
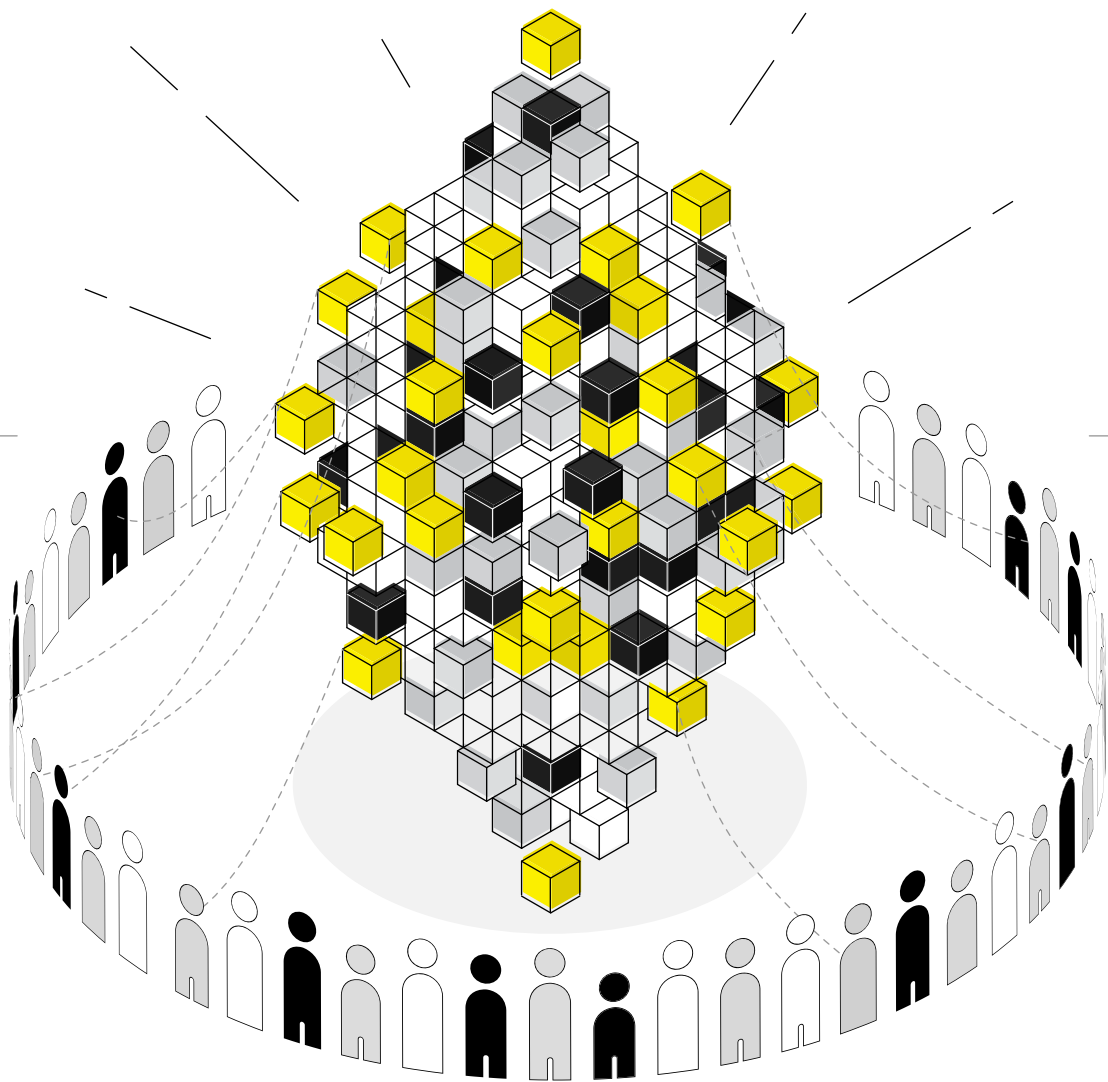


# ALIGNMENT ASSEMBLY ON AI AND THE COMMONS

*Outcomes and learnings*



**OPEN  
\_FUTURE**

**JUNE 2024**

Shannon Hong, Alek Tarkowski

# INTRODUCTION

Artificial intelligence shapes and affects the Digital Commons; however, there is no consensus on AI's specific impacts on the commons and how advocates and stewards of the Digital Commons should seek to manage this impact.

Generative AI is built on the digital infrastructure of the commons and uses the vast quantity of images, text, video, and rich data resources of the internet: open science research, open source code, and various sorts of training data that is either public or openly shared. Most importantly, AI developers train their models on large amounts of content and data shared by a multiplicity of collections and repositories.

Access to the Digital Commons enables innovation and the development of systems that could become the next general-purpose digital technology. But these developments are not without risks and challenges: from bias and lack of transparency to energy consumption and environmental footprint, from new concentrations of power to impacts on creative work – these are all challenges that can influence the commons and need to be addressed.

To this end, the Alignment Assembly on AI and the Commons aimed to answer the question: What do open movement activists, creators, and organizations think about regulating generative AI? An alignment assembly is a combination of a survey and a conversation designed to inform policy debates and align technology development with collective values. The methodology has been developed by the Collective Intelligence Project. Open Future, together with Creative Commons and Fundación Karisma, organized this conversation over six weeks between 13 February and 17 March 2024.

The Alignment Assembly on AI and the Commons followed up on previous work at the 2023 Creative Commons Summit. At this event, a group of 30 activists and experts discussed the regulation of AI in the context of Digital Commons. This resulted in a proposal for a set of principles that were further elaborated at an in-person alignment assembly. Our goal with this initiative was to reach a broader range of individuals and organizations from around the world.

The results of this process show that the emergence of generative AI is challenging established approaches to openness and the Digital Commons. We found consensus around the need to consider values beyond openness and the imperative of funding public infrastructure, and alternatives to dominant AI systems. The principles for regulating generative AI received broad support, although the assembly also revealed potential areas for refinement. We identified two groups with divergent perspectives that need to be reconciled: the Regulation Skeptics and the Interventionists. The differences in perspectives are, to some extent, regional, pointing to different dominant attitudes between North America, Europe, and other regions.

This report from the Alignment Assembly on AI and the Commons begins with an explanation of our methodology. We then review the results of the proposed principles for regulating generative AI, followed by an analysis of the key areas of consensus. We conclude with an analysis of the key differences between the two opinion groups.

## METHODOLOGY

Between 13 February and 18 March, Open Future and its partners hosted a virtual, asynchronous Alignment Assembly on AI and the Commons. We gathered over 260 respondents from more than 40 countries to discuss and explore principles and considerations for regulating generative AI from the perspective of the Digital Commons.

### *Previous work on principles for regulating generative AI*

The Alignment Assembly on AI and the Commons builds on work from the Creative Commons’s summit on AI and the Commons, which took place in October 2023. Open Future and Creative Commons hosted [a workshop on generative AI and its impact on the commons during the summit](#). The group agreed and released seven principles for regulating generative AI. After the Summit, the principles [were published](#) “for further community discussion and to help CC and the global community navigate uncharted waters in the face of generative AI and its impact on the commons.” We are providing a copy of these principles as an annex to this report.

We treated the principles as a starting point, and we were interested in revealing the degree of alignment between activists, creators, and stewards of the commons and among different subsections of the open movement.

### *Alignment Assembly methodology and Pol.is platform*

Alignment assemblies are experimental deliberative processes aimed at incorporating collective input into technology development processes. This methodology was pioneered by the Collective Intelligence Project (CIP), led by Divya Siddarth and Saffron Huang. Siddarth and Huang [describe the alignment assembly model](#) in the following way:

*“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”<sup>1</sup>*

Alignment assemblies are part of a broader trend aimed at increasing deliberation and participatory governance of digital technologies. Citizen panels are a related, more advanced form of deliberative process that is gaining popularity, with citizen panels on AI and data being organized in [Belgium](#) and the [United Kingdom](#).

Alignment assemblies typically take place online, although [the Creative Commons assembly on AI and the commons](#) is an example of one that took place in person. They are typically organized using [Pol.is](#), a survey platform developed by [the Computational Democracy Project](#). Pol.is is “a real-time system for gathering, analyzing and understanding what large groups of people think in their own words, enabled by advanced statistics and machine learning.” It is based on the

---

<sup>1</sup> CIP offered their feedback on the design of this Alignment Assembly. For more information about CIP’s alignment assemblies, see: <https://cip.org/alignmentassemblies>.

concept of Wikisurvey, a survey that is collectively developed by users, with the set of questions expanding through input from participants.

The key feature of Pol.is is its ability to map out differences in opinion, group individual respondents by their opinions, and identify consensus that holds across groups who otherwise disagree. Pol.is conducts this analysis in real time and makes the data available to all participants. This creates a feedback loop that encourages users to add statements that further explore the issue in finer-grained detail. Each opinion group that pol.is identifies is then defined by a representative set of statements. Groups are “differently different,” meaning that the representative statements for group A are not the same as for group B.

Pol.is is not necessarily intended for decision-making but rather “for discovering unrealized possibilities in complex, conflicted situations involving widely diverse perspectives.”<sup>2</sup> The pol.is report from our Alignment Assembly [is available on the Pol.is website](#).

### *Target group and survey process*

Our goal was to bring together people involved in building and supporting the Digital Commons. There are several terms used to describe this group, including open movement, Free Knowledge movement, or Free Culture movement.<sup>3</sup> For the purpose of the assembly, we identified three key groups of stakeholders within the movement:

1. **Activists and experts**, including digital rights advocates and legal experts
2. **Stewards**, people from organizations that steward collections that are part of the Digital Commons such as Wikimedia, Open Access repositories, and heritage collections
3. **Creators**, people who create works that form part of the Digital Commons, broadly: not only visual artists and musicians but also researchers involved with open science or open source AI programmers

Taken together, these groups represent key stakeholders that build and steward Digital Commons, and from this perspective engage with generative AI technologies and their impact. Our hypothesis was that these groups would also represent “ley lines” in the conversation so their individual opinions would line up with their group identity. This hypothesis is not supported by the results of our assembly, and we discuss this further below.

---

<sup>2</sup> Computational Democracy Project. “Polis,” <https://compdemocracy.org/Polis/>; Small et al, 2021, “Polis: Scaling deliberation by mapping high dimensional opinion spaces” <https://www.e-revistas.uji.es/index.php/recerca/article/view/5516/6558>.

<sup>3</sup> Our “Shifting Tides. The open movement at a turning point” study found that many of these people may not necessarily identify with the open movement, therefore, we avoid using the term to describe our target group. Read more: <https://openfuture.eu/wp-content/uploads/2023/06/Shiting-Tides-research-report.pdf>.

Furthermore, we asked participants to select fields of open that they are active in,<sup>4</sup> using a typology that we created based on mapping organizations and individuals active in the open movement. As we demonstrate below, some differences between representatives of these fields are visible. Respondents first answered a demographic questionnaire, which asked about their level of experience in AI, the field of open that they are active in, nationality, and stakeholder group (creator, activist, steward). They were then funneled to the assembly, hosted on Pol.is.

In Pol.is, participants submitted and voted on short text statements; vote options were “Agree,” “Disagree,” and “Unsure.” To start the conversation, we planted seed comments, which are supposed to “set the tone of the conversation and teach the initial participants how to write good comments.”<sup>5</sup> These included statements based on the above-mentioned principles for regulating AI models.<sup>6</sup> Overall, 265 participants cast 13,327 votes on 140 statements.

Over the course of the assembly, we moderated statements on two criteria. First, relevance to the conversation – we removed statements that were obvious mistakes, statements that were beyond the scope of our conversation, and spam statements. Second, a difference between the statement and existing statements: we were mindful of voting fatigue, and we wanted to surface statements that would be additive to the conversation. However, we kept many comments that address similar issues with different tonality and expression, following advice from the Computational Democracy Project, the creators of Polis, that the way something is said is just as important as what is said. We had 23 seed statements. Participants submitted 188 statements, and we rejected 71 of them, resulting in a total of 140 statements that were voted on.

We approached representativeness with the mindset that increased participation would allow us to surface interesting ideas and see novel patterns. We hoped to draw insights that will spark discussion and drive the momentum among those involved in building and supporting the Digital Commons in understanding AI. We based our understanding on the breadth and depth of the open movement from Open Future’s report on the different fields of open.<sup>7</sup> We hoped to gain insight into a broad swath of this Movement, and we wanted to ensure a distribution of perspectives across different fields of open to allow for more interesting possibilities to emerge. The results of the pol.is-based survey should not be understood as demographically representative in the way that traditional quantitative surveys are. The methodology provides qualitative insights into shared views and attitudes in the open movement.

We also faced limitations in analyzing incomplete data. We share these limitations of our methodology in the appendix.

---

<sup>4</sup> Open Future conducted a network based study on individuals involved in the open movement. We used this work to identify subfields within Open that individuals were involved in. Read more here: <https://openfuture.eu/publication/fields-of-open/>.

<sup>5</sup> Computational Democracy Project. “Seed Comments,” <https://compdemocracy.org/seed-comments/>.

<sup>6</sup> The full list of statements is available in the Pol.is report: <https://pol.is/report/r5mnkmeedi5nm8a4dw8ej>.

<sup>7</sup> Open Future. “Fields of Open. Mapping the open movement,” <https://openfuture.eu/publication/fields-of-open/>.

## DEMOGRAPHIC ANALYSIS OF PARTICIPANTS

In the Alignment Assembly, 265 individuals voted, but there were some data limitations: we have demographic data for 231 individuals who filled out the demographic form.<sup>8</sup>

126 activists, 68 creators, and 37 stewards were counted, and they worked across all fields of open. The field most represented was Open Education, which had 54 total participants, followed by Open Culture and Open Software which had 36 participants each.

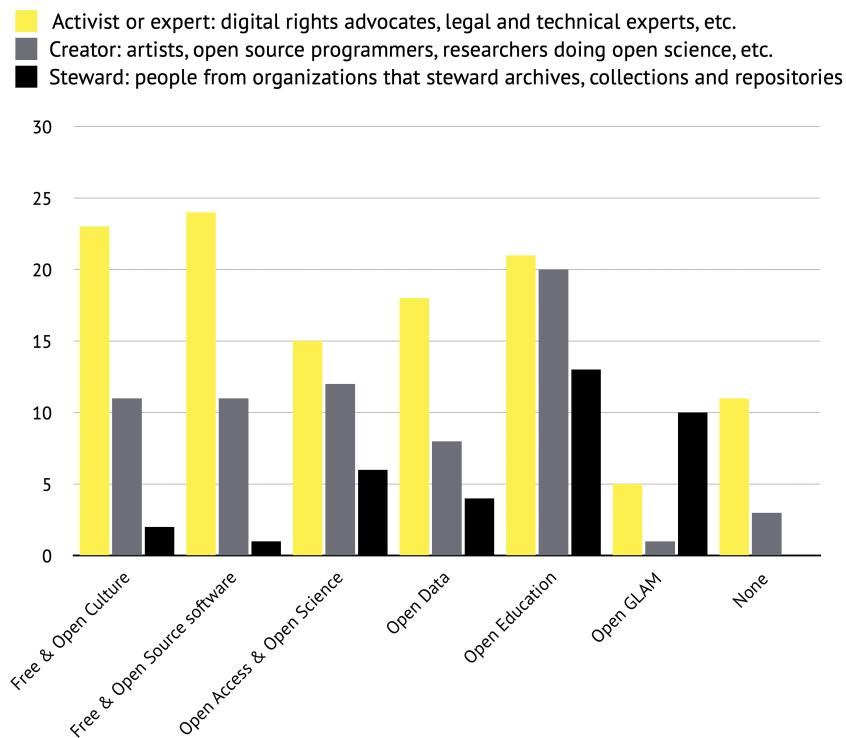


Figure 1: Fields of Open

A minimum understanding of AI is critical for an informed discussion on this topic, so we asked about respondents' AI understanding in the survey. While we did not specifically exclude people for their lack of knowledge, we found that our respondents overwhelmingly self-identified as individuals with expertise in AI (24%) or understanding of AI (71%), compared to those with limited or no experience with AI. We also asked respondents whether they use free or open licensing to share their works. 61% do so regularly, 30% occasionally, and 8% do not use open licenses.

Over 40 countries were represented, but our respondents were mostly clustered in the United States and Western Europe. The greatest number was based in the United States, the United Kingdom, Germany, Canada, and Italy. Latin American organizations like Derechos Digitales and

<sup>8</sup> Polis groups individuals who have voted 7 or more times in order to have enough data to make assumptions about an individual's votes.

Karisma, based in Mexico and Colombia, helped us translate and expand our reach in Latin America. This broader coalition was critical to understanding if there are regional differences in perspective. There were 107 respondents from Europe, 75 from North America, 19 from South and Central America, 13 from Africa, 13 from Asia, and 4 from Oceania.




Figure 2: Respondents by Country

# RESULTS

The starting point for our assembly was the principles for regulating generative AI models and so we begin the presentation of results with an analysis of the levels of support for the different principles. The open-ended and participatory nature of the pol.is survey means that the votes on additional statements can be treated as further exploration of the issues addressed in these principles.

## Examining the Principles

One of the main aims of the assembly was to verify [the principles](#) that were defined during the 2023 Creative Commons annual summit. We wanted to see to what extent there is broader alignment around these principles. Data from the assembly shows considerable agreement, with five out of seven principles being supported by more than 80% of the respondents<sup>9</sup>.

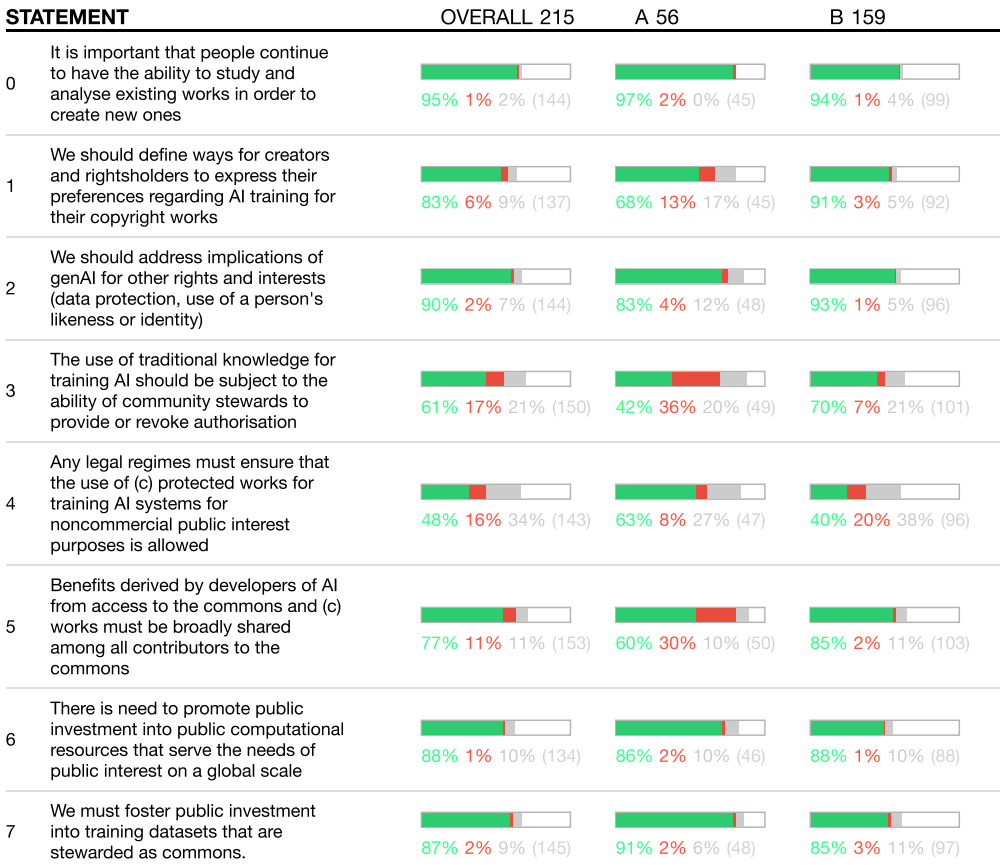


Figure 3: Principles defined during the Creative Commons Summit 2023

<sup>9</sup> For the purpose of reviewing the statements in the Alignment Assembly, we divided Principle No. 7. into two statements, resulting in 8 principles-derived statements overall. In addition, please note the discrepancy in numbering is due to the numbers in Pol.is starting at no. 0 not no. 1.



Almost unequivocal support (95%) for the statement “It is important to have the ability to study and analyze existing works in order to create new ones” (Principle no. 1) is not surprising, as it expresses one of core underlying principles for advocates of Free Knowledge, open science and free software.

More interesting is the high level of support, by 90% of participants, for the statement “We should address implications of genAI for other rights and interests” (Principle no. 3), as it signals a need for a more expansive approach for activists that traditionally have been focusing on copyright. Only slightly lower is the level of support (83%) for the open movement to engage in “defining ways for creators and rightsholders to express their preferences regarding AI training for their copyrighted works” (Principle no. 2).

Respondents also support measures supporting the creation of public AI systems: investments into public computational resources (88%) and public training datasets (87%), which were the two parts of Principle no. 7.

There was lower support (77%) for measures that ensure that benefits derived by AI developers must be broadly shared among contributors to the commons (Principle no. 6).

We saw greater disagreement and unsureness in two key statements. 61% of respondents agreed that “The use of traditional knowledge for training AI should be subject to the ability of community stewards to provide or revoke authorization” (Principle no. 4). And less than half of respondents (48%) agreed that “Any legal regimes must ensure that the use of copyright protected works for training AI systems for noncommercial public interest purpose is allowed” (Principle no. 5). This result, with 16% against and 34% uncertain, spells a major shift for open advocacy, as exceptions for text and data mining (a broader category of which AI training is part) has been one of the main goals of advocacy related to the right to research.

Overall, the pol.is conversation has confirmed support for the majority of the principles, while two of them should potentially be revised.

## *The two opinion groups*

Pol.is aggregates the votes and divides participants into opinion groups, based on an analysis of the combined responses. Opinion groups are made of participants who voted similarly to each other, and differently from other groups. Pol.is groups those respondents who provided seven or more responses – in our case, 211 individuals.

These respondents were divided by the pol.is algorithm into two opinion groups, Group A and Group B.

It is critical to note that groups are not divided on every statement; there is significant overlap in many statements. Thus, groups should be treated as two distinct factions that help understand internal divisions in the open movement on the debate on AI and the commons. Pol.is also does not provide a measure of how distinct or divided in opinions the various groups are, which may be an interesting area for future exploration.

The two groups are the Interventionists (Group B) and the Regulatory Skeptics (Group A), with the first being the dominant one. This group believes that AI regulation is needed to support the commons and that copyright should be used to address other concerns that can broadly be understood as ethical. The second minority group sees openness as facilitating innovation, is optimistic about new ways that AI tools can contribute to the commons, and believes that copyright is often not the right tool for regulating AI.

### **GROUP B: THE INTERVENTIONISTS (74%)**

Group B is the majority of our respondents, who see generative AI as exploiting creators and the commons and violating social norms and potentially even laws. They are also uncertain about the positive impact of AI on the commons and the need to use AI-based tools to create content in the commons. Therefore, they are Interventionists: in agreement with many policy proposals that limit AI development in order to limit corporate interests related to AI and to protect the commons.

### **GROUP A: REGULATORY SKEPTICS (26%)**

Group A is more optimistic than Group B about the positive benefits that AI can have on the commons. They are more inclined to believe that the use of existing works in AI is valuable and reasonable and fits with the values and goals of open sharing. They are also keen to explore how AI-generated content can be part of the commons. They agree with Group B that the emergence of AI technologies might require introducing some restrictions to openness but tend to believe that copyright law is not the right tool for this purpose. They are the Regulatory Skeptics, more critical of specific proposals for AI regulation, and they tend to have an aversion to regulatory overreach. Group A comprises 56 individuals, which is 26% of our participants.

The two groups should not be understood simply as AI optimists and AI pessimists – and it's worth noting that the two groups share a lot of views in common. In particular, they agree that AI technologies need to be regulated in some way and that approaches to openly sharing need to be modified due to the emergence of generative AI. Still, group A is closer to a traditional vision of open sharing and is critical of using copyright-based mechanisms to regulate AI. Group B, in turn, is interested in a wider variety of regulatory measures for managing the commons.

### *Group demographics*

We have both the demographic data and grouping data from pol.is for 172 respondents. Exploring the demographic information between Group A and Group B, we looked at the ratios of different categories of respondents present in each group. There were no significant differences between Group A and Group B with regard to the share of activists, stewards, or creators. This is a surprising takeaway, as we had hypothesized that this axis would be divisive.

On the level of expertise, experts are far more likely than non-experts to be in Group A. 50% of those who self-identify as experts were in Group A, which is far more than the 20% of those who self-identify with less expertise.

Taking into account the various fields of open, there is a higher-than-average ratio of Group A members among participants who work in the fields of Open GLAM, Open Software, and Open Culture. In turn, there are proportionally more members of Group B among those who work in the fields of Open Education, Data, and Science.

Finally, in examining country data, Group A contains proportionally more North Americans than Group B, with 29 out of 55 respondents from Canada and the United States. At the same time, all respondents from the UK and Germany (29 total) were in Group B. These samples were too small to allow for meaningful analysis of regional differences. Still, the results suggest possible differences in attitudes towards AI, the shape of policy debates, and acceptable solutions between different parts of the world.

There are also visible differences between the two groups regarding the more divisive of the seven principles. First, traditional knowledge and community authorization (Principle no. 4) was more controversial in Group A, where over 36% of the group disagreed and 20% were unsure. In turn, Principle no. 5, concerning the use of copyrighted works for noncommercial AI systems, was controversial for Group B members, where 20% disagreed and 38% were unsure.

	<b>GROUP A</b>	<b>GROUP B</b>
<b>STAKEHOLDER GROUP</b>		
Activists	29%	71%
Creators	29%	71%
Stewards	23%	77%
<b>FIELD OF OPEN</b>		
Free & Open Culture	30%	70%
Free & Open Source Software	36%	64%
Open Access and Open Science	27%	73%
Open Data	17%	83%
Open Education	23%	77%
Open GLAM	36%	64%
<b>EXPERTISE IN AI</b>		
I am an expert and I can teach someone else	54%	47%
I understand how AI works, on some level	20%	80%
I have heard, or read about AI	0%	100%

Figure 4: Groups’ demographics

# ON CONSENSUS AND DIVISION

One of the main goals of pol.is is to surface which statements are divisive for the participants. The algorithm clusters participants into distinct opinion groups based on these distinctions. In addition, Pol.is provides tools that allow users to add their own new statements – so that the issues can be further explored in fine grained detail.




Figure 5: Visual representation of the spectrum of consensus and divisive statements

In this section, we highlight key issues on which there is consensus and division. Where there is divisiveness, this paper will go deeper to tease out the threads of why people disagree, using comments that address related problems with different language and expression.

## CONSENSUS

This Alignment Assembly revealed three key areas of consensus. First, the prioritization of values beyond openness in considering the open movement’s policies towards AI (related to principle no. 3). Second, the need for public investment in AI (related to principle no. 7). And third, the call for the open movement to make education about AI and its impact, and public-facing communication on AI a priority.

### *Beyond openness*

Considering values beyond openness alone resulted in the emergence of a new key area of consensus – possibly the most interesting result of the assembly. Almost all participants agree that consideration of ethics in AI is just as important as openness of AI systems (statement 110) and that openness is not the only value relevant for activists, creators, and stewards in the commons (statement 18). More important is the result from statement 13, where 72% of respondents do not agree that “any restrictions to sharing, including ethical ones, are against the spirit of “open.” We read this result as a sign that the open movement believes that the emergence of AI technologies spells fundamental shifts in how openness is defined and what limitations are considered acceptable. That this group agreed that the “spirit of open” allows for ethical restrictions is in some ways, an affirmation of the changing nature of the open movement.

STATEMENT	OVERALL 215	A 56	B 159
13 Any restrictions to sharing, including ethical ones, are against the spirit of "open"	14% 72% 13% (118)	14% 73% 11% (42)	14% 71% 14% (76)
18 Openness is not the only value our community should care about.	90% 2% 7% (139)	90% 4% 4% (44)	90% 1% 8% (95)
110 Openness in AI systems should be aligned to ethical principles.	88% 1% 10% (96)	84% 2% 13% (38)	91% 0% 8% (58)

Figure 6: Beyond openness

## Public investment in AI

Public investment, public alternatives, and public good were a resounding consensus point for this community. Both Group A and Group B voted overwhelmingly for public investment that serves the public good, and both desired non-commercial public alternatives and public participation in AI. This suggests that advocates for openness are strongly aligned with those promoting ideas such as [Public Digital Infrastructure](#), AI systems as digital public goods, or public options for AI.

STATEMENT	OVERALL 215	A 56	B 159
6 There is need to promote public investment into public computational resources that serve the needs of public interest on a global scale	88% 1% 10% (134)	86% 2% 10% (46)	88% 1% 10% (88)
7 We must foster public investment into training datasets that are stewarded as commons.	87% 2% 9% (145)	91% 2% 6% (48)	85% 3% 11% (97)
47 We need alternative public good AI systems which the public can participate and opt in to.	80% 6% 12% (136)	76% 8% 14% (47)	83% 5% 11% (89)
85 There must be non-commercial open public alternatives to closed corporate AI systems.	89% 2% 8% (121)	87% 4% 7% (41)	90% 1% 8% (80)

Figure 7: Public investment in AI

## Thought leadership

Two statements that were ranked highly in group-informed consensus were on the open movement's thought leadership on AI. The first of them (statement 31) is a call to support citizens in developing skills to both understand and critique AI, and the second (statement 146) calls for the open movement to be involved in the public-facing conversation around AI. The open movement can be insular, and perhaps these statements indicate a desire for the movement to become more influential in guiding public opinion around AI.




Figure 8: Thought leadership

## DIVISION

The key areas of division between Group A and Group B were, first, the extent of AI’s exploitation of the commons; second, the role of AI in producing commons-based resources; and third, the use of copyright as a legal tool in discussing AI. As facilitators of this conversation for the open movement, we wish to hold these areas of division with reverence so as not to aggravate or cause further divides between individuals who hold these diverging opinions, but rather to help us more coherently work together with awareness of difference in our movement.

### Exploitation

Group B believes that Generative AI developers and providers, and the systems that they create, exploit creators and the commons for profit. Group B has a consistently negative impression of AI’s legality, ethics, and value. Group A mostly disagrees or is unsure about the extent to which AI is exploitative. Some statements around exploitation and generative AI are emotionally charged – see, for example, statement no. 23. Nevertheless, it’s interesting that nearly 30% of all respondents agreed that GenAI is an “effort by big tech to devalue creators’ labor” (statement no. 23). On the other hand, Group A is almost universally against this statement.

Overall, 63% of individuals in our participant pool believe that generative AI is exploitative of creators (statement 46), indicating that the open movement may need to devote more attention to concerns of exploitation and, more broadly, of AI-related impact on creative labor.




Figure 9: Exploitation

## The commons and AI tools

A significant discussion point was how AI should be used as a tool for creating the commons. 59% believed that there is a benefit to AI-generated contributions to the commons (statement no. 33), with this proportion being much higher in Group A than in Group B. However, there was significant disagreement and confusion on whether or not AI should be used to generate educational resources (statement no. 39, 48% said no) and a desire to steward verified, human-crafted reference points (statement no. 29, 68% want commons repositories to be human-crafted). The answers suggest that the distinction between human and synthetic (AI-generated) content is relevant for the commons, with possible policy considerations for both types of content.


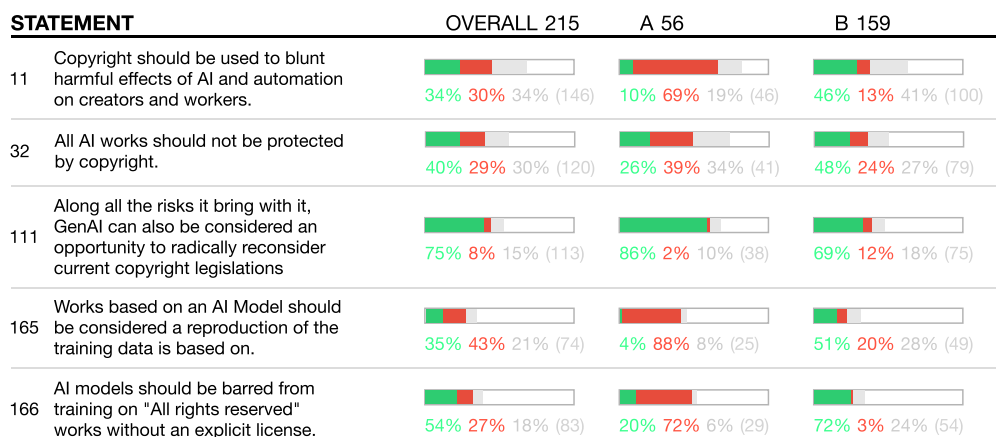


Figure 10: The commons and AI tools

## Copyright

While there is agreement that generative AI can be considered an opportunity to reconsider current copyright legislation, groups differed significantly in how this reconsideration should be done. Group A is generally against most suggestions that were offered on how copyright law should be used to shape the training of AI models and categorically against using copyright to specifically “blunt the harmful effects of AI and automation on creators and workers” (statement no. 11) While Group A wants to steward ethical AI broadly speaking, it also believes that copyright is the wrong tool to consider this mission.



Although participants see opportunities for new legislation and legal innovation, there is significant disagreement on how AI applies to current copyright law. Many are unsure whether AI works should be protected by copyright and, if so, how. Overall, there is significant confusion and uncertainty on the right path in copyright legislation and policy related to generative AI. It is worth noting that the two groups have opposing views on the statement, “AI models should be barred from training on “All rights reserved” works without an explicit license.

High levels of “unsure” responses for some statements (11, 32) signal a lack of clarity among participants on specific regulatory proposals – suggesting both the emergent nature of these debates and the opportunity for broader outreach and education on these issues. The Computational Democracy Project team offers that in pol.is conversations on technical topics, high levels of “pass” indicate individuals’ willingness to receive new information and to hold off forming their opinion until further equipped. These are positive qualities associated with a culture of learning.



## CONCLUSION

We would like to highlight four high-level conclusions from the Alignment Assembly.

**First, the emergence of generative AI challenges established approaches to openness, sharing, and the commons: there is a consensus that the open movement should consider values other than openness alone.** [Our research](#) shows that a large set of factors, beyond the emergence of generative AI, contribute to this change in attitude among movement members. Revisiting the principles and norms that underpin the various fields of open will be as relevant as establishing shared positions on regulating generative AI.

**Second, the seven principles on regulating generative AI that emerged from the Creative Commons summit received high support.** While further refinement is needed to secure as broad an endorsement as possible, this consensus indicates that the open movement will be able to reach a firm set of principles for regulating generative AI.

**Third, public investment, public accountability, and public infrastructure in AI are issues with high consensus among the participants.** Both groups believe in the benefits of public involvement in AI, which suggests a clear direction for advocacy and movement-building

**Fourth, the Assembly shows significant differences in views between participants from North America and the rest of the world, or at least from Europe.** While there is agreement on the principles, there are regional differences with regard to some of the key and most divisive statements. These differences can impact advocacy work and suggest that cross-regional dialogue is needed to explore these differences and seek shared advocacy positions.

This Pol.is-based Alignment Assembly is one step in defining a shared position and a unified response to changes related to AI technologies. It has had the value of being swift and asynchronous and, therefore, relatively inclusive. But we also know that many people whose voices should be heard and whose opinions matter were not present in this conversation. And so we acknowledge that pol.is conversations do not offer means for a deeper exploration of issues being discussed. We hope that the results of the Assembly will serve as a basis for further explorations of shared positions on regulating generative AI.



## ABOUT THE AUTHORS AND OPEN FUTURE

[Open Future](#) is a European think tank that develops new approaches to an open internet that maximize societal benefits of shared data, knowledge and culture.

[Shannon Hong](#) is a Research Fellow at Open Future. She is researching democratic inputs into AI. Previously, as a Product Manager at Scale AI, she led projects labeling for self-driving cars and creating a fairer marketplace of labor for digital work. Shannon studied economic development and data science at Berkeley.

[Dr Alek Tarkowski](#) is the Strategy Director at Open Future. He holds a PhD in sociology from the Polish Academy of Science. He has over 15 years of experience with public interest advocacy, movement building, and research into the intersection of society, culture, and digital technologies.

## ACKNOWLEDGMENTS

The Alignment Assembly was organized by Open Future together with Creative Commons and Fundación Karisma.

Thank you to our collaborators: Patrick Connolly coded the bespoke pol.is instance that we used for the conversation and provided admin support during the assembly. Alicja Peszkowska was responsible for outreach and engagement with assembly participants.

Thank you to Divya Siddarth, Saffron Huang and Flynn Devine of the Collective Intelligence Project, as well as Liz Barry of the Computational Democracy Project for advice throughout the project. Thank you to Paul Keller and Luis Villa for the feedback at various project stages. Gratitude to our thought partners Creative Commons and Karisma, particularly Anna Tumadóttir, Viviana Rangel, and Maria José Parra, for their contributions to the project. Thank you to our committed audience members, Open Knowledge Foundation, the Communia Association, and Derechos Digitales. Gratitude and appreciation to the thirty individuals who helped us with sense-making and understanding the results of our report.



This report is published under the terms of the [Creative Commons Attribution License](#).

# ANNEX 1: PRINCIPLES FOR REGULATING GENERATIVE AI MODELS

The following seven principles for regulating generative AI models were formulated during a workshop on AI, creators and the Commons organized by Open Future and Creative Commons. The workshop took place on 3 October 2023 in Mexico City, as a side event to the Creative Commons Summit 2023. The principles are meant to ensure that regulation of AI technologies serves to protect the interests of creators, people building on the commons (including through AI), and society's interests in the sustainability of the commons.

The original version of the principles can be found [on the Creative Commons site](#).

## *Background considerations*

1. Recognizing that around the globe, the legal status of using copyright-protected works for training generative AI systems raises many questions and that there is currently only a limited number of jurisdictions with relatively clear and actionable legal frameworks for such uses. We see the need for establishing a number of principles that address the position of creators, the people building and using machine learning (ML) systems, and the commons, under this emerging technological paradigm.
2. Noting that there are calls from organized rightholders to address the issues posed by the use of copyrighted works for training generative AI models, including based on the principles of credit, consent, and compensation.
3. Noting that the development and deployment of generative AI models can be capital intensive, and thus risks resembling for (or exacerbating) the concentration of markets, technology, and power in the hands of a small number of powerful for-profit entities largely concentrated in the United States and China, and that currently most of the (speculative) value accrues to these companies.
4. Further noting that, while the ability for everyone to build on the global information commons has many benefits, the extraction of value from the commons may also reinforce existing power imbalances and in fact can structurally resemble prior examples of colonialist accumulation.
  - A. Noting that this issue is especially urgent when it comes to the use of traditional knowledge materials as training data for AI models.
  - B. Noting that the development of generative AI reproduces patterns of the colonial era, with the countries of the Majority World being consumers of Minority World's algorithms and data providers.
5. Recognizing that some societal impacts and risks resulting from the emergence of generative AI technologies need to be addressed through public regulation other than

copyright, or through other means, such as the development of technical standards and norms. Private rightsholder concerns are just one of a number of societal concerns that have arisen in response to the emergence of AI.

6. Noting that the development of generative AI models offers new opportunities for creators, researchers, educators, and other practitioners acting in the public interest, as well as providing benefits to a wide range of activities across other sectors of society. Further noting that generative AI models are a tool that enables new ways of creation, and that history has shown that new technological capacities will inevitably be incorporated into artistic creation and information production.

## *Principles*

1. It is important that people continue to have the ability to study and analyse existing works in order to create new ones. The law should continue to leave room for people to do so, including through the use of machines, while addressing societal concerns arising from the emergence of generative AI.
2. All parties should work together to define ways for creators and rightsholders to express their preferences regarding AI training for their copyrighted works. In the context of an enforceable right, the ability to opt out from such uses must be considered the legislative ceiling, as opt-in and consent-based approaches would lock away large swaths of the commons due to the excessive length and scope of copyright protection, as well as the fact that most works are not actively managed in any way.
3. In addition, all parties must also work together to address implications for other rights and interests (e.g. data protection, use of a person's likeness or identity). This would likely involve interventions through frameworks other than copyright.
4. Special attention must be paid to the use of traditional knowledge materials for training AI systems including ways for community stewards to provide or revoke authorisation.
5. Any legal regime must ensure that the use of copyright protected works for training generative AI systems for noncommercial public interest purposes, including scientific research and education, are allowed.
6. Ensure that generative AI results in broadly shared economic prosperity – the benefits derived by developers of AI models from access to the commons and copyrighted works should be broadly shared among all contributors to the commons.
7. To counterbalance the current concentration of resources in the the hands of a small number of companies these measures need to be flanked by public investment into public computational infrastructures that serve the needs of public interest users of this technology on a global scale. In addition there also needs to be public investment into training data sets that respect the principles outlined above and are stewarded as commons.

## **ANNEX 2: LIMITATIONS OF THE POL.IS METHODOLOGY**

### 1. Incomplete demographic and voting data.

Our aim was to collect additional demographic data through a Typeform survey displayed on the conversation's website. The survey was filled by 230 individuals out of 265 who participated in the pol.is conversations. Furthermore, pol.is groups only those respondents who voted on at least seven statements. This meant that, in our case, only 211 individuals were grouped. Both issues ultimately reduce our sample size. All told, out of 292 interactions with either the Typeform form or the Pol.is survey, 172 respondents had both been grouped and had reconciled demographic information.

### 2. Challenges with iterative surveying

Pol.is methodology assumes that respondents will return to the survey multiple times in order to review new statements and add their own. While we saw some activity by returning respondents, overall we observed limited iterative engagement. This is addressed to some extent by the pol.is a system whose algorithm selects which questions are displayed to both new and returning respondents based on the relevance of these questions for establishing groups, consensus, and division.

### 3. Outreach and inclusion

We made an effort to make the conversations inclusive and, in particular, took care to conduct outreach in various languages and regional networks. The Pol. tool offers automatic translation of statements, and in addition, we translated all content, with the help of the Karisma Foundation, into Spanish. Nevertheless, the response rate from Global Majority countries was low.

### 4. Limited means for in-depth conversation

Experts who helped us with sense-making commented that they would want to understand more about the rationale of someone's vote or statement. Similarly, as we analyzed the results, we asked: how could we get the context or lived experience that informed the statements that participants contributed? What does it mean to name these groups and interpret their interest areas? How much further do we need to expand the conversation before coalescing on a movement-wide set of policy positions? More research into the positions and attitudes of activists, creators, and stewards from the commons is needed.

## **ANNEX 3: LIST OF STATEMENTS**

Below is the list of the 140 moderated statements participants voted on, including 23 seed statements and 117 submitted by participants. The screenshots in the report above show the original numbering from the pol.is system. Note that the numbers differ between the screenshots and this list because in its listings, Pol.is includes statements that were moderated out. The additional discrepancy in numbering is due to the numbers assigned by Pol.is starting at no.0 not no. 1.

### *Seed statements*

1. It is important that people continue to have the ability to study and analyse existing works in order to create new ones
2. We should define ways for creators and rightsholders to express their preferences regarding AI training for their copyright works
3. We should address implications of genAI for other rights and interests (data protection, use of a person's likeness or identity)
4. The use of traditional knowledge for training AI should be subject to the ability of community stewards to provide or revoke authorisation
5. Any legal regimes must ensure that the use of (c) protected works for training AI systems for noncommercial public interest purposes is allowed
6. Benefits derived by developers of AI from access to the commons and (c) works must be broadly shared among all contributors to the commons
7. There is need to promote public investment into public computational resources that serve the needs of public interest on a global scale
8. We must foster public investment into training datasets that are stewarded as commons.
9. Organizations that steward collections and repositories should label content as human or synthetic.
10. Transparency of training data should be a requirement for any AI model or system.
11. Without investment into broadly available training datasets, large companies will dominate AI.
12. Copyright should be used to blunt harmful effects of AI and automation on creators and workers.
13. Requiring that AI training is done only with content licensed by rightsholders will help large media companies, but won't help most artists and creators

14. Any restrictions to sharing, including ethical ones, are against the spirit of “open”
15. A project that shares resources only with known organizations and people can still be meaningfully open.
16. I am worried that AI will make it even harder for artists to make a living.
17. AI systems should have the same rights as humans to read and consume material.
18. AI reduces people's motivation to share works openly.
19. Openness is not the only value our community should care about.
20. AI training is just like any other use of openly licensed content.
21. Open repositories (OA journals, GLAM collections, etc.) need special strategies to cope with AI
22. Addressing negative environmental impacts of AI is part of the conversation about AI and the commons.
23. Governments should build their own large language models.

### ***Statements submitted by participants***

24. GenAI produces mostly lousy art and bad text, but is being sold as transformative in an effort by big tech to devalue creators' labor.
25. Like last bubble's killer app, paying ransomware with cryptocurrencies, the “best” of genAI is deepfake porn and political misinformation.
26. AI should enhance, rather than replace, human cognition & creativity. It follows that AI legislation should also be human-centric in design.
27. In light of AI advancements and data sovereignty issues, the “open” movement should consider a rebrand
28. All AI works should not be protected by copyright.
29. Non-copyrightable AI work can benefit open culture, for example Wikipedia.
30. AI-generated photos cannot replace human-taken photos and will mislead people.
31. Being open means being against extractive economies of all kinds from minerals through to data
32. AI has fundamentally changed what being “open” means
33. Open approaches in an AI era need new business models

34. Open must always be transparent
35. AI should not be used to generate educational resources (e.g. Wikipedia, Oxford dictionary)
36. The open movement is in crisis at present
37. Let's separate the idea of 'open' from the idea of 'business'
38. Big tech and openness are mutually exclusive
39. AI models contain no copyrightable human expression and so are not copyrightable.
40. All gen AI produced works are in the public domain constituting a new kind of synthetic commons.
41. Gen AI exploits creators and the commons for profit without permission, credit or compensation.
42. We need alternative public good AI systems which the public can participate and opt in to.
43. The open community must speak to differences in the AI hype/marketing vs actual uses
44. AI work may be copyrightable if there is a high degree of human involvement.
45. We need to limit AI to being used militarily (e.g. cyberattack).
46. AI generated content is only in the public sphere once a person has decided to publish it. The responsibility is still on individuals.
47. There is no "AI works". There is only work done by, and published by, people assisted with various tools.
48. Wikimedia can benefit from AI assistance, but its community will want to maintain editorial and publishing power to keep its validity.
49. AI brought forward the questions of ownership and fair use. We've benefitted from not having to worry about it until now.
50. Just like a low-effort snapshot photograph is copyrightable, so too should AI-assisted works with low human involvement be copyrightable.
51. By default LLM should be prevented from crawling the web; a developer specified flag should indicate whether it is acceptable to download.
52. AI systems must be able to identify when they are hallucinating or inventing information.



53. The "open science" community has already done great work to offer competitive models with limited external investment.
54. AI systems must always offer a means of exporting user data for the purposes of data portability.
55. There must be non-commercial open public alternatives to closed corporate AI systems.
56. AI systems should not undermine human autonomy or agency.
57. Users of AI systems must be granted ownership of outputs created through their interactions and inputs.
58. Users of AI systems should receive a share of profits if their data or usage trains or improves the AI's capabilities.
59. AI-generated articles should not be allowed on Wikipedia, even if the content is accurate.
60. GenAI violates norms and probably laws but is somehow protected by being "at scale" and by being associated with vast monied interests.
61. Commons repositories such as Wikimedia projects must remain human-crafted to provide a verified reference point
62. Support citizens in developing the basic skills needed to understand AI, GenAI and mainstream applications with a critical approach
63. Current open source licenses aren't applicable to AI models as models don't contain any source code.
64. The personification of AI serves to undervalue true human contributions to arts and sciences
65. Open Community has more to gain by focusing on Redress and Terms of Service rather than try to control the Hype cycle
66. A library or museum should be set up to collect AI-generated works.
67. AI systems must disclose how their continuous long-term use could influence behavior, habits, perceptions, or mental health over time.
68. Openness in AI systems should be aligned to ethical principles.
69. AI should not generate "fake" photos of real people.
70. Open GLAM will benefit from digging deeper into studying and examining models for open data, licensing and attribution from open science.
71. If open, AI can increase access to knowledge

72. Big AI can be curtailed by regulating Big Data.
73. For an AI model to be meaningfully considered open source, all training data must be public and re-usable for other open source AI models.
74. Large scale, automated use of openly licensed works is a mark of success. AI has issues, but it should not shape our definition of openness
75. The open movement has focused on freedoms and permissions at the expense of other norms. The issue around AI just highlights this.
76. Educational organizations need easy and direct access to safe, open, transparent and independent AI systems.
77. Along all the risks it brings with it, GenAI can also be considered an opportunity to radically reconsider current copyright legislations
78. For the importance of linguistic diversity, LLM Gen-IAs should be programmed in native languages to avoid language extinction
79. The open movement should be paying attention also to sovereignty related to traditional and Indigenous knowledges and cultural works.
80. Indigenous data and content raises particular ethical questions for openness
81. Basic info about how GenAI works should be included in every educational program, starting with primary school or even kindergarten, asap
82. The Open Source Definition's requirement that open source allow all uses is in tension with current or future AI laws.
83. Open source can't solve all the world's problems.
84. Predictability for users of openly-licensed works is more important than reducing legal liability for creators of openly-licensed works.
85. Open should focus on collaboration, leaving ethical issues to be solved through government.
86. The open movement needs a 'decolonial turn' to acknowledge its roots in the global north and explore what openness means for global south
87. Openness can be highly exploitative esp. when it concerns data that represents or is collected by/from people in underserved communities
88. Openness in AI systems is not binary (open vs. closed). We should rather think of a gradient of openness relating to different AI elements.

89. Advances in the generation of synthetic data will lessen or eliminate many commons-based concerns over generative AI systems.
90. Open-source models release model weights and code under an open-source license. A description of what data was used is sufficient to be open
91. Orgs & people involved in the Open movement need to be actively involved in public-facing conversations surrounding the use & training of AI
92. The principle that publicly funded research should be publicly available should be expanded to computed science.
93. Generative AI will necessitate the rethinking of copyright, attribution, and royalty structures.
94. Technological approaches to tracking provenance, trust (ex.distributed ledger-based solutions) are a promising way improve LLM data quality
95. The open movement should work on a open GenAI model.
96. AI can aid in generating educational resources but needs human supervision.
97. Copyright holders may not be the original creators of the work so focusing on copyright is not necessarily best to address artist job loss
98. Commercial use of AI models trained on commons based data should require contributing a % of revenues back to the commons.
99. The training & fine-tuning of AI by everyone around the world will only happen if they are contributing to a widely available open platform.
100. Democracies should consider regulating chatbot speech in election-related contexts. Otherwise, regulation should focus on use cases.
101. To save resources and keep things simple, we need a focus on AI-sufficiency: Does this system require AI at all?
102. Integrity in science we don't know what it refers to, such as the case of predatory journals that operate online (not always open-access)
103. AI should be able to make use of anything existing for its training, just like humans are permitted to do.
104. While copyright is helpful in regulating Big Tech short term, overemphasis is likely to cause further centralization of power long-term
105. Governance is integral to an understanding of Commons, and Commons is therefore a better concept than Open, which lacks the same emphasis.

106. An AI model should be considered a derivative work of all the data it was trained on.
107. Works based on an AI model should be considered a derivative work of all the training data it was based on.
108. Works based on an AI Model should be considered a reproduction of the training data is based on.
109. AI models should be barred from training on "All rights reserved" works without an explicit license.
110. The Commons should be cultivated to provide open training data for all purposes.
111. KI wird massiv zur gesellschaftlichen Zerstörung beitragen und sollte verboten oder stark reglementiert werden.
112. Si bien internet es una fuente pública de conocimiento, los modelos de IA deben pagar por entrenar sus bases de datos si buscan el lucro
113. Every source online should disclose whether or not it's content has been AI generated
114. I consider myself an activist.
115. The Commons should be cultivated to provide open training data to its members; big tech should get access by paying fees
116. Just because something is publicly available, it doesn't mean that it's ok to exploit that resource for private profit
117. I feel more respected when I'm informed whether something I'm reading is written by an AI vs a human.
118. Publicly available data should be accessible to AI developers, however, there should be opt out systems for creators.
119. Solutions to many challenges and concerns related to AI (ethics, labor, harm, copyright) begin with AI literacy across all education levels.
120. Data repositories should be funded to devise/implement interoperable approaches for data AI-readiness checks, metadata and distribution.
121. There are adequate technical tools to determine if an AI has been trained on a specified dataset.
122. There should be an agreed way of specifying the person legally responsible for the output of a generative AI.
123. Professional codes of conduct have a part to play in regulating generative AI.

124. Generative AI should have an independent legal personality, like a corporation.
125. No plantear criterios de transparencia en la IA podría afectar la sostenibilidad del patrimonio y conocimiento de comunidades ancestrales
126. Diseñar un sistema que permita a creadores excluir su contenido de los datos usados para entrenar GenIA es un arma de doble filo
127. AI-generated synthetic media presents critical new issues around trust and authenticity.
128. Generative AI requires from societies and citizens of the world to Rethink everything!  
We should not waste this crisis.
129. Open should also mean "efficient" for re-use and focus on proportionality
130. Indigenous communities should be involved in the development of training datasets in their mother tongue and their ancestral knowledge
131. Indigenous communities should be involved in the stewarding of training datasets, ethical requirements, alignments testing and AI governance
132. The biggest issue with GenAI is a lack of transparency and accountability.
133. Open licenses should be adapted to allow licensors to select different permissions for AI training
134. Using works to train AI should be based on opt-outs, not opt-ins. We should focus in creating standards for opting-out of AI training.
135. I don't think I understand enough to make a coherent statement or identify what is missing or needs to be added
136. AI systems should not be used for military purposes.
137. It depends on how the generated AI art is generated & it's artistic merits
138. There should be a global movement for ethical, open AI training data sets to lower competitive barriers and improve AI
139. A wikipoll like Pol.is is a good start but we need actual deliberation to collectively answer the complex questions this conversation poses.
140. Private commercial entities are likely to exploit issues of trust and authenticity around synthetic media to gain control over the commons.