

COACT – A Community-centered, Participatory and Actionable Roadmap for Equitable Language AI

Ulla Petti

Language Technology Lab, University of Cambridge, English Faculty Building, 9 West Road, Cambridge, CB3 9DA, United Kingdom, ump20@cam.ac.uk
<https://orcid.org/0000-0003-0203-9956>

Hannah M. Claus¹

Language Technology Lab, University of Cambridge, English Faculty Building, 9 West Road, Cambridge, CB3 9DA, United Kingdom, hmc78@cam.ac.uk
<https://orcid.org/0009-0006-1881-8689>

Anna Barford¹

Language Technology Lab, University of Cambridge, English Faculty Building, 9 West Road, Cambridge, CB3 9DA, United Kingdom, ab423@cam.ac.uk
<https://orcid.org/0000-0002-4250-7171>

Malak Sadek

The Centre for Human-Inspired Artificial Intelligence (CHIA), University of Cambridge, 16 Mill Lane, Cambridge, CB2 1SB, United Kingdom, mfzas2@cam.ac.uk
<https://orcid.org/0000-0001-8284-890X>

Roi Reichart

Technion Israel Institute of Technology, Haifa, 3200003, Israel, roireichart@gmail.com
<https://orcid.org/0000-0001-6918-0554>

Anna Korhonen

Language Technology Lab, University of Cambridge, English Faculty Building, 9 West Road, Cambridge, CB3 9DA, United Kingdom; and The Centre for Human-Inspired Artificial Intelligence (CHIA), University of Cambridge, 16 Mill Lane, Cambridge, CB2 1SB, United Kingdom alk23@cam.ac.uk
<https://orcid.org/0000-0002-3692-3144>

Language AI risks deepening global inequities, particularly for speakers of under-served and marginalized languages. Although participatory and community-centered approaches have been proposed across several disciplines, the guidance remains scattered and difficult to apply in a systematic way. This paper integrates insights from Natural Language Processing, Human-Computer Interaction, international development, and participatory research to propose a unified roadmap for equitable Language AI. We outline five core principles—centering community needs, supporting local capacity, ensuring sustainability, promoting transparency and fairness, and embedding continuous reflection—and translate them into a practical five-stage process that covers planning, fieldwork, development, deployment, and long-term maintenance. By drawing these elements into a clear and actionable structure, the roadmap provides concrete guidance for researchers and developers seeking to build Language AI that is contextually grounded, collaboratively shaped, and globally inclusive.

CCS CONCEPTS • Human-centered computing ~ Collaborative and social computing • Software and its engineering ~ Software creation and management ~ Collaboration in software development • General and reference → Surveys and overviews

¹ Equal contribution

Additional Keywords and Phrases: Participatory research, Co-research, Community-centered research, Natural Language Processing, Human Computer Interaction

1 INTRODUCTION

Language AI refers to a wide range of technologies - from machine translation to Large Language Models (LLMs) - that rely on Natural Language Processing (NLP), speech technologies, and related methods. These technologies now underpin many everyday applications that help people communicate, learn, and access information. Yet their benefits remain unevenly distributed: current Language AI disproportionately serves dominant, well-resourced languages and communities, leaving many linguistic groups globally underrepresented and reinforcing existing social inequities. In this paper, we introduce a participatory roadmap that brings together insights across disciplines to enable the development of Language AI that is globally equitable, accountable, and responsive to the needs and priorities of diverse communities.

The need for such an approach is increasingly recognized across the field. A growing body of work calls for Language AI that distributes benefits more fairly, better reflects real-world needs, and supports a wider range of languages and cultural contexts. For example, recent works highlight the opportunities for inclusive AI development [20, 44], make recommendations for socially responsible language data collection [62], and give suggestions for engaging with Indigenous communities [13, 21]. More structured approaches include frameworks for decolonial pathways in Human Computer Interaction (HCI) [5] and community-based [55] and participatory design-inspired NLP [16]. What has been missing is an overarching framework that integrates these insights into a unified approach that enables developers, organizations, policymakers, and communities to collaborate in ways that center the needs and priorities of intended users throughout the development of Language AI.

Centering the needs of the community and intended users in developing Language AI can have multiple benefits: (1) it can promote more equitable distribution of the benefits that language technology can bring, and serve the people whose needs and experiences might be overlooked in mainstream Language AI development (such as speakers of smaller non-dominant languages), contributing to increasing equality as a whole; (2) engaging the communities in the development process can boost the sense of ownership and empower people; (3) grounding the technology in local needs and contexts can increase its usability and help bridge the applicability gap.

These benefits can produce positive outcomes for many stakeholders. Participatory and inclusive approaches can increase the availability of accessible Language AI tools in users' native languages that are easy to use, responsive to real-world needs, culturally appropriate, and capable of capturing linguistic nuance. Such tools can help address challenges in areas like healthcare and education, benefitting individuals, governments, businesses, and those developing Language AI. Greater applicability can also expand reach and adoption of Language AI tools, strengthening economic value and creating business incentives for participatory practices. Human-centered and collaborative methods can deepen understanding of community needs and values, thereby increasing trust in decision-making - an outcome particularly valuable for governments and policymakers. Additionally, close engagement with local experts and intended users can reduce duplicated efforts across research teams and enable the creation of technology with greater real-world impact.

Although incentives to adopt participatory methods vary across stakeholder groups, and such approaches may not always align with short-term interests, it is essential to design systems in which all groups have clear reasons to participate. Advancing equity in Language AI requires contributions from all stakeholders, and a more equitable landscape strengthens the ecosystem for everyone. Centering communities therefore benefits all stakeholders.

Our roadmap is motivated by a central research question: How can Language AI be designed and implemented in an equitable way? By equitable, we refer to approaches that seek to meet the needs of everyone, including under-served

populations. To understand how equitable Language AI can be developed in practice, we reviewed literature from a range of disciplines - including HCI, NLP, and the social sciences - that address stakeholder engagement, fieldwork, and the practicalities of technology development, deployment, and maintenance. Although many researchers have noted the gap between Language AI's potential and its real-world usefulness, existing recommendations and best practices remain dispersed across fields and therefore often fragmented. We bring these insights together and introduce a roadmap that offers a holistic approach, supporting stakeholders through five stages: (1) planning, (2) fieldwork, (3) technology development, (4) deployment, and (5) maintenance.

This roadmap can help teams identify and plan for the key moments in a multi-step Language AI development process. Because specific needs vary across projects—and because the boundaries between stages are often fluid—it is intended to be used flexibly, prompting reflection on which activities are most appropriate at each point. In doing so, the roadmap supports our broader goal of fostering greater equity and enhancing the societal usefulness of Language AI. Beyond offering recommendations, we aim to contribute to shaping an ecosystem where the individual incentives of technology creators, mediators, and regulators pull in the same direction as the needs of users.

The outline of the paper is as follows: in Section 2 we share background literature on participatory development, placing our roadmap into wider context; Section 3 details the methodology we used to develop the roadmap; Section 4 presents COACT – A Community-centered, Participatory and Actionable Roadmap for Equitable Language AI and explains it in relation to the literature upon which it is built; Section 5 concludes the paper. Team positionality is addressed at the end of the paper.

2 BACKGROUND

Diverse disciplines have long examined equity and participation in social sciences and technology, offering insights relevant to equitable Language AI development. Participatory approaches to development, popularized in the 1980s-1990s, sought to empower disadvantaged groups by tailoring interventions to people's needs, culture, and aspirations, and countering top-down decision-making [18, 23, 37]. In computer-based interactive system design, an international standard under review at the time of writing [33], highlights the following benefits of human-centered and participatory methods: “Increasing the productivity of users and the operational efficiency of organizations; Being easier to understand and use, thus reducing training and support costs; Increasing usability (effectiveness, efficiency and satisfaction); Increasing accessibility (for people from a population with the widest range of user needs, characteristics and capabilities); Improving user experience; Reducing discomfort and stress; Providing a competitive advantage, for example by improving brand image; Contributing towards sustainability objectives” [33, p.5].

Anthropology and ethnography have greatly contributed to shaping participatory methods, emphasizing the knowledge of the communities and people's lived experiences, treating participants as co-experts and involving them in decision-making throughout the research process, the importance of in-depth engagement with the community and local context, building relationships, and co-designing interventions grounded in real world needs. These practices have informed the theoretical foundations and guided the practical approaches used in participatory research today.

Participation can come in many forms, as characterized by the ladder of participation [6]. In this ladder, the lowest rungs represent non-participation, such as manipulation and therapy; the middle rungs describe the more superficial and tokenistic forms of participation, such as informing, consultation and placation; the higher and more desirable rungs are partnership, delegated power and citizen control. Most participatory AI projects still “inform or consult, rather than partner with or delegate control to participants” [22]. This illustrates that meaningful participation can be slow and costly [46, 56], potentially leading to the use of more superficial approaches, which can lead to missed opportunities and the benefits of

participation being elusive [23]. For example, participatory work with young people in mental-health technology remains mostly short-term, consultative, and consumerist, with mismatched expectations between youth and designers [53].

Another challenge to inclusive participation is reaching marginalized communities, often resulting in engagement with intermediaries rather than those most affected [17, 19]. Careful consideration of how participation takes place is therefore essential from the outset [46, 56]. The design of M-PESA money transfer service operating in Africa illustrates the value of broad stakeholder input and locally grounded needs assessment, beginning with workshops in Nairobi (Kenya) and Dar es Salaam (Tanzania) engaging banks, micro-finance organizations, microcredit NGOs, telecoms and finance regulators, and technology suppliers. One of the project team reflected how “Sitting in a comfortable office in England and deciding what Africa needs is an approach doomed to failure.” [32, p.68]. With 51 million customers making \$314 billion transactions in a year [65], M-PESA is used widely.

HCI has long stressed designing technology around people’s social and cultural contexts. Recent work provides participatory guidelines for NLP, including community involvement in data collection, modelling and evaluation [16], culturally grounded frameworks such as METAL [40], and arguments for community-based NLP to address missing cultural context [55]. HCI also studies user experience with Language AI tools, examining trust, feedback and adaptability [27]. By evaluating tools in real context with qualitative studies and usability tests, HCI research helps identify gaps that algorithmic performance metrics overlook. HCI research on language learning and preservation show that early engagement with local speakers and prioritizing social and cultural context improves alignment with community needs, empowers users and supports wellbeing [1, 12, 41, 57].

NLP research has widely discussed representation and inclusion challenges [36, 63]. Recent work also increasingly highlights the importance of socially responsible, human-centered data practices [62]. Participatory approaches long used in linguistics and language documentation have emphasized collaboration with language communities and explored questions of data ownership – topics especially relevant for linguistically sensitive Language AI in underserved contexts.

Developing equitable Language AI requires integrating insights from all relevant disciplines — from best practices in project planning and fieldwork, to data collection and model development, deployment and long-term maintenance.

3 METHODS

To develop our roadmap, we first surveyed the state-of-the-art literature in participatory approaches to Language AI development. Our particular focus was on key challenges, emerging opportunities, and proposed recommendations. Our research question was: How can Language AI be designed and implemented in an equitable and inclusive manner? To this end, we draw on interdisciplinary approaches from the NLP and HCI communities, as well as from broader AI research and social sciences.

Our literature search encompassed the following academic databases: Web of Science, ACM Digital Library, Elsevier, and Google Scholar, along with major conference proceedings (ACL, EMNLP, NAACL, LREC, FAccT, CHI, NeurIPS). This search extended to the fields of NLP, HCI, broader AI research and social sciences, with a focus on participatory and responsible methods, decolonizing approaches, and ethics. The survey prioritized literature published between 2020 and 2025 to include the most recent developments including but not limited to LLMs, while also incorporating foundational work prior to 2020 that has significantly influenced current thinking.

Particular attention was given to including work by authors from the relevant language communities or potential user communities, recognizing the critical role of local agency in participatory design. This often required deliberate effort, as much of the existing literature in this field is produced by scholars in the Global Minority. We included geographic areas of under-served language communities in search terms (Africa, South America, Oceania), focused on authors affiliated

with local institutions, and included literature from local conferences and initiatives (such as Deep Learning Indaba, AfriCHI, AfricaNLP, SEACrowd, AmericasNLP). While we aimed to prioritize the work and perspectives of authors who are familiar with, or embedded in, the context they study or influence, we acknowledge an important limitation: the literature search was conducted exclusively in English, and only English-language publications were included in the review. This constraint likely restricted the number of contributions from the Global Majority that could be identified and incorporated into this study.

We extracted the recommendations and best practices that addressed the core challenges presented in the reviewed articles: (1) the availability and diversity of language data (including underrepresentation of languages, lack of diverse annotators and inclusive annotation standards, bias towards English, and economic and academic drive towards rich countries and languages); (2) the practical relevance and accessibility of AI applications (including lack of practical and real-life grounded applications, reinforcement of structural inequalities); (3) cultural and linguistic sensitivity (including cultural abstraction and the expectation to generalize, lack of inclusive evaluation metrics and culturally diverse benchmarks, limited inclusion of native speaker knowledge, local perspectives and localized development, treating language as data source); (4) ethical and sociopolitical implications (including consent for data collection and ownership, lack of actionable pathways to decolonizing, inconsistent government support, uneven distribution of risks and benefits of AI to underrepresented communities). We focused on the opportunities and practical advice to tackle these challenges, engage stakeholders and develop human-centered Language AI.

Through the review, we encountered a wide range of recommendations and frameworks that inform our roadmap. While significant research has been conducted in this area, our literature survey indicates that existing actionable solutions often remain fragmented, highlighting the need for a comprehensive and overarching roadmap that consolidates best practices and facilitates the transition of research into practice. This paper responds accordingly by developing A Community-Centered, Participatory and Actionable Roadmap for Equitable Language AI. Taking an iterative approach, the team have worked and reworked the roadmap, drawing upon learnings from the review combined with their own subject expertise and experiences. By bringing together perspectives from different disciplines and integrating them into one overarching and actionable roadmap, we aim to provide a practical tool to support equitable Language AI development.

4 COACT - A COMMUNITY-CENTERED, PARTICIPATORY AND ACTIONABLE ROADMAP FOR EQUITABLE LANGUAGE AI

In this section, we present the Community-Centered, Participatory and Actionable Roadmap for Equitable Language AI, bringing together recent recommendations and best practices from a variety of perspectives. See Appendices A.1 for the overview of the roadmap. In what follows, we present the roadmap and detail key insights from the literature survey which underpins it. We introduce five overarching principles that should be applied throughout the development process, after which we turn to the sequence of development, consisting of planning, fieldwork, development, deployment, and maintenance. Although the roadmap is presented in a linear sequence, it is intended to be applied flexibly to accommodate a variety of projects and contexts. The fluidity of stages in AI development and the non-linearity of participatory design in language technology development has also been described in previous work [16].

4.1 Principles

Five overarching principles that stem from the wider literature and are relevant throughout the entire technology development process include: (1) engaging the community and other stakeholders and focusing on their needs; (2)

promoting knowledge exchange; (3) prioritizing sustainability and managing benefits and risks; (4) promoting transparency, privacy and fairness; and (5) reflecting and documenting throughout the process.

4.1.1 Engage the community and other stakeholders throughout, focusing on their needs and the local context

The importance of engaging the community and focusing on their needs when developing Language AI has been stressed by many researchers [13, 14, 15, 16, 20, 21, 26, 40, 43, 44, 49, 50, 52, 55, 62, 64]. To start, it is useful to discuss the terms *community* and *stakeholders*, and what is meant by them.

The apparently uncontentious term *community* is worth problematizing. Since the 19th century, two distinct meanings have been attached to the term *community*. One referring to “natural local neighborhoods” with a focus on conformity and conservatism, and the other as a more progressive and democratic alternative [59]. The term *community* can be somewhat homogenizing given that most groups – be they defined by location, identity or interests – have internal heterogeneity, diverse concerns, and their own power structures. To some, community creates an identity, a sense of belonging, and depending on language, location, and culture, it can change its meaning. The term *community* can also influence the way people engage with problems and their solutions [47].

When planning to engage communities, it is important to identify the community and recognize heterogeneity and the diverse needs and priorities which may exist. It is also important to remember that the term *community* can have a different meaning to different stakeholder groups. Positioning the community at the center of each project and prioritizing their language, culture and needs will challenge the current approaches often centering the Global Minority, creating space for new solutions that have not had a chance to emerge thus far [43, 55], and contribute to creating tools that are meaningful and practically applicable.

Stakeholders is a term often used in business, to refer to all those who have “a stake” or an interest, including employees, workers in the supply chain, managers, customers and the society in which a business operates [9]. Regarding Language AI, stakeholders should be considered for the following non-exhaustive list of actors: the Government (local, regional or national), authority figures and gatekeepers (religious or spiritual leaders, elders, community leaders, influencers), those who provide the infrastructure and possibly funding for deployment (businesses, government, international organizations), experts and possible project partners (NGOs, academics, international organizations), intended users of the new technology, and members of the language community. Each project will need to identify its stakeholders. A detailed approach to stakeholder identification is described by [60].

Stakeholder collaboration can begin with the joint planning of the initiative, continue through shared decision-making during the development and deployment stages, and extend to sharing responsibility for maintenance. Adopting an inclusive and participatory approach, in which authority over decisions is shared, can foster a sense of ownership [42, 56]. The early and sustained involvement of individuals with lived experience, as well as local experts, enhances the relevance and contextual appropriateness of the work, increases trust [35], and can help avoid redundancy of efforts [62].

4.1.2 Prioritize sustainability, maximize benefits and manage risks

Language AI development should be guided by a thorough and ongoing evaluation of both potential benefits and risks, promoting environmental, social, and economic sustainability of the technology. Stakeholder engagement in risk assessment could include providing an overview of how NLP models are translated into real-world settings and discussion of the potential risks, benefits, and impacts of deployment [55]. Transparent communication is essential for informed decision-making and trust.

The ethical strategies aimed at promoting socially beneficial AI (outlined by [35]) include aligning AI with human values (see [39, 67] for how heterogeneous community values can be incorporated into algorithm design and governance), collaborating with those impacted by the technology, advancing scientific knowledge, decentralizing power structures, and using influence to promote human rights. Strategies to mitigate potential harms can include establishing clear legal accountability, addressing the root causes of harm, encouraging internal reporting of unethical practices (whistleblowing), fostering diversity within development teams, and integrating ethics education into formal training programs [35].

It is widely recognized that technological development should be balanced with sustainability and energy efficiency, with the aim of minimizing ecological impacts and avoiding disproportionate burdens on marginalized communities [10, 26, 30, 35, 62]. Contemporary Language AI systems can have a significant environmental footprint due to the high energy demands of large-scale computational infrastructure [8, 30]. One practical step to mitigate this is to measure and report the environmental impact of AI systems, for example, by tracking carbon emissions and energy consumption (see framework in [29]), assessing the environmental impact of NLP systems (see overview of tools in [8]), and improving reporting practices [58]. Environmental considerations should not be limited to the training phase alone; they should be addressed throughout the project lifecycle, including during planning, fieldwork (e.g., travel), deployment, and long-term maintenance. The importance of political and social engagement to reduce the carbon footprint of data-driven technologies is emphasized by [11].

4.1.3 Engage in knowledge exchange, support and extend existing projects and AI training

The value of developing local talent and contributing to ongoing, community-driven AI initiatives is widely acknowledged across literature [4, 14, 20, 26, 31, 43, 44, 52, 54, 62]. Local capacity-building strengthens the long-term impact of technological interventions and supports more inclusive, contextually grounded AI development [54]. Creating training opportunities at local and national levels and expanding existing initiatives can help build this capacity. Involving community members throughout the development process - through co-created tutorials, open-access materials, and culturally relevant resources - advances equity and inclusion, builds local expertise, and enables the transfer of day-to-day management to local teams [20, 43, 55]. These practices encourage knowledge exchange, support skills development aligned with local needs, and open pathways for locally driven AI projects, promoting local ownership, empowering those who may lack access to formal training, and enhancing the technology's long-term relevance and sustainability.

4.1.4 Promote transparency, privacy, and fairness

Transparency, privacy, and fairness are foundational principles to be upheld throughout all stages of Language AI development. Early-stage efforts can include raising public awareness of AI technologies, associated rights, and regulatory frameworks, and engaging in co-planning, co-design, and shared decision-making with a diverse range of stakeholders - particularly those directly impacted by the systems [35, 55]. Responsible development practices encompass transparent communication, informed and inclusive data collection, and safeguarding the interests of underrepresented groups when needed [44]. This also entails ensuring secure data storage and responsible use.

Technical mechanisms to protect privacy may include privacy by design, data minimization, and access controls [35]. From a regulatory perspective, legal compliance, adherence to AI-specific regulations, and the use of adaptive licensing models have been proposed to support ethical deployment [35, 52].

Transparency in model development is essential, particularly in identifying and mitigating biases that may compromise robustness, especially for underrepresented populations [44]. The use of diverse evaluation metrics that reflect the values and needs of affected communities can enhance fairness and utility. Furthermore, deployment should prioritize equitable

access and distribution of benefits. Where necessary, additional investment should be directed toward ensuring fairness, with a long-term view of societal benefit [44].

Finally, we recommend that AI developers actively engage in policy discourse to promote frameworks that support equitable outcomes. By centering transparency, privacy, and fairness, the development of Language AI can be more trustworthy, inclusive, and responsive to the needs of diverse communities.

4.1.5 Continuously reflect and document the process

As a final principle, we recommend incorporating ongoing reflection and thorough documentation throughout the development of Language AI systems. Engaging in continuous, collaborative reflection with stakeholders enables the identification of emerging opportunities and potential shortcomings, allowing developers to adapt to evolving societal needs, values, policy contexts, and technological landscapes. Critical evaluation of the intended purposes and impacts of the technology, underlying power dynamics, and past practices can foster a more inclusive and pluralistic development approach [5, 44]. Integrating "what-if" scenarios during development can contribute to understanding model behavior and enhance the fairness of outputs [55]. Meanwhile, systematic documentation not only supports reproducibility and transparency, but also facilitates ongoing maintenance and informs future projects. Thus, documentation can contribute to the long-term relevance and societal impact of the technology [10, 55].

4.2 Planning: understand local context, needs and AI readiness

When planning the development of Language AI, we recommend identifying and building connections with stakeholders, understanding the context in which the technology will be developed, used, and maintained, assessing the current role of technology, and obtaining necessary approvals.

4.2.1 Identify and build connections with stakeholders

Involving a diverse group from the start, including local experts, individuals with lived experience, and those whom the technology is intended to serve, supports the development of inclusive solutions that are responsive to local needs, fosters trust, broadens perspectives, and contributes to challenging narratives that prioritize Western knowledge systems [66]. This stage in participatory Language AI research has been described notably by [13], and is supported by other scholars [43, 62]. By approaching the project as equal partners, researchers can remain open to learning from the community and prioritize their perspectives and needs throughout the process.

It is necessary to acknowledge how colonial legacies continue to shape structures of advantage and disadvantage around the world, operating alongside the capitalist practices which create and reinforce new inequalities. It is in this context that more equitable Language AI is needed, yet this context also shapes the possibilities to achieve this. Current work to decolonize technology development involves working with non-dominant knowledge systems to reimagine the ideas and knowledge upon which our societies, economics and politics are based [24]. However, there is a risk that "the decolonization concept ... be emptied of its substance and instrumentalized by settler researchers and institutions." [7, p.645]. In terms of these risks, "Indigenous researchers are usually well aware of the risks of entering into research relationships, [however] it is often difficult for Settler researchers to fully grasp the depth, breadth, and long-lasting impacts of the harm done by research." [28, p.332]. When engaging stakeholders, including communities, it is important to recognize the inequities inherent within our relationships and promote equitable working relationships at every step.

4.2.2 Understand context: needs, priorities, AI readiness, power dynamics and the role of technology in society

Thorough research into the context in which the technology will be applied is needed before development begins. Such research would examine the stakeholders' specific needs and priorities [13, 38, 44], helping to identify the broad research focus. It is also relevant to assess AI readiness and existing power dynamics that may influence data collection and system design [5], and consider how and by whom the technology will be used, recognizing that its impact may vary across different communities [5, 44]. This contextual understanding can be developed through engagement with domain experts and local stakeholders, including community members and AI practitioners. Complementary desk research should draw on statistical data, academic studies, and grey literature to build a comprehensive understanding of the setting.

4.2.3 Assess existing technology and how it could be useful

Drawing on previous work [13, 14], we recommend that researchers identify existing technologies and datasets before developing new tools, and actively support local initiatives wherever possible. Furthermore, it is important to consider what communities are already motivated to achieve, how they exercise their agency, and how technology can meaningfully enhance and sustain these efforts.

4.2.4 Gain ethics and other approvals to undertake research

Researchers should ensure that all necessary approvals are secured before commencing development, including ethical clearance and any other required authorizations from relevant authorities, such as national ethics committees in both the host country and home institution. Local approval for community engagement should be sought to ensure alignment with local norms and expectations. When research involves human participants, appropriate documentation must be prepared to provide clear information and obtain informed consent in accordance with ethical standards.

4.3 Fieldwork: learn with and from the community

The goal of fieldwork stage is reaching a better understanding of local priority needs and the context for which new technology is intended. This stage also entails agreeing collaboratively on whether, what and how to develop, and on the project timeline, data ownership, and knowledge transfer. Learning with the community involves using participatory or co-research approaches to establish more equitable research relationships and a locally relevant research design [56]. Learning from the community requires the research team to foreground the knowledge, expertise and preferences of stakeholders to allow these to become central influences on the focus, design and uptake of new technologies.

4.3.1 Learn about local culture, economy, environment, and politics, and identify priority needs

This stage includes building understanding of local cultures and their history, to learn about the values, worldviews and traditions [43]. This in itself can be a multi-year undertaking, so we recommend also engaging with the writings of anthropologists, novelists and travel writers to learn from them. By better understanding people, their practices and their places, the team can learn how to address people appropriately and respectfully when collaborating on culturally meaningful tasks [14]. This cultural learning is also relevant for subsequent language-specific steps in this roadmap, as this early work builds understanding around the epistemology of language [*ibid.*].

Fieldwork offers an important opportunity to also engage in qualitative and ethnographic research, which complements and extends more data-driven approaches [62]. Useful qualitative methods include interviews and focus groups [55], and given the internal heterogeneity of communities, it is important to engage with diverse individuals. Such human-centered methods can adopt collaborative and participatory approaches [5], contributing to more equitable relationships with

stakeholders [56]. Crucially, fieldwork will enable the team to learn important information about the locality, including the main matters of concern [14]. To conduct fieldwork, it is valuable - if not essential - to involve local researchers, speakers of local languages, and government [43, 45, 46].

4.3.2 Decide whether to develop a technology and agree on roles and responsibilities

Findings from focus groups and interviews should be discussed collaboratively with all relevant stakeholders, including intended users and decision-makers. Based on these discussions, stakeholders should jointly determine whether and what to develop, and identify the types of data required for Language AI development. People with lived experience should have an active role in this process, which might involve voting on decisions about data sources [55] or consensus decision making. While defining roles and responsibilities before development begins can promote clarity, having fluidity in roles can also have its benefits [49]. Reaching an agreement is essential before embarking upon technology development [55].

4.3.3 Design an approach and timeline and plan how to sustain technology

Co-design how to approach technology development with stakeholders, including an agreement on the project's direction and discussion on potential outcomes [43]. Establishing a clear timeline and scheduling regular follow-up discussions can maintain important engagement with stakeholders throughout subsequent stages [55]. As part of this process, the team can plan the long-term sustainability and maintenance of the technology, including financing, so that preparations can be made with sufficient time.

4.3.4 Clarify data ownership, use and storage; and transfer of know-how

It is essential to discuss data ownership transparently with the community and to establish a clear, preferably written, agreement before proceeding with Language AI development. This should be prepared in the most suitable language for the stakeholders and done in a culturally sensitive and respectful way. A key principle here is that communities should have access to data and retain control over their use, supporting communities' right to self-determination [43]. The types of data to be collected, and its intended use and storage should be clearly communicated and mutually agreed [16, 55]. This process of enhancing community awareness of how personal data is used can help build trust [35], and care needs to be taken to explain this in a way that makes sense to stakeholders. Researchers should also prioritize transfer of knowledge and create meaningful learning opportunities and maintain clear agreements ensuring equitable collaboration [13].

4.4 Development: develop technology

Participatory methods throughout the technology development process can support shared decision-making, equitable data practices, and the creation of learning opportunities. When working with language data, it is essential to recognize and address the complexities and nuances, the cultural, historic, social, and political aspects inherent in language.

4.4.1 Collect data

Collected language data should be representative, diverse, and inclusive [3, 4, 31, 44, 61, 66]. Data collection, be it language or any other type of data, should respect individuals' privacy, anonymity, rights, consent, and data sovereignty [62]. Collaboration with communities, native speakers, linguists, and individuals from diverse demographic backgrounds can ensure that data practices are contextually grounded, representative, and ethical [44, 49, 64]. Implementing ethical crowdsourcing through the community offers alternative ways of engaging a diverse range of native speakers in the data collection process [2]. Employing datasheets for datasets and clear data licensing can support transparency and

accountability [44]. Researchers can draw valuable lessons from fields like archival studies, which have long addressed consent and privacy considerations [34]. Information about data sources and any potential biases should be shared with the community in an open and accessible manner [55]. Above all, the interests of the speech community take precedence over viewing a language merely as a data resource [14].

4.4.2 Prepare data

Data preparation may involve multiple processes. When annotation is required, we recommend adopting inclusive annotation standards and involving a diverse group of annotators who are familiar with the cultural and linguistic context to minimize potential biases [40, 44], while ensuring that their labor rights are protected [48]. Community members should be provided with clear information about the annotation process and its role in Language AI development [55]. Crucially, language data should be understood within their social and cultural context, rather than as a technical resource [14, 16, 62].

4.4.3 Develop technology

Involving community members in technology development can include providing an overview of different models, discussing their strengths and limitations collaboratively, and jointly deciding which model to implement [55]. Discussing the goals, expected outcomes, and plans for deployment can ensure the resulting system is usable and relevant [55]. It is important to account for multilingual and multicultural nuances in benchmarks and training [44], and develop and use models that work effectively on the specific data, considering the amount of available data as well as the linguistic and cultural detail and complexities.

4.4.4 Evaluate

Identifying appropriate and inclusive evaluation metrics that reflect the needs and values of those affected by the technology is crucial for a locally relevant form of evaluation [4, 44, 49]. This may include combining human evaluation with automated metrics, using culturally and linguistically diverse benchmarks, and applying fairness metrics suited to the specific context [44].

4.5 Deployment: deploy technology

Collaborating with stakeholders in technology deployment and dissemination strategy development can maximize success by securing varied forms of support from local partners and reaching the intended users, promote local agency and pave the way for future local maintenance and governance (if so agreed).

4.5.1 Develop dissemination strategy

Dissemination strategies might include ways to test and effectively share new technology with the intended users. Some elements of a dissemination strategy could include: an early phase involving a pilot roll out and sharing with advocates who might help to popularize the new technology; setting up implementation partnerships with stakeholders who could support with the physical infrastructure and hardware providers to share technology widely, or perhaps with the backing of organizations, such as Government, unions or NGOs to access key individuals or institutions according to the focus.

Close collaboration with key stakeholders and the user community will be useful for planning the roll-out and any associated training. Community members can be involved in discussions about how models will be applied in real-world contexts, including potential risks, benefits, and broader social impacts [55]. They may wish to take the lead in

disseminating the knowledge gained through the project, supported by the team as required. This can build local capacity, ultimately sustaining its benefits over the long term [55].

4.5.2 Make technology available to users, consider piloting first

Technology should be made accessible for community users to maximize its benefits and ensure that its value is equitably distributed [43]. It is also important to recognize and account for the heterogeneity of technology use and users across different cultural contexts, as patterns of adoption and impact may vary significantly [5]. A pilot or small-scale deployment can be used as a test-run with a willing group; this can be invaluable in identifying and addressing arising challenges, resulting in a smoother, larger-scale deployment.

4.6 Maintenance: sustain technology over time

The maintenance stage may involve securing necessary financial resources, engaging with policy processes, and regularly updating the technology to respond to societal and technical change. It is important to ensure that the technology can be sustained and maintained with local agency, feedback, and continued improvements. While users may not wish to have the responsibility of maintaining the technology, having the option to do so and technical support available is important for the technology's longevity. Adequate support for community members includes fostering the necessary skills to use, maintain, and take ownership of the tools. Leadership roles within this collaborative group can be such that local people are empowered to lead plans that secure long-term benefits for their community [55]. Reaching consensus on these plans and establishing mechanisms for accountability are essential for upholding commitments to the community [55].

4.6.1 Secure financial sustainability of technology

Establish a sustainable funding model that covers ongoing operational costs. Stakeholders might be in a position to fund or advocate for investment in AI technologies, especially if the project attends to social needs which fall under the remit of the Government [44]. Developing a business model - whether that be based on grant funding, a commercial model, or a hybrid - will be enabled if evidence is available to demonstrate the potential long-term social and economic benefits of the technology [44]. Any fundraising needs to carefully balance commercial interests with principles of fairness and equity to ensure contributing to broader societal good rather than reinforcing existing inequalities [44].

4.6.2 Contribute to policy debates, respond to changing policy landscape

Actively engaging with policymakers can enable the team to stay abreast of new policy developments. As policies shift, the technology may need to adapt accordingly. Engaging in policy discussions can also provide a pathway, where appropriate, to contribute to government policies and programs seeking to promote equitable AI development [3, 44]. Such engagement can involve contributing to the development of regulatory frameworks, sharing research evidence to inform policy decisions, and advocating for resource allocation to support marginalized communities. Working collaboratively with policymakers, researchers, and practitioners positions the team to contribute to AI governance while respecting national-level processes and recognizing when it is appropriate to engage.

4.6.3 Update technology according to changing needs, correct technical problems that arise

Being attentive to changing social, cultural, and technological contexts will help to keep the new technology relevant and useful. This can be achieved either through local management and updates [16], direct community engagement, or by maintaining strong relationships with key local contacts. This ongoing engagement enables timely identification of

emerging challenges, and raises the possibility of altering the technology in response. In parallel, maintaining a dedicated technical team is essential to ensure that adjustments or updates to the technology can be implemented promptly. Together, these practices can sustain the relevance, efficacy, and responsiveness of the technology over time, making all the preceding steps more worthwhile.

4.7 Discussion

We reviewed a wealth of literature on how to build participatory Language AI and synthesized best practices from a variety of research fields, including NLP, HCI, AI, and social sciences. While these fields differ in focus, a common conclusion emerged: designing equitable Language AI that genuinely benefits communities is greatly enabled by participatory methods. Our proposed roadmap, developed based on this literature, is governed by five overarching principles: stakeholder engagement; sustainability; knowledge transfer; transparency and fairness; and reflection and documentation. These principles are widely acknowledged across the surveyed literature as central to ethical and effective technology design, particularly in community-centered and cross-cultural contexts. Moving to the specifics of the roadmap, we present five stages of project development. While these are presented in a linear fashion, the roadmap can be flexibly applied to diverse projects in varied contexts. Despite this case-specific flexibility, we map out the entire process where the stages complement each other, to offer a non-dogmatic tool to guide AI development based on the best practice to date.

The first stage of the five-stage roadmap is desk-based planning to start to understand the issues, context and key actors, and gain ethics approval. This stage involves building trust and relationships between researchers and communities, understanding the sociocultural and technological context, and identifying existing tools and initiatives. As one of the main challenges of participatory AI development is reaching the marginalized communities, starting to build connections and relationships early on is crucial, and sufficient time and resources should be allocated to this process. For example, the Masakhane project [49] describes recruiting participants from meetups, universities and local machine learning school, making connections through Twitter (now X), press coverage, conference workshops, and research publications. While identifying and reaching relevant stakeholders can be time-consuming, it would lay a strong foundation for the next stages. Identifying which technology already exists and how it can support the project, and vice versa, is another important step - first, it helps to avoid duplication of work, making the process more efficient for AI developers, businesses and research labs, and the technology more relevant for the intended users; second, it contributes to the compliance with the principle of supporting local initiatives. From the literature that we reviewed, topics relevant to planning stage were most prominent in the work on decolonizing technology development, equitable language model design, and HCI, particularly in research involving Indigenous and marginalized communities.

The second stage of the roadmap is fieldwork, including interpersonal interaction with the stakeholders and community groups that were identified in the previous stage. This stage emphasizes deeper engagement with local culture, collaborative decision-making regarding technology design, and clear agreements on data ownership and knowledge transfer. Fieldwork lays the foundation for developing equitable Language AI in the next stage: engaging with local context, culture and language can help build culturally and linguistically sensitive Language AI, maximizing community ownership and technology uptake; understanding and centering community needs can help build Language AI that is grounded in real world needs, increasing its practical relevance and cultural applicability. Fieldwork practices were most often discussed in social sciences, ethics, socially responsible NLP and participatory HCI literature.

The third stage, language technology development, entails the collection and preparation of data with an iterative design process whereby evaluations and corrections are built in. According to the reviewed literature, the challenges that can be encountered include limited language resources, lack of diverse annotators and inclusive annotation standards, inclusive

evaluation metrics and culturally diverse benchmarks. Co-designing and -developing the technology with the community, including native speaker knowledge, local perspectives and expertise, building on the relationships built and knowledge gained in the planning and fieldwork stages can help tackle these challenges. These topics were primarily discussed in the NLP and machine learning literature, focusing on under-served languages where representativeness, fairness, and inclusivity concerns are particularly prominent.

Stage four involves the deployment of technology, highlighting the importance of dissemination strategies, with a view to maximizing the accessibility, usability and availability of the new Language AI to its intended users. This stage can effectively build on the previous stages, benefitting from the relationships built earlier, and from the technology design that is already informed by, and relevant and applicable to local needs and cultural context. These themes were most visible in international development literature and in ethical and community-centered approaches to technology implementation.

Often overlooked is stage five - the maintenance of Language AI - addressing the long-term sustainability of the new technologies, including financial considerations, participating in policy discourse and contributing to skills development. Investing time and resources into this stage helps ensure that tools remain responsive to societal and technological change and can be maintained and adapted independently by the community. These considerations have been explored primarily in AI ethics literature and emerging work in participatory NLP. Of course, without deployment and maintenance, the value of the earlier stages of this work is negated, so planning carefully for these later stages is worthwhile.

Our work demonstrates that while numerous disciplines have examined the challenges and opportunities associated with equitable Language AI development, their rich contributions tend to focus on specific stages of the development process. As a result, practical recommendations have remained fragmented and often lack integration across domains. The roadmap presented here addresses this gap by synthesizing insights from diverse fields and proposing an interdisciplinary approach to equitable Language AI development. This roadmap provides a practical tool for researchers and practitioners for designing, deploying, and maintaining AI technologies that prioritize community needs and maximize societal benefit. While this work is written from the perspective of outside people going into a community, it could usefully be adapted for people to use within their own communities, too.

5 CONCLUSION

This paper proposes an actionable tool to support equitable Language AI design and implementation, drawing on emerging best practices from multiple disciplines. We highlight five key principles and translate them into a five-stage roadmap, covering planning, fieldwork, technology development, deployment, and maintenance. Equitable Language AI is important for several reasons: it supports a fairer distribution of the benefits of these technologies, boosts the sense of ownership and empowers people, and increases the relevance and applicability of the resulting tools, making them more sustainable, more usable, and ultimately more valuable for society.

This increased applicability benefits not only intended users but also the wider ecosystem of stakeholders. For example, more applicable Language AI tools could help fulfill statutory and discretionary duties more cheaply or at a larger scale from a government perspective; increased trust and legitimacy, better understanding of people's needs and AI literacy may benefit policymakers, organizations and governments; access to customers and testing market, improved operating environment, and the increased satisfaction motivating people to use their applications may incentivize companies; meaningful real world impact could benefit AI developers and academic labs. Another incentive to develop participatory Language AI, be it in the public or private sector, could be compliance with regulations - while the regulations are constantly changing, increasing emphasis is on inclusion and diversity in AI development (for example, EU AI Act [25] or OECD AI Principles [51]). Reflecting on what may incentivize different stakeholder groups to use participatory methods

is important, as - while these methods can be costly and time-consuming - advancing equitable Language AI requires contributions from all stakeholders, and, as the roadmap demonstrates, all stakeholders will benefit from a more equitable landscape that can be achieved by centering the community.

As a synthesis-based contribution, the roadmap presented in this paper has limitations that point directly to important directions for future work. It is grounded in a wide and diverse body of prior work, including studies that involve community-based, participatory, and inclusive approaches to Language AI. It also draws on the authors' collective experience across language technology development, international development, and international research collaboration. At the same time, the roadmap itself was developed through a review- and synthesis-based process, rather than through direct participatory engagement with diverse affected communities. As such, the roadmap should be understood as a foundational structure rather than a finished prescription. Its value lies in making existing participatory insights more visible, systematic, and actionable for AI researchers and developers. An important next step is to apply, test, and refine the roadmap in real-world settings through engagement with diverse stakeholders, particularly those who are most affected by current inequities in Language AI.

We present this roadmap as a practical and adaptable tool to support teams in developing Language AI in ways that are more attentive to equity in both process and outcomes. More broadly, this work highlights the need to foreground equity not only in the impacts of Language AI, but in the decisions about how, where, why, and with whom such systems are designed, deployed, and maintained.

GENERATIVE AI USAGE STATEMENT

Generative AI (ChatGPT, version 5.1) was used to assist with grammar and fluency of writing in some paragraphs. Entire passages were not rewritten, but some suggestions for grammatical change and fluency were adopted.

ACKNOWLEDGMENTS

This research was supported by UK Research and Innovation (UKRI) Frontier Research Grant EP/Y031350/1 under the UK government's funding guarantee for ERC Advanced Grants for the project entitled 'Towards Globally Equitable Language Technologies (EQUATE)'. Hannah M. Claus is supported by Gates Cambridge Trust (Grant no. OPP1144 from the Bill & Melinda Gates Foundation).

POSITIONALITY

Much of the critique in this paper addresses the dominant role of the Global Minority in NLP development. At the same time, we acknowledge our own position as researchers based at a well-resourced and privileged university, which provides access to funding, publications, and research networks that shape and enable this work. Our experience in language technology, international development, and global research has shown us the value of collaboration and participatory approaches, and this paper combines a literature review with these insights. However, we also recognize our limited lived experience in colonized contexts and the privileges of mobility and choice that affect how research settings are experienced. For this reason, the roadmap is grounded in humility: encouraging researchers to listen, to collaborate equitably, to refrain from developing technology when it is not wanted, and to share outcomes fairly.

REFERENCES

- [1] Vanessa Abou-Khalil, Samer Helou, Brendan Flanagan, Norbert Pinkwart, and Hiroaki Ogata. 2019. Language learning tool for refugees: Identifying the language learning needs of Syrian refugees through participatory design. *Languages* 4, 3, Article 71. <https://doi.org/10.3390/languages4030071>
- [2] Basil Abraham, Danish Goel, Divya Siddarth, Kalika Bali, Manu Chopra, Monojit Choudhury, Pratik Joshi, Preethi Jyoti, Sunayana Sitaram, and Vivek Seshadri. 2020. Crowdsourcing speech data for low-resource languages from low-income workers. In *Proceedings of the Twelfth Language Resources and Evaluation Conference (LREC 2020)*, May 11-16, 2020, Marseille, France. European Language Resources Association, Paris, France, 2819–2826. <https://aclanthology.org/2020.lrec-1.343>
- [3] Ife Adebara, Hawau Olamide Toyin, Nahom Tesfu Ghebremichael, AbdelRahim A. Elmadany, and Muhammad Abdul-Mageed. 2025. Where are we? evaluating LLM performance on African languages. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 32704-32731. [10.18653/v1/2025.acl-long.1572](https://doi.org/10.18653/v1/2025.acl-long.1572)
- [4] Jesujoba Alabi, Michael A. Hedderich, David Ifeoluwa Adelani, and Dietrich Klakow. 2025. Charting the landscape of african nlp: Mapping progress and shaping the road ahead. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 27795-27829. [10.18653/v1/2025.emnlp-main.1414](https://doi.org/10.18653/v1/2025.emnlp-main.1414)
- [5] Adriana Alvarado García, Juan F. Maestre, Manuhua Barcham, Marilyn Iriarte, Marisol Wong-Villacres, Oscar A. Lemus, Palak Dudani, Pedro Reynolds-Cuéllar, Ruotong Wang, and Teresa Cerratto Pargman. 2024. Decolonial pathways: Our manifesto for a decolonizing agenda in HCI research and design. In *Extended abstracts of the 2021 CHI conference on human factors in computing systems*, pages 1-9, Yokohama, Japan. Association for Computing Machinery.
- [6] Sherry R. Arnstein. 1969. A ladder of citizen participation. *Journal of the American Institute of Planners*, 35(4): 216–24. <https://doi.org/10.1080/01944366908977225>
- [7] Hugo Asselin and Suzy Basile. 2018. Concrete ways to decolonize research. *ACME: An International Journal for Critical Geographies* 17, no. 3: 643–650.
- [8] Nesrine Bannour, Sahar Ghannay, Aurélie Névéol, and Anne-Laure Ligozat. 2021. Evaluating the carbon footprint of NLP methods: a survey and analysis of existing tools. In *Proceedings of the second workshop on simple and efficient natural language processing*, pages 11-21, Virtual. Association for Computational Linguistics.
- [9] Anna Barford, Annabel Beales, and Marina Zorila. 2025. An expert study of systemic influences on progress towards living wages: a key to unlock the Sustainable Development Goals. *Business Strategy & Development* 8, no. 1 (2025): e70048. <https://doi.org/10.1002/bsd2.70048>
- [10] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT 2021)*, March 3-10, 2021, Virtual Event, Canada. Association for Computing Machinery, New York, NY, USA, 610–623. <https://doi.org/10.1145/3442188.3445922>
- [11] Eleretta Bietti and Roxana Vatanparast. 2019. Data Waste. 61 *Harvard International Law Journal Online*. Available at SSRN: <https://ssrn.com/abstract=3584251>
- [12] Steven Bird. 2018. Designing mobile applications for endangered languages. In *The Oxford Handbook of Endangered Languages*. Oxford University Press.
- [13] Steven Bird. 2020. Decolonising speech and language technology. In *28th International Conference on Computational Linguistics, COLING 2020*, pages 3504-3519, Virtual, Spain. Association for Computational Linguistics.
- [14] Steven Bird and Dean Yibarbuk. 2024. Centering the speech community. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 826-839, St. Julian's, Malta. Association for Computational Linguistics.
- [15] Abeba Birhane, William Isaac, Vinodkumar Prabhakaran, Mark Diaz, Madeleine Clare Elish, Iason Gabriel, and Shakir Mohamed. 2022. Power to the people? Opportunities and challenges for participatory AI. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, pages 1-8, Arlington VA, USA. Association for Computing Machinery.
- [16] Tommaso Caselli, Roberto Cibin, Costanza Conforti, Enrique Encinas, and Maurizio Teli. 2021. Guiding principles for participatory design-inspired natural language processing. In *Proceedings of the 1st Workshop on NLP for Positive Impact*, 27–35. Association for Computational Linguistics, Online.
- [17] Robert Chambers. 1983. *Rural development: Putting the last first*. Longman.
- [18] Robert Chambers. 1994. The origins and practice of participatory rural appraisal. *World Development*, 22(7): 953–69. [https://doi.org/10.1016/0305-750X\(94\)90141-4](https://doi.org/10.1016/0305-750X(94)90141-4)
- [19] Robert Chambers. 2017. *Can We Know Better? Reflections for Development* (Rugby: Practical Action Publishing).
- [20] Alan Chan, Chinasa T. Okolo, Zachary Terner, and Angelina Wang. 2021. The limits of global inclusion in AI development. *Computers and Society*. arXiv preprint arXiv:2102.01265.
- [21] Ned Cooper, Courtney Heldreth, and Ben Hutchinson. 2024. “It’s how you do things that matters”: Attending to process to better serve Indigenous communities with language technologies. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, 204–211. Association for Computational Linguistics, St. Julian's, Malta.
- [22] Eric Corbett, Remi Denton, and Sheena Erete. Power and public participation in AI. 2023. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1-13. <https://doi.org/10.1145/3617694.3623228>
- [23] Andrea Cornwall. 2006. Historical perspectives on participation in development. *Commonwealth & Comparative Politics*, 44(1), pp.62-83.
- [24] Sara Duvisac. 2022. Decolonize! What does it mean? Oxfam. <https://oxfamilibrary.openrepository.com/bitstream/handle/10546/621456/tr-decolonizewhat-does-it-mean-151222-en.pdf?sequence=1>

- [25] The EU AI Act (Regulation (EU) 2024/1689)
- [26] Alia Faruqi, Muhammad Asif, and Zahra Khan. 2025. Speaking for Sustainability: AI-Driven Language Technologies in Achieving the SDGs. *The Critical Review of Social Sciences Studies* 3, no. 2 (2025), 1863-1869. <https://doi.org/10.59075/dahvc044>
- [27] Spence Green, Jeffrey Heer, and Christopher D. Manning. 2015. Natural language translation at the intersection of AI and HCI. *Commun. ACM* 58, 12, 90–99.
- [28] Michael A. Hart, Silvia Straka, and Gladys Rowe. 2016. Working across contexts: Practical considerations of doing Indigenist/anti-colonial research. *Qualitative Inquiry* 23, no. 5: 332-342. <https://doi.org/10.1177/1077800416659084>
- [29] Peter Henderson, Jieru Hu, Joshua Romoff, Emma Brunskill, Dan Jurafsky, and Joelle Pineau. 2020. Towards the systematic reporting of the energy and carbon footprints of machine learning. *Journal of Machine Learning Research* 21, no. 248 (2020), 1-43. <https://www.jmlr.org/papers/v21/20-312.html>
- [30] Daniel Hershovich, Nicolas Webersinke, Mathias Kraus, Julia Anna Bingler, and Markus Leippold. 2022. Towards climate awareness in NLP research. *Computation and Language*. arXiv preprint arXiv:2205.05071. <https://doi.org/10.48550/arXiv.2205.05071>
- [31] Alex Hradecky, Omar Jagne, and Maria Giulia Vitagliano. 2024. Equitable AI for Africa. United Nations Development Programme Blog. Retrieved 16 July, 2025 from <https://www.undp.org/blog/equitable-ai-africa>
- [32] Nick Hughes and Susie Lonie. 2007. M-PESA: Mobile Money for the “Unbanked” Turning Cellphones into 24-Hour Tellers in Kenya. *innovations / winter & spring 2007*, p.64-81. https://www.gsma.com/solutions-and-impact/connectivity-for-good/mobile-for-development/wp-content/uploads/2012/06/innovationsarticleonmpesa_0_d_14.pdf
- [33] ISO. 2019. 9241-210:2019(E). Ergonomics of human-system interaction — Part 210: Human-centred design for interactive systems. Second edition. Standard Preview. <https://cdn.standards.iteh.ai/samples/77520/8cac787a9e1549e1a7ffa0171dfa33e0/ISO-9241-210-2019.pdf>
- [34] Eun Seo Jo and Timnit Gebru. 2020. Lessons from archives: Strategies for collecting sociocultural data in machine learning. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 306-316, Barcelona, Spain. Association for Computing Machinery.
- [35] Anna Jobin, Marcello Lenca, and Effy Vayena. 2019. The global landscape of AI ethics guidelines. *Nature machine intelligence* 1, no. 9 (2019), 389-399. <https://doi.org/10.1038/s42256-019-0088-2>
- [36] Pratik Joshi, Sebastian Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. The state and fate of linguistic diversity and inclusion in the NLP world. arXiv:2004.09095. Retrieved from <https://arxiv.org/abs/2004.09095>
- [37] Uma Kothari & Bill Cooke. 2001. ‘The case for participation: the new tyranny?’, *Participation: the New Tyranny?* (London: Zed Books). 1–15
- [38] Shaimaa Lazem, Danilo Giglito, Makuochi Samuel Nkwo, Hafeni Mthoko, Jessica Upani, and Anicia Peters. 2022. Challenges and paradoxes in decolonising HCI: A critical discussion. *Computer Supported Cooperative Work (CSCW)* 31, no. 2. 159-196. <https://doi.org/10.1007/s10606-021-09398-0>
- [39] Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See et al. 2019. WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on human-computer interaction* 3, CSCW. 1-35. <https://doi.org/10.1145/3359283>
- [40] André Lewis. 2025. Unpacking cultural bias in AI language learning tools: An analysis of impacts and strategies for inclusion in diverse educational settings. *Int. J. Res. Innov. Soc. Sci.* 9, 1878–1892. DOI: <https://doi.org/10.47772/IJRISS.2025.9010151>
- [41] Areej Liaqat, Cosmin Munteanu, and Collin Demmans Epp. 2021. Collaborating with mature English language learners to combine peer and automated feedback: A user-centered approach to designing writing support. *Int. J. Artif. Intell. Educ.* 31, 638–679. DOI: <https://doi.org/10.1007/s40593-020-00204-4>
- [42] Jonathan K. London. 2007. Power and Pitfalls of Youth Participation in Community-Based Action Research. *Children, Youth and Environments*, 17(2): 406–32. <https://doi.org/10.1353/cye.2007.0078>
- [43] Manuel Mager, Elisabeth Mager, Katharina Kann, and Ngoc Thang Vu. 2023. Ethical considerations for machine translation of indigenous languages: Giving a voice to the speakers. *Computation and Language*. arXiv preprint arXiv:2305.19474.
- [44] Rada Mihalcea, Oana Ignat, Longju Bai, Angana Borah, Luis Chiruzzo, Zhijing Jin, Claude Kwizera, Joan Nwatu, Soujanya Poria, and Thamar Solorio. 2025. Why AI Is WEIRD and Shouldn't Be This Way: Towards AI for Everyone, with Everyone, by Everyone. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 27, pages 28657-28670, Philadelphia, Pennsylvania. Association for the Advancement of Artificial Intelligence.
- [45] Grace Mueller and Aashma Baniya. 2024. ‘Caring’ Through Translation: A Dialogue on Ethics and Inclusivity in Cross-Language Research. *International Journal of Qualitative Methods*, 23. <https://doi.org/10.1177/16094069241256280>
- [46] Grace Mueller, Anna Barford, Helen Osborne, Kaajal Pradhan, Rachel Proefke, Soniya Shrestha, and Andi M. Pratiwi. 2023. Disaster diaries: qualitative research at a distance. *International Journal of Qualitative Methods*, 22, p.16094069221147163.
- [47] Jacob R. Mugumbate and Admire Chereni. 2020. Now, the theory of Ubuntu has its space in social work. *African Journal of Social Work* 10, 1 (April 2020), 13 pages. <https://hdl.handle.net/10779/uow.27844983.v1>
- [48] James Muldoon, Callum Cant, Mark Graham, and Funda Ustek Spilda. 2025. The poverty of ethical AI: impact sourcing and AI supply chains. *AI & society* Volume 40, pages 529–543. <https://doi.org/10.1007/s00146-023-01824-9>
- [49] Wilhelmina Nekoto, Vukosi Marivate, Tshinondiwa Matsila, Timi Fasubaa, Tajudeen Kolawole, Taiwo Fagbohunbe, Solomon Oluwole Akinola et al. 2020. Participatory research for low-resourced machine translation: A case study in african languages. *Computation and Language*. arXiv preprint arXiv:2010.02353.
- [50] Fransisco Nunes, Joana Couto Silva, Beatriz Félix, Ricardo Melo, Heike Winschiers-Theophilus, Naveen Bagalkot, Nervo Verdezoto et al. 2023. African co-design: Past, present, and emerging. In *Proceedings of the 4th African Human Computer Interaction Conference*, 316-318. <https://doi.org/10.1145/3628096.3629080>

- [51] Organisation for Economic Co-operation and Development (OECD). *Recommendation of the Council on Artificial Intelligence (OECD AI Principles)*. May 2024. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
- [52] Chijoke Okorie and Vukosi Marivate. 2024. How African NLP Experts Are Navigating the Challenges of Copyright, Innovation, and Access. Carnegie Endowment for International Peace. Retrieved 16 July, 2025 from <https://carnegieendowment.org/research/2024/04/how-african-nlp-experts-are-navigating-the-challenges-of-copyright-innovation-and-access?lang=en>
- [53] Simone K. Orłowski, Sharon Lawn, Anthony Venning, Megan Winsall, Gabrielle M. Jones, Kaisha Wyld, Raechel A. Damarell, Gaston Antezana, Geoffrey Schrader, David Smith, Philippa Collin, and Niranjana Bidargaddi. 2015. Participatory research as one piece of the puzzle: a systematic review of consumer involvement in design of technology-based youth mental health and well-being interventions. *JMIR human factors*, 2(2), p.e4361.
- [54] Jane Palmer. 2024. How to harness AI's potential responsibly. *Nature*, 632, 1181. 1181-1183. [10.1038/d41586-024-02762-2](https://doi.org/10.1038/d41586-024-02762-2)
- [55] Meghana Pillai, Alondra C. Griffin, Carly A. Kronk, and Tiffany McCall. 2023. Toward community-based natural language processing (CBNLP): Cocreating with communities. *J. Med. Internet Res.* 25, Article e48498. DOI: <https://doi.org/10.2196/48498>
- [56] Rachel Proefke and Anna Barford. 2023. Creating spaces for co-research. *Journal of the British Academy*, 11, 19–42. <https://doi.org/10.5871/jba/011s3.019>
- [57] Michelle Richards, Claire Jones, Jennifer Lardy, Andrew Godden, William Godden, Sarah Bock, and Hannah Lardy. 2024. Warma Mangarrayi: Co-designing an app for learning Mangarrayi, an Indigenous language of Northern Australia. *Int. J. Hum.-Comput. Interact.* DOI: <https://doi.org/10.1080/10447318.2024.2381924>
- [58] Roy Schwartz, Jesse Dodge, Noah A. Smith, and Oren Etzioni. 2020. Green ai. *Communications of the ACM* 63, no. 12 (2020), 54-63. <https://doi.org/10.1145/3381831>
- [59] Mae Shaw. 2008. "Community development and the politics of community." *Community development journal* 43.1: 24-36.
- [60] Kathleen Siminyu, Jade Abbott, Kólá Túbòsún, Aremu Anuoluwapo, Blessing K. Sibanda, Kofi Yeboah, David Adelani, Masabata Mokgesi-Seling, Frederick R. Apina, Angela T. Mthembu, Arshath Ramkilwan, and Babatunde Oladimeji. 2023. Consultative engagement of stakeholders toward a roadmap for African language technologies. *Patterns* 4, 8 (August 2023), 100820. <https://doi.org/10.1016/j.patter.2023.100820>
- [61] Thapelo Andrew Sindane, Vukosi Marivate, and Michelle Terblanche. 2023. Point of pivot: cross-lingual embeddings calibration for Southern Nguni and Niger-Congo low-resourced languages. In *Deep Learning Indaba 2023*.
- [62] Andrew Smart, Ben Hutchinson, Lameck Mbangula Amugongo, Suzanne Dikker, Alex Zito, Amber Ebinama, Zara Wudiri et al. 2024. Socially responsible data for large multilingual language models. *Computation and Language*. arXiv preprint arXiv:2409.05247.
- [63] Anders Søgaard. 2022, December. Should We Ban English NLP for a Year? In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022)*, December 7-11, 2022, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics, Kerrville, TX, 5254–5260. <https://doi.org/10.18653/v1/2022.emnlp-main.351>
- [64] Atnafu Lambebo Tonja, Tadesse Destaw Belay, Israel Abebe Azime, Abinew Ali Ayele, Moges Ahmed Mehamed, Olga Kolesnikova, and Seid Muhie Yimam. 2024. Natural language processing in ethiopian languages: Current state, challenges, and opportunities. *Computation and Language*. arXiv preprint arXiv:2303.14406.
- [65] Vodafone. No date. Technology and Innovation: M-PESA. <https://www.vodafone.com/about-vodafone/what-we-do/m-pesa>
- [66] Nan Yan, and Cheng Xu. 2024. Decolonizing african NLP: A survey on power dynamics and data colonialism in tech development." In *5th Workshop on African Natural Language Processing*.
- [67] Haiyi Zhu, Bowen Yu, Aaron Halfaker, and Loren Terveen. 2018. Value-sensitive algorithm design: Method, case study, and lessons." *Proceedings of the ACM on human-computer interaction* 2, no. CSCW. 1-23. <https://doi.org/10.1145/3274463>

A APPENDICES

A.1 A Community-centered, Participatory and Actionable Roadmap for Equitable Language AI

A COMMUNITY-CENTERED, PARTICIPATORY AND ACTIONABLE ROADMAP FOR EQUITABLE LANGUAGE AI			
PRINCIPLES:			
1. Engage the community and other stakeholders throughout, focussing on their needs and the local context 2. Engage in knowledge exchange , support existing projects and create meaningful learning opportunities 3. Prioritise sustainability, maximise benefits and manage risks to the community, research team and environment 4. Promote transparency, privacy, and fairness 5. Continuously reflect and document the process			
	GOAL	ACTION POINTS	HOW?
PLANNING	understand context, needs and AI readiness	identify and build connections with stakeholders	reach out to local experts and gatekeepers, engage diverse stakeholders
		understand local context: needs, challenges and priorities, AI readiness, power dynamics and the role of technology in society	speak with domain experts and stakeholders, including local people and national AI researchers; desk research on statistics, academic papers, and grey literature
		assess existing technology and how it could be useful	apply for ethics approval from national ethics committees, home institution, and local authorities
FIELDWORK	learn with and from the community	gain ethics and other approvals to undertake research	apply for ethics approval from national ethics committees, home institution, and local authorities
		learn about the culture, economy, environment and politics, and identify priority needs	participatory research using surveys, focus groups, interviews with the community and other stakeholders and experts; engage social scientists for methodological rigour
		decide whether to develop a technology and agree on roles and responsibilities	present findings to the community, then transparently discuss and agree with stakeholders, including intended users and decision-makers
DEVELOPMENT	develop AI technology	design an approach and timeline for technology development, and plan how to sustain technology	clear, possibly written, agreement with the community, with a plan for next steps
		clarify AI data ownership, use and storage, and transfer of know-how	
		collect data	use participatory methods and engage community and stakeholders in every stage; consider specifics of working with language data; promote shared decision-making, share data, and create training opportunities
DEPLOYMENT	deploy technology	prepare data	
		develop technology	
		evaluate	
DEPLOYMENT	deploy technology	develop dissemination strategy	work with stakeholders and user community to plan roll-out and training as needed
		make technology available to users, consider piloting first	ensure adequate support, including the skills to use, maintain and own the tools are acquired
MAINTENANCE	sustain technology over time	secure financial sustainability of technology	establish a funding model to cover costs
		contribute to policy debates, respond to changing policy landscape	ongoing engagement with policy makers
MAINTENANCE	sustain technology over time	update technology according to changing needs, correct technical problems that arise	stay abreast with changing technological and sociopolitical contexts either directly or through key contacts, maintain a technical team to make prompt alterations as required

Figure 1: A Community-centered, Participatory and Actionable Roadmap for Equitable Language AI