



AI and the Commons: Outcomes from the 2023 CC Global Summit Alignment Assembly

Shannon Hong, Kat Walsh, Timid Robot Zehta, Nate Angell

This year's Creative Commons Global Summit theme was AI and the Commons, focused on supporting better sharing in a world with artificial intelligence — sharing that is contextual, inclusive, just, equitable, reciprocal, and sustainable. A team including Creative Commons (CC) General Counsel Kat Walsh, Director of Communications & Community Nate Angell, Director of Technology Timid Robot, and Tech Ethics Consultant Shannon Hong collaborated to use alignment assembly practices to engage the Summit community in thinking through a complex question: how should Creative Commons respond to the use of CC-licensed work in AI training? We identified concerns CC should consider in relation to works used in AI training and mapped out possible practical interventions CC might pursue to ensure a thriving commons in a world with AI.

This paper will discuss the purpose, methodology, process, results, and limitations of the alignment assembly run on 5 October 2023 at the Creative Commons Global Summit.

Purpose

There is significant debate in the Creative Commons community on how the organization should respond to the challenges and concerns around AI. [Community consultations](#) at conferences like MozFest, RightsCon, and Wikimania have revealed concerns on transparency, bias, fairness, and attribution. An additional challenge is the uncertainty around how different jurisdictions will consider copyright and AI. In a US legal context, many scholars consider AI training a fair use of copyrighted work, but lawsuits like *Silverman et al vs OpenAI* and *Authors Guild et al vs. OpenAI* challenge this premise. The upcoming EU AI act may require attribution of any copyrighted material used to develop AI systems. In Japan, laws explicitly permit developers to use copyrighted materials for commercial use.

Many CC community members have concerns about credit, consent, and public benefit in the use of their work in training AI, and CC would benefit from exploring more deeply what solutions the community would appreciate and use. The purpose of the CC Summit alignment assembly is to engage the community in thinking through how CC should respond to these challenges.

Methodology

On 5 October 2023, at the Creative Commons Global Summit, we gathered thirty creators, technologists, and community members in an two-hour alignment assembly to discuss the question: “how should Creative Commons respond to the use of CC-licensed work in AI training?”

An “alignment assembly” is an experiment in incorporating collective input at the ground level, developing new ways to determine what is good, and controlling structures that govern them. Alignment, so that we can bring technology into alignment with collective values. And assemblies, because they assemble regular people, online and across the country or the world, for a participant-guided conversation about their needs, preferences, hopes and fears regarding emerging AI. The feedback is then contributed back to AI labs and policymakers to design for collective good. This model is pioneered by the Collective Intelligence Project (CIP), led by Divya Siddarth, research director at Metagov and a research associate at the Ethics in AI Institute at Oxford, and Saffron Huang, previously a research engineer at Deepmind. CIP offered their feedback on the design of CC’s alignment assembly.

The alignment assembly model is a [Collective Response Process](#), a process in which participants both generate proposals and vote on them, following best practices in [participatory AI design](#). At the Summit, we used Pol.is, an open-source, real-time survey platform, for input and voting. In Pol.is, participants can submit and vote on short text statements; vote options are “Agree,” “Disagree,” and “Unsure.” In order to start the conversation, the facilitators can submit seed comments. [Seed comments](#) “set the tone of the conversation and teach the initial participants how to write good comments.” Participants will generally vote on these seed comments first, before writing their own comments and voting on their peers’ comments.

Creative Commons has previously gathered the community together for consultation at conferences like MozFest, RightsCon, and Wikimania. In these sessions, participants were asked to share their concerns and expectations around AI, and opportunities for the Commons to benefit from AI. While these conversations were productive and interesting,

there were two key issues: first, the ideas of individuals who are more outspoken were more likely to be featured than the ideas of those who spoke less, and second, it was difficult to formalize and understand the results of dialogue that was not necessarily captured in notes or voting.

This alignment assembly works to resolve those issues. With Pol.is, a synchronous voting platform, people who were unlikely to speak in group settings had more opportunity to contribute statements and vote on them. Furthermore, these preferences are explicitly captured and shareable — we will analyze those results in following sections.

Process


At the beginning of the two-hour session, CC General Counsel Kat Walsh framed the conversation, giving a brief introduction to the purpose of the alignment assembly. Participants then introduced themselves in pairs and discussed the question: “What’s one positive contribution or negative contribution of AI to the commons?” This question served to prime the conversation and enable participants to get to know each other. Afterwards, we posed two sets of questions for group discussion and voting, and held a reflection session. The structure of the workshop was as follows:

- 9:45-10:00: Introduction & Icebreakers
- 10:00-10:15: Pol.is 1
- 10:15-10:30: Large Group Discussion
- 10:30-11:00: Small Group Discussion
- 11:00-11:15: Pol.is 2
- 11:15-11:50: Large Group Discussion
- 11:50-12:00: Reflection

In the first Pol.is, we asked participants the question: “What would be important for CC to consider in its AI related policy?” We sourced seed considerations from previous community consultations, and a list of these seed considerations can be found in Appendix A.

Participants voted on the considerations and participated in a large group discussion, where individuals stood up to share context on their agreements and disagreements. For example, a participant shared that they were unsure about the statement on the right, because while they agreed with

People who contribute to the open commons are exploited when commercial services make money using AI trained on open content. CC should consider a strong ethical stance against creator exploitation.



taking a “strong ethical stance against creator exploitation,” they did not necessarily agree with the premise that creators were being exploited when commercial services make money using AI trained on open content.

After this discussion, participants separated in small groups based on types of intervention. The facilitators selected five types of interventions which were most commonly recommended in previous consultations. Participants were also invited to form new groups, if there were topics they believed were not covered, and one additional group was added. These groups can be found in Appendix B.

Participants then voted on the second Pol.is, which asked “How should Creative Commons respond to the use of CC-licensed work in AI training?” Participants added the interventions they discussed into the Pol.is, while also voting on five seed interventions, which were written by CC General Counsel Kat Walsh and Director of Technology Timid Robot Zehta and can be found in Appendix C. In a large group discussion, groups shared their interventions with the larger group and invited criticism and commentary.

The assembly ended with a reflection period, and a final ritual of shaking a fellow participant’s hand, and saying “Thank you for your brain.”

Results

The purpose of the assembly was to consider what actions CC should take around AI and the commons going forward. In that light, we turn to the second Pol.is, the culmination of two hours of discussion. 25 people voted in the final session, with 604 votes cast and 24.16 votes per voter on average, on over 33 statements, including both seed statements and statements provided by participants.

In the only instance of unanimity, all attendees disagreed with the statement: “CC should not engage with AI or AI policy.” This statement was a seed statement created by the facilitators in order to provide the option of doing nothing. The overwhelming rejection of this statement indicates a consensus for Creative Commons to take an active role in addressing the challenges of AI.

Opinion Groups

Pol.is aggregates the votes and divides participants into opinion groups. Opinion groups are made of participants who voted similarly to each other, and differently from other groups. There were three opinion groups that resulted from this conversation.

Group A: Moat Protectors

Group A comprises 16% of participants and is characterized by a desire to focus on Creative Commons' current expertise, specifically some relevant advocacy and the development of preference signaling. They uniquely support noncommercial public interest AI training, unlike B and C. This group is uniquely against additional changes like model licenses and strongly against political lobbying in the US.

Group B: AI Oversight Maximalists

Group B, the largest group with 36% of participants, strongly supports Creative Commons taking all actions possible to create oversight in AI, including new political lobbying actions or collaborations, AI teaching resources, model licenses, attribution laws, and preference signaling. This group uniquely supports political lobbying and new regulatory bodies.

Group C: Equitable Benefit Seekers

Group C, containing 32% of participants, is focused on protecting traditional knowledge, preserving the ability to choose where works can be used, and prioritizing equitable benefit from AI. This group strongly supports requiring authorization for using traditional knowledge in AI training and sharing the benefits of profits derived from the commons. Like group A, this group is against political lobbying in the US.

Conversation Divisiveness

This conversation produced significant consensus, with 17 of the 33 statements producing alignment between participants. Pol.is aggregates statements to show levels of divisiveness: “Statements (here as little circles) to the left were voted on the same way—either everyone agreed or everyone disagreed. Statements to the right were divisive—participants were split between agreement and disagreement.” Most statements are to the left and demonstrate consensus among participants. The most consensus driving statements are those in which different opinion groups vote together, and the most divisive statements are those in which opinion groups differ significantly.



Consensus statements

Divisive statements


Position Statement Analysis

In this section, we highlight specific positions and how the community voted. This section is not comprehensive over all statements, rather it is a subset of the most salient statements.

Preference Signaling

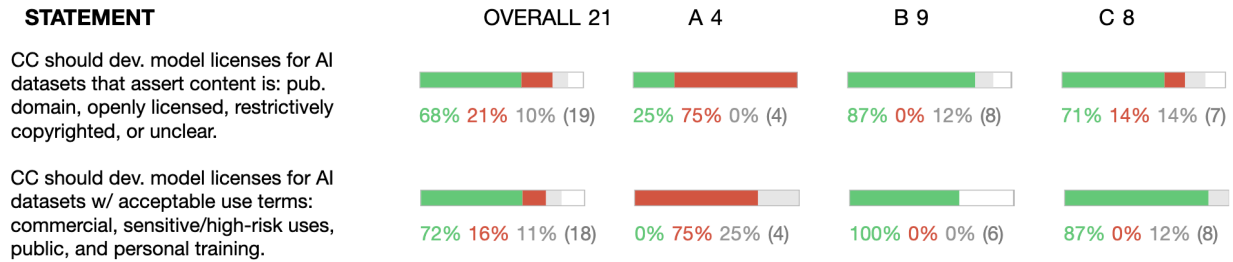
It's been clear in consultations with the community that we need a framework for preference signaling. [A recent blog post from CC](#) explores some of the existing methods of signaling preference and the challenges in developing preference signaling. The outcomes of the assembly further emphasize the appetite in the community for preference signaling, across groups.

The votes indicate that both copyrighted and CC-licensed works should be able to signal preference for use in AI training. More participants are unsure about adopting or endorsing existing mechanisms for preference signaling.



New Licenses

A major discussion point during the alignment assembly was how to create ways for AI developers to indicate that the training content in their models met particular standards. This is part of a larger conversation about licenses for datasets and models. During the small group discussion sections, a new group formed to discuss the possibility of licenses for datasets, indicating interest in exploring this topic. While groups B and C generally agreed with new licenses for models that indicate type of data and acceptable use terms, group A opposed these developments.

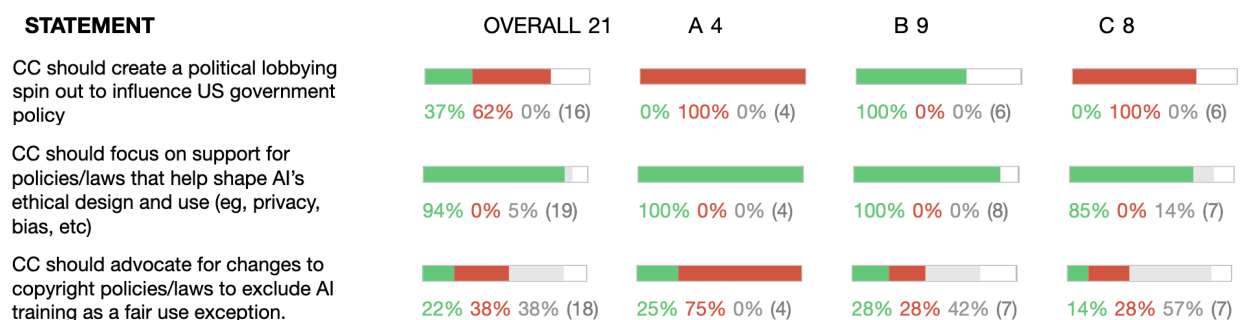


Political Advocacy

Political lobbying and different types of advocacy caused significant disagreement within the assembly. Lobbying is currently outside the scope of Creative Commons’ work, and most participants (62%) disagreed with creating a political lobbying spin out to influence US government policy. One participant shared that lobbying in the US might jeopardize CC’s fundraising position and advocacy for policy positions can include activities other than lobbying. Others may have disagreed with the US-centric approach inherent in this statement.

Participants overwhelmingly voted for Creative Commons to support policies that shape AI’s ethical design and use, indicating a desire for Creative Commons to lead in advocating for ethical AI. However, because “ethical AI” is broadly defined in the statement as “(eg, privacy, bias, etc)”, and the term “ethical” in the context of AI policy has been criticized as a mechanism to distract the conversation from specific policy/rights issues, this statement does not necessarily share actionable advice for Creative Commons beyond a broader desire to be ethical.

When discussing the copyright fair-use exception, participants were unsure (38%) or disagreed (38%) with advocating for AI to be excluded from the fair use exception. However, in conversations with respondents, we found differing interpretations of the statement itself: most believed the statement advocated for AI training to not be considered fair use, but some believed the opposite: that the statement advocated for AI training to be considered fair use. This confusion renders this statement’s results suspect, and further study is needed.

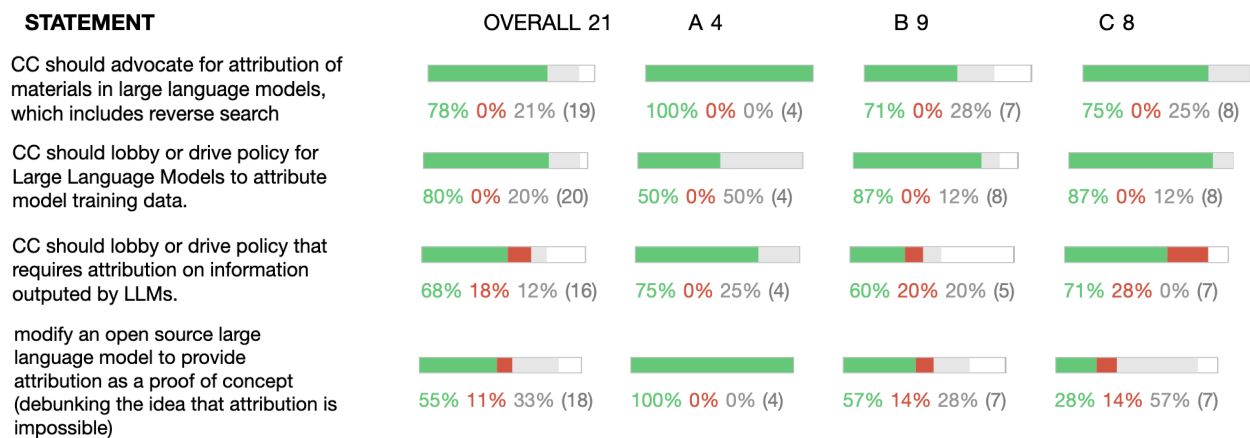


Attribution

Understanding what data is being used in AI training has been an important issue for the Creative Commons community. The general agreement that attribution is needed is reflected across groups. However, the varied amounts of “Unsure” responses indicate a lack of clarity on how this attribution should be provided. For example, the first statement “attribution of materials, [...] which includes reverse search” is an amalgamation of the second and third statement, yet received significantly more votes than the third statement, which required attribution on “information outputted [sic] by LLMs.”

It’s possible the use of the word “lobby or drive policy” had a small chilling effect on participants, as in the above section on Political Advocacy, we find that a majority of participants are against Creative Commons taking on a lobbying role.

There is an uncertain curiosity about the idea of Creative Commons developing its own LLM with attribution as a proof of concept.




Benefits Sharing


While the majority of participants agreed with statements around AI platforms sharing profits with creators of training materials, a significant portion were unsure about the mechanism through which profit sharing might occur. Releasing works under the Creative Commons licenses that do not specify “non-commercial” does relinquish rights and allow reproductions to be commercial. There is some tension between a perceived unfairness and the reality of the licenses offered.

STATEMENT

benefits derived by developers of AI from access to the commons and © works must be broadly shared among all contributors to the commons



New laws or frameworks that redistribute profits from AI to the “data sources”. Meaning the creators, the academics, etc..




Noncommercial Public Interest AI Training

This statement specifically asks participants if they think noncommercial public interest AI systems should be allowed to train on copyright protected work. Participants were divided on this subject; participants were possibly unsure about the implications of the statement, or perhaps this disagreement demonstrates a tension between wanting to honor creator preferences and also empower noncommercial public interest AI.

STATEMENT

Any legal regime must ensure that the use of © protected works for training AI systems for noncommercial public interest purposes is allowed




Traditional Knowledge

With 75% of participants agreeing with this statement, upholding community standards in the stewardship of traditional knowledge is an important value for Creative Commons. Access and use of traditional knowledge elements are often governed by rights, interests, protocols, customs and ownership structures unaccounted for under copyright law, which often casts them into the public domain. The Open Culture team at Creative Commons, in consultation with stewards of traditional knowledge, has found that open licenses such as CC licenses — operating solely within the copyright system — often fall short of expressing the whole range of permissions and/or restrictions with regard to traditional knowledge elements. This statement underlines that while copyright is one lens through which to assess authorization for AI training, it is not the only one, and community’s rights, needs and wishes must be taken into account.


STATEMENT

The use of traditional knowledge for training AI should be subject to the ability of community stewards to provide or revoke authorisation.



Teaching

Creative Commons currently offers courses for educators, librarians, and cultural institutions to understand the open ecosystem and how to use licenses. Participants in the room argued that having access to resources on how AI uses the creative works licensed by Creative Commons would enable technological literacy.



Limitations

There are two key limitations of this assembly: participant sample size and participant representativeness.

Participant Sample Size

There are over 22,000 members in the Creative Commons slack community, which is only a subset of the many more members of the CC community more broadly. Of these members, about 250 people attended the in-person Summit event in Mexico City. 30 people were present and active voting members of the assembly. While many participants were open movement leaders in their countries and represented the perspectives of more individuals, this sample is too small to have a complete picture of the CC community's desires.

Participant Representativeness

We did not perform a demographic survey of the room, but data from the overall conference suggests that American and European perspectives may be overrepresented in our assembly. There was criticism within the session itself that stated the language used to frame the discussion on fair-use was too US-centric. Furthermore, members who self-select to join an alignment assembly on internal AI policy are likely to be already overrepresented in the discussion, and the organization may wish to do more outreach to other groups who might be less likely to engage organically in such discussion.

Conclusion

This alignment assembly has given Creative Commons insight into the community's concerns and ideas for the future. In the future, we hope to run larger alignment assemblies that span time zones and continents to solicit more feedback from the community about how Creative Commons should respond to the challenge of AI.

This work is licensed via [CC Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/). Suggested attribution: "AI and the Commons: Outcomes from the 2023 CC Global Summit Alignment Assembly" by Shannon Hong, Kat Walsh, Timid Robot Zehta, and Nate Angell is licensed via [CC BY 4.0](https://creativecommons.org/licenses/by/4.0/).

Acknowledgement from Shannon Hong

Thank you to Carrie Morgridge and the Morgridge Family Foundation for the Summit Grant that made this project possible. To Saffron Huang and Divya Siddarth of the Collective Intelligence Project, thank you for welcoming me into the fold and giving much needed feedback. Aviv Ovadya and Paula Berman, I am so grateful for your advice. For your comments, support, and feedback, thank you to Brigitte Vézina, Paul Keller, Alek Tarkowski, Ash Gallegos, and Anna Tumadottir. Nate Angell, Kat Walsh, and Timid Robot Zehta, thank you for making it happen. My deepest gratitude to Catherine Stihler: your generosity of spirit and your clarity of vision has guided this project from beginning to end. Thank you.

Appendix A: Seed Considerations for Pol.is 1

1. AI is using existing content for training, but also evolving fast and finding new uses and may have unanticipated effects. CC should make policies that can be flexible and applicable to new changes.
2. Content creators can feel like they have little control over what AI may someday do with their content. CC should prioritize creator control over their work.
3. People who contribute to the open commons are exploited when commercial services make money using AI trained on open content. CC should consider a strong ethical stance against creator exploitation.
4. AI amplifies the biases of its training data. CC should consider policy that prioritizes training AI on diverse data, so that AI can reflect a more diverse and equitable future.
5. CC should make policy that ensures AI's ability to contribute to the commons.
6. Since we believe there is a fair use exception on AI work, CC should focus only on new non-copyright related interventions.
7. CC should be a leader in guiding policies on AI training and copyright globally.
8. Copyright may not be the right solution to AI training issues. CC should think beyond copyright.

Appendix B: Intervention Groups

1. Make Tools for Signaling Preferences of Use in AI
2. Endorse Existing Tools for Signaling Preference
3. Advocate for US exception to the "Fair Use Exception"
4. Advocate for New Laws
5. No Engagement with AI
6. Licenses for Datasets

Appendix C: Seed Interventions for Pol.is 2

1. CC should endorse existing mechanisms rightsholders can use to signal preferences about their works' use in AI training/inputs (eg, Responsible AI Licenses, TDM Reservation Protocol).
2. CC should release new mechanisms to signal rightsholders preferences about their works' use in AI training/inputs (eg, opt in/opt out/no preference).
3. CC should focus on support for policies/laws that help shape AI's ethical design and use (eg, privacy, bias, etc)
4. CC should continue to advocate for minimalist copyright policies/laws that enable diverse, noninfringing access to and reuse of copyrighted works (which would include AI training).

5. CC should advocate for changes to copyright policies/laws to exclude AI training as a fair use exception.
6. CC should not engage with AI or AI policy.