



# AI for Good, at Scale?

How can AI help address climate change, poverty, hunger and other sustainable development goals?



**Trent McConaghy**

Founder, Ocean Protocol

## Abstract

The growing “AI for Good” movement is about using AI (artificial intelligence) to catalyze sustainable development for the planet. This article explores how AI for Good might scale. It starts with community, including AI Commons. Then, it describes a funnel to track ideas from inception to deployment at scale. Finally, it describes the gap between problem owners and problem solvers, and how to close the gap.

## Introduction

In Berlin, every single Friday, thousands of students go on strike. They see that the older generations are passing climate debt and other debts on to them. They see that they’re getting a raw deal. And they’ve had enough. They demand a voice, but are too young to vote. But they can strike. So they are.

It’s not just Berlin. [Millions](#) of children in [thousands](#) of cities are now striking each Friday. They demand a say in the future. They’re right.



Opening rally of the FridaysForFuture demonstration on January 25, 2019 in Berlin (image by [Leonhard Lenz](#), [CC0](#))

We all need and want a future, as individuals. So do our children, our grandchildren, and so on. For that, we need the planet to be sustainable, while continuing development for all. **We need sustainable development.** Growth is not a deity; our children's future is.

Alas, the response in terms of political will has been tepid. Will we see support from elsewhere?

## The UN SDGs

There's a bright spot. A big one. In 2015, the United Nations agreed on 17 [Sustainable Development Goals \(SDGs\)](#). Dozens of UN agencies and affiliates are backing the SDGs, not to mention many NGOs, governments, and funding bodies. The SDGs are goals for the year 2030 that include "no poverty", "zero hunger" and "climate action". If we miss the climate SDG we'll be boiled frogs, and not just metaphorically speaking. The other ones really matter too. Let's try to hit them!

We can think of the 17 UN SDGs as [KPIs for civilization](#) (Key Performance Indicators). For once, we actually have targets to go for in terms of good for the planet. You can actually measure them: there are >180 different measures towards seeing how well we're doing across the SDGs. It's a **direction for good**

as opposed to, say, letting social media turning us all into [shallow-thinking zombies](#).

## AI for SDGs? (For Good?)

We know that artificial intelligence (AI) is a powerful set of tools. We can ask: can we use AI to drive towards the UN 17 Sustainable Development Goals? More compactly, let's pragmatically define "good" as "helps at least one SDG" or similar (see note 1). Then:

### - Can we use AI for Good?

It's a great question. People have started to ask it. And, we're getting some preliminary answers. For example, people are using satellite imagery to start to track how reforestation is progressing. They're tracking livestock as a means of predicting conflict. Talk about nonlinear interactions! People are using AI for smarter traffic signals, to reduce congestion & pollution in cities.

And overall, it's towards this long term "for good" vision—AI has great potential to bring us abundance, to reduce costs by not just 10 percent or 50 percent but by 90 or 99 percent. And what does it look like? If we're not careful, all the efficiency gains will go into the [hands of a few](#). But if we are thoughtful about it, then maybe we can channel that to **spread the benefits of AI to the whole planet**.

Let's summarize so far. To the question "can AI help for good?" And the answer is yes. I gave a few examples. There are many out there!

## Moment of Impact?

For all this initial excitement, we haven't seen massive impact yet. It's because **AI for Good hasn't scaled yet**.

How do we know **when** we've hit scale? How about this: 50 for-good AI projects deployed, each to 100 jurisdictions. That's 5000 deployments of AI for Good. That's scale. While the numbers are a bit arbitrary, it crystallizes our goal.

So how do we achieve that specific goal? Here are some answers:

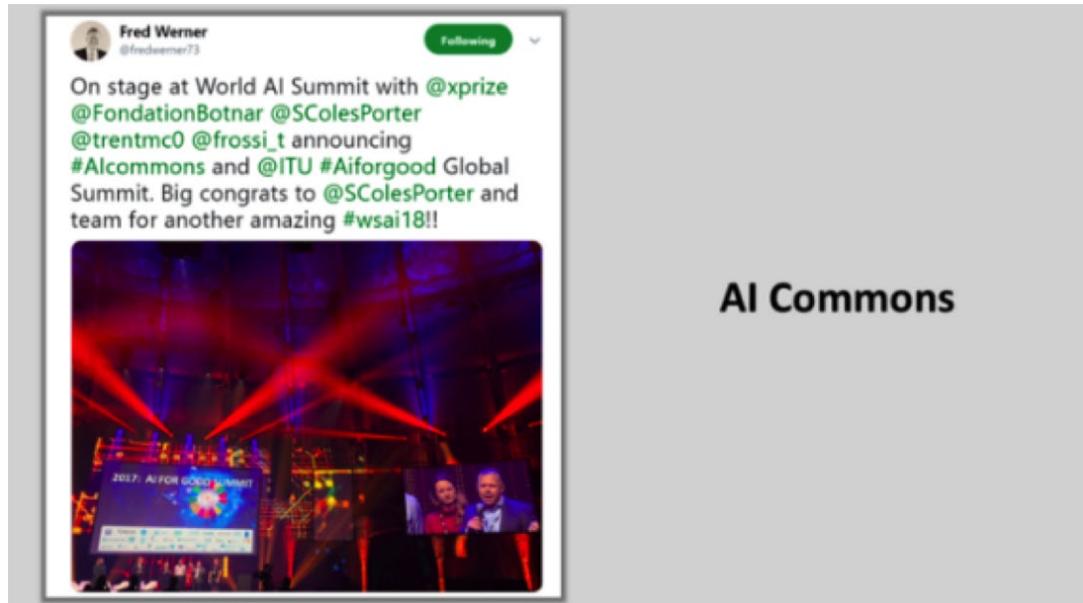
- first and foremost, a **community movement**
- working through a **funnel** that tracks each AI for Good project from inception to deployment at scale
- a key piece being to help **connect problem owners with problem solvers**

The rest of this essay elaborates on each of these.

## #AIforGood Community Movement

The "AI for Good" idea is spreading. There's now an [AI for Good Global Summit](#). It was initiated in 2016 by XPRIZE Foundation and ITU (the telecommunications arm of the UN), with the first annual summit was in May 2017. There are AI for good workshops, [like at NeurIPS](#). And there's an active hashtag [#AIforGood](#) ;) The question then arises: how do we coordinate action around the spread of this idea?

**The [AI Commons](#) is a new initiative and organization dedicated to problem solving with AI more democratic and accessible; and a framework for the coordination of scaling of AI for Good.** It was announced in May 2018 at the AI for Good Summit, by stakeholders from ITU, XPRIZE, Ocean and more announced (see image below). It's been quietly gathering steam since then. I plan to write more about this in the coming months.

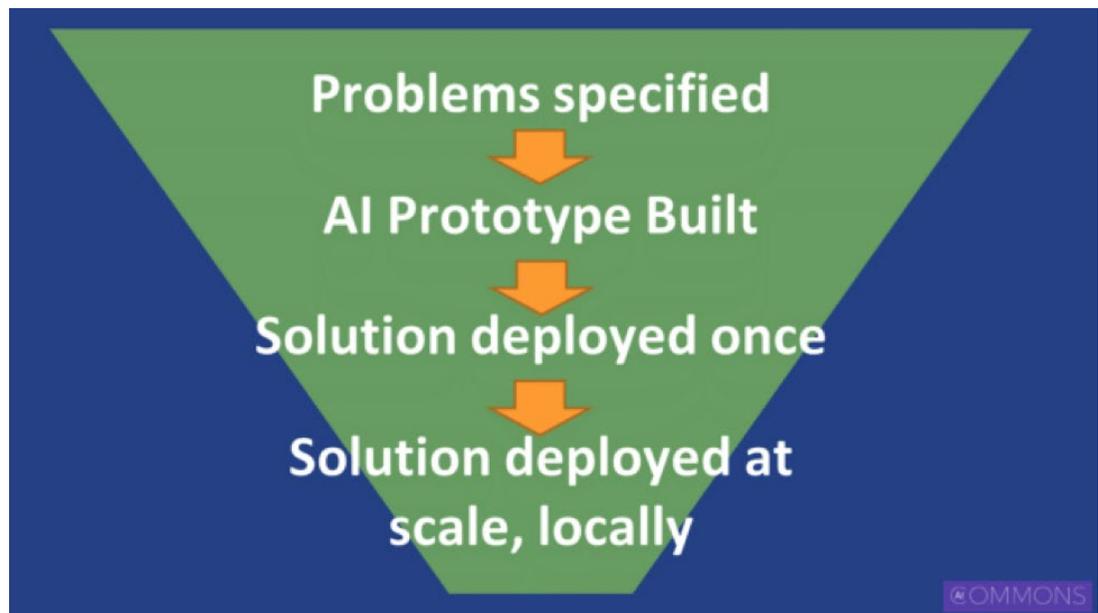


To help catalyze community action, the [2019 AI for Good Global Summit](#) (May 28–31 in Geneva) has a track dedicated to Scaling AI for Good. It’s my pleasure to co-organize this track with Alexandre Cadain, Masha McConaghy, and the ITU. At the track, we’ll be learning and sharing about efforts on scaling/capacity building and on AI for Good that are on the threshold of scaling, and about getting involved even further.

## The #AIForGood Scaling Funnel

The next question is: what are the precise ways in which diverse community members with diverse skillsets can get involved? We’ve found that if we frame the process as a funnel that tracks each AI for Good project from inception to deployment at scale, then the funnel naturally gives specific engagement points for each stakeholder. For example, problem owners have the most to say at the inception of the project, but need to be involved in every step to ensure that their problem is actually being solved.

Just as there are funnels for sales processes, hiring processes, and more, we can have an #AIForGood Scaling Funnel that walks solving a given impact problem from initial specification to final deployment at a scale. Here’s what it looks like:



The AI for Good Scaling Funnel (Image rights - see note 2)

Let's walk through each step.

- 1. Define Problem** - Problem owners work with AI researchers/data scientists to specify the problem. Specify objectives & constraints, gather data, and cast it as an AI problem (regression, classification, optimization, etc.)
- 2. Prototype AI Solution** - AI researchers/data scientists (working with problem owners) apply & invent AI technologies to come up with a prototype solution that meets the problem aims. For example, an AI model that meets the target error.
- 3. Deploy once & find product-market fit** - "Make it happen" business types (e.g. entrepreneurs) work with AI researchers/data scientists and problem owners to deploy the solution into the problem domain and iterate until the product-market fit is found. This is the other 95 percent of the solution, to make it truly be used, with success, in the field. In some cases, capacity building efforts will be critical for success at this step.
- 4. Deploy N times, for impact** - Once we have legitimate product-market fit in one deployment, we scale it up to into more deployments in other

jurisdictions, with problem owners for each jurisdiction. Often scale-up can include adding adjacent products around the initial product. It's not just about scaling the idea, it's about scaling for maximum impact in each jurisdiction (KPIs can be different in each jurisdiction).

Only when an AI project has completed the final step can it be considered "scaled up". By "working the funnel" we can nurture a wide variety of AI for good projects and help take them to scale.

The funnel gives stakeholders more precise ways to engage. For example, UN agencies are best suited to helping define problems, and especially, taking solutions to scale (steps 1 and 4). Entrepreneurs will focus on steps 3 and 4, but will draw ideas from steps 1 and 2. Data scientists and AI researchers will be needed at all four steps; but now they will have much more context versus traditional focus on step 2.

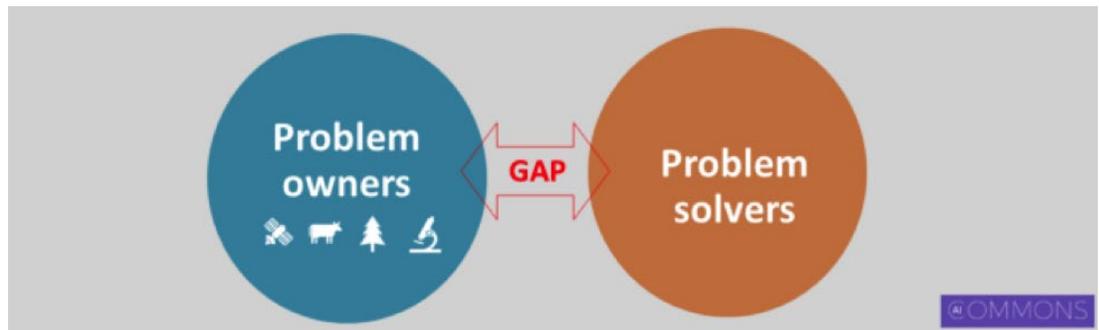
What if we could take 50 projects through the funnel, and deploy each to 100 jurisdictions? We'd have 5000 scaled AI deployments. That's scaling AI for Good.

## Connecting Problem Owners with Problem Solvers

### The Gap

If you're a first year AI PhD student, you're usually given a data/problem set to work with, such as ImageNET for computer vision, or UCI Repository, or OpenML. More often than not, you're not connected to the people with the actual problem. This was [even my experience](#), 20 years ago. And very rarely are you connected to people with impact problems like SDGs. This disconnect often doesn't disappear as students progress through their PhD and career.

In the context of the funnel, this issue is framed as linking steps 1 and 2. There's a **gap** between the people with the challenges (**problem owners**) and the people with the expertise to solve the problems (**problem solvers**).



The gap between problem owners and problem solvers. (Image rights - see note 2)

The problem owners have the AI-type problems to solve at step 1, such as tracking deforestation or cattle movement. The problem solvers are the data scientists and AI researchers, who can build the AI prototypes at step 2.

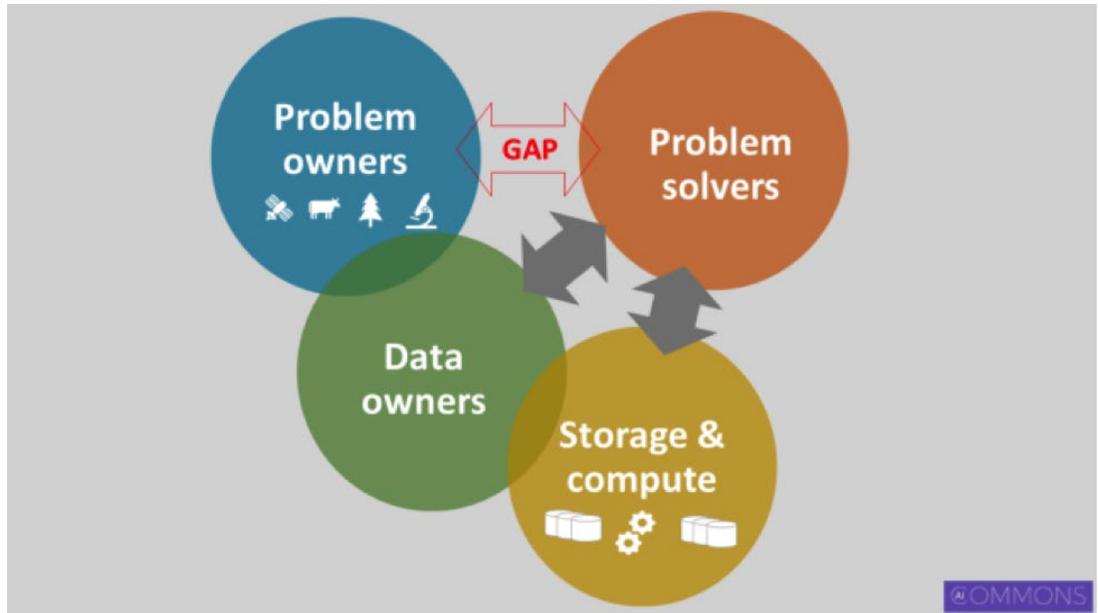
There are two complementary solutions to help solve this:

- **Human-based approaches**
- **Connective substrate technology**

I'm excited about both! The former is about AI researchers getting out of the office and into the field, for meaningful face-to-face interactions; along with systematic ways to capture the problem definition towards solving it. The "field" could be local to the AI researcher, or it could be farther away. I see that AI for Good opportunities can exist anywhere, so why not start local? The latter is about using technology as a global connective substrate. The next section elaborates.

## Connective Substrate Technology

First, we flesh out the ecosystem with other key actors. The image below illustrates.



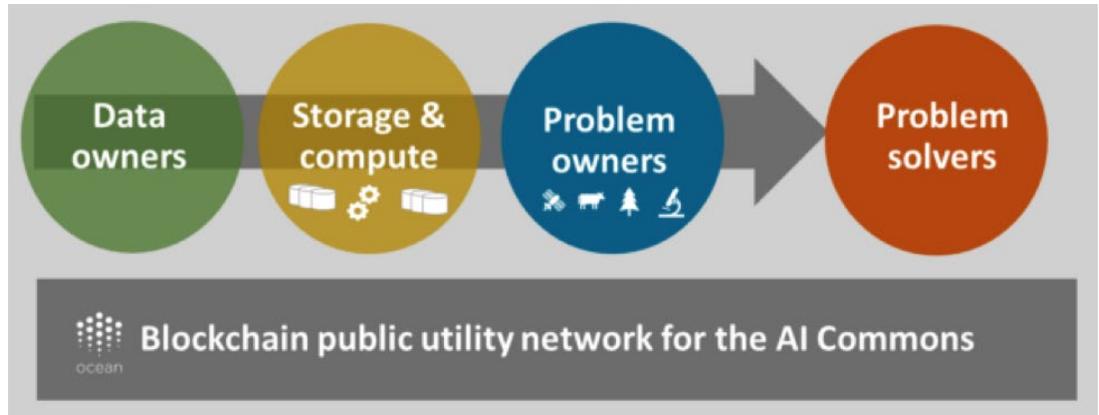
Ecosystem around the gap between problem owners and problem solvers. (Image rights - see note 2)

Most modern AI applications are about building classification and regression models. These models [love data](#). The data is held by the **data owners**. The data owners are usually talking to the problem owners, at least a bit.

Then there's also the **storage** of the data itself, and the **compute** by the problem solver to train the models. And the high-level goal remains the same: we're helping connect **problem solvers** with **problem owners** to solve the AI problems.

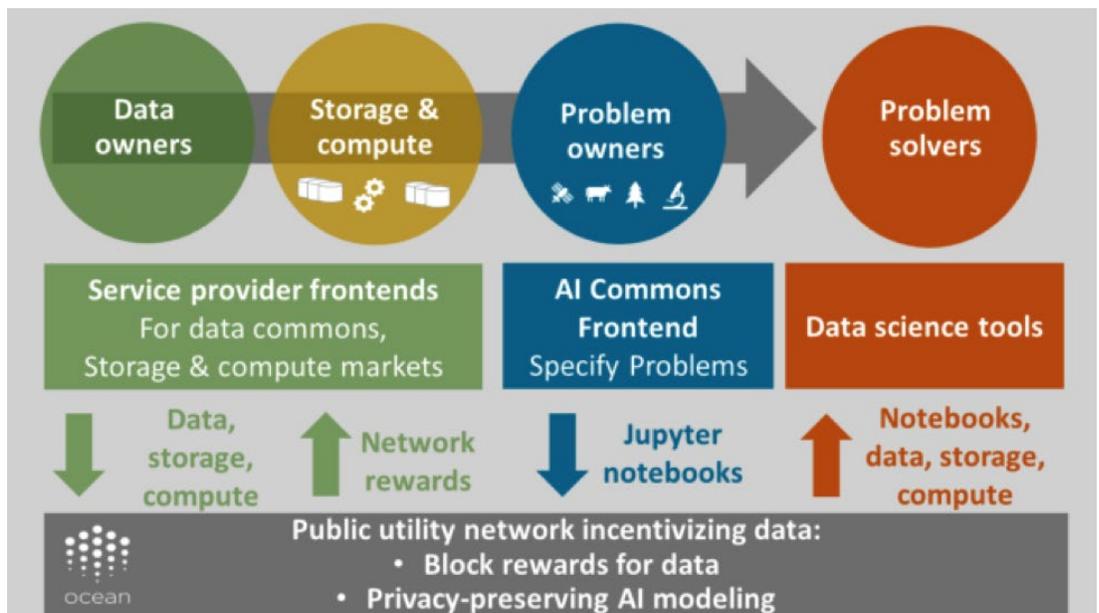
The idea is to create a connective substrate to make it easier for data owners, storage and compute, problem owners, and problem solvers to connect and to solve problems. It should be a public utility, not owned or controlled by anyone.

We've had public utility networks for a long time, like the public water system, the electric grid, and the internet. Some public utility networks sit on top of the internet base layer, including the World Wide Web and the Bitcoin blockchain network. Blockchain technology is the appropriate choice to implement our public utility network to connect problem owners, problem solvers, and the rest.



Connecting the ecosystem with a public utility network (Image rights - see note 2).

Let's further flesh out what this blockchain-based public utility network actually looks like.



Detail view of connecting the ecosystem with a public utility network (Image rights - see note 2).

**Data, storage, compute (top left).** Data owners interact with different service provider frontends which in turn connect to the public utility network. There are also suppliers of storage and compute services, interacting through markets to buy and sell the services. In return for data and services, they get paid by people on the demand side. That's for paid data. We also want to incentivize people to supply data to the commons, which in turn will spread the benefits of data and

therefore the benefits of AI. To implement this, the [network itself pays people](#) to supply data to the commons, using inflation.

Finally, we can unlock private data by bringing compute to the data. In this fashion, the private data never leaves the premises, yet AI models can be built from it.

**AI Commons frontend (top middle).** In the AI Commons frontend, problem owners specify the problems and work closely with problem solvers. There's a back and forth between solving the problem, and defining the problem with more clarity. Problem owners can define the problem in interactive notebooks such as Jupyter.

**Problem solvers (top right).** Problem solvers take as input a problem from the network (e.g. Jupyter) and start to solve it. They draw on data, compute and storage services. Their progress is recorded in the network, giving provenance to the compute and data. Their final result is also stored in the network, for example as a Jupyter notebook.

Problem specification and problem solving is an incremental process: do an initial specification of problem, do a first-cut solution (perhaps with [automated ML](#)), then take the learnings from the initial solution to refine or change the problem, do a second-round solution, and so on. This also implies that the boundary for tools for problem specification and problem solving will blur. For example, AI Commons frontends will evolve towards including solutions.

## Applications of scaling AI for Good

The previous three sections described a path for Scaling AI for Good: community, scaling funnel, and closing the gap between problem owners and problem solvers.

Here are some applications using some of the ideas above to help scale AI for good.

**Water Quality.** This is something that [IXO Foundation](#) is working on. In this case, there's forest restoration to help water supply. And you want to understand how much can forest restoration can help the water supply. Satellite data is super useful for that and ideally, the satellite data is out there in the commons.

Imagine that all the satellite data, as well as the forest data, and so on is actually all available for the whole planet to see and via basically a unified API. And also, there are block rewards for people who are supplying that data to make it available.



**Health.** Imagine building an AI model across 10,000 hospitals, rather than the status quo of a handful at best. You'll be able to predict diseases far better. [ConnectedLife](#) is aiming for this for Parkinson's disease. The data stays in each hospital. This means that make a single model across borders, across hospitals, across different jurisdictions, and promote research collaboration more broadly.

**Agriculture.** A World Economic Forum (WEF) spinoff called [Grow Asia](#) is about helping farmers to allocate fertilizer. I grew up on a farm in rural Canada and every year we had a fixed budget for fertilizer. Typically, because of the way water flows, you don't need as much fertilizer at the bottom of a hill as you do near the top. So, how do you actually figure out the optimal allocation? Grow Asia is addressing this need for farmers in Asia with the help of data and AI.

**Policymaking.** Policymakers are bumping up against a new tension. On one hand, you want your citizens to benefit from AI, which means making data

available. But you want to ensure good privacy for your citizens, too. These are conflicting goals. But there's a solution: build the models without letting the data escape, and instead bring the compute to the data. [Ocean Protocol](#) is working with the Government of Singapore's data authorities (IMDA) to flesh this out, to get the benefits of sharing data for AI modeling, at the same time as addressing privacy.

## Conclusion

If there were KPIs for civilization, perhaps it would be the UN Sustainable Development Goals. AI can help reach these SDGs. This is the aim of the growing #AIforGood movement.

This article explored how AI for Good can scale. It starts with community, including AI Commons. Then, we use a funnel as a systematic process to track ideas from inception to deployment at scale. Finally, we close the gap between problem owners and problem solvers with a combination of human-based approaches and connective substrate technology (a blockchain public utility network).

Want to get involved? Come to the "Scaling AI for Good" track at the [AI for Good Global Summit](#), May 28–31 in Geneva. [Here's the registration link](#).

## Acknowledgements

A huge thank you to Amir Banifatemi of AI XPRIZE and AI Commons, who is a key initiator of the #AIforGood movement and continues to help grow it. He's a contributor to many ideas in this essay. Moreover, he's how I (and my colleagues at Ocean) have become more involved in #AIforGood: educating us about the challenge of scaling up AI for Good, and collaborating towards solutions. Thanks also to Sean McGregor of AI XPRIZE who has helped in many ways; in particular for co-conceiving of the "scaling funnel" with Amir and myself.

Thanks to Masha McConaghy and Alexandre Cadain who are my co-organizers in the “Scaling AI for Good” track of the 2019 AI for Good Global Summit. The folks at ITU who have also been instrumental in starting and growing #AlforGood, hosting the Summit, and organizing the scaling track. These include Fred Werner, Bastiaan Quast, Ayda Dabiri, and Reinhard Scholl. Thanks to the rest of the AI XPRIZE team, as well as the broader XPRIZE family, who have helped to catalyze this movement. Thanks to Sheridan Johns and John Enevoldsen of Ocean Protocol, for helping get Ocean more involved in #AlforGood. Thanks to the rest of the Ocean team for helping create a commons platform for #AlforGood.

### **Notes**

[1] To be clear, “good” should not be constrained to “helps at least one SDG”. It is, however, a pragmatic starting point for conversation.

[2] Images referencing this note are © Ocean Protocol Foundation, 2019, with non-exclusive perpetual license to AI Commons Foundation to distribute, remix, tweak, and build upon these works without attribution.



## **Trent McConaghy**

Founder, Ocean Protocol

Trent McConaghy has 20 years of deep technology experience with a focus on machine learning, data visualization, and user experience. In 2017, Trent founded Ocean Protocol, a decentralized substrate for AI data and services. He is an ambassador to AI XPRIZE, and advises Estonia, Germany, and other countries about blockchain & AI strategies.

Trent was a researcher at the Canadian Department of Defense, and in 1999, he co-founded Analog Design Automation Inc. as CTO until its acquisition by Synopsys Inc. In 2004, he co-founded Solido Design Automation Inc., once again in the role of CTO.

Trent has written two critically acclaimed books on machine learning, creativity, and circuit design, and has authored and co-authored more than 40 papers and patents.

Trent has a PhD in Engineering from KU Leuven, Belgium, and Bachelor's degrees in Engineering and in Computer Science from the University of Saskatchewan, where he won awards for the top PhD thesis and top undergraduate thesis.

[@trentmc0](#) [@OceanProtocol](#)