ICICLE Educational Fellowship Final Report: Democratizing AI through Stakeholder Participation[1]

Tim Elmo Feiten & Collin Lucken

Contents:

1. Overview

1.1 Project Outline

This document is the final report of the ICICLE Education Fellowship project carried out by Tim Elmo Feiten and Collin Lucken. As stated in the article "Educational fellowship cohort works to democratize AI," our project was initially described as initiating a dialogue between ICICLE members and the Center for Public Engagement with Science to devise strategies for deepening stakeholder involvement in agricultural and ecological AI systems. The original timeline for this project was as follows. In August, we traveled to the Translational Data Analytics Institute at the Ohio State University for our ICICLE Thought Leader Workshop. In September, we held another workshop on stakeholder engagement at the UC institute for public engagement with science with ICICLE members. Between June and November, we held a series of one-on-one interviews with ICICLE members that focus on stakeholder engagement in their use-case domains. Finally, in November, we attended the ICICLE all-hands meeting, presenting the work we had developed as part of our fellowship. We also presented a poster on stakeholder engagement in ICICLE, and we held a brief meeting regarding ICICLE's stakeholder roundtable to be discussed below.

As the results of these workshops and events, we planned to have 5 documents as deliverables at the end of the project. The first three documents are documentation of our meetings: the ICICLE thought leader workshop, the UC PEWS workshop on stakeholder engagement, and brief summaries of each of our interviews with

[1] We would like to express our deep gratitude to everyone within ICICLE and beyond who discussed these topics with us and generously provided us with input and feedback, and to the ICICLE AI research institute for granting us the opportunity to conduct this research as educational fellows.

ICICLE members. This document, our executive summary of the project, is our fourth document. The fifth document is a draft of an academic research paper in philosophy based on our time in ICICLE as educational fellows.

<u>1.2 The 'Plale Diagram' and the status quo of ICICLE's stakeholder engagement</u>

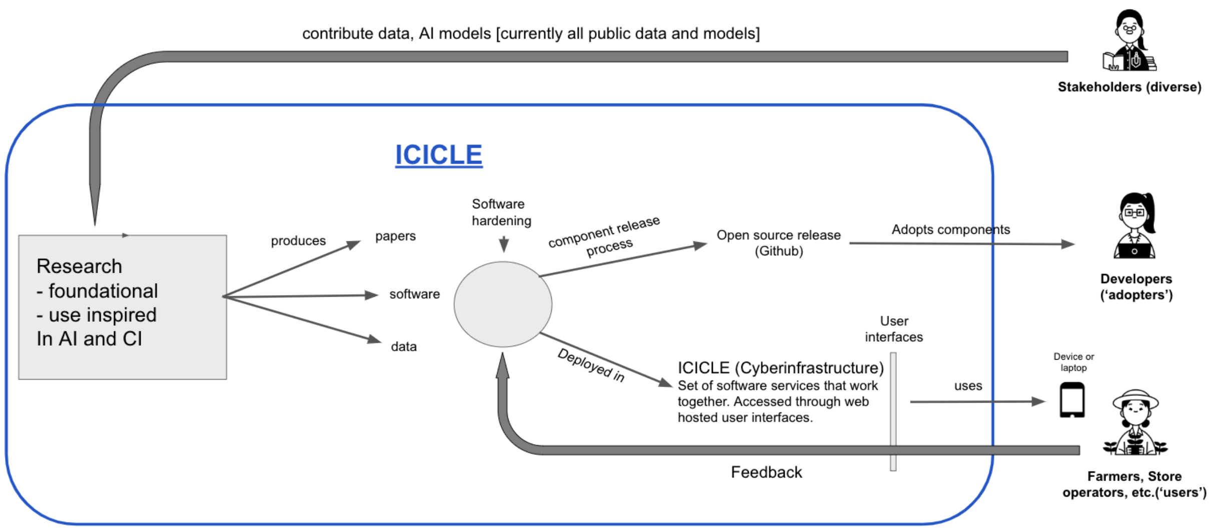

Figure 1. The 'Plale Diagram'.

Figure 1, above, is what we refer to as the 'Plale Diagram'. This diagram was originally drawn at the ICICLE Educational fellows kick-off cohort meeting in June 2023 by Beth Plale. Its purpose was to clarify the overall organizational structure of ICICLE for the educational fellows who had just been introduced to the overall project for the first time. Additionally, the diagram represents the relations between ICICLE's stakeholders and its internal research and development activities. As such, it will be an especially useful guide for representing the overall status quo of ICICLE's stakeholder engagement. On the far left of the diagram, a box representing the foundational and use inspired research efforts of ICICLE in AI and CI is depicted. Three types of products are produced by these research efforts: papers, software, and data. Software is then pushed to a hardening phase before bifurcating into the two distinct but potentially interrelated streams.

On top, the stream is described as beginning with the component release process where research products are published as open source releases on Github. From there, one group of stakeholders we have labeled Developers can adopt these components for use in their own systems. What's notable about this group of stakeholders is their relatively high degree of technical knowledge. This means the kinds of engagements that can be expected between ICICLE and developers will accordingly be of a more technical nature than those with the other, following group of stakeholders. The second stream, labeled on the bottom, following the software hardening phase is the deployment stream where ICICLE products are deployed as a set of software services that work together. This set of software services is accessed through web hosted user interfaces, and can be expected at times to be available on devices ready for use such as phones, laptops, etc. This stream leads to our second group of stakeholders, labeled Farmers, Store operators, etc. These are ICICLE's users, as opposed to its adopters. Unlike developers or adopters, these users cannot be expected to have the technical skill or knowledge necessary for adopting ICICLE's research products directly to their needs.

In addition to the two production streams, the Plale diagram also features feedback loops leading from users and adopters back to earlier stages of the research and development process. The notion conveyed by these feedback lines is that AI and CI research efforts ought to be continuously re-visited in light of the feedback obtained by ICICLE members from their users and adopters. These lines, then, are potentially the most important features of the diagram for understanding the stakeholder engagement situation of ICICLE. In what follows, we summarize the strategic situations of specific use-cases in the second section. Doing so will allow us to contextualize which of ICICLE's pre-existing resources might be best fit as institutional mechanisms for realizing the feedback lines of the Plale diagram. We suggest which of these resources seem best fit for that task in section 3. In section 4, we articulate some principles of democratization theory that are relevant to ICICLE's stakeholder participation situation. After explaining these principles, we illustrate how they are relevant to specific strategic situations and ICICLE stakeholder engagement resources. Finally, we end our report by explaining how, on the view we have developed as Educational Fellows at ICICLE, stakeholder engagement should not be viewed as a separate, distinct activity that AI and CI researchers must engage in in addition to their AI research. Instead, we insist that doing good AI research and democratizing AI are one and the same.

2. Strategic Situations for Stakeholder Engagement

2.1 Specifics matter

Since ICICLE is a large institute made up of teams working on very different projects, it also contains a large variety of different situations and contexts for stakeholder engagement. Different groups of stakeholders have varying degrees of technical expertise, varying access to different kinds of resources, and are affected by AI technologies in very different ways. They have their own goals and cultural identities, and their encounters with AI do not happen in a vacuum, but within social situations and environments that already exist. The specific conditions in which AI technologies get deployed crucially shape the potential effects these new technologies can have, so both the success of the ultimate goals of AI research and the processes of stakeholder participation which can help achieve these goals depend on an awareness of the relevant social, economic, and political constraints that already exist.

We cannot provide a comprehensive account of the landscape of stakeholders for every part of ICICLE here. Instead, we want to give three brief accounts of individual projects and the strategic challenges and opportunities for stakeholder engagement in their particular context. These examples can serve to highlight the kinds of questions that are most relevant when thinking about stakeholder engagement strategically, and thus provide a lens through which other projects within ICICLE or beyond can be investigated. Our examples concern the use of AI for computational ecology, smart foodsheds, and the development of an ICICLE reference architecture. The first two of these focus on external stakeholders and the third on internal stakeholders, where the first two could be thought of as 'users', while the reference architecture is mainly relevant for 'adopters' of components developed within ICICLE. These short vignettes illustrate the strategic challenges for stakeholder participation in AI research, demonstrate its potential to contribute to democratization, and highlight existing expertise and success stories within ICICLE.

2.2 Computational ecology

Leveraging AI technologies for research on animal ecology is a difficult task, but holds great potential for the conservation of biodiversity (cf. GPAI 2022). One central problem in such conservation work is a lack of different kinds of data. For many species, we have no reliable numbers about population changes, and this knowledge gap cannot be filled because lower-level data about the local presence and movement of animals is missing. AI solutions for image recognition could help track animals in the field by automatically analyzing data from a variety of visual sources, including satellites, drones, and camera traps, as well as bioacoustic sensors and on-animal trackers (Berger-Wolf 2023). This could significantly improve the availability of biodiversity data, but only if challenges relating to several distinct groups of stakeholders can be overcome.

First, the researchers developing machine-learning solutions for animal ecology need to account for the highly specific contexts in which their software will be deployed. These solutions are most often developed using benchmark image datasets that are too far removed from the actual data in the field, and this data undergoes distribution shifts once the system has been deployed. Additionally, the cyberinfrastructure available in locations where biodiversity data is gathered is very heterogeneous and most often very different from the computational constraints of most machine-learning development. Second, the local communities living in the areas where biodiversity is to be measured often have well-founded concerns about the introduction of surveillance technologies into their environment and what motivates the organizations driving these developments. Lastly, a large and heterogeneous group of different kinds of academic, government, and NGO stakeholders engaged in supporting conservation efforts has to navigate both of these groups if they want to leverage AI technologies, often without specific AI expertise of their own nor any particular interest in the technology beyond its utility as a tool for assessing biodiversity (Berger-Wolf, GPAI 2022).

Dr. Tanya Berger-Wolf emphasizes the importance of engaging different stakeholders by contrasting it with its opposite, the unfortunate implicit assumption that technology development can succeed with the mantra "build and they will come". However, many of the stakeholders for whom ICICLE is designing AI tools currently have nothing to do with AI and are not interested in it beyond the support it can provide to conservation and biodiversity research. In contrast, the challenges outlined above can only be addressed successfully by involving stakeholders in the development process, similar to the notion of participatory design. Stakeholders need to be engaged at different stages of development from design to deployment; the challenge consists in finding out when and in what way they can best participate. Different types of engagement include research collaborations, inclusive community-building, field work and data collection, tech-transfer, and tool development.

Because "environmental data ethics, particularly around data sovereignty and privacy, can be complex, especially in the context of local communities and indigenous groups", the GPAI suggests to follow not just the more operationally-oriented FAIR principles—"data should be findable, accessible, interoperable, and reusable"—, but also to adopt the CARE principles—"providing collective benefit, ensuring authority to control, responsibility and ethics" (GPAI 2022: 30). One key concern is how to gain the trust of local stakeholders. ICICLE seeks to address this challenge in three ways, by ensuring safety, "transparency" and "accountability", and by fostering an "inclusive, collaborative partnership" (Tanya-Berger Wolf 2023: 23). This last point is particularly relevant to us. The general plan for achieving this kind of participation involves diversifying AI and biodiversity, as well as changing the way solutions are created by building them "with whom they benefit, not for them" (Tanya-Berger Wolf 2023: 23). These measures are necessary in order to produce AI solutions that really work for their intended purpose, but they simultaneously align with the goal of democratizing AI.

2.3 Smart Foodsheds

Within the domain of use-inspired science called "smart foodsheds", there are several different projects. We focus first on the challenges surrounding stakeholder engagement from the perspective of Michelle Miller, an economic anthropologist who first started getting involved in AI because of problems around information asymmetry in the food system. Big business has access to AI in ways that small and medium sized actors and supply chains don't have. These smaller participants in the food system are very disadvantaged when it comes to accessing markets. The existing power imbalance between differently-sized actors within the food system got exacerbated when the large companies started adopting more powerful information technology, and part of Miller's mission now is to help develop AI technologies that enable the smaller players in the food system equitable access to markets and supply chains. The global food system is a huge challenge, because each individual person is active in the food system (9 billion participants). It gets complex very quickly and it has to be self-organizing. As we get more and more uncertainty in our food systems (primarily due to climate change but also other factors), there is a need for more communication to keep it self-organizing. Otherwise, the system will either break down or become very authoritarian. To avoid that, we need to keep the communication lines very open, and AI can help with this.

Miller started working on how to develop new supply chains for organic and sustainable food products, and now works beyond that on any kind of emerging supply chain in the food system area. For instance, she is working with tribal partners to get their traditional foods into supply chains and has done some work with African American communities, and more generally rural communities that don't have good food access. Miller sees AI as a very positive opportunity for improving access to food and food supply chains, but also as raising many challenges. Some of these challenges can be quite fundamental, such as a lack of internet access in rural areas. Even today, there is still a long way to go in solving this problem. Working in academia one gets used to having a computer with a screen—suddenly, having to use a smartphone to do more complex tasks is difficult and irritating, but a lot of people don't have access to a desktop monitor. Using federated data becomes a big issue, not all phones are able to work with the kind of data that is necessary to participate in supply chains. If someone wants to plug in data about a shipment that just came in, the system has to be able to take the input data from various different kinds of phones. These kinds of barriers to access can easily go unnoticed if research is done only in the lab or in a university office.

Miller's primary approach is participatory action research. This means working with communities, and turning the issues they express into a research question. One of these issues concerns transportation and wholesale markets: Farmers can't figure out how to access wholesale markets in Chicago. A possible way to approach this problem is to use transportation data and create shared transport systems that use AI to keep the supply chain informed all the way across. We can think of this as 'creating a strong information ecosystem'.

Miller learned about the methodology of participatory action research as an undergrad reading the journal *Practicing Anthropology,* this was during the AIDS crisis, and anthropologists working in public health were trying to figure out what the theory of disease was in specific cultures in Africa, to figure out the public health component. While it is likely that the public health community has done the most work using this strategy, it presents clear opportunities for democratizing AI research. The framework focuses on the co-generation of knowledge. One way to think is to assume that certain people are experts and others are just supposed to use the knowledge. This sort of authoritarian logic is faster, but also less democratic. Instead, participatory action research enables various stakeholders to weigh in on the research who normally do not have the privilege needed to access and participate in this kind of knowledge generation.

Working with different stakeholders in AI research for smart foodsheds comes with a variety of challenges. A central issue concerns data: On the one hand, smaller players in the food system have economic difficulty participating in the information ecosystem that enables large-scale Supply Chain Optimization, largely due to prohibitive fees. On the other hand, various communities are also concerned about privacy and data sovereignty. Some communities have already had NSF-funded research violate their boundaries by flying drones over their territory without permission or even identifying hidden burial sites using satellite imagery in direct violation of important cultural traditions. Because of these and other factors, getting the kind of data necessary to develop an AI solution that really addresses the existing needs can be very challenging, and even when the necessary information is available in principle, it often does not come in the form needed for AI research, but may only exist on paper or even in someone's head.

In some sense, AI solutions that enable smaller players in food supply chains better access to logistical processes are solving a problem that AI created—they address a current power-imbalance created by the adoption of AI-driven supply chain technology by large corporations, which has given them an advantage over local and regional producers and communities. The kind of work that Miller does also involves a lot of code-switching. As Miller puts it, "one of the really difficult ethical challenges is understanding that we don't understand each other". This is a problem when talking to external stakeholders, but also occurs within ICICLE itself. Terms like data, information, and knowledge have very specific meanings in the AI domain, but normal people often use them interchangeably. This can lead to a disconnect over what's being discussed. Additionally, tribal communities place a strong emphasis on keeping their languages alive.

Another great example of stakeholder participation expertise that is already present within ICICLE is the work of Dr. Ayaz Hyder and his collaborators. The core of his project within ICICLE has been to develop a conversational AI agent that provides users with easier access to the information provided by an agent-based model of a food system. The agent-based model had been developed by Hyder and his collaborators in a separate project prior to ICICLE, and this new AI agent allows users to ask their questions and have them answered in a conversation, rather than having to navigate the interface of the agent-based model and interpret its results themselves. The process through which the agent-based model was originally developed involves an intensive and structured methodology of stakeholder participation and could potentially be adopted as a valuable model for engaging stakeholders across a variety of contexts within ICICLE's many research thrusts.

When developing the original agent-based model, Dr. Hyder and his collaborators adopted the method of "group model building (GMB)" that had been developed for the study of "[c]ommunity-based system dynamics (CBSD)", "a participatory method for involving communities in the process of understanding and changing systems from the endogenous or feedback perspective of system dynamics" (Koh, Reno, and Hyder 2018: 278, Hovmand 2014: 1). In contrast to these investigations of high-level or global phenomena in a system, agent-based models simulate the effects of possible actions taken by individuals, such as participants in a food system reacting to a change in their infrastructural environment (such as a store closure). The use of group model building addresses a specific methodological shortcoming: "Since agent-based modeling has been often considered highly complicated to program and computationally intensive to run, ABM modelers typically have a greater role in constructing their models than stakeholders and often build their ABMs with the input of one or very few subject matter experts. This represents several inherent limitations in ABMs, such as issues of face validation, lack of uptake by stakeholders, and limited confidence in the model's inner workings." (ibid. 279). Introducing a method into the construction of agent-based models that centers the involvement of different groups of stakeholders in the entire process, from scoping, over designing, to evaluating the model, promises to enable not just the development of straightforwardly better models and higher adopting in the field, but also constitutes an instance of direct democratization of the research enterprise.

In the original formulation of group model building for the study of complex social system dynamics, Hovmand emphasizes the need to define both a community and the notion of participation. This shows that group model building from its inception centers around careful attention to factors that are essential if the method is to play a central and genuine role in efforts to democratize research is already present in the history of group model building. Hovmand describes five phases:

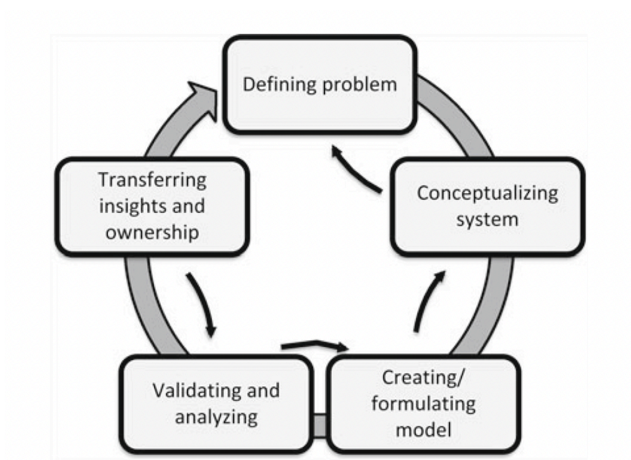

Fig. 1: the process of group model building, from Hovmand (2014), p. 13

Dr. Hyder and his collaborators adopted this method to conduct two workshops with relevant stakeholders from the school system, a public health department, NGOs, and academia. Over the course of these workshops two, modelers and stakeholders co-created a model based on the results of activities focusing on "hopes and fears", "main barriers", "mapping interactions", "key stakeholders", "intervention levers", and "validating models" (Koh, Reno, and Hyder 2018: 281-286). These exercises focus on establishing a systems-level view. Identifying the parts of the system and the connections between the parts is already a concept definition exercise. The next steps include identifying key stakeholders, and graphing over time which outcomes these stakeholders care about, which ones they hope for and which outcomes are feared, as well as policy recommendation exercises. After getting the broad parameters from a workshop, the modelers create a prototype and have another meeting with the stakeholders one or two months later, where they ask the stakeholders to evaluate the strengths and weaknesses of their prototype.

This kind of work is important as an example of democratization of research within ICICLE itself. A variety of stakeholders who will potentially benefit from and/or be affected by the final product get to participate in the development process from the very beginning, which helps ensure that the solution meets their needs, respects their priorities, identities, and boundaries, and can help reduce biases in the data and solution design. The fact that this success story and the expertise connected to it are available within ICICLE itself makes it both more relevant and more accessible for democratization efforts within ICICLE than appeal to external case studies and best practices would be.

2.4 Reference Architecture

Dr. Rajiv Ramnath is part of the ICICLE leadership team and is involved in two different thrusts: the AI for CI cluster, and the software thrust, a unifying and cross-cutting thrust that takes software developed in ICICLE as a basis for developing a reference architecture for ICICLE. A reference architecture provides a unified description of the components of ICICLE as a cyberinfrastructure and specifies e.g. interfaces, output formats, and standards to ensure that different components developed within ICICLE can work together, and can be understood and used in a homogeneous way. This can mean that a module used for authorizing a developed component, or different versions of this module, can be used across ICICLE research teams, and that training materials developed for the purpose of enabling adopting of ICICLE components and for broadening participation in computing more generally will be relevant and useful for understanding ICICLE across the board. Without a reference architecture, different teams might be developing components that follow different principles and are not interoperable, and training materials might explain concepts and structures that are relevant only to some components and not others.

The work of creating a reference architecture is relevant for the aim of democratization in two ways: First, the process involves researchers from different teams and thrusts across ICICLE and supports them in developing components that are more easily adoptable by external stakeholders, because they follow common standards and principles. Second, the RA itself helps external stakeholders adopt these components by making it clear how they work and what is required for existing systems to interface with them. By serving as a 'blueprint' for ICICLE, the RA can help remove barriers to access, making ICICLE components more widely available to technical audiences inside and outside of ICICLE, and thereby eventually also to non-technical end-users. This function of the RA is also crucial when the NSF is considered as a stakeholder: Since the NSF wants the work done at ICICLE to 'scale', an RA is crucial: The more standardized and well-documented the shared characteristics of ICICLE components are, the more feasible it becomes for them to be adopted and maintained by people other than their original developers. This feeds into the important question of 'sustainability', not in the sense of ecological impact, but in terms of producing technology that can be sustained and remain usable and useful for prolonged periods of time. This is doubly important for ICICLE because of its focus on cyberinfrastructure—which by its nature is required to be operational for on longer time-scales than the services that are built on top of it and can be active for varying amounts of time—and because of its status as an NSF-funded project with a five year runtime and definite horizons for renewal of funding. If ICICLE is to build cyberinfrastructure that helps democratize future AI technologies, it needs to be accessible and open to a wide range of future stakeholders that not only make use of ICICLE but also help maintain it. The RA, both in its role of standardizing ICICLE output and by enabling accessible educational materials and activities, is one of the central pieces for addressing this challenge.

The concrete challenges within ICICLE for developing and adopting the RA concern motivation, experience, and resources. The team for RA is small and its members also have other commitments. They work with different thrusts to unify their work into the RA, and the leadership team represents the RA project in its dialogue with the NSF. One of the largest groups of internal stakeholders is made up of students who do CS research and develop ICICLE components. The biggest challenge with respect to this group is that their experience is in doing CS research that yields something akin to proof-of-concept code products, rather than the kind of hardened, deployable software components that can be released and used by other stakeholders. As a project that ranges over many different thrusts within ICICLE and works to unify their work, the RA project is also in a similar strategic situation as the efforts to democratize AI: existing silos have to be broken up or connected in order to enable the kind of institute-wide transparency and shared language that is necessary for making research accessible and democratic.

3. Existing Resources within ICICLE

In the previous section, we reported on some of the expertise and success stories around stakeholder participation that are already present within ICICLE and could function as invaluable resources for future democratization efforts. In this section, we identify two further pre-existing resources relevant to ICICLE's stakeholder engagement situation, as it pertains to the goal of democratizing AI.

The first pre-existing resource is the plan for an ICICLE Stakeholder Roundtable. The stakeholder roundtable is described on the ICICLE website as "Organizations and individuals interested in using ICICLE AI tools and data services are encouraged to be part of the Stakeholder Roundtable. The Roundtable is a venue to stay connected with the Institute and have your voices heard. Membership is free." Early on in our fellowship, we identified the stakeholder roundtable as a concrete institutional mechanism through which the ICICLE's relevant stakeholders could have their voice heard by ICICLE members responsible for AI and CI research and development. Having identified the stakeholder roundtable as such a helpful mechanism, we held our workshop at the UC Center for Public Engagement with Science (UC PEWS) with the stakeholder roundtable as its theme. We asked attendees for recommendations regarding how the stakeholder roundtable could work, and what it should accomplish.

What we concluded at the UC PEWS workshop was that the stakeholder roundtable needs to be accompanied by an accountability mechanism if it is to play the role of "a venue to stay connected with the Institute and have your voices heard." This is partly what we communicated regarding the stakeholder roundtable at our working meeting on the all-hands agenda. At that working meeting, our aim was to identify who in ICICLE would be willing to continue working on the stakeholder roundtable after the cessation of our fellowship.

The second relevant pre-existing resource we identified was the AI ethics framework in progress by Sadia Khan. Sadia shared this framework with us during our bi-weekly meetings throughout the fellowship. The document describes its goal as "a living document, intended to guide the ICICLE team by means of a set of core principles, as institutional commitment, and operationalizable steps." While this aim is obviously laudable, it also clearly aligns with the AI democratization goals of our stakeholder engagement project. Likewise, its first listed AI Ethics Core Principle is accountability. The framework describes accountability as follows: "Accountability in ICICLE can be interpreted as a commitment to documentation trails that account for the actions, decision, and outcomes of development and deployment." Other principles listed are trustworthiness, fairness, privacy, and democratization. What we suggested in our workshop at the all-hands meeting, as well as in this document, is that accountability is necessary for the guarantee of these other principles. If nobody is accountable to lapses in fairness, privacy, and trustworthiness, then the final goal of democratization will not be achievable. In our political democracy, congress people, for example, are accountable to their constituents. When they fail to live up to their campaign promises, they lose their job. What kind of analogous mechanism exists for ICICLE? If the answer is none, then we can expect the goal of AI democratization and the Ethics Core Principles of the framework to ring hollow.

4. Conceptualizing Democratization and Participation

Our investigation into the complex research situation at ICICLE has yielded theoretical insights into the role of stakeholder participation for the aim of democratizing science that can be grouped into two broad categories. The first category conceptualizes participation through the lens of research into public engagement with science (PEWS) and the theoretical toolbox of radical embodied cognitive science (RECS). This line of investigation has resulted in the paper "Leveraging Participatory Sense-Making and Public Engagement to Democratize AI", which we have submitted to a special issue of *Studies in History and Philosophy of Science* on "Community Engagement as Scientific Practice". Section 4.1 gives a brief summary of the central points of this paper. The second category of theoretical insight concerns the relationship between the specific strategic situation within ICICLE, philosophical theories of democratization, and practices of AI research. These points are summarized in section 4.2.

4.1 research paper: “Leveraging Participatory Sense-Making and Public Engagement with Science for AI Democratization”

Our paper explores new potentials for productive dialogue between public engagement with science (PEWS) and radical embodied cognitive science (RECS). We establish a strong connection between the two fields by highlighting parallels between the views they reject: the ‘deficit model’ in science communication and the ‘information processing paradigm’ in cognitive science. Furthermore, we show that the positive visions of PEWS and RECS are similarly aligned: The concept of participatory sense-making from enactive cognitive science provides an account of why active, dialogical engagement in science communication is so effective. Conversely, processes in which affected communities actively engage developments in science and technology through contribution and contestation provide an invaluable case study for RECS accounts of emergent dynamics in techno-cultural systems. After establishing the connection between PEWS and RECS, we motivate the need for what we call ‘participatory cognitive strategies’. Finally, a brief case study shows the potential for these strategies in actively involving different groups of stakeholders throughout the development of large-scale AI systems, allowing us to make a conceptual contribution to ongoing debates about the meaning of ‘democratizing AI’ in this project and in the larger AI initiative of which it is a part.

There is a strong resonance between two fields of academic research and practice that have hitherto not been connected. These fields are public engagement with science (PEWS) and radical embodied cognitive science (RECS). Public engagement with science is an interdisciplinary effort on the part of scientists, science communicators, and other researchers to both better understand the public’s relationship to science and improve that relationship through strategic interventions such that the public is more thoroughly involved in science’s practices and aims. For example, by improving the quality of these relationships, PEWS practitioners aim to help the public become more trusting of science and also help science become more trustworthy. Radical embodied cognitive science is a rapidly developing research program incorporating several sub-fields of psychology including ecological psychology, enactivism, and embodied and extended cognition. These subfields are allied in virtue of their joint rejection of some of the most prevalent assumptions of the dominant paradigm of cognitive psychology. Ecological psychology and enactivism, for instance, reject accounts of cognition where intelligence ultimately amounts to the capacity to construct internal mental models of an external world.

Instead, they both conceive of intelligence as emerging through the embodied interactions of organisms with their environments.

Even though the two fields at first glance seem to investigate different phenomena using disparate methodologies, we claim that there is substantial overlap and resonance between the two. They both criticize views of human social and communicative interactions which model this rich and complex kind of process as a simple problem of transmitting information through a channel. Instead of investigating human interactions through a lens borrowed from electrical engineering, both PEWS and RECS emphasize that meaningful engagement happens in concrete physical environments and social contexts, between persons whose lives are shaped by different identities and cultures. This is important e.g. in technology development projects involving disciplinarily and/or geographically diverse teams across cultures and time zones, and also during early phases of a project when the end users and end uses of a technology are not yet fully specified, or indeed cannot be predicted in advance by the creators.

Exploring the similarities between these fields promises to be rewarding in multiple ways: Efforts to systematize the practice of PEWS and provide guidance for practitioners typically consist of collections of best practices with an occasional reference to one of several psychological theories, but without a coherent conceptual framework connecting these parts into a whole. We propose that RECS can provide just such a framework, explaining why it is that open, participatory strategies for engagement that make room for the active and creative participation of its 'audiences' are often more effective than unidirectional broadcasting of information by the 'experts'. Conversely, PEWS is a complex intellectual activity that happens across different time scales and in rich cultural contexts, providing RECS with an opportunity to demonstrate that it can tackle the 'scaling up problem', applying the fundamental concepts of embodied cognition beyond examples of motor-coordination and dyadic interactions involving little abstract symbolic thought. Furthermore, we propose that uniting insights from PEWS and RECS into an account of participatory cognitive strategies should yield a conceptual lens that can be fruitfully applied to other complex social interactions involving science and technology. To demonstrate this potential, we briefly sketch a case study from the field of computational ecology, at the intersection of AI research and citizen science.

We first highlight the similarities between the critiques of information-processing models developed by RECS and the rejection of the deficit model in PEWS. Instead of the computationalist orthodoxy, which places representations at the heart of cognitive science, RECS maintains that subjects directly perceive meaningful opportunities for action in their environments, so-called affordances. We discuss the PEWS critique of the 'deficit model' using the example of Cumbrian farmers affected by nuclear pollution following the Chernobyl disaster. Next, we trace resonances between the positive counter-proposals developed by RECS and PEWS. One concept from embodied cognition in particular, participatory sense-making, provides an elegant theoretical grounding for the recommendations given by PEWS experts of bidirectional participatory engagement conducted by experts and non-experts seeing eye-to-eye. We refer to the various best practices of PEWS which can be united under the conceptual framework of embodied cognition as 'participatory cognitive strategies'. Finally, we apply this theoretical lens to a case study of a large-scale AI research project which leverages AI in the service of use cases including smart food distribution and animal ecology and conservation efforts. The emphasis on participatory engagement helps us think through the question of what it would mean to "democratize AI", given the particularities of a large and diverse AI research project.

4.2 Democratization and Participation in AI Research

When asking what it means to 'democratize AI', we can distinguish between two basic lines of inquiry. One question concerns the target of democratization, the other question concerns the methods employed. The former question is about what exactly it is that should be democratized, the object of democratization. The latter question asks about the process by which this democratization is to be achieved.

The NSF has described a specific object for the democratization of cyberinfrastructure. Manish Parashar, who recently served as Office Director of NSF's Office of Advanced Cyberinfrastructure, has argued that democratization in the context of cyberinfrastructure ought to be understood essentially through the notion of access: "ensuring broad, fair, and equitable access to advanced cyberinfrastructure (CI), including computing, data, networking, software, and expertise, that is, democratizing access to CI, is essential to democratizing science" (Parashar 2022, 80). For ICICLE, Plale, Khan, and Morales (2023) follow Seger et al. in adopting a broader perspective, distinguishing between four different targets for democratization:

> *"Democratization of AI use* refers to a need for making AI technologies more accessible—easy to acquire, build, and operate by the general public. This can also be called democratization of access. *Democratization of AI development* encourages diversity in the design and development of AI. *Democratization of AI profits* is about "facilitating the broad and equitable distribution of value accrued to organizations that build and control advanced AI capabilities." *Democratization of AI governance* seeks to distribute influence over decision-making about the risk and benefit of an AI product 'to a wider community of stakeholders and impacted populations' [56]". (Plale, Khan, and Morales 2023: 6)

The definition advanced by the NSF specifically for cyberinfrastructure seems to be limited to the democratization of AI use. This makes sense, and it is important to take into account the scope of that definition, which is explicitly concerned with reducing inequalities between the degree of access that different groups of scientists have to AI technology to help them conduct their own research. For the topic of stakeholder participation in AI research conducted within ICICLE, the question of democratization must be broader. First, there are many different groups of stakeholders that are relevant to ICICLE's research—with two of the biggest and most clearly different groups being computer scientists and software developers who have the expertise to use ICICLE components to develop their own software products, or want to develop technologies that make use of ICICLE-enabled CI on the one hand and the various groups of stakeholders in the areas of use-inspired science who often have little technical background and resources, and may use ICICLE solutions in their own work, or may be affected by them without themselves being users. Second, the way in which these groups relate to ICICLE varies widely, so that different combinations of the four dimensions of democratization outlined above may be relevant for any given group of stakeholders and concrete situation.

With its variety of different research projects and types of stakeholders, ICICLE serves as a good illustration of the complexity and context-dependent nature of how AI research is related to the social environments to which it is directed and which are affected by its deployment. For efforts to democratize AI research within ICICLE and beyond, we think that it is important to be aware of the different dimensions that can be relevant to this endeavor, but not to attempt to reduce this complexity to any simple and unified definition of 'democratization' in advance. What really matters and how it can actually be achieved depends on the concrete situation and the needs and perspectives of the stakeholders involved, and cannot be settled in advance, but an awareness of the variety of different factors that can become relevant can make it easier to recognize what matters in any particular situation.

The second question concerns the methods by which democratization is to be achieved. A growing trend to involve stakeholders in the development of AI technologies has been dubbed the "participatory turn" and some of its features have been mapped out. Delgado et al. (2023) note that there are different ways to describe the varying levels of participation that are implemented, e.g. as "five levels of participatory mechanisms, from inform, consult, involve, collaborate, to empower", or as "four levels of participatory AI: consultation, contribution, collaboration, and co-creation" (2). The literature on participation in AI research draws on a wide range of academic disciplines and their methodologies called "user-centered design", "service design", participatory design", "co-design", "value-sensitive design", "participatory action research", "social choice theory and mechanism design", "participatory democracy and civic participation", and "deliberation theory" (Delgado et al. 2023: 2-3). This diversity of methods and levels of participation even corresponds to differences in how the overall aims of participation are conceived. Sometimes participation is rendered merely as a method for improving the technical solutions developed, but often participation is conceived of as an explicitly ethical and/or political question, arguing "that participation can enable AI systems to better reflect the values, preferences, and needs of users and other impacted stakeholders, or more broadly, that participation will empower stakeholders in shaping the design of AI systems" (Delgado et al. 2023: 1).

Despite the widespread call for participation and the impression that it "implicitly conveys" "implications for empowering stakeholders", Delgado et al. find that "the current state of participatory AI speaks to a largely unaccounted for and unaccountable heterogeneity that lacks a shared ethos or set of principles" (2). In addition to the conceptual heterogeneity of the literature, "AI researchers and practitioners struggl[e] with a tension between their participatory ambitions and the practical constraints they face on the ground" (Delgado et al. 2023: 12). Birhane et al. (2022) argue that in light of challenges like these, "a contextual and nuanced understanding of the term as well as consideration of who the primary beneficiaries of participatory activities ought to be constitute crucial factors to realizing the benefits and opportunities that participation brings" (ibid. 1). Since a key part of participation involves the "spread knowledge about technical systems and their impacts" (ibid. 6), "involv[ing] the knowledge of technical experts, but also the local knowledge embedded in the lived-experience of communities" (ibid. 7), what is needed are "participatory approaches that occur at different time frames, over various duration, and with different groups" (ibid. 7). Despite the potential for genuine progress and empowerment of stakeholders, Birhane et al. (2022) also note that "[p]articipatory activities convened by private actors or parallel institutions, cannot stand in for democratic politics, and participatory AI should not aspire to do so or be perceived to meet this function" (6).

The notion that "participation is not the best mechanism for decisions/values/norms that are better decided and codified by democratic institutions, governance and laws" (Birhane et al. 2022: 6) is echoed by Himmelreich (2023), who develops a general call "against broadening and deepening public participation in the governance of AI" (Himmelreich 2023: 1333). While the ultimate upshot of Himmelreich's argument is that "AI should be democratized not by broadening and deepening participation but by increasing the democratic quality of the administrative and executive elements of collective decision making" (1333), along the way three questions come into view that are useful for thinking about democratization in the context of ICICLE:

1. "the democratization of *what*?"
2. "*Why* should AI be democratized?" and
3. "*How* should AI be democratized?" (Himmelreich 2023: 1344)

Even though we do not share Himmelreich's pessimism about stakeholder participation as a tool for democratizing AI, we agree that these are the most important questions to ask. Similarly, some of his concerns about democratization are well-founded and apply to our project as well, for example his point that "democracy is simply expensive in terms of material, cognitive, cultural, informational, and social resources" (Himmelreich 2023: 1341). The costs of participation mean that in practice, not everyone has equal opportunity to participate, even when they are formally included and invited to participate. This issue is discussed in the literature on participation in two ways: First, there is a danger of conflating participation with inclusion, which is a problem because inclusion can never be complete and always leaves someone excluded: "Typically those excluded are the very worst-off, those with low literacy, those who do not have the time to seek out participatory opportunities, are not members of the right networks, etc." (Birhane et al. 2022: 6). Second, in practice the way to deal with the cost of participation is often to select individuals who are treated as representatives of larger groups. Problematically, "these approaches often rely on the assumption that stakeholders acting as a proxy may be able to adequately speak on behalf of the preferences for other stakeholders in a similar role (e.g., other teachers or other museum guides), or for stakeholders in roles they work closely with (e.g., students), or for other members of the same sociolinguistic or demographic group" (Delgado et al. 12). However, "participatory methods cannot rely on simplified assumptions about the reasons people have for engaging in a participatory process. This returns to the need to challenge uneven power distributions and oppressive social structures, as well as the ways that 'community' itself can hide power dynamics" (Birhane et al. 7).

Considerations such as these are important when designing participation processes, and seem particularly relevant for ICICLE when considering e.g. the concept of a Stakeholder Roundtable. To ensure that the roundtable can be an effective tool for democratization, the above three questions from Himmelreich need to be addressed at the outset, and in answering the third question, special attention has to be paid to who gets to participate in the roundtable and which groups they can be taken to represent to what degree. We believe that design practices from ML research itself can provide a useful analogy for conceptualizing this problem of selecting participants from larger groups of stakeholders, in particular different methods for data sampling.

The broader motivation for exploring synergies between questions of representation and participation in democratization and the different data sampling methods used in machine learning is our conviction that the activities involved in doing more democratic AI research and in doing better AI research overlap to a substantial degree. Instead of thinking about democratization as an external goal that creates further work which researchers are tasked with in addition to the AI research they are already doing, we should think of stakeholder participation as an integral part of the AI research itself. This view dovetails with the inclusion of ethical criteria in formal accounts of theory choice, where one feature of a "good theory" is that it minimizes "harm, the reckoning with how theory is forged in a fire of historical, if not ongoing, abuses—from past crimes against humanity, to current exploitation, turbocharged or hyped by machine learning, to historical and present internal academic marginalisation" (Guest 2023: 3). We believe that it may help integrate ethical concerns in engineering environments if the logic underlying such problems as selection of representative participants can be explained in analogy to known technical methods in the research domain.

For the comparison of sampling methods and participant selection, a discussion could begin by going through some varieties of "nonprobability sampling", such as "convenience sampling", "snowball sampling", "judgment sampling", and "quota sampling" (Huyen 2022: 83) and asking what each method means when applied to human participants in stakeholder engagement processes. For the ICICLE Stakeholder Roundtable or similar formats with constraints on the number of participants, the concepts of “stratified sampling” and “weighted sampling” can help us think about how different groups of stakeholders differ from each other, which groups are more important to include and what weight the concerns of different groups should carry—think for example of the difference between the people affected by a technology and the people who profit from its deployment (Huyen 2022: 84, 85).

5. Recommendations

Since one of the primary constraints on the democratization efforts within ICICLE is the time commitment required from researchers, we recommend to focus first on developing measures and resources that are already present within ICICLE or have already been planned. The two examples which we think are most important and want to highlight are the internal expertise and success stories in stakeholder engagement within ICICLE’s team and the Stakeholder Roundtable, which ICICLE has committed itself to establishing already but which is still in the planning stage at present.

5.1 Leverage Internal Expertise and Success Stories

Among the things we learned during our project, one of the most striking to us was the depth and breadth of expertise concerning stakeholder engagement that is already present within ICICLE. We believe that the methodologies that have been used in past projects (participatory action research, group model building) and informing the plans for future endeavors (participatory design) in the domains of smart foodsheds and computational ecology are invaluable resources that should be shared and propagated within ICICLE. One challenge that we see here is that the opportunities and avenues for stakeholder participation seem more obvious in the domains of use-inspired science (UIS) than they are for the majority of more foundational CI research, where the direct stakeholders of e.g. the development of a component by computer scientists are other computer scientists. Even though it may be less obvious how stakeholder participation plays a role in these projects, we believe that it is just as important there as it is for use-inspired science, and the development of a reference architecture for ICICLE illustrates why that is the case: Just as in the more obvious examples of UIS, the establishing of common standards and definitions for technical aspects of software components also involves communication between various groups that need to gain a better understanding of each other's perspective in order to coordinate their needs and capacities, and agree on a common meaning of key terms. In each case, the participatory methods for achieving these aims can contribute importantly to democratizing the research process and its outputs.

We believe that a lot of deep stakeholder engagement has already happened or is being planned within ICICLE, but that these efforts are not always visible from the outside or even from all vantage points within ICICLE itself. Our recommendation for leveraging existing resources and best practices is then to:

1. make visible and accessible existing expertise and success stories and
2. translate proven methods for stakeholder participation to new contexts.

The second point could take the form of 'pilot projects', which ideally become success stories of their own and feed back into the first point. One such translation project could be to conceive of the work necessary to define and communicate the reference architecture for ICICLE as a process involving various groups of stakeholders at different stages and to leverage some of the methodologies described above in order to facilitate this process.

5.2 The Stakeholder Roundtable

As described above in section 3, we identified the ICICLE Stakeholder Roundtable as an excellent potential mechanism for democratizing AI. In the rest of this subsection, we summarize our recommendations for what the stakeholder roundtable or roundtables should look like, but we have one overarching suggestion: make it happen!

Our first recommendation for the stakeholder roundtable is to establish organizational constraints. It should be decided where the roundtable should live within ICICLE. For instance, Is the Roundtable part of an existing thrust or project? Possibly, it could be organized within the broader impacts backbone network. Another important organization constraint that needs to be articulated is who will make the final decisions about the roundtable, and how will they go about making these decisions. Someone should also be established as the person who sends the emails, convenes the roundtables, and designs the meetings. We are thinking of this role as the Roundtable Organizer. One of our goals for the all-hands meeting was to identify who this person would be, but we have not yet had a volunteer.

We also have a list of recommendations for what we think the roundtable itself should look like given our experience during our fellowship, feedback from the UC PEWS workshop, and feedback from other ICICLE members at the all-hands meeting. Our first recommendation is that the mission of the roundtable should be clearly articulated. Here is our recommendation for this mission:

Mission: The ICICLE Stakeholder Roundtable is a vehicle for democratizing AI through specific instances of stakeholder participation with the goal of making ICICLE accountable to its stakeholders.

In addition to the mission of the roundtable, we also have recommendations for what its functions should be. First, it should identify concrete needs that stakeholders have. What are the real problems that they need ICICLE to address for them? Additionally, the roundtable should work to identify opportunities for the adoption of ICICLE's software services. ICICLE experts present at the roundtable could discover use-cases that stakeholder might not have recognized given their relative lack of understanding of ICICLE products' capabilities. Second, the roundtable could be used to match resources by identifying what different use cases could share implementations or solutions with each other. Without the roundtable, these potential cases of resource matching might not be visible. Third, the roundtable should also serve as a review for what was proposed at the last roundtable.

Given the structure of the Plale diagram, and useful feedback we received at the all-hands meeting, we also think it would be helpful to convene two separate Roundtables, one for adopters and one for users. Adopters, given their relatively high amount of technical knowledge regarding ICICLE's products, would be capable of asking hard, specific questions about the ICICLE software's architecture and capability. Discussion around and answers to these types of questions would likely not be useful to 'users' given their relative lack of technical knowledge.

We suggest, then, based on the two-roundtable design, the following features of the users roundtable. This roundtable should be populated mostly by external stakeholders and ICICLE members. We are thinking, for instance, of formal representatives of groups of stakeholders. This could be something like a representative from an association of farmers who plan to use ICICLE's digital agriculture products, but who are not themselves technical experts. At the same time, we recommend making sure that the roundtable's membership is representative of the diversity of users. It would not be ideal if all of the farmers at the round table happened to be dairy farmers, leaving out crop farmers, or big farm owners, leaving out small farmers. Additionally, the users roundtable could include representatives of entities with specific purposes for contributing to AI democratization. We are thinking, for example, of NGOs or similar advocacy groups.

In addition to a User's roundtable, we suggest a second roundtable tailored to the needs of the "adopters" group of stakeholders. For this group, we imagine members would be largely internal stakeholders, such as computational ecologists involved in the production of machine learning models. Ourselves being less technically apt, our suggestions are based on feedback we've received about the types of topics that would be discussed at an adopters roundtable. Some topics we imagine would be useful to discuss are reference architectures, and alpha, beta, and production priorities. Additionally, this roundtable would be well-suited for discussion about what training materials should be produced by ICICLE for its adopters or more technical users. In general, we see this roundtable as especially apt for discussion workforce development.

Having given these recommendations for what the stakeholder roundtable should look like, we end on some recommendations for what the next steps in putting together the roundtable could be.

1. Decide whether there will be one roundtable, or the two we have recommended, or more.
2. Decide what the missions and functions of the roundtables will be.
3. Identify an 'organizer' for each roundtable.
4. Gauge interest of candidate stakeholders and ask their preferred frequency of meetings.
5. Organizers collect feedback and devise meeting details to be signed off on.

References:

Berger-Wolf, Tanya (2023), "Trustworthy AI for Biodiversity: Challenges and Opportunities", presentation slides.

Birhane, Abeba, William Isaac, Vinodkumar Prabhakaran, Mark Diaz, Madeleine Clare Elish, Iason Gabriel, Shakir Mohamed (2022), "Power to the People? Opportunities and Challenges for Participatory AI" *EAAMO '22: Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization,* Boston, MA.

Delgado, Fernando, Stephen Yang, Michael Madaio, and Qian Yang (2023), "The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice", *EAAMO '23: Equity and Access in Algorithms, Mechanisms, and Optimization,* Boston, MA.

Koh, Keumseok, Rebecca Reno, and Ayaz Hyder (2018), "Designing an Agent-Based Model Using Group Model Building: Application to Food Insecurity Patterns in a U.S. Midwestern Metropolitan City", *Journal of Urban Health,* 95: 278–289.

Himmelreich, Johannes (2023). Against "Democratizing AI". *AI and Society* 38 (4):1333-1346.

Hollander, Allan D., Casey Hoy, Patrick R. Huber, Ayaz Hyder, Matthew C. Lange, Angela Latham, and James F. Quinn (2020), "Toward Smart Foodsheds: Using Stakeholder Engagement to Improve Informatics Frameworks for Regional Food Systems" in *Annals of the American Association of Geographers,* 110(2): 535–546.

Hovmand, Peter S. (2014). *Community Based System Dynamics,* New York, NY: Springer.

Huyen, Chip (2022). *Designing Machine Learning Systems.* Beijing: O'Reilly.

Plale, Beth, Sadia Khan, and Alfonso Morales (2023), "Democratization of AI: Challenges of AI Cyberinfrastructure and Software Research", *2023 IEEE 19th International Conference on e-Science (e-Science).*

Lafont, Cristina (2019). *Democracy without Shortcuts: A Participatory Conception of Deliberative Democracy.* Oxford: Oxford University Press.

Guest, Olivia (2023), "What Makes a Good Theory, and How Do We Make a Theory Good?", *Computational Brain & Behavior.*

Hyder, Ayaz, Angela Blatt, Allan D. Hollander, Casey Hoy, Patrick R. Huber, Matthew C. Lange, James F. Quinn, Courtney M. Riggle, Ruth Sloan, and Thomas P. Tomich (2022), "Design and Implementation of a Workshop for Evaluation of the Role of Power in Shaping and Solving Challenges in a Smart Foodshed", *Sustainability,* 14(5): 2642.

Parashar, Manish (2022). "Democratizing Science through Advanced Cyberinfrastructure", *Computer* (Long Beach, Calif.) 55 (9): 79-84.

The OSU (2023). "ICICLE: Intelligent CI with Computational Learning in the Environment". https://icicle.osu.edu/.