

Innovating AI for Diverse Ways of Knowing

A **MANIFESTO** FOR GENERATIVE AI IN A GLOBAL SOUTH CONTEXT

Version: March 2025

A PROJECT OF LIVING ARTS INTERNATIONAL

INTRODUCTION AND BACKGROUND

Introduction and Background			
Methodology			
The Document in Context			
Part One: The Manifesto in Brief	7		
Part Two: The Manifesto in Practice	9		
Phase 1: Onboarding and Design	9		
1.1 Onboarding Stakeholders	10		
1.2 Use cases and the end user	16		
1.3 Design or Choice of the Model	18		
Concluding Phase 1	21		
Phase 2: Developing the AI	22		
2.1 Data Collection	23		
2.2 Training / Fine Tuning the model	27		
Concluding Phase 2	30		
Phase 3: Evaluating and Growth	31		
3.1 Evaluation	32		
3.2 Data Governance and Growth	35		
Concluding Phase 3	38		
Conclusion	39		
Annex A: List of Resources Shared with the Working Group	42		
Annex B: Profiles of the Working Group	44		
Annex C: References	45		
	Introduction and Background Methodology The Document in Context Part One: The Manifesto in Brief Part Two: The Manifesto in Practice Phase 1: Onboarding and Design 1.1 Onboarding Stakeholders 1.2 Use cases and the end user 1.3 Design or Choice of the Model Concluding Phase 1 Phase 2: Developing the Al 2.1 Data Collection 2.2 Training / Fine Tuning the model Concluding Phase 2 Phase 3: Evaluating and Growth 3.1 Evaluation 3.2 Data Governance and Growth Concluding Phase 3 Conclusion Annex A: List of Resources Shared with the Working Group Annex B: Profiles of the Working Group		

INTRODUCTION AND BACKGROUND

In the past two years, the landscape of generative AI has shifted dramatically. From the initial explosion of excitement particularly in the western world around AI tools for daily use such as ChatGPT and image generation such as Midjourney and Dall-E, the conversation quickly pivoted to questions of regulation, data protection, and copyright infringement. But as use of these systems expanded leaps and bounds, it became clear that they represented a biased view of the world, one that favored white western narratives and western ways of knowing through written texts and high resource languages like English. Many companies, organizations, and activists turned their attention to the ethics and inclusiveness of generative AI, pointing out the damaging outcomes of relying on biased data.

Today the landscape is still evolving, still very much at the rapid pace as before. While Silicon Valley giants such as OpenAI and Microsoft continue to push ahead with huge LLMs, newcomer Deepseek challenges the narrative that building a true LLM costs millions of dollars. However, while the rivalry between Deepseek and Big Tech captures the attention of the world due to its connection with high profile geopolitics, small organizations in Africa are building highly functional small language models with low resource languages and developing "African-centric" approaches and methodologies (Awarri, n.d.; Masakhane, n.d.; Tonja et al., 2024). Indigenous groups in Australia and North America have articulated their own frameworks of data governance and data rights, seeing data as the next frontier of exploitable resources that are being taken from them (Hao, 2022; Li, 2024).

While generative AI technology races ahead as fast as possible, the cultural sphere is not as quick to follow the trends. The cultural discourse, which is generally more grounded in ethics and collaboration, tends to be fairly suspicious of AI, especially in the global south context where mainstream LLMs at best lack meaningful data and cultural sensitivity and at worst present stereotyped and damaging western-centric perspectives. Between this skepticism and Big Tech's focus on progress and innovation at all costs,

the two separate spheres rarely collaborate on solutions, despite their close connections (Kulesz, 2024).

One area which is particularly threatened by the increasing digital divide and bias in generative AI is cultural heritage, particularly in cultures in the global south which have traditionally been reliant on oral transmission. Despite the potential of technology to democratize access to cultural resources, digital heritage projects around global south communities which are created by and for the communities themselves are lacking. If digital scans exist, they are often the property of global north companies or museums, and are not accessible to the communities which they are about. Further, digital archives and training datasets for AI rely heavily on the existence of written records, a particularly western way of knowing, and do not honor or consider other ways of knowing. As such, if generative AI continues the way it started, the digital future risks perpetuating traditional colonial power dynamics, presenting a monocultural and exclusive vision of culture. Human-centered approaches and governance grounded in decoloniality and ethics are required to ensure the richness of culture can exist in the digital age (Artificial Intelligence and the Challenge for Global Governance | Chatham House - International Affairs Think Tank, 2024).

"An AI of Our Own: Innovating AI to Include Diverse Ways of Knowing" (AAOO) is an initiative of Living Arts International launched in spring 2024 in an attempt to leverage the potential and exponential growth in AI technology to create adaptive and responsive systems wherein living cultural heritage is preserved in a way that goes beyond traditional wikis, written datasets, and western methodologies of knowing and documenting. Rather than trying to fit the diverse heritage of the global south into a system which was not built for it, we propose to create "an AI of our own," which leverages existing technological expertise and advancements of the global north but invites them to go beyond and envision a multi-cultural, diverse, and inclusive digital future.

AAOO was created to explore innovative solutions to the stated issues. The initial concept note was developed in a collaborative workshop, drawing on research and conceptual work LAI has been conducting since 2021. Four main pillars of the project were identified: Connect, Nurture, Fund, and Co-create, which have guided activities thus far.

In summer 2024, LAI commissioned a research project to understand the current landscape of digital heritage and AI in Asia and Africa, culminating in a report.¹ The research highlighted the strengths of AAOO as an interdisciplinary project with an approach that balanced the separate tech and culture spheres as well as a practical, on the ground approach. However, it also illuminated the need for clear guidance on how AAOO would action its ways of working and further development of the methodology and approaches that would distinguish the project.

To respond to this need, in November 2024 LAI assembled a diverse working group of 6 individuals from different countries, backgrounds, and expertise, including data rights, computer science, technological storytelling, culture management, and philosophy.² The objective of the working group was to develop a manifesto that would define the values and methodology of AAOO on a practical level, a document that could be used to both assess projects for their adherence to AAOO's values as well as guide the actual development of a generative AI model.

This manifesto articulates a vision of a multicultural model that draws from the ways of knowing from different cultures with an intended goal for equitable AI. It focuses heavily on collaboration and stakeholder onboarding, and centers a value-sensitive design approach.

^{1.} To read the report, click here.

^{2.} To see the profiles of the working group, please see Appendix B.

METHODOLOGY

The six member working group was curated to balance both knowledge and practical experience in both technology and arts and culture. The working group convened in three online sessions: the first focused on AAOO as a whole, exploring ideas of how the manifesto fits into the larger vision and how the AI can match the project values. The second session focused on the design of the AI, touching on the choice of models, how communities are onboarded, and a 'bottom-up' approach that focuses on the end use before beginning. The final session focused on data, including sourcing, sharing, and protection.

Each session was organized around the central theme with four branching questions, and was conducted in an open discussion format. Resources and reading were shared before the session, providing context, background, or best practices as food for thought.³ Notes were gathered and shared after the session.

Following the initial sessions, the working group created a series of collaborative brainstorms, each dealing with a particular stage of designing and developing a generative AI model. These brainstorms were organized, followed by discussion and feedback. The initial draft was then developed, followed by another round of feedback and edits by the group.

^{3.} For the complete list of resources shared with the group and the questions for each session, please see Appendix A $\,$

THE DOCUMENT IN CONTEXT

Overall, this document proposes a methodology and approach for building a generative AI model that reflects a global south-centric world view, in order to guide the creation of models that are inclusive, representative, and both useful and accessible to the communities that contribute to them.

It is meant to be used as a guide for a range of potential projects, and we have endeavored to make it as diversely applicable as possible. As such, there are points we do not specifically clarify, such as the choice of base model to use. Rather, we present the values and considerations that should be taken into account in making the decision, such that it matches with the overall vision and ethic of the project.

Importantly, the Manifesto should be understood as a living document – one that may adapt and evolve as we learn and explore deeper. There are many stages and steps which can only be developed in a practical setting.

Internally, LAI aims to use this document as a practical resource as it launches a pilot case of building a generative AI centered around Cambodian culture, as a kind of hypothesis to be tested on the ground. However, we hope that it will be used by others for test cases in their own communities, to be able to examine the hypotheses and methodologies presented and suggest clearer options, models, best practices, and challenges. These will then be fed back into the Manifesto, which will inform new cases, and so on.

The current draft (March 2025) is organized in two main parts. Part One outlines the foundation of the Manifesto in its values, grounding the process of building the AI in AAOO's fundamental principles: Community-centered, Responsible, and Responsive.

Part Two then goes into much more depth, with the overall approach divided into three phases: onboarding and design, developing the AI, and evaluation and sharing. Each phase is further broken down into sections. The break down is illustrated in the table on the next page.

PHASES		SUB-CATEGORY		
Onboarding & Design	÷	Onboarding Stakeholders End Use/Use Case Design		VALUES
Developing the Al model	:	Sourcing Data Training the Model Fine Tuning the Model	•	ACTIONS
Evaluation & Sharing	:	Evaluation Data Governance and Sharing		OBSTACLES

VALUES

ACTIONS

Specific principles to be followed for that stage and guidance on how it would manifest in the context of a generative Al model. Concrete steps to carry out. In some cases, the actions are conjectures as they may change for feasibility or to match on-the-ground realities.

OBSTACLES

Possible challenges in actualizing the AI in line with the desired values, and complications to keep in mind.

Each section is explored through its values, actions, and obstacles.

The **values** are specific principles to be followed for that stage and guidance on how it would manifest in the context of a generative AI model. These provide the overarching guidance for the stage and should be referred back to in case of uncertainty in actions.

Actions are steps that will be required to be carried out. In this version, the actions are conjectures as they may change for feasibility or to match on-the-ground realities.

Finally, **obstacles** are possible challenges in actualizing the AI in line with the desired values, and complications to keep in mind. Where possible, ideas are shared as to how to address these anticipated challenges; however, they will be developed further throughout the pilot case.

PART 1: THE MANIFESTO IN BRIEF

This part articulates the values that underlie the AAOO approach

Community stakeholders are **onboarded** before any development work begins, in a process that is transparent, inclusive, and community driven/respectful. The role of the intermediaries between the development team and the community is clearly outlined and compensation is defined. Onboarding includes capacity sharing, training, and mentoring for both parties. It addresses and acknowledges the embedded power dynamics between developer and community, and establishes a diverse focus group of technical experts, developers, intermediaries, community members, and translators, that will make decisions throughout the process.

The **end user and practical use cases** are explored along with the community before development begins. Use cases are sensitive to the community situation around digital literacy and access to digital resources, and balance broader goals of increased representation and decreased bias with practical uses for the community. The end uses envision possibilities for community engagement and togetherness while avoiding excessive techno-solutionism. They remain flexible and alterable, reflecting the evolving nature of culture and technology. Transparency around goals, feasible outcomes.

The **design or choice** of the AI model ensures an outcome that is culturally sensitive, inclusive, multimodal, multicultural, and community driven. It is customizable for a wide range of possible use cases. It reflects on the mutual relationships of AI and the user, and their mutual influence on each other. It considers environmental damage and resources.

Data collection and sourcing is conducted with informed consent. Procedures for collecting data are outlined with the community, sensitive to local customs, and approached as a collaborative process with trust, accountability, and transparency as central principles. It is inclusive, both in the type of data collected and a focus on non-western ways of knowing and sharing information. In the case of creating a new model, the process of **training the model** is a collaborative and iterative process that respects different ways of knowing. Community stakeholders are central to defining bias in their context and desirable/undesirable outcomes, which is the case in **fine tuning** existing models as well. In both scenarios, the process and outcomes reflect the ever-evolving nature of culture and center community input and principles. It is connected to the end use while aiming to expand representation of the global south in big data/AI, and is grounded in inclusive principles.

Evaluation is an ongoing process throughout that is community driven and reflects a vision of the model that is ever evolving and learning. It is grounded in feedback loops and metrics defined and designed by the focus group.

Data governance and growth prioritizes community ownership and data protection while encouraging learning and expansion of the methodology and knowledge gained. The strategy is flexible and adaptable to represent multiple use cases and the different needs of different communities. As the model grows beyond a single community, broader systems of data governance are put in place to ensure communities continue to have a voice beyond their specific model.

PART 2: THE MANIFESTO IN PRACTICE

The part explores how the values established above can be understood and actioned in context.

PHASE I ONBOARDING AND DESIGN

This phase focuses on setting up the overall process, which centers the community stakeholders and establishes the collaborative frameworks. There is some overlap in the steps here, however, while onboarding and establishing/understanding the use case may happen concurrently, the design or choice of the model will happen after most of these stages are complete as it relies on the understanding between the developers and community around on-ground realities.



1.1 ONBOARDING STAKEHOLDERS

This is the most crucial step in the process and will lay the groundwork for everything to come. This step cannot be rushed or incomplete; until and unless a strong relationship has been established between the community, the project and development teams, the intermediaries/translators, and any other stakeholders, the AI cannot be built.



Transparency: Center informed decision making and ensure cultural stakeholders are substantially involved and informed with what is happening in development. Be open about the experimental and evolving nature of technology and the project.

Inclusivity: Interact with a diverse and inclusive range of stakeholders from the community. Consider that individual voices are as important as the community voice and all perspectives need to be heard, considered and discussed (as opposed to a simple majority). Center voices that are in a sense already representatives of the community's cultural stand/ambitions while also making sure to include those who may not already be representatives but who have a tech bent, who are vested in the evolving world and who understand that change is coming/inevitable and that it is necessary to accept this and work with it.

Community Driven: Follow community procedures/customs in onboarding. Be aware of embedded power dynamics in a situation where the development team approaches the community as 'experts,' or as someone from outside with a solution to an internal issue. Follow a learning and knowledge sharing mindset.



IDENTIFY INTERMEDIARIES FOR COMMUNITY

The intermediary role is key for the success of the project. These intermediaries, ideally members of the community themselves, can help explain and translate, acting as a mediator and guide to help the development team and community understand and collaborate effectively. They can travel to villages, talk to elderly members, and update us regularly.



The process for choosing the intermediaries draws on existing relations with the community (we do not recommend working with communities in which the developers, project team, or outside stakeholders have no connection). Trust is vital in this process, and working with established connections enhances this. However, an open call or application can help identify other intermediaries who represent an independent perspective, especially if they can speak for or with marginalized groups inside the community.

There could be multiple intermediaries that have different functions across the project lifecycle (see below), however, there should be at least one or two that are consistent throughout. Consider inclusivity in identifying the intermediaries to ensure that minority groups within the community are also represented, or that the chosen intermediaries have access to them.

DEFINE ROLE FOR INTERMEDIARIES AND STAKEHOLDERS AT DIFFERENT STAGES OF THE PROJECT, INCLUDING COMPENSATION

The intermediary role is deeply important, and it should be compensated appropriately. There are two possible models for compensation (and both may be used at different stages of the process): a payment model and a token-based model. The payment model is ideal for the prerelease stage, especially for those intermediaries that are with the project throughout its lifecycle.

Other intermediary roles are connected with data annotation, training, and evaluation stages. For example, outside data annotation professionals or developers cannot correctly interpret the community's data and annotate it. As such, the community has to be constantly consulted and involved to ensure that the data is annotated correctly. Hence, community involvement is necessary every step of the way.

For these later stage intermediaries, a token-based model may be considered. In this context, users of the model have a limited number of messages they can send daily, however, they can get



additional tokens by helping to annotate data. This organic system is sustainable, and ensures that annotators are from the pool of end-users. It would be important in this case to ensure that community members can learn/grapple/be trained in data annotation. It could be interesting to engage unemployed youth of the community in data annotation and/or collection so they earn a basic income while using the AI.

While this step is later on in the model's development, the clarification of roles, expectations, and compensation should be addressed at this opening stage, even if adjusted later on.

For the sake of the project's feasibility at the start, consider starting with a few intermediaries and scaling up as the model grows and token-based compensation becomes relevant.

It may be relevant at this stage to consider whether or not governing bodies of the country and/or region need to be involved from the get go, and if not, if there is likely to be any interference from them throughout the process.

Fig 2. Potential model of scaling up involvement from intermediaries and the associated payment model.





ONBOARDING INCLUDES THE TECHNICAL AND DEVELOPMENT TEAM

In the interests of dismantling standard power dynamics where the technical team holds the greater authority, the technical or development team must be in full alignment with the principles and methodologies of the Manifesto, and will be required to have their own onboarding. This will include understanding community principles and establishing collaborative practices. The technical team onboarding must be deeply sensitive to how the community identifies itself and their core values (even outside of technology). The focus on onboarding for the technical team will help ensure the AI mirrors the community principles of their everyday life, in its logic and reasoning as well as its output.

KNOWLEDGE SHARING BETWEEN TECHNICAL TEAM AND COMMUNITY

Organize capacity sharing and training workshops for the developers to understand community needs and for community members to understand AI and its possible uses. Taking a human-centered approach to AI, both community and developers need to be mindful of both their own and each other's limitations and practicalities. The training should be bi-directional and ongoing throughout the project. The development team would continue to learn community needs and procedures, and the community and intermediaries would continue to learn about the technology and its potential uses.

CONVENE THE FOCUS GROUP FOR FUTURE DECISION MAKING AND DECIDE ON FRAMEWORKS FOR CONSENSUS AND MANAGING CONFLICTS

The focus group would include technical experts, community leaders, the design team, and a cultural intermediar(y)(ies) who have a good grasp of technology and AI, aware of community engagement/issues, and sensitized to dealing with cultural assets and stakeholders. This group will form the core decision makers going forward. The decision-makers themselves will need to pre-define their own conduct of working together to help them achieve cohesive outcomes as they all have individual aspirations



for the project but to blend them together can cause differences of opinions and frustrations.

AGREE ON WHICH DECISIONS NEED TO BE MADE AT THIS STAGE

The focus group will eventually need to apply the decision making framework to decide on approaches and principles for data collection, basic principles for data protection, royalty distributions for datasets, evaluation metrics, and data governance. Depending on the level of understanding and staggered approach to onboarding stakeholders, not all these decisions may necessarily be made at the beginning, however, the focus group may choose to set out some guidelines for when they will be addressed.

ESTABLISH A CODE OF CONDUCT AND OVERARCHING PRINCIPLES FOR THE PROCESS

Consider what to do should any stakeholder, within or without the community, misuse the data along the way, and how to vouch for honest participation. A code of conduct or engagement agreement, mutually agreed to by the team and the community, could be a requisite for any involvement. This code of conduct can include principles and guidance for using the AI and the data within the community itself.

CONNECT ONBOARDING TO THE END USE

During the onboarding process, clearly communicate the stakes involved and the overall vision of the project - what we are trying to do and why. Embed the process of uncovering potential use cases in the onboarding and knowledge sharing and collaborative brainstorming.



MISMATCHES IN EXPECTATIONS AND FRAMING

It is important to set realistic expectations from the start with the community and ourselves. It needs to be clear that this is an experimental project and our goals are to ensure their culture is not forgotten and they are not left behind, while also co-creating something that they can use practically. However, these objectives



and goals may not be fully achieved as expected. Embed mechanisms across that cater to this transparency, such as regular communication channels where progress/challenges are shared with the community (intermediaries etc) so they know what to expect along the journey and are aware of the challenges (and can potentially suggest their own solutions). Approach the project not as providing a service, but as a co-created journey to explore new possibilities.

OVER-RELIANCE ON TECHNOSOLUTIONISM

Understand that technology may not be the solution to everything, and the community may not feel that technology is a good solution for them. Be flexible in onboarding to allow for the co-creation of the use cases as well as different understandings and openness to AI. Allow the community representatives to be an equal partner in developing the overall design of the project, guiding principles, and value non-technical knowledge and experience to inform and enhance the technical experience.

TENSION BETWEEN THE TECHNICAL TEAM AND THE COMMUNITY/INTERMEDIARIES

Emphasize the intangible outcomes of the process throughout the onboarding equally to the technical outcomes. Adopt an ecosystem approach where everyone, including the AI itself, are vital parts of the whole, and each aspect must be nurtured and cared for. In addition to the community, the project itself and all its participants will within themselves become a 'community' that sets its own principles.

1.2 USE CASES AND THE END USER

In a community driven approach, the end user is considered as an average member of the community and is prioritized in envisioning potential use cases. Wider use can be explored but the community user must not be left behind. Balancing innovation, understanding, and collaboration is key for this stage.



Flexibility and Adaptability: The AI should be flexible enough for a wide variety of use cases, with the community's particular needs or interests at the forefront. However, it should be responsive to wider goals of expanding visibility of the global south in generative AI and reducing bias as well.

Accessibility: The end use, use case, or final dissemination format of the AI should be planned in connection with the community's existing digital literacy, connectivity, and digital resources in mind. For example, if the community has access to smartphones and reliable internet access, the final version may be different than if the majority of users would work without (such as a minified version that can run locally). It should also be accessible in terms of language and literacy (ie, able to run on audio input if literacy rates are lower).

Transparency: There must be clear lines of communication and trust established to understand the nature and goals of the project and end use. The project/development team should be careful not to make overarching promises or rely excessively on technosolutionism, but approach the end use as an opportunity. From the project team, be open on the goal of visibility worldwide and the interest of ensuring the community's cultural systems and narratives are not forgotten, as well as the interest in providing interesting new ways for the community to share and create knowledge.

Community-driven: The end use should be something valuable to the community, and they must be able to access and easily use the final model. The end-use could have an element of gaming, radio, or anything that excites, piques and engages the whole

community; making them keen to use the portal. This will ensure more use, easier access, and hence better Al. The design can include ways to help the community members identify and recognize themselves and their ways of being in these cultural systems. It should consider how the end use case can help the community bond better in person, so a collective engagement rather than singular or individual isolated use.



ASSESS POTENTIAL USE CASES IN CONSIDERATION OF THE COMMUNITY'S SITUATION

Together with the intermediaries, assess the current situation of digital resources and literacy to understand what is possible or not in terms of end use. Assess what types of information communities are already retrieving/using. Embed visualization and brainstorming into training and capacity sharing sessions to help both developers and community envision how the AI could be used to facilitate existing or desired actions in daily life. This process should be sensitive to balance individual interest with community interest as well as the long and short term visions of the community. Be aware of putting the burden too much on the community to identify problems/solutions; collectively ideate and be open to new possibilities occurring throughout the process. For the sake of design and advancing the project, at least 1-2 use cases should be identified



Fig 3. Considerations for creating potential end uses cases.



DEFINE OVERARCHING GUIDELINES FOR THE END USE

It is important that the focus group defines what are not acceptable use cases and/or what they do not want the AI to respond to or generate. This also includes safeguards against actors within or without the community in case of misuse of the technology. In particular, content moderation rules should be discussed by the focus group of experts and community members. This discussion should be done with a wide range of stakeholders to ensure inclusivity in opinion, and in the case of conflicting opinions, the focus group should refer to established internal guidelines for making decisions and managing conflict.



BALANCING PRACTICAL USE CASES WITH WIDER GOALS

For communities with little digital literacy, finding practical use cases could be challenging, and it may be more practical to focus primarily on the wider narrative of ensuring their culture is not forgotten. On the other hand, more widespread representation may only happen with multiple communities and more data, so in the first cases it may be better to start with a practical solution and work towards wider goals as more communities join. Keeping the balance and continually checking in on this will be necessary for each new case that joins the project.

1.2 DESIGN OR CHOICE OF THE MODEL

Considering the availability of resources in the global south, it is expected that a foundational model that is already trained will be chosen, which will be then fine tuned. However, whether training from scratch or working with an existing model, the same considerations apply, bringing together the end user, the model's ability to enact the values, and parent company's values especially in data protection and environmental sustainability.



Meaningful Engagement: Although AI models are generally meant to be used by individuals, the overall design considers how the model or overall technical infrastructure included in the project can help drive meaningful engagement at the community level



rather than focusing only on individual users (i.e. intra-community forums, knowledge sharing, ability to share prompts and responses with other users, etc.)

Ecosystem Approach: The overall approach to creating or choosing a model is grounded in an ecosystem approach where all actors have an effect on each other. End users should be encouraged to understand how their words and tone are reflected in the AI and how the AI and human actors mutually influence each other. Environmental sustainability must be a consideration if choosing an existing model and the parent company's policy on resources.

Inclusive: Ideally the AI will be *multimodal* and *multicultural*. It accepts and processes all forms of input and output, vector, image, finger drawing recognition, voice input and recognition and voice to text and vice versa, audio, text, scanned text and handwriting, video, movement mapping, image to text prompt and vice versa. The maximum amount of multi modalism is involved so that as many people as possible can use it. If the AI grows to multiple communities, it will ideally be able to consider and take into account the differences in culture and context in its responses.

Culturally Sensitive: The responses and narratives portrayed and shared by the AI reflect how the community identifies itself in the present, including their own definition and perspectives of their histories. It provides multiple historical perspectives and reflects a variety of narratives and folk tales, such as the community's own origin story and understandings of pre and post-colonial histories.

Community Driven: The AI serves and reflects community needs and realities, and the design is customizable enough that end users may easily adapt it to suit their needs after initial use/testing. The end user has a level of control over the system outputs. The end user drives initial decision making on the model's capabilities.



DECIDE ON THE MODEL

Considering the large amount of resources required to train a model from scratch, it is likely more practical to choose an existing model, however, many factors must be taken into account, including the average end user's digital capability, local internet access, and device capacity. The focus group must consider the possibilities of using API based models or locally-run models. The use of APIs is generally easier to build given its centrality and puts less burden on the end user given that everything is going to be working in a backend server or cloud. However, the API approach has the disadvantage that people need to be connected to the internet to be able to use the API. On the other hand, a local model can work without internet access but it might put a heavy demand on the end-user resources.

Fig 4. Considerations in choosing a foundational model



The environmental and data protection policy of the parent company of any existing model must be considered as well. Companies that are known to exploit environmental resources especially in the global south should not be chosen, as well as companies where the use and dissemination of the data cannot be controlled by the end user. Should the project decide to work with any BigTech companies, the focus group could work to leverage their case through a mutually beneficial licensing agreement wherein the project helps correct bias in the larger models in exchange for data protection.

Small language models and models developed in and by global

south actors should be prioritized.

In the case that original models are chosen, the same considerations and design factors apply.



BALANCING CONVENIENCE WITH VALUES

In deciding on the model, it may be that larger, established models have a stronger existing base or can be trained with less resources, but the parent company's values do not align with the manifesto. It is important to balance ease with responsible decisions. It could be best to follow an MVP (minimum viable product) approach and aim to release a model with as few features as possible just to test it in the field than attempt to polish a product only to realize later on that the community needs are different. This allows for interactively incorporating the end-user's feedback into thedesign loop.

CONCLUDING PHASE 1

At the end of phase 1, a strong bond of trust and collaboration has been established between the project and development teams, community representatives, and the intermediaries. There is a clear path charted for what the model may be used for, and how decisions will be made. Overarching guidelines are established and all stakeholders are prepared to begin the build.

PHASE II DEVELOPING THE AI

This phase focuses on setting up the overall process, which centers the community stakeholders and establishes the collaborative frameworks. There is some overlap in the steps here, however, while onboarding and establishing/understanding the use case may happen concurrently, the design or choice of the model will happen after most of these stages are complete as it relies on the understanding between the developers and community around on-ground realities.



2.1 DATA COLLECTION

Along with onboarding, data collection is one of the most sensitive and critical stages in the process. It has loaded histories of extractive processes and embedded power dynamics wherein the community providing data has traditionally little visibility on what their data is used for and how, as well as little control on what happens to it. The balance of collecting enough good data to train the AI while maintaining respectful, transparent, and collaborative approaches is key here.



Consensual sharing of data: All data that is input and collected is done so with informed consent. The focus group will discuss the terms and conditions for consent before the process begins, as well as how much information/training is required for informed consent to be given. This will include guidance on existing global best practices and definitions to help provide context.

Community driven and culturally sensitive: The procedures for gathering data are agreed upon with the community and sensitive to local customs and ways of gathering information. The language barrier must be addressed, and data input goes beyond written text in order to be inclusive – ideally it should include audio, visual, and vector based inputs.

Collaborative and transparent: The development team is clear and open with the community and individuals how their data is going to be used. Methods for collecting data will be co-developed by the project team and the community.

Trust and accountability: Acknowledge the power dynamics in the process of gathering data and develop trust with the community through transparency and accountability, clearly communicate the benefits (and potential drawbacks) of sharing data. Allow the community the choice if the shared data should be anonymous or attributable to them - and if attributable to them, then what terms, what sharing classifications or license, would there be any payment



for usage of data, etc. Consider how asking for consent can affect how people share data, and if it has any effect on authenticity of data.

- <u>Tension between the young and old:</u> It could be that the vast majority of internet users are young. Thus, those who are relatively old may not produce sufficient training data for their voices to be heard (unless their opinions are assigned greater weights). Additionally, it could be that most end-users are young, which can cause a disconnect if stakeholders consist only of elderly cultural leaders, who may hold somewhat dated opinions that no longer reflect the youth's opinions.
- <u>Capturing the opinions of minorities:</u> Naturally, cities are bigger than villages, and so those who live in urban areas will produce more training data than those in rural areas. Thus, unless we increase the weight of data from rural areas, their perspective will be neglected in the grand scheme of the training data.
- <u>Misinformation/Disinformation</u>: suppose there are some community members with extreme opinions that do not represent their cultural background (e.g., someone who holds very negative opinions of elderly, etc.). Arguably, this person's opinion can be thought of as "misinformation". How can we detect, and mitigate, such data? We could have community members evaluate a random sample to detect such misinformation, then train a classifier and test it based on this data. Such a classifier can scale up compared to manual detection.
- <u>Translation</u>: if we collect data in the local language and/or in English, some text or data might need to be translated. We may consider evaluating the translation, and assigning lower weights to translated text. When we consider audio input/output, any analysis and/or transmission would be more efficient when the audio is converted into text format, however, it should be sensitive to cultural needs (e.g., a song from folklore, where the tune, and not just the lyrics, need to be preserved).



ESTABLISH PROCEDURES FOR DATA COLLECTION

Based on the community's customs and methodologies, the focus group will establish the primary procedures for data collection, including conduct and guidelines (i.e. when pictures are allowed, procedures for interviews, etc.)

ESTABLISH GUIDELINES FOR CONSENT

With a community facing model, consent includes a shared understanding of purpose along with other prerequisites for legitimate consent such as 1) addressing power imbalances in data collection, 2) data subjects knowing and agreeing to what their data will be used for and 3) special considerations for the most vulnerable of the data subjects. The focus group may consider the possibility of a communal consent model vs an individualistic one.

Fig 6. Consent in a community driven model



PURPOSE

Shared sense of purpose and agreement of overall goals



POWER DYNAMICS

Power imbalances are addressed and acknowledged



AGREEMENT TO USE

Communities are aware of and agree to how their data will be used



SPECIAL CONSIDERATIONS

Vulnerable groups may require more sensitive data handling or considerations.

IDENTIFY THE CHARACTERISTICS OF THE DATA REQUIRED

The data must be inclusive and culturally sensitive. They should include location, outfit, ethnicity, language, and dialect. They should reflect historical narratives, contemporary cultural narratives, and pre and post-colonial narratives, and value storytelling and lived experience. They should include environmental distribution (Climate zones, Ecosystems, Seasons, Natural landmarks, Built environments) and cultural distribution



(Traditional practices, Contemporary lifestyle, Technology adoption, Art forms, Social structures).

Technical requirements for any types of input must be clearly defined and communicated to any intermediary collecting data.

Consider how audio and non-written data can be incorporated into these requirements. Keep in mind potential future datasets and characteristics of data that will be required to ensure interoperability.

IDENTIFY EXISTING SOURCES OF DATA AND INTEROPERABILITY

Working with the community at the outset, the development team can identify existing sources of data, such as archives or social media, and assess how it can be gathered with consent. With these existing sources, assess interoperability and integration with the overall characteristics. Finding alternative avenues to gathering data will be necessary to have diverse examples and possibilities for the AI to draw on. Investigate ways to support startups and small businesses incubators through data collaboration.

VALIDATE THE ACCURACY OF COLLECTED DATA BY WORKING WITH THE COMMUNITY

Throughout the data collection process, the development team and the community must continually collaborate to assess the accuracy of the data; however, inclusivity should also be a consideration.



BRINGING IN EXISTING DATA

Some existing datasets might exist but not in a digital format. How can we integrate or digitize and what does consent look like in this case? Will the project have the resources available to do this? Another option might be to collaborate with content creators with some value-added or payment model that allows creators to see a benefit in participation.

LACK OF SAMPLE SIZE

Training AI models often require multiple examples of the same



thing, but it may be difficult to find multiple examples of cultural stories and practices, especially in small communities. In this case, the development team in collaboration with the focus group can designate golden data that will be weighted more heavily in the training process.

PROCESS FOR CULTURAL ACCURACY VERIFICATION

In working with intermediaries especially, how does one ensure intermediaries are not unknowingly altering data/adding their own impressions/interpretations? And if they are allowed to do so, how are we authenticating these to be true to the community data? One possibility is to work with multiple intermediaries and/or work with a wider group for data verification and annotation. Consider what accuracy means in the context of the community and the interplay of fact and elaboration in a storytelling/oral history context.

2.2 TRAINING/FINE TUNING THE MODEL

In the unlikely case of training a new model, training will be its own step, however, in both cases the overall values, actions, and obstacles remain very similar. In a training scenario, cost, resources, and limited data will be greater challenges, though they will still be challenges in the fine tuning scenario.



Culturally Sensitive: The model is attentive to nuance and different sensitivities. It could differentiate answers based on the user. It includes mechanisms for ongoing adaptation and evolution of culture.

Inclusivity: The fine tuning process ensures that the model data and responses are inclusive of gender, age, classes, education levels, socio-economics, professions, etc. It acknowledges and seeks to include people that don't use or don't have access to tech and/or are illiterate. Ensuring inclusivity may involve the physical collection of data, audio files, and stories (and has implications for the cost). It should be noted that convincing might be needed in some cases to help communities or people understand the value in inclusivity.



Additionally, the process of training/fine tuning accounts for different ways of knowing, allows for verification through non-written methods, and is grounded in community procedures. This is also reflected in different understandings of history and cultural narratives through community stories (i.e. the 'standard' and/or colonial version of history is not given more weight than community stories and narratives).

Adaptive and iterative: The approach to training/fine tuning is iterative with ongoing feedback loops and knowledge sharing.

Community driven: Community stakeholders are involved throughout the process in an integral way. They will define what is bias in their context and shape desirable/undesirable outcomes. They help annotate data and evaluate outcomes.



DEFINE BIAS AND DESIRED OUTCOMES

Clarify what counts as bias in a community-specific model. The definition should include all sections and classes within a community in the process and taboos, undesirable outcomes, etc. should be outlined by the community. The training algorithm can use moderators or guide rails to allow the model to refuse to answer or participate in certain queries. These should be present from the beginning but should be continually evaluated and updated to prevent as much as possible misuse of the data or model.

It is also important to be mindful of the digital with 'generative values' along with 'preservation values.'

CONNECT TO THE END USER/USE CASES

In order to arrive at a practical product that is connected with the originally envisioned use cases, the fine tuning process can learn from other models and potentially allow our model to cross-reference by drawing on relevant data from outside, and/or enlist similar models to be able to draw comparisons. Learn from the lessons and mistakes of others and the fine tuning should be focused on the specific context. However, if external data is used, questions



of data protection, copyrights, and bias must be addressed.

While the original use cases can drive initial fine tuning, as the community works more with the model, other uses may be discovered. The model should be adaptable enough to respond to these shifts in the short and long term.

ONGOING TRAINING AND LEARNING

The community must be continually involved at this stage; however, ensure the concepts are simplified and/or stakeholders receive the required training to understand how their input is reflected in shaping the AI. A collaborative knowledge sharing model approach can be best to ensure a more even balance of power, wherein both the development team and the community share knowledge and learn together.

Fig 7. Continual feedback and community responses correspond to ongoing fine tuning.





ISSUES AND LIMITATIONS WITH DATA

The quality of an AI relies heavily on the quality of data. Using existing data as a response to limited resources and/or data collection can expose biased data or data that is not accessible via language or other obstacles. Managing the model's capabilities and the availability of data will be challenging, especially in the requirement to support multiple languages and ways of knowing. Mitigating this will require an agile model and full community participation in collecting data.



COVERT BIASES IN EXISTING LLMS

In training an existing model on community specific data, the existing biases (explicit or covert, such as favoring western ways of knowing or western narratives.) Copyright infringement and faulty translations could also be an issue.

ACKNOWLEDGING THE EVOLVING NATURE OF CULTURE

In order to allow the model to continually evolve, it is important to consider the allocation of resources for ongoing fine-tuning, iterations, and evaluation. At minimum, the project must build feedback mechanisms such that the model can re-adapt to changes in the evolving culture (i.e. allow users to report if things do not align with the community culture.) The community should be empowered to continue adding data and input and frameworks for onboarding new users including the code of conduct and other guidance should be developed. We should be mindful as well of the continuity and longevity of AI in the rapidly evolving tech sphere.

LIMITATIONS DUE TO COST AND COMPUTATIONAL POWER

The sustainability and longevity of the project will rely on the availability of resources, both in money, data, and computational power. The project should aim to collaborate as much as possible with local developers already building local models and establish other collaborative methods of ongoing support and contributions.

CONCLUDING PHASE 2

At the end of phase 2, at least a minimally viable product has been developed and partially deployed, with mechanisms for ongoing feedback established. Ideally, the community through the intermediaries can be empowered to become the guardians of the AI and supervise its ongoing evolution. Long term sustainability can be an issue, however, as the AI grows to other communities the continual evaluation and maintenance of previous communities can be incorporated.

PHASE III EVALUATION AND GROWTH

This final section looks at the stage after the build is mostly complete; however, evaluation may begin as soon as testing of the fine tuning begins, and will be an iterative and ongoing process. Although data governance is treated in the last section here, it may be discussed as early as onboarding. Ideally it should be discussed in conjunction with data collection, as this may be part of the consent process. This section also touches on the post-project life of the AI and considerations of expanding the model to other communities.



3.1 EVALUATION

In order to ensure that the AI remains responsive and relevant, it must be evaluated from all aspects and stakeholders, especially end users on the community side. The focus group will decide on evaluation metrics together and importantly help create mechanisms for ongoing evaluation and iteration.



Representation: Evaluation methods and feedback opportunities are accessible to a wide number of users and responses, with various cultural nuances and perspectives. It should involve a core team of people with varied vested interests. An example of a diverse evaluation group might be: 1 evaluator from the community, 1 evaluator from the data bias of Global South perspective, 1 evaluator from the technical aspect, 1 evaluator who is not vested in any other motive other than for end-use-efficacy, 1 evaluator looking at the success of the project, financial or otherwise (In the onboarding phase, the focus group may define what success looks like.

Community driven: Evaluation metrics center the original end use and practicality of the AI. The community is integral in developing the metrics and evaluation questions.

Flexibility and continual evaluation: Evaluation needs to be constantly done regularly as part of the process, not at one time alone. At every step, the project must ensure that it has remained relevant to the needs of the community and pre-defined goals. Evaluation must consider the need for the model to keep evolving and growing.



BUILD A TECHNICAL DATA QUALITY TEAM

The evaluation team might not be the same as the development team, and should include the community.

DEFINE EVALUATION METRICS

(Note: some of these might be defined at the beginning with the end use.)



The metrics balance technical assessments and community assessments, and are based on the tasks, generating model outputs on the dataset, comparing those outputs to ground truth data using chosen metrics - both automated and human judgement. They assess the extent of hallucinations and inaccurate and/or insensitive output.

For community assessments, multiple outcome measures can be considered, i.e., measurable impact on the community (financially or in terms of facilitating daily tasks, etc.). Ideally, this goes beyond mere surveys and looks at randomized controlled trials in the field. Other outcome measures may quantify the efficiency of the model itself, e.g., by quantifying the degree to which it captures culture and tradition, etc. This could be by creating a test set, involving questions about the culture, etc.

User friendliness and interoperability of the platform is key, particularly whether or not the data are dynamic and useful and coming alive in the community in real time/life to solve a problem or add value to the community.

Other evaluation metrics that can be considered may connect to the AI's contribution to the broader goals of the project such as increased visibility of the culture, preservation or cultural value of the output, decreased bias in other models, etc.

CREATE FEEDBACK LOOPS AND INCENTIVIZE USE

Initial workshops can be conducted to use and explore the model and give feedback, involving the intermediaries, the community itself, and outside users to get a wide understanding of the model's capability. Ongoing participation and feedback can be incentivized by offering credits or tokens for use and moderation of the model.

Feedback loops should constantly circle back to training and fine tuning of the model, and include short and long term use cases to determine effectiveness. The exact protocol for testing can be decided later according to what is most appropriate, but could



include surveys, workshops, focus groups, independent feedback within the AI, randomized control trials, etc.

TEST THE LIMITS OF THE AI

Evaluation can include red teaming quality assurance, where the team conducts intentional jailbreaks to check if it's possible to make the model do things against what it was trained for, such as using certain prompts to make the AI answer in an offensive way or generate hate speech, etc.

EVALUATE THE PROCESS

Evaluation should also scrutinize the overall process and the AI against the guidance and methodology laid out in this Manifesto, in order to continually update this document. Each community can be considered as a test case for the Manifesto as a methodology and approach, and the focus group can contribute valuable feedback in how the values can more fully be enacted in practice. Eventually practical guides for different possible cases may be developed to accompany this more theoretical document.

Fig 9. The layers of evaluation from the generated output to the end user, expanding to users outside the community and the overall methodology and process.





CHALLENGES OF GATHERING ONGOING FEEDBACK

As evaluation will be an ongoing endeavor, it is important to explore how we can observe and evaluate interactions with the model and how we translate this into learnings and/or evolutions in the model. Automating some of the evaluation process, if possible, can help increase the sustainability of efforts.

3.2 DATA GOVERNANCE AND GROWTH

The final stage is the ongoing governance model of the collected datasets and how they can be responsibly shared, in a system that is devised and driven by the community. It is important to strike a balance between sharing data for the overall benefit of increased representation and decreased bias globally, while still allowing the communities to make clear decisions about how and where their data is used, and in what way they are compensated for it.



Community driven: Any model of data governance and sharing centers and prioritizes the community. Sharing the datasets and the benefits of wider circulation of the data should be clearly communicated to the decision makers as per the framework established in the Onboarding stage; however, the community should have the ability to control how and where their data is used, and receive some sort of compensation (monetary or otherwise) if the data is shared or sold.

Advocacy and learning: The approach to data governance and sharing should stand as a model for inclusive methodologies, AI that centers non western ways of knowing, and ethical representation of minority cultures in AI. This can include advocacy and marketing efforts, but should also include practical steps by putting learnings into practice, sharing research and code, and continually adapting this manifesto document to reflect this.

Flexibility and adaptability: If a single data governance strategy is articulated for the project overall, it will need to include a variety of usage situations, such as usage rights, commercial considerations,



and attribution. Different communities (and potentially different actors inside the same community) will have varying priorities, from monetization to ownership of the system and whatever way that looks, and the strategy needs to be adaptable to fit these different needs.



DEVELOP A COMPREHENSIVE DATA GOVERNANCE STRATEGY Apply the decision making framework outlined in the Onboarding stage to define how the data will be shared, acceptable/unacceptable scenarios for sharing, and models of compensation.

One model could be payment up front for data (which would be a conversation in the data collection stage), whereas another could be the community receives a percentage of future sales if the datasets are sold. It could be acknowledgement or attribution to the community, or any mix of the above.

The strategy will also need to address individual vs collective governance and interests, who has the right to speak/act/sell data on the community's behalf, and decision making frameworks outside of the focus group. It will address where the data is stored and who can access it - this could be a "secretariat" type of model appointed by the committee, or the community outlines their priorities and the project team is responsible for carrying these out in negotiations with other companies in sharing data.

It may be useful/required to align the strategy to national or regional data protection policies or regulations. It can make use of known data licenses, for example Creative Commons Attribution-NonCommercial-ShareAlike (CC-BY-NC-SA). Some ongoing nature or continuity of data protection contracts should be maintained for the scalability of the project.

SHARE THE MODEL CODE AS OPEN SOURCE AND PUBLISH RESEARCH

The development team can release the code as open source without the data attached as a way to share learnings from the

build and get feedback from outside. It would be advised as well to publish learning from the process and research around the methodology, technology, or overall approach may be conducted alongside the project. In this case, it is important that the researchers are onboarded into the project ways of working and approach the work with a community centric vision.



OVERALL GOVERNANCE OF THE WIDER MODEL AS MORE COMMUNITIES COME ON BOARD

As the project grows beyond one community, it might be necessary to adjust the models of data governance (as well as document this evolution and ensure that the original communities are kept in the loop.) One possibility is that each community that is onboarded can decide to join a global governance team that looks after the wider model. The global model may continue to grow while smaller models with higher customization could be rolled out for specific communities. Later peer-to-peer sharing and learning/collaborations between multiple communities that use the AI can be incorporated. Existing communities can help other / newer communities to join in and set up. Once 7-10 communities are engaged in the framework - a community-to-community networking could be established. A broader consortium as envisioned by AAOO could facilitate this growth, conversation, interaction and multiplicity.

It is also possible that different communities may have different enough needs to use separate foundational models, meaning that there may not be a single AI build that represents all communities. In this case, the uniting factor would be the methodology rather than the actual build; however, in the interests of sharing resources, working with one model is advised. This will be evaluated as more test cases are conducted.

MANAGING ONGOING MAINTENANCE

Once the focus group is disbanded and the model is 'handed over' to the community, there may still be some ongoing need for technical support. Cost and human resources could be an issue in this scenario. Sustainability and adaptability need to be built into all strategies to avoid obsolescence. Sharing data and ongoing learning can combat this; however, it must be done respectfully so as not to override or ignore the community-centered nature.

CONCLUDING PHASE 3

At the end of phase 3, the focus group and intermediaries conclude their work. The AI model continues to adapt and evolve using the established feedback and evaluation mechanisms, and the learnings from the process, both culturally and technically, are shared with the wider AI industry. If resources allow, new test cases and communities can expand the broader AAOO consortium; in any case, advocacy efforts can help expand awareness and growth.

CONCLUSION

The above represents an attempt to practically define a methodology for building a generative AI model that is responsible, responsive, and communitycentered particularly with a global south focus and intentional inclusion of diverse ways of knowing and sharing knowledge.

It is important to note that much of the above is still theoretical in the sense that deep understanding will come when it is put into practice via the test case(s). It is clear that many of the steps overlap; for example, decisions about evaluation may happen in the onboarding phase, while data collection, fine tuning, and evaluation may end up happening simultaneously. Centering transparency and shared decision making will be key to gauge how much information and training has to be shared along the way with communities, especially those with less technical literacy, so as to not place all the burden of learning and deciding at one time. The focus group will need to continually assess and discuss with intermediaries what needs to be addressed and decided at the time.

As the project grows and development progresses, it will be important to remember the living and evolving nature of culture and technology as it reflects in this document and methodology. Rather than attempting to codify each stage to an exact science, we will need to accept the experimental approach and swiftly changing landscapes around us. Rather than seeking to erase the dynamic tension of questions such as the balance between overarching AAOO goals and specific community needs, or the different needs for data governance and acknowledgement, we must create robust collaborative and transparent networks that respond to these questions as living things.

Further, we may find that the overarching goals of AAOO and the community could converge. One idea that emerged from the working group was to fully center audio and voice input, putting 'listening', 'oral histories', 'stories' at the forefront, an essential global south form of cultural preservation. A potential product of the AI could be to process the community audio input and conversations to create a weekly FM show in the local dialect that can be

broadcast anywhere in the world, which could collate speeches, prayers, announcements, sounds of the local environment, including topics of living and past, tangible and intangible, natural and built, imagined and real culture of the community. Radio as an end use, outside of complicated technology, could be easily understood and shared even in communities with less digital literacy.

Whatever the final format, model, or end use, the key point the manifesto brings up is a collaborative approach, where knowledge is shared freely, decisions are made jointly, and community needs are consistently respected. Traditional power dynamics are addressed and acknowledged, but the effort is continually made to center the community's knowledge and needs, leading to a responsive and practical model that benefits all.



ANNEX A: LIST OF RESOURCES SHARED WITH THE WORKING GROUP

Session 1: Overall Design and Methodology:

- 1. Is AAOO a way of doing or a way of being? How can the ethics/vision/approaches of AAOO be integrated from the beginning?
- 2. What are the overarching guidelines we want to maintain for the process? (ie do no harm)
- How does AAOO approach the overall ecosystem of the model (what methodology or approach do we adopt, i.e. kinship, ecosystem, stewardship, etc)?
- 4. What can we codify, and what can't we?
- 5. How can we respect pluralism as well as the particular needs of people not currently included in AI discourse (ie, global south, minorities, colonized or oppressed peoples, Indigenous groups in North America)?

AI Ethics and decoloniality, essays 4, 5, 6, and 9:

https://www.chathamhouse.org/2024/06/artificial-intelligence-and-challeng e-global-governance-OLD/about-authors_

The Copenhagen Manifesto on human centered generative AI:

https://www.researchgate.net/publication/381052824_Generative_Al_in_Sof tware_Engineering_Must_Be_Human-Centered_The_Copenhagen_Manifesto

AAOO Landscape Mapping Research

Session 2: Designing and Onboarding:

- 1. What stakeholders must be on board for the design, and what role do they play? Do they require any training to be involved?
- 2. Where does AAOO need to be explicitly multicultural/pluralistic, and where does it need to address the specific needs of the community that is working with it?
- 3. Does AAOO begin with the end use case? Is it necessary to have an outcome to begin? Does this affect the design of the model?

4. Does the existence (or lack thereof) of resources, in data or computing, affect how AAOO is designed? How does sustainability and resources affect the community-centering of the project?

Empirical ethics: Joint DSC/IAS Fellowship Call on Empirical Ethics and Data Science Methods - Data Science Centre - University of Amsterdam:

https://dsc.uva.nl/shared/subsites/uva-institute-for-advanced-studies/en/ne ws/2024/11/joint-dsc-ias-fellowship-call-on-empirical-ethics-and-data-scienc e-methods.html

The Manifesto for Teaching & Learning in a Time of Generative AI: A Critical Collective Stance to Navigate the Future:

https://www.researchgate.net/publication/381052824_Generative_Al_in_Sof tware_Engineering_Must_Be_Human-Centered_The_Copenhagen_Manifesto

Making Kin With the Machines:

https://www.researchgate.net/publication/326614247_Making_Kin_with_the _Machines

Bringing AI participation down to scale:

https://arxiv.org/abs/2407.11613

Session 3: Tackling the Data Question:

- 1. What does "consent" mean for data sourcing? What does consent mean in the context of group as well as that of individual consent?
- 2. What kind of data do we need to gather in a cultural heritage context?
- 3. How will AAOO protect the data of its communities?
- 4. What does data governance mean in AAOO, and how does it work practically?

An article on participatory data gathering in Zambia.

https://datavaluesdigest.substack.com/p/in-zambia-participation-is-creating

An article on data governance from the perspective of Indigenous language data in NZ.

https://www.context.news/ai/nz-us-indigenous-fear-colonisation-as-bots-lea rn-their-languages

The website for this Indigenous data sovereignty initiative: https://www.temanararaunga.maori.nz/

Project websites for initiatives that work closely with communities on data sourcing:

https://www.catamaborneo.com/about ; https://www.malaysiadesignarchive.org/geraioa/ ; https://digitalbenin.org_

ANNEX B: PROFILES OF THE WORKING GROUP

Arthur Gwaga - https://www.linkedin.com/in/arthurgwagwa/ Malik Afegbua - https://www.linkedin.com/in/malik-afegbua-04606b66/ Talal Rahwan - https://www.linkedin.com/in/talal-rahwan-a654385/ Tanya Dutt - https://www.linkedin.com/in/tanya-dutt-51382b336/ Sameer Raina - https://www.linkedin.com/in/samraina/ Bobina Zulfa - https://www.linkedin.com/in/bobina-zulfa-0820471a4/

Project Team:

Phloeun Prim – <u>https://www.linkedin.com/in/phloeunprim/</u> Gillian Rhodes – <u>https://www.linkedin.com/in/gillianrhodesofficial</u>

LAI website: www.livingartsinternational.org

ANNEX C: REFERENCES

Artificial intelligence and the challenge for global governance | Chatham House – International Affairs Think Tank. (2024, June 6). https://www.chathamhouse.org/2024/06/artificial-intelligence-and-challeng e-global-governance

Awarri. (n.d.). Whitepaper: Leveraging 'Native Intelligence' To Drive The Al Revolution In Africa. Awarri. Retrieved October 8, 2024, from https://www.awarri.com/whitepaper-leveraging-native-intelligence-to-drive-th e-ai-revolution-in-africa

Hao, K. (2022, April 22). A new vision of artificial intelligence for the people. MIT Technology Review.

https://www.technologyreview.com/2022/04/22/1050394/artificial-intelligen ce-for-the-people/

Kulesz, O. (2024). Artificial Intelligence and International Cultural Relations: Challenges and Opportunities for Cross-Sectoral Collaboration. In Ifa-Edition Kultur und Außenpolitik; 2024 (p. 1345 KB, 119 pages) [Application/pdf]. ifa (Institut für Auslandsbeziehungen e.V.). https://doi.org/10.17901/1203

Li, T. (2024, October 8). In data governance, good policies aren't enough [Substack newsletter]. The Data Values Digest.

https://datavaluesdigest.substack.com/p/in-data-governance-good-policies? publication_id=452561&utm_campaign=email-post-title&r=2ibr40&utm_medi um=email

Masakhane. (n.d.). Retrieved July 19, 2024, from https://www.masakhane.io/

Tonja, A. L., Dossou, B. F. P., Ojo, J., Rajab, J., Thior, F., Wairagala, E. P., Aremu, A., Moiloa, P., Abbott, J., Marivate, V., & Rosman, B. (2024). InkubaLM: A small language model for low-resource African languages (arXiv:2408.17024). arXiv. http://arxiv.org/abs/2408.17024