



Technical Report

Generative AI for Accelerated Competency-Based Teacher Training (GAI-ACT)

Authors

Ggaliwango Marvin, Joseph Tibakanya, Nakayiza Hellen, Rovincer Najjuma, Michael Gallagher

Institutions

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

Project Overview

The Generative AI for Accelerated Competency-Based Teacher Training in Crisis Contexts (GAI-ACT) initiative successfully concluded its 17-month implementation cycle, fundamentally transforming the integration of educational technology within Uganda's fragile and crisis-affected environments. Initially conceptualised as a straightforward technology deployment and teacher-training initiative, the approach underwent a profound, data-driven evolution guided by a rigorous Participatory Action Research (PAR) framework. By directly addressing the intense pedagogical friction at the intersection of Uganda's newly rolled-out Competency-Based Curriculum (CBC) and the highly compressed Accelerated Education Programme (AEP), the consortium shifted from merely importing foreign tools to establishing a deeply decolonial, system-wide capacity-building ecosystem. This project overview maps the original strategic objectives against our final, vastly expanded deliverables, demonstrating to stakeholders how agile project management has resulted in the exceptional optimisation of grant resources to secure sustainable digital sovereignty for marginalised educators.

Planned Project Activities

We engineered the original scope of this work around a direct, technology-driven intervention specifically targeting active information and communication technology (ICT) educators. The planned activities focused on building the technological and pedagogical capacities of four distinct cohorts of ICT teachers operating across 20 refugee camps and 34 refugee-hosting schools within three regions of Uganda. The core technological objective was to train these frontline educators in prompt engineering by utilising existing, globally dominant large language models (LLMs), such as ChatGPT and Gemini. To achieve rapid and cost-effective scale, the initiative proposed a highly structured Trainer-of-Trainers (ToT) cascading mechanism operating on a strict 1:9 trainer-to-trainee ratio. Through this architectural model, the primary beneficiaries were intended to act as "master trainers" who would mentor their peers, ultimately aiming to practically build the capacities of 612 secondary-level teachers across 9 to 11 subject areas. Under this initial blueprint, the planned deliverables to support the ToT scale-up were limited to a single accelerated competency-based training curriculum, standard prompt engineering manuals, and a corpus of prompts.

Completed Project Activities

The completed activities vastly exceeded the original scope in direct response to ground-truth diagnostic data and urgent stakeholder feedback, successfully transforming the intervention into a comprehensive digital infrastructure initiative. Prior to executing the capacity-building phase, the consortium developed and deployed the comprehensive GAI Teacher Profiling Tool to rigorously capture baseline data regarding the demographic contexts, infrastructural digital readiness, pedagogical practices, and existing AI attitudes of the educators. Recognising that a singular training manual was entirely insufficient to address both institutional safeguarding and the complex demands of fragile classrooms, the initiative co-created highly specialised open educational resources (OERs), encompassing higher education, administrative governance, and trauma-informed pedagogy. To ensure equitable distribution, we launched a custom, offline-friendly e-learning platform deliberately engineered for low-bandwidth environments, allowing teachers to download these dynamic OERs via

portable media. The idea evolved from simply responsible AI usage to actually co-creating two sovereign, localised generative AI prototypes after the team discovered that foreign LLMs exhibited severe epistemic erasure and hallucinations. We successfully developed and benchmarked the CBC AI Support Prototype—which has since been officially handed over to the National Curriculum Development Centre (NCDC)—as well as a dedicated AEP AI Support Prototype tailored specifically for curriculum compression and multilingual bridging.

Deviations from Planned Activities (The Strategic Pivots)

The original idea called for a linear deployment of existing generative AI tools directly to active ICT teachers. However, our Participatory Action Research (PAR) framework quickly showed that this top-down approach would miss important systemic problems in the Global South. Recognising that rigid adherence to the initial plan would ultimately fail to deliver sustainability, the consortium employed agile, data-driven project management to execute three major operational deviations. These strategic pivots transformed the intervention from a straightforward teacher training exercise into a highly contextualised, decolonised digital infrastructure initiative, ensuring that our final deliverables directly addressed the realities of Uganda's fragile educational ecosystems.

Deviation 1 (Target Audience): The first major deviation involved a profound shift in our target audience, necessitating a pause in direct classroom deployments to prioritise institutional governance. During initial focus group discussions, high-level stakeholders—including district education officers and policymakers—rightfully viewed generative AI as an abstract "buzzword" and withheld their support, citing severe ethical and regulatory risks regarding vulnerable refugee populations. Instead of forcing the technology into schools, we adopted a "governance-first" approach, pivoting to train these administrators and policymakers first. By developing and delivering specialised governance tracks, such as the GAI-ADM-101 course, we built their capacity for data privacy, ethical boundaries, and digital security, successfully transforming their apprehensions into informed regulatory support before any teacher training commenced.

Deviation 2 (Pipeline Focus): Our second deviation expanded the idea's pipeline focus upstream toward Teacher Training Institutions (TTIs), addressing a critical gap identified by our stakeholders. The consortium was warned that training only active, in-service teachers was merely a reactive "band-aid", as national universities were currently failing to produce "AI-ready" graduates, a reality that would inevitably widen the digital divide. Consequently, we deviated from our exclusive focus on refugee camps to strategically partner with Muni University in Northern Uganda. When university lecturers initially exhibited resistance due to their lack of AI exposure, we pivoted to upskill them first through targeted higher education courses (GAI-ED tracks), empowering these lecturers and their pre-service teachers to become active co-designers and securing a sustainable national teacher pipeline.

Deviation 3 (Technology Scope): The last and most important change in our technology scope was going from training teachers to use only foreign AI to actively co-developing local, sovereign AI prototypes. Originally, the initiative relied on training educators in prompt engineering using globally dominant large language models (LLMs) like ChatGPT and Gemini. However, our "Speculative Epistemic Auditing" exercises revealed a severe "hallucination crisis"; these foreign models enacted epistemic erasure by fabricating localised facts, defaulting to Euro-American cultural examples, and ignoring the structural realities of the Accelerated Education Programme (AEP). To protect against this data colonialism, we ceased our reliance on imported AI and instead used anonymised local datasets to build two sovereign generative AI prototypes specifically fine-tuned for Uganda's Competency-Based Curriculum (CBC) and AEP contexts, ultimately handing the CBC model over to the National Curriculum Development Centre (NCDC) for scaling.

Reasons for Deviations & Insights from Reflection Meetings

The deviations executed during this work were not arbitrary; they were direct, data-driven responses derived from rigorous feedback collected during our initial Focus Group Discussions (FGDs) and ongoing reflection meetings. By embedding a Participatory Action Research (PAR) framework into our methodology, the consortium treated initial operational friction and stakeholder hesitation not as barriers but as vital diagnostic data. This continuous cycle of reflection and feedback ensured that our interventions remained contextually grounded, ultimately revealing three profound insights that necessitated our strategic pivots to protect the educational sovereignty of fragile communities.

Reflection Insight 1 (The Buzzword Barrier): The first critical insight emerged from the "Buzzword Barrier" identified during early stakeholder engagements. In our FGDs, high-level district administrators and Monitoring, Evaluation, Accountability, and Learning (MEAL) officers forcefully articulated that generative AI was perceived merely as a high-risk "buzzword." Because they lacked a foundational understanding of the technology and its data mechanisms, they rightfully withheld their approval for direct classroom deployment, citing their mandate to protect vulnerable refugee populations. This reflection proved that deploying technology without securing institutional consent was reckless and would create unsanctioned, high-risk environments. It directly justified our deviation to adopt a "governance-first" approach, delaying classroom deployment to build stakeholder regulatory capacity through the specialised GAI-ADM-101 course.

Reflection Insight 2 (The TTI Gap): The second major reflection insight exposed the "Teacher Training Institution (TTI) Gap". Stakeholders consistently warned the consortium that focusing solely on active, in-service educators in refugee camps was a temporary solution to a deeper systemic crisis. Reflection meetings highlighted that national universities and TTIs were failing to produce "AI-ready" graduates, meaning any localised progress would quickly evaporate as untrained teachers continued to enter the workforce, inadvertently widening the digital divide between rural and urban contexts. This urgent feedback justified our deviation to expand the project's target audience upstream, resulting in the strategic pivot to secure the national teacher pipeline by upskilling university lecturers and pre-service educators at Muni University.

Reflection Insight 3 (The Hallucination Crisis): The final and most profound insight emerged during the "Speculative Epistemic Auditing" reflection sessions, where the "Hallucination Crisis" was exposed. As master trainers actively tested foreign large language models (LLMs), like ChatGPT and Gemini, against the Ugandan competency-based curriculum (CBC) and acceleration education program (AEP), they discovered that these tools exhibited severe "epistemic erasure". The imported models consistently fabricated local facts, ignored local languages, and defaulted to Western-centric pedagogical assumptions. This critical reflection demonstrated that dependence on foreign AI monopolies was not only technically incorrect but also pedagogically dangerous, jeopardising the education system's autonomy. This provided undeniable justification for our ultimate departure: ceasing our dependence on imported tools and co-creating our own sovereign, localised CBC and AEP generative AI prototypes.

Project Impact, Scalability, and Ethical Governance

The profound, statistically significant cognitive shifts among the participating educators serve as the true measure of this work's impact. Baseline diagnostics originally exposed a severe epistemic vulnerability, revealing that 17% of pre-service and in-service teachers possessed a "poor" or "very poor" ability to critically identify algorithmic biases and hallucinations in AI-generated content. Following our targeted capacity-building intervention, post-training analyses demonstrated a massive reduction in this vulnerability. The percentage of teachers rating their speculative auditing skills as "Good" or "Excellent" surged dramatically, systematically collapsing the "poor" baseline categories. Furthermore, the initiative's "governance-first" approach yielded exceptional dividends, with over 77% of participating educators reporting they now felt highly prepared to engage in critical The

technical sovereignty necessary to enforce trauma-informed pedagogical boundaries should support institutional AI ethics discussions.

The consortium set up a strict data governance system to protect very vulnerable refugee populations and ensure that this pedagogical empowerment was matched by unassailable ethical integrity. On April 10, 2025, War Child Canada and the GAI4Educ Consortium, which includes Makerere University, the National Curriculum Development Centre (NCDC), and the University of Edinburgh, signed a full Data Sharing Agreement (DSA). This DSA ensured strict compliance with Uganda's Data Protection and Privacy Act, the European Union's GDPR, and Canada's PIPEDA. By legally mandating data minimisation, pseudonymisation, and the use of secure, air-gapped AI sandboxes, the consortium permanently shielded learner data from extraction by foreign technology monopolies, guaranteeing that our sovereign AI models were developed without perpetuating digital colonialism or algorithmic bias.

The process was deliberately engineered for long-term scalability and organic growth, beyond immediate statistical success and ethical compliance. By abandoning centralised, parachute-style training in favour of a 1:9 Trainer-of-Trainers (ToT) cascading model, the initiative successfully transformed its initial cohort of 612 targeted educators into highly proficient "master trainers". Strategically anchoring this pipeline within teacher training institutions like Muni University ensures that a continuous stream of AI-ready, decolonially minded educators will enter the workforce. Equipped with friendly Open Educational Resources (OERs) and localised prompt libraries, these master trainers are perfectly positioned to cascade their sovereign competencies autonomously to unreached veteran teachers deep within rural and refugee-hosting ecosystems, sustaining the intervention's impact long after the initial grant cycle concludes.

Recent global recognition reflects the ultimate validation of this work's participatory, culturally grounded methodology. The rigorous, decolonial approach taken by the consortium yielded preliminary findings so robust that the project was awarded the prestigious Google Academic Research Award. This critical supplementary funding serves as a powerful endorsement of our strategic deviations, providing the resources necessary to rapidly scale the continuous hosting of our custom e-learning platform and further optimise our sovereign CBC and AEP Generative AI prototypes. By successfully shifting educators from passive consumers of foreign technology to sovereign co-creators of localised AI, the GAI-ACT initiative has established an empirically proven, scalable blueprint for securing educational equity and digital sovereignty across the Global South.

Bibliography

1. 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>
2. 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.
3. 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.
4. 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.
5. Google Academic Research Award Program recipients. (n.d.). Research.Google. Retrieved March 28, 2026, from <https://research.google/programs-and-events/google-academic-research-awards/google-academic-research-award-program-recipients/?filtartab=2024>