

Effective Altruism Foundation

<http://ea-foundation.org>

Closing remarks


**Workshop for Reliable Artificial Intelligence 2017,
ETH Zurich**

Max Daniel

Research Project Manager, Effective Altruism Foundation



Growth in AI Safety Spending



Who is doing technical AI safety research?

- Google DeepMind
- OpenAI
- CHAI at UC Berkeley
- MILA
- Percy Liang at Stanford
- MIRI
- FHI (Oxford)
- ...

Useful resources

- [80K career review: AI safety](#)
- [80K AI strategy guide](#)
- [Open Phil's AI Fellows Program](#)
- [Paul Christiano's blog](#)
- [Arxiv Sanity Preserver](#)
- [FB group: AI safety discussion](#)
- [80K AI safety syllabus](#)
- [CHAI bibliography](#)
- [Allan Dafoe, Reading Guide for the Global Politics of Artificial Intelligence](#)

Thank you.

max.daniel@ea-foundation.org

