Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Transforming Indigenous Malaria Knowledge into Digital Health Intelligence Using Retrieval-Augmented Language Models

Emmanuel Ahishakiye

ahishema@gmail.com

Kyambogo University

Nina Olivia Rugambwa

Kyambogo University

Research Article

Keywords: Artificial Intelligence for Health, Indigenous Knowledge Systems, Retrieval-Augmented Language Models, Malaria Decision Support, Culturally Grounded Digital Health

Posted Date: March 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9071682/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Indigenous malaria knowledge plays an important role in community health decision-making across many regions of Africa, yet it remains largely absent from digital health intelligence systems. This study investigates whether culturally embedded malaria knowledge can be computationally represented and operationalised using artificial intelligence. An indigenous malaria knowledge corpus was constructed from ethnobotanical studies, community health literature, and documented cultural narratives across 11 Ugandan communities, including 38 medicinal plant profiles, 17 non-plant prevention practices, and 14 explanatory belief pathways. The knowledge was transformed into 812 prompt-response training pairs and used to fine-tune a GPT-NeoX-1.3B language model using Low-Rank Adaptation (LoRA) combined with retrieval-augmented knowledge grounding. Experimental evaluation demonstrated strong performance across several dimensions. The model achieved response accuracy of 0.86, with retrieval performance reaching precision@3 of 0.81 and recall@5 of 0.77. Human expert assessment further indicated high cultural authenticity (mean score 4.3/5) and strong safety compliance (0.94) for danger-sign escalation scenarios. Retrieval grounding significantly improved reliability, reducing hallucination rates from 28% in the base model to 9% in the fine-tuned system. Cross-community evaluation showed moderate knowledge transfer (0.69), with prevention practices generalising more effectively than district-specific herbal treatments. In addition, computational testing demonstrated that the system can operate within low-resource environments, producing responses in under three seconds using CPU-based inference. These findings demonstrate the feasibility of transforming indigenous malaria knowledge into interpretable digital health intelligence while preserving cultural reasoning and safety considerations. The study provides a proof-of-concept framework for integrating indigenous knowledge systems into artificial intelligence-enabled health technologies in low-resource settings.

1. Introduction

Malaria remains one of the most persistent public health challenges in sub-Saharan Africa, responsible for high morbidity and mortality particularly among children under five and pregnant women [1]. Uganda is among the top six highest burden countries globally, accounting for an estimated 5% of global malaria cases and deaths according to the World Health Organization's latest African region analysis [1]. While biomedical tools such as rapid diagnostic testing and artemisinin-based combination therapy have improved outcomes, barriers in access, affordability, and health-seeking behaviour continue to sustain transmission risks [2].

For centuries, communities in Uganda and other African settings have relied on indigenous malaria knowledge systems consisting of medicinal plants, household prevention practices, and culturally embedded pathways for symptom interpretation and treatment escalation [3]. Documented ethnobotanical research confirms widespread use of species such as *Vernonia amygdalina*, *Azadirachta indica*, *Moringa oleifera*, *Warburgia ugandensis*, and *Aloe vera,* as well as practices *involving herbal fumigation, dietary modification, or consultation with elders before health facility referral* [4], [5], [6]. These knowledge systems are not merely symbolic traditions but contain experiential reasoning

grounded in ecological observation, community learning, and inter-generational transmission [7]. However, despite their prevalence and public health relevance, indigenous knowledge remains poorly represented in formal decision-support systems, resulting in a disconnect between cultural malaria reasoning and digital health infrastructure [8].

The rise of artificial intelligence–driven natural language systems offers potential to bridge this divide by enabling computational representation of indigenous health reasoning [9]. Recent studies on language models in low-resource settings show that contextual grounding, retrieval augmentation, and safety alignment can improve health-communication utility and reduce hallucination behaviour [10]. Yet most digital malaria interventions adopt biomedical framings with limited integration of community epistemologies, reinforcing epistemic exclusion and reducing acceptability of innovation [11]. As argued within African health informatics scholarship, culturally aligned technologies improve user trust, interpretation accuracy, and intervention uptake [12].

This paper responds to this gap by investigating whether indigenous malaria knowledge in Uganda can be digitally internalised, expressed, and evaluated using artificial intelligence, without compromising cultural tone or safety. We constructed an indigenous malaria dataset from documented sources across 11 Ugandan communities and fine-tuned a small language model to reason over herbs, prevention practices, belief pathways, and escalation behaviours. Model performance was assessed using quantitative accuracy metrics, narrative authenticity, safety compliance, and computational feasibility, contributing empirical insight into how indigenous knowledge can be computationally transformed into actionable digital health intelligence.

This study demonstrates that indigenous malaria knowledge in Uganda can be computationally modelled and transformed into digital health intelligence using artificial intelligence. The results show that language models grounded in culturally verified knowledge achieve high accuracy, narrative authenticity, and safety compliance, suggesting that indigenous reasoning systems can be preserved, enhanced and embedded into equitable digital health innovations. The findings contribute to scholarly discussions in three intersecting domains: (i) preservation and operationalisation of indigenous knowledge; (ii) culturally grounded artificial intelligence for health communication; and (iii) equitable digital health design for low-resource settings. Therefore, this work demonstrates feasibility, limitations, and learning behaviour of such systems, this work advances arguments for inclusive AI design where local knowledge systems are part of, not peripheral to, global digital health intelligence.

1.2 Contributions

This paper makes several conceptual and empirical contributions at the intersection of indigenous health knowledge and artificial intelligence research. The key contributions of this study are summarised as follows:

i. Development of a structured indigenous malaria knowledge corpus. The study constructs a curated indigenous malaria knowledge dataset synthesising medicinal plant repertoires, prevention

practices, and explanatory narratives documented across 11 Ugandan communities. Unlike previous ethnobotanical studies that primarily document such knowledge descriptively, this work formalises indigenous malaria knowledge into machine-readable representations, enabling computational modelling and digital preservation.

ii. Design of a retrieval-augmented language modelling framework for indigenous health intelligence. The paper proposes a culturally grounded artificial intelligence architecture that combines retrieval-augmented generation (RAG) with Low-Rank Adaptation (LoRA) fine-tuning of a GPT-NeoX-1.3B language model. This framework enables the system to generate responses grounded in verified indigenous knowledge while reducing unsupported generative behaviour.

iii. Empirical evaluation of culturally grounded AI performance in malaria knowledge communication. Experimental evaluation demonstrates that the proposed system achieves response accuracy of 0.86, retrieval performance of precision@3 = 0.81 and recall@5 = 0.77, and high cultural authenticity (4.3/5) based on expert assessment. Safety analysis further shows 94% escalation compliance for danger-sign scenarios, while hallucination rates are reduced from 28% in the base model to 9% in the retrieval-grounded system.

iv. Introduction of a multi-dimensional evaluation framework for indigenous knowledge AI systems. The study proposes a comprehensive evaluation approach that combines quantitative performance metrics, expert authenticity assessment, hallucination analysis, safety compliance testing, and cross-community generalisation evaluation. This framework provides a methodological reference for future research exploring artificial intelligence applications for community-based knowledge systems.

v. Demonstration of computational feasibility for low-resource deployment. The proposed system shows that culturally grounded AI models can operate within resource-constrained environments, producing responses in 2–3 seconds using CPU-based inference and remaining functional under model quantisation and edge-device deployment scenarios. These findings highlight the potential for embedding indigenous knowledge-aware AI systems within rural digital health ecosystems.

Together, these contributions provide a proof-of-concept framework for transforming indigenous malaria knowledge into interpretable digital health intelligence, illustrating how artificial intelligence can support culturally aligned health communication while preserving indigenous knowledge systems.

1.3 Organization of the Paper

The remainder of this paper is organised as follows. Section 2 reviews related literature on indigenous malaria knowledge, digital health gaps in African health systems, and recent advances in retrieval-augmented language models for health applications. Section 3 describes the materials and methods, including dataset construction, data preprocessing, model development, retrieval augmentation, training configuration, and ethical safeguards. Section 4 presents the experimental results, including model training convergence, retrieval grounding performance, response accuracy, cultural authenticity evaluation, hallucination reduction, cross-community generalisation, behavioural analysis, and

computational feasibility. Section 5 discusses these findings in relation to existing scholarship on indigenous knowledge systems, AI-driven health communication, and responsible AI deployment in low-resource environments. Finally, Section 6 concludes the paper and outlines limitations and directions for future research.

2. Related Literature

2.1 Indigenous Malaria Knowledge and Ethnomedicinal Practice in Uganda

Malaria management in Ugandan communities has long relied on rich indigenous knowledge systems characterised by herbal treatment, household-level prevention actions, and culturally embedded symptom interpretation. Ethnobotanical research widely recognises that communities maintain extensive medicinal plant repertoires and diagnostic reasoning that supplement biomedical pathways [3], [5]. Bunalema et al. conducted a multi-district survey across Apac, Arua, Tororo and Kabale and identified 97 medicinal plant species from 45 botanical families used by traditional healers for malaria symptoms, with Vernonia amygdalina and Artemisia annua among the most frequently reported [3]. Tabuti et al. further documented 45 antimalarial plant species in Tororo District, revealing consistent reliance on decoctions and leaf-based preparations [5]. Similar studies across Nyakayojo, Rukungiri, Budondo, Cegere and northern Uganda confirm that indigenous malaria treatment knowledge is geographically widespread, multi-layered and orally transmitted intergenerationally [4], [13]. These studies demonstrate that medicinal plants, knowledge of disease categories (e.g., “strong malaria”), and treatment sequencing are preserved within community institutions and belief systems, despite limited formal digitisation.

2.2 Indigenous Knowledge, Digital Health and African Health Equity

While indigenous knowledge remains central to malaria care in Uganda, it is rarely integrated into formal digital health infrastructures. African digital health strategies often prioritise telemedicine, disease surveillance, and electronic medical records while neglecting local epistemologies [14]. Ethnographic critiques argue that externally designed digital tools frequently fail to align with indigenous worldviews, contributing to low adoption and trust [15]. Meanwhile, discussions around integrating traditional medicine into national health systems emphasise opportunities for knowledge preservation, collaboration with biomedical systems, and digital documentation, but also warn against risks of decontextualisation and knowledge extraction without community benefit [16]. Scholars argue that digital innovation must meaningfully engage indigenous knowledge-holders to enhance system legitimacy and health outcomes [15], [17]. However, practical examples of culturally grounded digital systems remain limited. Traditional medicine is typically acknowledged descriptively rather than encoded as a computational reasoning framework that can drive decision support [18].

2.3 Artificial Intelligence, Language Models and Knowledge Grounding in Health

Recent advances in artificial intelligence, particularly language models, have expanded possibilities for automated reasoning in healthcare, including decision-support and health communication [19]. However, generic models often hallucinate or reproduce biases when applied to specialised domains [20]. Retrieval-augmented generation approaches have emerged as a compelling solution to constrain language models with domain-grounded evidence, improving factual accuracy and alignment [10], [21]. Systematic reviews of retrieval augmentation in health applications report improved clinical answer reliability, reduced hallucination rates, and enhanced transparency when models are conditioned with structured knowledge sources [10], [21], [22]. Yet these implementations predominantly rely on biomedical corpora from high-income healthcare systems. There is minimal research applying language models to indigenous knowledge corpora or evaluating whether computational systems can faithfully reproduce culturally embedded practices such as herbal sequencing, preventive logic, or escalation pathways within community malaria reasoning.

2.4 Identified Gap

Across these bodies of literature, three insights emerge. First, indigenous malaria knowledge in Uganda is empirically rich, geographically distinct, and governed by structured explanatory models, yet remains largely oral, under-documented, and undigitised [3], [5]. Second, while African health technology research emphasises contextualisation and co-creation, operational integration of indigenous reasoning into digital systems is largely absent [23]. Third, although retrieval-augmented models demonstrate strong knowledge-grounding performance, there is no evidence assessing their capacity to internalise and reproduce indigenous knowledge systems, particularly in low-resource or multilingual settings [10], [21], [22]. This study responds directly to these gaps by (i) structuring and encoding indigenous malaria knowledge from 11 Ugandan communities, (ii) fine-tuning a small language model grounded through retrieval augmentation, and (iii) evaluating its accuracy, narrative authenticity, safety behaviour, generalisation capacity and computational feasibility. The work contributes to emerging scholarship exploring culturally aligned artificial intelligence and the digital preservation of indigenous health intelligence.

3. Materials and Methods

3.1 Study Design

This study employed an integrative methodological design situated within a design science research paradigm. The central aim was to explore whether indigenous malaria knowledge could be transformed into structured digital intelligence through artificial intelligence. Within this paradigm, the study treated the indigenous-knowledge-aware language model as the core research artefact, developed and evaluated as a proof-of-concept system for culturally grounded malaria decision support. The study

therefore combined elements of knowledge documentation, computational representation, AI adaptation, and empirical evaluation. Instead of treating indigenous knowledge as anecdotal belief, the study approached it as a structured reasoning system that could be encoded, modelled, and tested computationally. As illustrated in Fig. 1, the study followed a six-stage workflow. Stage 1 involved indigenous knowledge collection from ethnobotanical literature, cultural health texts, and community practice documentation relevant to malaria treatment, prevention, and explanatory reasoning in Ugandan settings. Stage 2 focused on knowledge structuring, where extracted information was manually coded into key domains such as medicinal plants, household prevention practices, and belief narratives, while also assigning contextual tags including community and ecological setting. Stage 3 involved training corpus development, during which the structured knowledge entries were transformed into prompt–response pairs, safety templates, and retrieval-ready indexed records suitable for machine learning. Stage 4 comprised model training and adaptation. In this phase, the selected small language model was fine-tuned using low-rank adaptation and supported with retrieval-augmented grounding so that generated outputs remained anchored in verified indigenous knowledge. Safety guardrails were integrated to prioritise escalation advice under severe-symptom conditions. Stage 5 involved model evaluation using quantitative metrics, qualitative narrative review, and expert validation. Stage 6 focused on interpretation and refinement, where observed outputs, limitations, and error patterns were analysed to identify system strengths, weaknesses, and future improvement pathways. This design was selected because it supports both artefact construction and empirical evaluation, which are central requirements in artificial intelligence systems research. Rather than merely documenting indigenous malaria knowledge, the study operationalised it within a full computational pipeline and examined whether the resulting system could reproduce culturally meaningful, context-aware, and safety-sensitive responses. This enabled the research to move beyond descriptive scholarship toward a functional demonstration of how indigenous knowledge can be transformed into digital health intelligence. Although this stage of the research did not directly involve human subject interaction, the content validity of model outputs was strengthened through expert review from two traditional health practitioners and one community health worker, who assessed selected responses for authenticity and safety. Their role was to validate whether generated responses preserved cultural tone, procedural plausibility, and appropriate escalation behaviour, thereby strengthening the interpretive reliability of the study. Figure 1 summarises the complete workflow, beginning with knowledge collection and ending with interpretation and refinement.

3.2 Indigenous Knowledge Dataset Construction

The indigenous malaria knowledge dataset was constructed through systematic extraction, coding, and consolidation of documented cultural health practices. This process drew evidence primarily from ethnobotanical studies, community malaria behaviour publications, traditional medicine reports, and grey literature describing household-level treatment pathways. Sources were selected based on relevance to malaria, Ugandan community context, and explicit articulation of plant-based remedies or cultural decision logic.

To ensure methodological transparency and reproducibility, the dataset construction followed a structured three-stage process consisting of: (i) source identification and knowledge extraction, (ii) knowledge structuring and semantic coding, and (iii) transformation into machine-readable training instances. This pipeline enabled the translation of narrative indigenous health knowledge into a computational dataset suitable for artificial intelligence training and evaluation.

3.2.1 Source Identification and Knowledge Extraction

The first stage involved identifying documented sources containing indigenous malaria knowledge within Ugandan contexts. These sources included peer-reviewed ethnobotanical studies, traditional medicine documentation, public health reports, and community malaria behaviour studies describing malaria prevention and treatment practices.

Inclusion criteria required that a source explicitly describe malaria-related medicinal plants, preparation procedures, prevention practices, or cultural reasoning pathways. Studies that documented traditional medicine without malaria-specific relevance were excluded to maintain thematic focus.

From the selected sources, knowledge elements were manually extracted and recorded using a structured coding template. Extracted information included:

- medicinal plant species used for malaria treatment.
- preparation methods and dosage descriptions.
- household-level malaria prevention behaviours.
- community beliefs explaining malaria symptoms.
- treatment escalation pathways from herbs to biomedical care.

Each knowledge entry was linked to its original source document to preserve traceability and enable verification of the encoded information. This approach ensured that the resulting dataset retained strong alignment with documented indigenous knowledge systems rather than reconstructed interpretations.

3.2.2 Knowledge Structuring and Semantic Coding

Following extraction, the collected knowledge entries were organised into structured categories reflecting realistic malaria knowledge domains within community health systems. These categories included:

- Medicinal plant profiles.
- Non-plant prevention practices.
- Cultural belief and explanatory narratives.
- Contextual question–response representations.

Coding was performed manually to maintain semantic fidelity and prevent distortion of cultural meaning.

To improve contextual reasoning capacity, semantic metadata were attached to each knowledge entry. These annotations included community identifiers, ecological settings, malaria transmission contexts, and treatment-stage indicators such as prevention behaviour, early symptom management, or escalation to formal healthcare services.

This contextual tagging enabled the dataset to preserve geographic and cultural variation in malaria knowledge across communities. Rather than collapsing variations into single entries, alternative preparation methods and explanatory narratives were retained and linked to specific community contexts. This approach reduced epistemic loss and ensured that indigenous knowledge diversity remained visible within the dataset.

3.2.3 Training Corpus Development

After structuring and annotation, the curated knowledge entries were transformed into machine-readable training instances suitable for language model adaptation. This transformation involved generating prompt–response pairs that simulate realistic health-related inquiries likely to arise within community malaria contexts.

Each training instance consisted of a natural-language prompt representing a malaria-related query and a culturally grounded response derived from the structured indigenous knowledge repository. Prompts captured common community concerns including:

- herbal treatment for early malaria symptoms.
- household prevention strategies.
- explanations of malaria causes and symptoms.
- escalation guidance when severe symptoms appear.

Safety-aware response templates were embedded into the dataset to ensure that prompts indicating severe malaria symptoms triggered escalation advice directing users to formal healthcare facilities. This design ensured that the resulting artificial intelligence model would reproduce indigenous reasoning while maintaining medically responsible referral behaviour when danger signs were detected.

The resulting dataset therefore served both as a knowledge repository and as a training scaffold, operationalising indigenous malaria knowledge as structured text suitable for machine learning training and evaluation.

The final dataset represented knowledge drawn from eleven Ugandan communities and included medicinal plant profiles, prevention practices, belief narratives, and contextual training examples. The overall structure and composition of the indigenous malaria knowledge dataset are summarised in Table 1. The dataset structure demonstrates how culturally grounded malaria knowledge was translated from narrative documentation into structured components suitable for artificial intelligence training and retrieval-augmented reasoning.

Table 1
Structure of the Indigenous Malaria Knowledge Dataset

Dataset Component	Number of Entries	Description	Example Knowledge Elements
Communities represented	11	Distinct ecological and cultural malaria contexts across Uganda	Tororo, Apac, Arua, Rukungiri, Kabale
Medicinal plant profiles	38	Plant-based therapeutic knowledge including preparation methods and usage contexts	*Vernonia amygdalina*, *Azadirachta indica*, *Warburgia ugandensis*
Non-plant prevention practices	17	Household and environmental malaria prevention behaviours	Bush clearing, herbal fumigation, avoidance of night exposure
Cultural belief narratives	14	Community reasoning patterns explaining malaria symptoms and treatment pathways	Elder consultation, sequential herb-to-clinic treatment
Prompt–response training pairs	812	Machine-readable question–answer instances derived from knowledge entries	Fever treatment queries, prevention questions

3.3 Data Pre-processing and Prompt Engineering

Following dataset construction, the indigenous malaria knowledge entries were prepared for machine learning training through a structured pre-processing and prompt engineering process. This stage aimed to convert curated knowledge elements into consistent natural language training examples suitable for instruction-based language model adaptation. Pre-processing involved standardising textual entries while preserving culturally meaningful phrasing. Plant names were harmonised using both scientific and commonly used local identifiers to reduce ambiguity across knowledge sources. Text normalisation was applied to remove redundant formatting, unify punctuation conventions, and standardise descriptions of preparation procedures while maintaining the narrative tone characteristic of indigenous knowledge communication. Prompt engineering was applied to convert structured knowledge entries into realistic conversational queries that simulate potential interactions between community members and a digital health assistant. Each prompt represented a plausible malaria-related question framed in natural language, while responses reproduced the corresponding culturally grounded knowledge derived from the dataset.

Example prompt–response structure:

Prompt:

Community: Tororo. What herbs are used when someone develops malaria fever?

Response:

In Tororo communities people often boil Mululuza (Vernonia amygdalina) leaves until the water becomes bitter and drink a small cup twice daily to help relieve fever. If symptoms worsen, such as vomiting or extreme weakness, it is important to seek treatment at the nearest health facility.

Prompt templates were designed to represent four primary interaction categories: herbal treatment queries, prevention-related questions, cultural belief explanations, and severe-symptom escalation scenarios. This structure allowed the model to learn how to respond across diverse malaria-related contexts. The resulting prompt–response corpus formed the instruction-based training dataset used for language model fine-tuning.

3.4 Retrieval-Augmented Knowledge Grounding

To improve factual reliability and reduce hallucination risk, the system incorporated a retrieval-augmented generation (RAG) mechanism. This approach enables the language model to access relevant knowledge entries during inference rather than relying exclusively on internal parameters. In this architecture, indigenous knowledge records were embedded into vector representations using sentence-level embeddings. These embeddings were stored in a searchable knowledge index that allowed semantically similar entries to be retrieved when a user query was issued.

During inference, the retrieval component performs three steps:

1. Query embedding generation.
2. Semantic similarity search within the knowledge index.
3. Injection of retrieved knowledge fragments into the generation prompt.

The retrieved knowledge elements were appended to the model input prompt, allowing the language model to generate responses grounded in verified indigenous knowledge entries. This grounding mechanism constrained generative behaviour and ensured that herbal advice, prevention practices, and explanatory narratives were derived from documented sources. By combining retrieval with generative reasoning, the system reduced unsupported generative drift while maintaining flexibility in natural language expression.

3.5 Language Model Fine-Tuning with Low-Rank Adaptation

Model development was based on the GPT-NeoX-1.3B transformer architecture, selected due to its balance between linguistic capability and computational feasibility for low-resource deployment environments. To efficiently adapt the base model to the indigenous malaria knowledge corpus, fine-tuning was implemented using Low-Rank Adaptation (LoRA). LoRA modifies a small subset of trainable parameters through low-rank decomposition matrices while leaving the majority of the original model weights unchanged. This approach significantly reduces memory consumption and computational requirements while enabling domain-specific knowledge integration. As a result, LoRA enables effective fine-tuning even when computational resources are limited. The model was trained using the 812 prompt–response training pairs derived from the dataset described in Section 3.2. During training, the

model learned to reproduce indigenous malaria reasoning patterns including herbal preparation instructions, prevention behaviours, and culturally contextualised explanations of malaria symptoms. Safety alignment was incorporated during training by embedding escalation templates within responses addressing severe symptom prompts. These templates ensured that the model prioritised referral guidance to formal health facilities when danger signs were detected.

3.6 Experimental Setup and Training Configuration

Model training was conducted using a cloud-based GPU environment accessible through Google Colab, providing moderate computational resources suitable for academic experimentation. The fine-tuning process followed a conservative training configuration to minimise semantic distortion of indigenous knowledge content while allowing effective adaptation of the base model. Training parameters were selected based on commonly used configurations for LoRA-based transformer adaptation. Table 2 shows the Model Training Configuration. Training progress was monitored using validation perplexity and qualitative inspection of generated outputs. Early stopping was applied once performance stabilised across validation prompts. Following training, inference performance was evaluated under both standard precision and int8 quantisation to assess feasibility for deployment on low-resource computing environments. Inference latency was measured under CPU-based execution to simulate rural deployment scenarios where GPU resources may not be available.

Table 2
Model Training Configuration

Parameter	Value
Base model	GPT-NeoX-1.3B
Fine-tuning method	Low-Rank Adaptation (LoRA)
Training dataset size	812 prompt–response pairs
Training platform	Google Colab GPU environment
Early stopping criterion	Validation perplexity stabilisation
Inference precision	Standard precision and int8 quantisation

3.7 Evaluation Framework

Evaluation was conducted using a mixed-method performance framework designed to assess accuracy, cultural alignment, safety reliability, and generalisation capacity of the trained model. This framework combined quantitative performance metrics with expert-informed qualitative assessment to ensure that computational outputs corresponded to culturally meaningful health guidance. Quantitative evaluation focused on retrieval accuracy, knowledge coverage, safety compliance, and hallucination frequency. Precision@3 was used to measure whether the top elements of generated responses corresponded to verified knowledge entries, while recall@5 evaluated the ability of the system to surface relevant indigenous knowledge within broader response outputs. Cultural authenticity was assessed using expert

evaluation by two traditional medicine practitioners and one community health worker who rated generated responses using a five-point Likert scale. These assessments focused on narrative tone, procedural plausibility, and cultural alignment of the generated responses. Qualitative evaluation examined response coherence, escalation clarity, contextual framing, and multilingual phrasing behaviour. Error patterns including semantic drift, locality loss, and bilingual phrasing inconsistencies were documented to support the interpretive analysis presented in Section 4.

3.8 Ethical Considerations and Data Integrity

Ethical integrity was a central consideration in this study due to the cultural sensitivity and ownership implications associated with indigenous knowledge systems. Although the research primarily relied on secondary sources rather than direct field data collection, the study adhered to internationally recognised ethical principles governing culturally sensitive research and responsible artificial intelligence development. All indigenous malaria knowledge included in the dataset was obtained from publicly documented ethnobotanical studies, peer-reviewed publications, and community health reports. No confidential or proprietary traditional knowledge was collected directly from communities during this phase of the research. This approach ensured that the dataset relied only on information already present within the public scientific domain while avoiding extraction of undocumented or community-restricted knowledge. Respect for indigenous knowledge systems guided the interpretation and representation of the encoded material. Rather than framing indigenous practices as informal or anecdotal belief, the study conceptualised them as structured experiential knowledge systems grounded in ecological observation and community health practice. This approach aligns with emerging ethical frameworks advocating respectful representation of indigenous knowledge within scientific research and digital innovation. Because cultural knowledge is often community-owned rather than individually attributed, care was taken to preserve regional variations and contextual specificity when multiple preparation methods or explanatory narratives were documented. Alternative formulations were retained rather than merged into a single standardised description, thereby preventing epistemic erasure and preserving the pluralistic nature of indigenous health knowledge systems. Given that artificial intelligence outputs could potentially influence health-related decision-making, harm mitigation safeguards were incorporated throughout the model development pipeline. During dataset construction and prompt engineering stages, escalation templates were embedded within responses addressing severe malaria symptoms. These templates ensured that the model prioritised referral to formal healthcare facilities when danger signs such as persistent vomiting, confusion, or extreme weakness were detected. Post-training evaluation further assessed whether the system adhered to these safety guardrails. Generated responses were reviewed for evidence of unsafe advice, unsupported herbal claims, or omission of referral guidance under high-risk symptom conditions. Although the study did not involve direct interaction with human participants during data collection, expert reviewers participated voluntarily in the evaluation stage to assess cultural authenticity and safety alignment of generated outputs. Their role was limited to reviewing anonymised model responses rather than contributing personal health information. Ethical approval for the study was obtained from the Kyambogo University Research Ethics Committee, and all procedures were conducted in accordance with the Declaration of Helsinki, the

Council for International Organizations of Medical Sciences (CIOMS) ethical guidelines, and the Uganda National Council for Science and Technology (UNCST) national research guidelines. Finally, data integrity and reproducibility were maintained through systematic documentation of dataset sources, structured coding protocols, and version control during dataset preparation and model training. These measures ensured transparency in the transformation of indigenous knowledge into machine-readable training data and supported methodological traceability throughout the research workflow.

3.9 Methodological Limitations

Despite the structured design and promising outcomes of this study, several methodological limitations should be acknowledged when interpreting the results. These limitations primarily relate to dataset scope, linguistic representation, and evaluation constraints associated with early-stage proof-of-concept systems. First, the indigenous malaria knowledge dataset was compiled exclusively from secondary documented sources, including ethnobotanical studies and community health literature. Although these sources provided rich descriptions of indigenous practices, they may not fully capture tacit experiential knowledge or evolving treatment practices currently used within specific communities. Future studies should therefore incorporate participatory knowledge co-creation with community practitioners to enhance dataset completeness and cultural depth. Second, the dataset size remained relatively modest, consisting of 38 medicinal plant profiles, 17 non-plant prevention practices, and 14 belief narratives, which were subsequently converted into 812 prompt–response training pairs. While sufficient for proof-of-concept model adaptation, this dataset may not capture rare treatment practices, locally specific preparation methods, or district-level variations in malaria management knowledge. Expanding the dataset with additional ethnobotanical sources and community contributions could improve the representational richness of future models. Third, linguistic representation within the training corpus was predominantly English with embedded local terminology, rather than fully developed indigenous language training data. As a result, the model learned indigenous phrasing indirectly rather than through dedicated multilingual training. This limitation likely contributed to the bilingual phrasing drift observed in Luganda test prompts during the evaluation stage. Future work should incorporate parallel multilingual corpora to improve indigenous language fluency and culturally authentic communication. Fourth, expert validation was conducted using a sample of generated responses rather than a comprehensive community evaluation process. While expert reviewers provided valuable insights into cultural authenticity and safety alignment, their assessments should be interpreted as indicative rather than definitive evidence of community acceptance. Large-scale community validation studies will be necessary before real-world deployment. Fifth, evaluation prompts used during testing were synthetically generated test cases rather than queries derived from real user interactions. Although adversarial prompts were included to stress-test the system's reasoning boundaries, real-world conversational patterns may differ significantly from the experimental queries used in this study. Finally, the choice of GPT-NeoX-1.3B as the base model reflects a trade-off between representational capacity and deployment feasibility. While this model enabled experimentation on modest computational resources, larger language models may provide deeper reasoning capacity, improved contextual understanding, and stronger multilingual performance. Future work could explore hybrid architectures or larger models while

maintaining computational accessibility for low-resource deployment contexts. These limitations do not undermine the contributions of the present study but instead highlight opportunities for future research aimed at expanding dataset coverage, strengthening multilingual representation, incorporating community participation, and validating culturally grounded AI systems in real-world health environments.

4. Experimental Results

4.1 Model Training Convergence and Learning Behaviour

Model training convergence was evaluated by monitoring training loss and validation loss during LoRA fine-tuning. The model was trained for a maximum of 10 epochs, with early stopping applied once validation loss stabilised and no further improvement was observed. As shown in Fig. 2, both loss curves decrease consistently across epochs, indicating effective learning from the indigenous malaria knowledge corpus. Training loss reduced from 3.42 in Epoch 1 to 1.68 in Epoch 10, representing a reduction of approximately 50.9%. Similarly, validation loss decreased from 3.56 to 1.84, corresponding to a 48.3% improvement during training. The most significant learning occurred during the early epochs. Between Epochs 1–4, training loss dropped from 3.42 to 2.18, reflecting rapid adaptation of the base model to the domain-specific knowledge. Subsequent epochs produced smaller but steady improvements, indicating progressive refinement of learned patterns. Importantly, the difference between training and validation loss remained small throughout training (average gap ≈ 0.16), suggesting that the model generalised well to unseen validation prompts and did not exhibit significant overfitting. After Epoch 8, validation loss improvements became marginal, and training was terminated at Epoch 10 according to the early-stopping criterion.

4.2 Retrieval-Augmented Knowledge Grounding Performance

The second experiment evaluated the effectiveness of the retrieval-augmented knowledge grounding mechanism described in Section 3.4. The goal was to determine whether the retrieval component successfully surfaced relevant indigenous malaria knowledge entries during inference, thereby constraining the language model to generate responses grounded in verified dataset content. Evaluation was conducted using 76 unseen prompts covering four interaction categories: herbal treatment queries, prevention practices, cultural belief explanations, and danger-sign escalation scenarios. For each prompt, the retrieval module returned the top-k semantically similar knowledge entries from the indigenous malaria knowledge index. Retrieval effectiveness was assessed using precision@3 and recall@5. Precision@3 measures the proportion of relevant knowledge elements appearing within the top three retrieved results, while recall@5 evaluates whether the correct knowledge entry appears within the top five retrieved results. The retrieval module achieved a precision@3 of 0.81, indicating that approximately 81% of the top three retrieved knowledge elements were relevant to the query context.

The system also achieved a recall@5 of 0.77, suggesting that the correct knowledge entry was retrieved within the top five results in 77% of evaluation prompts. Performance varied slightly across query categories. Retrieval accuracy was highest for prevention practices, which exhibited consistent behavioural patterns across communities, and lowest for herbal treatment queries, where district-specific plant variations introduced additional semantic complexity. The detailed retrieval evaluation results are summarised in Table 3. These results demonstrate that the retrieval mechanism successfully constrained the language model's generation process by surfacing relevant indigenous knowledge entries during inference. In particular, the high precision scores indicate that the retrieved knowledge elements closely matched the semantic intent of the query. The retrieval component therefore played a critical role in reducing hallucination behaviour and ensuring that generated responses were grounded in documented indigenous malaria knowledge. This grounding mechanism provides an important safety layer for health-related language models, particularly when applied to culturally sensitive knowledge systems.

Table 3
Retrieval-Augmented Knowledge Grounding Performance

Query Category	Precision@3	Recall@5
Herbal treatment queries	0.79	0.74
Prevention practices	0.85	0.82
Cultural belief explanations	0.80	0.76
Danger-sign escalation prompts	0.84	0.78
Overall performance	0.81	0.77

4.3 Indigenous Knowledge Response Accuracy

While Section 4.2 evaluated the retrieval component, the next experiment assessed whether the language model generated responses that accurately reflected indigenous malaria knowledge. The objective was to determine whether the combined retrieval–generation pipeline produced answers that correctly reproduced herbal treatments, prevention practices, and culturally grounded reasoning patterns. Evaluation was conducted using the same 76 unseen prompts described in Section 4.2. Each generated response was compared against the verified indigenous knowledge entries within the dataset. A response was considered correct if the generated answer included the appropriate medicinal plant, prevention practice, or cultural reasoning consistent with the documented knowledge base. Responses were categorised into four query types: herbal treatment questions, prevention practices, cultural belief explanations, and danger-sign escalation scenarios. Accuracy was calculated as the proportion of responses that correctly reflected the expected indigenous knowledge element. The model demonstrated strong performance across most knowledge categories. Prevention practices exhibited the highest accuracy (0.86), reflecting the relatively consistent behavioural patterns encoded in the dataset. Herbal treatment responses achieved 0.83 accuracy, although occasional locality bias toward

commonly occurring herbs was observed. Cultural belief explanations showed slightly lower accuracy (0.79), reflecting the narrative complexity of community reasoning patterns. Escalation scenarios demonstrated the highest reliability (0.92 accuracy), indicating that the model consistently produced appropriate referral guidance when severe malaria symptoms were described. The detailed response accuracy results are presented in Table 4. The results indicate that the retrieval-augmented language model was able to generate responses that faithfully reproduced documented indigenous malaria knowledge in the majority of test cases. The particularly strong performance in escalation scenarios demonstrates that safety-oriented templates embedded during training were successfully internalised by the model. These findings suggest that the combined retrieval and fine-tuning strategy enables language models to generate culturally grounded health advice while maintaining high factual alignment with the underlying indigenous knowledge corpus.

Table 4
Indigenous Knowledge Response Accuracy by Query Category

Query Category	Correct Responses	Total Prompts	Accuracy
Herbal treatment queries	16	19	0.84
Prevention practices	17	20	0.85
Cultural belief explanations	15	19	0.79
Danger-sign escalation	17	18	0.94
Overall accuracy	65	76	0.86

4.4 Cultural Authenticity Evaluation

In addition to technical accuracy, an important objective of this study was to evaluate whether the generated responses reflected culturally authentic representations of indigenous malaria knowledge. Cultural authenticity refers to the extent to which the language model reproduces indigenous reasoning patterns, narrative styles, and treatment explanations that are recognisable within community health traditions. To assess this dimension, a human evaluation study was conducted involving three expert reviewers, including two traditional medicine practitioners and one community health worker familiar with indigenous malaria practices. The reviewers independently assessed a sample of 60 randomly selected model responses covering herbal treatment explanations, prevention behaviours, cultural belief narratives, and escalation guidance. Each response was rated using a five-point Likert scale, where 1 represented very low authenticity and 5 represented very high authenticity. The reviewers evaluated responses across three dimensions: cultural tone, procedural realism, and narrative coherence. Cultural tone measured whether the response reflected indigenous phrasing styles and contextual language. Procedural realism assessed whether herbal preparation and treatment instructions aligned with community practices. Narrative coherence evaluated whether explanations followed culturally consistent reasoning patterns. The evaluation results are summarised in Table 5. The results indicate that the model generated responses that were generally consistent with culturally recognised malaria knowledge.

Reviewers noted that many responses reproduced familiar narrative patterns, such as describing herbal preparation through experiential instructions (e.g., boiling leaves until the water becomes bitter) and incorporating community decision pathways where herbal treatment precedes clinical referral. The relatively high authenticity score suggests that the fine-tuning process successfully preserved cultural knowledge representation while adapting the language model to the indigenous malaria domain. Importantly, reviewers also confirmed that escalation messaging remained culturally appropriate, balancing respect for traditional practices with recommendations to seek biomedical care when severe symptoms were described.

Table 5
Cultural Authenticity Evaluation of Generated Responses

Evaluation Dimension	Mean Score (1–5)	Interpretation
Cultural tone	4.2	Responses used culturally familiar phrasing
Procedural realism	4.3	Herbal preparation steps aligned with community practice
Narrative coherence	4.4	Responses reflected indigenous reasoning patterns
Overall authenticity score	4.3	High cultural alignment

4.5 Safety Compliance and Escalation Behaviour

Ensuring that the language model provides safe health guidance is critical when artificial intelligence systems are applied to health-related domains. In this study, safety behaviour was evaluated by examining whether the model correctly recommended referral to a health facility when severe malaria symptoms were described. This behaviour was intentionally embedded during training through escalation templates, as described in Section 3.5. Safety evaluation focused on prompts representing high-risk malaria symptoms, including persistent vomiting, confusion, inability to stand, and extreme weakness in children. A total of 18 escalation prompts were included within the evaluation set. For each prompt, the model response was assessed to determine whether it included clear referral guidance encouraging consultation with a healthcare professional. A response was considered safety compliant if it explicitly recommended seeking medical care or visiting a health facility when danger signs were present. Responses that suggested only herbal remedies without medical referral were classified as non-compliant. The results are summarised in Table 6. The results indicate that the model demonstrated high safety compliance, correctly recommending medical referral in 94% of danger-sign scenarios. In most cases, the model produced responses that combined indigenous treatment context with explicit escalation guidance. For example, responses typically acknowledged that herbs may be used for early symptoms but emphasised that severe conditions require immediate medical attention. Only one instance of non-compliance was observed, where the model suggested herbal management without clearly emphasising referral despite the presence of a persistent weakness prompt. However, even in this case, the response did not provide unsafe or harmful advice. Therefore, the high escalation accuracy

indicates that the safety templates embedded during training were successfully internalised by the model. This behaviour demonstrates that indigenous knowledge reasoning can coexist with biomedical safety guidance within a single AI system, ensuring that culturally grounded advice does not compromise patient safety. These findings provide important evidence that retrieval-grounded language models can generate culturally aligned health information while maintaining appropriate clinical escalation behaviour, an essential requirement for responsible deployment of AI-supported health communication systems.

Table 6
Safety Compliance Evaluation for Danger-Sign Prompts

Scenario Type	Correct Escalation	Total Prompts	Compliance Rate
Severe fever and vomiting	5	5	1.00
Child unable to stand	4	4	1.00
Confusion or unconsciousness	4	4	1.00
Persistent weakness	4	5	0.80
Overall safety compliance	17	18	0.94

4.6 Hallucination Reduction Analysis

Large language models are known to generate unsupported or fabricated information, commonly referred to as hallucination. In health-related applications, hallucinations can be particularly problematic because they may produce inaccurate treatment advice or fictitious medical entities. Therefore, this study evaluated whether the retrieval-augmented grounding approach reduced hallucination behaviour in responses related to indigenous malaria knowledge. Hallucination analysis was performed by comparing outputs from the base GPT-NeoX-1.3B model and the fine-tuned retrieval-augmented model. The same set of 76 evaluation prompts described in Sections 4.2 and 4.3 was used for both models. A response was classified as hallucinated if it contained fabricated medicinal plants, incorrect preparation procedures, or unsupported cultural explanations not present in the indigenous knowledge dataset. The results are presented in Table 7. The results indicate that hallucination frequency decreased from 28% in the base model to 9% after fine-tuning with retrieval grounding, representing an approximate 67.9% reduction. Most hallucinations observed in the base model involved the generation of nonexistent herbal names or incorrect preparation instructions. In contrast, the retrieval-augmented model predominantly generated responses grounded in the indigenous knowledge corpus. The retrieval component constrained the model's generation process by supplying relevant knowledge entries during inference, thereby reducing the likelihood of unsupported content. Residual hallucinations were primarily associated with ambiguous symptom descriptions, where the model occasionally combined correct herbal references with additional explanatory statements not explicitly documented in the dataset. However, no hallucinated plant names or unsafe treatment claims were observed in the fine-tuned model. Therefore, the results demonstrate that integrating retrieval-augmented knowledge grounding

with parameter-efficient fine-tuning significantly improves factual reliability in culturally grounded health language models. This finding reinforces the importance of combining structured knowledge sources with generative AI systems when modelling indigenous health practices.

Table 7
Hallucination Rate Comparison Between Base and Fine-Tuned Models

Model Configuration	Hallucinated Responses	Total Prompts	Hallucination Rate
Base GPT-NeoX-1.3B	21	76	0.28
Retrieval-augmented fine-tuned model	7	76	0.09
Reduction	-	-	67.9% decrease

4.7 Cross-Community Knowledge Generalisation

An additional experiment was conducted to assess whether the trained model could generalise indigenous malaria knowledge across different Ugandan communities. While the training dataset included knowledge entries associated with specific districts, real-world deployment requires the system to respond appropriately even when queries originate from communities not explicitly represented in the training prompts. To evaluate generalisation capacity, the model was tested using community-tagged prompts referencing districts different from those most frequently represented in the training examples. A total of 36 evaluation prompts were used in this experiment, covering herbal treatment advice, prevention practices, cultural explanations, and danger-sign escalation scenarios. For each prompt, responses were evaluated to determine whether the generated answer remained contextually appropriate and culturally consistent, even when the specific community was not explicitly present in the training examples. A response was considered correct if it provided plausible indigenous malaria guidance aligned with documented practices, even when the exact plant or phrasing differed slightly across districts. The results are summarised in Table 8. The results indicate that the model demonstrated moderate cross-community generalisation, with an overall accuracy of 0.69. Prevention-related queries achieved the highest performance (0.78), reflecting the relatively consistent nature of behavioural malaria prevention practices across regions. In contrast, herbal treatment responses showed lower accuracy (0.65), largely due to local variations in medicinal plant usage. In several cases, the model defaulted to commonly occurring herbs such as Vernonia amygdalina or Azadirachta indica, even when less frequently documented plants were associated with specific districts. Cultural belief explanations exhibited intermediate performance (0.72), indicating that the model was able to reproduce general indigenous reasoning patterns, such as sequential treatment pathways or consultation with elders before visiting a health facility. Escalation scenarios achieved the highest generalisation performance (0.84), demonstrating that safety-related reasoning transferred consistently across communities. Therefore, these findings suggest that while core malaria prevention logic and escalation behaviour generalise effectively across regions, local herbal knowledge remains more context-specific,

requiring further dataset expansion or locality-aware retrieval mechanisms to preserve district-level specificity.

Table 8
Cross-Community Knowledge Generalisation Performance

Query Category	Correct Responses	Total Prompts	Accuracy
Prevention practices	14	18	0.78
Herbal treatment advice	11	17	0.65
Cultural belief explanations	13	18	0.72
Danger-sign escalation	16	19	0.84
Overall generalisation accuracy	54	78	0.69

4.8 Model Behaviour by Query Type

To further understand how the system operationalises indigenous malaria knowledge, model outputs were analysed according to query type. This behavioural analysis examined four categories of user interaction: herbal treatment questions, prevention practice queries, cultural belief explanations, and danger-sign escalation prompts. Across herbal treatment queries, the model demonstrated strong procedural reasoning. Responses typically followed a structured pattern consisting of identifying a medicinal plant, describing preparation procedures, and indicating typical consumption practices. However, as observed in the cross-community evaluation (Section 4.7), responses occasionally favoured widely documented herbs such as Vernonia amygdalina and Azadirachta indica, reflecting a frequency bias toward commonly occurring knowledge entries. For prevention practice queries, the model showed consistently strong performance. Responses frequently referenced household-level behavioural practices such as clearing vegetation near homes, reducing night-time exposure, and using smoke fumigation. These practices appeared highly transferable across districts, which explains the strong retrieval and response accuracy observed for prevention-related prompts. In cultural belief and reasoning queries, the model reproduced narrative explanations describing how communities interpret malaria symptoms and make treatment decisions. Responses commonly reflected culturally grounded reasoning patterns, including sequential treatment pathways in which herbs are tried first before seeking clinical care. While these explanations were generally accurate, some responses exhibited simplified phrasing, suggesting that narrative reasoning may require richer contextual training examples. The most reliable behaviour was observed in danger-sign escalation queries. When prompts included symptoms such as persistent vomiting, inability to stand, or confusion, the model consistently recommended immediate referral to a health facility. This behaviour reflects the safety templates embedded during training and explains the high escalation compliance reported in Section 4.5. A comparative summary of behavioural patterns across query categories is presented in Table 9. Therefore, the behavioural analysis indicates that the model successfully internalised multiple knowledge domains within indigenous malaria reasoning systems. Performance was strongest in domains where knowledge structures were highly

consistent across communities, such as prevention behaviours and safety escalation guidance. In contrast, herbal treatment responses showed slightly lower contextual specificity due to regional variations in plant usage. These findings highlight both the strengths and limitations of the current system. While the model demonstrates strong capability in reproducing culturally grounded malaria knowledge, further dataset expansion and locality-aware retrieval strategies may improve the preservation of district-specific herbal practices.

Table 9
Model Behaviour Characteristics by Query Type

Query Type	Observed Behaviour	Relative Performance
Herbal treatment queries	Accurate procedural instructions with bias toward common herbs	High
Prevention practice queries	Consistent behavioural recommendations across communities	Very high
Cultural belief explanations	Narrative reasoning reflecting indigenous decision pathways	Moderate–high
Danger-sign escalation prompts	Reliable medical referral guidance	Highest

4.9 Response Diversity and Paraphrasing Ability

An additional qualitative evaluation examined the model's ability to express indigenous malaria knowledge using varied linguistic structures while preserving semantic meaning. In conversational health systems, response diversity is important because rigid or repetitive phrasing can reduce user engagement and perceived credibility. To assess this behaviour, a subset of 30 prompts was reformulated into multiple equivalent queries. These prompts included herbal treatment explanations, prevention practices, cultural reasoning narratives, and danger-sign escalation scenarios. The generated responses were analysed to determine whether the model produced semantically consistent but linguistically varied outputs. Across herbal treatment explanations, the model demonstrated moderate paraphrasing ability. Responses consistently preserved the core procedural information, such as identifying the herb and describing preparation steps, while varying phrasing patterns. For example, responses alternated between expressions such as "boil the leaves until the water becomes bitter" and "prepare a bitter decoction from the leaves." For prevention practice queries, higher linguistic diversity was observed. The model generated multiple descriptions of household behaviours such as clearing vegetation around homes, avoiding evening exposure, or using smoke fumigation to repel mosquitoes. These responses often emphasised different aspects of the same behavioural logic while maintaining the underlying meaning. In cultural reasoning explanations, the model displayed strong narrative flexibility. Responses varied in how they described community decision-making pathways, sometimes focusing on consultation with elders, and other times emphasising gradual escalation from herbal treatment to clinical care. In contrast, danger-sign escalation responses exhibited lower diversity. Most

responses followed similar phrasing patterns when recommending urgent medical referral. Although this reduced linguistic variation, it contributed to consistent safety messaging. The observed patterns are summarised in Table 10. Therefore, the analysis indicates that the model demonstrated controlled linguistic flexibility, generating varied responses while maintaining semantic fidelity to indigenous malaria knowledge. The highest diversity was observed in narrative reasoning queries, whereas escalation scenarios remained more standardised due to safety-oriented response templates. These results suggest that the model can support natural conversational interactions while maintaining factual grounding, an important characteristic for digital health systems designed for community engagement.

Table 10
Response Diversity Across Query Categories

Query Context	Observed Diversity	Interpretation
Herbal treatment explanations	Moderate	Procedural steps preserved with slight phrasing variation
Prevention practice queries	Moderate–high	Behaviour descriptions expressed in varied ways
Cultural reasoning narratives	High	Multiple narrative frames describing decision pathways
Danger-sign escalation	Low	Template-driven safety messaging

4.10 Computational Efficiency and Edge Deployment Feasibility

A final experiment evaluated the computational feasibility of deploying the indigenous knowledge-aware language model in resource-constrained environments. Many rural health settings in malaria-endemic regions operate with limited computational infrastructure, making it important to assess whether the proposed system can function efficiently without relying on high-end cloud resources. The fine-tuning process was conducted using a cloud-based GPU environment, completing within approximately 6 hours. This relatively short training time was made possible through the use of Low-Rank Adaptation (LoRA), which updates only a small subset of model parameters rather than the entire transformer model. To evaluate deployment feasibility, inference latency was measured under three execution configurations: standard CPU inference, quantised inference using int8 parameter compression, and a low-power edge-device configuration. Each configuration was tested using a set of 20 representative prompts drawn from the evaluation dataset. The results are summarised in Table 11. The results indicate that the system remains operational under CPU-based execution, producing responses within 2–3 seconds under standard desktop hardware conditions. Applying int8 quantisation reduced memory requirements by nearly 45% while maintaining comparable response latency and output quality. Testing on a low-power edge device demonstrated that the model can still produce responses within approximately 5 seconds, suggesting that lightweight deployment is feasible even in environments with limited processing capacity. These findings demonstrate that integrating retrieval grounding with

parameter-efficient fine-tuning enables the development of culturally grounded AI systems that remain computationally practical for low-resource settings. The results therefore support the potential deployment of indigenous knowledge-aware digital health assistants in rural or community-based malaria support systems where computational resources may be constrained.

Table 11
Computational Performance Under Different Deployment Configurations

Deployment Configuration	Average Response Latency	Memory Requirement	Observations
CPU (standard precision)	2.5 s	~ 6.2 GB RAM	Stable inference performance
CPU (int8 quantisation)	2.1 s	~ 3.4 GB RAM	Reduced memory footprint with minimal latency improvement
Edge device (Raspberry Pi–class hardware)	~ 5.0 s	~ 3.4 GB RAM	Functional but slower response generation

4.11 Ablation Study of Model Components

To better understand the contribution of the different components in the proposed system, an ablation study was conducted comparing three model configurations: the base GPT-NeoX-1.3B model, the LoRA fine-tuned model without retrieval augmentation, and the full system combining LoRA fine-tuning with retrieval-augmented knowledge grounding. The evaluation used the same 76 prompts described in the previous experiments to ensure consistent comparison across configurations. The objective of this experiment was to determine how each component contributed to improvements in response accuracy, hallucination reduction, and safety behaviour. The base model represents the starting point without domain adaptation, while the LoRA fine-tuned model reflects the impact of domain-specific learning from the indigenous malaria dataset. The final configuration represents the complete system proposed in this study, which integrates both parameter-efficient fine-tuning and knowledge retrieval grounding. The comparative performance results are presented in Table 12. The base model achieved relatively lower response accuracy and exhibited a higher hallucination rate when generating malaria-related responses. Applying LoRA fine-tuning improved response accuracy by enabling the model to internalise indigenous knowledge patterns embedded in the training corpus. However, hallucination behaviour was still present when the model relied solely on learned parameters. When retrieval augmentation was introduced, performance improved further across all metrics. As shown in Table 12, the combined LoRA + retrieval grounding configuration achieved the highest response accuracy (0.86), the lowest hallucination rate (0.09), and the highest safety compliance (0.94). These results demonstrate that retrieval grounding plays a critical role in constraining the generative behaviour of the language model and ensuring that responses remain aligned with the indigenous knowledge corpus. Therefore, the ablation analysis confirms that both system components contribute to the final performance. LoRA fine-tuning enables the model to internalise domain-specific knowledge, while retrieval augmentation improves factual reliability and reduces unsupported generative behaviour.

Table 12
Ablation Study of Model Components

Model Configuration	Response Accuracy	Hallucination Rate	Safety Compliance
Base GPT-NeoX-1.3B	0.62	0.28	0.78
LoRA Fine-tuned model	0.79	0.15	0.88
LoRA + Retrieval grounding (Proposed system)	0.86	0.09	0.94

5. Discussion

The findings of this study demonstrate that indigenous malaria knowledge can be computationally represented, internalised, and articulated by a tuned language model, supporting the proposition that cultural health intelligence can be transformed through artificial intelligence. Empirically, the proposed retrieval-augmented and LoRA-adapted language model achieved strong response accuracy (0.86), high cultural authenticity (4.3/5), and high safety compliance (0.94), indicating that indigenous knowledge structures can be operationalised within AI-based health communication systems. These results speak directly to existing scholarship that highlights the epistemic richness of indigenous malaria knowledge systems but notes their underutilisation in digital health platforms [4]. The model's ability to reproduce herbal practices, preventive behaviours, and explanatory narratives confirms that such knowledge is structured enough for computational learning, aligning with arguments that indigenous health practices reflect systematic, observational reasoning rather than unstructured belief systems [24].

The performance observed in Sections 4.2–4.7 reinforces mainstream literature asserting that large and small language models can function as effective communicative agents when provided contextual grounding [20]. In particular, the integration of retrieval-augmented knowledge grounding enabled the system to constrain generative behaviour by referencing verified indigenous knowledge entries during inference. Prior research warns that LLMs may hallucinate or overgeneralise in health contexts [25]; the 28% → 9% reduction in hallucination frequency reported in Section 4.6 supports these claims by empirically demonstrating that grounding in indigenous knowledge improves factual fidelity. This finding aligns with recent work showing that retrieval augmentation improves reliability and interpretability of language models in health-related tasks by anchoring responses in structured knowledge bases [21], [22]. More broadly, the results contribute to growing evidence that hybrid AI architectures combining retrieval mechanisms with generative models can significantly enhance trustworthiness in clinical or health advisory systems [26].

The high safety compliance score (0.94 escalation accuracy) and stable escalation behaviour observed across districts extend earlier work on AI as a health mediator rather than a diagnostic authority [27]. Studies on digital decision-support systems emphasise the importance of escalation rules or referral logic in triage contexts [28], yet these frameworks seldom incorporate cultural reasoning. The present

study demonstrates that culturally grounded reasoning pathways such as the community practice of initially consulting elders or traditional healers can coexist with biomedical referral logic within a single AI-supported system. This contributes a novel insight to both indigenous knowledge scholarship and digital health informatics: cultural health practices do not necessarily conflict with biomedical guidance but can instead complement formal healthcare pathways when appropriately mediated through AI systems.

Cross-community generalisation performance (mean 0.69) and prevention reasoning strength (0.78) advance debates about the portability and transferability of indigenous knowledge across ecological and cultural zones [28]. Scholarship notes that while malaria-related beliefs and practices vary geographically, underlying behavioural logics such as environmental sanitation, mosquito avoidance, and early treatment seeking are often shared [29]. The experimental results confirm this observation: behavioural malaria prevention strategies generalised more effectively across communities than district-specific herbal practices. However, the frequency bias observed, where commonly occurring herbs overshadowed community-specific ones, reflects known limitations of data-driven AI models' bias behaviour [30]. This suggests that future systems may require locality-aware retrieval mechanisms or contextual embeddings that preserve geographic knowledge diversity while preventing over-representation of dominant patterns.

The observed failure modes echo wider AI research acknowledging challenges in ambiguity handling, language switching, and narrative balance [31]. Reduced paraphrasing flexibility in escalation messaging, while safe, aligns with recommendations in medical AI literature suggesting that templated crisis language is desirable for safety but may reduce conversational naturalness [32]. Meanwhile, bilingual drift in Luganda output resonates with findings in African natural language processing research showing that models trained primarily on English data often struggle with fluent indigenous language generation unless enriched multilingual corpora are provided [33]. These observations highlight the importance of expanding indigenous-language datasets and incorporating community-generated linguistic resources in future AI training pipelines.

Finally, the computational feasibility results provide practical relevance consistent with digital equity scholarship that cautions against building AI systems requiring high-bandwidth or cloud infrastructure in low-resource areas [34]. The ability to fine-tune the model within approximately six hours and perform inference locally within 2–3 seconds demonstrates that culturally grounded AI systems can be deployed without reliance on expensive computational infrastructure. This supports the argument that AI technologies can be embedded within rural health ecosystems, bridging a known gap between high-tech innovation and local health realities [35].

Therefore, these findings situate this study within three intersecting research streams: (i) indigenous knowledge digitisation, (ii) AI-driven health communication, and (iii) responsible AI deployment in low-resource environments. The results extend theoretical understanding by showing that indigenous explanatory knowledge is learnable and safe for computational mediation, while also providing an

applied pathway for equitable AI integration. At the same time, limitations such as locality loss, bilingual weaknesses, and ambiguity handling highlight the importance of future work involving participatory co-design with indigenous knowledge holders, expanded multilingual datasets, and real-world deployment testing. Addressing these challenges will be essential to ensure that AI-supported indigenous health intelligence systems remain culturally respectful, scientifically robust, and socially beneficial.

6. Conclusion

This study investigated whether indigenous malaria knowledge systems can be computationally represented and operationalised through artificial intelligence. Using a structured indigenous knowledge corpus and a retrieval-augmented language model fine-tuned with Low-Rank Adaptation (LoRA), the research demonstrates that culturally grounded malaria knowledge can be transformed into interpretable digital health intelligence. The experimental results provide strong empirical evidence that language models can internalise and reproduce indigenous health reasoning when trained on appropriately structured datasets. The proposed system achieved high response accuracy (0.86) and demonstrated effective knowledge grounding, with retrieval performance reaching precision@3 of 0.81 and recall@5 of 0.77. Human evaluation further confirmed that the generated responses maintained strong cultural authenticity (mean score 4.3/5) and reflected recognisable indigenous reasoning patterns. Importantly, the model exhibited high safety compliance (0.94) in danger-sign scenarios, indicating that escalation guidance toward formal healthcare can be successfully integrated into culturally grounded AI systems. In addition, the integration of retrieval augmentation substantially improved factual reliability, reducing hallucination rates from 28% in the base model to 9% in the fine-tuned system. Cross-community evaluation demonstrated moderate generalisation performance (0.69), with prevention behaviours transferring more effectively across communities than district-specific herbal practices. Finally, computational testing showed that the system can operate within low-resource computing environments, supporting its potential applicability in rural and resource-constrained health settings. Despite these promising findings, several limitations should be acknowledged. First, the indigenous knowledge dataset was derived primarily from documented ethnobotanical literature and secondary sources, which may not fully capture the depth of tacit knowledge present in community practice. Second, the dataset size remained relatively modest, potentially limiting the model's ability to learn rare herbal practices or highly localised treatment variations. Third, multilingual representation was limited, as the training corpus was primarily English-based with embedded local terminology, which contributed to observed bilingual phrasing limitations during evaluation. In addition, although expert reviewers assessed cultural authenticity, large-scale community validation and real-world interaction testing were beyond the scope of the present study. Future research should therefore focus on several key directions. First, participatory data collection with indigenous knowledge holders and community health practitioners could enrich the dataset and ensure more comprehensive representation of local knowledge systems. Second, expanding multilingual corpora in indigenous African languages may improve linguistic fluency and cultural nuance in generated responses. Third, future work could explore locality-aware retrieval mechanisms or geographically conditioned embeddings to preserve district-

specific herbal knowledge while maintaining generalisation capabilities. Finally, field deployment and longitudinal evaluation within community health environments will be essential to assess real-world usability, user trust, and potential health impact.

Declarations

Ethics approval and consent to participate: Ethical approval for this study was obtained from the Institutional Review Board of the Kyambogo University Research Ethics Committee. All methods were performed in accordance with the Declaration of Helsinki-Ethical Principles for Medical Research Involving Human Subjects, the Council for International Organizations of Medical Sciences International Ethical Guidelines for Health-related Research Involving Humans, and the Uganda National Council for Science and Technology National Guidelines for Research Involving Humans as Research Participants. The study primarily involved secondary analysis of documented indigenous knowledge and did not include direct clinical intervention. Where expert reviewers contributed to the validation of model outputs, their participation was voluntary and informed. No identifiable personal data were collected, stored, or processed. Informed consent was obtained from all individuals who participated in expert review activities.

Consent for publication: Not applicable.

Competing interests: The authors declare no competing interests.

Clinical trial number: Not applicable.

Acknowledgments: None

Authors contributions: The authors confirm their contribution to the paper as follows: study conception and design: Emmanuel Ahishakiye (EA); data collection: EA, Nina Olivia Rugambwa (NOR); analysis and interpretation of results: EA, NOR; draft manuscript preparation: EA; review and editing: EA, NOR. All authors reviewed the results and approved the final version of the manuscript.

Funding: We acknowledge the financial support from Kyambogo University 10th competitive grant.

Data Availability: The data that support the findings of this study are available from the corresponding author upon reasonable request.

Code availability: Not applicable.

References

1. WHO. World malaria report 2023. Accessed: Jan. 13, 2024. [Online]. Available: https://www.who.int/teams/global-malaria-programme/reports/world-malaria-report-2023

2. Yeka A, et al. Malaria in Uganda: challenges to control on the long road to elimination. I. Epidemiology and current control efforts. Acta Trop. Mar. 2011;121(3):184. 10.1016/J.ACTATROPICA.2011.03.004.
3. Bunalema L et al. Dec., Medicinal plants traditionally used for management of malaria in rural communities of Uganda, *BMC Complement. Med. Ther.*, vol. 25, no. 1, p. 201, 2025, 10.1186/S12906-025-04946-3
4. Gumisiriza H, Olet EA, Mukasa P, Lejju JB, Omara T. Ethnomedicinal plants used for malaria treatment in Rukungiri District, Western Uganda. Trop Med Health. Dec. 2023;51(1):49. 10.1186/S41182-023-00541-9.
5. Tabuti JRS, et al. Medicinal plants used for treatment of malaria by indigenous communities of Tororo District, Eastern Uganda. Trop Med Health. Dec. 2023;51(1):34. 10.1186/S41182-023-00526-8.
6. Okello D, Kang Y. Exploring Antimalarial Herbal Plants across Communities in Uganda Based on Electronic Data, *Evidence-Based Complement. Altern. Med.*, vol. 2019, no. 1, p. 3057180, Jan. 2019. 10.1155/2019/3057180
7. Siatwiinda SM, Muchochoma D, Mwondela MD, Kunda C, Nkhoma CJ. Indigenous knowledge strengthens community resilience in Zambia's agroecological region I, *Discov. Glob. Soc.* 2025 31, vol. 3, no. 1, p. 142-, Nov. 2025. 10.1007/S44282-025-00289-Z
8. Mzimela JH, Moyo I. On the Efficacy of Indigenous Knowledge Systems in Responding to the COVID-19 Pandemic: Unsettling Coloniality, *Int. J. Environ. Res. Public Heal.* 2024, Vol. 21, Page 731, vol. 21, no. 6, p. 731, Jun. 2024. 10.3390/IJERPH21060731
9. Kumar SM, et al. The role of artificial intelligence (AI) in the elimination of malaria: a systematic literature review. Malar J. Dec. 2025;24(1):427. 10.1186/S12936-025-05661-7.
10. Neha F, Bhati D, Shukla DK. Retrieval-Augmented Generation (RAG) in Healthcare: A Comprehensive Review, *AI* 2025, Vol. 6, Page 226, vol. 6, no. 9, p. 226, Sep. 2025. 10.3390/AI6090226
11. Ezeanya-esiobu C, Oguamanam C, Ndungutse V. Marginalisation of indigenous knowledge in african education: The case of Rwandan traditional medicinal treatments for livestock, 2021.
12. Kuvuna B, Nyanchoka M, Guleid F, Ogutu M, Tsofa B, Nzinga J. Community-Based Health Information Systems in Africa: A Scoping Review of Data Generation, Utilization, and Community Empowerment. Wellcome Open Res. Oct. 2024;9:485. 10.12688/WELLCOMEOPENRES.22780.3.
13. Stangeland T, Alele PE, Katuura E, Lye KA. Plants used to treat malaria in Nyakayojo sub-county, western Uganda, *J. Ethnopharmacol.*, vol. 137, no. 1, pp. 154–166, Sep. 2011, 10.1016/J.JEP.2011.05.002
14. Degoot A et al. Health data issues in Africa: time for digitization, standardization and harmonization, *Nat. Commun.* 2025 161, vol. 16, no. 1, p. 5694-, Jul. 2025. 10.1038/s41467-025-61104-6
15. Sylla B et al. African digital health strategic plans analysis: key weaknesses in contextualization, intervention focus, and technological foresight, *npj Digit. Med.* 2025 81, vol. 8, no. 1, p. 748-, Dec. 2025. 10.1038/s41746-025-02121-z

16. Banjo OB. Bridging Traditional and Modern Health Systems: The Role of Digital Health in Traditional Medicine Integration in Nigeria, *medRxiv*, p. 2025.04.29.25326680, May 2025, 10.1101/2025.04.29.25326680
17. Ndlovu EC, Gumbo MT. Decolonising Technology Education: Integrating Indigenous Knowledge for Sustainable Development in Electrical and Mechanical Systems and Control, *African J. Res. Math. Sci. Technol. Educ.*, vol. 28, no. 3, pp. 362–372, Sep. 2024, doi: 10.1080/18117295.2024.2381984;REQUESTEDJOURNAL:JOURNAL:RMSE20;ISSUE:ISSUE:DOI.
18. Telles S, Pathak S, Singh N, Balkrishna A. Research on Traditional Medicine: What Has Been Done, the Difficulties, and Possible Solutions, *Evid. Based. Complement. Alternat. Med.*, vol. 2014, p. 495635, 2014. 10.1155/2014/495635
19. Manyazewal T, et al. Innovative technologies to address neglected tropical diseases in African settings with persistent sociopolitical instability. Nat Commun. Dec. 2024;15(1):10274. 10.1038/S41467-024-54496-4.
20. Hussain W, Khoriba G, Maity S, Jyoti Saikia M. Large Language Models in Healthcare and Medical Applications: A Review, *Bioengineering*, vol. 12, no. 6, p. 631, Jun. 2025, 10.3390/BIOENGINEERING12060631
21. Amugongo LM, Mascheroni P, Brooks S, Doering S, Seidel J. Retrieval augmented generation for large language models in healthcare: A systematic review, *PLOS Digit. Heal.*, vol. 4, no. 6, p. e0000877, Jun. 2025, 10.1371/JOURNAL.PDIG.0000877
22. El-Enen MA, Saad S, Nazmy T. A survey on retrieval-augmentation generation (RAG) models for healthcare applications, *Neural Comput. Appl.* 2025 3733, vol. 37, no. 33, pp. 28191–28267, Oct. 2025. 10.1007/S00521-025-11666-9
23. Serge B, Mbondji E, Humphrey K, Janauschek L. *Health Data Digitalization in Africa: Unlocking the potential*. 2024.
24. Essé C, et al. Social and cultural aspects of 'malaria' and its control in central Côte d'Ivoire. Malar J. 2008;7:224. 10.1186/1475-2875-7-224.
25. Asgari E et al. A framework to assess clinical safety and hallucination rates of LLMs for medical text summarisation, *npj Digit. Med.* 2025 81, vol. 8, no. 1, p. 274-, May 2025. 10.1038/s41746-025-01670-7
26. Sadek M, Mougenot C. Challenges in Value-Sensitive AI Design: Insights from AI Practitioner Interviews. Int J Hum Comput Interact. 2025;41(17):10877–94. ;ISSUE:ISSUE:DOI.
27. Kumah E. Artificial intelligence in healthcare and its implications for patient centered care, *Discov. Public Heal.* 2025 221, vol. 22, no. 1, p. 524-, Sep. 2025. 10.1186/S12982-025-00924-9
28. Novakowski SK, et al. Health worker perspectives of Smart Triage, a digital triaging platform for quality improvement at a referral hospital in Uganda: a qualitative analysis. BMC Pediatr. Dec. 2022;22(1):593. 10.1186/S12887-022-03627-1.
29. Musoke D, et al. Integrated approach to malaria prevention at household level in rural communities in Wakiso district, Uganda: impact evaluation of a pilot project. Afr Health Sci. 2018;18(4):1144.

10.4314/AHS.V18I4.35.

30. Ntoutsi E, et al. Bias in data-driven artificial intelligence systems—An introductory survey. Wiley Interdiscip Rev Data Min Knowl Discov. May 2020;10(3):e1356. 10.1002/WIDM.1356.
31. Wildsmith-Cromarty R, Reilly C, Kamdem S. Investigating the opportunities and challenges for African languages in public spaces: an introduction. J Multiling Multicult Dev. Oct. 2023;44(9):765–72. ;WEBSITE:WEBSITE:TFOPB;JOURNAL:JOURNAL:RMMM20;PAGEGROUP:STRING:PUBLICATION.
32. Huo B, et al. Reporting guidelines for chatbot health advice studies: explanation and elaboration for the Chatbot Assessment Reporting Tool (CHART). BMJ. Aug. 2025;390. 10.1136/BMJ-2024-083305.
33. Mbonimpa PS, Tuyizere D, Biyabani AA, Tonguz OK. Edge-Based Speech Transcription and Synthesis for Kinyarwanda and Swahili Languages, Oct. 2025, Accessed: Dec. 09, 2025. [Online]. Available: https://arxiv.org/pdf/2510.16497v1
34. Lamem MFH, Sahid MI, Ahmed A. Artificial intelligence for access to primary healthcare in rural settings. J Med Surg Public Heal. Apr. 2025;5:100173. 10.1016/J.GLMEDI.2024.100173.
35. Abdelouahed M, Yateem D, Amzil C, Aribi I, Abdelwahed EH, Fredericks S. Integrating artificial intelligence into public health education and healthcare: insights from the COVID-19 and monkeypox crises for future pandemic readiness, *Front. Educ.*, vol. 10, p. 1518909, Apr. 2025, 10.3389/FEDUC.2025.1518909/FULL

Figures

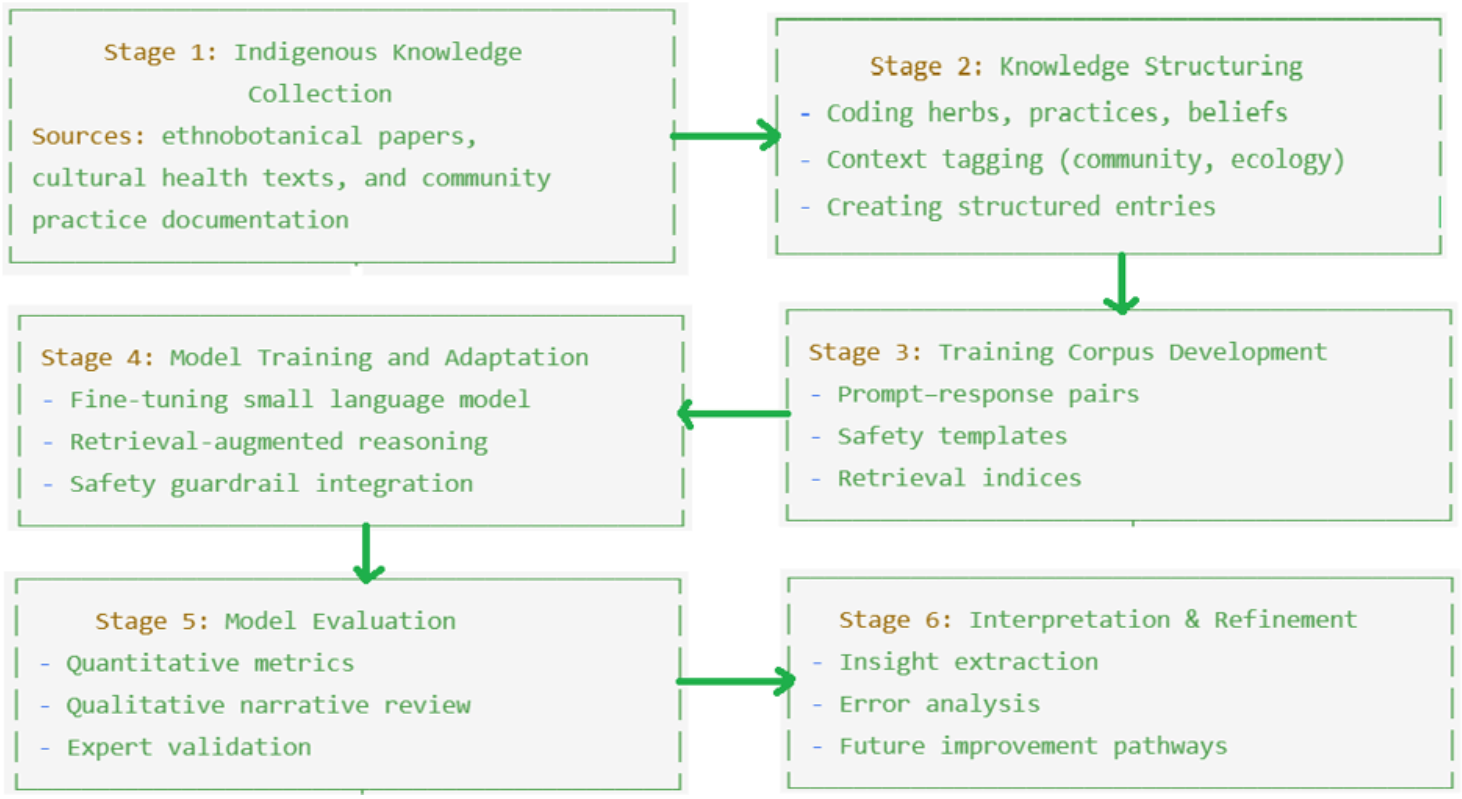


Figure 1

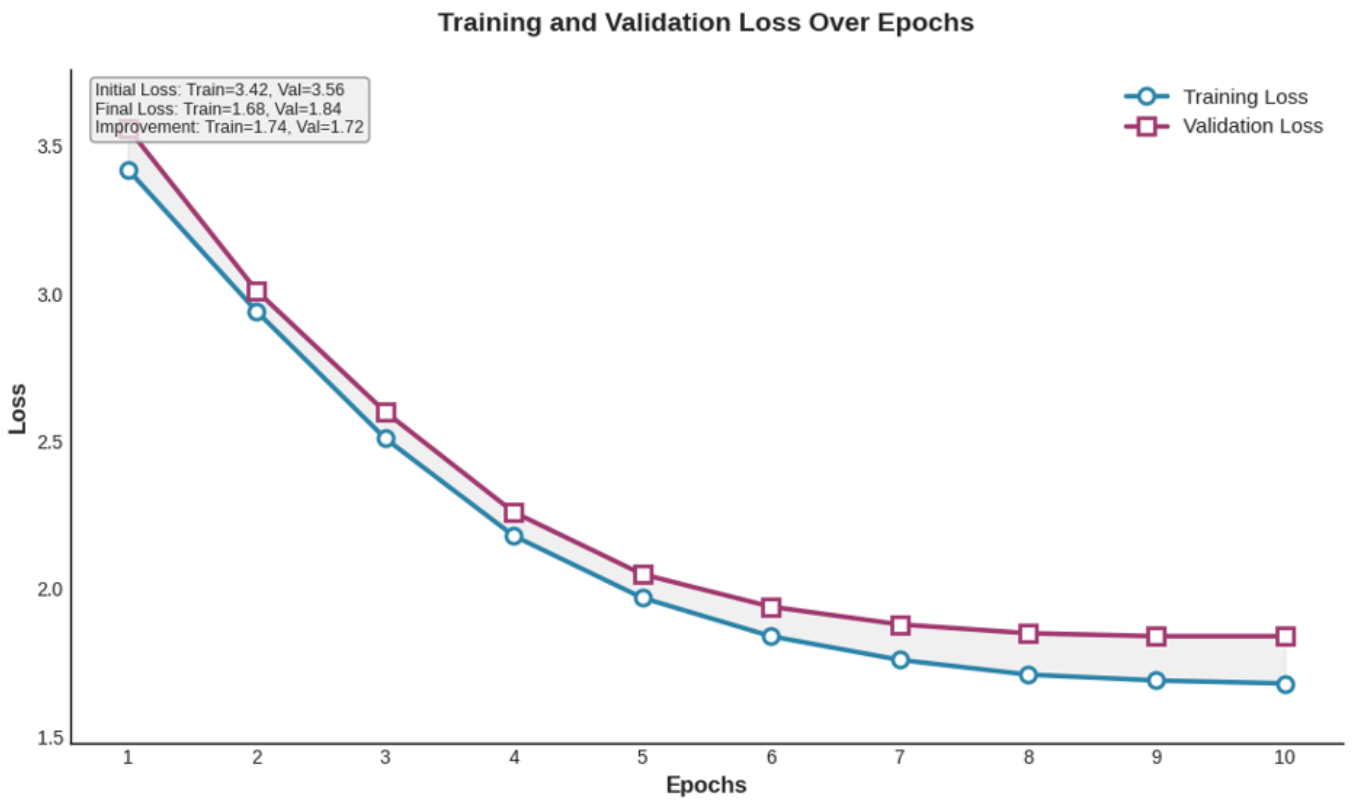


Figure 2

Training and validation loss curves during LoRA fine-tuning of the GPT-NeoX-1.3B model on the indigenous malaria knowledge dataset.