



Join the Team



Join the Team

[ERA](#) is a talent & research organisation based in Cambridge, UK for researchers, entrepreneurs, and policy-makers working to mitigate risks from frontier AI. Our mission is to advance the scientific & policy breakthroughs to mitigate catastrophic risks from frontier technologies.

In 2025, we ran one ERA:AI Summer Research Fellowship, **hosted 30+ speakers & workshops, organised a [Technical AI Governance Forum](#), and co-organised the 2025 [Vegas AI Security Forum](#)**. In 2026, we have already run four programmes, two cohorts each of the [AIxBiosecurity Fellowship](#) and the [ERA:AI Fellowship](#); we have launched a new [Frontier AI Security Residency](#) to help build the backbone of AI security & verification; and we are running another iteration of the [Technical AI Governance Forum](#). Alongside our core programmes, we are building stronger in-house research capacity, longer research positions for our top fellows/residents, and deeper collaborations with leading researchers and organisations across the field.

In 2027, we expect to scale up our work across AI safety, AI security, & biosecurity, including potentially expanding beyond research into more incubation support. Over the past 5 years, we have supported **over 200 researchers from 30+ countries** through our research programmes and conferences, leading to high counterfactual impact on their careers. We provide our fellows with mentorship from organisations such as [UK AISI](#), [RAND](#), [GovAI](#), [Google Deepmind](#) etc., and our alumni have gone on to work at some of the most impactful organisations in this space.

Open roles

We are currently hiring for the following positions:

- [ERA:AI Research Manager \(full-time\)](#)
- [ERA:AI Research Manager \(part-time\)](#)
- [AIxBio Research Manager \(full-time\)](#)
- [AIxBio Research Manager \(part-time\)](#)

The application deadline for all of the roles is Aug 28th, 2026.

Please contact hello@erafellowship.org if you have any questions.

*At ERA, you will help operate **one of the world's largest, longest-running** AI safety and governance talent programmes, and have significant autonomy to contribute to the strategic direction of a rapidly growing organisation working across some of the most important problems in frontier AI safety and security. We hope you consider joining us!*

Why Research Management at ERA?

At ERA, **Research Managers (RM)** provide the central intellectual backbone of all of our work. As a RM, you will have significant scope to shape our strategy and you will do much more than manage individual research projects. We care deeply about our mission, and for the right candidates, we are excited to give you substantial autonomy to think about the state of the world, form inside views on the most impactful research directions & interventions, etc., and drive it upward into our org-level strategy and thesis.

Alongside this, you will recruit and work with fellows who are highly motivated, intellectually ambitious and care deeply about their work. The role offers substantial scope to influence both 1) individual research projects with some of the most impactful orgs and mentors in the space, and 2) the research infrastructure for emerging fields which need to scale up significantly in a very short period of time.

Depending on your domain and strengths, you will have significant autonomy to:

- Help turn ambitious ideas into concrete, decision-relevant projects
- Identify neglected or unusually important research questions to help seed new directions
- Help exceptional fellows develop into independent researchers, research leads, founders or policy-makers
- Diagnose when projects are stuck, confused or aimed at the wrong target, and help redirect them
- Build relationships with mentors, researchers and organisations that can strengthen ERA's work
- Shape the intellectual direction of ERA's programmes and future initiatives

Application and Assessment Process

We recognise that this is a demanding process. We offer compensated work trials, and we will work with candidates to accommodate scheduling constraints.

1. **Initial Application:** due Aug 28th. The initial application stage only requires you to submit your CV/LinkedIn profile for our review, with a few other short questions about background and motivation.
 - *Applications will be reviewed on a rolling basis, so we encourage early applications.*
2. **Interview:** Following this, candidates progressing to the next stage will be invited to attend an interview with the ERA team on a rolling basis. We aim to send all interview invitations by the start of September, and complete all interviews in early September.
3. **(full-time roles only) Paid work test:** The final stage of the review process will be a paid work test, either a remote take-home task or an in-person work test at our Cambridge office. We will confirm the details after your interview.
4. **Reference Check & Offer:** In parallel with conducting the work tests, we will reach out to your references, and we aim to extend offers and final decisions by ~September 25th.

If you have any questions about the application process, please contact us at hello@erafellowship.org

Life at ERA

At ERA, we are a tight-knit team who all deeply care about our overarching mission of mitigating catastrophic risks from frontier AI. ERA has been thinking about AI Safety & Governance since before the launch of ChatGPT, and this makes ERA one of the world's largest, longest-running AI Safety & Governance talent programmes. You could help us 10x what we are doing, and we are really excited about bringing top talent onto all our teams.

ERA's [Culture & operating values](#) centre around 3 core ideas:

1. **Integrity:** Doing the right thing, especially when no one is watching.
2. **Intimacy:** Building genuine, psychologically safe relationships so we can tackle the hard problems together.
3. **Intensity:** Operating with focused urgency: because the stakes are high and the time is limited.

Benefits

Our benefits package currently includes:

- Competitive salary, dependent on experience
- Professional development budget of £5,000
- A generous productivity & equipment budget
- Employer pension contribution of 10%
- Four weeks of PTO recommended per year
- Private health insurance, including dental and vision
- A mission-driven, tight-knit team working on one of the most important challenges of our time

The start date is flexible, and we may be willing to wait for an extended period of time for the best candidate, though we'd prefer successful candidates to start as soon as possible after receiving an offer. **We are able to sponsor visas** and provide relocation support, so we strongly encourage international candidates to apply.

Diversity Statement

We encourage all qualified candidates to apply, even if you don't meet every listed requirement. We believe the strongest candidates often excel in some areas while showing potential for growth in others, and we value diverse perspectives and experiences in our hiring process.

Research indicates that candidates from underrepresented groups, including women, racial and ethnic minorities, and other marginalized communities, are statistically less likely to apply for positions unless they meet 100% of the qualifications. We actively encourage applications from candidates of all backgrounds, particularly those from historically underrepresented communities. We are committed to building an inclusive team and welcome applications from candidates who may bring unique perspectives, experiences, and skills to our organisation.

ERA:AI Research Manager (full-time)

[Apply Here](#)

Role Details

Start Date: 1st November 2026 (flexible, and prefer earlier start, e.g., 1st October)

Location: Cambridge, UK

Employment type: Full-time, in-person (visa sponsorship available)

Salary: £60,000 - £90,000 pro rata, depending on experience & scope of role

As an ERA:AI Research Manager, you will work closely with the ERA:AI programmes and research management team to support our flagship ERA:AI research programme. This includes but is not limited to: 1) guiding our fellows in finding & scoping impactful projects, 2) directing research across AI Safety and Governance, and 3) helping translate ambitious ideas into rigorous, decision-relevant work. You will have the opportunity to take up intellectual ownership of a portfolio of projects, supporting fellows develop credible theories of change, scope projects to an appropriate level of ambition, and navigate the many judgement calls required to produce genuinely impactful research. Building strong relationships with researchers, mentors and organisations across the field will also be a central part of the role. We expect this to be a high trust, high autonomy role.

For the right candidate, there's also an opportunity here to own the intellectual thesis underpinning our broader work in AI Safety and Governance. At ERA, we care deeply about mitigating catastrophic risks. [Frontier AI is moving from systems we supervise task-by-task towards agents that can operate over longer horizons, automate parts of AI R&D, and in controlled settings strategically evade oversight, turning alignment, control, security, verification and international coordination into increasingly concrete technical and governance problems.](#)

You can help ERA develop a clear and opinionated view of this rapidly changing landscape, and contribute to translating those views into concrete [programmes](#) and [initiatives](#), both within ERA and across the wider field, in the mission of mitigating risks from frontier AI.

Skills, Knowledge, and Expertise

We are interested in a wide range of candidate profiles and do not expect any one person to meet every criterion below. However, strong candidates will likely demonstrate many of the following:

Expected

- A **strong background in AI Safety & Governance**, alongside a deep understanding of the research landscape in this area, and connections with others working in this field
- **Commitment to supporting the development of early-career AI safety researchers**

- **Excellent research judgement**, including the ability to identify important questions, interrogate assumptions and distinguish interesting work from work that could meaningfully reduce catastrophic risk
- Experience conducting, mentoring, or evaluating research, including research proposals
- **A genuine commitment to the mission**
- Strong organisational and **time management skills**
- Excellent **written and verbal communication skills**
- Ability to work independently and as part of a team

Nice to Have

- Opinionated and well-calibrated views about which AI research agendas should be prioritised, and why, given rapidly improving AI capabilities
- The ability to develop clear and nuanced views on unfamiliar or rapidly changing technical and strategic questions
- Knowledge or experience spanning frontier AI
- Existing relationships across AI Safety, Governance, or adjacent research communities
- Experience directing researchers or managing a portfolio of projects, coupled with an instinct for identifying exceptional people and helping them develop into increasingly independent researchers and research leaders
- Familiarity with the challenges involved in conducting, communicating or governing sensitive and potentially dual-use research

Responsibilities

You will:

- **Review and evaluate applications to the ERA:AI Fellowship** and help identify candidates with unusual research ability, judgement, agency or potential.
- **Take intellectual ownership of a portfolio of fellows and research projects**, working closely with our mentors & institutional partners to steer impactful research by providing substantive feedback on research design, prioritisation, execution and final outputs.
- **Help fellows identify important questions and scope projects that fit their skills**, interests and intended pathways to impact.
- **Work with fellows to develop credible theories of change**, sharpen their research questions and select appropriate methods for research, as well as supporting dissemination strategy.
- **Diagnose when projects are blocked**, poorly scoped or aimed at the wrong theory of change, and help fellows redirect them.
- **Identify and recruit mentors with relevant expertise**, and build strong matches between mentors and fellows (building on ERA's existing network in this space).
- **Maintain trusted relationships** with fellows, mentors, researchers and organisations across the frontier AI safety communities.
- **Connect fellows with relevant experts**, collaborators, institutional partners and potential future opportunities, to contribute to our community building goals.

- Potentially take on a direct mentorship position for fellows
- Work with ERA's programme team to design workshops, seminars and other activities which strengthen fellows' research and professional development.
- **Identify recurring gaps or promising directions** across the field and help translate these into future research streams, programmes or collaborations.
- **Contribute to ERA's longer-term AI Safety and Governance strategy** and the development of new initiatives at the intersection of frontier AI and biosecurity.

If you think you may be a good fit but you're hesitant to apply because of certain preferences around the role, we'd be excited to start a conversation with you about how to make the role a strong mutual fit.

[Apply Now](#)

(applications due August 28th, 2026)

ERA:AI Research Manager (part-time)

[Apply Here](#)

Role Details

Start Date: 1st November 2026 (flexible, and prefer earlier start, e.g., 1st October)

Location: Remote or Cambridge, UK

Employment type: Part-time

Salary: £60,000-£75,000 pro rated, with some flexibility for exceptional candidates

As an ERA:AI Research Manager, you will work closely with the ERA:AI programmes and research management team to support our flagship ERA:AI research programme. This includes but is not limited to: 1) guiding our fellows in finding & scoping impactful projects, 2) directing research across AI Safety and Governance, and 3) helping translate ambitious ideas into rigorous, decision-relevant work. You will have the opportunity to take up intellectual ownership of a portfolio of projects, supporting fellows develop credible theories of change, scope projects to an appropriate level of ambition, and navigate the many judgement calls required to produce genuinely impactful research. Building strong relationships with researchers, mentors and organisations across the field will also be a central part of the role. We expect this to be a high trust, high autonomy role.

Skills, Knowledge, and Expertise

We are interested in a wide range of candidate profiles and do not expect any one person to meet every criterion below. However, strong candidates will likely demonstrate many of the following:

Expected

- A **strong background in AI Safety & Governance**, alongside a deep understanding of the research landscape in this area, and connections with others working in this field
- **Commitment to supporting the development of early-career AI safety researchers**
- **Excellent research judgement**, including the ability to identify important questions, interrogate assumptions and distinguish interesting work from work that could meaningfully reduce catastrophic risk
- Experience conducting, mentoring, or evaluating research, including research proposals
- **A genuine commitment to the mission**
- Strong organisational and **time management skills**
- Excellent **written and verbal communication skills**
- Ability to work independently and as part of a team

Nice to Have

- Opinionated and well-calibrated views about which AI research agendas should be prioritised, and why, given rapidly improving AI capabilities

- The ability to develop clear and nuanced views on unfamiliar or rapidly changing technical and strategic questions
- Knowledge or experience spanning frontier AI
- Existing relationships across AI Safety, Governance, or adjacent research communities
- Experience directing researchers or managing a portfolio of projects, coupled with an instinct for identifying exceptional people and helping them develop into increasingly independent researchers and research leaders
- Familiarity with the challenges involved in conducting, communicating or governing sensitive and potentially dual-use research

Responsibilities

You will:

- **Review and evaluate applications to the ERA:AI Fellowship** and help identify candidates with unusual research ability, judgement, agency or potential.
- **Take intellectual ownership of a portfolio of fellows and research projects**, working closely with our mentors & institutional partners to steer impactful research by providing substantive feedback on research design, prioritisation, execution and final outputs.
- **Help fellows identify important questions and scope projects that fit their skills**, interests and intended pathways to impact.
- **Work with fellows to develop credible theories of change**, sharpen their research questions and select appropriate methods for research, as well as supporting dissemination strategy.
- **Diagnose when projects are blocked**, poorly scoped or aimed at the wrong theory of change, and help fellows redirect them.
- **Identify and recruit mentors with relevant expertise**, and build strong matches between mentors and fellows (building on ERA's existing network in this space).
- **Maintain trusted relationships** with fellows, mentors, researchers and organisations across the frontier AI safety communities.
- **Connect fellows with relevant experts**, collaborators, institutional partners and potential future opportunities, to contribute to our community building goals.
- Potentially take on a direct mentorship position for fellows
- Work with ERA's programme team to design workshops, seminars and other activities which strengthen fellows' research and professional development.
- **Identify recurring gaps or promising directions** across the field and help translate these into future research streams, programmes or collaborations.
- **Contribute to ERA's longer-term AI Safety and Governance strategy** and the development of new initiatives at the intersection of frontier AI and biosecurity.

[**Apply Now**](#)

(applications due August 28th, 2026)

AlxBio Research Manager (full-time)

[Apply Here](#)

Role Details

Start Date: 1st November 2026 (flexible, can start earlier, e.g., 1st October)

Location: Cambridge, UK

Employment type: Full-time, in-person (visa sponsorship available)

Salary: £60,000 - £90,000, depending on experience

As ERA's AlxBio Research Manager, you will work closely with ERA's programmes and research management teams to support our AlxBio fellowship programme. This includes but is not limited to: 1) guiding our fellows in finding & scoping impactful projects, 2) directing research across AlxBiosecurity, and 3) helping translate ambitious ideas into rigorous, decision-relevant work. You will have the opportunity to take up intellectual ownership of a portfolio of projects, supporting fellows develop credible theories of change, scope projects to an appropriate level of ambition, and navigate the many judgement calls required to produce genuinely impactful research. Building strong relationships with researchers, mentors and organisations across the field will also be a central part of the role. We expect this to be a high trust, high autonomy role.

For the right candidate, there's also an opportunity here to own the intellectual thesis underpinning our emerging work at the intersection of frontier AI safety and biosecurity. At ERA, we care deeply about mitigating catastrophic risks, and [AI is beginning to compress both the pathway to biological harm and the pathway to biological defence](#). You could help ERA develop a clear and opinionated view of this rapidly changing landscape, and contribute to translating those views into concrete programmes and initiatives, both within ERA and across the wider field.

The role could therefore extend beyond the management of an individual fellowship cohort, and into shaping ERA's longer-term AlxBio strategy, identifying opportunities for new programmes and collaborations, and supporting exceptional individuals in building sustained, high-impact careers in the field.

Skills, knowledge and experience

We are interested in a wide range of candidate profiles and do not expect any one person to meet every criterion below. However, strong candidates will likely demonstrate many of the following:

Expected

- A **strong background in biosecurity or a closely related field**, alongside a deep understanding of at least one relevant research landscape
- **Excellent research judgement**, including the ability to identify important questions, interrogate assumptions and distinguish interesting work from work that could meaningfully reduce catastrophic risk

- Experience conducting, mentoring, or evaluating research, including research proposals
- **A genuine commitment to the mission**
- Strong organisational and **time management skills**
- Excellent **written and verbal communication skills**
- Ability to work independently and as part of a team

Nice to have

- Opinionated and well-calibrated views about which biosecurity research agendas should be prioritised, and why, especially given the rapidly improving AI capabilities.
- The ability to develop clear and nuanced views on unfamiliar or rapidly changing technical and strategic questions
- Knowledge or experience spanning both frontier AI and biosecurity.
- Existing relationships across the biosecurity, frontier AI safety, public-health or adjacent research communities
- Experience directing researchers or managing a portfolio of projects, coupled with an instinct for identifying exceptional people and helping them develop into increasingly independent researchers and research leaders
- Familiarity with the challenges involved in conducting, communicating or governing sensitive and potentially dual-use research

Responsibilities

You will:

- **Review and evaluate applications to ERA's AlxBio Fellowship** and help identify candidates with unusual research ability, judgement, agency or potential.
- **Take intellectual ownership of a portfolio of fellows and research projects**, working closely with our mentors & institutional partners to steer impactful research by providing substantive feedback on research design, prioritisation, execution and final outputs.
- **Help fellows identify important questions and scope projects that fit their skills**, interests and intended pathways to impact.
- **Work with fellows to develop credible theories of change**, sharpen their research questions and select appropriate methods for research, as well as supporting dissemination strategy.
- **Diagnose when projects are blocked**, poorly scoped or aimed at the wrong theory of change, and help fellows redirect them.
- **Identify and recruit mentors with relevant expertise**, and build strong matches between mentors and fellows (building on ERA's existing network in this space).
- **Maintain trusted relationships** with fellows, mentors, researchers and organisations across the biosecurity and frontier AI safety communities.
- **Connect fellows with relevant experts**, collaborators, institutional partners and potential future opportunities, to contribute to our community building goals.
- Potentially take on a direct mentorship position for fellows

- Work with ERA's programme team to design workshops, seminars and other activities which strengthen fellows' research and professional development.
- **Identify recurring gaps or promising directions** across the field and help translate these into future research streams, programmes or collaborations.
- **Contribute to ERA's longer-term AlxBio strategy** and the development of new initiatives at the intersection of frontier AI and biosecurity.

If you think you may be a good fit but you're hesitant to apply because of certain preferences around the role, we'd be excited to start a conversation with you about how to make the role a strong mutual fit.

Potential Focus Areas for Winter 2027

- Developing better computational tools for detecting novel pathogens in wastewater metagenomic data
- Developing practical know-your-customer tools for gene synthesis providers
- Building and improving guardrails for potentially dual-use biodesign tooling
- Building and improving robust evaluations of frontier models' CBRN capabilities
- Increasing adversarial robustness against dual-use model queries
- Building and improving machine unlearning methods for CBRN safety of frontier AI systems
- Other areas not mentioned here at the interface of frontier AI and Biosecurity

[Apply Now](#)

(applications due August 28th, 2026)

AlxBio Research Manager (part-time)

[Apply Here](#)

Role Details

Start Date: 1st November 2026 (flexible, can start earlier, e.g., 1st October)

Location: Remote or Cambridge, UK

Employment type: Part-time

Salary: £60,000-£75,000 pro rated, with some flexibility for exceptional candidates

As ERA's AlxBio Research Manager, you will work closely with ERA's programmes and research management teams to support our AlxBio fellowship programme. This includes but is not limited to: 1) guiding our fellows in finding & scoping impactful projects, 2) directing research across AlxBiosecurity, and 3) helping translate ambitious ideas into rigorous, decision-relevant work. You will have the opportunity to take up intellectual ownership of a portfolio of projects, supporting fellows develop credible theories of change, scope projects to an appropriate level of ambition, and navigate the many judgement calls required to produce genuinely impactful research.

Skills, knowledge and experience

We are interested in a wide range of candidate profiles and do not expect any one person to meet every criterion below. However, strong candidates will likely demonstrate many of the following:

Expected

- A **strong background in biosecurity or a closely related field**, alongside a deep understanding of at least one relevant research landscape
- **Excellent research judgement**, including the ability to identify important questions, interrogate assumptions and distinguish interesting work from work that could meaningfully reduce catastrophic risk
- Experience conducting, mentoring, or evaluating research, including research proposals
- **A genuine commitment to the mission**
- Strong organisational and **time management skills**
- Excellent **written and verbal communication skills**
- Ability to work independently and as part of a team

Nice to have

- Opinionated and well-calibrated views about which biosecurity research agendas should be prioritised, and why, especially given the rapidly improving AI capabilities.
- The ability to develop clear and nuanced views on unfamiliar or rapidly changing technical and strategic questions
- Knowledge or experience spanning both frontier AI and biosecurity.

- Existing relationships across the biosecurity, frontier AI safety, public-health or adjacent research communities
- Experience directing researchers or managing a portfolio of projects, coupled with an instinct for identifying exceptional people and helping them develop into increasingly independent researchers and research leaders
- Familiarity with the challenges involved in conducting, communicating or governing sensitive and potentially dual-use research

Responsibilities

You will:

- **Review and evaluate applications to ERA's AlxBio Fellowship** and help identify candidates with unusual research ability, judgement, agency or potential.
- **Take intellectual ownership of a portfolio of fellows and research projects**, working closely with our mentors & institutional partners to steer impactful research by providing substantive feedback on research design, prioritisation, execution and final outputs.
- **Help fellows identify important questions and scope projects that fit their skills**, interests and intended pathways to impact.
- **Work with fellows to develop credible theories of change**, sharpen their research questions and select appropriate methods for research, as well as supporting dissemination strategy.
- **Diagnose when projects are blocked**, poorly scoped or aimed at the wrong theory of change, and help fellows redirect them.
- **Identify and recruit mentors with relevant expertise**, and build strong matches between mentors and fellows (building on ERA's existing network in this space).
- **Maintain trusted relationships** with fellows, mentors, researchers and organisations across the biosecurity and frontier AI safety communities.
- **Connect fellows with relevant experts**, collaborators, institutional partners and potential future opportunities, to contribute to our community building goals.
- Potentially take on a direct mentorship position for fellows
- Work with ERA's programme team to design workshops, seminars and other activities which strengthen fellows' research and professional development.
- **Identify recurring gaps or promising directions** across the field and help translate these into future research streams, programmes or collaborations.
- **Contribute to ERA's longer-term AlxBio strategy** and the development of new initiatives at the intersection of frontier AI and biosecurity.

Apply Now

(applications due August 28th, 2026)