

Dependency Before Language: Ancestral Memory and the Real Question of African AI

Author: Divine Nwobodo Chidera

Abstract

This essay argues that Africa's AI conversation has been organized around a visible problem, language, when the deeper contest is about dependency. More than two thousand languages are spoken across the continent, and the world's most powerful models handle only a fraction of them; there is a huge gap and it deserves the attention it gets. Yet the models being built to close it sit, for the most part, on foundations owned elsewhere: base weights trained abroad, data priced and exported, compute rented in foreign currency. Building on ancestral memory as the framework introduced in this series, I read the language race as the visible front of a structural fight over who owns the means of remembering. Examining Nigeria's N-ATLAS, Lelapa AI's InkubaLM, and the GSMA's pan-African language initiatives, I argue that language work matters most when it is treated as a foothold toward sovereignty rather than a substitute for it. The essay closes with four principles for African AI: Ownership over Access, Community Data over Exported Data, Local Evaluation over Imported Benchmarks, and Dependency Recognition, and with the claim that a tradition which keeps knowledge distributed across the collective is the oldest available architecture for resisting a single center of control.

Keywords: African AI, digital sovereignty, ancestral memory, language models, dependency, data sovereignty, N-ATLAS, InkubaLM, GSMA, Afrofuturism

Introduction

Here's the question that has run through this series, simple and old: whose knowledge counts, and who gets to keep it? In the opening essays I proposed that ancestral memory (the distributed, communal keeping of meaning across generations) can serve as a compass for Black futures, and that this memory operates as a living technology of its own. This essay carries that compass toward the most concentrated site of contemporary knowledge production: artificial intelligence.

AI is now the most powerful memory technology we have ever built. It remembers the world on behalf of everyone, and in remembering, it decides what the world was and can be. For Africa, the problem with this new memory is not that it forgets the continent. It does worse. It remembers Africa poorly: in someone else's language, through someone

else's categories, on someone else's terms. The question I started this series with has therefore left the library and entered the training run.

What the Language Story Hides

By now the evidence is familiar. GSMA, the global mobile industry body, estimates that Africa is home to more than 30% of the world's languages, over 2,000 of them, and that the majority remain effectively absent from mainstream AI systems. The consequence is practical. Voice interfaces that could serve the majority of the continent's population remain out of reach, and the benefits of AI concentrate among the linguistically privileged. That is not a small thing.

This language gap is genuine, and it earns its place as the public face of African AI. It is also the layer where the continent can win. Building a language model for Yoruba, Hausa or Igbo is not the same problem as challenging GPT's scale; it rewards data collection, consent, and community effort rather than billions of dollars in compute. So it is no accident that the flagship African efforts are language efforts: Awarri's N-ATLAS, fine-tuned on Llama-3 8B to serve Yoruba, Hausa, Igbo and Nigerian-accented English; Lelapa AI's InkubaLM, a 0.4-billion-parameter model trained from scratch on five African languages; and the GSMA's G6 coalition of mobile operators, building African language models 'by Africa, for Africa.'

But here is the quiet fact underneath: most of this building sits on someone else's foundation. N-ATLAS, the most famous African LLM, is a fine-tune of weights trained abroad. That's the language story told honestly: a correction layered over another architecture. It is not incompetence; it is economics. Frontier training costs hundreds of millions to billions. The question is what happens when the correction is all the continent builds.

Three Faces of Dependency

Strip away the language story and three dependencies are doing the real work.

Base-model dependency. The continent's best-funded model work is fine-tuning open weights. It's cheap, rational, and fast, and it means the foreign lab controls the chassis while African builders add the trim. When the underlying model changes, the fine-tune changes with it. The griot composes within the tradition she inherits; the question is whether the tradition itself is hers.

Data dependency. The world's models are trained overwhelmingly on Western data, by Western teams, for Western contexts. Africa's knowledge (its history, its law, its informal markets, its storytelling) is either missing from the training runs or exported as raw

material for products whose terms the continent does not set. A community that loses control of its own record loses the thing that makes memory sovereign.

Infrastructure dependency. Compute is rented in hard currency, abroad. Every African team profiting from AI is also, in some ratio, renting back its own constraints from a foreign cloud. The budget that should grow domestic capacity instead flows outward.

Case Study: N-ATLAS, InkubaLM, and the Cooperative Turn

Three recent developments show the range of responses to dependency.

First, N-ATLAS, Nigeria's open-source, government-backed model, is the fine-tune as a digital public good. It is built on Llama-3 8B, localized with hundreds of millions of tokens of Nigerian data, and released on Hugging Face as a shared foundation. Its own maintainers frame it as reducing dependence on foreign black-box APIs. It's the bridge position, and it's defensible: a country that cannot yet train a frontier model can still own the layer of localization, the speech stack, the community of builders surrounding it.

Second, InkubaLM represents the from-scratch position at small scale. Lelapa AI trained a 0.4-billion-parameter model from scratch on 2.4 billion tokens across five African languages, and published the training datasets alongside it. It does not rival large models in scale; it rivals them in what it proves: that African teams can own the entire stack, from corpus to trained weights, when they design for constraint rather than against it.

Third, the cooperative turn. The G6 operators, joined by Awarri, Lelapa AI and others, moved the contest from national projects to a continental architecture; ATLAS Umoja carried that momentum into a platform backed by five governments and focused on sharing datasets, benchmarks and talent across countries. This is call-and-response at continental scale: many voices building on shared ground rather than a single broadcaster. It matters as much for what it prevents, a regional gatekeeper emerging in place of a foreign one, as for what it builds.

Read fairly, all three are legitimate. Fine-tuning is not naivety; it is the correct move at today's prices. The danger is not the bridge but permanence on it. If African builders fine-tune open weights forever, the continent becomes the world's most skilled adapter of other people's foundations, fluent in Yoruba, and still tenants.

Four Principles for African AI

If the essays in this series have proposed principles for Afrofuturist design, the same discipline can serve the politics of African AI. I want to propose four.

Ownership over Access: Access is the language of tenancy: granted, revocable, priced in dollars. Ownership is the language of a people and their record. African AI should be measured not by how many APIs the continent can call, but by what it holds: models it can inspect, datasets it controls, infrastructure it participates in building.

Community Data over Exported Data: Data collected from the continent should benefit the community that produced it, on terms that community sets. This is the memory-keeping contract made technical: provenance, consent, and the right to say how one's own story is used.

Local Evaluation over Imported Benchmarks: Whoever defines the test shapes the field. Western benchmarks measure a world the continent does not fully inhabit. Until Africa builds its own evaluations of its languages, its contexts, and its informal economies, no one can prove a model works here, and worse, no one has to.

Dependency Recognition: Name the arrangement before designing around it. A fine-tune is a dependency with a roadmap, or a dependency with a lifetime lease. The difference is whether the builder says so. Recognition is not pessimism; it is the condition for a bridge that actually gets crossed.

What the Fight That Matters Looks Like

This is what the fight that matters looks like: owned datasets, collected with consent and held in community; benchmarks built by African researchers for African realities; compute partnerships that treat the continent as a partner rather than a market; shared infrastructure of the kind the G6 and ATLAS Umoja gesture toward; and edge and efficient AI that runs on the low-cost devices most of the continent actually holds. The path Fastagger has begun walking with mobile operators. None of these require rivaling OpenAI's compute budget today. They require a different ambition: to own the layers dependency would leave to others, at whatever scale is real.

Research Directions and Open Questions

Several questions follow. Who builds and maintains African evaluation benchmarks, and what counts as legitimacy for them? What does sovereignty mean at 0.4 billion parameters, for a model too small to be powerful but built entirely in the community's hands? What ownership models can honor data contributors as co-authors rather than sources? And how should diaspora and continental teams divide the labor of keeping African AI African? These are research questions, but they are also design questions, which is precisely why they belong in a series about building.

A caution also belongs here. Not every gap is worth filling. A gap matters only if someone will use and pay for the work, and if the work can realistically succeed. Language and evaluation have that property; some regulated, high-stakes fields do not, for a builder without deep resources. The discipline of a small team is choosing the fight it can actually win.

Conclusion

Stories are technologies, and memory is the medium both of dependence and of freedom. A single center that remembers the world for everyone recreates, in silicon, the colonial archive: the record kept elsewhere, in the keeper's language, on the keeper's terms. The traditions this series has drawn on insist that knowledge can be held otherwise: distributed across a collective, verified by participation, revised by retelling, never owned by a single authority. That is not a metaphor. It is an architecture, and it is older than the training run.

Language is the foothold; dependency is the climb. African AI has made a reasonable beginning by building models that speak Yoruba, Hausa and Igbo. The task that remains is to make sure that when those models speak, the voice that answers for the continent is the continent's own, and that the foundation underneath it is something a people can eventually call home, not a ship they are renting.

The griot's art survived because it served the community's need to remember together. African AI will matter insofar as it does the same. We are not fighting over the alphabet. We are fighting over who owns the library, and every principle above is a shelf in it.

Select References

Awarri Technologies. N-ATLaS-LLM: A Multilingual African Language Model. Hugging Face, 2025.

Federal Ministry of Communications, Innovation and Digital Economy (Nigeria). "Nigeria Launches Landmark AI Model Powered by Awarri." 2025.

Lelapa AI, et al. "InkubaLM: A Small Language Model for Low-Resource African Languages." arXiv:2408.17024, 2024.

GSMA. "The African AI Languages Model Project." GSMA Foundry, 2025.

GSMA / ATLAS Umoja AI. Launch announcement, Abuja, July 2026.

MWC Kigali. "GSMA, Africa's Leading Mobile Operators and the AI Ecosystem Unite." 2025.

Nsabimana, S. "Oracle of the Baobab and Ancestral Memory." (prologue essay, series)

Vansina, J. Oral Tradition as History. University of Wisconsin Press, 1985.

Chidera, D. N. "Envisioning Black Futures through Ancestral Memory." Nubian Research, 2026.

Chidera, D. N. "Narrative Translation as Design Method." Nubian Research, 2026.