AI and Ethics (2026) 6:407
https://doi.org/10.1007/s43681-026-01242-9

RESEARCH

Curated information frameworks and the Governance of artificial intelligence

Emily LaRosa[1]

Received: 15 April 2026 / Accepted: 22 June 2026

Abstract

In this work, I introduce the Curated Information Framework (CIF) as a participatory model for governing artificial intelligence (AI) systems. I argue that meaningful AI governance must be understood as integrating practices of resistance to technological determinism, reclamation of community authority and sovereignty in technological use domains, and reimagining of a collaborative future rather than disclosure and use alone. Using agricultural AI as a central case, I show how prevailing transparency regimes often reinforce technocratic power by releasing information in forms that remain inaccessible or misaligned with the lived realities of those most affected. Drawing on participatory agricultural research and democratic epistemology, I reconceive transparency as epistemic curation: the role-specific, participatory, and iterative organization of information to support shared understanding across stakeholders. Through analysis of a case study involving an autonomous apple-picking system, I show how CIFs can redistribute interpretive authority among engineers, growers, farmworkers, and regulators, enabling communities not only to understand AI systems but to contest their terms of evaluation and deployment. I argue that such frameworks resist epistemic injustice, refuse extractive models of expertise, and reclaim governance authority for communities whose labor, land, and knowledge are entangled with AI systems. In doing so, I offer a practical and normative account of how AI governance can be re-imagined as a democratic, context-sensitive, and community-embedded process.

Keywords Artificial intelligence · Local knowledge · Governance framework · Knowledge reclamation

1 Introduction

As artificial intelligence (AI) systems become increasingly embedded in agricultural, ecological, and social life, the central question of governance is not simply how much information institutions disclose, but who has the power to define what counts as knowledge, whose understanding matters, and who is authorized to interpret technological risk. Both participatory agricultural research (PAR) and Curated Information Frameworks (CIFs) begin from the recognition that dominant, top-down models of expertise often reproduce epistemic injustice by treating affected communities as passive recipients of knowledge rather than as agents capable of shaping it. In agroecology, this problem appears in the traditional pipeline model of research, where scientific expertise is privileged over farmers' lived, local, and (often) intergenerational knowledge. PAR challenges that hierarchy by insisting that knowledge is situated, plural, and co-produced, and by directly involving farmers and local communities in defining research questions, methods, and applications. In doing so, it resists technocratic control and redistributes epistemic authority toward those most affected by agricultural change.

This same struggle over power animates the governance of current AI systems. In dynamic, embodied systems such as agricultural AI, naive transparency can itself become a form of exclusion when information is disclosed in ways that overwhelm, misalign with lived experience, or preserve expert control over interpretation. CIFs respond to this problem by reframing transparency as a participatory and democratic practice rather than a dump of technical information. Importantly, I am not trying to save transparency as a maximal ideal of disclosure; nor am I willing to throw away transparency in its entirety. Instead, I am rejecting transparency as a unitary, self-justifying ideal and

✉ Emily LaRosa
emily.larosa@uib.no

1 University of Bergen, Bergen, Norway

replacing it with more demanding questions that ought to be answered: what information should be shared, by whom, when, and for what purpose. The CIF is an attempt to displace naive transparency. I am arguing that while maximal disclosure is often useless or harmful, abandoning transparency altogether would give up on the important demand that institutions remain answerable to those affected by their systems. CIFs core structure focuses on role-specific information tiers, participatory construction, and iterative revision, and creates a mechanism through which disenfranchised communities can help determine what information matters, how it should be communicated, and how AI systems should be assessed over time. In that sense, CIFs do more than improve understanding: they function as a countermeasure to epistemic injustice by redistributing interpretive responsibility, respecting epistemic diversity, and enabling governance that is contextual, co-constructed, and oriented toward the reclamation of power by those whose lives are shaped by AI. Importantly, PAR provides the philosophical and political foundation for this project, while CIFs translate the needs of communities into a model of AI governance grounded in democratic accountability and shared authority over knowledge.

To develop this argument, I draw not only on PAR but also on democratic epistemology. Democratic epistemology helps clarify why inclusive and pluralistic knowledge practices are not merely ethically attractive but epistemically valuable: communities with diverse standpoints can identify risks, blind spots, and evaluative concerns that narrow technical perspectives may miss. These together reveal why prevailing AI governance models so often fail: they organize knowledge in ways that preserve existing asymmetries of authority and sustain patterned forms of ignorance about what communities actually need to know.

I develop these claims through two central moves. First, I underscore the philosophical foundations of PAR by exploring how agricultural knowledge can be organized in ways that challenge hierarchical expertise, foreground co-production, and treat communities as active participants in the governance of innovation, providing us a model for how participatory research in AI development can redistribute interpretive authority and reclaim knowledge sovereignty. Second, the paper applies these insights to agricultural AI through a case study involving creating a CIF related to an autonomous apple-picking system. This case illustrates how concepts such as reliability, error, success, and accountability do not carry a single self-evident meaning across stakeholder groups. Engineers may understand reliability in terms of measurable system performance, while growers may interpret it in relation to crop timing, economic predictability, and seasonal pressures; farmworkers may relate it to changing labor conditions, safety, and surveillance; regulators may frame it through compliance and liability. CIFs provide a way of mediating amongst these divergent but legitimate standpoints without reducing them to a single technical vocabulary and while respecting the situated knowledge of each knower.

The broader contribution of the paper is both practical and normative. Practically, it offers the CIF as a model for designing governance processes for dynamic AI systems that are more responsive to the realities of those who live and work with them. Importantly, CIFs are a tool in the toolbox towards a broader regulatory architecture. They are not themselves a complete regulatory methodology. Normatively, it argues that participatory information design can function as a form of resistance to epistemic injustice, a refusal of extractive models of expertise, and a means of reclaiming governance authority for affected communities. In this respect, the paper contributes to wider debates about how AI systems might be governed in ways that do not simply mitigate harm within existing structures of power, but instead transform the terms under which authority, accountability, and understanding are distributed.

Ultimately, meaningful AI governance ought to integrate practices of resistance to technological determinism as a path forward; to allow for reclamation of community authority and sovereignty in technological use domains; and to reimagine the future of AI governance as one that is collaborative and continuously iterative, in the hands of the communities which the evaluated system affects.

2 Participatory agricultural research: philosophical foundations

Participatory research is an approach to knowledge production that emphasizes collaboration, co-creation, and reflexivity. Rather than treating research subjects as a passive object of study, participatory research models involve them as co-creators and co-investigators in the research process. [9, 64] This approach is particularly well suited to agricultural practices and agronomy itself, where farming is deeply contextual, place-based, and dynamic. [6]

A core feature of participatory research is that it challenges the conventional distinction between experts and lay people. In traditional, agricultural scientific research paradigms, knowledge is assumed to be "created" in universities and laboratories, then passed along through a unidirectional pipeline until it reaches farmers in the field. This technology-transfer paradigm often discounts farmer contributions, focusing on technological adoption over situated understanding. Such an epistemic structure helps explain why these approaches can lead to compounding environmental challenges. When research and intervention design exclude farmer's experiential and ecological situated knowledge, agricultural challenges are more likely to be framed as discrete technical problems requiring scalable inputs (such as pesticides) rather than complex ecological dynamics. Resultingly, interventions may fail to account for local variation and long-term system effects. This has led to such challenges as the "pesticide treadmill." Discussed since the 1970 s, this treadmill is the phenomenon where insects evolve to become resistant to pesticides, leading to more application, higher concentration, and/or new chemicals being used, leading to still further resistance. [7] While the evolution of pests is a biological process, it is also shaped by governance structures which privilege standardized technological fixes over participatory, context-sensitive approaches. Farmers may possess early insights into resistance abilities and ecological imbalance, but in non-participatory research systems, such knowledge goes unincorporated.

In this sense, the persistence of the pesticide treadmill reflects not only ecological dynamics, but epistemic as well as institutional failure: a system of knowledge which marginalizes actors who may be the most attuned to the environments being managed. As Warner notes, "the [traditional] extension practice that helped create industrial agriculture is unlikely to be able to help it manage the environmental problems that result from this kind of farming." (Warner 2008:3) [65] This happens, in part, because it continues to rely on the same non-participatory assumptions that produced the initial problems.

This linear model is both epistemically flawed and ethically problematic, as it discounts the rich, context-dependent knowledge that farmers and growers accumulate through lived experience in the field. In contrast, participatory agricultural research complements agroecology by involving farmers and other stakeholders in every stage of the research process, engaging them in knowledge production. Recognizing farmers as vital custodians of agricultural systems, the PAR approach values local knowledge and integrates it with scientific research. Farmers collaborate with researchers to identify challenges, set research priorities, and participate in experimental design, implementation, and data collection, often conducted directly in their fields. This two-way exchange of knowledge allows for the opportunity to blend scientific insights with local knowledge, as well as potentially traditional ecological knowledge (TEK), equipping farmers with tools and skills to adapt and manage their systems effectively while valuing local and traditional knowledge as both legitimate and indispensable parts of the scientific process. [49, 67] While PAR does not necessarily have to involve TEK, it often benefits from incorporating it, depending on the context and goals of the research.

The iterative nature of this research ensures its relevance and practicality by incorporating continuous feedback from participants. [49] Key philosophical features of participatory agricultural research are (i) knowledge as a co-production process; (ii) democratization of expertise; and (iii) reflexivity in research practices through integrating and incorporating feedback from farmers and other stakeholders themselves. This differs from the extractive or hierarchical models we have discussed; in particular, the technology pipeline model. As we will uncover in Sect. 3, PAR allows us to have a much richer epistemic picture of agroecological findings.

The epistemic shift in PAR aligns with broader movements in feminist epistemology, indigenous studies, and science and technology studies (STS) that critique the supposed neutrality of, and universalizability of, scientific knowledge. [33] Donna Haraway (2013) argues that all knowledge is situated, meaning it is shaped by the social, historical and ecological contexts by which it emerges. [32] Participatory research embodies this idea by valuing the lived experiences of farmers, recognizing that the deep place-based knowledge they have offers insights which formal institutions often overlook. Haraway's account of situated knowledge supports this emphasis on partial, embodied and accountable perspectives, as it is objectively grounded in limited location and "views from somewhere". [32] Haraway also suggests that knowledge regimes can be studies through material-semiotic practices and field sites through which objects, boundaries and scientific authority are produced. [31] As such, site visits themselves become analytically rigorous and key to qualitative work, as is the case with this work.

Harding (1991) expands on this, advocating for epistemic pluralism, which suggests that multiple forms of knowing (including indigenous and experiential knowledge) ought to be integrated to achieve a more just and holistic understanding of the world. [33] In PAR, epistemic plurality challenges the conventional hierarchy that privileges scientific, over experiential, knowledge. The traditional pipeline model of agricultural extension, for example, treats farmers as passive recipients of technological innovation and scientific works; PAR replaces this with a framework that recognizes farmers as co-researchers and collaborators. [35] This shift acknowledges that the sustainability of farming systems relies upon not only technological advances but also the deep contextual knowledge of those working the land daily.

There is a stark contrast between objectivist epistemology, and this richer, situated, participatory knowing. Participatory research fosters new epistemic communities, where knowledge is co-produced; in the context of agriculture, participatory research foregrounds the knowledge of those most directly engaged with the land, acknowledging that farming expertise cannot be boiled down to controlled experiments or abstract models, but emerges through embodied, practical, and collective experiences. [52] As such, participatory research fosters new epistemic communities, where knowledge is co-produced, blending both lived experience and scientific inquiry. Such an approach to information organization is both a refusal of extractive informational asymmetry, as well as allowing the organization of information to act as a site of democratic engagement.

There is also the role of reflexivity to consider: participatory research is a continuous process, a cycle of learning, feedback and adaptation based on learning and feedback. This epistemically pluralistic framework not only enriches our understanding of agroecological processes, but also democratizes knowledge production, making it more inclusive and responsive to local needs.

The ethics of participatory research stem from its commitment to justice, autonomy, and empowerment of local and traditional actors; conventional agricultural research often perpetuates epistemic injustice, wherein the knowledge of farmers (in particular those from Indigenous communities, smallholders, and women) is systematically devalued. [65–67] PAR seeks to rectify this through ensuring that those most affected by agricultural challenges are actively engaged in defining research questions, methodologies, and applications, bringing an “emic (insider) perspective”. [53] In other words, PAR challenges these injustices directly, recognizing local knowledge as legitimate and valuable in the ethics of knowledge production.

PAR also addresses the role of power in research. Participatory models resist technocratic control, and encourage autonomy in agricultural communities. Such knowledge production is a collaborative effort where scientific, practical, and ecological concerns must align, requiring ongoing negotiation and adaptation to the specificities of place and practice. In following these calls for collaborative research, work, and respect for TEK, we may well make strides towards what Kyle Whyte (2016) calls respecting the “collective capacity” of indigenous peoples and landscapes. [66] In doing so, such participatory models of information-sharing grant authority over the meaning of information to groups other than the technocratic authority, making reclamation of governance authority possible (as will be discussed in Sect. 6).

Beyond epistemic justice, participatory research embodies an ethic of care: care for land, care for local communities, and for future generations. (Thompson 2017) By integrating ethical considerations into research design, participatory approaches align closely with the principles of ecological resilience, sustainability, and social equity. We ought not to treat governance and participatory research as something that can be captured by purely quantifiable outputs. This is a push against purely objectivist knowledge regimes, which are in tension with TEK.

PAR cannot survive in a vacuum. Institutional structures, policy frameworks, and long-term funding mechanisms are critical to sustaining such collaborative research initiatives. Warner (2008) highlights that, while participatory models empower farmers, their success depends on institutional backing that facilitates ongoing engagement and knowledge exchange. PAR plays a crucial role in shaping sustainable agriculture by fostering resilient farming systems based on local adaptation, biodiversity, and farmer autonomy. [28] Participatory models require institutional support, policy alignment, and frameworks that enable long-term farmer-researcher partnerships. [64]

Participatory Agricultural Research represents a paradigm shift in agricultural knowledge production, challenging traditional scientific hierarchies by incorporating farmers and stakeholders as co-creators of knowledge alongside researchers. This transition challenges the traditional paradigms of agricultural science, which have historically positioned researchers and scientific institutions as primary knowledge producers while relegating farmers and their communities to the role of passive recipients of scientific advancements, and gives us a foundational model for the refusal of extractive informational asymmetry in the face of AI adoption.

The broader significance of PAR stems from its ability to democratize knowledge, empower communities, and transform agricultural systems into more sustainable and just systems of production. Participatory agricultural research fosters agroecological knowledge, democratizes expertise, and embraces farmer-led innovation. By embracing epistemic plurality, fostering farmer involvement and autonomy, and integrating both scientific and experiential knowledge, PAR offers a pathway towards a more resilient agricultural system. Moving forward, a commitment to participatory research will be essential for building just and sustainable systems that prioritize collaboration, care, and ecological wisdom, particularly in the face of novel technological adoption such as Agricultural Artificial Intelligence. [62, 66]

3 Agricultural AI, or AgAI

Artificial Intelligence (AI) systems increasingly act not merely as informational tools but as dynamic systems embedded within physical and social environments. Such AIs adapt through retraining, respond to hyper-local conditions, and interact continuously with human users [50]. Additionally, they actively, iteratively reshape the environments they are designed to impact, sometimes to the detriment of those who engage with said environment. [19]

Agricultural AI (AgAI) provides a particularly revealing domain in which to examine these conditions and challenges. AgAI systems operate directly on land, crops, labor practices, and ecosystems. Their ethical and social impacts are highly context-dependent, varying by region, crop, and community. Moreover, these systems are embedded within complex social relations involving engineers, growers, farmworkers, regulators, and Indigenous or local knowledge holders. Local knowledge may be defined "as situated, practice-based expertise developed through sustained interaction with a particular ecological and/or social environment; it is often contrasted with generalized scientific knowledge in that it is hyper-specific, experiential, and context-bound" (LaRosa 2025:136; see also Thompson 2010; Ludwig et al. 2018). [41] [45, 60] As will be uncovered in the following section, AgAI requires coordination across diverse forms of expertise.

AgAI systems possess three features that distinguish them from many other non-physical AI applications, underscoring their dynamic nature and exemplifying the challenges that lay at the heart of attempting to govern such systems. Importantly, they are materially embedded in the world and co-evolve with it; are context sensitive; and are socially entangled. These features often overlap, and the presence of one feature can contribute to the presence of others. Current AgAI development practices ought to consider each of these features in turn, as well as collectively.

First, such systems co-evolve with their context of deployment. Unlike recommendation algorithms or large language models, AgAI systems act directly on physical environments, altering crops, soil, labor conditions, and ecosystems. However, the co-evolution here should not be understood solely as physical; the sociocultural becomes material and salient for for AgAI's co-evolution in the sense that traditions, institutions, and values become embedded in both our stewardship of land, and of technological systems themselves. Cultural norms and social policies leave a lasting impact on our AI systems, as well as the datasets they are trained on. In this way, AgAI systems exist in material contexts that are simultaneously physical and sociocultural, with their dynamic nature meaning they transform these contexts even as they themselves change. A governance challenge emerges from such materiality: as AgAI are co-evolving with their physical and sociocultural contexts, effects tend to have a spillover effect. Decisions made by an algorithm are influenced by, and in turn stand to reshape, policy and farming practices. Additionally, as AgAI becomes the norm in agricultural praxis, said systems can entrench particular labor or growing practices. This may generate path dependency, where (regardless of consequences of use) stopping the use of a dynamic system becomes costly and difficult. Governance approaches must therefore be able to anticipate and respond to downstream systemic consequences of current design and adoption decisions.

Second, such systems are context-sensitive. Performance and impact vary dramatically across regions, crops, and farming practices. A system optimized for one orchard may perform poorly, or dangerously, in another. Governance models that presuppose centralized, technological-expert-driven regulatory architecture becomes ill-fitted to the needs of AgAI, as such governance models lack the context-sensitivity needed to appropriately capture and integrate local knowledge and expertise. We have long seen the limitations of technocratic approaches called into reproach where lay knowledge and situated expertise are epistemically significant (see Longino, 2002; Bucchi, 2009; Friedman, 2019). [5, 23, 44] These limitations become increasingly important: should both local and hyperlocal conditions go unacknowledged, these systems risk technical as well as epistemic failure. For example, AgAI trained on generalized datasets risk misclassification of weeds or crops; mismanagement of irrigation for a unique block or orchard; or making recommendations inappropriate for the microclimate of the area. The issue is not merely technical failure in the familiar sense of "garbage in, garbage out"; rather, agriculture is a domain in which the conditions of successful action are partly constituted by situated ecological and practical knowledge. The adequacy of an AgAI system, therefore, cannot be determined solely at the level of generalized model performance; it must also be assessed in relation to the local conditions that define what competent performance is in the first place. AgaI exposes the limitations of technocratic governance: where local knowledge is constitutive rather than supplementary, centralized expertise alone is insufficient. AgAI is distinctive, not because it alone requires contextual calibration, but because the relevant context is both materially and operationally embedded in the object of governance itself.

Beyond technical failure, neglecting local knowledge and Traditional Ecological Knowledge (TEK) constitutes epistemic injustice. [67] [48] Miranda Fricker (2007) defines epistemic injustice as the wrong done to someone in their capacity as a knower, expanding on this in 2017 to encapsulate practices whereby marginalized groups are denied full participation in knowledge practices that impact them. [21, 22] In agricultural practices, when developers and regulators privilege abstract models over lived experience, discounting said lived experience as "anecdotal" due to prejudicial judgements, farmers and land stewards are denied credibility and authority in shaping the systems which impact their environments, leading to epistemic injustice. With both technical and epistemic failures in mind, a "one-size-fits-all" approach which ignores the context-sensitive needs of

AgAI (as it currently stands without just knowledge integration) may well lead to technical as well as epistemic failure events with catastrophic results. As such, granting authority over the meaning of information and making reclamation of governance authority possible becomes critical for just AI adoption.

Lastly, AgAI are socially entangled. These systems exist, not only in material systems, but also in sociopolitical relations, engaging with labor, land stewardship, and livelihoods. AgAI's context-dependency, coupled with their embeddedness, leads to distinct entanglement issues. Though this is an implication of the previous two points, social entanglement deserves to be acknowledged in its own right. AgAI systems operate within networks of growers, farmworkers, engineers, regulators, and communities whose livelihoods and knowledge systems are directly affected by technological changes. Through these systems' use, we see a blurring of public and private boundaries: though systems may operate on private farms, their visibility and knowledge about use shapes labor dynamics; their deployment may shape biodiversity; and said use may impact regulatory obligations related to emissions and environmental impacts. [47] [46] AgAI are not isolated technical tools, but are instead infrastructural actors, co-constituting landscapes, economies, and farming itself. Their public nature as actors creates an epistemic uncertainty: who is entitled to oversight, who bears responsibility, and who should participate in governance gets called into question. Fragmented oversight which ignores social entanglement lacks the epistemic and contextual sensitivity needed in dynamic, multi-actor systems (for further discussion on this point, see LaRosa (2025)). [41]

These three features, along with the above analysis, render one-size-fits-all governance models applied to AI systems inadequate at best. Local, user level governance is the focal point of the CIF; it is intended to help provide shared, locally-applied implementation rather than a sweeping regulatory framework. Governance must be capable of mediating amongst multiple epistemic communities with distinct priorities, vocabularies, and risk perceptions. It must also be materially reflexive, epistemically plural, and participatory: such governance ought to draw on local expertise as well as TEK alongside analytically rigorous methods to inform the approach to the system. Importantly, however, governance surrounding these systems must first address the concerns that arise when we attempt to have transparent governance of dynamic AI—and grapple with the challenges that arise when we have informational overload. Justifiable AI governance requires that any approach ought to be contextual and co-constructed. As such, we ought to turn to the concept of transparency, and to how curation of information can serve our goals.

4 Transparency beyond disclosure: curated information frameworks

Transparency is treated as an unqualified good in AI ethics: more transparency is assumed to yield more accountability, justifiable trustworthiness, and better governance. [20, 34, 42] This assumption underlies many regulatory proposals that emphasize too-broad an access to source code, datasets, or performance metrics. In another camp, we have those who argue that transparency is only sometimes a good; or, indeed, may be a red herring (see Creel (2020), Durán and Jongsma (2021), and Sullivan (2022)). [11, 15, 58] Similarly, we have many who have argued that explainability and interpretability of AI systems are not unqualified goods. [27, 37, 43] While these arguments are not necessarily couched in the language of transparency, they are largely making the same point that I am: naive transparency is neither achievable nor desirable in many AI contexts, and as such, we should be seeking something more.

Proprietary constraints, security risks, and sheer system complexity limit what can be disclosed. [24, 51] More importantly, even when disclosure is possible, it does not guarantee understanding. [3, 55] [69] Raw performance metrics, disclosures, or code repositories are often unintelligible to those most affected by AI systems, including workers and local communities. [29, 63] By providing all available information but withholding the tools necessary to interpret or understand said information, such transparency can foster a form of selective ignorance whereby the audience believes they have the information but lacks a coherent understanding of the subject itself (what Dotson might call third order epistemic oppression (Dotson 2014)) [17] [14] What is needed is not mere informational access, but qualitative forms of understanding, without which transparency can leave audiences formally informed yet substantively disoriented. [38]

Justifiable transparency, therefore, should not be understood as the maximal release of information, but as the provision of information that is adequate for the purposes of accountability, decision-making, and ethical evaluation. We ought to seek improved interpretive capacity by knowers. This shift reframes transparency as a relational and intentional practice, rather than an absolute ideal without additional context. Philosophical work on transparency emphasizes this point. Kevin Elliott's 2022 taxonomy identifies multiple dimensions of transparency (purpose, audience, content, mode, and danger), highlighting that transparency must be evaluated relative to who needs to know what, and why. [18] Elliott's taxonomy calls for audience-specific modes of disclosure while respecting interpretive capacity: regulators, community members, and operators of systems will need different kinds of information to make sense of the systems embedded in their material and cultural landscapes. Any form of rich knowledge, such as those affecting and being affected by dynamic systems, cannot be reduced to a single, universal model without risking what epistemic injustice, erasure or marginalization of ways of knowing embedded in material and cultural worlds. From this perspective, indiscriminate disclosure can function as a form of opacity, obscuring rather than clarifying ethical and practical stakes. [41]

Building on these insights, the Curated Information Framework (CIF) was created as a model for context-sensitive transparency in dynamic AI governance. CIF creation and implementation rejects the notion that transparency consists of exposing systems "as they are." Instead, this approach, as it is both a methodology and a governance tool, treats transparency as a practice of knowledge curation: the selective, structured, and audience-specific presentation of information designed to support meaningful understanding in a way that gives power back to the knower. Per Wilholt (2020), proper curation can facilitate knowledge acquisition, and enable a more informed decision. As such, we may see various forms of authority emerge, allowing for a more robust governance of AI systems. [68] CIFs operate via selective, purposeful information sharing, grounded in a commitment to democratic epistemology [13, 24, 61] Unlike naive or full transparency, curation (as proposed by Elabbar (2023)) emphasizes relevance, clarity, and usability that factors in a stakeholder's educational background and specific goals. [16]

A CIF is thus defined by three core features:

1. Role-specific information sharing, which aligns disclosures with the needs and capacities of distinct stakeholder groups;
2. Participatory construction, in which affected communities help determine what information matters to each stakeholder community and how it should be communicated;
3. Iterative revision, recognizing that both AI systems and social contexts evolve over time.

The CIF's strongest advantage is not *more participation*, but what happens *after participation*. It is better than existing methodologies, not because it is more participatory; but because it is better for AI adoption in that it converts participation into an ongoing communication infrastructure. Whereas participatory methods can surface lived experiences, local knowledge, and stakeholder values, the CIF goes further, asking how said knowledge becomes timely, role-specific, useable, and revisable guidance for those who are affected by and live alongside an AI system.

Participatory research already argues that communities ought to help define and analyze problems that affect them. Cornwall and Jewkes (1995) define participatory research as a shift away from extractive expert-led inquiry, while Chamber's (1994) Participatory Rural Appraisal (PRA) emphasizes methods that allow local communities to share, analyze, plan and act on their own knowledge. [8, 10] Participatory AI literature has identified a major weakness in current participatory practices: participation is often vague, uneven, or merely symbolic. Birhane et al. (2022) warn that participatory AI can suffer from unclear expectations, co-opting, and conflation with weaker forms of consultation; Delgado et al. (2023) argue that AI design lacks clarity about what meaningful stakeholder agency truly entails; and Groves et al. (2023) find that participation in commercial AI labs is often piecemeal and has little decision-making impact on the technology. [4] [12, 30] The CIF responds to these gaps by making participation legible as outputs and obligations: local lexicons, explanation profiles, public toolkits, interface prototypes, warnings, disclosure formats, and continuing revisions. Additionally, AI documentation frameworks such as datasheets for datasets may improve technical documentation, but it does not assist in public-facing AI adoption as thoroughly as the CIF: the CIF, rather than primarily documenting for technical or downstream users, asks how information should be curated for different audiences with different stakes, literacies, languages, and decision points. [26]

CIFs are explicitly grounded in a more democratic epistemology, which holds that knowledge relevant to collective decision-making is distributed across social groups rather than concentrated amongst technical experts alone. [29] By treating transparency as a curatorial, ongoing, community-embedded process, CIFs shift the focus from disclosure, to that of understanding. Transparency as a process through the implementation of a CIF therefore becomes a means of coordinating diverse forms of knowledge about a system, rather than merely a symbolic gesture of openness that evades responsible information sharing. Trustworthiness in AI governance becomes a relational achievement, accomplished via governance processes that affected and engaged communities see as responsive and revisable. Justifiable AI governance, should we seek to have such a thing, ought to be both contextual and co-constructed; as will be discussed in Sect. 6, CIFs allow for these criteria to manifest.

5 Case study: an autonomous apple-picking system

To illustrate the development and operation of CIFs in practice, we ought to examine a real-world 2024 case study conducted in partnership with AgAID. This section develops the Curated Information Framework (CIF) through an extended case study centered on the development and early-stage evaluation of an autonomous apple-picking system in the Pacific Northwest of the United States of America, developed by Dr. Joseph Davidson and his lab out of Oregon State University, and serves as a proof of concept for how curated, participatory transparency can function as an understanding mechanism for dynamic AI systems and provide people with the preconditions to enable just governance.

Traditional transparency approaches, such as releasing performance metrics or technical documentation, do little to bridge these interpretive gaps and stand to reinforce existing hierarchies of expertise. [40, 56] CIF development, and an in-vivo workshop, were therefore designed to mediate these differences.

The process unfolded in six phases:

1. Curated interviews to identify stakeholder concerns and informational priorities;
2. Pre-workshop surveys to assess baseline understanding;
3. Participatory lexicon-building, using a case study, around such contested terms as "failure," "risk," and "autonomy";
4. Collaborative, clear mapping of who needs to know what, when, and how, as well as potential dangers of sharing said information;
5. Publication of the resulting framework as an open, revisable document; and
6. Iterative follow-up as the system evolved.

The CIF methodology focuses on qualitative, rather than quantitative, output, as governance and participatory action cannot be effectively captured through quantitative measures alone. While the participatory orientation of the CIF aligns similarly to participatory design, the CIF is not intended as a design methodology. As a form of participatory governance, it operates adjacent to such methodology through using participatory tools, but its main function is to become an informational infrastructure for AI adoption. Where participatory methodologies focus on how knowledge is produced, the CIF focuses on how knowledge is communicated, updated, and made actionable during or after inquiry. The CIF translates stakeholder knowledge, technical system behavior, and institutional responsibilities into practical, role-specific guidance that can be reviewed as technologies, contexts, and user needs shift and change over time.[1]

The resulting CIF specified, for each stakeholder group, the form, content, and dissemination of relevant information. For example, farmworkers were determined to need safety-oriented explanations tied to observable system behaviors, whereas growers were determined to need economic and operational summaries.

5.1 Context and system description

The autonomous apple picker focused on by Davidson's lab integrates computer vision, robotic manipulation, and machine learning to identify, grasp, and harvest apples within commercial orchard environments. Unlike non-autonomous AI robotics, the highlighted dynamic system operates in highly variable field conditions, interacting with living plants, uneven terrain, and human laborers. Its performance is therefore shaped, not only by internal technical parameters, but by local ecological and social factors. During early evaluations, the following (hypothetical but realistic) failure scenario was presented to participants:

"You are working with an autonomous apple picker that has a dexterity and grip similar to that of a human apple picker. The apple picker is being used in an orchard block of jazz apples. The autonomous picker identifies an apple, attempts to grasp that apple, fails, and moves on to the next tree. The system left over half the ripe apples on the tree. The system does not offer a reason as to why it failed to pick the apples, nor a reason as to why it moved on to the next tree. A human picker is following the system, and picks the apples it does not. Given this situation, what would it mean for the autonomous apple picker to be/have X?"

Interestingly, the case study's system successfully navigates orchard rows and identifies ripe fruit, but fails to grasp more than half of the apples, leaving them on the tree without providing any explanation for its behavior. A human picker follows behind the machine, compensating for its missed harvest. Importantly, the human picker is not provided a rationale as to why the autonomous picker neglected to pick the fruit. This scenario deliberately foregrounds opacity at the level that matters most to end users: not algorithmic opacity per se, but the absence of interpretable reasons for action and failure. Early stakeholder assessments using this case study as a foil for contested terms revealed significant epistemic mismatches between stakeholders. One key example would be the difference in approaches to the notion of "reliability", as will be explored in Sect. 5.2.

[1] For a robust breakdown of the methodology, see Appendix (Curated information framework methodology).

5.2 Epistemic mismatches and the limits of Naive transparency

Initial stakeholder engagement revealed deep epistemic asymmetries across participant groups. This is thrown into sharp relief when examining definitional interpretations of system “reliability”, determined via anonymized interviews conducted by myself. Engineers tended to interpret system reliability in terms of measurable performance metrics such as success rates, or mechanical tolerances. One interviewee stated for a system to be reliable meant "performing at tolerance function". Growers framed reliability economically, focusing on yield consistency, labor substitution, and return on investment ("I mean, what is the ROI?" is a quote from one grower when questioned about reliability of the apple picker in question). Farmworkers, by contrast, understood reliability through embodied interaction: whether the machine moved predictably, respected spatial norms, and could be safely anticipated to act without harm to themselves or others ("Will it hurt me?" was a common refrain in the interview, applying not only to reliability but also trustworthiness). These tendencies were determined by myself.

Crucially, none of these interpretations were reducible to one another. It was predetermined that releasing raw performance data or technical documentation would not have reconciled these divergent understandings. As others who came before me have uncovered, such disclosure risks reinforcing agnotological failures, producing ignorance through overload, misalignment, or misplaced authority (see Bawden and Robinson (2009) and Schaffer et al. (2019)). [2] [54] This illustrates a central claim of the CIF approach: transparency without doing sufficient interpretive work is not only unsatisfactory but potentially harmful in deployment scenarios when it comes to socially entangled, dynamic AI systems.

5.3 CIF methodology: six phases of participatory information-sharing design

To address said epistemic mismatches, I implemented a six-phase CIF methodology, designed to elicit, organize, and iteratively refine role-specific informational needs. Stakeholder groups were determined via collaboration between myself and AgAID, as well as through a snowball sampling method.

Phase 1: Curated interviews

Structured interviews were conducted with stakeholders across the agricultural value chain, including engineers, growers, farm labor representatives, extension specialists, and researchers. Interview questions were co-developed with institutional partners and pared down to essential prompts to respect participants’ time while surfacing core informational priorities.

Phase 2: Pre-Workshop surveys

Participants completed surveys assessing baseline assumptions about who should share information about the system, at what stages of deployment, through which channels, and for which audiences. These surveys provided an initial map of epistemic expectations and revealed significant divergence across groups.

Phase 3: Participatory Lexicon-building

In a facilitated workshop, participants engaged in small-group discussions centered on the above apple picker failure scenario. Using a curated vocabulary list drawn from interview data, groups explored what terms such as “reliability,” “failure,” “autonomy,” and “accountability” meant in this specific context. The aim was not to fix definitions, but to surface how meanings shift across epistemic roles.

Phase 4: Collaborative information mapping

Participants then identified (i) what information each stakeholder group needs, (ii) when it should be shared, (iii) in what form, and (iv) by whom. Participants also (v) identified and articulated potential harms or dangers of sharing said information for each community group. This phase translated abstract concerns into concrete informational responsibilities, revealing that different groups require qualitatively different kinds of explanations rather than different quantities of the same data.

Phase 5: Framework synthesis and publication

The resulting CIF was collaboratively drafted in-vivo as an open, editable document specifying role-based information tiers. For example, farmworkers prioritized explanations of system behavior and safety protocols, growers emphasized operational reliability and economic impact, and regulators focused on environmental and labor implications. The framework was released for two weeks as an editable document to all participants, and then closed and published as an open-source resource to contributors through AgAID to encourage adaptation and iterative upkeep.

Phase 6: Iterative evaluation and revision

Post-workshop surveys mirrored the pre-workshop surveys, allowing the team to assess shifts in understanding and alignment. Although long-term iterative upkeep was limited by project timelines, the CIF methodology explicitly builds in protocols for ongoing revision as technologies and community tolerance for failure evolve, recommending an updating every three months with a reconvening of community parties and newly identified stakeholders annually.

5.4 Outcomes: situated transparency and epistemic infrastructure

The primary outcome of the CIF process was not consensus on system performance, but the construction of an epistemic infrastructure capable of supporting disagreement without misunderstanding. By epistemic infrastructure, I mean the organized social, technical, and communicative system that determines how knowledge becomes understandable, useable, and reviseable for affected participant communities. Participants reported greater confidence in their ability to interpret system behavior, articulate concerns, and identify when expectations had been violated. Importantly, the CIF did not eliminate opacity at the technical level. Instead, it redistributed epistemic responsibility: engineers remained accountable for system design, but the burden of interpretation no longer fell unjustly on non-technical stakeholders. In this sense, CIFs function as countermeasures to epistemic injustice, ensuring that those affected by AI systems are not excluded from meaningful understanding due to erroneous transparency practices. Critically, CIFs balance an openness of the information surrounding an AI system with intelligibility, rejecting the idea that disclosure without guidance is a justified governance practice. In this sense, CIFs function as countermeasures to epistemic injustice, ensuring that those affected by AI systems are not excluded from meaningful understanding due to erroneous transparency practices.

5.5 The apple picker as proof of concept

The autonomous apple picker case demonstrates that appropriate, understandable AI deployment depends less on maximal disclosure, and more on situated explicability. Information must be curated according to relevance, digestibility, and actionability, relative to the epistemic role of the recipient. When transparency is treated as a practice of democratic coordination rather than a full data release, it can support a Baierian form of trust between the deploying community and affected agents without demanding technical expertise from all participants (see Grasswick (2010) on earning epistemic trust with lay communities via knowledge sharing) [1, 29].This case therefore underscores that Curated Information Frameworks stand to provide a viable, ethically grounded mechanism for governing dynamic, embodied AI systems in context-sensitive domains such as agriculture. Situational transparency is transparency as a relational, contextual practice, where explanations must be tailored to the lived experience of stakeholders so they can understand, evaluate, question, and act on AI systems which affect them. By embedding appropriate, situated transparency within participatory processes, CIFs transform explanation from a technical artifact into a shared, contextual, co-constructed social practice of participatory information-sharing.

6 Participatory information-sharing design as resistance, refusal, and reclamation

Participatory, information-sharing design should be understood both as an improved method of transparency, and as a normative intervention into the distribution of epistemic authority in AI governance. It can function simultaneously as a resistance to epistemic injustice, a refusal of extractive models of expertise, and a means of reclaiming governance authority for affected communities. [14, 21, 29] In many governance settings, technical actors determine what counts as relevant information, which risks matter, and what forms of explanation are adequate, while affected communities are positioned as recipients of already-curated knowledge. Such an arrangement is not neutral; it privileges expert standpoints and reproduces epistemic injustice by denying communities standing as credible interpreters of the very conditions that shape their lives. [14, 21]

Participatory information-sharing design resists this structure by treating the organization of information itself as a site of democratic engagement. Rather than assuming experts first generate understanding, and then transmit it outward, it begins from the premise that meaningful understanding must be co-constructed across different social positions and forms of experience. This matters because transparency, by itself, does not guarantee understandability, uptake, nor allow for contestability. Institutions may disclose substantial information while preserving control over its framing and practical meaning. Under such conditions, affected groups may be told more while remaining unable to challenge what the information omits, what standards it encodes, or how success and failure are being defined. [41] Participatory information-sharing design resists this pattern by asking not only what information should be shared, but who gets to decide what is relevant, intelligible, and action-guiding.

This resistance is especially important in domains such as agriculture, where local, lived, experience-based knowledge has typically been subordinated to formal expertise. The CIF is intended to intervene at a local context, rather than attempting to change the regulatory landscape overall. The CIF is not intended as a tool to change the aforementioned landscape; like any tool, it will be subject to the limitations of its deployment and institutions that deploy said tool.

Growers, farmworkers, and local communities often possess detailed understandings of ecological variability, labor conditions, potential failures, and practical constraints that do not map neatly onto technical metrics. When AI systems are explained only through abstract performance measures, these forms of knowledge are not merely overlooked: they are displaced by narrower conceptions of what counts as legitimate understanding. Participatory information-sharing design counters this displacement by treating community knowledge as constitutive of governance rather than as optional feedback. In that sense, it acts as a resistance practice: it challenges informational orders that render affected communities visible only as end users, data sources, or risk bearers rather than as knowers.

At the same time, participatory information-sharing design functions as a refusal of extractive models of expertise. Extractive expertise draws on communities for labor, data, practical insight, or situational feedback without sharing corresponding authority over interpretation or decision-making. In such arrangements, community participation is valued instrumentally, insofar as it improves system performance, adoption, or legitimacy, but not insofar as it entitles participants to shape the goals, evaluative criteria, or governance structures of the system itself. [6, 65]. A participatory approach refuses this asymmetry. If communities are expected to contribute their situated knowledge to the development or assessment of AI systems, then they must also have a standing in shaping the informational environments through which those systems are understood and governed. Refusal here is structural, rather than merely oppositional: it rejects a model in which institutions absorb community knowledge while insulating themselves from community judgment.

Participatory information-sharing design also makes possible the reclamation of, or reassertion of a particular community's, governance authority. Governance authority is exercised not only in formal decision-making, but earlier, when categories are defined, risks prioritized, performance thresholds established, and informational interfaces structured. To reclaim governance authority, then, is not simply to gain a voice at the end of a process, but to participate in shaping the informational conditions under which judgments become possible in the first place. This is where the Curated Information Framework (CIF) becomes normatively significant. By organizing information in role-specific, participatory, and revisable ways, CIFs redistribute not only access to information but authority over its meaning. Engineers and regulators may still contribute technical and legal expertise, but they no longer monopolize the terms of interpretation. Growers, workers, local residents, and other stakeholders become participants in determining what counts as a meaningful explanation, which uncertainties deserve emphasis, what kinds of failure matter most, and how acceptable system behavior should be judged over time. [41]

CIFs do not, by themselves, democratize AI governance or redistribute institutional authority. Rather, they provide an epistemic and communicative infrastructure through which such possibilities can be made more practicable. In particular, the CIF translates complex technical information into usable, situated knowledge for affected agricultural communities; creates a structured process for interviews, surveys, workshops, lexicon-building, said CIF construction, and iterative updates; and helps stakeholders articulate what reliability, transparency, trustworthiness, safety, and accountability mean *to them* in relation to specific agricultural technologies.

CIFs should be understood as necessary but insufficient supports *when taken alone* for more accountable and participatory AI governance: though they can make agricultural technologies more legible, contestable, and responsive to affected communities, their capacity to alter institutional power depends on if they are embedded in procurement, design workflows, extension services, regulatory review, and ongoing accountability mechanisms. The limitation of the CIF is made clear within the broader governance constraints identified in digital agriculture scholarship, including adoption barriers, farmer identity and labor impacts, data ownership, privacy, ethics, agricultural knowledge systems, economics, and management (see Klerkx et al. 2019); fragmented and unclear agricultural data governance arrangements that may weaken farmer trust (see Jouanjean et al. 2020); and the need for responsible innovation approaches that emphasize anticipation, reflexivity, inclusion, and responsiveness without treating participation as a guaranteed democratic outcome while pushing for open deliberation (see both Stilgoe, Owen and Macnaghten 2020; and Gardezi et al. 2022). [25, 36, 39, 57] CIFs are therefore not a complete solution to the sociopolitical economy of AgAI, but a concrete mechanism for making transparency more situated, actionable, and institutionally usable.

The importance of the CIF approach is therefore not that they maximize data access, but that they generate a shared understanding of informational needs surrounding a system. Their primary outcome is not disclosure for its own sake, but rather improved interpretive capacity among participants. Stakeholders become better able to understand system behavior through understanding vocabulary used, better identify when performance diverges from local expectations, and raise concerns when said expectations are violated. This shift is normatively important because it moves governance away from passive reception and toward an active engagement. Communities are no longer asked simply to accept technical systems built elsewhere; they themselves participate in deciding what information matters, how it should be communicated, and how AI should be evaluated in a lived context.

CIFs produce this situated transparency by balancing openness with intelligibility. They do not reject disclosure; rather, they reject the assumption that disclosure alone is sufficient for trustworthy governance. They organize information in ways responsive to the differing roles, forms of knowledge, and practical stakes of those interacting with a system. In this respect, the CIF model offers a procedural pathway toward what NIST describes as "trustworthy AI," while grounding that aspiration in lived contexts rather than abstract standards alone. [59] Trustworthiness, on this view, is not merely a property verified through compliance, but a relational achievement made possible by governance processes that communities can recognize as responsive to their experiences and open to revision in light of their concerns.

The implications extend beyond agriculture. As AI systems increasingly mediate human relationships with their environments, regulators face a recurring problem: technical transparency without epistemic translation can erode rather than build trust. [16, 41] CIFs offer a replicable and ethically grounded response. They align the governance of intelligent systems with the situated interpretive capacities of those affected by them. In doing so, they show why justifiable and trustworthy AI governance must be contextual and co-constructed. It must resist epistemic injustice by challenging hierarchies of credibility and interpretation, refusing extractive models of expertise by demanding reciprocal authority, and reclaiming governance authority by enabling affected communities to shape the categories, standards, and informational structures through which AI systems are evaluated and contested.

7 Conclusion

Curated Information Frameworks offer a practical response to AI integration concerns by treating transparency as a participatory, role-specific, and iterative process of epistemic coordination. In doing so, CIFs do more than improve communication: they create conditions under which affected communities can contest, interpret, and help govern the systems that structure their environments and livelihoods. Read through the lens of Participatory Agricultural Research, CIFs can be understood as part of a broader effort to reclaim authority over situated knowledge from narrow technocratic institutions and redistribute it toward those most affected by technological change.

Their significance lies not only in making AI more intelligible, but in resisting epistemic injustice by affirming that farmers, workers, and local communities are not merely end users of technical systems, but knowers and governors of systems and their usage *in their own right*. CIFs provide an important mechanism within a larger democratic AI governance toolkit: one that centers local expertise, supports shared understanding, and makes room for ongoing revision as systems and communities evolve. The broader lesson of this work is therefore that trustworthy AI governance is not achieved by disclosure alone, but by building participatory institutions capable of aligning technical systems with the lived realities, values, and authority of the communities they affect. In pursuit of meeting these needs, CIFs are a uniquely well-situated tool in the governance toolbox of regulators and communities alike.

Curated information framework methodology

Creating inteviews

Step 1: collaborative question development with partner institutions

To ensure the interviews effectively address the specific informational needs of the partnering (redacted for peer review), we initiated the process by developing the interview questions collaboratively. This phase involved: Identifying Key Themes and Objectives: Working closely with the institution, we defined the core themes and objectives the interview seeks to address. By aligning our questions with these goals, we ensured that the interview results will yield actionable insights. Drafting Initial Questions: After defining objectives, we generated a list of potential questions covering the identified themes. These initial questions serve as a foundation for deeper refinement. Iterative Feedback and Refinement: We then shared the initial draft with institutional partners for feedback, refining the questions to ensure they are clear, relevant, and respectful of the interviewee's time and perspective. This collaborative approach ensured the questions were well suited to the needs of both the institution and the interviewee.

Step 2: paring down to essential questions

To keep interviews focused and efficient, we limited the question set to no more than 10 questions. This process involved: Prioritizing Core Questions: We identified the most essential questions, ensuring each aligns closely with the institution's objectives. Eliminating Redundancies: Any overlapping or tangential questions were removed, resulting in a streamlined set of questions that capture the necessary information concisely. Ensuring Flexibility: We structured questions to allow for elaboration, enabling participants to expand on their responses if they chose, while also accommodating time constraints.

Step 3: establishing initial contact and communication

Connecting with interviewees required a strategic approach, especially given the varied ways communities respond to outreach. We initiated contact through: Initial Outreach: We reached out to potential participants directly by phone and/or email, leveraging our partner institutions' many existing relationships with a variety of individuals. In a few cases where we could not rely on existing contacts, we also cold-called particular participant types. Snowball Sampling: Through our institutional network, and trusted community partners, we connected with community members who can vouch for the project. Snowball sampling allowed us to build rapport within communities where trust is paramount, such as farming networks.

For instance, the farming community required endorsement from a trusted insider before engaging with the interview process with (redacted for peer review). We acknowledge the support of key individuals, such as (redacted for peer review), as their assistance was instrumental in making interpersonal connections and fostering community trust.

Step 4: timing and interview duration

In designing the interview, we sought to be sensitive to participants' schedules and preferences. To maximize engagement while respecting participants' time, we implemented the following guidelines: Target Interview Length: We aimed for a 45 min interview duration. This length allowed for a thorough conversation without imposing too much on the participant's time. Flexibility for Longer Conversations: While the interview was designed to last no more than 45 min, we remained open to continuing as long as the interviewee is willing to share. This flexibility respected participants who may feel comfortable speaking at length on certain topics, ensuring they feel heard and valued.

Phase 2: conducting in-person interviews

Step 1: In-person interviewing for optimal engagement

To cultivate a conversational atmosphere, we prioritized face to face in person interviews wherever possible. This approach fostered a more relaxed environment, encouraging openness and engagement. Key considerations included: Zoom or Phone Alternatives: In cases where face to face interviews are not feasible, we offered Zoom or phone interviews. However, we acknowledged that remote formats may impact the conversational dynamic, and we adapted our approach accordingly. Adapting to Participant Needs: We were mindful of the individual needs of interviewees, understanding that factors such as familiarity with the interview topics and educational background may require adjustments. Translation services and simplified explanations were made available when necessary, ensuring the interview process was inclusive and accessible.

Step 2: documentation of responses

Accurate documentation is essential for capturing participants' insights authentically. Our approach to documentation includes:

Longhand notetaking: To remain fully engaged in the conversation, (redacted for peer review)documented responses manually. This approach allowed us to capture the interviewee's tone and mannerisms more naturally than a digital transcription might allow. Recording Direct Quotes: To preserve the authenticity of participants' voices, we included several direct quotes in our notes. This practice enhanced the richness of the data and provided the institution with an accurate reflection of participants' perspectives.

This methodology reflects a commitment to respectful, participant centered interviewing. By curating the process from question development to documentation, we ensured the interview experience was both informative for the institution and considerate of the participant's time and insights.

Phase 3: methodology for creating curated information framework

Step 1: initial framework outline

The process began with creating a rough outline for the information sharing framework. This outline provides a foundational structure for expansion and refinement based on qualitative input from motivators of the project (for our purposes, these were (redacted for peer review) as well as technologists). Key aspects to consider in this draft process ought to be:

Core areas of focus: We identified initial information sharing needs of the project, such as the partner agents that would be involved and affected by the technology, stages of technology development and deployment, and methods of communication for each partner group throughout the project life cycle. Guiding Principles: We addressed critical concerns regarding information sharing, including transparency, frequency of information sharing, and audience needs. Preliminary Structure: We outlined the roles that exist, as well as potential responsibilities various partners may have in the information sharing process. This created a flexible foundation that could accommodate adjustments based on stakeholder feedback. This initial outline served as a starting point, setting the stage for deeper refinement based on the input gathered through qualitative interviews (as were conducted in Phase 2).

Step 2: refining the framework through qualitative interviews

Building on the preliminary outline, the framework was expanded and adjusted according to the insights gained from informational interviews with key stakeholders. This phase involved:

Analyzing primary partner interviews: [2] We engaged stakeholders in interviews to gather their views on the framework's relevance and effectiveness. We then focused on open ended questions that allow participants to express their expectations, concerns, and priorities regarding information sharing in AI development (for more detail on this step, see Phase 2). Identifying Themes: We analyzed interview responses to identify recurring themes, patterns, and unique perspectives. For example, stakeholders may have specific insights into what information is critical at different development stages, or who should be responsible for dissemination. Incorporating Feedback: We integrated this qualitative feedback into the drafted Curated Information Framework, refining its structure and content to better align with initial stakeholders' expectations and the realities of the deployment context.

This step ensured that the framework reflected the collective expertise of the stakeholders, enhancing its practicality and relevance as a draft document prior to in-vivo editing.

Step 3: pre-workshop survey

To establish a baseline understanding of participants' perspectives on information sharing, a survey was sent to all interviewees prior to the workshop. This survey addressed core questions that explore various aspects of information dissemination:

Survey focus: The questions revisited core information sharing topics, including: (i) who should share information about AI; (ii) the stages at which information should be shared; and (iii) the target audiences for information, and (iv) preferred methods of dissemination. Survey Questions: Participants were asked to answer questions on: (i) who should be responsible for sharing information about AI; (ii) at what stage of development and deployment information should be shared; (iii) with whom the information should be shared; and (iv) how it should be disseminated. Baseline Data Collection: The survey responses provided valuable data on participants' initial perceptions and expectations, which can be compared to postworkshop results to assess changes in understanding or perspective.

The pre workshop survey results help tailor workshop content to participant needs and inform discussions on information sharing best practices.

Step 4: workshop facilitation and engagement

We conducted a collaborative workshop to engage all participants as partners in the framework's development. The workshop consisted of two primary parts:

Part I: developing a working lexicon through case study evaluation

Case study analysis: Participants were presented with a case study relevant to the AI technology being researched. This exercise served as a practical and concrete basis for discussing the types of information that should be shared, the appropriate terminology, and any existing gaps in communication. Lexicon Development: Through discussion, participants collaboratively developed a working lexicon, defining key terms and concepts that are essential to effective information sharing. This shared language enhanced communication clarity in the framework and fostered a unified understanding among all partners.

Part II: real-time framework development

Collaborative refinement: Participants actively contributed to refining the framework in real time, suggesting modifications based on their expertise and the insights gathered from the case study discussion. This collaborative process ensured the framework was well aligned with practical needs. Feedback and Adjustments: Facilitators captured real time feedback from participants, documenting suggestions and integrating them into the evolving framework. This iterative approach allowed for immediate testing of ideas and adaptations to better meet the needs of the group.

The workshop emphasized the role of each participant as a co-creator of the framework, fostering a sense of ownership and partnership in the project, as well as within the AI system. This was critical, as it may alleviate shelfware concerns for such systems in the future.

Step 5: post-workshop survey

Following the workshop, a second survey was sent to all participants to gauge whether the workshop influenced their understanding or perspectives on information sharing. This postworkshop survey was identical to the pre-workshop survey, allowing for a direct comparison of responses. Key aspects of the survey included:

Survey focus: The questions revisited core information sharing topics, including: (i) who should share information about AI; (ii) the stages at which information should be shared; (iii) the target audiences for information; and (iv) preferred methods of dissemination. Note: this focus is the same as the pre-workshop survey. Evaluating Impact: By comparing pre- and postworkshop responses, we assessed whether participants' understanding or attitudes toward information sharing evolved. This data will highlight the workshop's effectiveness in fostering alignment and enhancing the clarity of information sharing practices.

The survey results will inform any necessary adjustments to the framework, ensuring it remains responsive to participant needs and expectations.

Step 6: open source sharing and ongoing updates

The finalized framework was shared as an open source resource, accessible to all partners and relevant stakeholders. The goal is to provide a flexible, practical guide for information sharing that can be adapted to various AI deployment contexts. Ongoing updates ought to be planned to ensure the framework remains relevant as technology and community needs evolve:

Continuous revision: As the AI technology matures and becomes more embedded in the community, the framework can be periodically reviewed and updated using the same methodology, incorporating new feedback and addressing emerging challenges. Monitoring Tolerance Levels for Failure: Regular updates will take into account the evolving tolerance level for failure within the community, ensuring that the framework's guidance aligns with both stakeholder expectations and ethical standards for deployment. Open Access: By making the framework open source, we encourage wider adoption, adaptation, and feedback from a broader community, promoting a culture of transparency and collaboration in AI information sharing.

This methodology provides a structured, collaborative approach to developing and refining an information sharing framework that is both practical and adaptable, supporting responsible AI deployment through curated, meaningful communication.

Results: autonomous apple pickers: a case example of the curated framework

The effort to create a shared lexicon or dictionary of terms was accomplished through the use of the following case study:

You are working with an autonomous apple picker that has a dexterity and grip similar to that of a human apple picker. The apple picker is being used in an orchard block of jazz apples. The autonomous picker identifies an apple, attempts to grasp that apple, fails, and moves on to the next tree. The system left over half the ripe apples on the tree. The system does not offer a reason as to why it failed to pick the apple, nor a reason as to why it moved on to the next tree. A human picker is following the system, and picks the apples it does not. Given this situation, what would it mean for the autonomous apple picker to be/have X ?

Groups of five, with a total of 36 participants, began with brief introductions and identified the moderator; they then assigned a group reporter to summarize the discussion, as well as a notetaker. They read the case study above, and reviewed an assigned vocabulary list, pulled from the most commonly-used terms found in the pre-meeting interviews. Participants were then encouraged to discuss what each word meant to them personally. Importantly, this exercise is about exploring context and making terms more concrete, not solving a problem or conclusively defining terms. There are no right or wrong answers; rather, the aim is to foster participants to think deeply about the terms and what they mean in their own usage.

Participants then analyzed the words, reflecting on how each word's meaning might shift when applied to the case study. Participants were then encouraged to define each word in the context of the case. Groups were given 25 min to complete this work in small groups, and then encouraged to share their insights with the group and report back over the course of 45 min.

From here, the meeting moved on to crafting the final, group-approved iteration of the Curated Information Framework. Working off of the drafted CIF, created in response to both the interviews as well as pre meeting survey, groups of five worked together to (i) identify their audience category, and (ii) complete the further categories as they saw applicable. The work was done collaboratively in a small group, and then after 35 min, pulled back together to the larger group for in-depth analysis and discussion. The latter half of the discussion took 45 min to complete.

The meeting closed with all participants agreeing to the CIF as it stood, with an understanding that the final product would be sent out for assessment and commentary prior to any dissemination external to the working group. This process has now concluded, and the most contemporary version of the CIF as it relates to autonomous apple pickers can be seen in (redacted for peer review).

Acknowledgments I would like to acknowledge the support of AgAID for my core research. The AgAID Institute is a multi-institutional consortium of Washington State University (lead), Oregon State University, University of California Merced, University of Virginia, Carnegie Mellon University, Heritage University, Wenatchee Valley College, Kansas State University, and innov8.ag. AgAID is supported by USDA-NIFA by the AI Research Institutes program, under award No. 2021-67021-35344.

Author contributions Emily LaRosa is a the sole author of the attached manuscript. All authors reviewed the manuscript.

Funding Open access funding provided by University of Bergen (incl Haukeland University Hospital)

Data availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors declare no conflict of interest.

References

1. Baier, A.: Trust and antitrust. Ethics **96**(2), 231–260 (1986)
2. Bawden, D., Robinson, L.: The dark side of information: overload, anxiety and other paradoxes and pathologies. J. Inf. Sci. **35**(2), 180–191 (2009)
3. Ben-Shahar, O., Schneider, C.E.: More than you wanted to know: the failure of mandated disclosure. Princeton University Press, China (2014)
4. Birhane, A., Isaac, W., Prabhakaran, V.: Power to the people? Opportunities and challenges for participatory ai. In: Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization, pp. 1–8. (2022)
5. Bucchi, M.: Beyond technocracy: science, politics and citizens. Springer Science & Business Media, New York (2009)
6. Bursten, J.R., Kendig, C.: Growing knowledge: epistemic objects in agricultural extension work. Stud. Hist. Philos. Sci. **88**, 85–91 (2021)
7. Carolan, M.: The sociology of food and agriculture. Routledge, Oxfordshire (2016)
8. Chambers, R.: Participatory rural appraisal (pra): analysis of experience. World Dev. **22**(9), 1253–1268 (1994)
9. Chambers, R.: Rural development: putting the last first. Routledge, Oxfordshire (2014)
10. Cornwall, A., Jewkes, R.: What is participatory research? Soc. Sci. Med. **41**(12), 1667–1676 (1995)
11. Creel, K.A.: Transparency in complex computational systems. Philosophy Sci. **87**(4), 568–589 (2020)
12. Delgado, F., Yang, S., Madaio, M.: The participatory turn in ai design: theoretical foundations and the current state of practice. In: Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization, 1–23. (2023)
13. Dewey, J.: Democracy and education: an introduction to the philosophy of education, vol. 8. Macmillan, New York (1930)
14. Dotson, K.: Conceptualizing epistemic oppression. Soc. Epistemol. **28**(2), 115–138 (2014)
15. Durán, J.M., Jongsma, K.R.: Who is afraid of black box algorithms? On the epistemological and ethical basis of trust in medical ai. J. Med. Ethics **47**(5), 329–335 (2021)
16. Elabbar, A.: The curatorial view of assessment and the ethics of scientific advice: Beyond decisional autonomy towards distributive epistemic justice. International Union for the History and Philosophy of Science and Technology (IUHPST) (Essay Prize in the History and Philosophy of Science). Available online at iuhpst. org/media/pdf/Elabbar_IUHPST_Prize_2023. pdf (2023).
17. Elliott, K.C.: Selective ignorance and agricultural research. Sci. Technol. Human Values **38**(3), 328–350 (2013)
18. Elliott, K.C.: A taxonomy of transparency in science. Can. J. Philos. **52**(3), 342–355 (2022)
19. Fairfield, J.: Clean data: recursion as pollution in environmental AI. Oxford Intersections: AI in Society, Germany (2025).
20. Felzmann, H., Fosch-Villaronga, E., Lutz, C., et al.: Towards transparency by design for artificial intelligence. Sci. Eng. Ethics **26**(6), 3333–3361 (2020)
21. Fricker, M.: Epistemic injustice: power and the ethics of knowing. Oxford University Press, Oxford (2007)
22. Fricker, M.: Evolving concepts of epistemic injustice. In: Miranda Fricker (eds) The Routledge handbook of epistemic injustice. Routledge, Oxfordshire 53–60 (2017)
23. Friedman, J.: Power without knowledge: a critique of technocracy. Oxford University Press, Oxford (2019)
24. Galison, P., et al.: Removing knowledge: the logic of modern censorship. Agnotol: Making Unmaking Ignorance. 37–54 (2008)
25. Gardezi, M., Adereti, D.T., Stock, R., et al.: In pursuit of responsible innovation for precision agriculture technologies. J. Respons. Innovat. **9**(2), 224–247 (2022)
26. Gebru, T., Morgenstern, J., Vecchione, B., et al.: Datasheets for datasets. Commun. ACM **64**(12), 86–92 (2021). https://doi.org/10.1145/3458723
27. Ghassemi, M., Oakden-Rayner, L., Beam, A.L.: The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit. Health **3**(11), e745–e750 (2021)
28. Gliessman, S.: Agroecology and social transformation. Agroecology and Sustainable Food Systems, 38(10), 1125-1126 (2014)
29. Grasswick, H.E.: Scientific and lay communities: earning epistemic trust through knowledge sharing. Synthese **177**(3), 387–409 (2010)
30. Groves, L., Peppin, A., Strait, A.: Going public: the role of public participation approaches in commercial ai labs. In: Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency, 1162–1173. (2023)
31. Haraway, D.: Teddy bear patriarchy: taxidermy in the garden of eden, New York city, 1908–1936. Soc. Text **11**, 20–64 (1984)
32. Haraway, D.: Situated knowledges: the science question in feminism and the privilege of partial perspective 1. In: Women, science, and technology. Ed.: Wyer, M., Barbercheck, M., Cookmeyer, D., Ozturk, H., Wayne, M. Routledge, New York p 455–472 (2013)
33. Harding, S.: Whose science? Whose knowledge?: Thinking from women's lives. Cornell University Press, New York (1991)
34. Haresamudram, K., Larsson, S., Heintz, F.: Three levels of ai transparency. Computer **56**(2), 93–100 (2023). https://doi.org/10.1109/MC.2022.3213181.

35. Henke, C.R.: Making a place for science: the field trial. Soc. Stud. Sci. **30**(4), 483–511 (2000). https://doi.org/10.1177/030631200030004001
36. Jouanjean, M.A., Casalini, F., Wiseman, L., et al.: Issues around data governance in the digital transformation of agriculture: the farmers' perspective. In: OECD Food, Agriculture and Fisheries Papers (2020)
37. Kaur, H., Nori, H., Jenkins, S.: Interpreting interpretability: understanding data scientists' use of interpretability tools for machine learning. In: Proceedings of the 2020 CHI conference on human factors in computing systems, pp. 1–14. (2020)
38. Khalili, M.: Against the opacity, and for a qualitative understanding, of artificially intelligent technologies. AI and Ethics **4**(4), 1013–1021 (2024). https://doi.org/10.1007/s43681-023-00332-2
39. Klerkx, L., Jakku, E., Labarthe, P.: A review of social science on digital agriculture, smart farming and agriculture 4.0: new contributions and a future research agenda. NJAS: Wageningen J. Life Sci. **90**(1):1–16 (2019)
40. Kolkman, D.: The (in)credibility of algorithmic models to non-experts. Inf. Communicat. Soc. **25**(1), 93–109 (2022). https://doi.org/10.1080/1369118X.2020.1761860
41. LaRosa, E.: Just and trustworthy applications of dynamic artificial intelligence. In: Phd thesis, Michigan State University, East Lansing (2025)
42. Larsson, S., Heintz, F.: Transparency in artificial intelligence. Int. Policy Rev. **9**(2), 1–16 (2020)
43. Lipton, Z.C.: The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. Queue **16**(3), 31–57 (2018)
44. Longino, H.: The social dimensions of scientific knowledge. In: Zalta, E.N., Nodelman, U. (eds.) The Stanford Encyclopedia of Philosophy, Spring, 2026th edn. Stanford University, Metaphysics Research Lab, California (2026)
45. Ludwig, D., Poliseli, L.: Relating traditional and academic ecological knowledge: mechanistic and holistic epistemologies across cultures. Biol. Philos. **33**(5–6), 43 (2018). https://doi.org/10.1007/s10539-018-9655-x
46. Michaliszyn-Gabryś, B., Bronder, J., Jarosz, W., et al.: Potential of eco-weeding with high-power laser adoption from the farmers' perspective. Sustainability **16**(6), 2353 (2024). https://doi.org/10.3390/su16062353
47. Michaliszyn-Gabryś, B., Bronder, J., Krupanek, J.: Social life cycle assessment of laser weed control system: a case study. Sustainability **16**(6), 2590 (2024). https://doi.org/10.3390/su16062590. (https://www.mdpi.com/2071-1050/16/6/2590)
48. Noll, S., Murdock, E.G.: Whose justice is it anyway? Mitigating the tensions between food security and food sovereignty. J. Agric. Environ. Ethics **33**(1), 1–14 (2020). https://doi.org/10.1007/s10806-019-09809-9
49. Pagliarino, E., Orlando, F., Vaglia, V., et al.: Participatory research for sustainable agriculture: the case of the italian agroecological rice network. European J. Futu. Res. **8**(1), 7 (2020). https://doi.org/10.1186/s40309-020-00166-9
50. Peterson, J.: Context sensitive frames and ai alignment. In: 2025 IEEE Conference on Artificial Intelligence (CAI). IEEE, Santa Clara, CA, USA, p 1251–1254, (2025) https://doi.org/10.1109/CAI64502.2025.00218, https://ieeexplore.ieee.org/document/11050601/
51. Pieters, W.: Explanation and trust: what to tell the user in security and ai? Ethics Inf. Technol. **13**(1), 53–64 (2011). https://doi.org/10.1007/s10676-010-9253-3
52. Pound, B., Snapp, S., McDougall, C., et al.: Managing natural resources for sustainable livelihoods. Earthscan, London (2003)
53. Dey de Pryck, J., Elias, M.: Promoting inclusive facilitation of participatory agricultural research for development. Dev. Pract. **33**(1), 122–127 (2023)
54. Schaffer, J., O'Donovan, J., Michaelis, J., et al.: I can do better than your ai: expertise and explanations. In: Proceedings of the 24th international conference on intelligent user interfaces, pp 240–251 (2019)
55. Schnackenberg, A.K., Tomlinson, E., Coen, C.: The dimensional structure of transparency: a construct validation of transparency as disclosure, clarity, and accuracy in organizations. Human Relat. **74**(10), 1628–1660 (2021). https://doi.org/10.1177/0018726720933317
56. Sloane, M., Solano-Kamaiko, I.R., Yuan, J., et al.: Introducing contextual transparency for automated decision systems. Nat. Mach. Intell. **5**(3), 187–195 (2023). https://doi.org/10.1038/s42256-023-00623-7
57. Stilgoe, J., Owen, R., Macnaghten, P.: Developing a framework for responsible innovation. In: The ethics of nanotechnology, geoengineering, and clean energy. Routledge, p 347–359 (2020)
58. Sullivan, E.: Inductive risk, understanding, and opaque machine learning models. Philosophy Sci. **89**(5), 1065–1074 (2022). https://doi.org/10.1017/psa.2022.62.
59. Tabassi, E.: Artificial intelligence risk management framework. (AI RMF 1.0) (2023).
60. Thompson, P.B.: The agrarian vision: sustainability and environmental ethics. University Press of Kentucky, Lexington (2010).
61. Thompson, P.B.: From field to fork: food ethics for everyone. Oxford University Press, Oxford (2015)
62. Thompson, P.B.: The spirit of the soil: agriculture and environmental ethics, 2nd edn. Routledge, New York (2017) https://doi.org/10.4324/9781315559971
63. Waddell, T.F., Auriemma, J.R., Sundar, S.S.: Make it simple, or force users to read?: Paraphrased design improves comprehension of end user license agreements. In: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems. ACM, San Jose California USA, p 5252–5256, (2016)https://doi.org/10.1145/2858036.2858149,
64. Warner, K.D.: The quality of sustainability: agroecological partnerships and the geographic branding of california winegrapes. J. Rural. Stud. **23**(2), 142–155 (2007). https://doi.org/10.1016/j.jrurstud.2006.09.009.
65. Warner, K.D.: Agroecology as participatory science: emerging alternatives to technology transfer extension practice. Sci. Technol. Human Values **33**(6), 754–777 (2008). https://doi.org/10.1177/0162243907309851.
66. Whyte, K.: Indigenous experience, environmental justice and settler colonialism. Environmental Justice and Settler Colonialism (April 25, 2016) (2016)
67. Whyte, K.P.: On the role of traditional ecological knowledge as a collaborative concept: a philosophical study. Ecol. Process. **2**(1), 7 (2013). https://doi.org/10.1186/2192-1709-2-7
68. Wilholdt, T.: On knowing what one does not know: ignorance and the aims of research, 1st edn. MIT Press, Massachusetts p 195–218 (2020)
69. Xuan, Y., Small, E., Sokol, K., et al.: Comprehension is a double-edged sword: over-interpreting unspecified information in intelligible machine learning explanations. Int. J. Hum Comput Stud. **193**, 103376 (2025). https://doi.org/10.1016/j.ijhcs.2024.103376.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.