

Transcript of Open @ UNGA81

<https://creativecommons.org/event/open-unga-81-a-conversation-on-collaboration/>

Amanda

I think the question around the concept of even decentralizing the importance of models in a modern tech stack is something I've been asking really since 2022, when things started gaining more attention—in my personal life and in my work life. What's new? What's different next?

I think the conversation was different in 2023 around models because the assumption was that models would replace the entire technology stack. They would replace data extraction and consumption. They would move into feature extraction. The model was basically doing everything. The model itself was the product. It was not integrated into a product, it was not a feature, it was not an enhancer, it was not a flavor modifier. It was the actual product in and of itself.

Three years later, I think you can find that that's actually not the case. I mean, it is a market in itself, but it is not the market in and of itself.

So I'm really encouraged by the way—and thank you so much, Duane, for my creative comments—you've structured the conversation today around governance and infrastructure and data and hardware and so many other pieces of what it takes to build the modern infrastructure stack.

I want to double down on the question of where we should be investing when we think about equalizing modern tech stacks through software.

I really appreciate Mitchell kicking off with the idea of interoperability and protocols and standards in global governance models, because I think in a modern tech stack, this is an area—and I think Chris brought this up as well when we were talking about protocols—where innovation and technology can change in a way that we currently cannot predict or commodify right now.

So, if we're looking at one specific area of the stack and trying to see how we can bridge that gap or break that mold, what we're not identifying is where constraints are happening and where we can allow people to speak a common language, both with technology as well as with human trust, that allows them to move in ways and come up with new ideas that we're not currently capturing through VC money.

Right now, I think we are in an immature state when it comes to the modern AI ecosystem. We can see that in the discussion that Yan brought earlier, talking around, for example, Robust.txt and how that's a protocol.

A protocol is not the same as a standard. There are no remedies associated with breaking a protocol versus breaking a standard. You feel them through the remedies of how people are acting, but it's not the same as having to go through remedies for regulatory things.

We see that developing now, too. There are some protocols and standards on the books around AI that have been donated to vendor-neutral foundations as a way of developing them, to allow for new markets to exist.

And I looked at them and I said, "These are Markdown files. I don't understand how this is a standard. This is a data structure. This is a metadata structure. This is not a standard."

But the reality was that there needed to be a vendor-neutral place where different groups could come together and talk about how their products could interact and work with each other. That was not going to be waiting on standards bodies, quite frankly, because standards bodies just don't move at the same speed.

So I think if I was looking for where we should be concentrating our attention, I actually do think there needs to be increased investment from the different groups into standards bodies, into protocols, into vendor-neutral organizations like OSI that are actually working to try to make sure that access is not to the few, but to the many, and that whatever level you're working on, you're able to access that.

Ilan (moderator)

Very interesting.

Have you drilled down to some of these specific protocols or standards, like MCP, and thought about what's enabled and what the limitations mean? So if we can get into some of these specifics next.

Madiha?

Madiha

I'm an academic. I'm currently a postdoc at the Data Science Institute at Columbia. I don't have an answer for MCP, but I do have something to maybe add to Amanda's point.

Fundamentally, I don't think that power moves around randomly. I think it moves on purpose.

This idea of commodification as a complement, right? Google didn't open up Android and Chromium for fun. It was in service of search and apps. The Llama family of models were

released as open source as a competitive advantage to OpenAI. And then we see this happening over and over again.

So this idea that opening a layer in the stack will fix everything, I don't necessarily think so. I think this is a move that's more of a complement, in service of some kind of competition.

So this question around where should we concentrate or focus our efforts—you know, I'm an academic, I study communities and user groups—and so I think the question is: Where do layers connect? Where do these layers connect? Where does participation actually happen?

For AI, that's model hubs. That's ModelScope. That's Hugging Face. And I think one of the most interesting, but also somewhat terrifying, research studies that I've taken up over the last few years has been on Hugging Face and what community participation in the open-source model ecosystem actually looks like.

And what it actually looks like is this, right? It doesn't look like what we previously observed in open-source software ecosystems. There's a significant Pareto distribution in open-source model ecosystems in terms of collaboration and community efforts.

However, ModelScope looks a little bit different, and the innovation coming out of China's communities looks very different.

So maybe I'll stop a little bit there, but I think when you scrutinize and think more critically about model hubs, there's another thing I'd add.

On Hugging Face, there is a page that lists licenses that a community can add to their repository, right? It's just a static webpage of licenses. You cannot click into them. The licenses are there, but you can't click them. There's no definition. There's no documentation.

So what does that actually mean?

GitHub looks a little bit different. But in my case, this is something I observed both two years ago and as of last night. You can see this.

Ilan (Moderator)

Interesting. And so, what is a model hub? What is a model hub?

Madiha

A model hub is just a term for where open-source models are being built and shared—like GitHub, but for open-source AI and model development.

Ilan (Moderator)

Okay, that's good to know.

We're all on the same page.

Any other panelist want to add anything on that?

When you get asked about MCP—

Ben

I'll just say a brief word about MCP.

MCP is the Model Context Protocol. It's an effective way for applications to connect to other sources of data and applications. It's fantastic if you're a developer.

But just speaking to the limitations of an open standard or open protocol to create competition: It's called Model Context Protocol because built into it is an assumption that we need to make it easier to bring data to models.

In fact, there's nothing like a marketplace, like you have even with the App Store. There's plenty of criticism of closed gardens and so on. MCP has not created the marketplace on ChatGPT or Claude.

If I want to bring a new application to market on top of those things as a platform, I can't.

And so you get more into the question of market structure and what's possible. MCP is a great example of a very neat protocol, but it's not doing anything to really create new competition.

Ilan (Moderator)

Okay, MCP.

So it's like, you know, with short code, because of MCP I can now get emails from my Gmail, potentially with read-write access. So it's like these universal, adaptive APIs.

But doesn't this mean you don't need anyone to create an application then? Potentially you can integrate your AI or digital application in the enterprise market.

Ben

Yeah, but in the enterprise market, for sure.

In the consumer market, MCP doesn't do anything to challenge Claude or ChatGPT because those platforms would love to have access to your email and your documents and so on, but they're not doing so in a way that creates the possibility for new consumer entrants to build value on top of those platforms yet.

That might change.

Ilan (Moderator)

But you still rely on the same platforms to call those applications, right?

So, thinking about this, doesn't it make it easier for you to make your own clients?

Instead of me—if you think about an effective AI agent today—imagine it wasn't a standard identity. I might need a bespoke deal to make my agent relevant. I would need a deal with Google, maybe, or I would need to hard-code an integration to Google's API and then to Slack's API.

Isn't it easier in some ways to make these harnesses and spin them up because you're not starting from scratch when it comes to integrating your harness with these tools? So that lowers the barrier to entering that layer, at least for us.

Amanda

I think one thing that we should make sure we're adding in here is: When you're talking about ease, who is that ease for?

In the past few years, one of the concepts that has fundamentally changed when we're talking about an open internet is who is having an interaction, and what are we assuming that that entity is?

Where this protocol sits in the stack conflates the idea of a probabilistic versus deterministic process.

Part of the problem we're having now when we talk about trust—little-t trust—and humans and machines and the ecosystem we're creating is that we can't conflate that with deterministic trust, which modern security is built on.

By having machines act as users in order to gain access to these systems and to be able to gain access to APIs, we don't currently have processes where we can understand, when something comes forward to a machine, is this a bot? Is this an agent? Is this a machine? Is this a human?

We're catching up there.

So the ease, and for whom, and whether or not you're an enterprise or whether or not you're a consumer—right now a lot of those distinctions have broken down in a way that is combining and conflating those assumptions, and we haven't fixed that yet.

Ilan (Moderator)

What you're trying to do about MCP—but I feel like we have not talked about this enough—is the concept of trust on the internet. It looks fundamentally different now than it did even a year and a half ago.

Are you saying that, in effect, this is stopping markets from arising?

Markets could be private. You could run an app store that at least potentially allows broader participation and value, right? Same as how countries monetize on YouTube.

You're also saying that these open standards, in comparison to these private markets, are also being impeded from arising to create these distributed layers of the market because there is a lack of trust, security, identification, and verification.

Duane

Well, this is such an interesting conversation because, as I just said, at the dawn of LLMs we thought, "Oh my God, these are going to be unit systems. Everything's going to run through an AI model."

A couple of years later, that still may be true in the long run, but we can see that it's a much more complicated situation.

There's a question of: Is the power in the market from the model, or is it from the point of distribution where a customer or user interacts with the model?

If I'm in the consumer market, a lot of the value comes from the harness or the client that connects to the model.

Madiha made a great point—or at least raised a wish-list item—in the last panel, saying, "Wow, the open-source world should have a go-to-market strategy for an open alternative to buying a ChatGPT subscription," and so on.

The models themselves are becoming a contest. In fact, you have many choices for a model that is 80%, 90%, maybe even 99% good enough for the things that most people want.

Whether it's an OpenAI model, an Anthropic model, a Google model, or the long tail of open models from companies you've never heard of, there are open-source models that run on your computer.

The model is no longer really the unique thing for most people. It is driving toward a commodity.

And yet, in any market, the ability to distribute the application is very important.

If we're thinking about the consumer market, even though technically there are open alternatives, they're not yet distributed. The distribution advantage, the advertising, the data, and the agreements that they have access to are different, right?

That begins to lead into all the other discussions that we've had.

So I applaud you for setting such a broad market-structure conversation, because we're at a moment where, one way or the other, all of these markets—consumer and enterprise—are being questioned.

Ilan (Moderator)

And you're talking about distribution still being especially important for consumers.

My mom, or other people—you go on Google, you type into Google Search, and there's an AI. It's embedded. Same with Meta.

Again, it uses distribution. That seems like a really important dimension.

Duane

Yeah. We fixed it, right?

Actually, the thing I want to talk about is a theme I hear coming up in some of the questions and responses: "Is open going to do this? Is open preventing that?" and so on.

Open by itself does very little. It enables a lot, right?

So if we're looking at open source just as a methodology to create markets, to push protocols that everyone can contribute to, that's not enough.

I think that's sort of the point I've been driving to. There's no particular go-to-market strategy, and that's not really what open is about.

You can have transparency. You can have reproducibility and so on. But especially those of us who are close to the space and love open and embrace it as our ethos, we lose sight sometimes of the fact that it's not enough by itself.

Ilan (Moderator)

So what do you think? Other than "it enables," what does that mean?

In Europe, let's say, does that mean that in Europe and America we should see a flourishing of innovation across the stack because we have the models that will be enabling distributed innovation down the stack, while things are still missing?

Duane

I think you can see it, but yes, I think there are things that are missing.

I think Linux is a really good example to cite for a lot of reasons.

It didn't happen organically. It happened as a result of significant, long-term, sustained investment.

I'm sure 30 years ago it was inconceivable that anybody could get to the level where they could compete with the incumbent market. Just as it was inconceivable at one point that anybody could disrupt Unix. It was inconceivable at one point that anybody could unseat Internet Explorer.

We forget, especially—and it's so hard right now because the space is changing so quickly and the big players are so big—that it can feel sometimes like it's just not possible.

There's an adage that the best time to plant a tree is 30 years ago, 40 years ago, 10 years ago.

If we want these kinds of markets to flourish in Europe or anywhere, we have to invest in creating those markets.

Ilan (Moderator)

I'll take five minutes just to see if we can drill deeper into the relationship between open protocols as open standards for communication on the network and open-source software.

With an open standard, as long as you are speaking that language, you can be using private software or public open-source software. It's agnostic.

I can connect to the internet with almost any kind of hardware and most kinds of software, as long as I'm speaking those specific internet protocols.

So it allows for continued innovation potentially on the backend, as long as the standard isn't set in too rigid a way.

I thought we'd see which of the panelists wants to tackle this. Go from left to right and, for the AI stack, where do you fit the most important interventions in the balance between open source and open protocols?

Do you need both?

And maybe, if there's a concrete example you want to discuss—whether it's a protocol in AI or a model in AI—that would be great.

You don't have to answer this one, though. We're going to interact.

Duane

I think it's constructive to talk about a stack.

At the same time, I think we should also step back a little bit, because there are discussions of what protocols are needed for the agent internet, for instance—the whole set of them.

There are discussions not just around trust or identification, which can be protocolized. You've got not just MCP, you've got agent-to-agent protocols, you have agent payment protocols, and you have flourishing protocols. They're very special-purpose, and the enterprise space is working them out.

I feel like that's a very healthy space.

At the same time, if you think about the market for AI as, "How do consumers, citizens, and people access AI?" that's a slightly different framework.

So this is sort of: How do you build a stack where you can enable lots of the knowledge needed to deliver a really advanced system? How is enough of that knowledge public and actionable?

The course of innovation will take its course, whatever that may be.

The stack is a very helpful organic framework because you want enough knowledge in the commons to exist about how you train a model. You want exemplary models. You do want open-weight models that you can pick up and play with. You want harnesses as applications that you can build on. You do want protocols.

You need a whole stack to enable openness.

I think you unite a group like this around that.

But it's also interesting because the internet is a very constrained way to think about what's coming.

We said that if you're born after January 1, 1970, you've lived your whole life in the internet era. For 50-plus years, we've all had these mental models for the governance of technology, government, and all that.

But the people building AI in the early days of the LLM revolution were saying, "This is bigger than the internet. This is like electricity or fire or the written word."

So I think when we talk about openness, there's a degree of value alignment here that's bigger than how do we keep it open on the internet.

Yes, I'd like AI to be more like email and the internet, not an oil company. The latter creates a lot of inequality.

I think we need to think more broadly about how openness across the stack enables choice, competition, and even freedom, if you'll permit me.

Protocols are fantastic, but I always think the stack is the right way to think about this. Knowledge tends to be what gives you the ability to create something that's amazing.

Madiha

I know you prefaced that question by hoping that we would get specific. I'm going to disappoint you, I'm sorry.

I'm going to answer the question through a slightly different mindset.

We're talking about where we can pick the right place in part of the stack, or in protocols, in order for this to work.

I don't think it's about where. I think it's about who.

Who can you mobilize for collective action and collaboration?

The most successful approach is not about finding a leverage point that's the most likely thing to be open. It's about creating that leverage point through collaboration.

We all agree that if the harness is the place where there's the most leverage, and nobody comes to work on it, it doesn't matter.

We have to find a place that everyone will collaborate on.

And that seems to be a very open-source-forward answer, as far as I can tell.

Ilan (Moderator)

In some ways, open source is a different kind of institutional framework in terms of how people work across and how they collaborate.

Madiha

Yes. Open source isn't free, as we mentioned earlier. Even if it can be free to distribute, it can require lots of money.

Where that money comes from—you know, third-party firms or private firms wanting to disrupt their incumbents—like that was the case when we wanted to develop this MCP standard to disrupt OpenAI's incumbents and their standards.

So I mean, open source as an organizing principle, but also as a kind of structural form: Is that then the right way to think about it in terms of how people can mobilize?

Or how, you know, little countries can band together?

But again, open source is global and cross-border, so it's still hard for me to square this cross-border internationalism with states and these kinds of goals.

Ilan (Moderator)

Well, I'll pick up on your question: Is open source still a helpful concept?

Duane

Well, of course it is.

But it's also under some pressure, I think, because there's a definition of open-source AI which is very robust, and a lot of what we call open-source AI isn't actually open source under that definition.

I don't want to get into that because I think that open-weight AI—that is, releasing the model so you can use it in a lot of cases—gets you a ton of the value of open source even if it's not open source all the way down to the steps.

I think that, at least when I come to these discussions, I think about how we have an open market and how we have open innovation. I think we do need to think a bit more broadly about "open" and not just "open source."

I'll give you another example. Distillation is a hot topic.

Distillation is when you train another model from the outputs of a teacher model.

Now, if you're Anthropic, you complain that all of these Chinese groups are building these models by stealing your model outputs by doing this on platforms and things like that.

And I think that's awfully rich if you trained your model on everything on the internet, right?

Then you go off and say, "Oh, but you can't train your model on my model."

As an open person, my construction of openness cuts both ways.

I think this is not an open-source discussion, but in spirit, to me it is.

I feel that we need the knowledge, the technology, and the ability for people to study and inspect and create things with knowledge, data, and systems to be preserved in the AI era, so we don't recreate these black boxes that are highly governed.

AI safety is a big topic, but if AI safety means that we need to turn away from open source, I think that's actually not a good thing. That's something that we want to avoid.

I always go back to something that predates open source.

There are people here who will laugh at me for bringing this up, but before open-source software was free software.

Open source is sort of version two of free software.

We may be interested in where we need version three that encompasses things like what's the right framework around distillation, what reasonable people can do with models, and the right to be adversarial and reverse-engineer things.

Madiha

On protocols versus—

A couple of points I just wanted to agree with and complement what you were highlighting.

The first is the Open Source AI Definition.

One of the challenges that academics have been having with the definition, and with researching the space—interviewing people, looking into how startups are using open models, or what the licensing structures look like—is figuring out what's complicated here and what is actually okay for us to use.

The definition of Open Source AI, I think, is on its way, but there is one part that is a little bit difficult to follow.

We are trying to apply the ideology of open-source software to a technology that is fundamentally materially different from what open-source software looked like.

From its development to its implementation, these are different artifacts that we're dealing with.

So when we have this Open Source AI Definition that is, I think, focused on the availability of training data, code, and reproducibility, but also on downstream use, we're forcing together two parts of the definition into one.

How do we reproduce the model, and how do we allow it to be reused?

Those are two different ideologies that we need to find a way to make work together.

The point about distillation, I think, is a great one.

Just this morning there was a really great blog post by Nathan Lambert that talks about distillation as well.

What he says is that distillation is not a great explanation for all of the success that the Chinese model families are being associated with. At best, it may explain a performance gap, but that's about it.

And again, on the comprehensive research and academia side, these model families are a foundational layer of research in machine learning communities.

The communities I'm part of are benchmarking on all of these models. You will not get your work accepted to a conference unless you show your work against these models and these benchmarks, because they are that relevant.

Ilan (Moderator)

Yeah, I think those are some important points that help also differentiate parts of the stack once you get inside of technology.

Amanda, I'm curious to hear from you. Maybe your answer to how you hear what you're working on at Google in terms of open source and audience is also interesting.

Amanda

Sure.

I think first, to address the question of fundamental change and that question around open-weight models and whether or not that is what's necessary to distribute value or market power—whether it's protocols or other aspects—I think one of the things we haven't talked about enough, and I think Madiha actually first started addressing this, was the concept of labor.

Right now there is a super-high concentration and an assumption about the necessary labor and technical skills and the kinds of education that are necessary to build technology stacks like this.

I would challenge whether or not that's true.

As somebody who, again—you may be surprised to know that I am not always a hit at parties discussing technology—especially when it comes down to the question of when people say things to me like pre-training and post-training and fine-tuning, I look at them very clearly and I say, "We're talking about data engineering."

This is the aspect of information science we're discussing.

And the fact is that this is, quite frankly, the most hated class of computer science education.

What they're really talking about is taking data and putting it into a way that statistical processes can be able to let your eyes pass through.

And they don't like that because it turns the value of their labor into something that is worth much less on the current market.

Quite frankly, you earn much less money as a data engineer doing traditional ML processes in companies than you do as an AI engineer or somebody who focuses on pre-training and post-training.

So companies who have spent a lot of money on their labor, their intellectual property, and their knowledge in that area have a reason to make sure that that investment is protected.

Part of that is by making sure that that's highly valued as part of your IT stack, as opposed to discussing the question of, "Well, what if all this stuff is going on and a group of open-source developers or people or engineers who are familiar with technology stacks can take a pile of data, use some open-source or open-access data, use some open-source software, put it through commodity hardware—and it turns out you can shrink it down, you can link some laptops together, things that are already existing—and make something that is high quality to solve good-enough problem spaces on good-enough hardware?"

Does it actually solve the problem you need, and what's that worth?

I think then the cost of problem-solving comes much farther down, which is probably frightening for some people who have invested a lot into that space.

But I think the question is: Are we at the point where, as a society, we're asking how much should be invested here in getting back from that, and what's the right size for a use case on good-enough hardware?

I think we're approaching that more rapidly than I think is showing up in the headlines.

Ilan (Moderator)

It seems like you're touching on the architecture of the system, right?

If you can run this thing on a laptop, that's kind of a different architecture from a centralized model.

We were going to discuss two other topics, which I thought were interesting points, like the cloud. Basically, over time we expect insurance costs to grow even if we're still using laptops and mobile, and that would be a source of profit-oriented attraction in society that grows over time.

And, well, if you may wave your magic wand of industrial policy, what would you consider?

But instead, let's go to the audience for questions and comments, and then you guys can post something like that where the system goes.

Duane

Yeah, while the audience is queuing up for questions, the Open Source AI Definition came up twice in conversation, and I didn't want to jump in one by one responding.

I also forgot to introduce myself when I started talking.

I'm Duane O'Brien. I'm the Executive Director of the Open Source Initiative.

We maintain the Open Source Definition, which describes whether a software license is considered open source or not.

Four years ago, we embarked on a process to craft an Open Source AI Definition, which we published two years ago as version 1.0.

It's not without criticism, and a lot has changed since then.

It took something like 10 years for the Open Source Definition to stabilize into the set of principles it has today. This is a continuous and ongoing process.

So if you want to be involved as a participant in this realm of discussions, yes, please, let's talk. We're going to begin those conversations very soon.

Ilan (Moderator)

Thanks for the invitation.

Okay, so any questions? Or if you want me to comment on some of the things we've been discussing, I'll give you—

Audience Member

I'm from the nonprofit sector.

A lot of discussion has been about open-weight models.

A lot of people I talk to in the sector are kind of terrified of Chinese open-source models and other models, as they have no idea what's in there—whether it's security risks or whether people's data could be infiltrated.

How would you know?

Why should we be worried about this? Or maybe we should be, and I just don't hear about it that much.

I get the impression that almost all of the industry is around 40% open.

So what's going on here? Why shouldn't that be working?

Ilan (Moderator)

Good question.

Let's read it quickly and see if we can get one more question together.

Audience Member

From Lab India.

You would have heard this: A lot of countries are now talking about model sovereignty and AI sovereignty.

How do you see open source and AI sovereignty coming together? Because open source fundamentally means a level playing field and more collaboration.

How do we balance that out?

Thank you.

Ilan (Moderator)

Okay, let's take those two questions.

Duane

On the question of people openly questioning models, I would instead ask: Why are you not worried about the things you can inspect?

Not all of this entire process involves offloading or building a certain amount of risk.

If you say, "I don't have to worry about this because Anthropic will handle it," you've offloaded all the risk to them.

If you do that without considering your own risk posture—where you deploy things, how you deploy things, how you integrate things, how you make sure that your customers are protected—you're also making an error.

Open-weight models give you some ability to tune and tweak and so forth.

The more open technologies are, I would say, the more it matters what we call open source.

The more open these things are, the more you can, if you're unsure about something—access to your email or credit cards, for example—start the process of constraining those things.

Madiha

I think I would just add that this morning there was a panel where one of the speakers said we have to be really careful about the narratives we hear and think about who's saying them and what incentives they have for us to kind of buy into them.

I think a lot of the rhetoric around the models coming out of China is unfounded, at least in research communities and academic circles.

We are all using the open models from China. We are researching with them. We are testing with them.

I think a lot of academics and researchers would argue that these model families are leaders in open weights as it currently stands.

They may be two months behind the frontier models coming out of the US.

So there are these narratives that American frontier labs would like to push toward us to kind of win the so-called AI race. I think there is a lot of interesting rhetoric out there.

And particularly if you look into the small-model space, small-model communities, there's a great place to go to read about what people are doing with hardware, extra compute, GPUs, gaming cafes.

People are sitting in communities and fine-tuning models and discussing constraints around both Chinese and American models.

So I would say the small-model space is also one that might give you some answers and help bust some of these myths.

Duane

And maybe I'll take that opportunity to bring up and speak to the sovereignty question.

If you're a nonprofit CTO, you can use trans—if you work with security or other companies that will handle the difference for you. Check out their security. They all serve you GLM or Qwen or any of those things. No problem.

Now, if you're worried about sovereignty, right? If you're a government, maybe you have some different strategic objectives.

An open model is not going to be a sufficient investment strategy, right? You may want a sovereign model.

Let that be part of your export policies. That might be part of your regional influence policies.

Indonesia has this incredible SEA-LION project that's trying to equip all Southeast Asian dialects with a foundation model.

There are groups like AI Singapore. You can expect things down the road.

And if you are thinking about US-China competition and you're a middle power, you don't want your future to belong either to access to a frontier model from China or the US, then yes, you may want to think in a way that says that open weights are not sufficient.

But I think if you can distinguish between the needs of a smaller organization looking at cost—"Well, I have some generalized tasks and these models are 100 times less expensive"—I think you can feel pretty good about open weights most of the time.

No, that's not the right answer if your consideration is national sovereignty.

I think there's something like: Is your destiny depending on some actor that may or may not share your values? That is different.

Amanda

If I was to bring these two questions together as well, I think that fundamentally this is a question of: Who on the internet can I trust?

If that's not me, then—

I feel like, asking about some of the work at Google, one of the remedies or solutions to this that I have been pushing back on, and we collectively have been taking a stance against, is the concept that you must have real identity on the internet in order to contribute to open source.

That's something we have taken a hard stand on, especially as open-source security has developed.

Before the AI initiatives, there was the open-source security push that came through in 2020, 2021, 2022, 2023.

The concept of who you are on the internet does matter—who you are on the internet that's contributing to open source.

There has been and still is a question around attestations and build materials and sovereignty, and around high-risk systems: Who has access to that and who's contributing to that?

But I think what we, as a community, as scholars, and as industry have to assess is: Who gets to say who somebody's identity is? Who holds that power? Who gets to give identity documents? And who gets to hold that trust of who identity is?

Part of the openness that makes the internet work is that trust is not built off the concept of true identity that's held by any one particular governmental organization.

We have to build trust mechanisms in modern ways, and we have worked for decades now to be able to do that.

So when it comes down to, "How do I know who I can trust? How do I build systems that I can trust?" there are ways that we work together to do that.

But trying to make sure that we're making things more secure by locking things down as individuals and putting trust in certain vectors—you should continue to challenge whether or not that's going to solve the problems we're trying to solve or create more problems.

It can reduce access, reduce approachability, reduce the ability of the total staff to participate, reduce who is allowed to be on the internet, and reduce access to information that they can participate in.

Duane

I love that you brought this up for so many reasons.

This conversation about hard identity on the internet—it's not about trust. It's about who am I mad at, right? Who do I get to be mad at? Who am I going to chase down and try to control?

Ultimately, if any development process requires me to personally know who you are in order to invest trust in you as a person, there are fundamental problems in the rest of the process.

You have to be able to trust that if code comes in, no matter where it came from, you have enough eyes on it and enough processes in place that you can find those issues, that it's inspected well, and that everything else is working.

And we see a reaction as a result of generated contributions across open-source communities, where contributions are going back to just people that they know, which is an understandable reaction, but it also highlights the same problem.

Ilan (Moderator)

Trust all the way through, because this is an open-source discussion.

Wow.

AI is really challenging the economics of open-source systems.

There's the question of, "I have to deal with all this slop contributions." I've got a bunch of mediocre people who have generated a lot of code and pull requests, and now I have to review 100 pull requests and only five of them are actually useful.

That's where AI is changing the economics of open source: putting more pressure on maintainers.

Another one I think is really interesting is reimplementation.

In the past, and again there is disagreement about the maintainers' ambitions for the future, you could say, "If you bring other people along with you, now you kind of control the project."

But now I can point at the project and say, "How do you reimplement this thing?"

So in some cases now there are questions of: Can I then apply a new license to a thing that's inspired by a research project, where the license was not chosen by me?

The cumulative decade, let's say, of work that's gone into that project is useful.

Now somebody can reimplement it by rewiring it, call it a new project, and give it a new license.

So I think we need to begin to scratch the surface of the ways that AI can change the economics of open source.

Some deeper questions for sure.

That's about one minute.

Ilan (Moderator)

Those are the closing comments.

Okay, time for closing comments.

And, Madiha, you were responding to the question about sovereignty?

Madiha

Just to say, again, David Ricardo's work on machinery—he was a British accountant, and he said that in international markets you can be sovereign but have no autonomy because you are part of international markets.

You have to be able to live within them.

Sovereignty and autonomy are much easier when technology is multiplied and commodified, right?

You get to deliver water, and you can have some kind of infrastructure.

So in markets where things are commodified, there's room for sovereignty, I would argue.

With the cloud and things like that, it comes off together.

But otherwise, most sovereign projects fail because they struggle to remain innovative.

So it's about whether you're focusing on a commodified or a non-commodified part of the stack.

Ilan (Moderator)

Specific closing thoughts?

Madiha

Yeah.

I just want to add to something that Ben and Duane were saying about the contributor problem.

This is a huge, huge problem.

One of the ways that I would tie it together with what I started with is governance.

One of the reasons that maintainers are struggling so much on GitHub with AI-generated contributions is because GitHub has not provided tools for maintainers to adequately manage all of the AI contributions.

There are norms that are emerging on GitHub. For example, people are adding the little lobster to say, "This was coded by AI," without saying it in so many words.

Or you have contribution guidelines, or `CONTRIBUTING.md` or `AGENTS.md` files emerging.

Whereas in the past you would see license files or READMEs that would have that kind of policy around contributions articulated, these are now emerging in very disorganized ways because they don't know what else to do.

I started by talking about Hugging Face and the kind of mess of licenses, but all of this really does come down to: What does contribution look like? What does participation and collaboration look like?

And in service of innovation, that really does mean that things are breaking in this really concerning way.

Duane

Yes.

There's never been more open knowledge. There's never been more open collaboration. There's never been more open innovation.

At the same time, there's never been the same potential for closure. There's never been the same potential for powerful gatekeepers.

And there's huge potential for getting involved across all these things.

So we're at a moment of incredible opportunity and challenge, and groups like Creative Commons are incredibly crucial to that.

Amanda

I don't know that I would say "never."

If we go back to Victorian times and what was happening in the first industrialization, it was significantly different from what we are now.

And the reason I'm not picking on you, but the reason I'm picking on that point, is because if we internalize the narrative that this is our only chance to do something, we undercut our ability to think positively about it in the future.

I've had many conversations with people before I took the role at OSI who were very stressed out about the process of the Open Source Definition because of the moment in time that we were in.

And those moments keep coming back.

Yes, there's a lot happening. Yes, it's changing faster than any of us can really keep up with. And yes, the big players feel enormous.

But we still have the ability to enact these changes, and they just keep coming back over time.

So let's not say "never," because it's important, but don't sell ourselves short for the future.

Amanda

I think one thing I really hope that this group especially can keep in mind is that there's this terrible communication pattern a lot of technology has gotten into where they just throw something at you—what they call now a "stack," although this is not a stack—and they're like, "What do you think about this?"

And it's very overwhelming, especially as you get as many ones as we've been getting.

A lot of my work now is to take a breath and sit down and talk to engineers and teams and say, "So, what problem are you trying to solve?"

And to really move from the question of, "Here's this thing I made; what do you do with it?" to, "Is this problem something that is valuable? Is it still worth solving? And is this something that can be solved with information-system techniques?"

So I would say that right now, again, I'm talking about where it's possible to get distracted.

Ilan (Moderator)

Thank you.