

Tab 1



ML4Good

Governance Curriculum Developer and Lead Instructor

Role Description

Summary

ML4Good runs in-person AI safety upskilling bootcamps. There have been 11 bootcamps so far, with the first camp outside of Europe taking place last summer in Brazil.

Our goals for the coming year are ambitious; we aim to bring ML4Good to many new regions, continue running camps in Europe, develop a governance bootcamp, and work with established organisations that could benefit from a bootcamp.

We're looking for someone to develop the curriculum for our first governance camp, based on our current technical curriculum and under the supervision of Charbel-Raphael Ségerie and Jonathan Claybrough. Approximately 40% of the content for the course will be rewritten for this camp, covering communication, governance and strategic thinking. The goal is to upskill participants in communicating AI risks effectively, developing strategic insights and engaging in policy-oriented actions. This will include working with people with specialist knowledge to develop and teach individual sessions on their area of expertise. This person will also act as the lead instructor for the bootcamp, supported by the specialists leading their sections as well as teaching assistants.

This is a part-time role while developing the curriculum, followed by full-time attendance as the lead instructor at the bootcamp June 16-26.

To apply for this role, please [complete this form](#). Applications will be assessed on a rolling basis until February 26th.

About Us

About ML4Good

ML4Good was started in 2022 by the AI Unit of EffiSciences (now the French Centre for AI Safety), and operated as a network of bootcamps organised by proactive individuals. We have recently set up an organisation to allow for full-time staff and longer-term funding. We've run bootcamps funded by Erasmus+ from the European Commission. Our goals for the coming year are ambitious; we aim to bring ML4Good to many new regions, continue running camps in Europe, develop a governance bootcamp, and work with established organisations that could benefit from our curriculum.

The first bootcamps were part of fieldbuilding research by EffiSciences, who concluded that ML4Good was the most effective fieldbuilding activity out of the many things they tried. Our alumni have gone on to work in the European AI Office, the French Centre for AI Safety (CeSIA), the Observatorio de Riesgos Catastróficos Globales (ORCG), SaferAI, and interning at Redwood Research, Timaeus, and MILA. Many have also gone on to participate in further programmes such as MATS and Pivotal.

Our goal is to transform participants' motivation into insights and concrete steps in 10 days. For example, during ML4Good Brazil, we created and reactivated 4 different AI Safety field-building groups in Latin America: Argentina, Colombia, AI safety Brazil and ImpactRio, and ImpactUSP (University of Sao Paulo).

We're proud to say we receive great feedback. Participants consistently rate us above 9/10 for recommending this bootcamp. And, without exaggeration, some participants told us ML4Good was the best experience in their life.

The Bootcamps

The bootcamp is a 10-day in-person intensive programme aiming to upskill researchers in the AI safety landscape, but also cover the basics of AI governance and strategic topics. Each day features a diverse set of activities such as technical coursework, presentations, group discussions, projects, and social activities in the evenings. We emphasize a hands-on application with a literature review project on one day, and a 2-day project at the end of the camp.

Our curriculum is continuously evolving to keep pace with the rapidly advancing field of AI safety. For an example, you can see the curriculum of our latest bootcamp [here](#).

Responsibilities

1. Curriculum design and development
 - Adapt existing ML4Good technical curriculum to a non-technical model focused on strategic communication and governance.
 - Around 60% of the curriculum can stay the same and 40% will need to be replaced or adapted.
 - Replace technical coding modules with hands-on workshops on communication, AI risk articulation, governance, and policy strategy.
 - Design modules covering:
 - Risk articulation and scenario planning
 - Strategic response and decision-making frameworks
 - Governance basics, policy consultation, and standards drafting
 - Practical communication for media, public speaking, and debate
2. Coordination with subject experts
 - Identify and recruit external speakers or instructors, including those with expertise in policy, governance, and strategic communication.
 - Collaborate with experienced instructors in AI safety, communication, and governance (both from our network and yours!) to deliver high-impact sessions.
3. Instruction and facilitation for the first camp
 - Lead instructional sessions on AI safety communication, governance strategy, and professional skills.
 - Facilitate dynamic workshops and interactive sessions designed to empower participants with skills in public speaking, media writing, and policy advocacy.
 - Provide one-on-one feedback and mentor participants to refine their presentations, pitches, and strategic insights.
 - Supported by specialists, TAs and the organisational team.
4. Program management and optimisation
 - Continually assess curriculum effectiveness through feedback and participant outcomes to refine and enhance the program.
 - Contribute to the overall strategy, ensuring program objectives align with ML4Good's broader mission to inform policy and reduce AI-related risks.

What we are looking for

We are particularly excited about working with individuals who:

- Are interested in developing the field of AI safety, and motivated by our vision and mission
- Have a background in AI safety and governance, policy, communication or a related field
- Are familiar with policy and governance issues in AI
- Are able to teach complex topics in AI safety and governance in an accessible way
- Are able to design workshops and exercises that cultivate strategic thinking
- Are organised and proactive - while you will be managed by the developer of the technical curriculum, much of the work will be independent

Logistics

- Contract length: 3 months, with full-time attendance at the bootcamp June 16-26th
- Hours: Part-time
 - This is flexible and we are open to multiple arrangements
 - We expect this to take 1 to 2 days a week until the camp, and then will require full-time in-person attendance at the bootcamp from June 16-26th
- Compensation: £7-9k, proportional to experience
- Location: Remote, able to travel to attend the bootcamp in France June 16-26th

Application

If you are interested in this role, please submit [this form](#) by February 26th.