

Civic empowerment in the development and deployment of AI systems



Final Report of FAccT CRAFT Workshop

Cite as

Renée Sieber and Ana Brandusescu. Final Report. FAcct CRAFT: Civic empowerment in the development and deployment of AI Systems. Critiquing and Rethinking Accountability, Fairness and Transparency Workshop at ACM's Fairness, Accountability and Transparency in Machine Learning. Montreal, Canada. June 30, 2021.

Table of contents

Acknowledgements	5
Executive Summary	6
1. Introduction	7
2. Speaker presentations	7
2.1 Renee Sieber on a Brief & Deep Dive into Roots of Participatory Theory	8
2.2 Mich Spieler on Community Activism in Tech and AI	12
2.3 Bianca Wylie on Community Activism in Tech and AI	14
2.4 Jean-Noé Landry on How a Nonprofit Builds Public Participation and Empowerment in Other Organizations	17
3. Discussions on how key terms are defined	19
3.1 If we're talking about citizen participation & civic empowerment, who could be participating/empowered?	19
3.2 What role do individuals, groups, organizations play & what responsibility do they have in the development and deployment of AI systems?	21
3.3 What words do you use to define civic empowerment and how would you “operationalize” those?	22
4. Assessment questionnaire	24
4.1 Existential questions	25
4.2 Epistemic questions:	25
4.2.1 Defining publics and engagement	25
4.2.2 Defining and Enforcing Accountability	26
4.2.3 Defining and Ensuring Empowerment	27
4.3 Process questions (once you’ve decided on the above...)	28
4.4 AIA-type questions (to ask the developer)	29
Appendix A: Resources and Documents	31
Appendix B: Workshop Part II Breakout session	32
Appendix C: What would be contained in a questionnaire?	39

Acknowledgements

In alphabetical order

Organizers: [Ana Brandusescu](#) and [Renée Sieber](#)

Facilitators: Abigail Adu-Daako, Rob Davidson, Sam Lumley, Suthee (Peck) Sangiambut and Zhibin (Allen) Zheng

Speakers: Jean-Noé Landry, Mich Spieler and Bianca Wylie

Attendees: Letitia Adu-Sarkodie, Erin Arnesen, Renata Avila, Jacqui Ayling, Jumana Baghabra, Rafael Becerril Arreola, Amara Bains, Solon Barocas, Marguerite Barry, Natalie Byrom, Madeleine Chang, Sarah Chu, Tessa Darbyshire, Emily Denton, Mark Díaz, Wren Dougherty, Eshan, Clàudia Figueras, Bernease Herman, Maria Heuss, Abigail Jacobs, Rachit Jain, Julia, Mike Katell, Joshua Kroll, Alexander Laufer, Elisabeth Lehagre, Ilse van der Linden, Christina Lu, Michael Madaio, Andrei Mircea, Natalia Modjeska, Midas Nouwens, Ioanna Papageorgiou, Tina M. Park, Ted Pedersen, Manuel Portela, Bogdana Rakova, Angie Raymond, David Robinson, Mayra Russo, Levent Sagun, Wells Santo, Jaclyn Sawyer, Matt Scherer, Ranjit Singh, Hongsup Shin, Wineke Sloos, Robert Soden, Matthias Spielkamp, Luke Stark, Eric Sublett, Mayesha Tasnim, Tracey, Laura Weidinger, Michael Yang, Mike Zajko and Janet Zhen-Rong Gan

Canada's Social Sciences and Humanities Research Council for funding the research project, AI for the Rest of Us.

Image credit: Dan-Cristian Pădureț, [Unsplash](#).

Executive Summary

Our Critiquing and Rethinking Accountability, Fairness and Transparency interactive workshop at ACM's Fairness, Accountability and Transparency in Machine Learning on March 5, 2021, provided workshop attendees with information on what constitutes meaningful civic empowerment in artificial intelligence (AI) systems' development and deployment. The workshop offered a primer on participatory theory from the social sciences and had speakers from nonprofits discuss the role of tech-enabled community activism and consider political power dynamics in using tech to influence AI development and policy.

We used digital whiteboards and breakout sessions to have attendees reflect on what they heard and supply their own definitions of actors, roles, and empowerment, which was summarized by one attendee's provocative question, "who would be participating if we weren't having this intentional conversation about participation?" In terms of the 'who', individuals and groups directly impacted by algorithms should definitely be considered participants. After that, the attendees offered a plethora of potentially impacted peoples, which ranged from future generations, journalists, and the undocumented to non-human actors like environmental ecosystems. Often the reflection on who was implicated and who was harmed by AI was as important as an invitation to a predetermined or convenient stakeholder to participate. In terms of the 'what', we asked for roles and responsibilities of all the members involved in the development and deployment of an AI system. Many attendees described the role of individuals, groups, and organizations as providing feedback and inputs on the development process and its outcomes, on the participants' part as well as ensuring transparency, accountability and fairness from the AI developers.

Defining and operationalizing civic empowerment proved the most difficult for workshop attendees. Definitions often centred around self-empowerment as personal satisfaction and community quality of life. Empowerment often was circular: a willingness to participate more or greater representation by marginalized communities. When attached to political power, empowerment varied from providing feedback to giving participants the power to halt an algorithm or change laws and algorithmic decisions.

We sought to co-develop an assessment questionnaire that will allow AI developers, government, and tech-enabled civil society organizations (CSOs) to gauge the level of civic empowerment in their AI system(s) (see Section 4). We categorized the contributions from attendees in terms of existential, epistemic, process-based and a few ready-to-go assessment questions. Our hope was that this set of questions could eventually augment AI certification, audit, and risk assessment tools. At this stage, instead of generating a simple list of best practices, our assessment questionnaire serves to "open Pandora's box," explicating the diversity of positions, confronting the challenge in synthesis and implementation, and moving towards meaningful and not merely performative empowerment.

1. Introduction

The goals of this workshop was to challenge the standard of civic or public participation that undergirds algorithmic impact assessments (AIAs), audits, risk assessments, participatory design, and methods to engage the public-at-large (e.g., citizen juries). We provided attendees with an overview on participatory theory; they also heard from civil society organizations (CSOs). We then worked with attendees to co-create a new method of assessing meaningful participation (i.e., meaningful to civil society and impacted individuals).

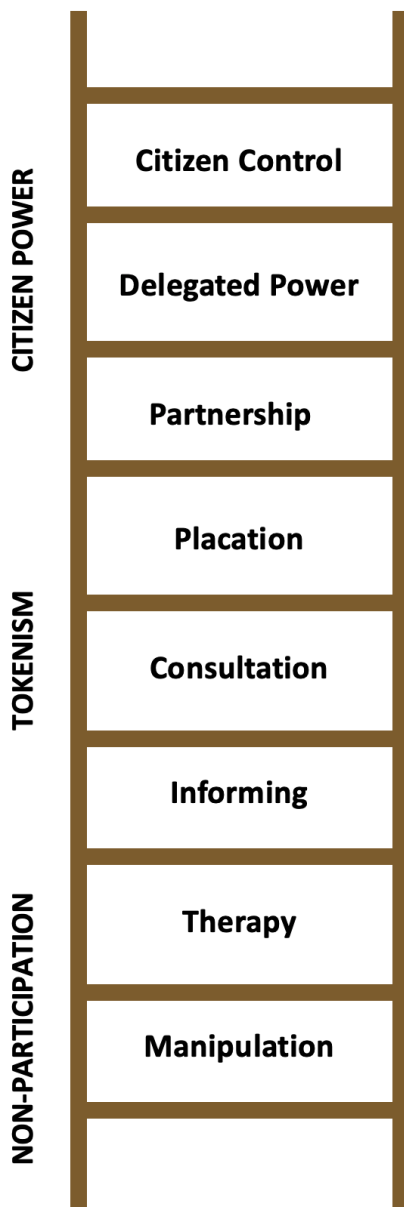
A large assumption underlying civic participation in artificial intelligence (AI) is opacity. Because AI/ML is labelled opaque and unknowable to the overwhelming percentage of the population, it can be rendered “beyond scrutiny” for any meaningful civic engagement. Current methods in AI ethics/responsible AI and within the FAccT/Machine Learning community have envisioned citizen participation as largely stakeholder engagement, which can be highly reductionist because it can substitute representation for influence. How do we move beyond stakeholder engagement or public provision of feedback? It is relatively easy to get a member of a CSO onto an AI oversight board. It is far more challenging for citizens to leverage political power to exert influence over AI development and deployment. For example, how do we actively engage civil society in influencing AI decisions and technically empower them to possibly even use AI to counter such policy? These are the questions we addressed in the workshop.

In addition to the above, the workshop used a digital whiteboard and guided breakout rooms, where participants had the ability to answer a series of questions and reflect on the discussions. A preliminary document of the workshop’s section on the Definitions and Characterizations of Participation and Empowerment was sent out to attendees in May, which is further refined and detailed in this document.

Our hope was that, by the end of the workshop, attendees would learn about the components of civic empowerment and actively include tech-enabled CSOs as AI systems are developed and deployed. And that attendees would share their expertise in co-creating an assessment questionnaire that would center civic empowerment in the development and deployment of AI systems.

2. Speaker presentations

To begin the workshop, we heard from four speakers, who provided perspectives on participatory theory, spoke on CSOs that can act as participation intermediaries, and advocated for positions on tech policy. Speakers came from academia, CSOs, and consultancies to the government on tech policy.



2. 1 Renee Sieber on a Brief & Deep Dive into Roots of Participatory Theory

Renee Sieber introduced the workshop with her presentation, “[A Brief & Deep Dive into Roots of Participatory Theory](#).” The presentation started by drawing on Polletta’s concern that, in the general state of civic participation, the people speak but they have no control over the final decisions that are made. Indeed, critics argue, grand spectacles of public participation may make it that much easier for back room decision-making to carry on as usual, unscrutinized and unchallenged. Such spectacles provide a satisfying experience of democracy in the moment, but they are no substitute for mechanisms of democratic accountability.¹

Polletta’s critique mirrors concerns in the AI/ML community of “participation-washing”.² Sieber described links between participation and democratic theory and how they do not always produce a participatory democracy. Greene, Hoffmann and Stark examine value statements like ones made in the Montreal Declaration for the Development of Responsible AI: “Fixes are best addressed through technical and design expertise. There is little sense from these documents that AI/ML can be limited or constrained.”³ Instead of a participatory democracy, many in government and the private sector may advocate for a technocracy, where the techies and the algorithms manage democracy and infer what the public wants.

Sieber then discussed the challenge in defining participation, since almost any interaction with the state or with firms could be labelled participatory. Cornwall calls civic participation, “an infinitely malleable concept..., it can easily be reframed to meet almost any demand made of it.”⁴ Who should be participating also is challenging, not only because there’s concern about whether the AI developer/deployer has chosen the right person or

¹ Polletta, F. 2016. Participatory Enthusiasms: A Recent History of Citizen Engagement Initiatives. *Journal of Civil Society* 12(3), 231-246, p. 234.

² Sloane, M., E. Moss, O. Awomolo, & L. Forlano. 2020. Participation is not a Design Fix for Machine Learning. Proceedings of the 37th International Conference on Machine Learning, Vienna, Austria, PMLR 119.

³ Greene, D., A. L. Hoffmann, & L. Stark. 2019. Better, Nicer, Clearer, Fairer: A Critical Assessment of the Movement for Ethical Artificial Intelligence and Machine Learning. Proceedings of the 52nd Hawaii International Conference on System Sciences.

⁴ Cornwall, A. 2008. Unpacking ‘Participation’: Models, Meanings and Practices. *Community Development Journal* 43(3), 269-283, p. 269.

group of people. Are these targets of an AI tool? Journalists? Or other developers (in case a government has initiated an AIA)? It's also a unit of analysis question: should one invite a single individual to participate or a collective like an entire neighbourhood. The individual usually stands in as a representative, but for what or for whom? Individuals could represent an identity-based group, a self-appointed community leader, or a representative accepted by a CSO. This is supported by Carole Pateman's argument that a participatory democracy is best served by self-appointed representatives of established CSOs.⁵

“stakeholder engagement processes flatten the power relations of participants and rob civil society of its power base”

Morley et al. reviewed 84 Applied AI Ethics tools and found they emphasized engagement by an individual over the collective.⁶ Participation operates within scale and efficiency constraints: “it is practically impossible to expect that participation can be expressed by all, everywhere, on all matters and constantly.”⁷ Hence the push towards inviting individuals who represent an identity group on very targeted and narrow issues (e.g., Black and Indigenous People of Colour considered by a government to be adversely impacted by predictive policing). Because of efficiency considerations, participation by individuals serves a market ideology as the best expression of morality/outcomes, with the public's input often reduced to consumer choice (e.g., ‘boycott this app’, ‘don't order from this online company for one week’). Another driver of the individual participant, for instance through random selections of members of the public, is that it supposedly generates a more authentic and disinterested participation.⁸ A self-appointed community leader could be perceived as too professionalized or as an agitator. Often, we see a preference in participation related to AI, for a “calculated public,” “the process of algorithmic presentation of publics back to themselves, and how this shapes a public's sense of itself.”⁹ This, Sieber argues, represents a reinvention of community organizations, even though there already are community organizations who could send representatives.

⁵ Pateman, C. 2012. Participatory democracy revisited. *Perspectives on Politics* 10(1), 7-19.

⁶ Morley, J., L. Floridi, L. Kinsey, A. Elhalal. 2019. A Typology of AI Ethics Tools, Methods and Research to Translate Principles into Practices. Presentation at AI for Social Good workshop at NeurIPS, Vancouver, Canada.

⁷ Floridia, A. 2013. Participatory Democracy versus Deliberative Democracy: Elements for a Possible Theoretical Genealogy. Two Histories, Some Intersections. 7th ECPR General Conference Sciences Po, Bordeaux, 4-7 September, p. 12.

⁸ Polletta, 2016.

⁹ Gillespie, T. 2013. The Relevance of Algorithms. In *Media Technologies: Essays on Communication, Materiality, and Society*. Tarleton Gillespie, Pablo Boczkowski, and Kirsten Foot (Ed.). Cambridge, MA: MIT Press, p. 3.

Sieber discussed the practice of civic participation. She does not advocate for a stakeholder engagement process, the predominant method in responsible AI initiatives, because it flattens the power relations of participants. Representatives do not have a power base, a constituency upon which to draw. At minimum, resource accessibility is key: “Public participants should have access to the appropriate resources to enable them to successfully fulfill their brief.”¹⁰ Many extant processes can ignore the roles of the various actors, such as the sponsor/host, who can prioritize expertise (e.g., process becomes a digital literacy project since the goal is awareness and not necessarily influence) and constrain terms of reference (e.g., acceptable topics, tone of discourse, level of influence). Neither should the facilitator’s role and agenda be ignored. Public engagement in many places has become professionalized, with a “respected facilitator—such as a newscaster.”¹¹ Professionalization supposedly produces a neutral process, where process is emphasized over outcomes. Participants may have different metrics of success as compared to facilitators or sponsors. The former may expect the process to result in policy change; for the latter two upholding a code of conduct during a meeting may mean success or ensuring the output does not “reduce the control and influence of the sponsoring organization”.¹² For sponsors and facilitators, civic engagement in AI can become a kind of containment of citizen voices or a performative activity, which limits the scope of that engagement while still being able to check off a box.

Whether large or small, inviting someone, usually a non-expert, into the design and development of the system cedes some control over that process. The process of participation should therefore be linked to that measure of control, which is usually referred to as empowerment. Like participation, empowerment is a vague term, variously defined as internal (self-empowerment) or external (political power). Zimmerman offered an all encompassing definition:

self-actualization, an internal state of enabling an individual's agency to reach their self-defined goals) or the capacity of individuals, groups and/or communities to take control of their circumstances, exercise power and achieve their own goals, and the process by which, individually and collectively, they are able to help themselves and others to maximize the quality of their lives.¹³

Zimmerman argued that empowerment overlaps with power but the two are not identical. Often there is too much emphasis on the former (awareness, literacy, i.e., an “an individual-level construct”) and too little on the latter. Sieber ended her talk by describing

¹⁰ Pateman, 2012.

¹¹ Rowe, G., & L. J. Frewer. 2000. Public Participation Methods: A Framework for Evaluation. *Science, Technology and Human Values* 25(1), 3-29.

¹² Rowe & Frewer. 2000.

¹³ Zimmerman, M.A. 2000. Empowerment Theory: Psychological, Organizational and Community Levels of Analysis. In J. Rappaport & E. Seidman (eds). *Handbook of Community Psychology*. Springer, 43-63.

Sherry Arnstein's ladder of citizen participation (see Figure above).¹⁴ Arnstein's model, more than 50 years old, is the most widely used model in the world and many models are built atop this ladder.¹⁵ It is notable because it explicitly links political power to public participation, with the highest level of engagement being citizen control. Arnstein also laid out in great detail how lower rungs of the ladder actually constituted forms of fake participation. Sieber asked the workshop attendees to consider the implications of Arnstein's ladder for the extant models of civic engagement in AI development and deployment.

In the Q&A, attendees had the following questions for Sieber:

Is participation always better for participants? Actually, it can often be used to benefit the sponsors, private firms, big tech and government. For instance, participation can act as a form of exhaustion, a 'vent your frustrations' exercise that results in no policy change.

Have scholars looked at non-human, non-algorithmic participants in civic participation? In my participatory GIS (geographic information systems) research, I've seen it. For AI? Not to my knowledge.

Is there work specifically on participatory governance of software? Depending on your definition of governance and who's a part of it, governance is participatory by design. We need to distinguish civic empowerment from governance models since governance doesn't necessarily include differentials in political power and resources.

The FOSS (Free and Open Source Software) or FLOSS (Free/Libre and Open Source Software) movement is a kind of participatory governance of software. Github is a good example of a participatory governance platform for software. These represent a narrow definition of participation as data sharing however.

What research has been done evaluating the impact of different models of participation? There really isn't research in AI ethics that sufficiently evaluates AI impacts. For instance, participation is very hard to operationalize. There is a saying in participation practice, "bums in seats", which refers to counting the number of people attending a meeting or on Zoom calls. The count is used because it's quantitative and easy to measure. That's different from evaluating the outcome. Outcome measures, like public influence on decision making, are very difficult because of the qualitative nature and the longitudinally. We have so few examples of public influence in relation to AI. The most prominent is companies' and communities' moratoria on facial recognition technology, although the details suggest less of a stop than implied. People point to the New York City ADS Task Force, which established

¹⁴ Arnstein, S. 1969. A Ladder of Citizen Participation. *Journal of the American Planning Association* 35(4), 216-224.

¹⁵ e.g., Fung, A. 2006. Varieties of Participation in Complex Governance. *Public Administration Review* 66, 66-75.

actionable items, but the 18 month process didn't mean there was transparency in NYC's use of algorithms or that any recommendations were enacted.¹⁶

2.2 Mich Spieler on Community Activism in Tech and AI

The workshop also included presentations on the role of CSOs and community activists in enabling participation vis-à-vis tech policy and tech capacity in communities.

Mich Spieler works at the intersection of technology and social justice. They spoke about their work as a Community Technology Co-coordinator at COCo ([Centre for Community Organizations](#)), a nonprofit supporting the health and well-being of community organizations in Québec. COCo consults with nonprofits that have a social justice mandate by augmenting capacity, for example in organizational health, anti-oppression work, and also some parts of technology (e.g., customer relation management systems, online collaboration tools and more recently, data management). Spieler points out that debates around civic participation have been ongoing for as long as the term exists. Some consider sending a tweet as a form of political participation. For others, only voting and holding office should count as political participation. What you count, what you don't, and what you want to happen depends on where individuals and groups are speaking from. Spieler's presentation is rooted in Quebec's autonomous community sector. The autonomous community sector is not non-governmental organizations, per se, but organizations with grassroots perspectives that have a mission of social transformation. Approximately 20 years ago, the government issued a policy on autonomous community action with dedicated funding that recognizes the unique role this sector plays in solving social problems and accords legitimacy to the sector as a result.

Spieler reminded us that we tend to think of AI as software, but AI also has a physical presence. We see the arrival of tech workers and companies with higher than average salaries. Development and rent prices go up that mainly impact racialized communities who get pushed out of their traditional areas. Montreal has experienced considerable gentrification due to the growth of the AI sector.

In their work, Spieler found many tech workers who wanted to do good and offer their services as volunteers for CSOs. This aligns with the increasing number of data/tech/AI for good initiatives in the area that offer technology tools to CSOs, promising that AI will solve their problems. There are issues with this type of 'tech philanthropy'. The time engagement on the tech participants side is fleeting yet demanding on the community groups. One-weekend hackathons demand a lot from participants, leaving them with 'proofs-of-concept' but no concrete path of following up. Another issue they saw is the 'move fast break things' approach, which is not all compatible with a general CSO culture of individual and collective care aimed at transforming living conditions and respecting rights

¹⁶ Lecher, C. 2019. NYC's algorithm task force was 'a waste,' member says. The Verge, November 11, 2019. <https://www.theverge.com/2019/11/20/20974379/nyc-algorithm-task-force-report-de-blasio>

and livelihoods. What Spieler has witnessed is an ongoing pattern of more interest in technical solutions than in social issues that need solving. These are technical solutions looking for problems to solve, and not the converse. CSOs are asked to change how they work, to make technical solutions work. But there is no equal openness, no reciprocity in knowledge and data sharing.

“tech philanthropy has a literacy problem”

Spieler talked about philanthropy’s “literacy problem”. This unfolds in two ways. First, this occurred in representatives from the tech sector to help CSOs address their problems. For example, which CRM (Customer Relationship Management system) would be a good fit for them and how would it be implemented? We cannot rely on vendors to equip community organizations with the technical knowledge they need to adequately understand CRMs. COCo’s tech program exists because of this type of vendor-customer gap and also the fact that consultants usually are business focused and don’t understand community organization realities, because of this type of vendor-customer gap. There’s little recognition that this gap exists in AI, especially among the tech workers and the philanthropists in these AI for good initiatives. If the recognition does exist then it manifests as a lack of knowledge on the "customer" side that these initiatives often seek to bridge or, cynically, make disappear.

Second, this can set up a conflict in the philanthropy sector’s support for the tech sector helping CSOs. Spieler brought forth the example of a collaboration set up by a philanthropist organization between CSOs and representatives from the tech sector who were proposing to combine various datasets to draw conclusions about CSOs impact. There were all kinds of problems with such a suggestion because of what the data was actually measuring versus what some people thought the data could answer. Ultimately the tech reps were introducing an economic demand logic to a complex social problem, an approach that the CSOs were not looking for nor what they would find useful.

Capacity building in community organizations around tech is needed, which requires time and money. It is important to remember that numerous community organizations are still in ‘the digital turn’, moving away from paper reports to digitized records. In some instances, participants may have no computers or an email address. There is as great a need for capacity building in the tech philanthropy sector (hence, a dual literacy problem for tech philanthropy). People from the tech sector need to become more literate in CSOs’ theories of change and their approaches in operating on the ground.

Spieler concluded that many participation models originated at a time when the nation state was the clearly delineated locus where political decisions were made. “We need to look deeper than just the surface distinction between government, private companies and NGOs.” Spieler emphasized that policy and governance discussions often only consider only two options: private sector or government. “Instead we should look to approaches that leave

room for different options that incorporate ideas like the Commons or look to other governance models for inspiration. For example, we could look at participation in rural cooperatives in the societies in the global South or traditional governance models practiced by Indigenous peoples.” Reinforcing Sieber’s point, in conversations communities often are reduced to stakeholder status. But they are so much more than that. “Communities are subjects, actors and protagonists and we need to respect them as such. We need to talk about knowledge without falling into an expert discourse. When we talk about AI and empowerment, we also must talk about AI and disempowerment. We must ensure that tech, in particular AI, does not lead to disempowerment in terms of decision making, as political decisions get taken out of the political realm and transferred in the technological and commercial realm.”

2.3 Bianca Wylie on Community Activism in Tech and AI

Bianca Wylie is an open government advocate with a dual background in technology and public engagement. In her talk she posed the question: “How does political accountability map to technical and AI accountability?” She drew on examples from her work in open data, civic technology, the tech sector, advocacy and tech policy. “Political accountability implicates democracy, which has so many different meanings. As an example of political accountability in Canada, we’re going through a truth and reconciliation process. This draws on treaties going back hundreds of years, that bolster holding each other accountable in relationships via contracts. Can we talk about modern accountability if we are failing to uphold the oldest version of accountability in colonial Canada?”

“Political accountability is mapped onto jurisdiction and place and falls apart where it concerns technology governance. We can see that in Toronto’s failed smart city project. In 2017, Sidewalk Labs entered into partnership with several government entities to set up an innovative city on the waterfront neighbourhood.” Wylie and many other people actively pushed back on the project and found lessons for AI. Not only was Sidewalk responsible for enacting the smart city, it was in charge of facilitating the public consultation. This shifted public consultation, which contained no public accountability, onto a corporate actor that had no accountability to residents. She argued that no one in Toronto consented to outsourcing this kind of activity, which was legal then and still would be considered legal now.

Wylie argued that this shift represents a slippage in governance that can happen during the procurement process. How decisions are made about public services (e.g., energy, water, waste, transportation) requires accountability; it’s what distinguishes the public from private sectors. When you shift those decisions, for example with AI, you can make it worse for the public to retain that accountability. If decisions made about public resources have shifted towards the private realm, residents can no longer hold governments accountable for how public resources are distributed. That results in a privatization of governance, in which the

government outsources itself and we don't even know if the government knows it is outsourcing itself. Privacy and privatization are fundamentally different issues. Wylie argued that privacy is a sub-genre of the privatization of governance. When we privatize accountability out of the state, privacy can fall victim. Once this governance goes private it is very difficult to get it back.

The public doesn't know that this is the problem either. Democracy relies on journalists, yet the number of journalists who understand the problem is low; for instance, they work for editorial boards that do not allow them to write critically about AI or governance and the associated political power. We may have great journalists in North America, but we also have status quo mediocrity and white supremacists. So editorial boards may have very low interest in getting into these power conversations. What they do talk about is the magic of AI. We saw that in the Toronto project. Ultimately Sidewalk Toronto didn't build anything; they abandoned the project. However, what they did build was an incredible narrative: beautiful renderings, an economic development plan, a beautiful story. The journalist and government investment in that story was the headline, not any construction or critique.

“when we privatize accountability out of the state, privacy can fall victim”

She argues that we're at litigation time again. Journalists cannot scare the politicians to do the right thing. Politicians aren't doing the right thing because AI, especially in Canada with its heavy investments, appears good for economic development. Every single point on this democratic circle is captured. So Wylie argues it's lawsuit time. Laws, especially common law, are supposed to be elastic and adaptive so you have to bring pressure to the courts. Strategic litigation is fundamentally becoming the preeminent tool. Foxglove Legal is one example. But who has the money and the time? Who are the legal people willing to get into this when the narrative is what it is?

Wylie's central question was: How do we bring the problem of accountability to the surface?

During her talk, Wylie recommended three books, which can be found in Appendix A.

In the Q&A, attendees had the following questions about Spieler's and Wylie's talks:

Do AI researchers have the responsibility to raise public awareness about technology issues of the tech they are creating?

Spieler: Doesn't everybody have that responsibility? Shouldn't AI researchers be people living in a community? You can't extract yourself from the community in general. I would start thinking about that: how your work contributes to the wellbeing of your community. Take yourself off of a pedestal.

Wylie: Everybody has a choice to make. If it makes sense to you then please do it. There are various ways to interact, like visiting a political office whether it's city, provincial, federal, state, however you are organized. In political offices there is more negligence than malice right now. These officials do not understand what is going on. In some cases there are councillors that I have worked with who want to come in on a problem but their staff and themselves do not feel foundationally secure enough to argue a point. There is an opportune moment to critically engage. Everybody has to pick how they want to raise awareness. So many people who do the most important work--you would never know who they are and how they do it. We always need to be mindful of visibility and also the positionality of our power and which version of it makes sense.

What kinds of organizing frameworks do you think could empower civic journalism/participation that gathers artifacts? Artifacts that could serve as evidence that is going to help policymakers and regulators.

Wylie: One thing is requesting Freedom of Information Access (FOIA; in Canada, Access to Information Act requests or ATIA). People are wildly unaware they are in a surveillance society. The things that people write in an email sometimes.... Journalists should try to really collaborate on FOIA and share how to do it. Because competitive journalism means that people do not like to share their source. By the time they share, the FOIA is not germane to the moment.

Spieler: Support community based journalism and help it take off in places it hasn't taken off. Community based journalism doesn't sound as professional--the equipment isn't as good--but it's important. Support local radios, newspapers, and journalism that bring out information in a different way. We've seen it in social movements. For example, the women's movement in Quebec is still focused on big media but it is proving to be less and less true. How do we not depend on big company journalism?

Wylie: I learned from the years of starting a conversation that we should not start with data, technology or AI. It is not a good narrative or a good entry point. It has to be wrapped in issues. No one feels confident or comfortable when it's data, tech or AI. Talk about privatization instead of privacy. Seda Gürses talks on weaponization of privacy: She speaks about Google and Apple influencing public health policy globally saying "We will protect your privacy better than the state" even though this will affect the impact and efficacy of what we are trying to do.¹⁷ Privacy to me is already being weaponized as a narrative against the state.

How do you engage participation effectively to bring about change? Is there anything you've written to help us employ? To bring people together around these issues effectively? I work with a legal empowerment foundation in the UK funding grassroots organizations to use strategic litigation to challenge government procurement, use of AI.

¹⁷ <https://www.radicalai.org/e5-seda-gurses>

Wylie shared her writing on the tactic of asking questions: Teach people to ask questions. You don't have to have the answer. It's a critical part of civic tech journalism. What are you doing about this or that? With Sidewalk Labs, we opened a document and generated hundreds of questions that the community crowdsourced together. We made a confident space to ask questions. Point out the absence of any of the protective mechanisms (e.g., municipal regulations) or how they are completely falling apart. You can have all the rights in the world but if you can't access them then they're not useful. We can build on a lot of access to justice and access to legal protection work. When you are in the middle of it, you see all the diversity of tactics. We also need to be unlikely bedfellows in this. We need to jam it and slow it down.

If technological know-how is concentrated in private industry, is there a way to reconcile civic accountability with the deployment of technological solutions that can only be privatized or is privatization fundamentally incompatible with accountability?

Spieler: What is our understanding of private and public? Or private and not private? There is private industry but there also are concepts where the state has a role and private industry has a role. We have to decide who will be responsible and for what. On accountability, from a community standpoint, we are at the point where we don't know who we would talk to. How many steps does it take to reach the person that makes the decision? As Wylie explained, if accountability is gone, we also lose transparency.

2.4 Jean-Noé Landry on How a Nonprofit Builds Public Participation and Empowerment in Other Organizations

Jean-Noé Landry, social entrepreneur and former executive director of Open North presented on how a nonprofit builds public participation and empowerment in other organizations (see [slides](#)). Founded in 2011, Open North is a Montréal-based non-profit organization with roots in civic technology development. Over the years, it's evolved into conducting applied research and engaging in capacity building and advisory service, all to advance the responsible and effective use of data and technology. He is a part of the Canada Mortgage and Housing Corporation's Housing Supply Challenge citizen jury. He also helped establish the Canadian Multi-stakeholder Forum on Open Government.

Landry has two degrees in philosophy so he recommends Machiavelli's *The Prince* as a way to think about power and influence. Before doing tech nonprofit work he worked in the political development spaces, which included pro-democracy movement and civic networks in different parts of the world.

Landry considered how to structure projects to ensure engagement and empowerment. Some of Open North's key concerns are related to strategy and approaches. As we engage with different stakeholders across the board, who are the key players? Those with a stake and those with/without political power in the conversation we are involved with. What do

stakeholders influence? Is it policy, strategy or operations? Who stands to gain? Who are the champions and how do we make allies of them? Who is at risk? Who is potentially going to be exposed? Put on the spot? What existing tensions underlie inter- and intra-organizational relations that will influence the conversation? How do we create an open, welcoming environment for people to engage? We need to plan for group dynamics, venue, participant numbers. How do we ensure that core issues/problems will still be heard? What is the group's composition? How do we set expectations? We need clarity in expectations of participation and outcomes.

He draws on examples that Open North has facilitated: City of Toronto Open Data Master Plan, Ville de Montréal (VdM) Plan d'action données ouvertes, and VdM's smart cities challenge public consultation. In those consultations, we need to ensure the right voices are heard. Through focus groups and group dynamics we try to get broad participation. We make that mandatory if necessary. We triangulate, by looking for the target participation at multiple levels of hierarchy from the same unit and across multiple units. We navigate organizational hierarchies and power structures. The focus on inclusion requires we listen across the entire spectrum, from analysts, to managers, to elected officials. Our models focus on feedback and opinions; we capture those as qualitative data. It's not just a listening exercise, it's also a data collection exercise. We need strong, qualitative data collection methods. Data capture in a large group is difficult so it is important to try to form smaller groups.

How do you feel about the professionalization of participation? And potentially advocating positions arising from these forums? You already take normative stances (e.g., you want transparent open government data--that's prescriptive). Do you worry that you will be criticized or will be labelled as no longer objective?

Landry: We should ensure inclusion and/or empowerment of stakeholders, including the representation of voices across stakeholders. Focus on feedback loops: make sure people know they have been listened to, that their feedback has been read and analyzed. Data and a report are often the end goal for us.

Sieber finished this part of the workshop with a few examples of how civic participation is evaluated in a responsible AI risk assessment tool that she had some small hand in crafting. One of the most explicit references is to whether or not the AI developer or deployer solicited public input, in the form of feedback from user surveys. Indeed, "participants" are generally the *end users* or targets of a tool. Another reference to civic participation is a governance board. Sieber reminded the attendees that representation on a governance board doesn't equal empowerment. Neither does it account for the duration or timing of this governance/oversight (e.g., long after the AI tool has been completed).

Many AIAs focus on data privacy. It is important to account for privacy protection in AI tool development but privacy is frequently framed as an individual concern. A counterexample in

Canada is the way that First Nations groups conceive of privacy and control: as a collective concern (see OCAP™--ownership, control, access, and possession).¹⁸ There frequently is no opportunity for participants to effectively challenge certain kinds of data being collected nor may there be accountability mechanisms in place, should a government/firm fail to comply with its privacy protection.

There also is a focus on transparency and explainability in AIAs. A lack of transparency is not necessarily the issue; it's that public participation can be reduced to explanation. A review of 84 Applied AI Ethics tools found “an over-reliance on ‘explicability,’” which aligns well with a computational approach to conceiving of responsibility and is overwhelmingly applied at the end stage or “post-hoc explanations”.¹⁹ According to Arnstein, this is civic participation as informing, a low rung on the ladder and an instance of non-participation.

Last, accountability: Many AIAs focus on internal accountability and self-reporting (e.g., “Were all labour laws followed?”). This is a possible place to expand an assessment, for instance, to include more explicit support for formation of unions and protection of whistleblowers. Protecting whistleblowers also aids transparency efforts.

3. Discussions on how key terms are defined

This section focuses on discussions in the four breakout rooms during the workshop. Each subsection represents a question we asked the attendees. See Appendix B for the raw data produced by attendees on the digital white boards (Jamboards).

3.1 If we're talking about citizen participation & civic empowerment, who could be participating/empowered?

It was quite remarkable how challenging it was to answer a question like this, which could be dismissed as “Of course we know who should be invited.” Discussions centred around trying to find common definitions for ‘citizen’, ‘community’, and ‘people impacted by the system’. There were concerns raised by attendees about how these terms and their descriptions could be exclusionary, especially if some people who are equally impacted by the system fall outside the description bracket or jurisdiction. Underrepresented and historically disenfranchised communities were mentioned as important participants by several attendees. Others commented that the definition of citizen should be broad enough to include all ages and future generations as well as groups like formerly incarcerated, undocumented and stateless people.

Some cases might fall outside even a broad definition of citizen, or citizens’ group, such as companies that fund political candidates or industry commentators on government regulation. Do we talk about individuals or groups that can be affected or do we invite those

¹⁸ <https://fnigc.ca/ocap-training/>

¹⁹ Morley et al. 2019, p. 3.

individuals or groups? It was interesting that the term ‘stakeholder’ never appeared on the digital whiteboard.

The underrepresented could include workers from industry who were non-engineers as they are often cast as invisible labour (e.g., cafeteria workers). The rise of tech unionism at the same time as the state of Nevada pushing legislation to allow tech/innovation companies to create their own governments is consistent with Wylie’s point on privatization of public resources. Of the impacted individuals and groups, unions may be where the power is right now.

Additionally, it was fascinating to see comments related to land, water and other environmental factors mentioned in relation to communities. Several related questions were raised: Do we need to have domain-based categories (e.g., climate change, medicine)? Should non-humans be included, and if so, how?

Should participants be determined by jurisdiction, that is, the physical community? Indeed, what is the role of geography or physicality--the spatial logic--in participation, as much of public participation is jurisdictional. There are also non-spatial forms of participation that go beyond physical connection. We no longer fully exist in a physical space, where physical borders are not clearly defined. People outside the jurisdictions are still implicated. Most regulations are not global but are jurisdictionally bound. However, policies within the global North can affect people outside their jurisdiction. Consider the role of the US Patriot Act in other countries. Even if Canada had stricter privacy measures than the US, Canadians can be subjected to the US Patriot Act (national privacy regulations are even less protected under the United States-Mexico-Canada Agreement). What are these extra-jurisdictional impacts? Are there differences between physical persons and digital persons (smart city research on digital twins may assist with this)?²⁰

Insufficient work is being done on communities of practices, defined as individuals organized or identified around an activity or interest. Communities of practice address the non-geographic nature of tech. However, communities of practice usually align with the ‘citizen as a consumer’, which is very much a market-based framing where participation centers on user satisfaction vis-à-vis service delivery and can place the private sector as the best place for AI governance. This was seen by some as the most effective site for activism. Some argued that in a globalized, digital space of tech, AI producers should convene these meetings and not governments.

There was interest in AI developers being part of the participation and in their being able to see themselves as active participants, rather than passive actors. Developers may have

²⁰ Dembski, Fabian, Uwe Wössner, Mike Letzgus, Michael Ruddat, and Claudia Yamu. 2020. Urban Digital Twins for Smart Cities and Citizens: The Case Study of Herrenberg, Germany. *Sustainability* 12, 2307, doi:10.3390/su12062307.

multiple identities, for instance, being a citizen or member of the community. It's important to help find ways to help them position themselves in that way.

The discussion suggested that we may have to accommodate numerous deficits even before we begin assessing AI, like a lack of access to government services, education or the resources to volunteer for an impact assessment. The irony is that individuals/groups with the least access to these services (and knowledges) may be the most vulnerable to AI systems.

Last, a point of reflection. One attendee suggested that the term 'participation' could be corrective to some pre-existing baseline of how engagement should take place, asking "who would be participating if we weren't having this intentional conversation about participation?" Without reflection, it is simple to tailor participants to those who we might anticipate impacts or participants that are convenient to a particular AI developer.

3.2 What role do individuals, groups, organizations play & what responsibility do they have in the development and deployment of AI systems?

Many attendees described the role of individuals, groups, and organizations as providing feedback and inputs on the development process and its outcomes, as well as ensuring transparency, accountability and fairness. Comments implied a need to clarify the meaning and intent of these words. Attendees also discussed the nuances of distinguishing the responsibilities of individuals, groups and organizations and the possibility of conflicting roles, as well as how those would be managed. Most roles mentioned focused on individuals, which raises concerns about how to evenly distribute roles to ensure that we are not overburdening/overwhelming the people who are also being disproportionately impacted by these systems, as well as ensuring they have the necessary resources (e.g., education on AI tools and processes) to help them execute their responsibilities.

Because it was difficult to tell whether participants could be those impacted by an algorithm, those developing the algorithm, or those deploying the algorithm (e.g., in the public sector/government), it was challenging for workshop attendees to assign responsibility for certain parts of the development, deployment or assessment process.

Recall that we are looking for responsibilities of all parties, including conveners. Notable responsibilities included ensuring representation of the diversity of participants, determining environmental impacts, and ensuring the fair distribution of influence over decision making throughout development lifecycles. These raise important questions, such as: How much power individuals, groups, organizations should have? Who defines and regulates these roles and responsibilities? What happens when roles/responsibilities overlap among groups as well as what happens when roles are relinquished? How much power do developers have to exercise ethical responsibility if it comes into conflict with their professional responsibility or requirements? Or, if there's a legal responsibility, at what stage

in the process does that come in? Who decides that? An attendee observed that they “like the idea of people having roles and responsibilities, but in practice the power structures dictate who gets to have what role. That's a challenge.”

For attendees, timing is critical. Having participants at the beginning to define the goals is important for responsibility in the process. An attendee provided an example of a police force in England that wanted to build AI to predict violent crime. “This is not Minority Report. It's not sensible or possible to do that. What are we doing this for? How useful is it going to be? We need to go through the whole process and really evaluate it; now that we've built it, is it really working in the way it is supposed to work? Sometimes it just gets built and not properly evaluated. How do we do that?” Another attendee stressed the need for continual assessment for bias.

One attendee gave an example from New York to highlight that “We start with the tech and then talk about the applications in terms of how government and law enforcement work. But we don't start with the problem and then figure out whether or not tech is the right solution and if there are better solutions.” Starting with the problem and investing in non tech solutions, instead of investing in surveillance cameras and drones, could have the additional benefit of improving community trust. There also could be a responsibility for AI developers and deployers to be vulnerable. In Hackney, London police did a stop and search of two young black men. They did a social media outreach and proposed community town halls to discuss what they did wrong and how they can improve.

One attendee suggested that, when we incorporate human rights and liberties of people in our AI considerations, we should take into account non-Western frameworks, particularly where they are adopted by non-Western communities that are being affected. For organizations that are setting up participation frameworks, we need mechanisms to help participants navigate power imbalances and access benefits of participation.

3.3 What words do you use to define civic empowerment and how would you “operationalize” those?

Unlike participation, characterizations of empowerment proved far more difficult and reaffirmed our need to continue. As one attendee asked, “what do we need to even know or understand to define civic empowerment?” That would be a kind of citizen empowerment: being able to ask the right questions.

Workshop attendees' definitions of civic empowerment often centred around empowerment as personal satisfaction and individual and community quality of life. There was a circularity to many attendees' characterizations. Empowerment could be considered as an increased future willingness to participate, greater representation, particularly by marginalized communities, a willingness to participate in the future or a broadening of participation.

Others reflected back on the originating procedural goals of FAccT; it was hoped these desires on the part of developers and deployers would ensure empowerment. This included a commitment to transparency and also called for AI developers or third parties to gain an accurate assessment of participants' concerns. Developers should ensure fairness and create an iterative process in which participants could revisit prior decisions. However, who gets to decide what constitutes empowerment?

The extent of political empowerment varied, from the provision of feedback to giving participants the power to alter or halt an algorithm. Empowerment might encompass the ability to refuse to abide by the outcomes of an algorithm as well as the results of an algorithmic impact assessment (AIA), the ability to interrogate past policies and laws, or the power to change regulations. Some attendees reflected on Arnstein's ladder of participation, which explicitly connects participation to political power as certain characterizations (e.g., citizen control of algorithm development) corresponded with higher levels of empowerment.

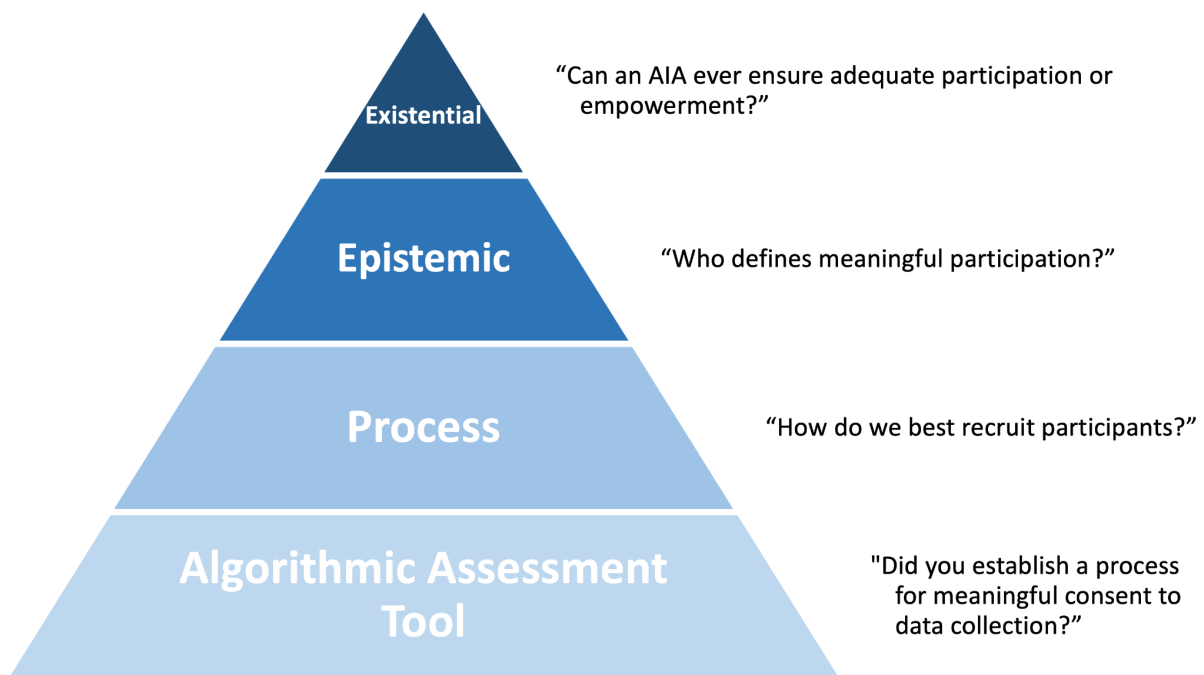
Some talked about informed consent as a way to empower civic participation. Another attendee commented that this could be challenging to implement. For example, what if people wanted to opt out of being included in language models? Would that even be possible? In addition, what impact does opt out have on someone's ability to use a platform - usually it would mean not being able to use it. Opt out strategies are highly divided by class, such as with the European Union's General Data Protection Regulation (GDPR). Often, the rich are more able to be forgotten than the poor or marginalized.

Many comments echoed this challenge in operationalizing empowerment. Some comments significantly challenged existing AIAs, what should be contained in them and what should be the outcomes of an AIA. One attendee noted that Canada's AIA, which was developed by the Treasury Board of Canada Secretariat, indicates the Canadian context. The AIA is supposed to be implemented at the start of system deployment or before the algorithm is put into production. Ensuring consistent engagement and granting the participants any influence becomes challenging, whether dealing with a third party or government bureaucracy. Governments are not monoliths: they contain multiple departments and agencies, with different levels of bureaucracy for AIAs and compliance depends on senior executive buy-ins. AI firms also may require multiple stages of approval, for instance through privacy people and technical people. It can take considerably long and make people averse to an extra process. These processes are meant to serve as a check to slow things down but time doesn't always fit well with the rush to innovation. How do we prevent people from ignoring these lengthy processes?

One attendee gave an example of rewriting the algorithmic rules in the US for allocating donors to transplant recipients. It was a multi-stakeholder forum that took 10 years. There were conflicts where balancing different ethical considerations. There appeared little chance of finding a consensus because every rule change would lead to someone's death. But in this process, people's hard edges of their original views were gradually worn off. Participants

came closer to shared understanding and that participation shaped the final opinions. This consensus building is a frequent outcome of participation, especially lengthy engagement. In its nefarious form, it's called co-optation. Note that consensus is a benefit to the government or firm, or in this case presumably society at large. This is not necessarily a benefit to the participants or the communities they may represent.

Attendees also commented on the jurisdictional and geopolitical challenges of building and enforcing laws. Tech might be developed in Country A and deployed in Country B. Country A's population might be empowered to influence an AI system's development, but not residents of Country B.



4. Assessment questionnaire

The goal of the workshop was to co-create a questionnaire to assess if civic participation and empowerment were occurring in the development and deployment of AI. This is similar to an audit/risk assessment tool, an inspection or examination of an AI system by an independent party. We found, however, considerable challenges in providing a simple tool to guide policy makers or AI firms in a single best practice. Rather, the questions served as a way for us to critically reflect on whether our current paths are suitable to the goals of civic participation and civic empowerment, both as defined by participants.

The assessment questionnaire was created based on workshop discussions (see Appendix C for raw data). We then categorized them into a pyramid of questions, some of which interrogate whether an assessment can ever ensure civic empowerment and some more prosaic how-to's. We have organized the questions as follows:

4.1 Existential questions

- Who decided that an AIA tool needed to exist? To what extent did people/groups impacted by an AI system want this?
- Can an AIA ever ensure adequate participation or empowerment?
- What normative assumptions are we building into an assessment? If we're using words like 'just' or 'risk' or 'responsible,' then we need to define terms. We can't assume the terms mean the same things to all people (e.g., in tech spaces). If we're using words like 'risk' then is it risk to the firm or risk to the people harmed?
- What structural inequalities intersect with the AI system and how are we accounting for the way the AI system may entrench or worsen those inequalities?
- If a particular AI system represents the latest manifestation of systemic bias, if it amplifies bias, then shouldn't we look past the AI, to more root causes? Should we instead be expanding rights (e.g., for workers), citizenship (e.g., for undocumented), resources (e.g., compensation) and regulations?
- Which forms of participation are valid and which are invalid (e.g., tokenism)?
- Which forms of empowerment are valid and which are invalid or potentially disempowering?
- Who should be able to review the decisions made by the AIA?
- Should we build a national or global AI accountability index (i.e., investigating the legal regime in which AI firms do their development, looking at history of compliance with other compacts, such as treaties)?
- What can we learn from alternate participation models (e.g., rural cooperatives in the global South) or governance models practiced by Indigenous peoples?
- How much responsibility do AI researchers and developers have to raise public awareness about technology issues of the tech they are creating?

4.2 Epistemic questions:

4.2.1 Defining publics and engagement

- Who
 - Who would be participating if we weren't having this intentional conversation about participation?
 - To what extent have we selected the easy participants (e.g., those with the resources to participate) and excluded the inconvenient ones (e.g., formerly incarcerated, undocumented and stateless people, uneducated or disruptive people)?
 - Should we focus on who is impacted? Who is missing? Should we be intergenerational? Include refugees? Journalists? AI developers? Industry lobbyists? Humans and non-humans (e.g., ecosystems)?
 - Should 'we' be making this decision on who should be included (e.g., as opposed to CSOs recommending participants)?

- What aggregations/types of participants are we considering (members of the public-at-large, representatives of marginalized groups, representatives of CSOs)?
- If we are engaging the public then whom have we defined as being the public? Is this a manufactured or “calculated public”?²¹
- How and What
 - How do we identify who should be included (e.g., terms/descriptions could be exclusionary if there may be some people who are equally impacted by the system but fall outside the description bracket/jurisdiction)?
 - Does the participatory process we choose match the participants (e.g., participatory design, trust surveys, use cases, oversight boards, stakeholder engagement)? Case in point: surveys encouraging trust in predictive policing aimed at ethnic minorities.
 - What is the influence of the individuals and groups on the other side of the participation ‘ladder’ (the AI developer, the government, the facilitator, any philanthropic organizations, nonprofits or CSO intermediaries)?
 - Is there undue influence on the participants (either positive or negative)? For example, might a participant lose a job, be shut out of an app, be subject to greater facial recognition technology? Conversely, might a participant get a better job or another form of compensation?
 - Given the diversity of participation activities and empowerment outcomes, who decides what constitutes meaningful participation and appropriate empowerment?
 - How much credence do we accord an AIA if it doesn’t provide sufficient empowerment to participants?

4.2.2 Defining and Enforcing Accountability

- How does one define accountability and who defines it? Is it internal or external (e.g., civilian oversight)?
- How do we balance participation with accountability as transparency?
 - Should we compel developers and deployers to report AI errors/failures?
 - What kinds of information (e.g., trade secrets, legal clauses, non-disclosure agreements) should still be included even if they’re considered by the organization to risk its intellectual property or competitive advantage?
- Mechanism
 - What accountability is in place (e.g., external sanctions, internal codes of practice)?
 - What is the chain of accountability? Multiple entities may be involved who bring forward different agendas, like a local government, a university

²¹ Gillespie, 2013.

researcher, or a vendor. How is the bureaucracy of government or layers of private sector governance navigated to ensure compliance? Does it require securing upper level buy-in?

- Who is assigned culpability?
- Who does the burden of proof of harm lie with (e.g., is it up to the harmed people to claim harm or is it up to the firm to prove that they don't harm)?
- What is the 'state' of legal mechanisms in a country or region? Is there a mechanism for statutory damages? Is there a possibility for a mass torts or a class action lawsuit? Can civil society prevent a merger, a mashup? Does a country or region support freedom of/access to information requests?
- Is due process followed? What is the due process for AI failures? For instance, what happens if AI system developers fail to comply with recommendations regarding participation and empowerment from a governing body's policy, such as an AIA? What if they refuse to fill out a mandated AIA?
- Are alternatives to legal routes warranted (e.g., Truth and Reconciliation Commission of Canada)?
- What is the accountability of the institution or organization sponsoring the process? What happens when the sponsorship changes (e.g., if government has outsourced participation to a public-private partnership)?

4.2.3 Defining and Ensuring Empowerment

- What exactly do participants influence (e.g., policy, strategy, operations, product development)?
- How much empowerment/power is granted to the participants (e.g., ability to contribute feedback or informing, ability to halt production)?
- Who decides that sufficient empowerment takes place? What time period or milestones are used?
- If empowerment is informing (e.g., raising awareness) or increasing literacy, what mechanisms have we used?
- How are power dynamics within or among groups of participants (e.g., representatives of various marginalized communities and industry representatives) made visible or mitigated? Many instances of stakeholder engagement are zero-sum, in which some stakeholders gain and some lose.
- How are changes to the level of civic empowerment evaluated?
- Once empowerment goals are determined, how are they put into practice?
- How does the empowerment strategy advocate for solidarity among (underrepresented) communities?
- What is the impact of the AI system on mental well-being (i.e., self- or individual empowerment)? Is personal satisfaction in the participation process experienced by the participant (e.g., civic duty, compensation, direct effect on policy)?

4.3 Process questions (once you've decided on the above...)

- Recruitment & retention
 - How are we recruiting for participation? What channels are we using to reach people, for an AIA or a survey?
 - If the participation method is representation-based then how are we ensuring the user base is representative of the population?
 - If the participation method is representation or stakeholder-based, then what proportions of each stakeholder group, including the public have been included?
 - Whose and which participation is the AI system designed to audit/assess (e.g., negatively impacted individuals or end users; the level of collaboration achieved or the degree of literacy gained)?
- Participation Design (not participatory design)
 - How well do the facilitators or the sponsors understand the AI system about which they're engaging participants?
 - What is the role of the facilitator? How independent is the facilitator from the developer/deployer (e.g., who hired them)? Will the facilitator advocate for the positions taken by the participants?
 - At what stage do participants participate in the development/ conceptualization of an AI system so that they can exert influence? For example, has the AI system already been publicized/done PR? Did the PR come before public participation in the project or after?
 - For how long are the participants involved? Do they only engage at the impact stage of an already deployed AI system?
 - What engagement method best suits the participants (e.g., citizen jury, participatory design)?
 - What activities did we do to ensure public participation (e.g., model cards)?
 - Who can audit the AI system in detail?
 - Can participants expand their role? Can more people be invited?
 - How do people who don't have access to or can afford the Internet participate?
 - What support have sponsors provided so participants can sufficiently engage in assessing the AI system (e.g., compensation, cf., industry association who pays its participants)? How much capacity has been built in already that reaches beyond literacy?
 - Can participants change the rules? Can they object to an algorithm or dataset being used?
 - How will conflicting roles be managed? Who defines and regulates these roles and responsibilities? What happens when roles and/or responsibilities overlap among groups as well as what happens when these responsibilities are relinquished?

- Should participants know any of the internal politics of the process (e.g., is the process transparent)? For instance, who are the champions for civic empowerment on the sponsor side and how do we make allies of them? Which champion is at risk? Who is potentially going to be exposed? What existing tensions underlie inter- and intra-organizational relations that will influence the conversation?
- How are we ensuring there's no tone-policing? Accommodation for uncomfortable dissent? Relatively free from social desirability biases?
- Outcome
 - What questions did the participants raise? Or what is done with the results of a questionnaire? What decisions/policy were made as a result?
 - Who should be able to review the decisions made by the AI system (e.g., Automatic Decision Making Systems)?
 - How do we ensure empowerment is fairly distributed (cf., rich people and the right to be forgotten)? More generally, how are we analyzing fairness²² and how are we acting on the insights?
 - What are our measurement tools (i.e., how did we determine that we have sufficient participation and empowerment)?
 - Can there be a human-guided method to collect and analyze the data in a survey (if that is the method) to ensure the participation isn't automated?
- Ongoing
 - How are participants monitored? What is the ongoing relationship and maintenance of trust with those who were/are participating?
 - Is there any transparency about who participated/is participating in the process?
 - How do we ensure participants haven't been overwhelmed?
 - How often are we revisiting the assessment to update it as the AI system evolves?

4.4 AIA-type questions (to ask the developer)

- Have you user-tested your product? (i.e., who do you consider your 'users' to be?)
- How is feedback/adjustment incorporated into the AI system?
- Is your AI system accessible? How do you define accessibility? Did you test for accessibility?
- Have you complied with RACI for data collection, retention, and disposition?
- Did you establish a process for meaningful consent to data collection? What alternatives were in place for those opting out?
- Can participants object to the AI system?
- Is there a possibility of refraining from designing the AI system at all?

²² <https://diginomica.com/fairness-ai-practical-possibility-new-angle-designing-ethical-systems>

- How did you test and correct for data bias?
- If you had one wish, what would be absolutely ideal to get from the AI system (e.g., if you had complete oversight of a predictive policing algorithm)?
- What data, information, and education do you need your participants/users to have?
- With regards to data, how do we protect groups and communities and not just individuals within data protection mechanisms?
- Who defined the problem the technology is supposed to solve?
- What decision(s) is the AI system replacing and how and by who were they made before the system was deployed?
- Do you use other AI systems (e.g., XAI, adversarial testing) to investigate how the AI system arrived at a specific decision?

Appendix A: Resources and Documents

This is a list of resources and documents that were referenced/mentioned by speakers as well as attendees. If we missed anything below, please feel free to contact us so we can add it to the list.

Bianca's three book recommendations:

- [*Freedom Is an Endless Meeting: Democracy in American Social Movements*](#) by Francesa Polletta, 2004
- [*Unsettling Canada: A National Wake-up Call*](#) by Arthur Manuel and Ronald Derrickson, 2015
- [*Undoing the Demos: Neoliberalism's Stealth Revolution*](#) by Wendy Brown, 2017

[*The Prince*](#) by Nicolo Machiavelli

[*GirlsInAI2021 Hackathon*](#)

[*Procurement as Policy*](#) by Mulligan and Bamberger, 2019

[*What do They Know*](#) - tool for making information requests from U.K. governments

[*Women and leadership*](#) in news media, Reuters, 2021

[*Zuboff's Cycle of Dispossession*](#) by Bianca Wylie, 2020

[*Progressive International homepage*](#) - movement empowering progressive activists

[*Community Technology — Narrative Adventures and Adaptations*](#) by Bianca Wylie, 2020

[*The First Nations Principles of OCAP®*](#)

Appendix B: Workshop Part II Breakout session

Question 1 - What words describe the individuals who could participate but might be excluded (e.g., citizen, representative of the Public, consumer) and how do we define those words? Are individuals the right unit of analysis (e.g., CSOs, community groups)?

- Public workers
- Government officials
- AI decision-makers and policymakers
- Educators, who work on making tech transparent and understandable to the public
- Educators, who work *for* citizen participation
- Community leaders, especially for underrepresented or disenfranchised communities
- Local influencers
- An advocate, working on behalf of someone (with explicit understanding of advocacy efforts)
- Also, someone who isn't normally the one developing/deploying/controlling
- Members of particular linguistic communities/speakers of particular language varieties
- Content moderators
- Journalists
- Youth
- Commuters, those who inhabit a jurisdiction while they work
- The inhabitant of a jurisdiction
- Data subjects (i.e., people represented by or contained in a dataset)
- The people impacted?
- People who are invisible/excluded in AI systems
- People directly impacted by system
- Occupants of various digital landscapes - new forms of internet existence in modernity challenges the idea of physical borders
- Homeless people who walk into a smart city
- Not just people, communities, also the land, the waters
- Social units around the citizens-communities, families
- Communities
- Underrepresented communities
- Marginalized communities
- The communities most negatively affected by the system

- Members of affected communities - those most likely to experience harm
- Future people
- Future generations?
- Those who are already in vulnerable situations in terms of exercising their rights (e.g. prison inmates)
- The folks who will build systems for participation
- Resident: individual who lives in a specific political jurisdiction
- “Affected individuals and groups”
- Not sure how to define, but would the focus be on people who have direct experience at the target of oppression?
- Anyone who has access to, or is impacted by, a service
- As broad as possible, including people from all ages
- People who exist in data gaps
- People without smartphone access
- People outside the imperial core/global North who are nevertheless affected by first world policies
- Refugees who are being used as a testing ground for biometric systems
- Public workers
- Government officials
- AI decision-makers and policymakers
- Educators, who work on making tech transparent and understandable to the public
- Educators, who work *for* citizen participation
- Community leaders, especially for underrepresented or disenfranchised communities
- Local influencers
- An advocate, working on behalf of someone (with explicit understanding of advocacy efforts)
- Also, someone who isn't normally the one developing/deploying/controlling
- Members of particular linguistic communities/speakers of particular language varieties
- Content moderators
- Journalists
- Youth
- Commuters, those who inhabit a jurisdiction while they work
- The inhabitant of a jurisdiction
- Data subjects (i.e., people represented by or contained in a dataset)
- The people impacted?

- People who are invisible/excluded in AI systems
- People directly impacted by system
- Occupants of various digital landscapes - new forms of internet existence in modernity challenges the idea of physical borders
- Homeless people who walk into a smart city
- Not just people, communities, also the land, the waters
- Social units around the citizens-communities, families
- Communities
- Underrepresented communities
- Marginalized communities
- The communities most negatively affected by the system
- Members of affected communities - those most likely to experience harm
- Future people
- Future generations?
- Those who are already in vulnerable situations in terms of exercising their rights (e.g. prison inmates)
- The folks who will build systems for participation
- Resident: individual who lives in a specific political jurisdiction
- “Affected individuals and groups”
- Not sure how to define, but would the focus be on people who have direct experience at the target of oppression?
- Anyone who has access to, or is impacted by, a service
- As broad as possible, including people from all ages
- People who exist in data gaps
- People without smartphone access
- People outside the imperial core/global North who are nevertheless affected by first world policies
- Refugees who are being used as a testing ground for biometric systems
- Not-for-profit organizations that work on ground
- Non-profits and advocacy groups
- Refusniks
- Media? Media can shape public opinion in ways that are orthogonal to accountability
- Business firms? In the US they are treated as persons
- Data gap: a dataset that is not included in the project as a result of bias
- A way to get diverse input, so the entire ecosystem of those who interact with X in any way or form

- Opportunity is to challenge participation designers to consider the types of data they don't collect, as a result of deliberate or non-deliberate decisions

What organizers added in advance:

- Public and private industries
- AI developers
- Tech union organizer
- People outside a country
- Resident: individual who lives in a specific political jurisdiction
- Non-experts (i.e., people who are not AI developers or have specific AI/ML experience)

Question 2 - What role do individuals, groups, organizations play & what responsibility do they have?

- Define the problem that should be solved in the first place
- Define the goal of the tech
- Push for anti-racist applications of tech
- To describe / contribute knowledge that is unknown / shielded in traditional decision making
- Input on design, starting with what systems are developed in the first place
- Identify what is missing from the system
- They need to provide transparency
- Demand transparency and fairness
- Raise concerns
- Disclose abuses
- Free prior and informed consent
- Individuals: Provide feedback on participation "methods" - did they feel involved?
- Stewardship
- Be transparent with the public on their role and assessments
- Oversight regarding respect of rights and access
- Setting requirements - like additional metrics, testing, etc
- Evaluate the diversity and inclusion in defining the pool of participants
- Conduct impact assessment before implementation
- Tech companies should be subordinate to community assessments of impact
- Continual input as the output goes through stages - feedback on MVP, beta, final, etc.

- Determine environmental impacts
- Follow up over meaningful period of time, to promote accountability on part of developers
- Ensure an adequate possibility of redress
- Evaluate potential negative consequences
- Determine what the system is meant to accomplish
- Articulation of what participants themselves see as the benefit of the system (since it may not always be in alignment with developers' vision)
- Organization: Facilitate translation between participants and technical aspects
- Organisations: Design people-centric tools to gather participation
- Organizations and groups are made of individuals, do they need additional representation beyond that of its members?
- Organizations - collective action that represents the group
- Organizations: Facilitate public audits
- Organisers should prescreen for conflicts of interest. Kick people with a conflict of interest out of the citizen processes
- Groups/people in power should bring in those with smaller platforms
- Groups: Publicise opportunities for participation
- Groups: Hold organisations as well as individuals accountable for their participation
- Mobilizing collective action
- Coalitions can organize around legislation to constrain harmful technologies
- Raising critical questions
- Those organising it should limit the power of those who can hijack the process - balancing
- it seems like it would depend on if the participants are asking for something (like a city park) or if they are being consulted about how to do something that is already planned (building a road)...
- Identifying, challenging and verifying ground truths
- People represented in data: contextualizing data
- Educate and get educated/informed
- Informing constituents and recruiting participants
- I think everybody, regardless of power, has a moral obligation to avoid harming and probably an obligation to benefit
- Who defines the responsibilities? Legal, professional, ethical...
- There is something performative about roles - which relates to how we define 'participant', 'citizen' etc.

- Is it not more about power rather than role?
- Challenging existing power structures and considering dismantling tech if needed
- Government has the responsibility to regulate; individual representatives have the responsibility to be literate in these issues
- Provide platforms for participation
- Clarifying intentions and goals of a community
- Provide their perspective - we don't all know all possible instances

What organizers added in advance:

- Decide which products go forward & which don't
- Provide feedback about outcomes (e.g., classifications)

Question 3 - What words do you use to define civic empowerment and how would you “operationalize” those?

attendees defined “civic empowerment” as:

- Equal opportunity metrics
- Meaningful community feedback
- Ability to revisit past policy, laws, decisions
- Create new laws
- Ability to enact policy change
- Ability to meaningfully refuse
- Representation from “people with lived experience”
- Material analysis of (marginalized) community impact
- Accountability and Transparency
- Independent mechanisms for accountability
- Problem-focussed user research
- Can people see processes and outcomes?
- Operationalize assessments through Beyond GDP work i.e. UN SDGs OECD and other sustainability frameworks as well?
- Well-being of communities
- I liked the point about power vs. role and roles/responsibilities being performative. Coming from a tech background, Arnstein's ladder was really interesting/novel for me, esp. with regards to tokenism which I think ties in to the previous point. That said, I think definitions should emphasize this tokenism/empowering distinction
- Tech power to the people
- Are participants compensated?

- Ongoing engagement over time (i.e., the resulting algorithm goes through iterative participatory/feedback sessions)
- Timeliness
- Participants feel that the process has been “fair” (the process has been justified at each stage with opportunities for recourse and flexibility)
- Recognising the "beyond borders" nature of the tech developed in country X and deployed in country B. When and who should say something about it?
- Checklists, standard operating procedures, frameworks that are transparent
- This reminds me of the veil of ignorance (Rawls!!)
- the ones offering a solution understand the actual problem and those affected by it understand the implications
- participants should see some impact in whatever results from the process - they should be able to identify their impact.
- Affected individuals' rating along the line of "do I find this generally net beneficial?"
- Improved Quality of Life: Has the outcome improved (even incrementally) the QoL for some folks?
- Increased number of citizen voices: Is the empowerment helping folks raise their voices, concerns and opinions?
- Increased involvement of marginalized/historically disempowered communities
- Future willingness for participation in similar opportunities
- Operationalize: regularly scheduled, periodic assessments on part of developers to gather feedback
- Escalation channel, for rapid response to urgent issues
- Third party assessment of "was input from the original participatory process adequately incorporated?"
- Increased efficiency of organisational mechanisms

What organizers added in advance:

- Do they feel good about the resulting algorithm?
- A set of metrics to assess how well do they understand the algorithm
- How can the government communicate with the civic groups just how they're being “heard”?

Appendix C: What would be contained in a questionnaire?

- What normative assumptions are you making? If using words like 'just', it's important to define terms - often in tech spaces people have different ideas about what these words mean
- Who decided that this tool needed to exist? To what extent did the public want this?
- Whose participation does the tool audit?
- Who have you defined as being the public? What support have they been provided to understand the technology product or service that is being built?
- What proportions of public have been included?
- What activities did you do to ensure public participation?
- Have you already done PR on this tool? Did this come before public participation in the project or after?
- What questions have these public participants raised? What was the response to their questions?
- What decisions were made through the design process as a result of civic participation?
- How are participants recruited or brought into the participation process?
- How did you reach out to and select participants?
- How are participants monitored? Is there any transparency about who participated in the process?
- What is the ongoing relationship and maintenance of trust with those who were/are participating?
- Can there be a human-guided method to collect the data in a survey to make sure people are even weighing in properly?
- At what stage do participants participate in the development/conceptualization of a system?
- How much tech/other literacy should be built in?
- What kind of trade secrets, legal clauses, etc should not be included especially if it will be risky to the organization's business?
- If participants object to an algorithm or dataset being used can that be eliminated?
- How are we garnering participation? What channels are we using to reach people?
- What is the outreach of the tool/survey? Who is missing?
- How do you identify who you're surveying, how do you build the frame of what you're measuring?

- What fairness analysis are you using and how are you acting on the insights?
- How do you evaluate changes in the degree of civic empowerment?
- Who defined the problem the technology is supposed to solve?
- How did you inform populations/communities affected by the technology on how it works? (Literacy)
- What decisions is the tool replacing and how and by who were they made before the tool was deployed?
- What data went into it? How does feedback/adjustment work?
- Who can audit the tool in detail?
- What is the ongoing relationship and maintenance of trust with those who were/are participating?
- What structural inequalities intersect with the system and how are you accounting for the way your system may entrench or worsen those inequalities?
- Can participants object to an algorithm or dataset being used?
- How did you test and correct for data bias?
- Did you establish a process for meaningful consent to data collection? What alternatives were in place for those opting out?
- What is the RACI for data collection, retention, and disposition?
- How are power dynamics within or between groups of participants made visible or mitigated?
- In the case of conflicts between groups of participants and their preferences, how are those conflicts resolved? For instance, How much power is granted to developers to exercise ethical responsibility if those ethics conflict with their professional responsibility or requirements?
- Who is filling out this assessment?
- Who should be able to review the decisions made by the tool?
- How often do you revisit this assessment to update it based on how the system is evolving?
- Can you ensure responses are free from social desirability biases?
- How do people who don't have Internet participate?
- Is your tool accessible? How did you test for accessibility?
- Are you ensuring your user base is representative of the population?
- What is the impact of the tool on mental well-being?
- Have you thought about the user testing of your product? (i.e. who do you consider your 'users' to be?)
- Who does the burden of proof lie with?

- Do you include tools to investigate how the system arrived at a specific decision?
- What is their TechUtopia: what would be absolutely ideal to tell what you really want from the tools?
- How does the tool advocate for some sort of solidarity among (underrepresented) communities?
- What ACCOUNTABILITY is in place?
- Is there a possibility of refraining from designing the system at all?
- How does the tool know I'm a human?
- How many cases of AI errors/failures have been reported?
- What is the due process for AI failures?

What organizers added in advance:

- What types of people are you inviting/including (CSOs, stakeholders?)
- What public participation model are you using?