

Do Responsible AI Artifacts Advance Stakeholder Goals? Four Key Barriers Perceived by Legal and Civil Stakeholders

Anna Kawakami^{1*}, Daricia Wilkinson², Alexandra Chouldechova³

¹Carnegie Mellon University

²Arizona State University

³Microsoft Research

akawakam@andrew.cmu.edu, daricia.wilkinson@asu.edu, alexandrac@microsoft.com

Abstract

The responsible AI (RAI) community has introduced numerous processes and artifacts—such as Model Cards, Transparency Notes, and Data Cards—to facilitate transparency and support the governance of AI systems. While originally designed to scaffold and document AI development processes in technology companies, these artifacts are becoming central components of regulatory compliance under recent regulations such as the EU AI Act. Much of the existing literature has focussed on the design of new RAI artifacts or their use by practitioners within technology companies. However, as RAI artifacts begin to play key roles in enabling external oversight, it becomes critical to understand how stakeholders—particularly stakeholders situated *outside* of technology companies who govern and audit industry AI deployments—perceive the efficacy of RAI artifacts. In this study, we conduct semi-structured interviews and design activities with 19 government, legal, and civil society stakeholders who inform policy and advocacy around responsible AI efforts. While participants believe that RAI artifacts are a valuable contribution to the broader AI governance ecosystem, many have concerns around their potential unintended and longer-term impacts on actors outside of technology companies (e.g., downstream end-users, policymakers, civil society stakeholders). We organized these beliefs into four barriers that help explain how RAI artifacts may (inadvertently) reconfigure power relations across civil society, government, and industry, impeding civil society and legal stakeholders’ ability to protect downstream end-users from potential AI harms. Participants envision how structural changes, along with changes in how RAI artifacts are designed, used, and governed, could help re-direct the role of artifacts to support more collaborative and proactive external oversight of AI systems. We discuss research and policy implications for RAI artifacts.

Introduction

Technology companies are now developing and deploying artificial intelligence (AI) at unprecedented speeds and scales (Muthukrishnan et al. 2020). In turn, the responsible AI (RAI) community has rapidly innovated new RAI artifacts—processes, tools, and other resources designed to support the ethical creation and use of AI (Reisman et al. 2018; Mitchell et al. 2019; Madaio et al. 2020; Metcalf et al. 2021;

Gebru et al. 2021; Rakova and Dobbe 2023). The research community’s contributions around RAI Artifacts have so far focused on introducing new responsible AI artifacts (Adkins et al. 2022; Hupont et al. 2024), examining how they are used within technology companies (Deng et al. 2022; Balayn et al. 2023; Okolo and Lin 2024)), and interrogating the values embedded in their language (Sadek et al. 2024).

A subset of these artifacts, which we refer to as *Multi-Actor Responsible AI (RAI) Artifacts*, intend to support different stakeholders, beyond those typically represented in an AI development team, in understanding AI-related information. In some cases, these artifacts intend to support *direct* communication (e.g., dynamic conversation between or amongst different stakeholders), as in the case with compliance documentation such as Hupont et al.’s Use Case Cards¹. Other times, they intend to support *indirect* communication (e.g., communication of information to different stakeholders in isolation from each other), for example, in Mitchell et al.’s Model Cards². While a growing body of literature has emphasized the importance of multi-actor collaboration to support responsible AI design (Kaminski 2018), we have little insight into how well existing Multi-Actor RAI Artifacts foster the cross-sector communication essential for collaborative governance.

In this paper, we center regulatory and civil society stakeholders to explore the perceived efficacy of Multi-Actor RAI Artifacts. Despite their central roles in governing the deployment of AI systems, their perspectives have so far been under-examined in studies on the design and use of Multi-Actor RAI Artifacts. As regulators begin to repurpose RAI artifacts as compliance tools, and civil society rely on transparent model documentation to examine industry practices, we anticipate that understanding these stakeholders’ perspectives will become increasingly important. We therefore ask: **RQ1.** *How do legal/regulatory and civil society stakeholders perceive the efficacy of Multi-Actor Responsible AI Artifacts?* and **RQ2.** *What are implications for improving*

¹Use Case Cards are designed for policymakers, researchers, and practitioners to assess risk level and document use cases (Hupont et al. 2024)

²Model Cards are promoted as an artifact that supports communication of model-related information to AI practitioners and audiences outside of the technology industry (e.g. [journalistshttps://modelcards.withgoogle.com/about](https://modelcards.withgoogle.com/about))

*The author completed this work while at Microsoft Research. Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

the design and use of Multi-Actor Responsible AI Artifacts?

To address these questions, we conduct semi-structured interviews and design activities with 19 participants involved in informing policymaking and advocacy around AI systems deployed by technology companies. During the activities, participants had an opportunity to directly examine existing filled-in examples of existing Multi-Actor RAI Artifacts to ground their reflections. The participants are from two focal stakeholder groups : (i) legal experts with prior or current engagement in the scholarship or practice of AI regulation; and (ii) individuals working for civil society organizations concerned with the impact of AI on society. Participants held beliefs about the role that RAI artifacts should play in society (i.e., protecting end-users from AI harms) that inherently conflicted with the goals they perceived the technology industry have for using RAI artifacts (i.e., promoting self-regulation and shaping societal standards for AI). Participants' concerns surfaced how this underlying misalignment may ripple down to shape industry practices around the design and use of Multi-Actor RAI Artifacts. We organized participants' concerns into four barriers: (1) end-users being positioned as second-order priority, (2) selective showcasing of laudable AI models versus truly high risk models, (3) over-reliance on transparency as a change mechanism, and (4) offloading duties for responsible use to end-users. Our findings describe how each of these barriers map to different motivating factors, artifact design and use practices, and downstream impacts on the broader society.

Overall, participants were concerned that, if left unaddressed, these barriers may (inadvertently) reconfigure power relations across civil society, government, and industry, impeding RAI governance efforts from those outside of the technology industry. Participants desired RAI artifacts that could, instead, foster more proactive and collaborative cross-sector AI governance efforts. We extend existing calls for broader community engagement in directing the role Multi-Actor RAI Artifacts play, particularly in supporting responsible AI efforts between and within the regulatory, civil society, and industry sectors.

Background

The Responsible AI community has made fundamental strides in establishing core principles to guide the responsible development of AI. To instantiate these principles to practice, the community has developed new tools, processes, and resources. Throughout this paper, we use *artifacts* as an umbrella term for these innovations. In this section, we review prior work on the evolving role of RAI artifacts in supporting multi-actor communication. We then review prior work on the design and use of RAI artifacts.

The Evolution of Artifacts as a Communication Stream for Industry, Regulation & Civil Society

Artifacts have a longstanding history as vehicles to promote adherence to established values and principles critical to oversight, accountability, and overall governance—often in an inter-stakeholder capacity. In heavily regulated industries like healthcare, energy, manufacturing, aviation, and phar-

maceuticals, artifacts such as checklists, risk assessments, and standard operating procedures facilitate communication and information flow among multiple actors (e.g. workers, managers, and external auditors) (Weske et al. 2019, 2018; Ywakaim Andrawes, Faisal Fakhry, and Abd El Azeem 2019; Kagan, Gunningham, and Thornton 2011). These industries all have a shared history of adverse impacts affecting people, communities, and society-at-large. The impacts have materialized as inequitable resource allocation (Blodgett et al. 2022), the perpetuation of systemic bias against social groups (Mehrabi et al. 2021), and the unethical exploitation of vulnerable people to gain resources for the further enrichment of industry leaders (Noble 2018).

In response, public pressure and advocacy across sectors have forced reflection within companies, pushed civil society members to drive major advocacy campaigns, and prompted a need for policymakers to introduce industry regulation to manage risk and error. Science historian Lorraine Daston posits that the identification, reduction, and elimination of errors has been a steady driving force of modern science influencing collective sensemaking and overtime shaping errors “as sites for collaborative, collective, and coordination work” (Lin and Jackson 2023). To materialize this shared goal for risk reduction, artifacts have served as communication streams that encourage debate about power dynamics shaping the authority to define risks, influence public perceptions and normative standards for acceptable levels of risk, and forge new forms of partnerships by including multiple actors. In healthcare, checklists are used to indicate compliance with industry standards for patient safety and to produce documentation that allows third party auditing to affirm regulatory compliance needed for institutions to remain operational. Checklists are also used in architecture to avoid critical life-endangering errors (Bierska, Buhnova, and Bangui 2022). In manufacturing, risk assessments are commonplace and often required by law promoting both short-term and long-term reflection on adverse impacts. The artifacts encourage communication among teams to discuss how to identify risks, they assist civil society advocates in holding companies accountable, and they communicate a level of compliance with enforced regulation.

Likewise, social computing scholarship has reflected many parallels. Researchers have extensively reported on AI risks, errors, and failures (Domínguez Hernández et al. 2023; Ntoutsis et al. 2020; Dolata, Feuerriegel, and Schwabe 2022) and in response have created an abundance of RAI artifacts for assessment and mitigation, and to foster better communication between non-industry audiences who contend with the governance of AI. Unlike other industries, RAI artifacts have largely pre-dated AI regulation. Efforts to support the development of RAI artifacts have largely been centered on practitioners' needs. Although RAI artifacts are commonly promoted for non-industry audiences, there is limited work incorporating the views of non-industry actors involved in AI governance. Some industry stakeholders have asserted that their technical expertise and domain knowledge uniquely positions them to foresee risks and build guardrails to mitigate those risks (Chatterjee and Bordelon 2023; Levin and Downes 2023). Meanwhile, proponents of regulation

have been jostling to propose stricter government-led compliance and accountability to ensure industry stakeholders deliver on their promises of positive transformation and its ability to self-regulate (Ferretti 2022). In a similar vein, Civil Society Organizations (CSO) have driven impactful change by leading grassroots efforts with impacted communities, grounding policy agendas in lived realities, and influencing AI practice (Stjernfelt et al. 2020). However, between total government regulation and total self-regulation, prior work suggests that the consensus will more than likely operate from a collaborative stance to take advantage of (rather than compete with) the varying expertise and abilities from the three sectors (Kaminski 2018; Cohen 2023).

The Design and Use of Multi-Actor Artifacts in Responsible AI

In the past half decade, a growing community of responsible AI researchers and practitioners has advanced a range of RAI artifacts to help promote governance and increase transparency over development practices. For example, Accenture created the Fairness Tool to identify and mitigate bias; Google created Model Cards to support more transparent AI documentation; and Microsoft created Responsible AI Impact Assessments to help more proactively envision and mitigate downstream harms. This set of artifacts is expanding quickly, gaining widespread attention and use across AI application settings.

The intended purposes and audiences of RAI artifacts differ. A subset of artifacts—like FairLearn (Bird et al. 2020) and InterpretML (Nori et al. 2019)—are designed to support AI practitioners in quantitatively evaluating and improving AI systems along dimensions of common principles (e.g., fairness, explainability, accuracy). Another subset of artifacts—like Model Cards (Mitchell et al. 2019), Datasheets for Datasets (Gebru et al. 2021), and Responsible AI Impact Assessments (Metcalf et al. 2021)—are also designed to support the improved design of AI systems, but focus on scaffolding *documentation* of model information and *reflection* on potential risks. For instance, Model Cards (Mitchell et al. 2019) supports AI teams in documenting model properties including training parameters and model evaluation results, and Datasheets for Datasets prompts documentation of dataset details, including the motivation, composition, intended uses of the dataset (Gebru et al. 2021). This latter class of artifacts serve broader functions of organizational, and increasingly, regulatory and societal governance of AI systems. In fact, based on documented information presented with these artifacts (e.g., (Mitchell et al. 2019)), many of them are designed with an intention to be used by multiple different stakeholders, including auditors, policy-makers, end-users, and other impacted AI actors who typically may not be part of the AI team developing the model. We refer to this specific class of RAI artifacts—those intended to support (direct or indirect) multi-actor communication—as *Multi-Actor RAI Artifacts*.

As Multi-Actor RAI Artifacts begin to play an increased role in supporting the governance of AI systems, their uses within the RAI ecosystem have also expanded and evolved. For instance, recently introduced regulatory regimes have

pushed for the adoption of RAI artifacts by making completion of the artifact-supported process mandatory. The EU-AI Act and Algorithmic Accountability Act in the United States require that AI deployers complete an impact assessment prior to having a high-risk AI system enter the market (Yoo and Lai 2020; Novelli et al. 2023). State legislatures in the United States have followed, proposing new AI bills that would require the developers and deployers of AI systems to complete an RAI artifact. The repurposing of RAI artifacts in these efforts are increasingly supported by RAI researchers, who have suggested that RAI artifacts are well-suited to support regulatory compliance (Pistilli et al. 2023).

With increasing multi-stakeholder interest and reliance on Multi-Actor RAI Artifacts, it becomes critical to understand how these cross-sector stakeholders perceive their efficacy. Existing work studying RAI artifacts have examined how they are used by industry AI practitioners (Deng et al. 2022; Okolo and Lin 2024; Balayn et al. 2023), what normative assumptions about “AI ethics” work they embed (Wong, Madaio, and Merrill 2023), and the need to center *effectiveness* evaluations—that is, an evaluation that examines causal relationships between the desired outcome of an artifact and their actual impacts in a given context (Berman, Goyal, and Madaio 2024). However, minimal work has stepped back to consider what the “desired outcome” for RAI artifacts should be from the perspectives of cross-sector AI actors (e.g., regulatory and CSO actors), and whether and to what extent the existing design and use of artifacts help achieve these outcomes. Our study aims to address these gaps.

Methods

In this study, our goal was to understand how legal/regulatory and civil society organization (CSO) stakeholders perceive the efficacy of Multi-Actor Responsible AI Artifacts, including their desires for improving artifact design and use. Towards this goal, we conducted semi-structured interviews and design activities with 19 participants across these stakeholder groups. Our study was approved by the institutional review board (IRB) at Microsoft. Below we discuss details about our steps we took to prepare for the study, our recruitment procedure, study protocol, and analysis approach.

Study Preparation

Prior to beginning our study, we aimed to contextualize our existing understanding of the broader landscape of RAI Artifacts, to inform the development of our study protocol. We therefore recruited and conducted formative interviews with four industry responsible AI practitioners who have experience informing the design of responsible AI artifacts. The formative interview focused on (i) gathering input on the goals and intended uses of existing RAI artifacts and (ii) learning about industry practitioners’ behind-the-scenes practices for ensuring their artifacts meet regulatory expectations and community needs. Knowledge shared during these formative interviews informed the creation of specific topics and questions in our study protocol, as well as our

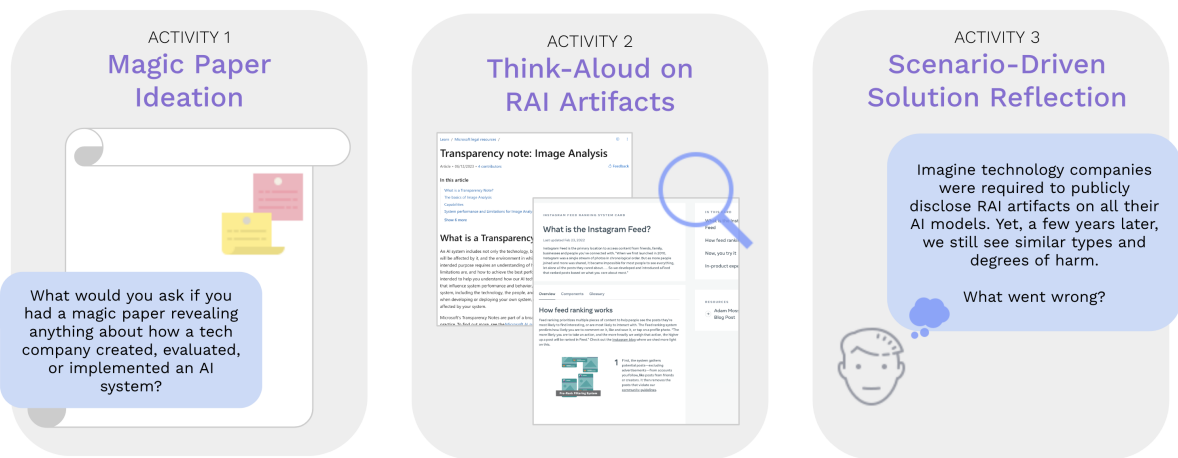


Figure 1: Overview of the three activities in our study protocol

list of RAI artifacts to include in the design activity³ The list includes a subset of publicly available RAI artifacts that were (i) created by technology companies and (ii) designed to be used by audiences from our stakeholder groups, based on publicly available documentation (e.g., research papers, company websites). We also completed pilot studies with three RAI researchers to help us identify additional points of improvement (Majid et al. 2017).

Recruitment and Participant Background

We aimed to recruit individuals who had experience working in or with regulatory bodies and CSOs on topics related to governing the responsible development and use of AI systems. We first searched for government agencies and CSOs that are heavily involved in examining the use and impacts of deployed AI systems. After identifying a list of potential organizations, we looked on the organizations' websites for employed individuals working on relevant topics (e.g., technology policy or community advocacy). We also searched for individuals employed in academic institutions with experience consulting for government agencies on AI regulation-related topics. We invited 34 individuals to the study via email. While conducting each study session, we recruited additional individuals using snowball sampling. Overall, 23 individuals responded to our recruitment emails; we scheduled 19 participants for interviews.

All 19 participants indicated having prior experiences informing regulation and public policy around AI systems in government or civil society, as well as having experiences investigating the downstream societal or legal impacts of industry-created AI systems. We report on participants' stakeholder grouping based on their current organizational affiliation. Our sample represented stakeholders across thirteen organizations of varying sizes, residing in both CSOs and regulatory bodies, representing four global locations

³The list includes: Data Cards (Google), Datasheets for Datasets (Microsoft), Model Cards (Google), Transparency Notes (Microsoft), System Cards (Meta).

(U.S., E.U., Mexico, and Brazil) with 47% identifying as men, 36% identifying as women, and 17% preferring not to disclose. We include participant demographics and details of their background in Table 1.

After conducting studies with 19 participants, no new themes emerged, indicating that a point of saturation was achieved (Saunders et al. 2018). The remuneration for the study was \$50 USD (or equivalent amount in the participant's local currency) which was distributed via an online gift card. All interviews were conducted online using a video conferencing platform between July and October 2023.

Study Overview

We conducted 60 minute study sessions with each participant, including a semi-structured interview and design activity. We began each study session by asking the participant to contextualize their question responses within their own past experiences, role-assigned responsibilities, and goals around responsible AI. By tying their responses to their background as a regulatory or civil society stakeholder, we aimed to better understand how their goals, beliefs, and experiences shaped their perceptions of RAI artifacts. After asking related questions about their RAI goals and work, we went through three core activities: (1) Magic Paper Ideation, (2) Think-Aloud on RAI Artifacts, and (3) Scenario-Driven Solution Reflection. Figure 1 includes an overview of the three activities and the Appendix⁴ includes the question guide.

Activity 1: Magic Paper Ideation Our goal for the Magic Paper Ideation Activity was to learn about participants' information needs and desires for RAI artifacts. In particular, we elicited each participant's ideas about the types of model-related questions they wished an "ideal" RAI artifact would address. Figure 1 includes an exemplar question. The participant's ideation work completed during this activity helped prepare them for the next activity, in which they were asked to examine an example of an existing RAI artifact.

⁴The appendix can be found here: <https://arxiv.org/pdf/2408.12047> (pages 13-14).

Participant ID	Professional Role	Stakeholder Group	Years of Experience	Location
P01	Researcher	Civil Society	1- 3 Years	U.S.
P02	Attorney	Law/Regulation	>7 Years	U.S.
P03	Attorney	Law/Regulation	<1 Year	U.S.
P04	Researcher	Civil Society	1- 3 Years	E.U.
P05	Researcher	Civil Society	1- 3 Years	U.S.
P06	Attorney; Community Advocate	Civil Society	<1 Year	U.S.
P07	Researcher	Civil Society	3 - 5 Years	U.S.
P08	Attorney; Community Advocate	Civil Society	>7 years	U.S.
P09	Researcher; Community Advocate	Civil Society	5 - 7 Years	U.S.
P10	Researcher	Civil Society	1- 3 Years	U.S.
P11	Community Advocate	Civil Society	1- 3 Years	U.S.
P12	Civil Society / Civil Rights Technologist	Civil Society	1- 3 Years	U.S.
P13	Researcher; Community Advocate	Civil Society	5 - 7 Years	U.S.
P14	Researcher; Analyst	Civil Society	<1 Year	U.S.
P15	Attorney	Law/Regulation	3 - 5 Years	U.S.
P16	Researcher	Civil Society	>7 Years	Mexico
P17	Attorney; Researcher	Law/Regulation	5 - 7 Years	U.S.
P18	Attorney; Researcher	Law/Regulation	>7 Years	U.S.
P19	Attorney	Law/Regulation	>5 - 7 years	Brazil

Table 1: Demographics of study participants

Activity 2: Think-Aloud on RAI Artifacts This activity helped ground participant reflections and assessments of artifacts using actual examples. At the start of this activity, each participant was shown a list of Multi-Actor RAI Artifacts (see Study Preparation). After reviewing the artifact list, the participant was asked to select one to delve deeper into. We then provided the participant a link to a filled-in example of the artifact they had selected. We asked the participant to think aloud their thoughts and reactions to the artifact, as they inspected it in their own browser. To better understand *what* specific components of the RAI artifact led to their reaction, *why*, and improvements they desired, we asked follow-up questions as needed. Figure 1 includes two examples of filled-in-artifacts.

Activity 3: Scenario-Driven Solution Reflection In this activity, we presented hypothetical scenarios to help prompt reflection around alternative artifact designs and processes that participants believed could help technology companies be accountable for their AI systems. We first asked participants to reflect on our prior discussions about the limitations they believed existing artifacts had. We then presented the scenarios to help scaffold ideation.

Data Collection and Analysis

After asking participants for their consent, we audio recorded and took notes from the design activity. All 19 participants consented to both. After transcribing the recordings, we followed a reflexive thematic analysis approach (Braun and Clarke 2019). To assist in the coding process, we used HeyMarvin⁵ to streamline collaborative coding among the three authors. Overall, we did multiple, iterative rounds of coding and affinity diagramming.

⁵HeyMarvin (<https://heymarvin.com/>) is an online platform that offers analytical tools for qualitative data with features to code collaboratively as a research team.

To code the transcripts, the first author familiarized themselves with the data and identified initial lower level codes for a subset of the transcripts. The first author then discussed the codes with the second and third authors to discuss disagreements and calibrate on the coding style and granularity. Any disagreements were resolved through synchronous discussions. The lower level codes were then transferred to a shared online board for collaborative affinity diagramming (Beyer and Holtzblatt 1999) using Figma, a web application where multiple users can synchronously generate and arrange sticky notes. The three authors iteratively grouped the first-level codes into successively higher level codes (i.e., second and third level codes). After completing each additional level of coding, the authors collaboratively discussed emerging themes and relationships between codes (Braun and Clarke 2006; Lazar, Feng, and Hochheiser 2017).

Limitations

We acknowledge the methodological limitations of our study. Our sample was skewed toward US-based participants. To diversify the participant sample, we recruited participants broadly across stakeholder groups, demographics, and location. Throughout the recruitment process, we monitored the participant demographics represented in our participant sample to guide decisions about which populations to refocus recruitment efforts towards. Second, it is possible that, because we are employees within a technology company, our participant sample is skewed towards those who are supportive and optimistic of industry responsible AI efforts. We believe this because we had recruitment contacts who objected to study participation due to moral discomfort with engaging in research funded by the technology company. We respect their right to express that stance. We encourage future research that extends this line of inquiry in different sectorial and geographic contexts, where expectations around responsible AI goals and needs differ.

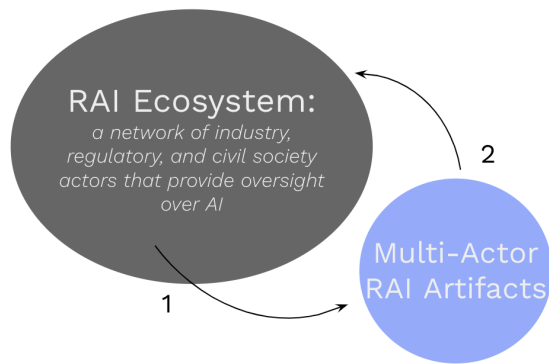


Figure 2: Participants described how (1) the technology industry exists within a multi-actor RAI ecosystem, which shape their goals and incentives for RAI governance. These goals and incentives (2) shape how the technology industry designs and uses Multi-Actor RAI Artifact; these decisions have rippling impacts on other actors within the RAI ecosystem. A more detailed breakdown of the properties and impacts of this dynamic are described in Figure 3.

Results

While participants recognized the importance of artifacts in the RAI ecosystem, our findings focus on communicating their concerns around the potential role and impacts of RAI artifacts. As described in Figure 3, participants’ reflections on and assessments of existing Multi-Actor RAI Artifacts surfaced four barriers they believed impeded their ability to protect end-users from downstream AI harms: (1) end-users⁶ being positioned as a second-order priority in the design and use of RAI artifacts, (2) intentionally showcasing the RAI artifacts for a select set of laudable AI models, (3) over-relying on transparency as a change mechanism when designing RAI artifacts, and (4) designing RAI artifacts to delegate duties for responsible use to end-users. We organize the findings subsections around these four barriers, with each subsection describing participants’ perspectives about the technology industry’s goals and incentives that contribute to the creation of the barrier, the design and use practices of RAI artifacts that make the barrier visible, the downstream impacts that barrier has on the broader RAI ecosystem, and corresponding desires participants had for improving the role, design, and use of RAI artifacts.

End-Users as Second Priority Stakeholder

A key theme that emerged revolved around how Multi-Actor RAI Artifacts (perhaps unintentionally) position end-users as secondary stakeholders. This can occur, for instance, when information documented in RAI artifacts is shaped in ways that are ultimately *inaccessible* and *inactionable* to impacted end-users. While RAI artifacts are currently mar-

⁶Here and throughout the paper, “end-users” refers to stakeholders who may be impacted by the downstream use of an AI system, including those who may use the AI model or those who may be impacted by decisions made with the AI model.

keted as being created for a wide range of stakeholders⁷, including impacted end-users and non-technical stakeholders, participants observed that the *type or presentation of content* is tailored towards technology companies’ needs rather than the needs of impacted community members. Participants acknowledged the challenges with designing “one size fit all” RAI artifacts and, instead, advocated for greater effort invested in designing RAI artifacts that center impacted end-users’ domain-specific informational needs.

For example, participants perceived a misalignment between the type of content presented in RAI artifacts versus the type of content that would be valuable for protecting community members from AI harms. In some cases, RAI artifacts appeared to omit details that would be critical for impacted community members. As one participant articulated, RAI artifacts like Model Cards that center accuracy-based performance measures overlooks other harms that may arise from the downstream use of the AI model:

“I’m so glad that they are addressing some of the harms that were articulated in Gender Shades. However, I think for a lot of community-based organizations that I’ve talked to, it would be a frightening outcome if the harms that they’re worried about or simply considered [only includes] accuracy[. . .] A more accurate Big Brother isn’t safer or better.” (P05, Researcher, Civil Society)

Participants also observed that many RAI artifacts actively discussed in industry RAI efforts provide upstream information about the AI model (e.g., dataset properties, performance on training data). However, they believed that civil society and impacted communities would find information contextualized to the AI model’s downstream use cases more meaningful. As one participant articulated, the current design of some RAI artifacts “divorces the conversation from the uses of the model, which is where the harms happen” (P02, reflecting on Model Card on Face Recognition Algorithm). Examining Facebook’s System Cards for Instagram Feed algorithm, another participant noticed that the artifact only showed a subset of their rules for how the algorithm selects posts to show. The participant wondered whether the company strategically decided to provide examples of rules that are the least problematic: “I love that they gave an example, but. . . I want to know what are all the rules. . . usually they select a very mundane [rule] that’s super nice and, you know, [one that is] simple and wouldn’t raise any red flags” (P10, Researcher, Civil Society reflecting on System Cards for the Instagram Feed Algorithm). The participant further elaborated that they desired complete information on the AI model to more critically examine the potential societal implications of the algorithm’s design decisions: “I want to see the full list of rules, so I can say, ‘OK, these rules are great. This rule might cause these harms for these and these

⁷For example, Model Cards is marketed towards broad audiences (both “experts and non-experts”) including developers, journalists, and industry analysts (<https://modelcards.withgoogle.com/about>). Datasheets for Datasets is marketed towards a wide range of actors who contend with datasets including practitioners and academics.

groups.” (P10, Researcher, Civil Society, reflecting on System Cards for Instagram Feed Algorithm). Participants were concerned because this information omission prevented end-users from accessing details that could empower them to better reason about or protect themselves from potential harms.

Moreover, several participants emphasized that this barrier was an inherent byproduct of technology companies’ motivations to leverage RAI artifacts as a tool to shape attitudes and industry norms. As one participant who informs technology policy and regulation at a CSO described, Multi-Actor RAI Artifacts are used tactically to promote self-regulatory practices and avoid stricter AI-related regulation: “[RAI artifacts] help head off stronger regulation. You know companies saying, ‘oh hey, we already have these standards. Just look at these standards.’ . . . I think it is very much driven by a desire to head off stronger regulations, perhaps” (P08, Attorney, Civil Society). As a result, even if RAI artifacts are designed with an intention to support model understanding for multiple actors, including end-users, there are disincentives embedded within industry norms that prevents prioritizing user needs in the design of RAI artifacts.

Selective Showcasing of Laudable AI Models

Participants were concerned that technology companies strategically showcased RAI artifacts for well-developed, often lower risk AI models—in part, as a tactic to improve the company’s public image—while avoiding showing artifacts for models that would benefit the most from external scrutiny and feedback. Participants emphasized the importance of showing external stakeholders RAI artifacts for potentially controversial or “work in progress” AI models, so they can use them to identify and surface problems that may arise. Participants pointed out that this created a paradox in efforts to develop Multi-actor RAI Artifacts: Even though Multi-actor RAI Artifacts are marketed as a mechanism to promote cross-stakeholder communication, technology companies are least likely to disclose RAI artifacts for the AI Models that would benefit the most from this form of external information sharing. As one participant articulated:

“I could also see a company releasing artifacts[. . .] where it makes them look good and they don’t see problems. But then the ones that are problematic are not going to be put forth publicly, and you might distract civil society and researchers by focusing on the things that you have put out there[. . .] sort of direct people’s energies away from the ones that are a work in progress.” (P02, Attorney, Regulation)

Participants acknowledged that companies may be hesitant to disclose RAI artifacts for all their models, in large part, because there are currently no safe mechanisms for companies to do so. Participants reflected on how conflicting theories of change between civil society and the technology industry pose natural tensions to realizing disclosure policies favorable to civil society. For example, some participants described how, in some countries (e.g., United States), civil society relies on raising “fire alarms” by publicly shaming or otherwise pressuring technology companies into improving their AI development practices. As one participant

reflected, RAI artifacts failing to prevent downstream harms from AI systems (by failing to showcase work-in-progress AI models) may be a “feature of the [US] system” (P08). Others pointed out that legal teams at technology companies are unlikely to allow disclosure of artifacts that surface potential model harms. As one participant acknowledged: “companies are really disincentivized from self disclosing these kinds of risks because their own legal departments are worried about liability if they claim to be aware of a certain kind of harm.” (P05, Researcher, Civil Society)

Reflecting on the information types included in Model Cards, another participant further elaborated that these liability challenges may disincentivize companies from documenting information about AI models that could help end-users protect themselves from potential harms:

“[harm use cases] is just not information that they want to connect to their technology [. . .] because [even if] they don’t intend this use [. . . if] people are using it in this way [. . .] it both shows acknowledgement of harmful uses and maybe an acknowledgement of unintended uses” (P02, Model Cards on Face Detection Algorithm)

Because of these dynamics, participants were concerned that technology companies may divert public attention from the limitations and risks of their AI systems by omitting information that may promote conversations about AI harms. Other participants were concerned that this may, in turn, weaken the role that civil society can play in identifying and mitigating these harms. Many civil society actors therefore acknowledged the importance of utilizing artifacts to study how downstream harms affect communities but they were eager to “*study up*” (Nader 1972) by learning more information about upstream power dynamics, structures and decision-making that contribute to more systemic challenges in AI. In particular, participants wanted Multi-Actor RAI Artifacts to document the *social and organizational context* surrounding technical decisions and *value-based judgments, assumptions, and policy decisions* made during AI development. For instance, participants pointed out that information around *who* was involved in model development (from problem formulation to evaluation), how organizational incentives may have impacted AI development decisions (e.g., metrics used for promotion), or even how much data labelers were paid could help surface organizations’ underlying values and priorities around AI development.

Other participants pointed out that, while the RAI Artifact they examined documented technical details about the AI model, they were more interested in “what policy choices are embedded in seemingly technical decisions” (P02). Participants described the value of documenting decisions from earlier on in the AI development process, including who framed the problem space that the AI model intends to address. As one participant elaborated:

“Just knowing what the purveyors of the system or the developers of the system thought that they were solving for is illuminating. Because again, sometimes it might just be, ‘Well, there wasn’t really a problem, but we have the capacity to do this. It’s efficient. We

think of tech as better always and so we just built it.”

(P02, Attorney, Regulation)

Some participants also believed that, in disclosing RAI artifacts for work-in-progress models, artifacts can help promote more proactive collaboration between civil society and the technology industry. Participants recalled that they sometimes provide consulting services to technology companies to share early-stage feedback for their AI models, but that these efforts also suffer from the same paradox: Companies are more likely to consult civil society on AI models that they know are already well-developed and may be disincentivized from engaging them in conversations around AI models relying on disagreeable problem formulations. Therefore, participants advocated for new processes and policies that can provide safe pathways for technology companies to solicit feedback from civil society earlier on in their development process.

Over-Reliance on Transparency as a Change Mechanism

In reflecting on why participants found it critical for civil society stakeholders to use RAI artifacts to “raise alarm bells” that prompt real-world change, participants described that transparency in itself typically does not help mitigate or address harm. However, they also acknowledged that proposals for greater transparency were a safe “political” first step that could help reach a middle ground across civil society, government, and industry actors. As one participant reflected:

“[Data transparency] is, in some sense, the least controversial thing that policymakers [and] civil society can ask for, that will piss the least people off, and that, frankly, [is] the least way to get companies to actually do anything right [...] You end up pissing the least people off politically when you are calling for transparency because it’s something that everyone can agree upon.” (P07, Researcher, Civil Society)

Even though it may be more challenging to push forth governance mechanisms “with teeth”, many participants emphasized the importance of exploring these next steps. As the participant further elaborated: “it’s not always clear what, then, gets done with that transparency. There’s a lot of, ‘sunlight will fix the problems.’” (P07, Researcher, Civil Society). Instead, one participant described that simply having potential harms documented may normalize their existence: “The interesting thing [...] I tend to feel about transparency [...] looking at the information that I have in front of me, and like not caring about it because it’s there, you know?” (P07)

Similarly, participants were particularly concerned about the continued of RAI Artifacts to help placate the image of AI models and products that they viewed as being fundamentally problematic. For example, examining the Transparency Notes for a sentiment analysis model, one participant described:

“It’s such a problematic product and you know it. It’s like, ‘oh, here are some key limitations to consider.’ And those limitations are so major and yet the product is being sold. You know, so I don’t think if you’re

selling a product that can be inherently problematic [...] a statement of some limitations is really sufficient.” (P15, Transparency Notes for Sentiment Analysis feature)

Highlighting patterns of inappropriate use and unanticipated impact, some participants believed that RAI artifacts may, in fact, negatively contribute to concerning trends in the types of AI systems that are deployed. In particular, participants conceptualized how RAI artifacts may serve as buoys for technology companies, allowing them to keep harmful AI systems afloat without addressing their harms. As one participant hypothesized, these dynamics may also expedite the deployment of new AI systems:

“Really, really powerful technology is going out the door with very few guardrails except for this like quite thin amount of documentation [...] think it’s a pretense that like a lot of systems have been expedited out the door with this in place. [It gets] the boxes checked.” (P05, Researcher, Civil Society reflecting on Model Cards for Face Detection Algorithm)

As such, participants described that they desired RAI artifacts that empower external stakeholders to not just relay, but *respond* to and help mitigate potential harms.

Offloading Duties of Responsibility to End-Users

Examining Multi-Actor RAI Artifacts that center the disclosure of harms, participants further expressed that technology companies may be leveraging artifacts to set expectations for responsible use and offload the burden for responsibility on to end-users. Participants described ways in which the content design of the RAI Artifact appeared to protect the company from liability risks, at the expense of protecting users from potential harms.

For example, participants expressed concern that, even in Transparency Notes, which is designed exclusively for impacted end-users, the information shared about the AI model negatively shapes expectations over *who* is responsible for potential downstream harms. Reading the starting paragraph of a Transparency Notes document⁸, one participant reacted:

“I think the note in the second paragraph, right away, raises a little bit of a red flag for me [...] Basically like [it’s saying, ‘if] you’re going to use our system, [it] is your responsibility to make smart choices.’ As opposed to saying, ‘here are the choices we made that will then impact what you can or cannot do with this system.’” (P13, Researcher/Advocate, Civil Society)

Similarly, another participant was surprised by the title “Transparency Notes”, expressing: “my initial response to Transparency Note [was] ‘Oh, Microsoft is being transparent about something.’ But what it feels like they’re doing here is putting the responsibility on someone else who is

⁸“Notes are intended to help you understand how our AI technology works, the choices system owners can make that influence system performance and behavior, and the importance of thinking about the whole system, including the technology, the people, and the environment.”

going to be using the technology” (P10, Researcher, Civil Society).

Another participant examining Transparency Notes pointed out that the example the artifact used to convey model limitations⁹ appeared to communicate limitations in an ideal case scenario, rather than the limitations that end-users may experience in practice. The participant elaborated that it would instead be helpful to provide examples that more closely mirrored the real-world limitations in the model that users may encounter:

“If I say ‘you’re such a boss a** b**ch’, that’s a good thing, right? But if I say you’re such a b**ch, that’s a bad thing. And that could easily be an example that’s included here. It doesn’t sound very professional, but it really illustrate extent to which it’s so sensitive [...] it just makes it clear that this [limitation is] kind of idealistic, and it won’t really be able to measure things.” (P16, Researcher, Civil Society, reflecting on Transparency Notes on Sentiment Analysis feature)

As P16 described, and as Section elaborates, participants generally desired RAI artifacts that prompted technology companies to disclose realistic model-related information that would could be more easily interpreted by end-users, even if this means that the model creator (e.g., the technology company or AI team) may not look as reputable.

Discussion

Participants’ reflections on existing RAI artifacts surface how Multi-Actor RAI Artifacts reside within a broader AI governance ecosystem, shaping their design and use. Figure 2 shows how the technology industry’s goals, incentives, and theories of change shape how they design and use RAI artifacts. Participants perceived their own RAI goals (i.e., protecting end-users from AI harm) as conflicting with that of industry (i.e., promoting self-regulation and shaping societal standards for AI), leading to concerns about the broader downstream impacts of these practices.

We organize participants’ concerns into four key barriers that help break down participants’ beliefs about the current overlooked impacts of Multi-Actor RAI Artifacts: (1) End-users as second-order priority, (2) Selective showcasing of laudable AI models, (3) Over-reliance on transparency as a change mechanism, and (4) Offloading duties for responsible use to end-users. To overcome these barriers, participants desired structural changes to the RAI governance ecosystem, coupled with improvements in how RAI artifacts are designed and used. Drawing on these findings, we discuss implications for research, design, and policy.

Challenges from Conflicting Theories of Change

Analogues to the kinds of processes and artifacts we study in the AI context have existed in mature industries like

⁹The participant was referring to the following text on the Transparency Note: “For example, that was awesome, could be either positive or negative depending on the context, tone of voice, facial expression, author and the audience.”

transportation, agriculture and pharmaceuticals for decades. Indeed many RAI artifacts take direct inspiration from consumer- and practitioner-facing documentation such as nutrition labels and checklists. However, whereas in the more mature industries artifacts generally arose *in response to* existing regulation, and were de facto instruments for compliance (Boon, Lichtenstein, and Wartella 2010), RAI artifacts largely predate regulatory guidance on AI. As our participants noted, RAI artifacts play a major role in *shaping* expectations and standards around AI in the first place.

Within these dynamics, we can observe how actors across sectors (e.g., regulatory, civil society, industry) each have their own—often implicit—desired outcomes, incentives, and strategies they use to advance their vision for what responsibly governing AI should or could look like. In our study, participants drew on these conflicting theories of change to both explain their own assessments of the efficacy of artifacts, as well as explain why technology companies currently leverage artifacts. In particular, participants described how their judgements about the technology industry’s underlying goals for RAI mediated their perceptions of the efficacy of Multi-Actor RAI Artifacts. Participants believed that their own desired end-goals for RAI artifacts (i.e., protecting end-users from AI harms) deviated from technology companies’ underlying end-goals (i.e., promoting self-regulatory practices and shaping societal expectations around AI). They identified how these clashing goals lead to the manifestation of four barriers that impede the effective design and use of RAI artifacts, where “efficacy” is judged with respect to their own end-goals.

At first glance, these tensions pose significant challenges. Many of the challenges we discussed are a natural consequence of the technology industry operating off of a self-governing model under a capitalist value system, indicating they may be nearly impossible to mitigate without substantial societal changes. For example, technology companies will tend to err on the side of not disclosing harmful model-related information, for concern of backlash and profit-related harms, even if this may be counter-productive to the goal of proactively identifying and mitigating AI harms to impacted stakeholders. However, we argue that these barriers are not *inevitable* in broader societal efforts to govern AI. Future work should investigate structural changes that could be made through new regulation and advocacy efforts, that could support the need for *meaningful disclosure* amongst regulatory and CSO stakeholders. We continue this discussion in the next section on design and policy implications.

Towards Facilitating Multi-Actor Communication

In our findings, we observe barriers to the responsible design and use of RAI artifacts that hinder multi-actor efforts around AI governance. These barriers ultimately reduce legal, CSO, and impacted end-users’ ability to proactively anticipate or mitigate AI-related harms. Borrowing from the social sciences, this can be described as reducing *social capacities*—the “abilities, skills and competences that help people to better prepare for, respond to, recover from or adapt to the negative impacts of [...] hazard events” (Höppner et al. 2012; Kuhlicke and Steinführer 2010). Höppner et al. argued

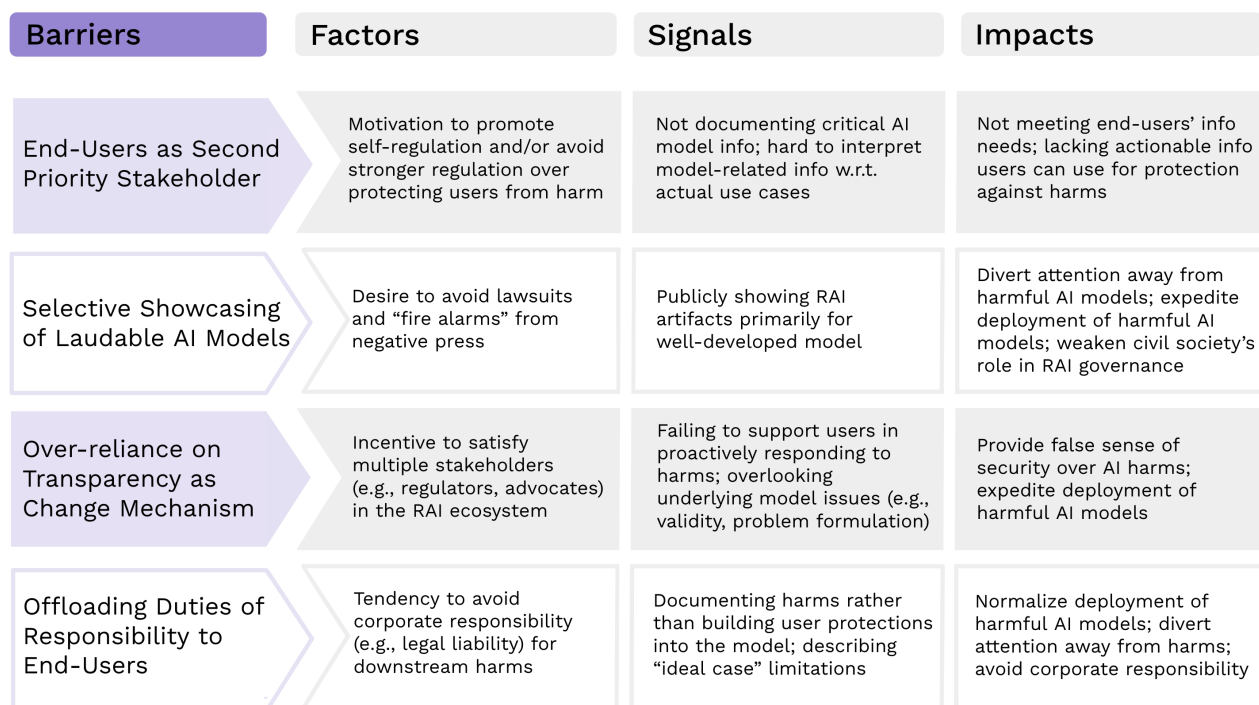


Figure 3: A summary of the four barriers to Multi-Actor RAI Artifacts surfaced by participants, each mapping to one subsection of the findings. For each barrier, we describe (1) Factors: the incentives, goals, and theories of change of the technology industry that contribute to the barrier, (2) Signals: ways in which the barriers become visible through design and use practices around RAI artifacts, and (3) Impacts: the broader downstream impacts the barriers have on the responsible AI ecosystem and society.

that improved risk communication is vital for crafting environments that help build up communities' social capacities (Höppner et al. 2012). In practice, risk communication has been a hallmark feature embedded in design and policy around Multi-Actor RAI Artifacts; it is also increasingly a fiduciary responsibility for high-risk use cases of AI with the emergence of newer regulatory frameworks (Calvi and Kotzinos 2023). However, as our findings suggest, industry-external stakeholders have been overlooked in efforts to communicate AI-related risks from RAI artifacts. This paper emphasizes the importance of questioning *for whom* RAI artifacts are designed, to in turn deliberate *which risks* and *how* these risks are (not) communicated.

Towards centering these considerations, we extend existing calls to increase stakeholder participation (particularly those residing outside of the technology industry) in the design of RAI artifacts. As our findings describe, treating these stakeholders as "secondary" stakeholders risks overlooking their unique needs and desired uses of RAI artifacts. For example, participants described how use cases and limitations for AI models should be grounded in what they experience in their day-to-day, to help them anticipate certain harms. Moreover, these information needs may differ depending on the type of AI model covered in the artifact. Future efforts that use RAI artifacts for governance should account for the type of AI model, like the associated risk level (e.g., as in the EU AI Act (Madiaga 2023)) or level of interaction with end-

users, to assess whether and when certain RAI artifacts can support governance. This may require scoping down the intended uses of a given RAI artifact, countering the industrial tendency to design for scalability (Hanna and Park 2020). Besides designing RAI artifacts to support improved communication of AI model information to stakeholders, participants wanted artifacts to better facilitate *collaboration* amongst industry, regulatory, and CSO stakeholders. For example, CSO participants described that accessing RAI artifacts for "work-in-progress" AI models could allow them to play a bigger role in risk identification and mitigation, before models are deployed. Future work could explore how policy interventions can incentivize technology companies to regularly involve CSO stakeholders as third party auditors of their AI systems, especially in earlier stages of AI design.

Conclusion

Multi-Actor Responsible AI Artifacts can play a pivotal role in fostering collaborative AI governance efforts amongst industry, legal, and civil society actors. This study advances our understanding of how to realize this vision by addressing an outstanding gap in literature: Understanding legal and civil society stakeholders' perspectives around the efficacy of existing Multi-Actor RAI Artifacts.

Acknowledgements

This research was conducted at Microsoft Research. We thank our participants for sharing their valuable time, insights, and experiences for the study, and the reviewers for sharing constructive feedback. We are also very grateful for the feedback shared by Ryan Steed, Seyun Kim, Alicia DeVrio, Mihaela Vorvoreanu, and those in the Microsoft Research FATE group and Office of Responsible AI.

Reflection on Ethics and Social Impact

Positionality and Adverse Impact Statement

In this work, our positionality stems from a multidisciplinary perspective at the intersection of technology, ethics, and society. As researchers, we acknowledge our own subjectivities, backgrounds, and positionalities, which inevitably shape our interpretations, methodologies, and contributions to the study.

Our team comprises individuals with diverse academic and professional backgrounds, including computer science, human computer interaction (HCI), and responsible AI. This diversity enriches our research by providing varied perspectives and insights into the complex issues under consideration in our study. While not paid explicitly to complete this project, the authors were supported by a large technology company in the United States. We recognize the inherent power dynamics embedded within the research process and opted to mitigate biases and adopt inclusive guidelines throughout our methodology to ensure participants were recruited respectfully while being mindful of potential privileges.

Our goal in conducting this research was to incorporate legal and civil perspectives around the use and design of Multi-Actor RAI Artifacts. However, in adopting this framing it is possible that criticisms surrounding existing artifacts could be misconstrued as advocating for their removal or limited use. Instead, we emphasize the important and fundamental role of current Multi-Actor RAI Artifacts and envision utilizing the insights presented to charter a wave of improvements to facilitate improved communication among RAI actors.

In presenting our findings, we acknowledge the limitations of our positionality and strive to be transparent about our methodologies, assumptions, and potential biases. We invite readers to critically engage with our work, interrogate our perspectives, and contribute to ongoing conversations about the role of Multi-Actor RAI Artifacts in the field.

Participant Safety and Ethical Considerations

Throughout the study, we reminded stakeholders that their participation was completely voluntary and that their responses would be kept anonymous. To minimize risks of identification, we report on demographics from an aggregate level, assign unique identifiers in our reporting, and opt to report on stakeholder groupings opposed to reporting on participant affiliations. For the study sessions, participants were given the agency to consent to their audio being recorded and data from the design activities being saved for note-taking and analytical purposes.

References

- Adkins, D.; Alsallakh, B.; Cheema, A.; Kokhlikyan, N.; McReynolds, E.; Mishra, P.; Procope, C.; Sawruk, J.; Wang, E.; and Zvyagina, P. 2022. Method cards for prescriptive machine-learning transparency. In *Proceedings of the 1st International Conference on AI Engineering: Software Engineering for AI*, 90–100.
- Balayn, A.; Yurrita, M.; Yang, J.; and Gadiraju, U. 2023. “Fairness Toolkits, A Checkbox Culture?” On the Factors that Fragment Developer Practices in Handling Algorithmic Harms. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 482–495.
- Berman, G.; Goyal, N.; and Madaio, M. 2024. A Scoping Study of Evaluation Practices for Responsible AI Tools: Steps Towards Effectiveness Evaluations. *arXiv preprint arXiv:2401.17486*.
- Beyer, H.; and Holtzblatt, K. 1999. Contextual design. *interactions*, 6(1): 32–42.
- Bierska, A.; Buhnova, B.; and Bangui, H. 2022. An Integrated Checklist for Architecture Design of Critical Software Systems. In *FedCSIS (Position Papers)*, 133–140.
- Bird, S.; Dudík, M.; Edgar, R.; Horn, B.; Lutz, R.; Milan, V.; Sameki, M.; Wallach, H.; and Walker, K. 2020. Fairlearn: A toolkit for assessing and improving fairness in AI. *Microsoft, Tech. Rep. MSR-TR-2020-32*.
- Blodgett, S. L.; Liao, Q. V.; Olteanu, A.; Mihalcea, R.; Muller, M.; Scheuerman, M. K.; Tan, C.; and Yang, Q. 2022. Responsible language technologies: Foreseeing and mitigating harms. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, 1–3.
- Boon, C. S.; Lichtenstein, A. H.; and Wartella, E. A. 2010. *Front-of-package nutrition rating systems and symbols: Phase I report*. National Academies Press.
- Braun, V.; and Clarke, V. 2006. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2): 77–101.
- Braun, V.; and Clarke, V. 2019. Reflecting on reflexive thematic analysis. *Qualitative research in sport, exercise and health*, 11(4): 589–597.
- Calvi, A.; and Kotzinos, D. 2023. Enhancing AI fairness through impact assessment in the European Union: a legal and computer science perspective. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 1229–1245.
- Chatterjee, M.; and Bordelon, B. 2023. Washington watches as Big Tech pitches its own rules for AI.
- Cohen, M. 2023. Why waiting for A.I. laws, regulations from government could be a catastrophic mistake. *CNBC*.
- Deng, W. H.; Nagireddy, M.; Lee, M. S. A.; Singh, J.; Wu, Z. S.; Holstein, K.; and Zhu, H. 2022. Exploring how machine learning practitioners (try to) use fairness toolkits. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 473–484.
- Dolata, M.; Feuerriegel, S.; and Schwabe, G. 2022. A sociotechnical view of algorithmic fairness. *Information Systems Journal*, 32(4): 754–818.

- Domínguez Hernández, A.; Ramokapane, K. M.; Das Chowdhury, P.; Michalec, O.; Johnstone, E.; Godwin, E.; Cork, A. G.; and Rashid, A. 2023. Co-creating a Transdisciplinary Map of Technology-mediated Harms, Risks and Vulnerabilities: Challenges, ambivalences and opportunities. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW2): 1–21.
- Ferretti, T. 2022. An institutionalist approach to AI ethics: justifying the priority of government regulation over self-regulation. *Moral Philosophy and Politics*, 9(2): 239–265.
- Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12): 86–92.
- Hanna, A.; and Park, T. M. 2020. Against scale: Provocations and resistances to scale thinking. *arXiv preprint arXiv:2010.08850*.
- Höppner, C.; Whittle, R.; Bründl, M.; and Buchecker, M. 2012. Linking social capacities and risk communication in Europe: a gap between theory and practice? *Natural hazards*, 64: 1753–1778.
- Hupont, I.; Fernández-Llorca, D.; Baldassarri, S.; and Gómez, E. 2024. Use case cards: A use case reporting framework inspired by the european AI act. *Ethics and Information Technology*, 26(2): 19.
- Kagan, R. A.; Gunningham, N.; and Thornton, D. 2011. Fear, duty, and regulatory compliance: lessons from three research projects. *Explaining compliance: Business responses to regulation*, 37–58.
- Kaminski, M. E. 2018. Binary governance: Lessons from the GDPR’s approach to algorithmic accountability. *S. Cal. L. Rev.*, 92: 1529.
- Kuhlicke, C.; and Steinführer, A. 2010. Knowledge Inventory: State of the Art of Natural Hazards Research in the Social Sciences and Further Research Needs for Social Capacity Building. *Rapport du projet CapHaz-Net Social Capacity Building for Natural Hazards—Toward More Resilient Societies*.
- Lazar, J.; Feng, J. H.; and Hochheiser, H. 2017. Chapter 8-Interviews and focus groups. *Research methods in human computer interaction*, 187–228.
- Levin, B.; and Downes, L. 2023. Who is going to regulate AI?
- Lin, C. K.; and Jackson, S. J. 2023. From bias to repair: Error as a site of collaboration and negotiation in applied data science work. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW1): 1–32.
- Madaio, M. A.; Stark, L.; Wortman Vaughan, J.; and Wallach, H. 2020. Co-designing checklists to understand organizational challenges and opportunities around fairness in AI. In *Proceedings of the 2020 CHI conference on human factors in computing systems*, 1–14.
- Madiega, T. 2023. Artificial intelligence act. *European Parliament: European Parliamentary Research Service*.
- Majid, M. A. A.; Othman, M.; Mohamad, S. F.; Lim, S. A. H.; Yusof, A.; et al. 2017. Piloting for interviews in qualitative research: Operationalization and lessons learnt. *International Journal of Academic Research in Business and Social Sciences*, 7(4): 1073–1080.
- Mehrabani, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2021. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6): 1–35.
- Metcalfe, J.; Moss, E.; Watkins, E. A.; Singh, R.; and Elish, M. C. 2021. Algorithmic impact assessments and accountability: The co-construction of impacts. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 735–746.
- Mitchell, M.; Wu, S.; Zaldivar, A.; Barnes, P.; Vasserman, L.; Hutchinson, B.; Spitzer, E.; Raji, I. D.; and Gebru, T. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, 220–229.
- Muthukrishnan, N.; Maleki, F.; Ovens, K.; Reinhold, C.; Forghani, B.; Forghani, R.; et al. 2020. Brief history of artificial intelligence. *Neuroimaging Clinics of North America*, 30(4): 393–399.
- Nader, L. 1972. Up the anthropologist: Perspectives gained from studying up.
- Noble, S. U. 2018. Algorithms of oppression: How search engines reinforce racism. In *Algorithms of oppression*. New York university press.
- Nori, H.; Jenkins, S.; Koch, P.; and Caruana, R. 2019. InterpretML: A Unified Framework for Machine Learning Interpretability. *arXiv preprint arXiv:1909.09223*.
- Novelli, C.; Casolari, F.; Rotolo, A.; Taddeo, M.; and Floridi, L. 2023. Taking AI risks seriously: a new assessment model for the AI Act. *AI & SOCIETY*, 1–5.
- Ntoutsis, E.; Fafalios, P.; Gadiraju, U.; Iosifidis, V.; Nejdli, W.; Vidal, M.-E.; Ruggieri, S.; Turini, F.; Papadopoulos, S.; Krasanakis, E.; et al. 2020. Bias in data-driven artificial intelligence systems—An introductory survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 10(3): e1356.
- Okolo, C. T.; and Lin, H. 2024. ” You can’t build what you don’t understand”: Practitioner Perspectives on Explainable AI in the Global South.
- Pistilli, G.; Muñoz Ferrandis, C.; Jernite, Y.; and Mitchell, M. 2023. Stronger together: on the articulation of ethical charters, legal tools, and technical documentation in ML. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 343–354.
- Rakova, B.; and Dobbe, R. 2023. Algorithms as Social-Ecological-Technological Systems: an Environmental Justice Lens on Algorithmic Audits. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 491–491.
- Reisman, D.; Schultz, J.; Crawford, K.; and Whittaker, M. 2018. Algorithmic impact assessments: A practical framework for public agency. *AI Now*, 9.
- Sadek, M.; Constantinides, M.; Quercia, D.; and Mougnot, C. 2024. Guidelines for Integrating Value Sensitive Design in Responsible AI Toolkits. *arXiv preprint arXiv:2403.00145*.

- Saunders, B.; Sim, J.; Kingstone, T.; Baker, S.; Waterfield, J.; Bartlam, B.; Burroughs, H.; and Jinks, C. 2018. Saturation in qualitative research: exploring its conceptualization and operationalization. *Quality & quantity*, 52: 1893–1907.
- Stjernfelt, F.; Lauritzen, A. M.; Stjernfelt, F.; and Lauritzen, A. M. 2020. The Role of Civil Society. *Your Post has been Removed: Tech Giants and Freedom of Speech*, 241–260.
- Weske, U.; Boselie, P.; van Rensen, E.; and Schneider, M. 2019. Physician compliance with quality and patient safety regulations: The role of perceived enforcement approaches and commitment. *Health services management research*, 32(2): 103–112.
- Weske, U.; Boselie, P.; Van Rensen, E. L.; and Schneider, M. M. 2018. Using regulatory enforcement theory to explain compliance with quality and patient safety regulations: the case of internal audits. *BMC health services research*, 18: 1–6.
- Wong, R. Y.; Madaio, M. A.; and Merrill, N. 2023. Seeing like a toolkit: How toolkits envision the work of AI ethics. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW1): 1–27.
- Yoo, C. S.; and Lai, A. 2020. Regulation of algorithmic tools in the United States. *Journal of Law & Economic Regulation*.
- Ywakaim Andrawes, S.; Faisal Fakhry, S.; and Abd El Azeem, H. 2019. Nurses' Compliance with Patient Safety Standards in an Accredited Hospital. *Egyptian Journal of Health Care*, 10(1): 65–79.