## The problem

Multiple organizations produce overviews of problem areas with breakdowns describing sub-problems. This includes the Open Philanthropy Project, 80,000 Hours, Giving What We Can, Sentience Politics, MIRI, GiveWell, D-Prize, Copenhagen Consensus, and the Good Technology Project.

This creates problems for both researchers and users of such content (of course, these are often the same people).

For the researchers, it means that work is duplicated, prior work is hard to reuse and integrate into new work, and related organisations don't default to collaborating on common problems.

For users, the content is scattered across multiple websites, which makes it inconsistent and awkward to consume. In addition, existing work isn't specific enough for some use-cases. For organizations that make practical recommendations (such as 80,000 Hours and the Good Technology Project), it is necessary to make *actionable* suggestions, which often means focussing on a specific sub-problem. An entrepreneur will have trouble taking action on "pandemic risk" but they could get inspired by browsing through a list of concrete sub-problems like "outbreak data visualisation tools for public health managers are hard-to-use" or "we have bad tools for monitoring the crossover of diseases from animal hosts".

## Our proposed solution

We should create a central problem repository that many researchers could contribute to. This would contain descriptions of problem areas, and a hierarchical structure of problem areas and their sub-problems. This should be a systematic answer to the question "what are the major problems in the world?".

For example, we might have a "Global health" problem area, which describes the general problem of improving global health, why there's usually a focus on the poor, etc. This might have a sequence of sub-problems: "Disease burden", "Malaria", "Inadequate access to bednets" (the concrete problem addressed by AMF), perhaps going even more granular (e.g. if "Inadequate production of bednets" were a real issue then that might be a sub-problem).

Additional information could be layered on top of this: for example, a prioritization organization might want to prioritize between the problems at each level, or globally. Having the prioritization information be separate allows organizations or individuals with different prioritization methods to share the problem breakdown.

For researchers, this prevents duplication of work between different organisations, improves research collaboration and specialisation, and exposes the research to a wider audience (especially if they aren't part of an established organisation).

For readers, this is a central place for them to find information about important cause areas and means they don't have to spend time searching across multiple websites for information worrying that they've missed something.

For organizations that make practical recommendations, this is especially useful since it provides problem area recommendations at whatever level of granularity is appropriate. A student might use the higher level breakdown to decide what courses to study, whereas a technology entrepreneur might be interested in quite granular problems that they might be able to solve.

It also makes it easier to start "technique specialist" recommendation organizations that focus on solving problems with particular techniques. The Good Technology Project fills this role for technology, but we could have similar organizations for policy etc. All such organizations will need a repository of problems to direct people to work on, and having this provided would allow them to focus on the specifics of how to solve problems with their technique.

## Unsolved problems

- Platform:
    - Where should this information be stored?
    - How should it be displayed?
    - How do we integrate information from other web sources?
- Structure:
    - How do we organise this information?
    - Should it be a tree or a network?
    - What information should be in the structure, as opposed to outside it?
- Collaboration:
    - Who should be able to contribute?
    - How should they contribute?
    - How do we deal with disagreement?
- Organization:
    - Should the project be "owned" by some organization? CEA?
    - How do we get experts involved to make sure the content is high-quality?
- Interest/Demand:
    - What would make research organisations really want this?
    - What would encourage experts outside of EA to contribute to this?
    - What would make readers and practitioners really want this?
- Prior attempts
    - Multiple people have tried this before. Why did they fail?

# Prior art

- People and organizations who've worked on similar problems
    - Cause prioritization wiki
        - Unstructured
        - Insufficient content
        - Not granular enough
    - Cause prioritization orgs (OPP, GPP)
        - Not shared
        - Not granular enough
    - Samuel Hargestam
- People and orgs who have thought about how to make the underlying platform (often referred to as knowledge mapping or argument mapping)
    - Arbital
    - Jacob Cole
    - Workflowy
    - Smallest Federated Wiki