



Teaching and Learning in Fragile Contexts (TLFC) Policy Brief



Co-Creating Safe, AI-Supported Learning Spaces for Teachers in Crisis Contexts

A Policy Brief on the Generative AI for Accelerated Competency-Based Teacher Training
(GAI-ACT) in Fragile Contexts

Ggaliwango Marvin
Joseph Tibakanya
Nakayiza Hellen
Michael Gallagher
Rovincer Najjuma

Makerere University, National Curriculum Development Centre, Muni University, The University
of Edinburgh



MAKERERE UNIVERSITY



THE UNIVERSITY of EDINBURGH
Moray House School of
Education and Sport



Executive Summary

The Generative AI for Accelerated Competency-Based Teacher Training in Crisis Contexts (GAI-ACT) project, a 17-month Participatory Action Research initiative funded by CODE and TLFC, addresses a critical educational gap in Uganda's refugee-hosting districts. The project was originally meant to teach active Information and Communication Technology (ICT) teachers how to use large foreign language models (LLMs) like ChatGPT to do prompt engineering. However, during its baseline profiling, it found serious systemic weaknesses. Groundtruth data revealed that unregulated AI adoption was already rampant, yet educators lacked the institutional safeguarding and critical auditing skills required to protect traumatised refugee learners from algorithmic bias and Western-centric "epistemic erasure."

In response to these fragile realities and direct stakeholder demands, the GAI4Educ consortium strategically pivoted from a top-down technological deployment to a localised, "government-first" capacity-building ecosystem. The project successfully secured the national teacher pipeline by upskilling university lecturers and pre-service educators at Muni University, addressing the systemic failure of Teacher Training Institutions (TTIs) to produce "AI-ready" graduates. Utilising a highly scalable 1:9 Trainer-of-Trainers (ToT) cascading model, the initiative empowered grassroots educators to transition from passive consumers of foreign technology to active "critical auditors" of algorithmic bias through a methodology termed "Speculative Epistemic Auditing."

To ensure sustainable, low-bandwidth access, the consortium co-created specialised Open Educational Resources (OERs)—ranging from administrative governance to accelerated education programming (AEP) and trauma-informed pedagogy—which are hosted on a sovereign, offline-friendly e-learning platform. To definitively resolve the hallucination crisis inherent in foreign LLMs, the project developed two sovereign, localised generative AI prototypes that were explicitly fine-tuned for Uganda's competency-based curriculum (CBC) and AEP contexts. By advancing Sustainable Development Goal 4 (SDG 4) and Uganda's Vision 2040, this initiative proves that safe, culturally relevant AI integration in the Global South requires participatory co-design, institutional safeguarding, and sovereign digital infrastructure.

Key Findings & Statistics:

1. **Unregulated Baseline Risks:** Prior to the intervention, 63.7% of educators used AI tools, yet 17% possessed a "poor" or "very poor" ability to critically identify algorithmic biases and hallucinations.
2. **Scalable Capacity Building:** The 1:9 ToT cascading ratio was engineered to practically build the capacities of 612 secondary-level teachers across 20 refugee camps and 34 host schools.

3. **Cognitive and Ethical Shifts:** Post-training, over 77% of participating educators felt highly prepared to engage in high-level institutional AI ethics discussions.
4. **Sovereign Decolonial Solutions:** Specialised friendly OERs and two sovereign, localised AI prototypes were successfully co-created to combat digital colonialism and epistemic erasure.

Context

Uganda currently hosts over 1.5 million refugees, making it the largest host country in sub-Saharan Africa. Within these displacement contexts, thousands of young people and adults have had their education severely interrupted by conflict. To provide a viable pathway back to formal schooling, the Accelerated Education Program (AEP) condenses years of conventional schooling for overage, out-of-school learners aged 16 to 45 years.

However, the educators serving these highly vulnerable populations face immense pedagogical friction. They are tasked with delivering Uganda's newly rolled-out, highly interactive Competency-Based Curriculum (CBC) within a highly compressed AEP timeline, all while navigating drastically overcrowded classrooms, severe language barriers, and the trauma-induced mental health challenges of their learners.



Figure 1. Staff reflecting and other LLM outputs after realising the tools were not as perfect as they expected during the training.



Figure 2. Gathering Structural and Safe Guarding Feedback from AEP implementers (MEAL officers)

Furthermore, baseline profiling of these teachers exposed universal infrastructural precarity, including a widespread lack of reliable electricity, limited personal digital devices, and severely restricted internet bandwidth. These extreme structural deficits, coupled with profound pedagogical isolation and a lack of consistent mentoring, make the traditional top-down deployment of digital tools both ineffective and hazardous. The research participants urgently needed localised, friendly AI technologies and trauma-sensitive pedagogical training to safely manage these complex, crisis-affected educational ecologies.

Methodology

To ensure that the integration of generative AI was both scalable and sustainably rooted in local educational ecosystems, the project bypassed traditional, top-down deployments in favour of a participatory action research (PAR) framework. This methodology actively transitioned educators from being passive consumers of imported technology to becoming co-researchers and critical sovereign designers of their digital teaching environments.

The capacity-building architecture used a robust 1:9 Trainer-of-Trainers (ToT) cascading model. The project targeted four cohorts of information and communication technology (ICT) teachers across 20 refugee camps and 34 host schools to serve as "master trainers.". Prior to any technical training, the consortium deployed the comprehensive GAI Teacher Profiling Tool to capture vital baseline data regarding demographic contexts, infrastructural access, pedagogical practices, and existing AI attitudes.

To empirically evaluate the behaviour of foreign large language models (LLMs) in Ugandan contexts, the project utilised a controlled, within-subject experimental design termed "Speculative Epistemic Auditing." Educators acted as their own controls, comparing AI outputs generated under a "Baseline Condition" (unstructured prompting) against an "Intervention Condition" (culturally nuanced prompt engineering). The resulting outputs were rigorously graded using a standard output scoring rubric on a 5-point scale, measuring factual accuracy, cultural appropriateness, epistemic integrity, and hallucination frequency.

Key Findings

The rigorous baseline profiling and post-intervention evaluations yielded profound empirical insights regarding digital precarity, algorithmic bias, and the efficacy of localised capacity-building in fragile settings.

1. **Prior to any formal training, 63.7% of surveyed educators already used** generative AI tools, overwhelmingly relying on ChatGPT for research and lesson ideation. However, this unregulated adoption presented a severe epistemic risk. Baseline diagnostics revealed that 17% of these educators possessed a "poor" or "very poor" ability to critically identify algorithmic biases or hallucinations within AI-generated content.

"If Ugandan teachers became dependent on foreign, proprietary tech monopolies to deliver their national curriculum, it would lock the education system into a perpetual cycle of subscription fees and epistemic dependency."

2. **The Hallucination Crisis and Epistemic Erasure** During the initial auditing phases, the project discovered that foreign LLMs exhibited profound Western biases and a structural inability to reflect the realities of Ugandan classrooms.

When prompted for localized content aligned with the Accelerated In the Education Programme (AEP), these models produced high rates of "hallucinations"—generating factually incorrect educational material and defaulting to Euro-American cultural examples, a phenomenon identified by the consortium as "epistemic erasure".

3. **The Universality of Infrastructural Precarity** baseline data exposed severe digital equity barriers across all education levels. Educators possessing personal digital devices exhibited substantially higher median ICT confidence than those without.




Figure 3. Stacked bar chart illustrating Limiting Factors in ICT/AI Use by Education Level

4. **The Teacher Training Institution (TTI) Capacity Gap** The data validated urgent stakeholder warnings that national Teacher Training Institutions (TTIs) were failing to produce "AI-ready" graduates. Formal academic attainment did not guarantee technological literacy, indicating that deploying AI solely to in-service teachers would be a reactive solution. This finding justified a major strategic pivot to secure the teacher pipeline upstream by first training university lecturers and pre-service educators at Muni University.
5. **Intervention Efficacy:** Mastery of Critical Epistemic Auditing: Following the ToT capacity-building programme, chi-squared tests for independence revealed statistically significant cognitive enhancements ($p < 0.05$).




Figure 4. Post-training, there is a clear improvement in the ability to critically assess AI-generated results, with a shift from "Average" to "Good" and "Excellent" ratings, indicating enhanced critical thinking skills. The pre-service and in-service teachers' newfound mastery of speculative epistemic auditing most critically observes this statistical shift. Baseline data indicates a dangerous epistemic vulnerability, with a considerable portion of these educators rating their ability to identify algorithmic biases in AI-generated content as "poor" or "very poor."

- The Necessity of "Governance-First" Institutional Safeguarding** The project found that introducing AI to teachers without the informed consent of administrators created unsanctioned, high-risk environments. By implementing a "governance-first" approach through the specialized GAI-ADM-101 course, administrators were trained to establish safe regulatory frameworks. *Deploying Generative AI to teachers without the informed consent and regulatory backing of their administrators would create unsanctioned, high-risk technological environments."*




Figure 5. The analytical results show that while there was a significant level of concern about ethical issues related to generative AI before training, the post-training understanding of ethical implications has largely shifted toward "good" and "excellent". This suggests the training was effective in not only addressing pre-existing concerns but also in deepening teachers' knowledge and awareness of AI ethics

This approach yielded massive dividends: post-training surveys confirmed that over 77% of participating educators felt highly prepared to engage in high-level institutional AI ethics discussions, supported by clear data privacy and trauma-informed pedagogical boundaries.

7. **Evolving from AI Consumers to Sovereign AI Producers** Faced with the profound risks of digital colonialism and epistemic erasure, the consortium concluded that true educational sustainability required localised digital infrastructure. Leveraging anonymised, locally curated datasets, the project co-created two distinct, sovereign generative AI prototypes. The CBC AI Support Prototype was fine-tuned explicitly on the Ugandan Competency-Based Curriculum, while the AEP AI Support Prototype was engineered specifically for curriculum compression and multilingual bridging.

"True educational innovation and sustainability require EdTech stakeholders to invest in local data infrastructure, participatory co-design, and sovereign AI model development, ensuring that the technology serves the educators, rather than the educators serving it."



Figure 6. Hands on training for Stakeholders to build their AI capacity to inform our Research and AI development approaches for Fragile contexts



Figure 7. Co-designing a Realistic Safe Institutionally Acceptable AI Approach with well-informed stakeholders

Key Recommendations

To sustainably scale Generative AI in fragile educational contexts: Policymakers and EdTech practitioners must do the following:

1. **Adopt a "Governance-First" Strategy:** Prioritise institutional capacity building and safeguard training for educational administrators before deploying any AI tools directly into classrooms, ensuring technology is legally regulated and trauma-informed.
2. **Invest in Sovereign, localised AI Infrastructure:** Fund the co-creation of offline-friendly, open-source AI models fine-tuned to national syllabi (like the Competency-Based Curriculum) to permanently prevent digital colonialism, algorithmic bias, and epistemic erasure.
3. **Institutionalise "AI-Readiness" Upstream:** Integrate foundational AI literacy and speculative epistemic auditing directly into national Teacher Training Institutions (TTIs) to ensure all graduating pre-service teachers are critically prepared.

4. **Expand the 1:9 ToT Cascading Model:** scale localised capacity by empowering peer-mentored "Master Trainers" to reach isolated in-service educators utilising low-bandwidth, portable Open Educational Resources (OERs).
5. **Champion Open-Source Public Goods:** Ensure that AI prototypes developed for marginalised environments remain open-source and free from commercial paywalls, protecting the data sovereignty of vulnerable learners.

Acknowledgements

The GAI4Educ consortium extends its profound gratitude to CODE, the Technology and Literacy for Children (TLFC) initiative, and the Government of Canada for their vital funding and strategic support. We sincerely thank our institutional partners: Makerere University, Muni University, the University of Edinburgh, the National Curriculum Development Centre (NCDC), and War Child Canada. Most importantly, we honour the tireless work of the teachers in fragile contexts, the district educational administrators, university lecturers, and frontline ICT educators who actively co-designed these solutions to secure a safer, more equitable digital future for Uganda's refugee learners.

The views expressed in this work are those of the creators and do not necessarily represent those of CODE, Global Affairs Canada (GAC) or the Government of Canada. This work was carried out with the aid of a grant from Global Affairs Canada.

References

- Joseph, T., Maria, N., Moses, T., & Marvin, G. (2026). Beyond hype: Will Ugandan educators continue using large language models for learning content creation? IEEE Access: Practical Innovations, Open Solutions, PP(99), 1–1. <https://doi.org/10.1109/access.2026.3675523>
- Joseph, T., Hellen, N., & Marvin, G. (2025). Retrieval-augmented LLMs for culturally sensitive learning content in refugee education. *Lecture Notes in Networks and Systems* (pp. 425–442). Springer Nature Singapore.
- Hellen, N., Marvin, G., Balikuddembe, J. K., & Tulinayo, F. P. (2024). Strategic utilization of ChatGPT in teaching and learning. *Lecture Notes in Networks and Systems* (pp. 129–139). Springer Nature Singapore.
- Marvin, G., Hellen, N., Jjingo, D., & Nakatumba-Nabende, J. (2024). Prompt engineering in large language models. *Algorithms for Intelligent Systems* (pp. 387–402). Springer Nature Singapore.

Download the full research report at: <https://code.ngo/tlfcresearch/>.



Figure 8. Designing Courses for the AEP Teachers



Figure 9. Peer feedback sessions among course



Figure 10. Building the capacity of Staff at Teacher training Institutions (Muni University)



Figure 11. Engagement Teacher trainers into critical discussions about Generative AI for Teacher Training in Fragile contexts



Figure 12. Lack of reliability/accuracy and ethical concerns are the most frequently cited challenges when using generative AI tools, indicating critical areas that training programs should address.

© 2026 Ggaliwango Marvin.

This report was produced under the Teaching and Learning in Fragile Contexts project and published by CODE with permission from the grantee. This report is licensed under the Creative Commons Attribution-Noncommercial 4.0 International Licence (CC BY-NC 4.0). Users may share and adapt the report for non-commercial purposes, provided appropriate credit is given to the author(s), and CODE.