakdown of anticipated costs, ensuring clear justification and alignment with the project’s scope and objectives.
-

**Submission of an EOI does not guarantee progression to the full application stage.*

Application Process | Stage 1 - Eligibility Statement

The first part of the application form requires applicants to confirm the Eligibility Statement and applicants are unable to proceed to the EOI form unless they complete the Eligibility Statement.

To be eligible for The Alignment Project the applicants must confirm that:

1. The **organisation** named in the application is one of the following:

Academic institution: defined as an institution that is eligible to receive funding from the research funding agency of the country in which it is based (e.g. eligible for Research Council funding in the UK, eligible for NSF/NIH funding in the USA, eligible for ERC funding in the EU, etc.)

Relevant non-profit: defined as an organisation that is a legal entity operated for a collective, public or social benefit and is registered as such in the country in which it is based (e.g. registered with the Charities Commission in the UK, registered as a 501(c)(3) in the USA, etc.).

Relevant for-profit: defined as an organisation that is primarily a commercial entity. These applications will be subject to additional terms according to the requirements of our funding partners.

2. The named person on the application has the authority to begin the application process on behalf of the affiliated organisation.
3. The named person on the application will provide AISI and the funding partners with all information required to undertake due diligence checks. AISI and the funding partners reserve the right to remove applications from the process if the information provided does not satisfy the due diligence requirements.
4. The named person on the application gives AISI and the funding partners permission to process data in this form in compliance with GDPR.

Application Process | Stage 1 - Expression of Interest

The Expression of Interest (EOI) form is a preliminary application designed to gather essential information about the applicant and their proposed project. It helps identify proposals that align with the objectives of The Alignment Project before proceeding to the full application stage. The EOI also allows applicants to outline their initial ideas, which may later be refined with input from the Research Sponsors. Below outlines the key information required in the EOI form and guidelines.

Contact Name	▪ Provide full name of the person responsible for the application.
Contact Title/Role	▪ Indicate your role in relation to the organisation.
Organisation	▪ State the name of the organisation you are affiliated with.
Organisation Country Location	▪ Enter the country where the organisation is based.
Relevant Regulatory Authority	▪ Provide the equivalent relevant regulatory authority (e.g. charity or research funding agency) in which your organisation is registered in your country.
Legal Registration Number	▪ Provide the registration number for the organisation / institution in the country in which it is based.
Email Address	▪ Enter a valid and active email address.
Orcid ID / Google Scholar / LinkedIn Profile Link	▪ To tell us about your career history, provide a link to your Orcid ID, Google Scholar, and/or LinkedIn profile.
Additional Collaborators	▪ List any key collaborators, including their names, roles, and their affiliations.
Relevant Research Outputs	▪ List your and your collaborators' top 5 previous research outputs of relevance to the proposed area of research, e.g. publications, patents, preprints, products. Include links to these publications (if applicable).
Project Title	▪ Provide a concise and descriptive title of your research in 1-2 sentences.
Problem Statement	▪ Outline your proposed project. You may wish to include: - What problem are you aiming to solve? - Why is this problem important and relevant to our Research Priorities? - How would you solve this problem? - How would this work be novel? The statement should not exceed 500 words.
Research Priorities Area	▪ Choose which area of the Research Priorities your proposal is mostly focused on.

Deliverables and Timeline	<ul style="list-style-type: none"> ▪ Outline the expected deliverables including timeline. This statement should not exceed 500 words.
Budget Narrative	<ul style="list-style-type: none"> ▪ Outline the budget and justification for the proposed research. This statement should not exceed 200 words. If shortlisted, you will be required to submit a more comprehensive budget based on the cost guidance (PFD available on website).
Additional Information	<ul style="list-style-type: none"> ▪ Share any additional information relevant to your application.

Application Process | Stage 2 – Full Application

Shortlisted applicants will be invited to submit a full application using a separate DSIT Funding Portal link and to collaborate with an assigned Research Sponsor. The full application consists of 3 sections and must be submitted on the DSIT Funding Portal.

1. Eligibility Declaration

- Confirm that there have been no changes since your EOI that would affect your eligibility. If there are changes, you must declare them and provide updated details in your Stage 2 submission.

2. Full Research Proposal

- The Full Research Proposal is an opportunity for you to further develop and refine initial ideas from the EOI stage.
- A full research proposal template will be provided in your Stage 2 application portal.

3. Budget

- This should include a breakdown of anticipated costs, ensuring clear justification and alignment with the project's scope and objectives.
- The budget must be completed in adherence to the cost guidance (PDF available on the website).

Selection Process | Overview

The selection process consists of two stages to assess eligibility, and the overall strength of each application score the quality against the Alignment Project Research Priorities.

Selection Process | Stage 1 - Initial Screening and EOI review

- Applications will undergo an **initial screening** to ensure they meet the basic eligibility requirements and pass initial due diligence checks.
- Eligible EOIs will then go through a **quality review** based on pre-defined selection criteria.
- Following this initial assessment, applicants will be notified of the outcome.

Selection Process | Stage 1 - Assessment Criteria

EOIs will be assessed based on the following five key criteria. Applicants should ensure their EOIs and full applications clearly address each area to strengthen their submission.

1: Is the proposal *relevant*?

- The proposed research must align with the objectives of the Alignment Project.
- We particularly welcome proposals that align to the Research Priorities.
- The problem statement should be well-defined, demonstrating a clear understanding of the challenges it seeks to address.
- The proposal should outline how the research contributes to advancing the Research Priorities, filling gaps in existing knowledge, or providing actionable insights for policymakers and practitioners.

2: Is the proposal *feasible*?

- The research plan should have clearly defined objectives, methodology, and expected outcomes.
- The proposal should demonstrate a realistic approach to achieving its goals within the proposed timeline and budget.
- The research should be technically and logistically feasible, with consideration of potential risks and mitigation strategies.

3: Is the proposal *innovative*?

- The project should introduce novel approaches, methodologies, or perspectives that advance research on safe AI development and/or deployment.
- The proposal should demonstrate originality and go beyond incremental improvements to existing work.
- Where applicable, interdisciplinary or cross-sector collaboration that enhances research impact is encouraged.

4: Is the proposal *actionable*?

- The research should have clear, tangible outputs in line with the Research Priorities.
- The proposal should outline pathways for impact, including how findings will be disseminated, shared with key stakeholders, and/or applied in practice.
- Consideration should be given to how the research aligns with industry or regulatory needs in order to safely develop and/or deploy AI systems.

5: Are the team and its requested resources *likely to deliver*?

- The research team should demonstrate the necessary expertise, experience, and track record to successfully execute the project.
- The proposal should clearly outline roles and responsibilities, including any key collaborators.
- The budget should be reasonable and well-justified, ensuring requested resources align with the scope of the project and are essential for its success.

Selection Process | Stage 2 - Full Application Assessment

- Shortlisted applicants will then **collaborate with their assigned Research Sponsors** to refine and develop their full application.
- Once the full application has been submitted, they will undergo a **final assessment** to be undertaken by a panel, independent of assigned Research Sponsors, against pre-defined selection criteria.
- **Each funder in The Alignment Project** coalition may review and score application according to their own criteria.
- Based on this review(s), final funding decisions will be made.
- Applicants will be informed of the outcome following the completion of the final assessment.

What happens next?

Once your application is submitted, it will undergo assessment as outlined in this document. If successful, the following are the expected next steps. Further guidance will be provided to fund recipient upon signing of agreement.

1. Funding Agreement and Project Kick-off meeting

- If your full application is successful, you will be issued a formal Funding Agreement (GFA) outlining the terms and conditions of the funding.
- You must review, sign, and return the agreement *within 2 weeks* upon receipt.
- The project is expected to kick-off *within 2 weeks* of agreement signed.

2. Performance Monitoring and Reporting

- The cadence of post-award performance monitoring will be confirmed as part of the proposal iteration stage, based on the level of input and support required from the Research Sponsor.
- We will provide reporting templates to ensure consistency and clarity.
- Monitoring may also include check-in meetings or reviews with the Alignment Project team and Research Sponsors.

3. Payment

- Payments will be made monthly in arrears, subject to submission of milestone-based reporting.
- In exceptional circumstances, depending on need, we may consider one payment in advance of activity. This will be decided on a case-by-case basis and reflected in the Funding Agreement.

4. Final Evaluation Report Submission

- Upon project completion, you must submit a final evaluation report summarising the research findings, deliverables, and overall impact.
- The final report will be reviewed by the Alignment Project team to assess whether all agreed outcomes have been met.

Further support

Webinar: An informational webinar will be held 11TH August, where applicants can learn more and ask questions. Register for the webinar here:

<https://events.teams.microsoft.com/event/fef1667f-5d7f-4e10-9f51-25ce024c9c30@cbac7005-02c1-43eb-b497-e6492d1b2dd8>

Clarification Questions (CQ): For clarifications, please refer to our CQ document on the website or email us at aisalignmentproject@dsit.gov.uk. To ensure fairness and transparency, we do not respond to individual emails but will update the CQ document if the question is relevant.