



# **PAUSEAI UK**

---

Donor Prospectus  
May 2026

# Contents

---

<b>Executive Summary</b>	<b>3</b>
<b>Foreword</b>	<b>5</b>
<b>Track Record</b>	<b>6</b>
<b>Theory of Change and Strategy</b>	<b>11</b>
<b>Past Expenses</b>	<b>18</b>
<b>Funding Targets</b>	<b>19</b>
<b>Team</b>	<b>22</b>

# Executive Summary

---

## Track Record

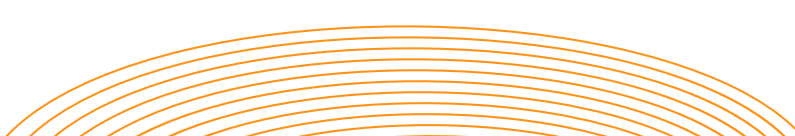
In the past year PauseAI UK has:

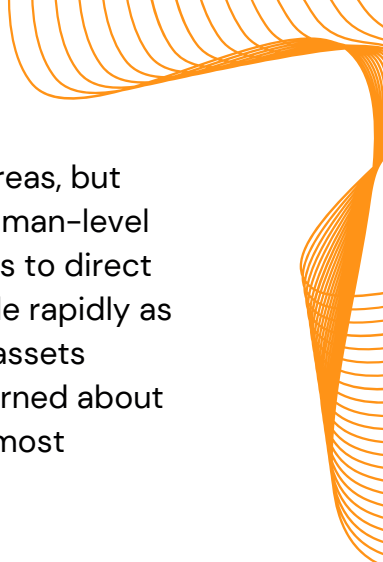
- **Delivered two conferences**, on behalf of PauseAI Global, bringing together volunteers from around the world, strengthening the global PauseAI community.
- **Written an open letter to Demis Hassabis signed by over 60 UK politicians**, establishing our political credibility and holding DeepMind to account.
- **Arranged a conference in the European Parliament** with speakers including **Brando Benifei**, the key architect of the EU AI Act and **Stuart Russell**, one of the world's top AI scientists, discussing concrete steps to reduce AI extinction risk.
- **Co-organised the largest AI protest in the world**, with around 300 volunteers marching to the London headquarters of Google DeepMind, boosting our visibility and building towards our long-term vision of a mass movement.
- **Hired a brilliant Deputy Director**, Matilda da Rui, and gained many new talented and dedicated volunteers to help execute on more projects.
- **Supported several volunteers in founding local groups** across the UK, laying the foundations for the infrastructure that will enable us to scale.

## Theory of change and strategy

We think that organising large numbers of citizens to boldly advocate for an AI pause will robustly help make the future go better. Public pressure for serious action on AI risks increases the likelihood of useful legislation and might be the only way that humanity avoids extinction.

We advocate for a global pause on AI development until there is broad consensus that superhuman general AI can be developed safely, with the benefits being widely shared. We are also in favour of other AI safety regulations, such as stronger liability for developers for AI-enabled harms.





We collaborate with other AI advocacy groups with different focus areas, but PauseAI UK positions itself as a movement focused on the risks of human-level and superhuman AI, rather than the current harms of AI. This allows us to direct our efforts towards the most severe issues, while also helping us scale rapidly as AI improves. PauseAI's strong SEO and name recognition are crucial assets because we automatically benefit when more people become concerned about AI risk. This turns AI companies and the progress of AI itself into our most effective marketing tool.

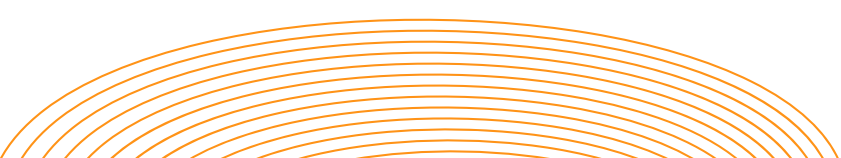
The UK is a valuable country in which to build a strong PauseAI movement because it has soft power with which it can help to convene discussions between the US and China, as with the first AI Safety Summit in Bletchley Park. London has the world's largest AI security institute and is a hub of AI research. There are more PauseAI members in London than any other city in the world. We have consistently held the largest AI protests since PauseAI began in May 2023, and we have seen a roughly exponential growth in the size of our demonstrations, with the number of attendees doubling approximately every 7 months.

We have recently launched our volunteering and project management platform, [Catalyse PauseAI UK](#). In the next year we will strengthen our local volunteer groups and organise a protest with over 500 participants. We will nurture our political connections, advocate for a global pause and support legislation that moves us meaningfully closer to real safety guarantees.

## Funding

The total cost for all of PauseAI UK's staff and activities is currently around £100k per year. To date, PauseAI Global has paid all of PauseAI UK's staff costs and other expenses. However, PauseAI is adopting a federated model in which national chapters operate as distinct legal entities and raise funding independently. Global is able to fund PauseAI UK until the end of Q2 2026.

We are seeking £425k in our main funding scenario which would allow us to support two additional full-time staff (for a total of four) for one year, alongside other part-time staff and consultants, while also providing an adequate budget for our operational costs and a major public event.



# Foreword

---

## One year of PauseAI UK

About one year ago, I started spending most of my time organising PauseAI UK. Before then PauseAI was only a project I did on the side. At that time our largest protest had seen fewer than 50 attendees, no prominent politicians or scientists were associated with PauseAI and I largely ran the UK chapter by myself.

In the past year PauseAI UK has delivered two conferences, written an open letter signed by over 60 UK politicians, arranged a conference in the European Parliament and co-organised the largest AI protest in the world. And we now have a very strong team, with Matilda joining as Deputy Director and several highly dedicated volunteers taking on substantial responsibility and launching their own local groups around the UK.

I'm proud of our track record and excited about the trajectory we are on. As AI capabilities improve exponentially, the number of people aware of the risks and motivated to take action increases commensurately. I believe we can harness this energy and turn it into real impact that actually improves humanity's chance of a positive future.

Joseph Miller  
Director of PauseAI UK

# Track Record

## June 2025 – PauseCon London

We delivered the first PauseAI conference, PauseCon, on behalf of PauseAI Global, bringing together around 60 volunteers from around the world for the first time and training them to be better organisers and communicators. We welcomed a range of excellent guest speakers from the AI safety community, including **Connor Leahy**, **Rob Miles**, **David Krueger** and **Kat Woods**.

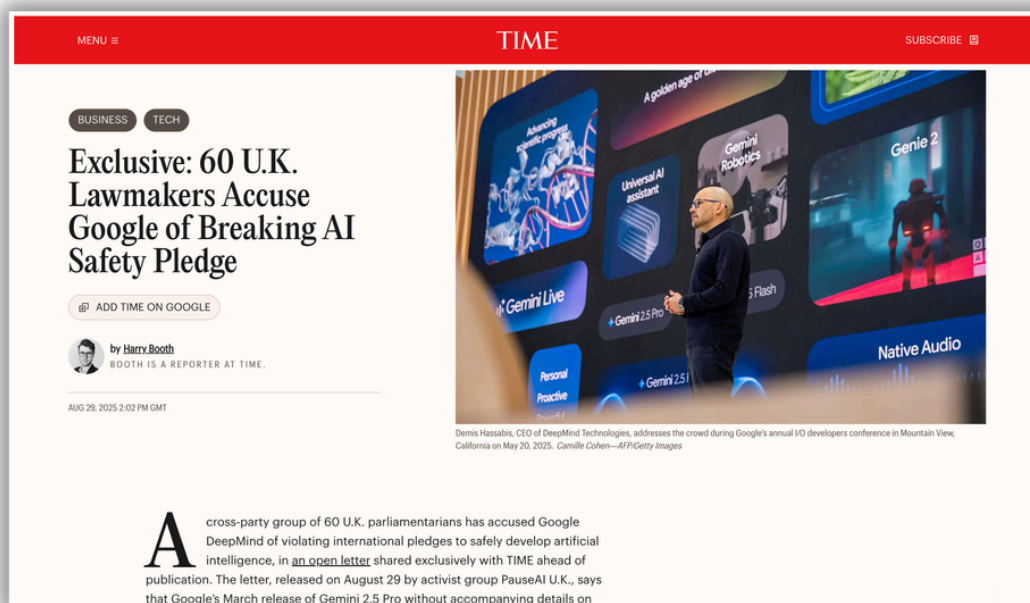
PauseAI Germany, among others, came away from the event with renewed purpose and went on to organise a petition signed by 150 German professors. One volunteer, Didier Coeurnelle, was inspired to initiate and fund the next PauseCon in Brussels.



## August 2025 – Open Letter to Demis Hassabis

In August we published an open letter signed by over 60 UK politicians, in response to Google DeepMind failing to uphold its AI safety commitments. Several of the MPs who signed later spoke in the Westminster Hall debate that we helped to organise in December (see below).

The article in TIME that broke the story established that Google DeepMind did not provide the UK AI Security Institute (AIS) with pre-deployment access to Gemini 2.5 Pro. Notably, Google did provide AIS with pre-deployment access to Gemini 3 Pro a couple of months after the letter was published.




## September 2025 – Book launch party

We held social events throughout the year, strengthening the sense of community that keeps people actively involved in PauseAI for months and years. One highlight was the book launch party for *If Anyone Builds It, Everyone Dies* in September.




## October 2025 – Documentary Screening in Parliament

In October we held a screening in the UK Parliament of filmmaker Michaël Trazzi's documentary about SB-1047, the proposed California AI legislation. This helped to inform MPs and Peers about the kinds of AI legislation that could be in a UK AI bill, and the battle with Big Tech that they should expect to face.



## December 2025 – Westminster Hall Debate

We proposed and helped to organise a Westminster Hall debate in Parliament on AI Safety. We wrote a memo which was sent to all MPs prior to the debate and drafted some of the speeches, putting us in a strong position to work with those MPs when proposing amendments to the Cyber Security and Resilience Bill.



## February 2026 – PauseCon Brussels

We delivered the next PauseCon in Brussels on behalf of Global with another two days of training workshops for PauseAI organisers from around the world.

The final day included a public conference in the European Parliament, featuring several prominent speakers, including:

- **Professor Stuart Russell**, author of the authoritative textbook on AI.
- **Brando Benifei MEP**, primary architect of the EU AI Act.
- **Victor Negrescu MEP**, Vice-President of the European Parliament.
- **Risto Uuk**, Head of European Policy at the Future of Life Institute.



Brando Benifei discussed the strengths and limitations of the EU AI Act candidly and argued that the Act is not merely a product regulation, but that the code of practice can be extended to cover internal deployment within AI companies. We hope that PauseAI will be able to work with Mr Benifei to help see such changes implemented.

Many volunteer projects were initiated over the weekend and several attendees have since held meetings with their own MEPs to follow up on the issues discussed.



## February 2026 – March for AI Safety

We co-organised a march past the offices of OpenAI and other Big Tech companies in King's Cross, London. **It was the largest ever protest focused exclusively on the risks of AI**, with around 300 people marching and media coverage in MIT Technology Review, The Independent, The Wall Street Journal and others.



The other organisers included *Pull the Plug*, a new group focused on the existing harms of AI. We consider the march a great success of coalition building between the historically opposed AI ethics and AI safety interests, with PauseAI and Pull the Plug represented in equal numbers.



# Theory of Change and Strategy



---

## Creating the political momentum for a pause


Organising large numbers of citizens to boldly advocate for an AI pause will robustly help make the future go better. Public pressure for serious action on AI risks increases the likelihood of useful legislation and might be the only way that humanity avoids extinction.

PauseAI UK exists to transform loose public concern into a focused political force in the UK, and to hold that pressure in place long enough to matter. Deep buy-in across the public is necessary to overcome industry lobbying. The work of converting awareness into durable political will is the community organising work that PauseAI UK specialises in.

## Our mission

The [proposal](#) on PauseAI Global's website outlines our primary policy goal. In brief, we are aiming for a global pause on AI development regulated by an international AI Safety Agency (AISA) that is responsible for determining when more powerful AI systems can be safely developed. Any sufficiently large group of countries would be empowered to veto the deployment of a superhuman AI system to ensure that, if some countries feel that they will be excluded from the benefits of AI, they have a strong negotiating position with which to demand their fair share. Or, if a group of democratic countries believe that an authoritarian country will deploy AI to oppress its own people, they can push for a deployment that empowers all citizens in every country.

Before safe AI is technically feasible, it is in the interest of all major powers to enforce the treaty globally. Once AI alignment is solved, AISA will control any superhuman AI prior to deployment and be able to use it to enforce the agreement.



Such a treaty is possible to enforce due to the highly centralised AI chip supply chain. Writing highly detailed policy proposals is not our comparative advantage, so we generally defer to other organisations such as MIRI for draft treaty texts and the precise specifications of enforcement mechanisms.

Having said that, it is very valuable for PauseAI staff to have a strong working knowledge of AI legislation and governance proposals in order to be credible in our discussions with politicians. In one instance, we wrote a summary of existing AI safety legislation for British MPs.

As mentioned above, we are in favour of other AI safety regulations, such as stronger liability for developers for AI-enabled harms. We may sometimes explicitly push for such policies both because they increase AI safety directly and they can be instrumental in increasing PauseAI's influence or credibility. For example, our open letter signed by 60 UK lawmakers criticised Google DeepMind for violating the Frontier AI Safety Commitments and helped to establish our voice in British AI policy.

## Positive outcomes for PauseAI UK

We cannot tell an exact story of what the path to a pause will look like, but we sketch below **two potential scenarios** where PauseAI UK has a positive impact. We are currently working towards realising both scenarios and in many ways they are complementary. At some point we may narrow our focus towards just one of these outcomes:

### Scenario 1: PauseAI as a special interest group

PauseAI UK has 10,000 highly dedicated volunteers who act as a dominant lobbying force on AI policy matters. Whenever a politician touches AI, they receive a policy document with PauseAI's view on the matter and constant communications from PauseAI volunteers, in the vein of the US sugar lobby:

*My phone did not stop ringing for the next five weeks. ... I had no idea how many people in my district were connected to the sugar industry. People were calling all day, telling me they made pumps or plugs or boxes or some other such part used in sugar production and I was threatening their job. Mayors called to tell me about employers their towns depended on who would be hurt by a sugar downturn. It was the most organized effort I had ever seen.*

Each MP has 20 PauseAI volunteers in their constituency who will send emails to their office and request meetings in which all 20 constituents will show up to express their views. PauseAI UK uses its Catalyse platform to coordinate its network to push the government to introduce an AI bill and ensure that it has the backing of every MP.

In the wake of an AI warning shot, PauseAI UK's volunteers contact every major British newspaper to ensure that journalists mention the idea of a global pause treaty in every major article about the incident. Protests are held outside Downing Street and any event the prime minister attends every day until they initiate negotiations for a global pause treaty.

## **Scenario 2: PauseAI as a mass movement**

PauseAI protests double in size every 7 months as AI capability itself improves exponentially. New PauseAI chapters are founded in every major UK city and many volunteers regularly put on talks in their local community to explain the risks of AI and recruit more volunteers for the movement.

At some point a significant, but not existential, AI catastrophe thrusts AI risks into the public consciousness and highlights the imminence of superhuman AI. Millions of British citizens become viscerally aware of the looming threat to their lives. PauseAI UK immediately announces a new protest and volunteers spread the sign-up page in their networks.

PauseAI UK organises a march in Westminster with 1 million attendees and dominates headlines in the British press. The prime minister is obliged to respond and commits to opening negotiations for a global pause treaty.

## **High-level strategy**

### **Brand and messaging**

PauseAI UK positions itself as a movement focused on the risks of human-level and superhuman AI, rather than the current harms of AI. This allows us to direct our efforts towards the most severe issues, while also letting us scale faster than movements focused on the existing harms of AI. PauseAI's strong SEO and name recognition are crucial assets because we automatically grow when more people become concerned about AI risk. This turns AI companies and the progress of AI itself into our most effective marketing tool.

A large fraction of our members have never been involved in grassroots advocacy before and we see this as a strength. It makes our protests more interesting to the media and makes the organisation more appealing to the silent majority who are not very politically active — unless, perhaps, they feel their lives are directly threatened.

We reflect our relatively moderate demographic in our messaging. We adopt a more measured tone than a typical advocacy group. Our imagery is positive and inspiring. We emphasise that we are taking the moral high ground and we represent universal, common-sense human values as part of a historic cause. This also reinforces our absolute commitment to non-violence.



Within the range of concerns around AGI, we encompass a broad set of risks. Many people will be more motivated by the threat of job automation or autonomous killer robots than extinction, because these risks are already becoming tangible and are easier to conceptualise. They are very important concerns in their own right and they are a good stepping stone towards confronting the risk of extinction. We present different AI risks as part of a single spectrum which we move along as AI becomes more powerful.

We are cognisant that building an AI movement in a context where many people have an incomplete understanding of the most severe risks requires caution and continual shaping of our message. Having our primary policy demand built into our name is a good safeguard against harmful distortions of our goals. In many contexts there are large short-term incentives to water down our demands and message, but we think this would reduce our long-term impact by moving our focus away from the most severe risks, so we are glad to have a name that commits us to a strong stance. We remain strictly non-partisan by focusing exclusively on our single issue and using politically neutral language.

### **Why the UK?**

PauseAI is starting chapters in every country and we think that having many different countries bought into the seriousness of AI risk will be critical to the success of a global treaty. But the UK is particularly valuable compared to other middle powers because it is a centre of AI research, including the headquarters of Google DeepMind and the second-largest offices of OpenAI and Anthropic. This soft power was demonstrated with the first AI Safety Summit in Bletchley Park.

Moreover, London is a hub of AI safety, with hundreds of AI safety researchers, the largest AI security institute and dozens of related organisations. The British public and political class are more aware of the risks of AI than those in comparable nations. Correspondingly, London has more PauseAI members and has consistently hosted larger protests than any other city in the world. These protests raise the bar for AI protests everywhere and can inspire others around the world to run bigger protests themselves.

## Growth

The fundamental bet of PauseAI UK is that there can be a very large and influential social movement dedicated to preventing the risks of advanced AI. Within PauseAI we already see evidence in our conversations with new members that a rapidly growing proportion of the population is truly grappling with the unprecedented danger that humanity is facing.


Global New Member Form Joins



We model the population as a bell curve with respect to the level of evidence that each person requires to become concerned about superhuman AI. As AI improves, we expect the fraction of the curve that has crossed the threshold of concern to increase accordingly. If capabilities continue to progress exponentially, the number of people worried about the situation will also grow commensurately. However, that concern does not automatically translate into well-coordinated action. Our job is to provide the infrastructure and guidance to turn that energy into impact.

We do not think that convincing more of the public to be concerned about AI risks is our comparative advantage at the moment. This is both because other organisations are already dedicating significant resources to mass communications and because we think that AI progress itself will be the primary driver of our growth. We benefit from being the largest AI protest organisation and positioning ourselves as focused on the risks of future AI, which naturally funnels people concerned about those risks into our ranks.

Instead, we see our role as maximising the utility of whatever level of concern already exists in the population at any given time, so that we can get to a pause as early as possible. This means always organising the biggest protest possible, providing excellent infrastructure, onboarding and support for individual volunteers and local chapter leaders, and planning our campaigns carefully.



Since PauseAI UK began, we have seen a (very) roughly exponential growth in the size of our protests, with the number of attendees doubling approximately every 7 months. New members register every day and chapters are popping up across the UK. If we can continue and accelerate this trend, then we expect to make substantial progress towards our goals in a relatively short span of time.

### Short-term plans

Two key players in the advancement of safe AI governance, Brando Benifei and Stuart Russell, spoke at our conference in the European Parliament in February. We want to organise an even more ambitious event in London in the next six months. As we have done before, we will use this event to bootstrap a protest held around the same time. It is generally much easier to get initial sign-ups for a conference or speaker event, and we can direct those people to also register for a protest at the same time. The aim is to hold a protest at least twice as large as our march through King's Cross in February.

We have recently launched our volunteering and project management platform, Catalyse PauseAI UK. This is allowing us to activate new and existing volunteers more easily by presenting them with a set of actions that they can take and a clear path towards contributing to bigger and more involved projects that make a significant impact. It will enable us to better coordinate grassroots lobbying efforts and empower highly motivated individuals to launch their own projects to which they can recruit other volunteers.

When cities hit a critical mass of enthusiastic volunteers, we launch new local groups in those cities. There are three core activities for local groups to engage in:

1. **Deliver a standard talk** designed to persuade people of the importance of AI risk. This can be given over and over at events for different audiences and in different venues.
2. **Lobby local politicians** to support our campaigns. Seek meetings, preferably in person, and have as many members as possible meet their MP and explain their concerns.
3. **Advertise protests** and help organise transport to get there.

We are currently working on proposing amendments to the UK's Cyber Security and Resilience Bill. One amendment that we think would be low cost and potentially impactful is to introduce a reporting pathway to the UK AI Security Institute so that they are informed about cyberattacks which use AI in novel ways. Another potential amendment would require critical cyber infrastructure to be capable of operating safely without AI (this is part of building the "pause button"). We have submitted written evidence to the parliamentary committee for the bill and we are currently contacting MPs and Peers who would introduce and support our proposed amendments.

We will continue to grow our social media presence and respond to outreach from journalists, where we expect this to help to grow our brand or credibility. We would strongly benefit from being able to hire someone to lead our content creation. It is particularly important for us to invest in improving our website with high-quality articles and A/B message testing, because organic search engine results are one of the main pathways through which new people join PauseAI.

We will also be putting significant effort into fundraising. From the end of Q2, PauseAI Global will no longer be able to fund our operations. We will add pages to our website outlining our theory of change and track record. We will ask for donations from our community, as well as applying to large philanthropic funds.

# Past Expenses

Calculating our historical costs is somewhat ambiguous as we have received significant assistance from PauseAI Global staff and to date we have officially operated as a department of PauseAI Global. Below are conservative estimates of our total costs during the past year, including the time PauseAI Global staff spent helping with events that we delivered on behalf of Global.

## 2025 Expenses

### PauseAI UK staff

€0

### PauseAI Global staff [approx.]

½ FTE for 4 months at €63,000 p.a. (inc. tax) - €10,050

### PauseCon London 2025

€18,830

### Other events [estimate]

€1,000

### Total for 2025 (rounded)

€29,900 EUR = \$34,500 USD = £26,100 GBP

## Q1 2026 Expenses

### PauseAI UK staff

Director - €5,000

Deputy Director - €10,000

### PauseAI Global staff [approx.]

½ FTE for 2 months at €81,000 p.a. (inc. tax and fees) - €6,750

### PauseCon Brussels 2026

€7,350

### Feb 2026 London Protest and other events

€1,830

### Total for Q1 2026 (rounded)

€30,900 EUR = \$35,600 USD = £27,000 GBP

### Annualised rate (rounded)

€123,600 EUR = \$142,400 USD = £108,000 GBP

# Funding Targets

To date, PauseAI Global has paid all of PauseAI UK's staff costs and other expenses. However, PauseAI is adopting a federated model in which national chapters operate as distinct legal entities and raise funding independently. **Global is able to fund PauseAI UK until the end of Q2 2026.**

Currency conversions are approximate.

## Minimum funding scenario (covers 1 year)

Our minimum scenario pays for our basic expenses and one small PauseCon event. Some travel budget is essential because we have local groups all over the UK and it is important for staff to be able to attend events outside their local area to strengthen the community.

### Staff

Director - €20,000 [N.B. Joseph also earns a PhD stipend]

Deputy Director - €50,000

Part-time administrative assistant - €10,000

Part-time social media specialist - €10,000

**Fundraising consultant** - €10,000

### Events

1x PauseCon event - €8,000 [N.B. PauseAI Global may provide additional funding for this event.]

Protests and local events - €2,000

**Travel** - €5,000

**Operational costs and buffer** - €5,000

### Total

€120,000 EUR = \$141,000 USD = £104,000 GBP

## Main funding scenario (covers 1 year)

Our main funding scenario would give us the resources we need to scale our local chapters and headline events faster than ever before. We have had to de-prioritise traditional and social media work while we have only two staff, but we are keenly aware of the potential for growth in this area, so we would like to hire a communications specialist. Similarly, a dedicated policy consultant with strong experience and connections in politics would help us reach our policy goals faster. Half-time local group leads would strengthen our volunteer community and help scale our grassroots actions.

### Staff

Executive Director - €40,000

Community organiser - €70,000

Communications specialist - €70,000

Administrative assistant - €40,000

2x half-time local group leads - €50,000

**Policy consultant** - €90,000

**Fundraising consultant** - €30,000

### Events

1x large PauseCon event - €40,000 [N.B. PauseAI Global may provide additional funding for this event.]

Protest and local events - €10,000

**Travel** - €10,000

**Coworking space** - €10,000

**Operational costs and buffer** - €30,000

### Total

€490,000 EUR = \$575,000 USD = £425,000 GBP

## Ambitious funding scenario (covers 2 years)

Our ambitious funding scenario would allow us to take on many more of the promising opportunities that we currently see for PauseAI UK. Offering higher salaries would make it easier for us to attract top-tier talent and by hiring more specialised content creators and allocating a substantial budget for video post-production, we would be able to grow our social media channels and create new funnels for recruiting volunteers. A dedicated community organiser would be responsible for assisting local groups and cultivating the social bonds that keep our volunteers consistently engaged.

A larger event budget would allow us both to hold larger conferences and to spend less time organising them (for example, by hiring a conference centre). With a large advertising budget we would hire a marketing agency to help us run highly impactful campaigns that catch the public imagination.

### Staff

Executive Director - €50,000

Comms director - €80,000

Community organiser - €80,000

Video content creator - €80,000

Events and operations director - €80,000

Administrative assistant - €60,000

5x half-time group leads - €125,000

**Policy consultancy hire - €200,000**

**Video editing and post-production - €50,000**

**Graphic design consultant - €20,000**

**PR and advertising agency hire - €250,000**

### Events

2x large PauseCon events - €150,000 [N.B. PauseAI Global may provide additional funding for one of these events.]

Protests and local events - €30,000

**Travel (for staff and external speakers) - €30,000**

**Coworking space - €40,000**

**Operational costs and buffer - €50,000**

### Total per year

€1,375,000 EUR = \$1,615,000 USD = £1,195,000 GBP

### Total for two years

€2,750,000 EUR = \$3,230,000 USD = £2,390,000 GBP

# Team

## Joseph Miller Director

Joseph is a researcher with over four years of experience in technical AI safety. He has been organising for PauseAI UK since May 2023 when he joined the very first protest. He is currently pursuing his PhD at Oxford University alongside his work for PauseAI UK. In 2021, he created one of the first text-to-image websites, which gained over one million users in less than six months. Before that he was a software engineer and studied mathematics.



## Matilda da Rui Deputy Director

Matilda's experience includes three years contributing to the formal AI alignment ecosystem, over ten years of entrepreneurial work, and direct experience in social movement strategy, nonprofit organising, interdisciplinary research, systems-oriented design, and operations.

