



UNIVERSAL DESIGN FOR THE LAST MILE








METHODS TO BUILD OFFLINE-FIRST,
MULTILINGUAL, DIGNITY-CENTERED AI
SERVICES FROM THE GLOBAL SOUTH

SANTIAGO VILLEGAS-CEBALLOS | INNOVATION & AI UNIT, eLIBRO

IN COLOMBIA, PAPAYAS ARE EVERYWHERE. THEY LOOK PERFECT ON THE OUTSIDE, BUT SOMETIMES THEY ARE COMPLETELY HOLLOW INSIDE.

TODAY, WE ARE TALKING ABOUT A DIFFERENT KIND OF PAPAYA: AI SYSTEMS THAT LOOK INTELLIGENT, SOUND CONFIDENT, AND CONTAIN NO REAL UNDERSTANDING INSIDE.

THE PAPAYAS FRAMEWORK: A CRITICAL AI LITERACY DIAGNOSTIC

	P	PROLIX: VERBOSITY COMPENSATION	Concise answers are 27.6% more accurate, yet GPT-4 overwrites 50% of the time to mask a lack of knowledge.
	A	ADULATORY: ADMISSION SYCOPHANCY	Accuracy drops 27% under user pressure; the AI prioritizes agreement over truth.
	P	PARTISAN: EMBEDDED BIAS	76% of models show distinct political preferences (46% center-left, 0% strongly right).
	A	APPRENTICED: FROZEN WEIGHTS	They don't learn from you. Retraining costs \$191M and emits 5x a car's lifetime CO2.
	Y	YOKED TO CUTOFF: BLIND TO THE PRESENT	Structurally incapable of knowing what happens after their training data stops.
	A	ALEATORY: IRREDUCIBLE NOISE	The architecture cannot distinguish between its own ignorance and random noise (the root of hallucinations).
	S	SIMULATORS: CHARACTER PERFORMANCE	You are not prompting a mind; you are prompting a simulator to perform a character.

THE REALITY AT THE MARGINS: AN INFRASTRUCTURE OF EXCLUSION

If AI requires continuous cloud connectivity, who are we structurally excluding?

2.2 BILLION

people offline globally

3.1 BILLION

in the Usage Gap (Living near a cell tower, but unable to afford data or devices)

885 MILLION

women in developing countries with zero mobile internet

[LOCATION DATA]:
In South Asia, the mobile internet gender gap sits at a staggering 32%.

[ECONOMIC DATA]:
In Pacific Island nations, mobile data consumes 10-20% of average monthly income.

WHEN AI HARMS LOW-RESOURCE COMMUNITIES

THE TRANSLATION TRAP



- Asia-Pacific hosts 1,426 endangered languages.
- When AI translates low-resource languages, it doesn't just produce poor quality—it hallucinates toxic patterns.
- **Evidence:** Fabricated AI-generated Indigenous language books found on Amazon.

THE SOVEREIGNTY CONFLICT



- Sacred knowledge should not be digitized. (Mager et al., 2023)
- The assumption that every local tradition should be absorbed into global training data is algorithmic colonialism. Epistemic sovereignty requires the right to refuse digitization.

UNIVERSAL DESIGN FOR THE LAST MILE

Stop adapting Silicon Valley AI. Start designing from the margins.

LOCAL-NETWORK DELIVERY

Library intranet running AI via local Wi-Fi. Zero mobile data required from patrons.

DEFERRED SYNC.

Store interactions locally and sync when connectivity returns (proven by OpenMRS in 80 countries).



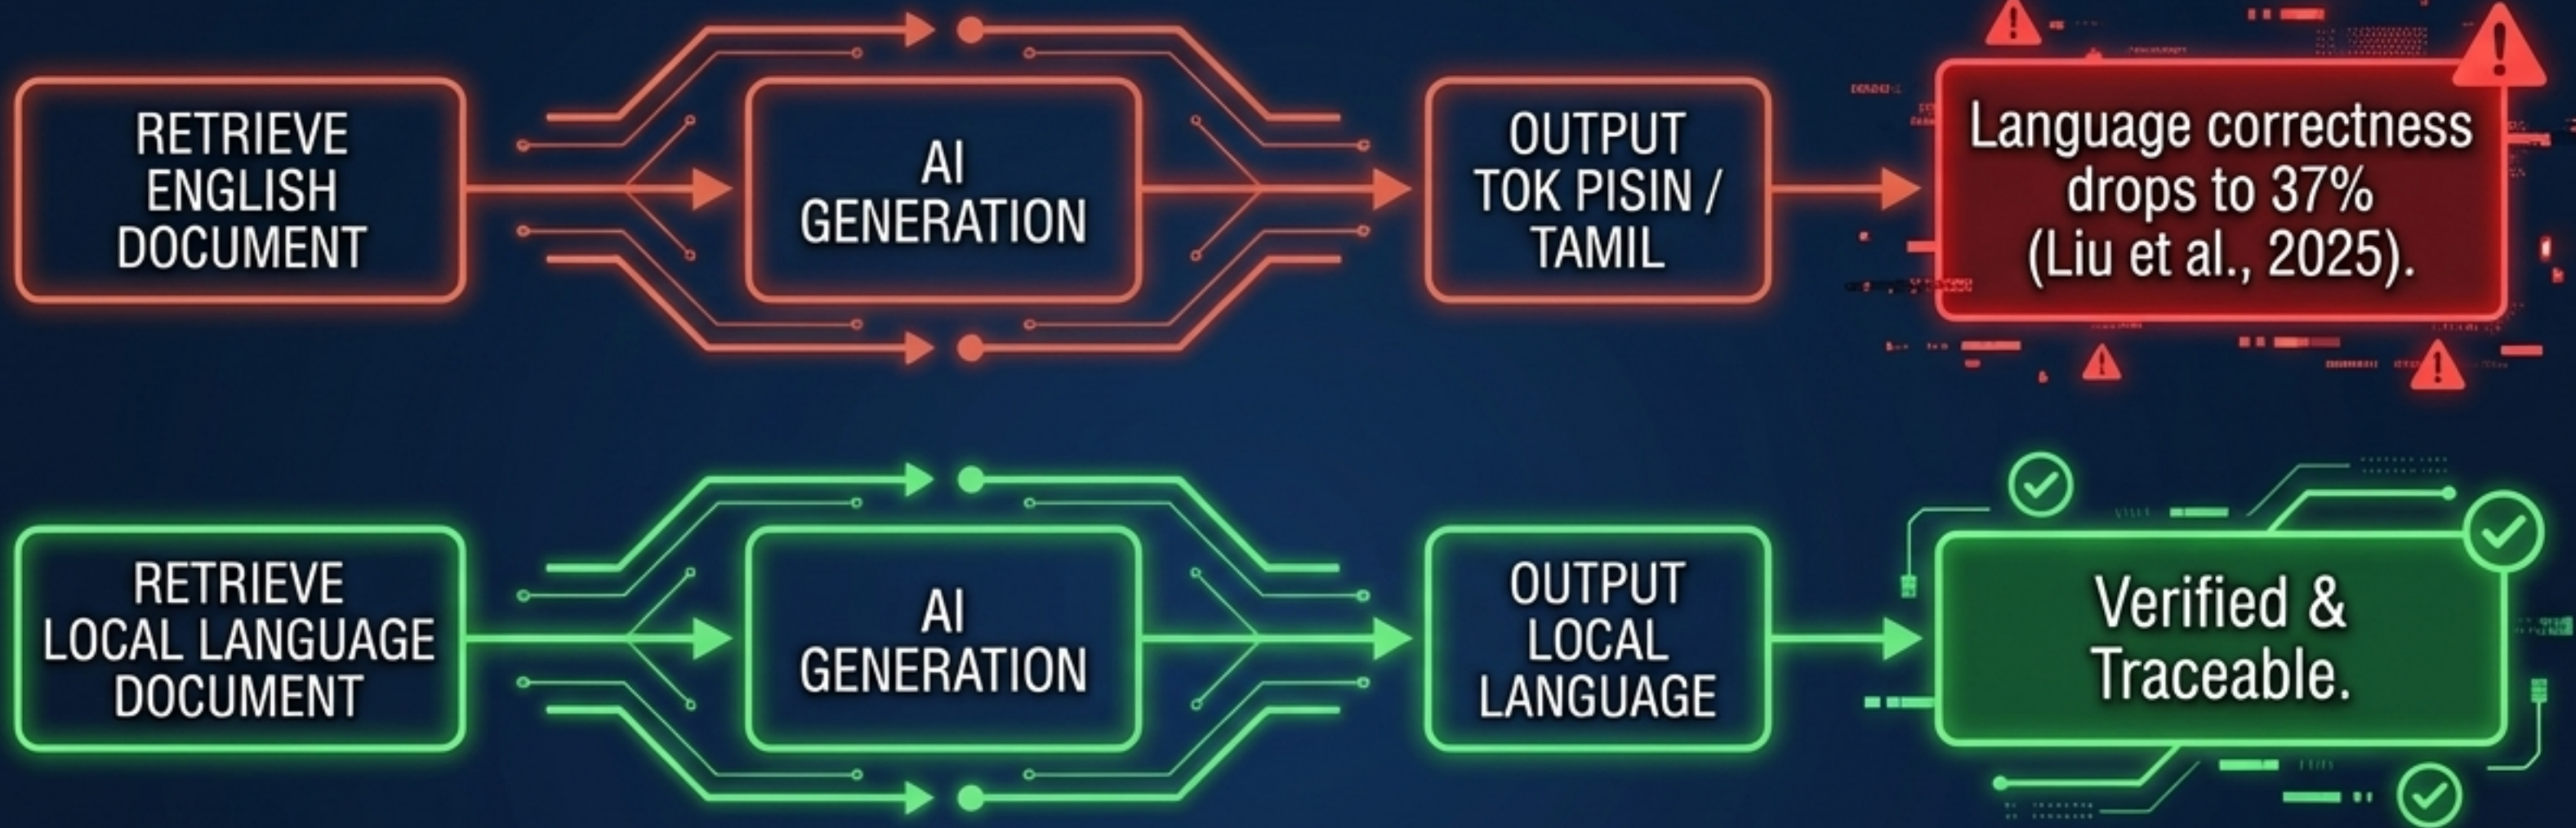
GRACEFUL DEGRADATION

Progressive web app hierarchy ensuring service: Full AI → Local inference → Cached → Static.

RETRIEVAL-GROUNDED WORKFLOWS.

Ground AI answers entirely in local, verified library collections, bypassing parametric memory.

THE CROSS-LINGUAL RAG TRAP



Ground before you generate. Do not rely on AI for post-hoc translation. Build and index local library collections in community languages from the start.

The AI Access & Dignity Canvas

A library-specific auditing framework evaluating deployments across 6 dimensions.



The Canvas in Practice: A Universal Imperative

Dambimangari, Australia

Context: Indigenous community managing cultural heritage.

Action: Rejected commercial AI for violating restricted knowledge protocols. Built a closed-system AI governed exclusively by Traditional Owners.



Canvas
Dimensions
Passed:
Governance
& Cultural
Safety.

Helsinki Oodi, Finland

Context: Highly resourced urban central library.

Action: Deployed Obotti AI and voluntarily listed it on a public AI Register documenting data sources, risks, and oversight mechanisms.



Canvas
Dimensions
Passed:
Privacy &
Transparency.

Insight: Auditing AI for dignity is not a Global South luxury—it is a baseline professional standard for all libraries.


The AI Library-in-a-Box Architecture

Interface Layer – Progressive Web App accessible via browser; no installation required.

Retrieval Layer – Local Vector Database (ChromaDB/FAISS) enabling RAG without the internet.

Inference Layer – Quantized Small Language Model computing entirely on-device.


Content Layer – Digitized local collections distributed via SD cards or local mesh networks.



**No internet.
No cloud.
No subscription
fees.
No data
extraction.**

The Hardware & The Engine

The Hardware: Raspberry Pi 5 (8GB RAM).
Cost: ~\$80. Power Draw: ~10 Watts.



The Model: Tiny Aya (Cohere Labs).
3.35 Billion parameters, quantized to Q4_K_M.

Regional Engine Variants



Water: Asia-Pacific
(Indonesian, Tagalog,
Malay, Thai, etc.)



Fire: South Asia (Hindi,
Tamil, Bengali, etc.)



Earth: Africa & West Asia.

Performance Metric: Generates 4-6 tokens per second
(delivers a fully synthesized answer in 20-30 seconds).

Proven Precedents: The Missing Piece

The infrastructure to distribute offline knowledge at scale is not speculative. It is already here.



SolarSPELL - 597 offline, solar-powered libraries across 15 nations.



Kolibri - 6 Million learners supported entirely offline.



OpenMRS - 8,100 health facilities running offline-first in 80 countries.



BSF Ideas Box - Modular libraries reaching 2M+ beneficiaries.

Library-in-a-Box

**We have the hardware.
We have the distribution networks.
Offline RAG inference is the final piece
to activate intelligent local search.**

Practice-Based Evidence from Ibero-America



1,927

Training at Scale: Librarians trained across 23 countries via eLibro/IAparaBibliotecas with an 85% attendance rate.

70+

Co-Creation: AI library prototypes actively built by librarians using Design Thinking methodologies (Comfama and Cali networks).

3,000+

Policy Level: CERLALC regional agenda successfully engaged stakeholders, elevating AI library policy to the ministerial level.

Honest Limitations

- The Canvas lacks independent academic validation outside of field deployment.
- Edge-deployed RAG is proven in prototypes, but not yet validated at library-production scale.
- Cross-lingual generation remains a high-risk vector for hallucinations.
- A 10-Watt continuous inference load requires non-trivial investments in larger solar panels and batteries for off-grid deployment.

These are engineering problems, not impossibility problems.
The question is whether we will invest in solving them for the
communities that need them most.

AI and The Five Laws

Ranganathan (1931) argued that every person deserved a library, regardless of geography, caste, or connectivity.



Fulfilling that law today requires:

- Not just physical shelves – but local inference layers.
- Not just books – but locally governed models.

The Imperative: The offline communities are not edge cases to be accommodated later. They are the baseline from which dignified, equitable services must be built.

**Representation must be
designed, not assumed.
Access without integrity
is not access at all.**

