
AI literacy is AI Safety

Karabo Mokoena 1 Zwakele Mbanjwa 2
AI Safety South Africa BigAfrica

With
Apart Research

AI policy in Africa Track

Abstract

Only 1 in 3 South Africans knows what artificial intelligence (AI) is (Ipso, 2025). Despite spending more than nine hours daily online, just 4% of citizens learn about AI formally and only 2% through their workplaces. This awareness gap leaves the population vulnerable to misinformation, exploitation, and missed opportunities. We argue that AI literacy is an AI safety imperative: an AI-literate population can question harmful applications, participate in governance, and use AI to solve problems in their own communities.

This paper proposes and demonstrates a six-step pipeline that converts peer-reviewed South African AI research into multilingual, safety-centred social media content. The pipeline uses Lelapa AI's Vulavula API to translate content into South Africa's 11 official languages. The content will be distributed by content creators and influencer partners. For funding we propose using private sector sponsors who are already engaging in literacy programmes and who need an AI literate population and workforce. This removes dependence on government budget cycles and incentivises creators by providing an income stream for their posts.

We run the pipeline on two South African papers. The example on algorithmic bias produces creator formats which include a skit, a storytime reel, and a WhatsApp voice note. These formats can build awareness of how AI outputs can contain bias. The second, on isiZulu natural language processing, produces a tutorial that shows audiences how to use a free South African AI tool in their home language. Together, the examples show that the same pipeline can serve both risk awareness and capability building, at scale, across diverse contexts.

We use South Africa as the primary case study and design the model for replication across the Global South.

1. Introduction

For the Global South, the AI transformation presents risks and opportunities that must be navigated in context. South Africa is a case in point that already embeds AI in hiring systems, credit decisions, and public services, but most South Africans do not know what it is. AI is also already capable of filling resource and knowledge gaps that South Africans can leverage to produce solutions to local problems, but these are missed due to knowledge gaps (CSIS, 2026).

Ng et al. (2021) define AI literacy as four competencies: knowing and understanding AI, using and applying AI, evaluating and creating AI, and engaging ethically with AI. We use this framework as our conceptual foundation, as it has a holistic definition of AI literacy that encompasses awareness and capacity building.

We contend that AI literacy is a foundational layer of AI safety and an enabler for leveraging AI benefits. The lack of AI literacy creates a national strategic vulnerability and safety risks (Setsoafia & Ng, 2025). Specifically, as a middle power, South Africa can not presently dictate how AI tools are developed, but can diffuse AI so as to leverage the benefits, but should do so safely. A society that understands what AI is, what it is capable of doing and what its risks it presents is better positioned to hold stakeholders accountable and use it to innovate safely. By putting safety knowledge in the hands of the population through channels they already use, people will better understand AI's capabilities and limitations and can innovate to solve local problems while guarding against local risks.

This claim rests on an explicit theory of change that situates the work within AI safety and societal resilience, as an advantage for middle powers who cannot dictate development but can guide diffusion of AI for their benefit. Each of the four AI literacy competencies in the Ng et al. (2021) framework maps to a distinct safety and national interest function. Knowing and understanding AI lets people recognise when a decision is automated and what it can get wrong, the precondition for any individual safeguard. Using and applying AI lets people adopt local-language tools safely instead of depending on systems trained on data that excludes them. Evaluating AI lets people question biased or harmful applications. Engaging ethically with AI builds the public constituency that holds developers and deployers to account. This is the causal pathway from literacy to safety that the pipeline is designed to activate.

Our main contributions are:

1. A people-centred AI literacy pipeline that grounds AI safety in public understanding, designed for South Africa's linguistic and socio-economic diversity and adaptable to other Global South contexts.
2. A distributed content production model that connects academic research to social media audiences through research liaisons, AI-powered translation tools, and independent content creators.
3. A sustainable funding model that routes private sector sponsorship directly to creators, removing government from the financial relationship and protecting editorial independence.

2. Related Work

We reviewed the AI strategy and policy landscape in Africa through the AI Policy Hub database. Currently, 39 countries are developing or implementing national AI strategies. Countries with active policies or draft frameworks include: Kenya, Tanzania, Rwanda, Uganda, and Ethiopia in East Africa; Nigeria, Ghana, Benin, Cote d'Ivoire, and Senegal in West Africa; Egypt and Tunisia in North Africa; and South Africa, Zambia, Mauritius, Namibia, Lesotho, and Zimbabwe in Southern Africa (AI Policy Hub, 2024).

A subset of nations has identified that professional literacy and technical empowerment are essential for safe AI implementation. Tanzania developed a Draft National Digital Education Strategy integrating digital competencies into the national curriculum (Ministry of Education, Science and Technology, 2024). Kenya published the Artificial Intelligence Practitioners' Guide, providing technical literacy and standards for AI developers and deployers (Kenya Ministry of Information, Communications and The Digital Economy, 2025). Nigeria released the NBA Guidelines for the Use of Artificial Intelligence in the Nigerian Legal Profession, establishing sector-specific AI literacy within the legal profession (Nigerian Bar Association, 2024). Rwanda uses Digital Adoption Implementation Guidelines to support government and societal integration of AI (Rwanda Information Society Authority, 2023). At the continental level, the African Union's Child Online Safety and Empowerment Policy provides a foundational framework for digital literacy and safety for younger populations (African Union, 2024).

Of the nations with dedicated AI strategies, most AI literacy programmes focus on capability building in formal education settings or with the pretext of digital literacy, and fewer address safety as the organising frame. Specifically, almost none deploy social media influencer networks as the primary distribution channel at scale.

Existing initiatives and their limits

We investigated existing AI literacy initiatives that are similar to what we are proposing. The focus was on African initiatives with a narrow look at some notable international approaches.

TikTok's AI Media Literacy Fund committed \$200,000 at the third Sub-Saharan Africa Safer Internet Summit in Nairobi, partnering with Africa Check, Mtoto News, and CJID (TikTok, 2026). The fund focuses on AI-generated misinformation and deepfakes, not general AI literacy. Mtoto News produces content that turns young people into digital detectives who can identify manipulated video. This is the closest existing African precedent for using a social media platform as both funder and distribution channel.

THRAETS, backed by the Africa Digital Rights Fund under CIPESA, built Spot the Fakes, a gamified quiz on synthetic media, and Community Fakes, a crowdsourced incident database where citizens can submit and analyse suspected deepfakes (CIPESA, 2024). This demonstrates demand for participatory, citizen-led AI safety tools in Africa. However, crowdsourcing relies on a population with low literacy rates and the method failed to reach a wide audience as it relied on crowdsourced incident data base and organised formal setting for quiz shows which misses a large part of the population.

AIfluence, a Kenya-based platform, uses AI-powered influencer networks for social awareness campaigns including health literacy and legal aid education across East Africa. It demonstrates that influencer-based public education reaches audiences that formal channels miss. It does not address AI literacy or safety as its subject matter, and lacks the mechanism to ground awareness in context, specifically for diverse populations like in South Africa.

The African AI literacy landscape clusters into three types: counter-disinformation tools targeting specific harms like deepfakes (TikTok, THRAETS); formal education and children's programmes in schools (Aphabet, ITU/Google School Clubs, RobotsMali); and technical community building for practitioners (Deep Learning Indaba, Masakhane, AI Saturdays). No existing African initiative combines a social media influencer distribution model with broad public AI safety literacy across all age groups and socio-economic contexts.

Internationally, Finland's Elements of AI (University of Helsinki and Reaktor, 2018) reached 1% of Finland's population in its first year and has since expanded to 170+ countries through a free online course. It frames AI as a skill to acquire, not a risk to navigate, and it launched in a high-ICT-skill, low-linguistic-diversity context that does not transfer directly to South Africa. This is the closest precedent for influencer-mediated public education, but it operates in a high-literacy, English-language context and focuses on online safety broadly, not AI safety.

What gap this work addresses

The project reframes AI literacy as a prerequisite for AI safety governance. It is the first intervention to apply structured social media influence methodology, proven in commercial and health communication contexts, to AI safety education in a multilingual African context. No existing programme combines:

1. social media influencer networks as the primary distribution channel,

2. AI safety as the organising frame rather than AI skills or consumer protection, and
3. design for a linguistically diverse, low-ICT-skill, high-social-media-engagement population in the Global South.

South Africa's existing strengths make the project feasible and strategically important. The country's success in using broadcast media and peer-led communication for public education on HIV and COVID-19 demonstrates that complex, sensitive topics can reach mass audiences through trusted voices (SANAC, 2021). This social media influencer model reaches audiences where they already are by using the formats they already consume. This method will work to promote an AI literate population since 73% of South Africans barely register the term 'AI' but get most of their information from social media (Ipsos, 2025).

3. Methods

This project proposes a six-step pipeline that converts South African AI research into multilingual, safety-centred social media content distributed through a network of content creators and influencer partners. The pipeline is described in sufficient detail for replication in comparable Global South contexts.

Research ingestion and curation

The pipeline begins with a research curation process which can involve research liaisons drawn from postgraduate students at South African higher learning institutions¹ or an automated AI tool such as SciSpace. The goal is to identify peer-reviewed findings that meet two criteria: (a) the finding has a direct implication for how AI affects everyday South Africans; and (b) it can be communicated without requiring technical background knowledge, as these will help make the content viral (Berger & Milkman, 2013). The idea is to use local research outputs, to take advantage of the 185% increase in AI research output in South Africa (UNESCO, 2025) and also ground content in local realities. Because all sourced material is already peer-reviewed, the curation role is not to validate scientific quality but to assess public relevance.

AI-Powered content translation

Curated research findings pass to a translation layer that converts academic language into content briefs designed for social sharing.

This step uses a learning and summary AI tool, for example NotebookLM to make summaries of the paper. The summaries can be in the form of a podcast, a mindmap, a video or simple text. This summary should then be inputted into a frontier LLM which is prompted with a standardised instruction set. The prompt should have the following characteristics:

¹ For example the Wits AI Institute, UCT's Information Systems department, and Stellenbosch University's Machine Learning Group,

1. The prompt must require the LLM to identify a moment of recognition in the finding, a scenario a South African would encounter in daily life and immediately see themselves in.
2. The prompt must surface the gap between what the audience assumes is true and what the research reveals. A finding becomes shareable when it disrupts an assumption.
3. The output must give the reader something they would want to pass on, a piece of knowledge that makes the sharer look informed and aware to their peers.
4. The LLM must be prescribed to structure the output into separate formats. For example, a spot the fake challenge, a relatable skit script, a single-stat graphic, a short video or animation, or even a podcast like output.

Multilingual translation into South Africa's 11 official languages uses two locally built tools: Lelapa AI's Vulavula API, covering isiZulu, isiXhosa, Sesotho, Northern Sotho, Setswana, Tsonga, Swati, and Afrikaans; and Bothale AI for speech-to-text and text-to-speech in audio formats. Using locally built tools ensures the pipeline supports South Africa's own AI ecosystem rather than routing language processing through foreign models trained on data that underrepresents African languages. All translated outputs are reviewed by native-speaking content creators to ensure cultural resonance.

Content production model

Creators receive the brief and have full creative freedom on format, style, and tone. The brief specifies the emotional hook and recommended formats, but does not dictate the output. A content creator in Durban making a 60-second isiZulu video and a creator in Cape Town making an Afrikaans infographic series are both working from the same brief but producing content that speaks to their own audience in their own voice. AI safety content that feels nationally produced travels less far than content that feels like it came from someone the audience already knows and trusts.

For a government adopting this model, content creators can be sourced through existing invitations to creators, competitions and auditions or creator networks already operating in the country.

Platform selection and distribution architecture

Platform choices should be driven by national usage data rather than global platform rankings. The relevant questions are: which platforms have the highest daily active users across income levels, which support multilingual content natively, which function on low-bandwidth connections, and which have the longest content shelf life. Platforms that meet all four criteria can become primary channels, while those that meet two or three can serve secondary distribution roles.

In South Africa, this principle produces the following architecture. YouTube serves as the primary channel because it has high viewership across age groups and income levels (92% penetration,

usage equal across demographics) (Helloyes, 2025), and provides a permanent searchable archive that extends content life beyond the initial posting cycle. TikTok and Instagram Reels serve as secondary channels for short-form content, reaching the 18-35 demographic (TikTok: 52.5% of adults 18+, Instagram: 64% aged 18-34) (Helloyes, 2025). WhatsApp serves community-level distribution where it is the most used platform (94% of internet users) (Statista, 2025). WhatsApp content is typically produced in compressed formats to function on low-bandwidth connections, though WhatsApp supports videos up to 16 MB.

Any other African country applying this model would run the same selection process against its own usage data and arrive at a platform mix suited to its population.

Influencer seeding strategy

Each content piece is seeded to a network of micro-influencers with 10,000 to 100,000 followers. The choice of micro-influencers over macro-influencers or celebrity endorsers follows evidence from influencer marketing research showing that micro-influencers generate higher engagement rates and are perceived as more credible and trustworthy by niche audiences (Hudders et al., 2021). In the context of AI literacy, perceived credibility is a critical variable as an audience already sceptical of AI is more likely to engage with a relatable peer than a celebrity.

Creators are incentivised through direct commercial relationships with private sector sponsors, drawn from sectors with a direct stake in public understanding of AI, including financial services, telecoms, and technology companies. Sponsors fund content themes rather than individual pieces: a bank might sponsor an 'AI and your money' series, a telecoms company an 'AI and your data' series. Payment flows directly between sponsor and creator. The private sector has a direct interest in having an AI literate staff, and organisations have already launched free AI literacy material. Beyond monetary value, creators retain their audience, receive co-creation credit on a government-endorsed initiative, and gain access to new audiences.

Measurement framework

The programme tracks impact across three outcomes. The first is sharing and community conversation, measured through share velocity, the rate at which content spreads beyond the initial influencer seed, and the ratio of organic to seeded reach. Content that travels on its own is content that has connected. Comment analysis on YouTube, TikTok, and Instagram can track whether audiences are actively discussing AI, not just consuming content about it. The second outcome is the number of clicks a secondary more formal AI literacy programme receives from the specific content. If there is an uptake of the already existing formal AI literacy training available that was accessed via a social media post that was linked to in that specific content. The third outcome is a shift in public confidence around AI.

Design alternatives considered and rejected

Two design decisions were reached after evaluating alternatives that were ultimately rejected. The first was the choice between SABC as the primary distribution channel. SABC has national reach and government credibility, and Microsoft's AI skills drive on SABC Plus demonstrated that it can deliver educational content to mass audiences (Microsoft, 2026). However, its audience skews older, its digital platform requires data access that excludes the most disconnected communities, and it cannot support the creator-led, multilingual, shareable content model this programme depends on. The second was the choice between government managing sponsor payments and a direct sponsor-creator financial relationship. Government administration of payments would give the programme more oversight and control, but it would slow the pipeline, create procurement bureaucracy, and tie creator income to government budget cycles.

4. Results

Appendices A and B demonstrate the six-step pipeline running end-to-end on two South African peer-reviewed papers. Together, they establish three things: that the pipeline produces content that is platform-native, that multilingual translation is technically achievable using existing local tools, and that the same methodology can serve two distinct audience outcomes depending on the research finding selected.

Appendix A uses Ngcobo et al. (2024) on algorithmic bias. The pipeline produced a skit, a storytime reel, and a WhatsApp voice note brief, all built around a single relatable scenario: a job rejection with no explanation. The finding enters the audience's world through a story rather than a statistic. This content is designed to shift awareness, not direct behaviour, which reflects how population-level attitude change works in health communication research. (Horsfield, 2013)

Appendix B uses Marivate and Sefara (2020) on isiZulu natural language processing. The pipeline produced a skit, a tutorial reel, and a WhatsApp voice note brief built around a different kind of moment: a small business owner realising that AI already speaks his language and that he can use it today. The call to action is practical and immediate, directing the audience to a free tool built by South African researchers. This content is designed to build capability, not just awareness.

The two examples together show that the pipeline is not fixed to a single emotional register. It can move between risk literacy and capability building depending on what the source research offers.

The translation step was run end to end against the live Lelapa AI Vulavula API. Across both briefs and all eight target languages, the pipeline issued 144 segment-level translation calls, of which 101 succeeded and 43 failed when a free-tier rate limit interrupted the second brief.

Appendix A returned complete translations in all eight languages. Appendix B returned isiZulu in full and isiXhosa in part before the limit stopped the remaining languages. The API produced usable first-draft text in every language it returned, which confirms that linguistic reach across South Africa's official languages is a present technical capability rather than a future aspiration.

Translation quality varied with the resource level of each language. Afrikaans, the highest resource language in the set, was the most fluent, while lower-resource languages showed more code-switching and untranslated English.

Three issues recurred across the output: whole English phrases left untranslated in harder sentences, inconsistent rendering of the acronym AI, and corruption of the POPIA acronym into variants such as P_O_P_A and P_P_I_A. The raw output is therefore a usable first draft that requires native-speaker post-editing, which is why the pipeline routes every translation through native-speaking content creators before publication.

The content formats in both appendices were designed to match what creators actually produce on South African social media. The skit format mirrors the two-person scene content that dominates TikTok in South Africa. The storytime format matches the first-person confessional style with high organic reach among 18 to 35 year-olds.

5. Discussion and Limitations

The pipeline proof of concept also reveals that South African researchers are producing work with direct implications for how AI affects everyday life. That work largely remains inside academic publishing and technical communities. The gap between what researchers know and what the public knows is a translation gap. The research exists, but no systematic mechanism converts it into content that can travel at the speed and scale of social media. Any Global South government with an active social media population can adopt this pipeline with contextual adjustments.

Limitations

Both appendix examples are illustrative, they don't show what the pipeline achieves. No empirical data exists yet on whether skit or storytime formats drive measurable shifts in AI awareness or capability in South African audiences. Implementation will generate the first real evidence on this question. The pipeline's quality also depends on the right research selection. The two papers chosen for this demonstration, Ngcobo et al. (2024) on algorithmic bias and Marivate and Sefara (2020) on isiZulu NLP, are well-suited to public communication because they have clear everyday implications and accessible hooks. Not all peer-reviewed AI research has these properties. For this model to work, there needs to be a functioning creator economy. The Appendix B tutorial format works because creators with audiences already exist in South Africa. In countries where that infrastructure is less developed, the influencer seeding step requires more investment in creator development before the pipeline can operate at scale. We also encountered issues with translation, as the native languages available are limited on LelapaAI.

Importantly, share velocity, comment volume and traffic from social media content to formal AI training courses are proxies, not direct measures of AI literacy. The shift from 'I saw a skit about AI bias' to 'I now think critically about AI decisions that affect me' is difficult to capture without validated assessment instruments.

Dual use and misuse

The same components that make AI literacy go viral are also a turnkey disinformation and influence engine. Repurposed, this machinery could spread false AI narratives, create a manufactured consensus, or push political messaging through a government-endorsed channel, and it would do so most effectively in low-resource languages where few independent fact-checkers can audit the output. The sponsor model adds a capture risk. Financial services, telecoms, and technology firms, are the sectors deploying the AI systems the public needs to scrutinise. This can create an incentive to soften criticism of those sponsors' own products, even with payment routed directly to creators.

Several safeguards reduce this exposure.

1. Every sponsored and seeded piece is labelled as such, so audiences can see the commercial and distribution relationship behind it.
2. A provenance standard restricts source material to peerreviewed research with the citation surfaced in the content, which makes fabricated claims harder to launder through the pipeline.
3. Native-speaker review of every translation, already part of the model, doubles as a check against biased or distorted framing at scale.
4. Independent oversight of theme selection, together with a public record of sponsors, keeps the programme accountable and protects the public trust the intervention depends on.

Future Work

A concrete first test would seed matched content sets, AI-literacy briefs against a neutral control, to comparable micro-influencer audiences, then compare share velocity, organic-to-seeded reach ratio, comment sentiment, and click-through to a linked formal course over a fixed two-week window. Pairing these behavioural proxies with a short pre- and post- survey on AI awareness in a sampled audience would give the first calibrated estimate of how far engagement translates into literacy.

6. Conclusion

This paper proposes a pipeline that connects South Africa's AI research findings, creators who are producing content that reaches the public and AI tools that speak the public's languages in order to mitigate against AI risks and equip the population with the thinking to start innovating towards local solutions.

The first example shows how a finding about algorithmic bias becomes a skit about a job rejection and a storytime about a loan that never came through. The second shows how a finding about isiZulu NLP becomes a tutorial that puts a free South African AI tool in a small business owner's hands. This model can be repeated at scale, across languages, through voices people already trust. AI literacy can start to reach the people who need it most if it is packed in the right way. It is also

what AI safety looks like when it moves beyond the technical community. An AI-literate population can recognise AI risks, question harmful applications, and use AI to build things that serve their own communities. Building that population should be an imperative for AI safety work.

Code and Data

This project is a programmatic intervention design. A proof-of-concept Python script integrating the Lelapa AI Vulavula translation API is provided in Appendix A. The script translates the plain-language brief from Appendix A into six South African languages and generates the three content format outputs. It is available to run locally with a free Vulavula API token from vulavula.lelapa.ai.

LLM Usage Statement

Claude (Anthropic) was used to assist in structuring and drafting sections of this report from the author's original presentation materials, research, and notes. All arguments, data points, and strategic recommendations originate from the author's research. All claims were independently verified against source materials before inclusion.

References

1. African Union. (2024). AU Child Online Safety and Empowerment Policy. Addis Ababa: African Union Commission.
2. African Union. (2024). Continental artificial intelligence strategy. Addis Ababa: African Union Commission.
3. AI Policy Hub. (2024). National AI strategies and policies in Africa. Retrieved from <https://aipolicyhub.oecd.org>
4. Berger, W., & Milkman, A. (2013). Why viral videos go viral: The role of psychological triggers in online content sharing. *Journal of Marketing Research*, 50(3), 317-332. <https://doi.org/10.1509/jmkr.50.3.317>
5. CIPESA. (2024). Introducing Community Fakes: A crowdsourcing platform to combat AI-generated deepfakes. Collaboration on International ICT Policy for East and Southern Africa. <https://cipesa.org/2024/11/community-fakes-by-thraets/>
6. CSIS. (2026). From Divide to Delivery: How AI Can Serve the Global South. Center for Strategic and International Studies. <https://www.csis.org/analysis/divide-delivery-how-ai-can-serve-global-south>
7. Helloyes. (2025). Digital Statistics & Usage South Africa 2025 Insights. <https://helloyes.co.za/digital-statistics-and-usage-south-africa-2025/>

8. Horsfield, B. (2013). The sciences of science communication. *Proceedings of the National Academy of Sciences*, 110(Supplement 3), 14033-14039.
<https://doi.org/10.1073/pnas.1213273110>
9. Hudders, L., De Jans, S., & De Veirman, M. (2019). The commercialisation of social media stars: A literature review and conceptual framework on the strategic use of social media influencers. *International Journal of Advertising*, 40(3), 327-375.
10. Ipsos. (2025). Global AI survey 2025: South Africa module (Google-Ipsos "Our Life with AI"). London: Ipsos.
11. Kenya Ministry of Information, Communications and The Digital Economy. (2025). *Artificial Intelligence Practitioners' Guide*. Nairobi: Government of Kenya.
12. Marivate, V., & Sefara, T. (2020). Improving low-resource languages in NLP: Observations from isiZulu. arXiv:2005.00300. University of Pretoria / CSIR.
13. Microsoft. (2026). Microsoft and SABC launch national AI skills initiative on SABC Plus. Johannesburg: Microsoft South Africa.
14. Ministry of Education, Science and Technology. (2024). Draft national digital education strategy 2024/25 - 2029/30. United Republic of Tanzania.
https://www.moe.go.tz/sites/default/files/DIGITAL%20NATIONAL%20STRATEGY%202025_0.pdf
15. Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, 2, 100041. <https://doi.org/10.1016/j.caeai.2021.100041>
16. Ngcobo, S., et al. (2024). Navigating algorithm bias in AI: ensuring fairness and trust in Africa. *Frontiers in Research Metrics and Analytics*.
<https://doi.org/10.3389/frma.2024.1486600>
17. Nigerian Bar Association. (2024). *NBA Guidelines for the Use of Artificial Intelligence in the Nigerian Legal Profession*. Lagos: Nigerian Bar Association.
18. Republic of Rwanda. (2024). *Digital Adoption Implementation Guidelines*. Kigali: Rwanda Information Society Authority (RISA).
19. SANAC. (2021). *South Africa national HIV, TB and STIs communication strategy 2021-2026*. Pretoria: South African National AIDS Council.
20. Statista. (2025). *South Africa: social media users by platform 2024*.
<https://www.statista.com/statistics/1312488/social-media-users-by-platform-in-south-africa/>

21. Setsoafia, P., & Ng, D. T. K. (2025). Can We Be AI-Literate? The State of GenAI Literacy in African Countries: Ghana as a Case Study. In D. T. K. Ng, S. K. W. Chu, X. O'Dea, & J. K. Lok Leung (Eds.), *From AI Literacy to Generative AI Literacy: Bridging Human and Machine: Future Education with Intelligence* (pp. 203-220). Springer, Singapore. https://doi.org/10.1007/978-981-95-0950-8_12
22. TikTok. (2026). TikTok announces additional \$200,000 investment in AI media literacy for Sub-Saharan Africa. Nairobi, Kenya: TikTok Sub-Saharan Africa (at Safer Internet Summit). <https://techbuild.africa/tiktok-ai-media-literacy-sub-saharan-africa-200k/>
23. UNESCO. (2025). South Africa: Artificial intelligence readiness assessment report (Document Code: SHS/BIO/2025/AI-RAM/SA/1). UNESCO.
24. University of Helsinki & MinnaLearn. (2018). *Elements of AI*. Helsinki: MinnaLearn. <https://www.elementsofai.com>

Appendix A

This appendix demonstrates the six-step pipeline in action using a single peer-reviewed source paper. It shows how the pipeline transforms academic research into a content brief, multilingual translations, and three ready-to-produce content formats.

Step 1: Source Paper and Selection Rationale

Source paper: Ngcobo, S. et al. (2024). Navigating algorithm bias in AI: ensuring fairness and trust in Africa. *Frontiers in Research Metrics and Analytics*.

<https://doi.org/10.3389/frma.2024.1486600>

Selection rationale: The research liaison panel selected this paper for three reasons. First, it directly implicates everyday South African experiences: job applications, loan decisions, and housing assessments. Second, it names a specific and measurable harm, facial recognition error rates up to 35% higher for dark-skinned faces, that can be communicated without technical background.

Step 2: Plain-Language Brief (English)

The following brief was produced by prompting Claude with the four influence principles: relatability, tension and surprise, social currency, and format fit. The brief is the input given to content creators and influencer partners.

AI systems can be unfair — and most South Africans don't know it.

When a bank uses AI to decide whether to give you a loan, or when a company uses AI to shortlist job applications, the AI can make mistakes that are not random. They can be systematic. They can unfairly disadvantage people based on their race, gender, or where they live — not because someone chose to discriminate, but because the AI learned from historical data that already contained those inequalities.

This is called algorithmic bias. A 2024 study published in *Frontiers in Research Metrics* found that AI systems trained on data that does not represent African populations are particularly likely to produce biased outcomes. In South Africa, where most AI tools are built using data from the Global North, this risk is significant.

Facial recognition systems misidentify dark-skinned faces at error rates up to 35% higher than light-skinned faces. If a South African employer or law enforcement agency uses such a system, the consequences for the person misidentified can be serious — from rejected job applications to wrongful accusations.

You have a right to know. Under POPIA, any organisation using your personal data in an automated decision must be able to tell you how that decision was made. If they cannot explain it, they may be breaking the law.

Influence principle applied: Tension and surprise (the AI disadvantaging you is not random — it was trained that way). Format prescribed: bold single-stat graphic for social media, plus a 90-second explainer video.

Step 3: Multilingual Translations (Vulavula API)

The plain-language brief is passed to Lelapa AI's Vulavula API, which translates it into South Africa's official languages. The table below shows the opening line of the brief translated into four languages. Full translations of the complete brief are generated programmatically and distributed to creators in each language community.

Language	Opening line (Vulavula API output)
isiZulu	Izinhlelo ze-AI zingaba nokungalungi — futhi abantu abaningi baseNingizimu Afrika abakwazi lokhu.
isiXhosa	Iinkqubo ze-AI zinokungalungisi — kwaye uninzi lwabantu baseMzantsi Afrika abayazi oku.
Sesotho	Lits'ebetso tsa AI li ka se loke — 'me batho ba bangata ba Afrika Borwa ha ba tsebe sena.
Setswana	Dithulaganyo tsa AI di ka nna tsa se siame — mme batho ba bantsi ba Afrika Borwa ga ba itse seno.

Table A1. Opening line of the plain-language brief translated into four official South African languages via the Lelapa AI Vulavula API. Full-length translations of the complete brief are generated for all eight supported languages: isiZulu, isiXhosa, Sesotho, Northern Sotho, Setswana, Tsonga, Swati, and Afrikaans.

Step 4: Content Formats

The translation step produces three format prescriptions tied to the finding. Each uses a relatable everyday scenario rather than a legal or institutional frame. The goal is to leave the audience thinking differently about AI, not directing them toward a specific action or organisation. Each format is presented as a production-ready brief that a creator receives alongside the plain-language brief.

Format A: Skit (60 seconds — TikTok / Instagram Reels)

Creator notes: Two characters. Film it like a scene you stumbled into. Keep it light until the reveal. The punchline is the moment the audience realises the system, not the person, made the call.

SCENE: A friend group chat. Thabo just got rejected for a job he was perfect for.

Naledi: *"But you had five years experience, how did they say no?"*

Thabo: *"They said my CV did not meet the requirements. I do not even know which ones."*

Naledi: *"Did a human even read it?"*

Thabo: *"I do not know. Probably not."*

[Pause. Naledi looks at camera.]

Naledi: *"So an AI said no. And it was trained on data that was never really about people like us."*

[Beat.]

Thabo: *"Wait, is that actually how it works?"*

[Creator turns to camera] *"More than you think. Start asking questions."*

Format B: Storytime Reel (90 seconds — TikTok / YouTube Shorts)

Creator notes: First person. Sit down, gather round energy. Tell it like it happened to someone you know. Pause at the reveal. End by looking directly at the camera.

"Okay so let me tell you about my cousin. She applied for a home loan. Good job, stable income, no debt. Rejected. No reason given."

"She tried again at a different bank. Rejected again. Same profile. Same numbers."

"What she did not know, what most of us do not know, is that those banks use AI to make that call. And that AI was trained on decades of financial data. Data from a country that did not look like her, did not have her name, did not come from where she came from."

"A 2024 study showed that AI systems built on data from the Global North consistently produce worse outcomes for people like us. That is not a conspiracy. It is just the wrong data doing the maths."

"So before you accept that outcome, ask yourself: was it a person who decided? Or was it an algorithm? And do you know what it was trained to value?"

Format C: WhatsApp Voice Note Brief

Creator notes: Record in your home language. Keep it under two minutes. Warm and direct, like a message to someone you trust. This is the shape of the message, not a word-for-word script.

Open: Start with something that happened to you or someone nearby. A job rejection, a loan that did not go through, something that felt unfair but had no explanation.

The reveal: These decisions are increasingly made by AI systems, not people. And those systems were built using data that mostly does not include people like us. Researchers have shown that this produces real, measurable bias. It is not random. It is systematic.

The nudge: You do not need to be a tech person to understand this. You just need to start asking: when I get an outcome that does not make sense, was a machine involved? And what was it trained to value?

Close: Forward this to someone who has ever been told no without a reason. They deserve to know how it works.

Steps 5 and 6: Distribution and Measurement

The three content formats above are seeded to micro-influencers in the creator network. Each influencer receives the brief and presents the content in their own voice. A financial services sponsor funds the 'AI and your money' content theme, which includes this piece.

Measurement for this piece tracks: (1) share velocity from initial influencer seed across TikTok, Instagram, and WhatsApp; (2) comment analysis tracking whether audiences ask questions about AI decision-making rather than just reacting to the content; (3) sentiment in social listening data for terms including 'AI bias', 'algorithm', and 'how does AI decide' during the two-week window after publication; and (4) comparative reach of the isiZulu, isiXhosa, Sesotho, and Setswana versions against the English version to measure equitable language distribution.

Appendix B

This appendix demonstrates a second pipeline run using a South African NLP research paper. Where Appendix A aimed to raise awareness of AI risk, this example aims to build capability: showing audiences that AI already speaks their language and that they can use it to build something real.

Step 1: Source Paper and Selection Rationale

Source paper: Marivate, V., & Sefara, T. (2020). Improving low-resource languages in NLP: Observations from isiZulu. arXiv:2005.00300. University of Pretoria / Council for Scientific and Industrial Research (CSIR).

Selection rationale: The research liaison panel selected this paper for three reasons. First, it is grounded in South African research by South African scientists, making it directly relevant and attributable to local expertise. Second, it shows that AI can already understand and generate text in isiZulu, a finding with immediate practical implications for anyone who communicates in a South African language. Third, the tools it describes are accessible: free APIs now exist that any creator, small business owner, or community organiser can use today without writing a single line of code.

Step 2: Plain-Language Brief (English)

The following brief was produced by prompting Claude with the four influence principles: relatability, tension and surprise, social currency, and format fit. The prescribed format for this finding is a tutorial or demonstration, because the best way to show that AI speaks your language is to show it, not explain it.

AI already speaks isiZulu. And you can use it right now.

South African researchers have spent years teaching AI to understand, read, and write in isiZulu and other South African languages. The work started at institutions like CSIR and the University of Pretoria and has now produced tools that anyone can access for free.

This matters because most AI tools were built in English, for English speakers. If you run a business, serve a community, or want to reach people who think and feel most clearly in isiZulu, Sesotho, or Setswana, you no longer need to communicate in a language that is not yours.

Tools like Lelapa AI's Vulavula API let you translate, transcribe, and generate content in South African languages. You can use them to write product descriptions your gogo can read, build a WhatsApp auto-reply that speaks to your customers in their language, or create social media content that does not lose anything in translation.

The researchers built the tools. Now it is your turn to use them.

Influence principle applied: Social currency — sharing this makes the creator look ahead of the curve. Format prescribed: tutorial demonstration, showing the tool working in real time.

Step 3: Multilingual Translations (Vulavula API)

The opening line of the brief is translated below into four languages via the Lelapa AI Vulavula API. Note that for this example the brief itself is about AI language tools, so the translation is both the demonstration and the message.

Language	Opening line (Vulavula API output)
isiZulu	I-AI ikhuluma isiZulu kakade. Ungayisebenzisa manje.
isiXhosa	I-AI ithetha isiXhosa ngoku. Unokuyisebenzisa ngoku.
Sesotho	AI e bua Sesotho. Mme o ka e sebedisa hona joale.
Setswana	AI e bua Setswana. O ka e dirisa jaanong.

Table B1. Opening line of the plain-language brief translated into four South African languages via the Lelapa AI Vulavula API. The translated line is itself a demonstration of the finding: the brief about AI language tools is delivered in the languages it describes.

Step 4: Content Formats

The three formats below are designed for a creator who is demonstrating AI, not just describing it. Each puts a tool in the audience's hands rather than explaining it from a distance.

Format A: Skit (60 seconds — TikTok / Instagram Reels)

Creator notes: Two characters. One is sceptical. The other shows them something on a phone. The reveal is that the tool just worked in isiZulu. Keep it natural, do not over-explain.

SCENE: Two friends. Sipho is setting up his spaza shop WhatsApp line. Lebo wants to help.

Sipho: *"My customers speak isiZulu but I only know how to type in English on my phone."*

Lebo: *"What if I told you there is a South African AI that speaks isiZulu?"*

Sipho: *"AI? For isiZulu? Where?"*

[Lebo types a message in English into the Vulavula tool. It comes back in isiZulu. They both look at the screen.]

Sipho: *"That is actually isiZulu. That is correct."*

Lebo: *"South African researchers built this. It was trained on our languages."*

*[Siph*o picks up the phone. Starts typing his first isiZulu product listing.]

Format B: Tutorial Reel (90 seconds — TikTok / YouTube Shorts)

Creator notes: Screen record your phone or laptop while you talk over it. Show the actual tool working. You are teaching, not performing. Keep the energy curious, not hype.

"I want to show you something that South African researchers built, and that you can use right now for free."

"It is called Vulavula. It is an AI tool made by a South African company called Lelapa AI. It was trained specifically on African languages, including isiZulu, isiXhosa, Sesotho, Setswana, and more."

"Watch what happens when I type this in English." [Creator types: 'Welcome to my shop. Today we have fresh bread and amasi.'] "Now I press translate to isiZulu." [Shows result on screen.]

"That took three seconds. And it was built here, by people who understood that most AI was not made for us."

"You can use this for your business, your community group, your content. Go to vulavula.lelapa.ai. Free to try. South African-made."

"The researchers did the hard work. You just have to start using it."

Format C: WhatsApp Voice Note Brief

Creator notes: Record in your home language. Under two minutes. Talk to your audience like you are telling a friend about something you just discovered. Show enthusiasm, keep it grounded.

Open: Tell them you found something made here, in South Africa, by South African scientists. Set the pride frame first.

The demo: Describe what the tool does in plain terms. It reads your language. It writes your language. You type in English, it comes back in isiZulu or Sesotho or whatever language your people use. You can use it to talk to your customers, your community, your family, in their language.

The practical push: Give them one specific thing to do. Write one message in English. Translate it. Send it to one person who does not usually get messages in their language. Watch what happens.

Close: The tool address is vulavula.lelapa.ai. Share this note with someone who runs a small business or a community group. Let them try it.

Steps 5 and 6: Distribution and Measurement

These formats are seeded to micro-influencers with audiences in small business, community organising, and digital skills spaces. A telecoms sponsor funds the 'AI and your language' content theme. Payment flows directly between sponsor and creator.

Measurement tracks: (1) share velocity and saves, since tutorial content that people save to return to signals genuine intent to use the tool; (2) click-through to vulavula.lelapa.ai as a direct signal of audience action; (3) comment analysis tracking whether people share their own results, which indicates the tutorial converted viewers into users; and (4) reach of the isiZulu, isiXhosa, Sesotho, and Setswana versions relative to English, to confirm the message reached the communities it was made for.

Appendix C

Full Raw Machine-Translation Output (Verbatim)

This appendix reproduces the complete, unedited output of the live Lelapa AI Vulavula run, exactly as the API returned it. It is included for transparency and reproducibility. No post-editing has been applied, so the artefacts discussed in the Results and Limitations sections are visible here in full: the acronym AI rendered as A_I or I_A, the POPIA acronym corrupted into P_O_P_A, P_O_B_I, and P_P_I_A, untranslated English fragments, and cross-language bleeding. Of 144 segment-level calls, 101 returned successfully and 43 failed when a free-tier rate limit interrupted the second brief.

C.1 Brief 1 (algorithmic bias): all eight languages returned

isiZulu (zul_Latn): izinhlelo ze-A_I zingenza okungalingile, kanti abantu abaningi baseNingizimu Africa abazi. Uma ibhenki seliyisebenzisa i-A_I ukuthi linqume ukuthi likunikeza isikweletu, noma uma inkampani isebenzisa i-A_I uku-shortlist-a izicelo zomsebenzi, i-A_I iyakwazi ukwenza amaphutha okungaze kwenzeke ngama-random. Zingahlehlela kahle. Basengaphazamisa abantu ngendlela engafanele ngokobuhlanga, ubulili, noma la ahlala khona, hhayi ngoba umuntu ukhethe ukucwasisa, kodwa ngoba i-A_I ifike yafunda kwi-data yomlando ebese iqukethe lezo ukungalingani. Lokhu kuthiwa i-algorithmic bias. Ucwango lwango-twenty four lwaveza ukuthi ama-A_I system aqeqeshwe kwi-data engamele abantu base-Africa banamathuba kakhulukazi okukhiqiza imiphumela enobandlululo. E-South Africa, la ama-A_I amathuluzi amaningi akhiwa khona kusetshenziswa i-data from i-Global North, um ingozi enkulu le. Ama-facial recognition system aqala ukuhlukanisa ubuso obumnyama ngama-mistake rate up ukuya ku-thirty five percent ngaphezulu kunezinto ezikhanyayo. Unelungelo lokwazi. Ngaphansi kwe-P_O_B_I, noma iyiphi inhlangotho esebenzisa i-personal data yakho kwi i-decision e-automatised kumele ikwazi ukukutshela ukuthi leyo decision yenziwe kanjani. Uma behluleka ukuyichaza kungenzeka ukuthi baphula umthetho.

isiXhosa (xho_Latn): I-A_I system ingenza okungalingile, and ama-South Africa amaningi awayazi. Ibhenki uma isebenzisa i-A_I ukuthi ithathe isinqumo sokukwenzela isikweletu, okanye ukuba inkampani isebenzisa i-A_I uku-shortlist-a izicelo zomsebenzi, i-A_I iyakwazi ukwenza amaphutha okungaqondanga. Zingahlehlela. Basenokuxhoba abantu ngokungekho sikweni ngokobuhlanga babo, ngesini sabo, okanye la ahlala khona, kungabi ngenxa yokuthi umuntu ukhethe ukucalucalula, kodwa ngenxa yokuba i-A_I ifike yafunda kwiidatha zomlando ebezikade ziqukethe lezo ukungalingani. Oku kuthiwa i-algorithmic bias. Uphando olwenziwe ngo-twenty twenty-four lwaveza ukuba ama-A_I system aqeqeshwe kwi-data engamele ama-African people ngokukodwa kunokubakhona ukuzala iziphumela ezi-biased. eMzansi Africa, indawo ekwakhiwa khona ama-tool amaningi e-A_I kusetshenziswa i-data evela kwi-Global North, um ingozi enkulu. I-facial ID ama-facial recognition system aqala ukuva ubuso obumnyama ngama-rate e-wrong up ukuya ku-thirty five percent ngaphezulu kunezinto ezikhanyayo. Unelungelo lokwazi. Phantsi kwe-P_O_P_A, nayiphina inhlangotho esebenzisa i-personal data yakho kwi i-decision

e-automatised kumele ikwazi ukukutshela ukuthi kwenziwa njani leyo decision. Ukuba abakwazi ukuyichaza, banokuphula umthetho.

Sesotho (sot_Latn): Sistimi sa A_I se ka hleka-hlaka, 'me ma-South Africa a maholo ha a e-know. Ha bankeng e sebedisa A_I ho decide whether to give you a loan, or when a company uses A_I to shortlist job applications, I_I e khona ho phosa liphoso eseng ka bo-random. Li ka hlopheka. Ba ka sotla batho ka ho sa lokang ka baka la morabe wa bona, Bong ba bona, or where they live, not because someone chose to discriminate, but because i-A_I ile ya learn from the history data that already contained those inequalities. Sona sebizwa ka-algorithmic bias. Phuputsong ka-twenty twenty-four ho ile ha sibolwa ukuthi ama-A_I system a koetliswe ku-data e sa emelang batho ba Africa ho ka 'na ha e-ba le phello e nang leeme. Naheng ya South Africa, la ho ahwang disebediswa eziningi ze-A_I kusebenzisana ne-data from i-Global North, ho na le likoloto. Li-facial recognition system ha li sa lemosa batho ba nang le skin e ntso ka error rates up ho isa ho liperesente tse ka bang thirty-five higher than face e nang le skin e khanyang. U na le right to know. Tlas'a P_O_P_A, ho na le mekhatlo e meng e sebedisang li-personal data za wena in a li-decision se ikentshise ka boomo se be se khona ho u joyela how that decision was taken. Ha ba sa khona ho e-explain-a, ba ka 'na broke the law.

Northern Sotho (nso_Latn): Dikanegelo za I_A di ka se be lekalelekane, le ma-South Africa a maholo ga a o rate. Ge e ba bankeng go thongwa ka A_I ya go-decide-a whether to give you a loan, or when a company uses I_A to shortlist job applications, I_I e ka phoxa e sa i-random. A ka ba a hlophekile. Ba ka sotla batho ka go se lekalekanye go ya ka morafe wa bona, Bong, okanye la ahlala khona, hhayi ngoba umuntu akhethe ukucwasisa, but because i-A_I ile ya learn from the history data that already contained those inequality. Se sebizwa ka-algorithmic bias. Nyakilala ya ka-twenty four e ile ya utolla ukuthi ama-A_I system a-train-wa ku-data se sa emelele batho ba ma-Africa go ka noba le kutlwisiso e fetogileng ya phelelo. Ka Aforika-Dike, la go agiwago didiriswa ka kakati za I_A kusetshenziswa i-data from le-Global North, go na le sekoloto se. Dikanegelo za go-recogniza-a ma-facial do-mistidentify-a o-faces a-skin-e omnyama at error rates up go ya ku-thirty five percent a godile go na le bo-skinny ba maso. O na leelungelo lokwazi. Ka-under P_O_P_A, mekgatlo le ge e ka ba mang le mang yo a sebedisayo i-personal data yakho kwi se-otomatiki sa decision kumele sikwazi ukukutshela ukuthi leso decision sithathwe kanjani. Ge e ba ba sa go hlalolele, ba ka senywa ke umthetho.

Setswana (tsn_Latn): A_I system e ka se nne ya lekalelekane, le ma-South Africa a le mang ga ba o itsile. Fa banka e sebedisa I_A go decide whether to give you a loan, or when a company uses A_I to shortlist job applications, A_I e ka iphosa e seng ka go ijwa. A ka nnela gohlelwa. Ba ka sotla batho ka go sa ikanyege go ya ka lotso, Bong, or kwa ba leng gone, not because someone chose to discriminate, but because i-A_I yo o ithutileng from the history data e already contained those inequalities. Se sejwaelwa go ba-algorithmic bias. Patlelo ya ka-twenty four e fitlhile fa mekgahlelo ya A_I ba thapisiwe ku-data se sa emele batho ba Africa go kaenzeka ka kakhulukhu ukuthi sivelise ama-output a nang neelo. Kwa South Africa, kwa go agiwang didiriswa di le dintsinzana ka A_I using data from i-Global North, go na le lobelo lo logolo. Dikanegelo za go lemogiwa ga matshwao a maso di sa lemogeng batho ba le batsho ka sekhahla se se kwa

godingwana go ya ka phesenti di le thantshana ka bo-thirty-five go na le batho ba ba nang le skin e e phatsima. O na leelungelo lokwazi. Ka fa go P_O_P_A, mekgatlho le fa e sebedisa personal data ya lona in a se-automatised decision kumele sikwazi ukukutshela ukuthi leso sinqumo sithathwe kanjani. Fa ba sa ka ke ba e explanate, they might be breaking the law.

Tsonga (tso_Latn): Sisiteme swa I_I swi nga ha va swi nga fambi kahle, ni ma-South Africa vo tala a va yi ti. loko bank yi use I_A to decide whether to give you a loan, or loko khampani yi use I_I to shortlist job applications, I-A_I yi nga ha swi kota ku hoxisa, a hi mhlasimbe. Ma nga hlongwa kahle. Va nga ha va vaxopaxopile hi ndlela yo ka yi nga n'wi fanelanga abantu ku ya hi mpango wabo, i-gender yabo, or la ahlala khona, hhayi ngoba umuntu ukhethe uku-discriminate, kambe ngoba i-A_I yifunde kwi-historical data levele iqukethe lezo ukungalingani. Loku ku thyiwa i-algorithmic bias. Nkambisiso wa le ka-twenty twenty-four wu kume ukuthi i-A_I system a yi leteriwe ku ya hi data le handle ko yimela ma-African people ku nga ha va ni mbuyelo yo ka yi nga yimi kahle. La e-South Africa, la switirho swa I_A swi akeke kona hi ku use data from le-Global North, ku ni mhangu yo chavisa. Ma-facial recognition system a ma va lemuki kahle ubuso obumnyama at error rates up ku ya fika eka thirty-five percent higher than ama-faces okhala ne skins ezikhanyayo. U nielungelo lokwazi. Ehansi ka P_O_P_A, i nhlango yin'wana ni yin'wana bhuti yi nga use-a um i-personal data yakho kwi i-decision ya othomathike kumele ikwazi ukukutshela ukuthi leyo decision yenzeke kanjani. loko va nga swi kota ku yichaza, va nga ha va va phhula nawu.

Swati (ssw_Latn): I-A_I system ingabi nabulungiswa, and ama-South Africa amaningi awayazi. Uma ibhenki lisebenzisa i-A_I ekuthatheni isinqumo sokukwenzela isikweletu, noma uma inkampani isebenzisa i-A_I uku-shortlist-a ama-job application, i-A_I iyakhona ukwenza amaphutha angavumelani. Angakhona ukuba-systematic. Basengaba-disadvantag-a kabi abantu ngesisekelo se-race yabo, i-gender labo, noma la ahlala khona, hhayi ngoba umuntu ukhethe ukucwasisa, kodwa ngoba i-A_I ifike yafunda kwi-data yasemlandweni le ebevele iqukethe lezo ukungalingani. Loku kuthiwa i-algorithmic bias. Uwaningo lwango-twenty twenty-four lwaveza ukuthi ama-A_I system aqeqeshwe nge-data um longamele abantu base-Africa banamathuba kakhulukazi okukhipha imiphumela le-biased. La e-South Africa, la ama-A_I amathuluzi amaningi akhiwa khona kusetshenziswa i-data from i-Global North, um ingozi ibalulekile. Ema-facial recognition system a-misidentify ama-faces ane-skin e-black at error rates up kuze kube u-thirty five percent ngaphezulu kunama-faces ane skins akhanyayo. Unelungelo lokwazi. Ngaphansi kwe-P_O_B_I, noma iyiphi inhlango lesebenzisa imininingwane yakho kumuntu i-decision le-otomatiki kumele ikwazi ukukutshela ukuthi leso sinqumo sithathwe kanjani. Uma behluleka ukuyichaza bangaphula umthetho.

Afrikaans (afr_Latn): A_I-stelsels kan oneerlik word, en die meeste Suid-Afrikaners ken dit nie. As 'n bank met behulp van A_I beslist of om u 'n lening toewys, of as 'n company A_I to shortlist job applications, die A_I kan take wat nie-random is nie. Hulle kan systematies word. Hulle kan onregtelik ten gunste van people op basis van their race, sex, of die area where they live, not because someone chose to discriminate, maar want die A_I-s leer van historiese data wat al die ongehoorhede beheld. Dit word 'n algoritmiese vooroordeel. In 'n studie in 2024 is bevindt dat

A_I-stelsels op data opgerig is wat nie Afrika-bevolkings representeer nie, is dit veral geneig om biase resultate te produceer. In South Africa, waartekort die meeste A_I-tools word opgerig met behulp van data from the Global North, is die risiko beduidend. Weerkeningsstelsels misidentify-a dark skinned face at error rates up om thirty-five percent hoer as die helder skimmedgesig. U has the right to know. Onder P_P_I_A, is 'n organisasie wat u personal data in 'n outomatiese decision-makers must be able to tell you how that decision was made. As they cannot explain it, they may be breaking the law.

C.2 Brief 2 (isiZulu NLP): isiZulu complete, isiXhosa partial, remaining languages not returned

isiZulu (zul_Latn): I-A_I vele ikhuluma isiZulu. And ungakwazi ukuyisebenzisa njengamanje. Abacwaningi baseNingizimu Africa sebecithe iminyaka befundisa i-A_I ukuthi bazozwisisa, ufunde, ubhale isiZulu nezinye izilimi zaseNingizimu Africa. Umsebenzi uqale ezikhungweni ezifana ne-C_S_I_R naseNyuvesi of Pretoria and manje usukhiphile amathuluzi lawo wonke umuntu angakwazi uku-access-a mahhala. Lokhu kubalulekile ngoba ama-tool amaningi e-A_I akhiwe ngesingisi, abalimi lesiNgisi. Uma uqhuba ibhizinisi, ukhonza umphakathi, noma ufuna ukufinyelela kubantu abacabangayo nabanomuzwa obala kakhulu ngesiZulu, Sesotho, noma iSetswana akusadingeki ukuthi ukhulumisana ulimi olungesiwo wenu. Amathuluzi afana ne-Lelapa A_I's Vulavula ARP akuvumela ukuthi u-translate, transcribe, and ukudala oku-content ngezilimi zaseNingizimu Africa. Abacwaningi bakha amathuluzi. Manje sekuwukuthi uwena osebenzisa wona.

isiXhosa (xho_Latn): I-A_I sele ikhuluma isiZulu. And ungakwazi ukuyisebenzisa njengamanje. Abaphandi baseMzansi Africa bachitha iminyaka befundisa i-A_I ukuqonda, ufunde, ubhale isiZulu nezinye izilimi zaseMzansi Africa. Umsebenzi waqala kumaziko afana ne-C_S_I_R naseYunivesithi lasePitoli and manje selikhiphile amathuluzi lawo wonke umuntu akwazi uku-access-a mahhala. Oku kubalulekile kuba uninzi lwezi-A_I zakhiwe ngesingisi, aba-English speaking. Ukuba uqhuba ibhizinisi, ukhonza uluntu, okanye ufuna ukufikelela kubantu abacabangayo nabazizwayo ngokucace bha ngesiZulu, Sesotho, okanye iSetswana akusadingeki ukuncokola kwi ulwimi olungalikho. Ama-tool afana neLelapa A_I's Vulavula A_P allows you to translate, transcribe, and generate content in South African languages. [Final segment not returned: rate limit]

Sesotho, Northern Sotho, Setswana, Tsonga, Swati, Afrikaans (sot_Latn, nso_Latn, tsn_Latn, tso_Latn, ssw_Latn, afr_Latn): Not returned. The free-tier rate limit halted the run after the isiXhosa call. These languages are pending a re-run under a higher rate allowance.