Cause-specific fellowships

Description

A “cause-specific fellowship” (CSF) is a reading and discussion group which has features that make it similar to regular fellowships (intro and in-depth), but instead of being centred around the core principles or ideas of EA, they’re centred around a particular cause that is a candidate top priority cause for EAs.

The features they share with fellowships might include:

- Having weekly meetings for a set length of time (sometimes 8 weeks)
- Having small (~3-8) groups of participants with one facilitator per group
- Having a set ‘curriculum’ which is decided on beforehand
- Dividing the curriculum by topic, with a view to providing a coherent overview of the area (especially aimed at career relevance)
- Covering topics in an order that “makes sense”

Not all of the projects that I think fall under this umbrella are explicitly marketing themselves as fellowships, although some are.

Motivation

I’ve noticed that there’s at least several group leaders working on projects that fit this description. In general, I think co-ordination and communication across groups in the CB space is a really good thing that we should be trying to do more. In this specific case, I think co-ordinating about our decisions, tactics, and findings, could be valuable.

In part, this is because I think this might be a relatively new phenomenon (?) and so it could be useful for others considering running a project like this to have a centralised resource which sets out the case for and against this type of project - certainly, I think no such resource currently exists, and I’d have loved to have one before I ran my own CSF.

List of extant projects (AFAIK)

- Wild Animal Welfare (Bella Forristal, EA Oxford)
- Animal Advocacy (Sean Richardson, UCLA)
- Intro to AI Safety Workshops (Lewis Hammond and Ondrej Bajgar, EA Oxford/FHI)
- AI Governance Intro Course (AIGIC) (Anne le Roux, FHI)
- AGI Safety (Dewi Erwan & Richard Ngo, Cambridge)
- Effective Giving & Global Poverty (Dewi Erwan, Cambridge)
- AI (Stanford EA and UCLA)
- China and Global Cooperation (Dewi Erwan)
- Social sciences and X-risk (Stanford)
- Technical AGI Safety (Mark Xu)
- Value Alignment (Berkeley, in development)

Aims

Get better info on key uncertainties about this kind of project:

- Is it useful for attendees? How useful?
    - Does it inform career/donation/cause prio?
    - Does it create a cause-specific community? Is that a good thing?
- What are the major pitfalls to avoid? What caused the most friction?
- What were the most successful aspects? What worked unusually well?
- Would we recommend that other groups think about running CSFs?
    - Are they valuable in some causes, but not others? Why?

Tentative suggestions for next steps

Agree on a mechanism to co-ordinate

- Slack channel on EA Groups?
- Group chat?
- Shared Google Drive folder?
- Google group?

Co-ordinate on how we gather feedback

- Agree on a few quant questions we could all ask that would allow us to aggregate
    - E.g. Ask participants for likelihood to recommend

Work together to write something up about our findings once (some of) the projects are done

- Post it on the Forum?
- Or just on EA Hubs/the groups Slack?

Other programs we might be excited to see

- Nuclear weapons policy or governance
- Cellular agriculture (e.g. following GFI's syllabus, we're planning to get this started at some point)
- Biosecurity