

Wahala for Who No Sabi: GRACE and the Cultural Translation of AI Governance

IBIYEMISI GBENEBOR, Brown University, USA

1 Introduction

1.1 The Rationalist Root of AI systems

Artificial Intelligence (AI) is deployed at scale. On a lower scale, consider: A doorbell at your neighbor's house that lets them know you're there, a self-driving car that is in the lane next to you, or a conversation with your favorite chatbot. These are examples of AI systems that use data to make informed decisions or perform automated tasks. People become data, but the social context that gives meaning to personhood, referred to as relational personhood, is systematically excluded from datasets [4, 17]. The social context that is stripped can include a new neighbor whose unfamiliar face flags the doorbell camera as a trespasser rather than someone searching for a misdelivered package. By encoding humans as objects in a probabilistic model, autonomous vehicle systems enact a logical reduction of personhood in which individuals become interchangeable units of risk within an optimization framework. In each case, the system captures the data point but not the relational web that makes that person who they are— the relationships, histories, and obligations that constitute personhood [17]. The data has become quantitative because the builders of AI systems operate from a rationalist framework, reflecting a particular logic about whose definitions of data and personhood get to count.

As AI scales globally, so does the reach of embedded assumptions born of Western Enlightenment-era ideologies that are far from universal, yet are imposed on the global majority as if they were [4, 17]. The traditional Western conception of personhood privileges rationality, positioning the individual as an autonomous, self-contained subject whose essence can be known through formal logic [4]. If personhood is defined by rationality, and rationality is understood as formal logic and computation, then only what can be calculated counts as knowledge. Relational context, the obligations, histories, and communal ties that constitute personhood, cannot be quantified for computation, and are not missing from these systems by accident but by design [17]. Furthermore, the rationalist worldview, rooted in Cartesian logic, treats knowledge as a distant act (abstract, contextless, and raceless), privileging isolation and clear binaries over connectedness and dynamic relation [4]. Such reductive logic is not new, nor are its dehumanizing effects. The dominance of Western rationalist epistemologies has historically enabled forms of violence under the guise of reason. The transatlantic enslavement of Africans and colonial expansion were justified through claims that certain populations lacked rationality and therefore required governance, discipline, or 'civilization' [17]. This logic denied the full humanity of entire peoples by subordinating their ways of being to a singular, exclusionary standard of human worth. Almost 150 years after the end of the transatlantic slave trade, the United Nations adopted a resolution recognizing the transatlantic slave trade as "the gravest crime against humanity" [29].

This belated recognition reflects a broader truth: the harms produced by rationalist logic did not end with formal colonialism but persist today through Neocolonialism, which has allowed structures of power and dominance to endure while shaping global technological systems. Colonialism and capitalism have always been twin processes, and the wealth accumulated through centuries

of extraction has largely consolidated within Global North institutions that now fund AI development [17]. Building and deploying AI at scale requires vast infrastructure, and only the most powerful corporations, governments, and elite universities have the resources to do so [11]. The goals these institutions pursue: science, surveillance, and selling, shape the end products of the design process and whose values get encoded [11].

The impact of neocolonialism is evident in cases such as the deployment of the Gojek application. In Indonesia, ojek refers to an informal motorcycle taxi system organized through localized, relationship-based labor and negotiated fares. Gojek is an app-based platform, similar to Uber, that overlays this system with a Western-modeled algorithmic infrastructure, displacing community-based practices in favor of opaque data-driven governance that workers have little power to contest or shape [18]. This pattern extends beyond platform labor into broader regimes of surveillance and governance. This algorithmic opacity displaces prior entrepreneurial practices, where workers could negotiate pay and set terms, replacing them with platform-controlled pricing mechanisms that limit their agency and allows the Gojek platform to control the flow of money. In another example, AI-enabled surveillance infrastructures have been deployed across South African urban spaces through extensive camera networks and facial recognition systems, producing what has been described as a form of “digital apartheid” [13]. These systems disproportionately subject already marginalized communities to heightened monitoring and policing, while offering little transparency or recourse to those most affected. Although framed as tools for security and urban management, they reproduce longstanding racial and spatial inequalities shaped by apartheid and its afterlives, embedding them within contemporary data infrastructures [13].

Across both of these cases in Indonesia and South Africa, AI systems disproportionately affect communities whose way of being are incompatible with data-driven logics. These dynamics operate through infrastructural systems that reorganize social and economic life through data extraction, classification, and control. Such systems concentrate power within the institutions that design and govern them, enabling those institutions to define what counts as knowledge. As a result, diverse knowledge traditions are not only excluded from representation but displaced as sources of authority in shaping how life is organized and governed. Other epistemologies exist, and they demand recognition.

1.2 Relational Foundations for Data Sovereignty

Alternative understandings of personhood, knowledge, and relation have long existed outside the Western rationalist tradition, offering fundamentally different foundations for thinking about what AI systems are and what they should do. Mhlambi’s notion of relational personhood challenges the rationalist view of personhood by emphasizing that personhood is constituted through social relationships, interdependence, and cultural context rather than isolated cognition [17]. Rooted in the Zulu notion of Ubuntu, Mhlambi’s view aligns with other Global South frameworks, including the Yoruba concept of “iwa” (character or moral behavior) and the Akan concept of “onipa” (personhood) [1]. Together these offer alternative ways of understanding personhood in technology and data governance because they are grounded in relational personhood, opposed to rational [1]. These frameworks are not just philosophical alternatives, they are expressions of cultural models of self that shape how communities understand identity, agency, and human relationships, including interactions with technology.

Consistent with this view, human computer interaction research reinforces that culture does fundamentally shape what people want from AI and how such systems should function [3, 12]. Barrett et al. demonstrate this through a proposed African data ethics framework built around six major principles: challenging power asymmetries, asserting data self-determination, investing in local data institutions, utilizing communalist practices, centering marginalized communities

and upholding the common good [1]. Each principle reflects deeply held African cultural values of community, reciprocity, and collective dignity. Together they offer a coherent and legitimate alternative framework for governing technology and data. However, technologies continue to be designed and deployed in African contexts without these commitments at their foundation. These commitments also shape how consent and data governance are understood, informing emerging frameworks of data sovereignty through which communities and nations assert the right to govern their data according to their own priorities.

Data sovereignty offers a critical intervention into dominant AI frameworks by asserting the right of polities to govern how their data is collected, interpreted, and used. Rather than treating data as an abstract, freely extractable resource, it reframes data as embedded within social, cultural, and political life [1]. This principle extends into indigenous contexts through indigenous data sovereignty, exemplified by the Māori Data Sovereignty Network. The Māori people center concepts such as harmony with others, ecological understanding, and kinship networks as foundational to how data should be governed [20]. Building on these foundations, Munn adapts the five tests originally developed by Māori scholar Sir Hirini Moko Mead for technology governance. Mead originally designed these tests as a method for Māori communities to evaluate novel and contentious issues through their own values and knowledge systems, arriving at a Māori position rather than deferring to external frameworks. Munn extends this method to AI, using it as a rubric for assessing whether data-driven technologies align with Māori values of sacredness, life force, reciprocity, communal well-being, and ecological sustainability [20].

This approach to indigenous data governance finds global expression through the Global Indigenous Data Alliance, whose CARE principles: collective benefit, authority to control, responsibility, and ethics, establish a shared framework through which indigenous peoples worldwide assert rights over their data on their own cultural and ethical terms [8]. Importantly, these frameworks are not interchangeable or universally transferable, as each is deeply rooted in the place-based philosophies and ancestral knowledge systems of particular peoples, further underscoring that no single governance model can speak for all [8]. Across these indigenous frameworks, a common thread emerges: cultural values are not secondary considerations in data governance but its very foundation, inseparable from questions of ownership, meaning, and community benefit.

1.3 Participatory Design as a Pathway to Data Sovereignty

Asserting the right to govern data and technology is only realized if communities have concrete mechanisms to shape how those systems are designed from the ground up. This is where participatory design methods become essential. Participatory methods provide a means to include wider publics into the design, development and deployment of AI systems [5]. By expanding who has a seat at the table, participatory design does not merely diversify input but aims to fundamentally redistribute the power to define what AI systems are built for and whose values they reflect [3, 5, 17, 18]. Te Hiku Media, a Māori non-profit, offers a concrete example of how participatory design can directly serve indigenous data sovereignty. Members of the Māori community were actively involved in recording and transcribing audio to create a dataset of their own language, ensuring that the data was produced by and for Māori people rather than extracted by external researchers or technology companies [8]. When an American technology company later approached Te Hiku with an offer to purchase the dataset, the organization rejected it, asserting that indigenous knowledge and language data must remain under indigenous control [8]. The Masakhane project also offers a compelling example of participatory research in practice. Bringing together over 400 participants from across Africa without formal research prerequisites, the project used a collaborative, community-driven approach to develop machine translation benchmarks for over 30 African languages that had been systemically left behind by mainstream NLP research, centering

African communities [21]. Participatory design has also been used to center community values in AI literacy education. Escape et al. conducted co-design sessions with high school girls and non-binary youth of color, a group historically underrepresented in computing, to design games that teach AI ethics and fairness concepts. A key finding is that youth drew on their own situated identities and cultural experiences to critically engage with and challenge AI systems [15].

These examples illustrate that participatory design is not simply a methodological choice but a political one, actively countering the power asymmetries that have historically excluded marginalized communities from shaping the technologies that affect their lives. When communities are included as co-designers rather than passive subjects, they bring situated knowledge, cultural values, and lived experience that technical experts alone cannot provide, producing AI systems that are more accountable and culturally grounded [5, 15, 21]. Implemented meaningfully, participatory design is a practical mechanism for offering stakeholders most affected by AI systems a more active role in defining how these systems should operate.

Informed participation in AI design processes presupposes a baseline understanding of how these systems work, and for many marginalized communities this represents a significant barrier in itself. The underrepresentation of AI experts in the Global South is not incidental but reflects the legacy of neocolonialism due to structural inequalities in access to education and technical resources, making AI literacy a critical foundation for enabling authentic community participation in AI design processes [24, 26]. The Masakhane project is also instructive here specifically for its approach to AI literacy, deliberately lowering the threshold for participation by requiring no formal prerequisites and providing accessible tutorials for training machine translation models [21]. Furthermore, mutual mentorship between international researchers and African community members enabled reciprocal knowledge exchange, ensuring that technical expertise flowed into the community while cultural and linguistic knowledge informed the research itself [21].

In Aotearoa New Zealand, the Eyes on Tāmaki museum installation at Auckland Museum offers a compelling example of participatory public engagement with surveillance technology, using an augmented reality experience centered on facial recognition to educate over 100,000 visitors about data surveillance while prompting collective reflection on its governance [24]. Grounded in the context of Māori data sovereignty, the exhibit deliberately reframed surveillance as a communal issue rather than an individual burden, inviting visitors to consider how technologies like facial recognition should be governed at the level of the city and community [24].

Both the Eyes on Tāmaki installation and the Masakhane project demonstrate how AI literacy incorporated within participatory methods can substantively lower the barrier between non-experts and complex AI systems, equipping communities with the knowledge and agency to engage in informed decision-making about technologies that affect their lives. This is particularly critical in the Global South, where nations are often characterized as low-resource not by choice but as a consequence of historical and structural inequalities. In these contexts, ensuring that non-experts have meaningful and accessible pathways to participation is a foundational requirement for building AI systems that reflect the values and priorities of the communities they are meant to serve [1, 3, 12, 14, 17, 24].

1.4 Design Tensions in Participatory AI Governance Methods

Participation alone is insufficient. Enabling non-experts to critically engage with AI systems and make informed governance decisions requires more than inclusion; it requires forms of cultural translation that preserve the knowledge systems communities bring, rather than flattening them in the process of integration into technical systems. For communities in the Global South, such engagement further demands tools that are accessible, replicable, and capable of centering local

knowledge systems while supporting meaningful decision-making about data sovereignty. Existing participatory frameworks, however, tend to excel at one dimension while falling short on the other. Those that are most culturally grounded and community-centered are often difficult to replicate at scale, while those that are easily deployable across contexts frequently sacrifice the cultural depth needed to surface genuinely diverse perspectives.

The STELA process, developed by Google DeepMind, offers a comprehensive framework using expert consultation, focus groups, and qualitative coding to elicit community-specific norms and translate them into rules governing how AI chatbots should behave toward marginalized groups [2]. However, its reliance on professional moderators, subject matter experts, and multiple stages of review makes it resource-intensive and difficult to replicate outside of well-funded institutional contexts. The RAD framework (Recognize, Analyze, Deliberate) offers a contrasting approach, providing a structured and accessible method for guiding non-experts through sociotechnical critique of AI by building awareness of harms, analyzing power dynamics, and deliberating on alternative futures [26]. A key strength of RAD lies in its use of narrative-based scenarios drawn from widely available sources, including news media, social platforms, and facilitator-constructed cases. These materials can be readily sourced or created without specialized expertise or institutional resources, making them accessible entry points for engaging communities in sociotechnical reasoning. While RAD enables broader participation across educational and community settings, it does not provide mechanisms for surfacing or embedding the specific cultural values of the communities it engages, leaving the work of cultural translation underdeveloped.

Together, these frameworks point to a set of core methodological requirements for participatory AI governance that supports non-expert engagement under real-world constraints. First, accessible entry points into sociotechnical reasoning are critical, enabled through the use of narratives and scenarios that can be readily sourced or constructed from widely available materials. Second, structured mechanisms for analysis, such as stakeholder mapping and harm classification, are necessary to support rigorous, informed critique rather than surface-level participation. Third, participatory frameworks must include explicit mechanisms for embedding culturally situated values throughout the governance process, as illustrated by STELA's emphasis on community-specific norm development.

Despite these advances, no existing framework integrates all three requirements: accessible engagement, structured sociotechnical analysis, and culturally grounded value elicitation, into a single, replicable approach suitable for non-expert communities operating in low-resource contexts. This gap limits the ability of such communities not only to critique AI systems, but also to translate those critiques into governance frameworks that reflect their own epistemologies. This paper introduces the GRACE framework to address this gap, evaluated here through a case study grounded in the Yoruba value system. Specifically, this work asks: 1) Does GRACE enable non-experts to critique AI systems and make informed decisions about governance frameworks? 2) Do the governance frameworks produced through GRACE meaningfully diverge from Western models centered on individualism and capitalist logics? and 3) Can GRACE function as an accessible and replicable toolkit for culturally grounded AI governance in low-resource contexts?

2 The GRACE Framework

The GRACE framework (*Ground, Recognize, Analyze, Center Culture, Engage*) is a participatory approach for culturally grounded AI governance. Through a structured sequence of participatory stages, it enables non-experts to critically engage with AI systems and collaboratively develop governance-oriented outputs through sociotechnical analysis and cultural translation. Each stage of the framework is defined in Table 1.

Table 1. GRACE Stages

Stage	Description
Ground	Facilitator defines the sociotechnical system, governance scope, stakeholders, and cultural value framework that structure the subsequent participatory stages.
Recognize	Participants surface perceived harms, tensions, and points of concern through interpretation of situated narratives or scenarios.
Analyze	Participants examine how potential harms emerge from power asymmetries between stakeholders and their roles within the system.
Center Culture	Participants use cultural values to reinterpret harms in the narratives, initiating the process of cultural translation.
Engage	Participants translate culturally grounded values and prior analysis into governance outputs, articulating actionable rights and protections.

3 Methods and Case Study

3.1 Research Approach: Qualitative Case Study Design rooted in Omoluabi

This study employs a qualitative demonstrative case study design to evaluate GRACE as a participatory AI governance framework. A case study approach is necessary because GRACE operates as a sociotechnical intervention embedded within a specific cultural context, therefore GRACE must be assessed within a specific cultural context rather than in the abstract.

The Yoruba cultural context was selected as a theoretically motivated critical case. The Yoruba are one of the largest ethnic groups in West Africa, numbering over 40 million people predominantly located in southwestern Nigeria. Among the oldest continuous civilizations on the continent, Yoruba city-states such as Ife and Oyo were major centres of political power, trade, and artistic production centuries before European colonial contact. Nigeria today is the most populous nation in Africa and its largest economy; and a global centre of cultural production whose influence spans music, film, and literature. Many Yoruba people now live in the United States, having relocated from Nigeria or other diaspora communities, carrying with them a cultural tradition anchored in the values of *èniyàn* and *Omólúàbí*.

The Yoruba conception of personhood, *èniyàn*, is constituted through relationships, community, and moral conduct, while *Omólúàbí* represents the ethical ideal of a person of good character that emerges from this conception. Together these concepts constitute an established, well-documented Global South epistemology that is substantively distinct from Western individualist frameworks, providing the conditions necessary to assess whether GRACE can produce AI governance frameworks that are not merely culturally inflected but represent a genuine alternative to dominant Western models.

The lead researcher completed the Ground stage and subsequently designed a two-session workshop to operationalize and evaluate GRACE: Session 1 comprised the *Recognize* and *Analyze* stages, and Session 2 comprised the *Center Culture* and *Engage* stages.

3.2 Participants

The research team consisted of the first author as lead facilitator and an additional co-facilitator who led Session 2. The lead facilitator completed the Grounding Stage, and led Session 1 (*Recognize and Analyze*). The lead facilitator observed Session 2 (*Center Culture and Engage*), and the co-facilitator led the Session 2. Both members of the research team are Yoruba scholars whose work in human-computer interaction focuses on ethics in sociotechnical systems.

Ten Yoruba participants were recruited to take part in the study. Participants were recruited through three Nigerian-led community organizations in the Greater Boston area, including a social group and two registered non-profit organizations. Recruitment was conducted exclusively through each organization's private online communities and mailing lists, ensuring that participants were drawn from within established Yoruba community networks rather than through open public channels.

All ten community participants identified as Yoruba, and had grown up in a Yoruba household, either in Nigeria or outside Nigeria. All had repeatedly visited or lived in Nigeria at some point in their lives. Participants ranged in age from 18 to 58. The group comprised five males and five females, spanning a range of professional backgrounds. The sole inclusion criterion beyond Yoruba identity was a self-reported familiarity with *Omólúàbí* values of at least 3 on a 5-point scale, ensuring that participants possessed a meaningful grounding in Yoruba cultural values without requiring expert-level knowledge. Each participant was required to complete a pre-screening questionnaire to ensure that they met the inclusion criteria.

Participants were required to be non-experts, defined as individuals who do not work professionally with AI systems in a technical or development capacity. Within this boundary, participants ranged from those with little to no knowledge of AI systems ($n=4$), to those who use AI tools functionally without technical understanding ($n=3$), to those with some technical knowledge of AI ($n=2$). One participant worked as a product manager in the AI industry; as their role was non-technical and their professional experience was outside the domain of facial recognition examined in this study, they were considered to meet the non-expert criterion. This definition is consistent with GRACE's design goal of accessibility for non-expert communities.

One individual served as a peer facilitator, while the remaining nine engaged as workshop participants contributing to the activities without facilitation responsibilities. The peer facilitator was familiar with the workshop goals and was present across both sessions to support small group ideation, providing guidance when needed without leading discussions. The peer facilitator's specific role in small-group ideation was limited to providing clarification on task direction or encouraging deeper engagement with emerging ideas when needed. The remaining 9 Yoruba people participated in the two-session workshop series (*Recognize, Analyze, Center Culture, and Engage*) to collectively produce a data rights charter.

3.3 Workshop Design and Procedure

This section explains how each stage of GRACE was implemented.

3.3.1 Ground. The *Ground* stage was completed by the researcher prior to participant engagement, and included soliciting feedback from Yoruba cultural experts within the organizations used for recruitment to ensure cultural accuracy and authenticity in the workshop materials and narratives. A prerequisite to this stage is the identification of the focal cultural group whose value system will anchor the application of GRACE. To complete this stage, the researcher identified five foundational elements:

- (1) **Technology Governance Scope:** the specific AI technology and context to be examined

- (2) **Stakeholders:** who has interest in, created, or is impacted by the system, and represent who should be reflected in narratives
- (3) **Governance Dimensions:** what aspects of the AI system should be governed in relation to those stakeholders
- (4) **Cultural Values:** the specific cultural values that will ground the governance process
- (5) **Governance Output:** the artifact or deliverable participants will produce through the workshop

The AI system examined in the context of this case study was Amazon Ring, a home surveillance technology that uses facial recognition to monitor and record in public and private spaces. Ring's integration into formal and informal policing networks has been well documented, with documented occurrences including racial profiling and over-surveillance of Black and brown communities [7, 9]. This gap between the technology's intended use (e.g. home security) and its actual deployment reflects what human-computer interaction scholars term deployment bias, the phenomenon whereby a technology's harms emerge not from the system in isolation but from the social, racial, and institutional contexts into which it is deployed" [28]. The potential impact of deployment bias is directly relevant to Yoruba people living in the United States, who are racially categorized as Black and therefore subject to the same consequences of deployment bias. However, their specific cultural identity, rooted in Eniyàn and Ọmọ̀lúàbí, is rendered invisible by both the surveillance infrastructure and the governance frameworks that currently regulate it. This is precisely the gap GRACE is designed to address, and provides a solid AI Governance scope for this case study.

Before narratives are created, the facilitator must identify stakeholders with a meaningful role in the AI governance scope, and at least one stakeholder group should directly represent the focal cultural community whose values the GRACE framework will aim to translate. The selected stakeholders represent who should be depicted in the narratives. For this study, stakeholder selection was informed by existing literature on Ring's harms to Black and brown communities, yielding five stakeholders: the target(s) of surveillance, the camera owner, platform users, developers, and law enforcement [7]. The target(s) of surveillance encompasses anyone captured within Ring's camera frame and serves as the stakeholder group representing the focal Yoruba people in this research.

The AI governance dimensions also emerged from the literature, which highlights specific mechanisms that contribute to harms in AI-powered surveillance systems. Based on this, the lead researcher defined the scope as data access, data storage, data usage, and consent.

After identifying relevant stakeholders and the governance dimensions, the lead researcher reviewed various media sources illustrating deployment bias harms associated with Ring Doorbell. The primary goal was to provide participants with narratives that would elicit intuitive reactions and discussion. Storytelling was particularly useful for this purpose because it allowed non-experts to engage with abstract AI governance issues through familiar social situations. Given the specific cultural context of Yoruba people, suitable real-world narratives for this group were difficult to identify. As a result, the researcher developed two fictional narratives (see Appendix C). The first narrative features Nigerian Immigrant Bolaji, who is misidentified as suspicious by neighbors after being recorded speaking Yoruba on a doorbell camera, leading to racialized suspicion and police involvement. The second narrative features Funke, a Nigerian mother whose joyful cultural celebration is misinterpreted by surveillance systems and law enforcement, resulting in her being wrongly flagged for potential child welfare concerns. These narratives and involved stakeholders will be instrumental in the *Recognize* and *Analyze* stages.

Ọmọ̀lúàbí encompasses a broad constellation of values. Drawing on a process analogous to Munn's adaptation of Mead's five tests for evaluating technology through Māori principles [16],

the researcher identified seven Yoruba values from within the *Omọlúàbí* framework as most directly relevant to the harms and governance questions raised by Ring’s surveillance practices: *Òtítọ́* (Truthfulness), *Ogbón/Ìmọ́* (Wisdom and Knowledge), *Ìtẹ̀ríbá* (Respect), *Ìwà Pẹ̀lẹ̀* (Good Character), *Ìnú Rere* (Good Mind Towards Others), *Ojuse Àwùjọ́* (Social Responsibility), and *Sùúrù* (Patience). The aforementioned values guide ideation during the *Center Culture* and *Engage* stages of GRACE. Importantly, these are living values that shape how Yoruba communities understand moral obligation and dignity in relation to data and surveillance. Therefore, these seven Yoruba values provided a lens to operationalize GRACE and evaluate its efficacy. To validate the selection and accuracy of translations, the researcher consulted with members of the Yoruba community organizations used for recruitment, confirming that the chosen values and their English translations accurately reflected their cultural meaning and significance.

Finally, having established the preceding four foundational elements (technical scope, stakeholder selection, governance dimensions, and cultural values), the researcher identified the governance output the artifact participants would collectively produce through the workshop series. The selected output was a Data Rights Charter articulating protections for Yoruba people against the harms of Ring camera surveillance. This output was chosen because it was feasible within the constraints of a single-day workshop, requiring non-experts to apply the AI literacy foundation and cultural value analysis established across both sessions to produce a practical and actionable governance artifact.

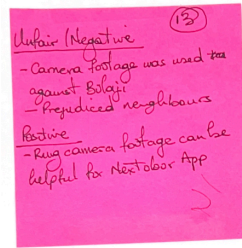


Fig. 1. Narrative 1 Annotation, Participant 3

3.3.2 Recognize. Session 1 began with the *Recognize* stage of the GRACE framework, and was led by the lead facilitator with the 9 community participants engaging with the GRACE framework for the first time.

This session began with an AI literacy component to educate the non-experts about the governance scope and understand how facial recognition works in Amazon Ring Doorbells. To start the AI literacy component, the lead facilitator highlighted 4 important definitions: consent, mass surveillance, facial recognition, and privacy. The lead facilitator selected these four terms because they represented foundational concepts that non-technical participants would need in order to meaningfully engage with the narratives and governance discussions. Community participants then watched a video that explained how facial recognition is used in Amazon Ring Doorbells and how the data it captured can be used. After the video played, the lead facilitator addressed key points about Ring that were expressed in the video: data storage, continuous monitoring, and consent; to reinforce the common themes in the video to allow participants to ask questions and address any gaps in understanding. This segment was essential so that non-experts understand the AI domain and device that they would be working on, to then be able to critique the Ring device from a shared baseline of knowledge.

After the AI literacy component, community participants began Activity 1: Narrative Analysis. Each participant individually read both narratives and recorded their annotations on sticky notes. Specifically, they were asked to make note of what they perceived as fair, unfair, or uncertain. Figure 1 shows an example of an annotation.

Once done individually, participants discussed what they noticed in small groups and moved sticky notes to unified chart paper. This activity surfaced potential social harms associated with Ring, and participants successfully identified critical points within the narratives that disproportionately affected different groups. Since the community partners are non-experts, it was important that the lead facilitator taught them the correct terms for the harms identified after initial critique was performed. The lead facilitator explained that these harms (racial profiling and over-surveillance) are associated with deployment bias and how informal and formal policing were vehicles for these harms against specific stakeholders present in the narratives. Since the narrative was fictional, the lead facilitator also shared two real-world examples of new articles that reported the same types of harms so that participants can understand that these are real world issues to address.

Community Participants were then given a 10 minute break before proceeding with the next stage of GRACE: *Analyze*.

3.3.3 Analyze. Following the group discussion of power and stakeholders within the narrative context, the lead facilitator began Activity 2: Power Mapping Analysis. Power mapping serves to make visible the relationships of influence, control, and vulnerability within a socio-technical system. Participants individually completed this activity, ordering the five identified stakeholders from least to most powerful to engage relationally with how power operates within Ring's surveillance infrastructure. Figure 2 shows an example of a completed Activity 2 worksheet.




Fig. 2. Power Mapping Analysis, Participant 13

After completing the power mapping, the lead facilitator opened a whole group discussion focused on: Is there a relationship between power and harm? How is that evident in your power mapping? Each participant agreed that the Target(s) of Surveillance had the least power, but there were mixed mappings on the most powerful within: Developers, Law Enforcement, and the Camera Owner. Stakeholder power rankings varied across participants, with law enforcement, platform users, and camera owners each positioned at the higher end of the power spectrum depending on the harm context under consideration. Notably, participants converged on placing the target of surveillance as the least powerful stakeholder—a shared recognition that primed them for Session 2, where they would produce governance rights centered on protecting that same target.

The Analyze stage was the final part of Session 1, which was one hour in length. Participants were given a 30 minute break after Session 1.

3.3.4 Center Culture. The *Center Culture* stage launched Session 2. Session 2 was led by the co-facilitator and observed by the lead facilitator. This stage opened with a brief recap of Session 1 before proceeding to its two core components: a culturally situated icebreaker designed to ground participants in shared Yoruba identity, and a cultural analysis activity in which participants examined how elements of the narratives, from the *Recognize* stage, aligned with or contradicted specific Yoruba values. This stage was designed by the lead researcher to begin the cultural translation process.

The icebreaker was modeled after a Restorative Justice Circle, a practice rooted in Yoruba communal tradition that centers collective sharing, relational dialogue, and community accountability. Restorative Justice Circles predate and are distinct from Western legal frameworks that later adopted similar structures [22]. Participants were asked to share one Yoruba tradition or value that deeply resonates with them, and to explain what it means to them. Participants shared Yoruba proverbs and cultural traditions that held personal significance to them, grounding the group in a sense of shared cultural identity and collective empowerment before transitioning into governance production. The selection of a restorative justice circle was intentional since community participants spent Session 1 identifying harms and power asymmetries that disproportionately affect Yoruba surveillance targets, the icebreaker served to reorient participants toward the strength and richness of their cultural tradition as the foundation from which governance rights would be produced.

The co-facilitator then presented Activity 3: Assessing Ring Cameras from a Yoruba perspective. In small groups, community participants selected a Yoruba value from the provided word bank and returned to the narratives from Session 1, identifying specific moments where the depicted treatment of Yoruba surveillance targets aligned with or contradicted the ethical standard set by their chosen value. The values chosen across the three groups include: *Ojuse Awùjo* (Social Responsibility), *Inú Rere* (Good Mind Towards Others), and *Ìwà Pẹ̀lẹ̀* (Good Character), with participants in two groups also incorporating *Ìtẹ̀riba* (Respect) into their deliberations. Figure 3 shows an example of small group notes produced during Activity 3.




Fig. 3. *Ojuse Awùjo* Group Activity 3 Assessment Notes

Following small group deliberation, the co-facilitator led a whole group discussion in which each group shared their findings, surfacing areas of alignment and contradiction between Yoruba values and Ring’s surveillance practices. This collective synthesis reoriented the harm and power analysis from Session 1 through a cultural lens, ensuring all participants entered the *Engage* stage with a shared Yoruba framework for critiquing the AI system from which governance rights would be produced.

Community Participants then received a 10 minute break before ending Session 2 with the final stage of the GRACE framework– *Engage*.

3.3.5 Engage. The final stage of the GRACE Framework *Engage*, focused on the community participants developing a governance output together. The co-facilitator led participants through Activity 4: Deliberation for Data Rights charter. Activity 4 consisted of 3 deliberation phases to create a Data Rights Charter focused on protecting Yoruba people from the potential harms of Ring Doorbells. These three phases included: individual deliberation prep, small group convergence, and whole group deliberation.

The co-facilitator explained the governance scope to community participants to guide deliberation. The intended scope focused on data rights for Yoruba people as targets of surveillance data collection anywhere in the world. To anchor the discussion, the co-facilitator posed a key guiding question: *How should that data be handled and protected?* This question was scoped across four domains: data access, data storage, data usage, and consent.

After establishing the context for the data rights charter, the co-facilitator presented a formula for constructing data rights: Because of [Yoruba Value] ✕ Yoruba people deserve [this protection]. This formula was recommended as a guide for rights creation; however, participants were not required to adhere strictly to it. Participants were instructed to develop “good” data rights that met three criteria: (1) Rooted in Your Value— the connection to the selected Yoruba value should be clear and explicit; (2) Specific and Actionable— the right should be implementable in a real-world context; and (3) Protective— the right should meaningfully prevent harms, such as those experienced by Bolaji and Funke in the scenarios. Finally, the co-facilitator provided an example to illustrate the expected output: *Because of Sùúrù (patience/caution), biometric data captured of any person must not be shared with law enforcement for at least 72 hours, allowing time for verification and human review before action is taken.*

Each community participant was then given a worksheet that communicated the intended scope and right creation guidance provided by the facilitator. The worksheet was designed to be fillable, requiring participants to individually propose one to two data rights and explicitly identify the Yoruba value they were translating. An example worksheet is shown in Figure 4.

★ TASK
Write down 1 – 2 rights stemming from your chosen cultural value(s). Specify the chosen value and then design a data right in the lines below.

Chosen value: Inu re
Because of inu re (caution of good intention), biometric
data or footage should be thoroughly inspected, including
employing language translation where applicable, allowing
for verification before any action are taken.

Chosen value: Inu re
Because of inu re, Yoruba people who are subjects
of biometric data or footage should be approached under
the assumption of innocence, to allow them view said
footage(s) and discuss important contexts before any
inference are made.

Fig. 4. Data Right Worksheet, Participant 2

This activity served as preparation for small group deliberation. During the small group deliberation, participants worked within their groups to merge their individual contributions and collectively settle on a core set of one to two data rights that reflected their shared value.

The *Engage* stage culminated in the production of the Data Rights Charter, the primary governance output of the two-part workshop series. Each small group presented their proposed rights to the full group, explaining the Yoruba values and reasoning behind each right. Community participants critically engaged with one another's proposals, affirming shared principles, challenging language, and building on each other's ideas throughout the facilitated discussion. The final charter consisted of seven collective data rights (see Appendix A). Once the charter was finalized, the co-facilitator led a closing reflection inviting participants to consider what changes Ring would need to make to its platform and policies to comply with the rights they had produced.

3.4 Researcher Positionality

As a Yoruba-American researcher raised outside Nigeria, the first author occupies both insider and outsider positions in relation to this study. Their Yoruba identity and membership in each of the three community organizations from which participants were recruited provided cultural familiarity and facilitated trust, grounding the research in existing community relationships. At the same time, their upbringing outside Nigeria and training within Western academic and technical institutions shape their perspective in ways that may diverge from participants raised within Nigeria. To address this, the first author engaged Yoruba cultural experts during the research design process to support the cultural grounding of the narratives and workshop materials.

For Session 2, the first author stepped back from facilitation and a co-facilitator from the research team, who is Yoruba and raised in Nigeria, led the session to center community voice in the workshop process and to assess whether GRACE could be operationalized without the first author's direct facilitation. Session 2 contained the main subject of study: cultural translation. Due to this, the first author's role as designer and lead facilitator of Session 1 and observer in Session 2 was intentional, enabling assessment of the framework's operation beyond the first author's direct facilitation.

This study was approved by Brown University's Institutional Review Board. Participants provided informed consent prior to participation and received gift card compensation upon completion.

3.5 Analysis

The dataset consists of multi-phase artifacts from the *Engage* stage of Workshop 2, which include:

- Phase 1: individual participant worksheets articulating one to two data rights grounded in each group's assigned Yoruba value
- Phase 2: data rights articulated in small groups based on individual preparation
- Phase 3: full group deliberation

Written artifacts were digitally scanned and imported into NVivo 15 for analysis, while the full group deliberation audio was recorded using a smartphone and transcribed using Otter.ai before import. We selected these three artifacts because together they constitute the rights-production process. Analyzing the primary source documents allows us to trace how participants translated cultural values into governance language after engaging with the *Recognize*, *Analyze*, and *Center Culture* stages of GRACE.

Codebook development followed a hybrid inductive–deductive approach informed by Braun and Clarke's reflexive thematic analysis [6]. A deductive baseline was first established through

review of Ring’s privacy policy and relevant Western privacy literature on data governance. Inductive open coding was then conducted across all artifacts through three iterative passes in NVivo, allowing Yoruba governance codes to emerge from participant data. Emergent codes were reviewed against existing African data ethics scholarship, including Barrett et al. ’s framework and Oyeshile’s work on Yoruba moral philosophy, to ensure alignment with documented non-Western value systems rather than relying solely on researcher interpretation [1, 23]. This process produced nine dominant Yoruba governance codes, presented in Table 2.

Table 2. Codebook

Code	Definition
Western Legal Rights	Participant use of American legal frameworks and vocabulary to articulate data protections for surveillance targets, including rights to privacy and due process.
Reciprocal Cultural Exchange	The assertion that Yoruba people should have the right to share their cultural values with other groups, and that governance frameworks should enable mutual exchange of values across cultures rather than imposing a single standard.
Target of Surveillance Control over Data	The right of surveillance targets to control the deletion, storage, and use of their own audio, visual, and biometric data, including the ability to withhold consent for use of their data in AI model training.
Collective Duty to Protect	The principle that responsibility for repairing harm caused by surveillance data falls collectively on all stakeholders who hold that data, rather than on the accused individual, including the obligation to destroy data upon a finding of innocence.
Reputation Protection	The right of Yoruba persons to be protected from reputational harm caused by surveillance data, requiring that any accusation be grounded in proper evidence and that data tied to those found innocent be destroyed.
Presumption of Innocence	The principle that Yoruba persons who are subjects of biometric data or footage should be approached with the assumption of good intentions, and afforded the opportunity to view footage and provide context before any inferences or actions are taken.
Cultural Identity Recognition	The right of Yoruba people to be identified as Yoruba by AI systems, treated as a foundational precondition for all other culturally grounded data rights to be enacted.
Surveillance Target Notification	The right of surveillance targets to be informed that they are being recorded, placing accountability on camera owners to acknowledge the personhood of those within a recorded space.
Cultural Representation in Design	The principle that AI systems should embed Yoruba cultural nuances, including language and gesture, directly into design controls so that the system can accurately interpret and contextualize the behavior of Yoruba people.

Code	Definition
Cultural Context Evaluation	The principle that data collected on Yoruba persons must be evaluated within the Yoruba value system, requiring that evaluators either understand Yoruba culture or that a culturally competent person be present during any assessment or investigation.

Following coding, quantitative analysis was conducted to examine patterns of code distribution across participants and the values that anchored their small group in the Engage stage. A matrix coding query in NVivo mapped each code’s frequency across the Engage stage artifacts and identified which participants contributed each reference. Individual worksheets and audio-transcript references were attributable to specific participants directly. References from the Data Rights Charter, a collectively refined product, were attributed to the participants in the small group from which each right originated, preserving the connection between articulated rights and their value-grounded source.

Finally, to surface patterns across participants, codes, and value groups, we produced two complementary visualizations: a bipartite network mapping participants to the codes they contributed to (color-coded by assigned value), and heatmaps showing reference frequencies at both the individual and group level.

4 Findings

This section will present findings from the Yoruba case study, a two-session workshop series, conducted with 10 participants, facilitated by two facilitators. None of the participants were AI experts, and one participant served as a peer facilitator. The main findings are organized around the three research questions:

- RQ1. Does GRACE enable non-experts to critique AI systems and make informed decisions about governance frameworks? (Section 4.1)
- RQ2. Do the governance frameworks produced through GRACE meaningfully diverge from Western models centered on individualism and capitalist logics? (Section 4.2)
- RQ3. Can GRACE function as an accessible and replicable toolkit for culturally grounded AI governance in low-resource contexts? (Section 4.3)

These findings draw on artifacts from the Activity 3: Assessing Ring Cameras from a Yoruba perspective and Activity 4: Deliberation for Data Rights charter, analyzed through thematic analysis to produce the codebook presented in Table 2. Network visualizations and participant heat maps were used to further interpret the emergent themes. Definitions of Yoruba values that guided deliberation can be found in Appendix E.

4.1 Critical Engagement and Informed Decision-Making

In the Recognize and Analyze stages, participants were introduced to key components of the socio-technical system surrounding Amazon Ring cameras. This included foundational AI literacy on consent, mass surveillance, facial recognition, and privacy; an explanation of how Ring’s facial recognition operates and how captured data can be used; and a guided introduction to the stakeholders involved (Target of Surveillance, Camera Owner, Platform Users, Developers, and Law Enforcement) and the harms that can arise, including racial profiling and over-surveillance as forms of deployment bias. Following these stages, the *Center Culture* and *Engage* stages required participants to make decisions about data rights to rebalance power and address harms to the Target(s) of Surveillance. As shown in Figure 5, two themes emerged as dominant across all groups: Cultural

Representation in Design and Cultural Context Evaluation. These two themes represent two sides of the same coin and together reflect participants’ understanding of the underlying system and the stakeholder power dynamics surfaced through power mapping. See Appendix D for the isolated network visualization for each group.




Fig. 5. This visualization maps each participant’s references, with edges linking them to the labeled codes. Compound participants drew on respect as an additional value alongside their group’s.

4.1.1 Cultural Representation in Design. Cultural Representation in Design covers culturally informed AI tools that enable the technical system itself to account for cultural context, as well as the inclusion of cultural experts in the design process. Participant 5 suggested that developers could build LLMs grounded in Yoruba culture, to identify “language and certain gestures of the tribe” and shape how the AI system interprets surveillance data. Two groups contributed data rights tied directly to culturally informed AI controls. The *Ìwà Pẹ̀lẹ̀* (good character) group built on Participant 5’s individual data right, contributing a right focused on the AI system understanding the gestures and language of Yoruba people, serving as a guide for stakeholders and users. The *Inú Rere* (good mind towards others) group contributed a right stating that biometric data or outages should be thoroughly inspected, including through language translation where applicable, allowing for verification before any action is taken. These rights can be found in Appendix A. In deliberation, Participant 11 emphasized that developing culturally informed AI tools requires incorporating Yoruba people in the design process before models are built, calling out developers for excluding them (see Appendix B.4.1).

These rights demonstrate informed decision-making because participants did not stop at identifying harm; they traced it to its source in the Ring camera design pipeline and proposed targeted interventions at both the model level (training on Yoruba language and gesture) and the process level (including Yoruba people as co-designers). Non-experts articulated, in their own terms, that bias is produced downstream of deployment and must therefore be addressed there.

4.1.2 Cultural Context Evaluation. Cultural Context Evaluation covers what should happen after data is accessed from the system and used by stakeholders to take action. This directly addressed the upstream of deployment. Participants emphasized that evaluation should provide cultural context to those reviewing footage so they can make better decisions about what is occurring. The culturally informed AI tools surfaced in Cultural Representation in Design are critical here, as they directly support the cultural competency of evaluators. The *Ojuse Awujo* (social responsibility) group contributed a data right explicitly stating that “data collected must be evaluated within

the Yoruba value system.” In deliberation, Participant 3 argued that someone from the Yoruba culture should be present to aid evaluation so that cultural context is understood. Participant 3 extended the right further to call for Yoruba representation within the institutions that access the data, such as law enforcement (see Appendix B.4.3).

This theme reflects informed decision-making because participants identified evaluators as a distinct point of failure in the socio-technical pipeline, separate from developers and from camera owners. They reasoned that even a culturally aware AI system can produce harm if the humans interpreting its outputs lack cultural competency. Their proposed rights therefore intervene at the level of who performs evaluation within powerful institutions, calling for Yoruba people to occupy evaluator roles in entities such as law enforcement that access and act on the data.

Together, these two themes show that participants engaged with Ring not as a single device but as a socio-technical system with multiple intervention points. By producing rights that target design, evaluation, and institutional representation, participants demonstrated they could trace harms relationally across stakeholders, applying what they learned in power mapping to identify where each stakeholder’s power needed to be rebalanced. This level of analysis is significant because it shows that non-experts, when supported by GRACE’s scaffolding, are capable of engaging with and critiquing AI systems at a depth typically reserved for technical or policy specialists.

4.2 Cultural Translation to diverge from Western Governance Models

Each cultural value group anchored its deliberation in one of three Yoruba values: *Ìwà Pẹ̀lẹ́* (good character), *Inú Rere* (good mind towards others), or *Ojuse Áwùjọ* (social responsibility). As shown in Figure 6, the participant heat map surfaces a dominant code within each group, representing the data right most strongly tied to that value’s logic. The dominant codes were Reputation Protection for *Ìwà Pẹ̀lẹ́*, the Right to Yoruba Identity for *Inú Rere*, and Surveillance Target Notification for *Ojuse Áwùjọ*. The following sections analyze each in turn.

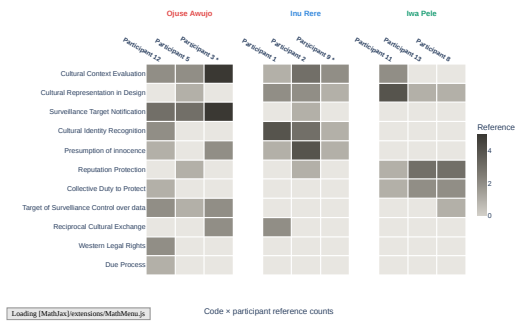


Fig. 6. This heatmap depicts code by participant Reference Counts, with participants organized by their chosen value.

4.2.1 Target Of Surveillance Notification. The *Ojuse Áwùjọ* (Social Responsibility) group produced a data right centered on notification; targets of surveillance should be made aware they are being recorded (see Appendix A). This right emerged consistently across both deliberation and individual responses (see Appendix B.1.1). Participant 5 framed notification as a corrective measure to the harms experienced by Funke and Bolaji in the case studies, arguing that awareness of being recorded would, “...inform how they behave” and prevent downstream harm. Participant 3 noted that without notification, recorded actions can be stripped of their intended meaning, leading to

misinterpretation—for instance, practices like prayer may be incorrectly flagged as “suspicious” behavior. The same participant distilled this into a principle of respect, “Out of respect, people should know they are being recorded.”

In the current design of Ring cameras, only the camera owner knows who is being recorded and retains unilateral control over how that data is shared. The proposed right reframes this dynamic, placing accountability on the camera owner to uphold their duty to society, by informing others that they are within a recorded space. This shifts surveillance from a one-sided act of observation into a relational practice, holding the camera owner morally accountable for acknowledging the personhood of those being recorded, rather than treating them as passive data subjects. Furthermore, this right also reflects an informed deliberation about harm. Participants recognized that notification serves a protective function for both parties: the recorded individual can adjust their behavior or contextualize their actions, while the camera owner avoids contributing to wrongful accusations rooted in cultural misreading.

To assess whether this translation produced a genuinely non-Western data ethic, we validated the resulting data right against Barrett et al.’s codebook of African data ethics principles. The resulting data right maps directly to the principle of **Preserve the Dignity of Data Contributors** from the African Data ethics, which holds that every person involved in the data lifecycle should be respected, that individuals should not be treated as a means to execute data work, and that collective agreements and communal well-being must be honored [1]. This alignment between Barrett et al.’s emphasis on preserving dignity and the *Ojuṣe Àwùjò* group’s articulation of target-of-surveillance notification suggests that this data right to notification is rooted in African data ethics and reflects a non-Western approach to data governance.

4.2.2 Reputation Protection. A code that emerged distinctly from the *Ìwà Pẹ̀lẹ̀* group concerned the protection of Yoruba persons’ reputations from harm caused by surveillance data. The group’s data right states: “Any accusation made by a stakeholder against a Yoruba person needs to be grounded in proper evidence to avoid accidentally tarnishing the reputation of Yoruba persons” (see Appendix A). This right was synthesized from both group and individual responses (see Appendix B.2.1). The group’s cultural narrative analysis articulated the underlying logic directly: “Reputation is slandered when the narrative violates this standard” (see Appendix A). Participant 2, agreed with this claim, refusing to be reduced to a single recorded moment: “I have my own culture, and you will not just assume the worst of me based on what I display at that point in time.” Participant 13 extended the right to its logical end, arguing that data tied to those found innocent should be destroyed because it otherwise persists on their record and continues to shape how they are perceived. In Yoruba philosophy, reputation and good character are deeply intertwined: *Ìwà Pẹ̀lẹ̀* is not a private virtue but one demonstrated and affirmed through how a person is known within their community. Surveillance data that distorts this knowing constitutes a direct harm to *Ìwà Pẹ̀lẹ̀* itself.

These participants rejected the rational reduction of personhood to data points, attending instead to the deployment bias and reputational harm that surface when surveillance is enacted without good character. Furthermore, the data right produced placed accountability on stakeholders making accusations to substantiate them with evidence, and on the data system itself to remove records that no longer serve a legitimate purpose. Participants therefore positioned reputation not as a private concern but as a communal good worth protecting through the design of the surveillance system. Like the *Ojuṣe Àwùjò* (Social Responsibility) group, these participants oppose Ring’s existing policy and further increase data agency for targets of surveillance.

We confirmed that this right represents a non-Western ideology, as it aligns with a principle in Barrett et al.’s African data ethics codebook: **Mend Data-Driven Harms**. Mend Data-Drive harms

calls for reconciliation processes that place disenfranchised communities in positions of power and hold perpetrating parties accountable [1]. Barrett et al. emphasize accountability and repair as central to ethical data practice, particularly for communities historically subject to misrepresentation through extractive data systems. The study participants' reasoning reflects this principle by holding accusers accountable and repairing harm by destroying data that could continue to damage an innocent person's reputation. This cultural translation of *Ìwà Pẹ̀lẹ́* therefore surfaces African data ethics by recognizing that data carries the power to harm communities and must be governed accordingly.

4.2.3 Cultural Identity Recognition. The *Inú Rere* (Good mind towards others) group produced a data right stating that Yoruba people deserve to be identified as Yoruba within the system (see Appendix A). Participants framed this right as a precondition for all others. Participant 1 reasoned that if a system already differentiates rights by categories like age or gender, then "in order to actually guarantee Yoruba people any rights, the system needs to be aware of the fact that they're Yoruba." Participant 12, from the *Ìwà Pẹ̀lẹ́* group, agreed during deliberation and extended this logic, explaining that recognition signals respect: "you're being acknowledged... you even know who Yoruba people are." See Appendix B.3.1 for excerpts from these participants.

These responses reframe identification not as a technical labeling problem but as a cultural and ethical prerequisite. Without knowing who is being observed, a stakeholder cannot act with *Inú Rere*, because good intentions are culturally situated. Good intentions must consider what counts as care, respect, or harm in the community in which a person is known. The right therefore positions cultural recognition as the foundation on which any subsequent governance rights become meaningful.

This data right aligns with Barrett et al.'s principle of **Design in Solidarity**, which calls for embracing diverse data ethics standards and resisting the imposition of dominant frameworks on communities at the margins [1]. Participants insisted that Yoruba identity be legible within the system rather than absorbed into a rational, and implicitly Western, user category. This move asserts a non-Western ethic in which cultural specificity is treated as the starting point of governance, not an afterthought.

4.3 Accessibility and Replicability in Low-Resource Contexts

This research question assesses the implementation of the GRACE framework for this case study, and whether that methodology is feasible for communities operating in low-resource contexts. A low-resource context, as used in this work, refers to a setting that lacks the institutional infrastructure often assumed by participatory AI governance research, including dedicated funding, large facilitation teams, and access to technical experts or digital research tools. Two features of the Yoruba case study reflect the feasibility of the GRACE framework: the limited expertise required to design and run the workshop, and the low barrier to participation for community members.

4.3.1 Limited expertise required to facilitate. The lead facilitator, who is Yoruba, developed the case-study narratives independently and refined them with informal feedback from Yoruba people in their personal community. This positioning is important: the facilitator must come from the culture being centered, both to ground the narratives in lived cultural knowledge and to ensure that the workshop is accountable to the community it serves. No team of AI subject-matter experts (SME) was needed beyond this. GRACE narratives can therefore be built with as few as one SME on the AI system being governed, provided that SME is also a cultural insider — or paired with one. Each activity in the workshop can be written by the facilitator in advance, requiring no additional resources beyond the facilitator's own preparation time.

4.3.2 Low barrier to participation. Participants did not need internet access, laptops, or any technical infrastructure of their own; the facilitator carried the responsibility of sourcing materials and preparing the workshop. Participants needed only to be able to read the narratives and to know their own cultural values. As demonstrated in Sections 4.1 and 4.2, this was sufficient for non-experts to translate cultural values into data rights. The cultural value chosen (*Ìwà Pẹ̀lẹ́, Inú Rere, Ojuse Àwùjọ*) did not introduce new content for participants to learn; rather, they provided shared terminology for traditions and moral concepts participants already lived by, ensuring that everyone in a group was reasoning from the same named reference point within Yoruba thought.

4.3.3 Replicability through Ground stage inputs. GRACE is replicable across contexts because the five foundational elements established during the *Ground* stage function as inputs that can be swapped out for different communities and AI systems. In this case study, the lead facilitator selected Amazon Ring as the technology governance scope, identified five stakeholders informed by literature on Ring's harms to black and brown communities, defined four governance dimensions (data access, data storage, data usage, and consent), drew seven values from the *Ọmọlúàbí* framework, and chose a Data Rights Charter as the governance output. None of these selections are intrinsic to GRACE itself. A different facilitator could anchor the workshop in a different AI system (predictive policing, hiring algorithms, generative AI), identify the stakeholders relevant to that system to populate the narratives, define governance dimensions specific to its harms, draw cultural values from a different epistemology, and select a governance output suited to their community's needs. The *Ground* stage therefore operates as a configuration layer: the structure of GRACE remains constant while its content adapts to the cultural and technological context being governed.

These features of the Yoruba case study show that GRACE can function as a replicable toolkit in low-resource contexts: the design work concentrates expertise and resource access in a small facilitation team rooted in the culture being centered, while participation remains open to any community member with cultural knowledge and basic literacy.

5 Discussion

5.1 Contributions

This work contributes to decolonial AI, data sovereignty, and participatory AI governance by introducing GRACE, a framework that enables technical non-experts to make informed AI governance decisions through cultural translation. Existing AI literacy programs aimed at non-experts tend to be resource-intensive, relying on technical experts, one-on-one mentorship, paid crowd-sourcing infrastructure, or capital to develop digital educational products [10, 15, 21, 24, 25]. GRACE addresses this gap by demonstrating that meaningful AI governance participation can be achieved with a single facilitator drawn from the community itself and minimal infrastructure. This matters because the everyday people most affected by AI systems should have the conditions to understand how their data is used and to shape its governance. In the case study used for this evaluation, these are the people recorded by Ring cameras, or otherwise positioned as data subjects.

A second contribution is empowerment through cultural translation. Participants produced an inaugural data rights charter rooted in *èniyàn*, the Yoruba conception of personhood as constituted through relationships, community, and moral conduct, and grounded in the moral obligations of *Ọmọlúàbí*, the ethical ideal of a person of good character. Few cultural translations from Global South epistemologies exist in the AI governance literature [1, 12, 14, 20], and this study contributes a concrete account of what people from one such epistemology want from AI, irrespective of what the institutions deploying these technologies want. The findings presented in Section 4.2 show this translation was substantive rather than surface-level. The rights produced aligned with principles

in Barrett et al.'s African Data Ethics codebook, including Mend Data-Driven Harms, Preserve the Dignity of Data Contributors, and Design in Solidarity, indicating that participants did not merely inflect Western frameworks but produced governance grounded in a distinct ethical tradition.

The *Center Culture* stage in particular surfaced a dynamic worth naming explicitly. As participants moved into cultural translation, participants began code-switching into Pidgin English, drew on jokes legible only within Yoruba culture, and expressed audible pride in being Yoruba. This was not captured in the written artifacts but was consistent across the workshop. It suggests that cultural translation is not only an analytic method but a generative one, creating space for participants to value their own epistemology and to see it as legitimate ground for governing AI. In doing so, GRACE pushes back on the colonial legacy embedded in dominant AI governance frameworks, which tend to treat Western individualist and capitalist logics as the default starting point.

Finally, GRACE is extensible by design. While this study focused on Yoruba cultural values, the framework is structured to accommodate any group with a shared ethical or moral tradition, including communities, institutions, activist groups, or other cultural collectives, paired with a relevant governance scope. This positions GRACE as infrastructure for surfacing alternative governance frameworks across contexts, not a single-use intervention.

5.2 Limitations

Several limitations bound the claims made here. The workshop was condensed from its intended four hours and fifteen minutes to three hours and thirty minutes because participants arrived late, which limited the depth of stakeholder analysis and technical engagement with the AI system. GRACE is likely to produce richer results when run within groups that meet regularly, such as workplace departments, neighborhood associations, religious congregations, or activist groups. Such groups can scale the workshop across multiple sessions and build toward more technical outputs like design requirements, consent protocols, or data strategies, rather than a single charter. AI literacy studies conducted over weeks or months consistently produce more developed participant outputs [15, 26, 27], and GRACE would benefit from similar extension.

The study was also conducted in English, though the community's primary language is Yoruba. Yoruba is a highly expressive language built on phrasal meaning rather than single-word denotation, and conducting the workshop in English likely flattened some of the cultural reasoning participants could otherwise have articulated. Pidgin English, which participants used, carries Yoruba influence and cultural meaning. Running the workshop in Yoruba or Pidgin English could produce insights more deeply aligned with African data ethics and Yoruba thought, since participants would be reasoning in the linguistic register through which their values are most fully expressed.

5.3 Future Work

We hope this work invites researchers to run GRACE with communities grounded in other Global South epistemologies, so that nations and groups around the world can begin shaping AI in ways that align with their own ethics rather than absorbing imported defaults. Comparative work across multiple GRACE deployments could surface patterns in how different cultural traditions approach AI governance. It could also press the field to develop AI/ML methods that incorporate cultural specificity rather than deploying one-size-fits-all models [1, 14, 17–19]. Indigenous data sovereignty projects have already demonstrated what becomes possible when communities reshape technology around their own values [8, 20], and continued work in this direction could meaningfully expand who AI is built to serve. Ultimately, this research aims to support communities long silenced by rationalist dominance in reshaping what counts as "safe" AI for their own contexts.

Acknowledgments

I would like to express my deepest gratitude to my advisor, Dr. Julia Netter, for being an inspirational and impactful during this journey. Her guidance, support, and unwavering belief in this project made all the difference. I am also profoundly grateful to my mother, Auntie Toun, and Grandma for their invaluable cultural insights and wisdom. This work would not have been possible without their generosity in sharing their knowledge and lived experiences with me. I would also like to acknowledge the use of Claude (Anthropic) to assist in adapting my writing for an academic audience. This support was necessary due to a short-term disability, as the lead researcher was experiencing post-concussion syndrome during the entire case study and writing process. All ideas, arguments, and analysis presented in this work are entirely my own and were written prior to being refined with AI assistance for academic tone and clarity.

References

- [1] Teanna Barrett, Chinasa T. Okolo, B. Biira, Eman Sherif, Amy Zhang, and Leilani Battle. 2025. African Data Ethics: A Discursive Framework for Black Decolonial AI. In *The 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*. ACM, Athens, Greece, 1–16. <https://doi.org/10.1145/3715275.3732023>
- [2] Stevie Bergman, Nahema Marchal, John Mellor, Shakir Mohamed, Iason Gabriel, and William Isaac. 2024. STELA: a community-centred approach to norm elicitation for AI alignment. *Scientific Reports* 14 (March 2024), 6616. <https://doi.org/10.1038/s41598-024-56648-4>
- [3] D. Bertram. 2022. Accounting for Culture in Policy Transfer: A Blueprint for Research and Practice. *Political Studies Review* 20, 1 (2022), 83–100.
- [4] Abeba Birhane. 2021. Algorithmic injustice: a relational ethics approach. *Patterns* 2, 2 (February 2021), 100205. <https://doi.org/10.1016/j.patter.2021.100205>
- [5] Abeba Birhane, William Isaac, Vinodkumar Prabhakaran, Mark Díaz, Madeleine Clare Elish, Iason Gabriel, and Shakir Mohamed. 2022. Power to the People? Opportunities and Challenges for Participatory AI. In *Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '22)*. ACM, Arlington, VA, USA, Article 6, 8 pages. <https://doi.org/10.1145/3551624.3555290>
- [6] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- [7] Dan Calacci, Jeffrey J. Shen, and Alex Pentland. 2022. The Cop In Your Neighbor's Doorbell: Amazon Ring and the Spread of Participatory Mass Surveillance. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–47.
- [8] Stephanie Russo Carroll, Marisa Duarte, and Max Liboiron. 2024. Indigenous Data Sovereignty. In *Keywords of the Datafied State*, Burrell, Singh, and Davison (Eds.). Data & Society. <https://datasociety.net/library/keywords-of-the-datafied-state/>
- [9] Madiha Z. Choksi, Marianne Aubin Le Quéré, Travis Lloyd, Ruojia Tao, James Grimmelman, and Mor Naaman. 2024. Under the (Neighbor)hood: Hyperlocal Surveillance on Nextdoor. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, Honolulu, HI, USA, 1–22. <https://doi.org/10.1145/3613904.3641967>
- [10] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2023. The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. In *Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '23)*. ACM, Boston, MA, USA, 1–23. <https://doi.org/10.1145/3617694.3623261>
- [11] Catherine D'Ignazio and Lauren Klein. 2020. The Power Chapter. In *Data Feminism*. MIT Press. <https://datafeminism.mitpress.mit.edu/pub/vi8obxbh7>
- [12] Xiao Ge, Chunchen Xu, Daigo Misaki, Hazel Rose Markus, and Jeanne L. Tsai. 2024. How Culture Shapes What People Want From AI. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, Honolulu, HI, USA, Article 95, 15 pages. <https://doi.org/10.1145/3613904.3642660>
- [13] K. Hao and H. Swart. 2022. South Africa's private surveillance machine is fueling a digital apartheid. *MIT Technology Review* (April 2022).
- [14] Lilly Irani, Janet Vertesi, Paul Dourish, Kavita Philip, and Rebecca E. Grinter. 2010. Postcolonial computing: a lens on design and development. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 1311–1320. <https://doi.org/10.1145/1753326.1753522>
- [15] Jared Ordoña Lim, Grace Barkhuff, Jane Awuah, Sophie Clyde, Riya Sogani, Christina Gardner-McCune, David Touretzky, and Judith Odili Uchidiuno. 2025. Escape or D13: Understanding Youth Perspectives of AI through Educational Game Co-design. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM, Yokohama, Japan, 1–17.

- <https://doi.org/10.1145/3706598.3714037>
- [16] Hirini Moko Mead. 2016. *Tikanga Maori: Living by Maori Values*. Huia Publishers, Wellington.
 - [17] Sabelo Mhlambi. 2020. *From Rationality to Relationality: Ubuntu as an Ethical and Human Rights Framework for Artificial Intelligence Governance*. Technical Report. Carr Center for Human Rights Policy, Harvard Kennedy School, Harvard University. <https://data-ethics.jonreeve.com/static/papers/mhlambi2020rationality.pdf>
 - [18] MIT Technology Review. n.d.. AI Colonialism. MIT Technology Review. <https://www.technologyreview.com/supertopic/ai-colonialism-supertopic/>
 - [19] Shakir Mohamed, Marie-Therese Png, and William Isaac. 2020. Decolonial AI: Decolonial Theory as Sociotechnical Foresight in Artificial Intelligence. *Philosophy & Technology* 33, 4 (2020), 659–684. <https://doi.org/10.1007/s13347-020-00405-8>
 - [20] Luke Munn. 2024. The five tests: designing and evaluating AI according to indigenous Māori principles. *AI & Society* 39 (2024), 1673–1681. <https://doi.org/10.1007/s00146-023-01636-x>
 - [21] Wilhelmina Nekoto, Vukosi Marivate, Tshinondiwa Matsila, Timi Fasubaa, Taiwo Fagbohunge, Solomon Oluwole Akinola, Shamsuddeen Muhammad, Salomon Kabongo Kabenamualu, Salomey Osei, Freshia Sackey, Rubungo Andre Niyongabo, Ricky Macharm, Perez Ogayo, Orevaoghene Ahia, Musie Meressa Berhe, Mofetoluwa Adeyemi, Masabata Mokgesi-Seling, Lawrence Okegbemi, Laura Martinus, Kolawole Tajudeen, Kevin Degila, Kelechi Ogueji, Kathleen Siminyu, Julia Kreutzer, Jason Webster, Jamiil Toure Ali, Jade Abbott, Iroro Orife, Ignatius Ezeani, Idris Abdulkadir Dangana, Herman Kamper, Hady Elshahar, Goodness Duru, Ghollah Kioko, Murhabazi Espoir, Elan van Biljon, Daniel Whitenack, Christopher Onyefuluchi, Chris Chinenye Emezue, Bonaventure F. P. Dossou, Blessing Sibanda, Blessing Bassey, Ayodele Olabiya, Arshath Ramkilwan, Alp Öktem, Adewale Akinfaderin, and Abdallah Bashir. 2020. Participatory Research for Low-resourced Machine Translation: A Case Study in African Languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, 2144–2160. <https://doi.org/10.18653/v1/2020.findings-emnlp.195>
 - [22] Don John Omale Phd. 2006. JUSTICE IN HISTORY: AN EXAMINATION OF ‘AFRICAN RESTORATIVE TRADITIONS’ AND THE EMERGING ‘RESTORATIVE JUSTICE’ PARADIGM. 2 (11 2006).
 - [23] Olatunji Oyeshile. 2021. Yoruba Philosophy of Existence, Iwa (Character) and Contemporary Socio-Political Order. *Philosophia: International Journal of Philosophy* 22, 1 (2021), 1–18. <https://doi.org/10.46992/pijp.22.1.a.1>
 - [24] Ethan Plaut, Kiri West, Fabio Morreale, Maya Gibson, Grace Thompson, Kara Woodward, and Danielle Lottridge. 2025. Surveillance on Exhibit: Using Problematic Technology To Teach About Problematic Technology. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM, Yokohama, Japan, 1–19. <https://doi.org/10.1145/3706598.3713710>
 - [25] Garrett Smith, Sarah Carson, Rhea Vengurlekar, Stephanie Morales, Yun-Chieh Tsai, Rachel George, Josh Bedwell, Trevor Jones, Mainack Mondal, Brian Smith, Norman Makoto Su, Bart Knijnenburg, and Xinru Page. 2024. “I Know I’m Being Observed”: Video Interventions to Educate Users about Targeted Advertising on Facebook. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, Honolulu, HI, USA, 1–27. <https://doi.org/10.1145/3613904.3642885>
 - [26] Jaemarie Solyst, Emily Amspoker, Ellia Yang, Motahhare Eslami, Jessica Hammer, and Amy Ogan. 2025. RAD: A Framework to Support Youth in Critiquing AI. In *Proceedings of the 56th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE TS 2025)*. ACM, Pittsburgh, PA, USA, 1–7. <https://doi.org/10.1145/3641554.3701966>
 - [27] Yu Sun, Qichao Dong, and Fang Tang. 2025. Integrating Socially Responsible Computing through Direct Community Engagement in CS2 to Promote Latinx Student Retention. In *Proceedings of the 56th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE TS 2025)*. ACM, Pittsburgh, PA, USA, 1–7. <https://doi.org/10.1145/3641554.3701895>
 - [28] Harini Suresh and John Gutttag. 2021. A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle. In *Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '21)*. ACM, New York, NY, USA, 1–9. <https://doi.org/10.1145/3465416.3483305>
 - [29] United Nations General Assembly. 2026. Declaration of the trafficking of enslaved Africans and racialized chattel enslavement of Africans as the gravest crime against humanity (A/80/L.48). <https://docs.un.org/en/A/80/L.48>

APPENDIX

A Data Rights Charter

The following Data Rights Charter was collaboratively developed by participants during the final *Engage* stage of the GRACE framework. The charter represents a culturally grounded governance artifact that translates Yoruba values into data rights to govern AI-powered surveillance systems and protect targets of surveillance.

- (1) Because of *Inú Rere*, Yoruba people who are subjects of biometric data or surveillance footage should be approached under the presumption of innocence. Individuals should be allowed to view the footage and discuss relevant contextual information before any inferences or actions are made.
- (2) Because of *Ojuṣe Àwùjọ*, Yoruba subjects of surveillance should know when they are being recorded and how their data is being stored, used, and shared.
- (3) Because of *Ojuṣe Àwùjọ*, data collected through surveillance systems must be evaluated within the Yoruba value system.
- (4) Because of *Iwà Pèlè*, any accusation made by a stakeholder against a Yoruba person must be grounded in proper evidence to avoid unjustly tarnishing the reputation of Yoruba individuals. Proper evidence should be corroborated by two or more eyewitnesses. If the accused person is found innocent, all related visual, audio, and biometric data should be destroyed by all parties who received or stored the data, including first-, second-, and third-party stakeholders. Responsibility for data removal should not fall on the accused individual.
- (5) Because of *Iwà Pèlè*, developers may build localized intelligence models (LIMs) grounded in Yoruba culture and identity, including culturally specific language patterns, gestures, and social nuances, in order to provide contextualized guidance and behavioral interpretation for users and stakeholders.
- (6) Because of *Inú Rere*, biometric data and surveillance footage should be thoroughly inspected and contextualized prior to any action being taken, including the use of language translation where applicable.
- (7) Because of *Inú Rere*, Yoruba people deserve the right to be recognized and identified as Yoruba people within technological systems.

B Participant Artifact Excerpts Informing the Data Rights Charter

This section contains excerpts from artifacts produced during the *Center Culture* and *Engage* stages that directly informed each right in the data rights charter. Entries are organized by cultural value group and mapped to the themes they shaped. The final section includes excerpts from codes that converged across all values.

B.1 Ojuṣe Àwùjọ

B.1.1 Surveillance Target Notification.

“Out of Respect, people should know they are being recorded”

(Participant 3, Individual Data Right for Iteriba)

“I don’t think anyone is more important than the other, because I think that the right to know is, is very important. I don’t want to be doing something, you know. I think, Oh, this is very cool with me, just to find somebody to, you know, and eventually know that somebody takes an offense with it like that’s why the right, the right to know that I’m being surveilled, is important. So I’m not maybe playing or maybe praying in a certain way I feel is, you know, allowed in my culture, and somebody records me and say he’s acting suspicious.”

(Participant 3, Deliberation Audio)

“I think that the first right directly answers the question in the context of Funke and Bolaji, they were unaware that they were being recorded. So making it known to them that, hey, you have been in a space where you know you’re being recorded,

that would save us a lot of headache. At the end of those case studies. So yeah, I think it would [...] inform how they ... inform how they behave,”

(Participant 5, Deliberation Audio)

B.1.2 Target of Surveillance Control over Data.

“[T]hat decision should be on the individual in question, not on any other person... if you have been found innocent, you should decide whether you want to use the data to train AI in the future or delete it outrightly.”

(Participant 12, Deliberation Audio)

“and also education, yeah, because all this thing is happening in almost a silo within this context. But I think it’s so important for consumers, or just people, to be because people who buy these things don’t know, who have been surveilled don’t know ... that they even have any protections. So yes, education is [gestures towards facilitators] Thank you. Education is paramount. Thank you. Thank you so much.”

(Participant 12, Deliberation Audio)

B.1.3 Western Legal Rights.

“Right to privacy: notice, written caution (similar to “beware of dog”) w/ such a device; notice/warning should be visible to public, so that public is aware”

(Participant 12, Individual Data Rights)

“Right to due process – fairness beyond visual and/or audio, evidence must present including subject (of surveillance) testimony without prejudice based on culture, nationality, ethnicity, and language”

(Participant 12, Individual Data Rights)

B.2 Iwà Pèlé

B.2.1 Reputation Protection.

“Reputation is slandered when the narrative violates this standard.”

(Cultural Narrative Analysis)

“I have my own culture, and you will not just assume the worst of me based on what I display at that point in time.”

(Participant 2, Deliberation Audio)

“So if somebody’s found innocent, it should be destroyed [...] because it affects people’s records. Even though, like, they’ve been found guilty, this thing [can] still [be] on their record.”

(Participant 13, Deliberation Audio)

B.2.2 Collective Duty to Protect.

“Having witnesses to good/moral character.”

(Cultural Narrative Analysis, Iwà Pèlé)

“Coming together to protect each other and neighborhood”

(Cultural Narrative Analysis, Ojuṣe Àwùṣò)

“It should not fall on the responsibility of the accused person to have to destroy their data. For the second third parties, all the first, second and third parties who have it should be in charge of destroying it, because now that creates a burden on the person who’s accused. And I’ll follow up with every single agency. Say, Hey, is my agent? Is my data with you? I want to destroy it.”

(Participant 8, Deliberation Audio)

B.3 Inú Rere

B.3.1 Cultural Identity Recognition.

“[Being identified as Yoruba] definitely shows that you know you’re being acknowledged, you’re being respected. It’s built in - like you’re respecting, acknowledging us, so you even know who Yoruba people are.”

(Participant 12, Deliberation Audio)

“If the system can identify someone as a male or female or young or old, I think having and ... give those different categories or identifiers certain rights, then I think in order to actually guarantee Yoruba people any rights, the system needs to be aware of the fact that they’re Yoruba, and so that as a baseline, I just thought that was really, really important.”

(Participant 1, Deliberation Audio)

B.3.2 Presumption of Innocence.

“In the first scenario, the neighbors likely had ‘Inú Rere’ which made them release the footage. On the other hand, law enforcement did not, which led to the unfortunate outcome.”

(Cultural Narrative Analysis, Inú Rere)

“It also kind of translates to the assumption of innocence which wasn’t the case in either scenario.”

(Cultural Narrative Analysis, Inú Rere)

“Bringing concerns or accusations to the accused as an engaging party instead of disregarding.”

(Cultural Narrative Analysis, Ojùṣe Àwùjọ)

“Because of Inú Rere (Good mind towards others), biometric data captured of any Yoruba person should be assumed to be of good intentions, until otherwise disputed by additional evidence from other sources.”

(Participant 2, Individual Data Rights)

B.4 Convergent Codes

B.4.1 Cultural Representation in Design.

“And I guess mine was more building... it’s like, as we said, into the design process controls. So when, so you ... when a real video developing the next iteration or software, you’re forced to put into consider... into consideration the nuances of Yoruba people. Wait so, so if for So, the user, any user across the world... can make any false assumptions. So those controls are already built into the app. So it’s like, Ah, this person is committing to crime, the use case, or the app can say, maybe prompt or give, like a warning or some type of statement saying ... so the same

way you're saying person walking right now, the ring I can show you ... if it sees me walking like boy walking on the corridor or [person] was walking. It's gonna say, Yoruba man playing with this child. It's not gonna say, Oh, Yoruba man is being or Yoruba woman is being suspicious, being suspicious or being manhandling the child. So it provides, so because it has that built into the app. It provides actual, true, better context, and it's safer, right? That's all."

(Participant 12, Deliberation Audio)

"Developers can build LLM based on culture, identifying nuances of Yoruba culture; with regards to language; and certain gestures of the tribe to provide controls; when observing behavior & and can serve as a guide for users."

(Participant 11, Individual Data Rights)

"Because of ojuse awujo, the developers should create a feature that translates Yoruba language to understandable computer language for the A.I to translate Yoruba and understand Yoruba gestures."

(Participant 5, Individual Data Rights)

"So having us as part of the research before building these models is critical, because they're so lazy when thinking of Africans or black people in general. So just them doing extra due diligence and understanding the nuances of black folk in general, I think, is critical. So hopefully that changes in the future, or now it's already too late."

(Participant 12, Deliberation Audio)

B.4.2 *Reciprocal Cultural Exchange.*

"And so I think it's only fair that if certain values are granted to Yoruba people, perhaps exclusively other groups of people. I think Yoruba people should have the right to share their values with other groups of people as well, and perhaps vice versa. Maybe there's a certain value that Chinese people have that maybe Yoruba people don't have it would be nice to be able to combine values to guarantee the best outcomes for your people, but also all people in general"

(Participant 1, Deliberation Audio)

"Some of these ideas doesn't necessarily apply only to Yoruba people. It would also be good ... good, good, good. I mean, it would be good social responsibility to also look at the same thing towards other people. I mean towards other cultures, you know, like we're saying that Yoruba Igbo or the hausas or the ibibios or the ethnic people too, should know some of these things they should ... if you're going to investigate or accuse or judge them, it should be done within the context of your culture, so that you presume them innocent before until proving guilty."

(Participant 3, Deliberation Audio)

B.4.3 *Cultural Context Evaluation.*

"the idea is that if you are going to investigate or throw accusations or want to understand, you need to do it to somebody that understands the culture, or, yeah, somebody that does understand the culture must be present, or must be the one who get evaluated."

(Participant 3, Deliberation Audio)

C Case Study Narratives

This appendix contains the two case-study narratives used in the workshop. The narratives were introduced during the *Recognize* stage and re-examined during the *Center Culture* stage.

C.1 Narrative 1: Bolaji

Bolaji was walking through his neighborhood in San Francisco on a lovely Fall afternoon. The sun was shining and the sky was blue as he made his way home from the corner store. Bolaji had moved to California six months ago from the UK, where he worked as a legal representative for a pharmaceutical company. His new job allowed him to work remotely, and he'd recently rented a home in this quiet residential area.

Feeling homesick, he called his mother in Lagos, Nigeria.

"Mummy mo ti de ile!"¹ he exclaimed cheerfully.

He recounted his day as he walked down the tree-lined street, asking after her health and general wellbeing.

"Se alafia le wa, Mummy?"² he asked warmly.

Continuing to speak animatedly to his mother, he paused in front of a neighbor's house to look at his phone for directions. His mother kept saying, "Soro soke now! Mi o gbo nkan ti o n'so!"³ because of the connection delay.

Bolaji raised his voice, speaking more loudly in Yoruba as he gestured with his free hand, checking the map on his phone. He didn't notice the Ring camera mounted beside the front door recording him.

The next day, a video appeared on the neighborhood watch app Nextdoor. It showed Bolaji standing in front of the house, speaking loudly in a foreign language and gesturing. The caption read: "ALERT: Suspicious man casing homes on Maple Street. Speaking in unknown language, acting agitated, appeared to be planning something. Has anyone else seen him? Should we call the police?"

Within hours, multiple neighbors commented: "I saw him too—definitely suspicious," "Call the cops before something happens," "Better safe than sorry."

That evening, as Bolaji walked to check his mail, three neighbors approached him.

"Excuse me, we need to talk to you," one said firmly, blocking his path. "We've seen you wandering around here, and we want to know what you're doing in our neighborhood."

Bolaji, confused, responded in his British accent, "I live here. Is there a problem?"

"We've got you on camera acting suspicious. You don't look like you belong here," another neighbor said, pulling out her phone. "We're calling the police."

Before Bolaji could explain, one neighbor was already dialing 911.

C.2 Narrative 2: Funke Taiwo

"E don happen Ooooo!"⁴ Funke was in her front yard, jumping up and down joyously. She danced from side to side, real "ko mole"⁵ style, as her 3-year-old son laughed and hopped around with her. "I don get alert, God Win! Aaaaaah!"⁶

Looking at her phone, she re-read the message that changed everything—after five years of waiting, she had been approved for U.S. citizenship!

¹Mum, I've arrived home!

²Mummy, I hope you are well.

³Speak up! I can't hear what you are saying!

⁴Finally! It has happened!

⁵Drop it like it's hot (Nigerian slang for dancing energetically)

⁶I've received the notification, God has won!




Fig. 8. Ojuṣe Àwùjọ




Fig. 9. Inú Rere

E Yoruba Value Descriptions

The following section provides definitions of the three Yoruba values selected by participants during the *Center Culture* stage of the GRACE framework: *Ojuṣe Àwùjọ* (Social Responsibility), *Ìwà Pẹ̀lẹ̀* (Good Character), and *Inú Rere* (Good Mind Towards Others).

E.1 Ojuṣe Àwùjọ

A core value of *omolúàbí* is *Ojuṣe Àwùjọ*, which translates literally to “duty to society” or social responsibility. In this context, social responsibility entails that individuals, groups, or businesses have a moral obligation to act in ways that benefit society and protect the environment. In order to uphold social responsibility, an individual’s choices support the well-being of others, not just oneself. Key components of social responsibility in Yoruba context aim to: foster the creation of safer and supportive communities, foster trust and respect, promote fairness, and encourage sustainable development. In practice, this can look like volunteering, helping neighbors in need, or businesses supporting the local community through donations or projects.

E.2 Ìwà Pẹ̀lẹ̀

Ìwà Pẹ̀lẹ̀ is another core value of *omolúàbí*, translating to “good character”, which reflects virtues of gentleness, humility, and patience. In Yoruba culture, there is a proverb that states “Character is beauty”. This emphasizes the importance of good character in the Yoruba cultural context, stating that good character is the greatest beauty a person can have. A person that emulates good character should not be hard or aggressive, speaks to others with wisdom, is tolerant and is patient.

Additionally, good character also encompasses showing respect to elders. This value is specifically outward, reflecting how an individual should act towards others.

E.3 Inú Rere

Inú Rere is another core value of ọmọ́lúabí that works alongside *Ìwà Pẹ̀lẹ́* (good character) to shape a person who is honorable, peaceful, and loved by society. *Inú Rere* translates to "good mind towards others," and a person who embodies this value carries a kind and good heart filled with compassion, love, and goodwill toward others. Unlike good character, which is outward, a good mind toward others is inward, focusing on the inner intentions of a person. Examples of *Inú Rere* include not harboring wickedness, desiring good things for others, showing mercy, and offering genuine care.