



An AI-Powered, Multilingual Pedagogical Support Framework for Integrated Science Teacher Education in Resource-Constrained Nigerian Contexts

ISMAIL ADO FUNTUA, MAHMUD MUHAMMAD YAHAYA
ALIYU DAHIRU KAROFI
Isa Kaita College of Education, Dutsin-Ma, Katsina State, Nigeria

Abstract. Integrated Science teacher education is central to developing scientifically literate learners capable of understanding and applying concepts across biology, chemistry, physics, and environmental science. However, teacher education programmes in many Nigerian contexts face persistent challenges, including limited instructional resources, inadequate pedagogical support, linguistic diversity, and uneven access to digital technologies. This paper presents the architecture and Python reference implementation of an adaptive, multilingual AI-powered pedagogical support framework for Integrated Science teacher education, designed for deployment in Nigerian colleges of education. The framework organizes classroom interaction as a three-stage input, processing and output pipeline. A Bloom's-Taxonomy-tagged instructional corpus feeds a Low-Rank Adaptation (LoRA)-fine-tuned LLaMA language model, teacher voice queries in Hausa are transcribed with a fine-tuned XLSR-Wav2Vec2 speech recognizer and translated for processing, and generated instructional content is delivered back to teachers as text and synthesized Hausa speech. Through voice-based and local-language interaction, the system promotes inclusive and context-responsive science teaching while embedding ethical safeguards. We situate this design against recent advances in African-language natural language processing and low-resource speech recognition, provide worked code for every pipeline stage, and outline an evaluation plan combining automatic speech-recognition and translation metrics with teacher-facing usability assessment across partner colleges of education. The framework is offered as a reusable, extensible architecture for AI-supported, multilingual teacher education.

Keywords: Integrated Science teacher education; artificial intelligence in education; low-resource natural language processing; Hausa speech recognition; Bloom's Taxonomy; multilingual pedagogy; Africa

1. Introduction

Integrated Science is designed to provide learners with a coherent understanding of scientific concepts by linking ideas across biology, chemistry, physics, and environmental science (Yakubu, 2024; Zhai et al., 2021). Effective teaching of Integrated Science requires teachers who possess strong content knowledge, the ability to connect concepts across disciplines, and the pedagogical skills to facilitate inquiry-based and learner-centred instruction (Holmes et al., 2024). Teacher education institutions therefore play a critical role in preparing pre-service teachers to meet these demands.

In Nigeria, Integrated Science teacher education faces systemic challenges. These include shortages of qualified instructors, limited access to up-to-date teaching materials, overcrowded classrooms, and the use of languages of instruction that may not align with learners' linguistic backgrounds (UNESCO, 2023). As a result, pre-service teachers often struggle to develop confidence in explaining abstract scientific concepts, integrating multiple science domains, and engaging learners effectively (Cukurova, Luckin, & Kent, 2020).

Recent advances in large language models (LLMs) and speech technologies have begun to narrow this gap. Large-scale multilingual evaluations show that even the strongest open and commercial LLMs still under-perform on African languages when compared

to English (Ojo et al., 2024), and dedicated adaptation efforts such as Luga-Llama have only recently started to bridge this divide for languages including Hausa, Yoruba, Swahili, and Zulu by combining curated African-language text with high-quality English educational material during continual pretraining (Buzaaba, Wettig, Adelani, & Fellbaum, 2025). On the speech side, self-supervised architectures such as wav2vec 2.0 (Baeviski, Zhou, Mohamed, & Auli, 2020) and its cross-lingual variant XLSR (Conneau, Baeviski, Collobert, Mohamed, & Auli, 2021) have made low-resource automatic speech recognition (ASR) increasingly viable, and recent work fine-tuning these models on Hausa corpora, augmented with LLM-generated synthetic speech, has pushed word-error rates for African languages into a practically usable range (DeRenzi, Dixon, Farhi, & Resch, 2025).

Artificial Intelligence (AI) systems can provide adaptive instructional support, generate pedagogically relevant content, and facilitate interactive learning experiences (Zawacki-Richter, Marín, Bond, & Gouverneur, 2019; Bonner, Lege, & Frazier, 2023). When aligned with science education principles and local classroom realities, such technologies can complement traditional teacher training rather than replace it.

This study adopts an interdisciplinary research approach to explore how AI can be used as a pedagogical support tool for Integrated Science teacher education. By combining insights from science education, cognitive learning theories, linguistics, and computer science, the study proposes a framework that addresses both instructional and contextual challenges within Nigerian teacher education programmes.

The remainder of the article is divided as follows. Section 2 situates the framework within the current literature; Section 3 describes the system architecture and provides a Python reference implementation of the core technical contribution, i.e., LoRA fine-tuning of LLaMA on the Integrated Science corpus, as well as the output stage i.e., **Instructional response generation and Hausa text-to-speech** pipeline stage; Section 4 demonstrates a worked classroom scenario; Section 5 outlines the evaluation plan and expected outcomes; Section 6 addresses ethical considerations and privacy; and Section 7 discusses the implications and limitations of the framework.

2. Related Work

2.1 Integrated Science Pedagogy and Teacher Preparation

Research in science education consistently emphasizes the importance of integrated approaches to teaching science, where learners are supported to make meaningful connections across scientific disciplines rather than studying biology, chemistry, and physics in isolation (Zhai, He, Liu, & Kuo, 2021). Integrated Science teaching has been shown to enhance conceptual understanding, scientific reasoning, and the application of knowledge to real-world problems, particularly when inquiry-based pedagogies are employed (Cukurova et al., 2020).

Studies in teacher education highlight persistent challenges in preparing Integrated Science teachers, especially in low-resource contexts. These challenges include limited exposure to practical teaching strategies, inadequate instructional materials, and insufficient opportunities for reflective practice during training (Yakubu, 2024; UNESCO, 2023). Without adequate pedagogical support, pre-service teachers often struggle to integrate scientific concepts and manage learner-centred classroom activities effectively.

2.2 AI and Large Language Models in Teacher Education

The growing body of research on AI in education demonstrates its potential in intelligent tutoring systems, adaptive learning platforms, and automated feedback mechanisms to support teaching and learning (Zawacki-Richter et al., 2019). LLMs, in particular, have shown promise in generating curriculum-aligned explanations, assessment questions, and instructional prompts that support higher-order thinking (Bonner et al., 2023; Holmes, Porayska-Pomsta, & Holstein, 2024). However, much of this research is concentrated in high-resource contexts and does not explicitly address Integrated Science teacher education for low-resource languages.

Elkins, Kochmar, Cheung, and Serban (2024) show that quiz items generated by LLMs are rated as pedagogically sound by teachers only when the intended Bloom's-Taxonomy level is made explicit in the prompt; when left unconstrained, generated activities tend to cluster at the lower, "remember/understand" end of the taxonomy. This finding directly motivates embedding Bloom's-level metadata into the training corpus and generation prompts of the present framework (Section 3.2), rather than relying on an unconstrained LLM to infer cognitive demand. Luckin, Moss, and Holmes (2022) further argue that teacher-facing AI tools succeed only when paired with structured AI-literacy support for the

teachers who use them, a consideration reflected in the deployment plan in Section 5.

2.3 Low-Resource, Multilingual NLP for African Languages

The performance gap between English and low-resource African languages in LLMs is still very wide. AfroBench, a large-scale evaluation covering a number of African languages and multiple task categories, found that both open models such as LLaMA and Gemma, and commercial systems such as GPT-4o and Gemini perform poorly compared to their English counterparts, with the gap widest for generation tasks and languages using non-Latin orthographies (Ojo et al., 2024). Targeted continual-pretraining efforts are narrowing this gap. For example, Lugha-Llama adapts a LLaMA-3.1 base model (Touvron et al., 2023) to sixteen African languages by mixing curated African-language text with high-quality English educational documents, improving performance on the IrokoBench and AfriQA benchmarks (Buzaaba et al., 2025). Adelani et al. (2022) established an earlier foundation for this line of work, showing that pre-trained multilingual models can be adapted to African languages with only a few thousand parallel sentences, a data setting realistic for a single teacher-education programme to reproduce for domain-specific content.

On the speech side, wav2vec 2.0 (Baevski et al., 2020) and its multilingual variant XLSR (Conneau et al., 2021) remain dominant self-supervised backbones for low-resource ASR; XLSR is adopted in the present framework because cross-lingual pretraining substantially outperforms monolingual training when labelled Hausa speech data is scarce. DeRenzi et al. (2025) recently showed that LLM-generated synthetic text, converted to speech and used to augment ASR fine-tuning, can match real-data-only baselines for Hausa at well under one percent of the cost of collecting equivalent real data, an approach directly applicable to extending this framework's speech pipeline once a Bloom's-tagged Integrated Science text corpus exists (Section 3.2).

2.4 Multilingual Pedagogy and Ethical Considerations

Multilingual education research underscores the importance of language flexibility in science classrooms, especially in African contexts characterized by linguistic diversity. Translanguaging practices, where teachers and learners draw on multiple languages, have been shown to enhance comprehension and learner participation in science lessons (García & Wei, 2014; Makalela, 2018). Advances in multilingual AI and speech technologies offer opportunities to support such practices, yet ethical concerns related to bias, data privacy and cultural relevance remain serious issues (Birhane, 2021; UNESCO, 2023).

Overall, the literature reveals a clear gap in research that combines AI, science education, linguistics, and ethical governance to support Integrated Science teacher education in African contexts. Addressing this gap requires the development of pedagogically grounded, multilingual, and context-aware AI frameworks tailored to the realities of teacher education programmes.

The framework presented below responds directly to this gap by aligning a domain-specific, Bloom's instructional corpus with African-language speech and translation components, and by specifying the reference implementation in Python, for the two core sections of the model, so that other Integrated Science programmes can adapt it to their own languages and curricula.

3. The Proposed Framework

The proposed framework conceptualizes classroom interaction as a structured pipeline consisting of three interrelated stages: input, processing, and output (See Figure 1). The design emphasizes pedagogical grounding, multilingual accessibility, and ethical responsibility within resource-constrained educational contexts.




Figure 1. Three Stage Classroom Interaction Pipeline Diagram (With a cross-cutting ethical and privacy safeguard layer).

3.1 Input Stage: Corpus Curation and Bloom's-Taxonomy Tagging

At the input stage, the foundational intelligence of the system is established through the curation of pedagogically relevant Integrated Science datasets. These include Integrated Science curricula, lesson plans, textbooks, and teacher education guides covering biology, chemistry, physics, and environmental science, sourced from colleges and universities of education in Nigeria. The collected materials are digitized, cleaned, and transformed into a structured XML format to ensure machine readability and compatibility with the Llama LLM ingestion pipeline.

To strengthen instructional relevance, the dataset is enriched with pedagogical metadata, including learning objectives, common student misconceptions, suggested teaching strategies, and inquiry-based activities. Each content unit is annotated with its corresponding Bloom's Taxonomy level. This metadata-driven approach ensures that AI-generated outputs align with intended cognitive demands and support the development of scientific reasoning and conceptual integration central to Integrated Science teaching. In addition to text-based inputs, the framework supports teacher voice input through Speech-to-Text mechanisms, enabling real-time interaction in local languages.

3.2 Processing Stage: Domain Fine-Tuning and Hausa Speech Understanding

The processing stage is based on a transformer-based LLM, specifically the LLaMA architecture, fine-tuned

on Integrated Science educational data to enhance domain relevance and pedagogical accuracy. Prompt-engineering strategies guide the model to generate localized instructional outputs, such as inquiry-oriented questions, explanations, and classroom prompts tailored to Integrated Science topics.

Rather than full-parameter fine-tuning, which is computationally difficult for most Nigerian teacher-training institutions, the reference implementation uses Low-Rank Adaptation (LoRA), which freezes the pretrained LLaMA-2 weights (Touvron et al., 2023) and trains small rank-decomposition update matrices injected into the attention projections (Hu et al., 2022). This reduces the number of trainable parameters by several orders of magnitude and permits fine-tuning on a single consumer-grade GPU, a practical requirement given the infrastructure constraints noted in Section 1. Listing 1 shows the fine-tuning routine, including the instruction-formatting template that conditions each generation on subject strand and Bloom's level.

Listing 1. LoRA fine-tuning of LLaMA on the Integrated Science corpus (Python)

```

from datasets import load_dataset
from transformers import (AutoModelForCausalLM,
                          AutoTokenizer,
                          TrainingArguments, Trainer,
                          DataCollatorForLanguageModeling)
from peft import LoraConfig, get_peft_model, TaskType

```

```

BASE_MODEL = "meta-llama/Llama-2-7b-hf" #
Touvron et al., 2023

```

```

def build_lora_model(base_model_name: str =
BASE_MODEL):

```

```

tokenizer = AutoTokenizer.from_pretrained(base_model_name)
tokenizer.pad_token = tokenizer.eos_token
model = AutoModelForCausalLM.from_pretrained(
    base_model_name, load_in_8bit=True,
    device_map="auto")

lora_cfg = LoraConfig( # Hu et al.,
2022
    r=16, lora_alpha=32, lora_dropout=0.05,
    target_modules=["q_proj", "v_proj"],
    task_type=TaskType.CAUSAL_LM,
)
model = get_peft_model(model, lora_cfg)
model.print_trainable_parameters()
return model, tokenizer

def format_example(example, tokenizer, max_len: int
= 512):
    prompt = (
        f"### Subject: {example['subject_strand']}\n"
        f"### Bloom Level:
{example['bloom_level']}\n"
        f"### Instruction: Generate an inquiry-based
classroom "
        f"prompt on '{example['topic']}'.\n"
        f"### Response: {example['text_en']}"
    )
    return tokenizer(prompt, truncation=True,
max_length=max_len,
padding="max_length")

def fine_tune(corpus_path: str, output_dir: str = "isc-
llama-lora"):
    dataset = load_dataset("json",
data_files=corpus_path)["train"]
    model, tokenizer = build_lora_model()
    tokenized = dataset.map(lambda ex:
format_example(ex, tokenizer))

    args = TrainingArguments(
        output_dir=output_dir,
        per_device_train_batch_size=4,
        gradient_accumulation_steps=4,
        num_train_epochs=3,
        learning_rate=2e-4, fp16=True,
        logging_steps=20,
        save_strategy="epoch",
    )
    collator = DataCollatorForLanguageModeling(tokenizer,
mlm=False)
    trainer = Trainer(model=model, args=args,
train_dataset=tokenized,
data_collator=collator)
    
```

```

trainer.train()
model.save_pretrained(output_dir)
return model, tokenizer
    
```

To support multilingual interaction, teacher voice queries in Hausa are processed using a fine-tuned XLSR-Wav2Vec 2.0 self-supervised learning model for speech recognition. This model is selected for its robustness in low-resource and noisy environments and its ability to perform effective cross-lingual transfer when trained on curated Hausa speech corpora.

The section implements this transcription step, building on the pretrained multilingual checkpoint of Conneau et al. (2021), fine-tuned on Hausa audio. DeRenzi et al. (2025) demonstrate that this checkpoint family can be improved further with modest amounts of LLM-generated synthetic speech, offering a lower-cost path to expanding the training corpus beyond what any single institution could record on its own. The transcribed Hausa input is translated into English for LLM processing, where Natural Language Understanding modules extract pedagogical intent, instructional preferences, and cognitive level requirements. This intent extraction mechanism enables personalization by aligning generated content with specific classroom contexts and instructional goals.

The translation step will adopt a general-purpose multilingual MT model as a placeholder. Adelani et al. (2022) show that adapting such models with only a few thousand in-domain parallel sentences meaningfully improves translation quality for African languages; this is the intended upgrade path once an Integrated-Science-specific Hausa-English parallel corpus is collected during the pilot described in Section 5.

Following content generation, the system translates outputs back into Hausa to maximize accessibility. Although this multi-stage machine translation process may introduce latency and minor translation errors, it represents a pragmatic interim solution pending the deployment of fully end-to-end translanguaging LLMs capable of direct pedagogical generation in African languages.

3.3 Output Stage: Instructional Resource Generation and Speech Synthesis

At the output stage, the system delivers structured instructional resources to pre-service Integrated Science teachers. These outputs include curriculum-aligned questions with model answers, explanatory notes, and adaptive pedagogical recommendations

such as learner grouping strategies, use of visual aids, or inquiry-based classroom activities. Text-to-Speech technologies convert textual outputs into audible formats in both English and Hausa, addressing literacy barriers and supporting flexible classroom use.

Listing 2 implements response generation from the fine-tuned LLM together with Hausa speech synthesis using a massively multilingual TTS checkpoint. The orchestration layer coordinates all modules into a coherent, real-time workflow that is transparent to the end user that realizes the full voice-in/voice-out interaction loop described conceptually above and shown in Figure 1.

Listing 2. Instructional response generation and Hausa text-to-speech (Python)

```
import torch, torchaudio
from transformers import VitsModel, AutoTokenizer
```

```
HAUSA_TTS_CHECKPOINT = "facebook/mms-tts-hau" # Massively Multilingual Speech TTS
```

```
def generate_instructional_response(model, tokenizer,
intent: dict) -> str:
    prompt = (
        f"### Subject: {intent['strand']}\n"
        f"### Bloom Level: {intent['bloom_level']}\n"
        f"### Instruction: {intent['query']}\n###"
    )
    Response:"
    inputs = tokenizer(prompt, return_tensors="pt")
    output_ids = model.generate(**inputs,
max_new_tokens=200, temperature=0.7)
    return tokenizer.decode(output_ids[0],
skip_special_tokens=True)
```

```
class HausaTextToSpeech:
    def __init__(self, checkpoint: str =
HAUSA_TTS_CHECKPOINT):
        self.tokenizer =
AutoTokenizer.from_pretrained(checkpoint)
        self.model =
VitsModel.from_pretrained(checkpoint)

    def synthesize(self, hausu_text: str, out_path: str =
"response.wav") -> str:
        inputs = self.tokenizer(hausu_text,
return_tensors="pt")
        with torch.no_grad():
            waveform = self.model(**inputs).waveform
            torchaudio.save(out_path, waveform,
self.model.config.sampling_rate)
        return out_path
```

4. Illustrative Walkthrough

To ground the architecture in a concrete classroom scenario, consider a pre-service teacher preparing a lesson on the water cycle, a topic that spans physics (evaporation and energy transfer), chemistry (phase change), and environmental science (the hydrological cycle's ecological role), making it a natural test of the framework's conceptual-integration goal. The teacher poses a question aloud in Hausa; Table 1 traces representative inputs and outputs through the pipeline at four Bloom's-Taxonomy levels, illustrating how the same underlying topic yields cognitively differentiated instructional prompts.

Table 1. Example of Contextual Questions & Answers for AI in Education (*Illustrative, hand-authored outputs intended to demonstrate the pipeline design; not yet generated by a running system.*)

Bloom Level	Teacher Query (Hausa)	English Gloss	AI-Generated Instructional Prompt
<i>Understand</i>	Ta yaya ake bayyana yanayin zagayowar ruwa?	How is the water cycle explained?	Ask learners to describe, in their own words, how evaporation and condensation work together to move water between the land, ocean, and atmosphere.
<i>Apply</i>	Ta yaya za a nuna zagayowar ruwa a aji?	How can the water cycle be demonstrated in class?	Guide learners through a simple evaporation-condensation demonstration using a covered basin of warm water, then have them label each stage of the cycle on the resulting condensation.
<i>Analyze</i>	Wace alaka ke tsakanin zafi da yawan tururi?	What is the relationship between heat and rate of evaporation?	Have learners compare evaporation rates of water samples at two temperatures and analyze which variables (surface area, temperature, humidity) most affect the rate.
<i>Evaluate</i>	Ta yaya sauyin yanayi zai shafi zagayowar ruwa?	How might climate change affect the water cycle?	Ask learners to evaluate competing predictions about regional rainfall change under warming, citing evidence for and against each claim.

Internally, the Hausa query is transcribed by Listing 3, translated and intent-tagged by Listing 4, passed to the fine-tuned LLM via the prompt template in Listing 2, and the English response is back-translated and voiced in Hausa via Listing 5, the full loop implemented in Listing 6. Because the Bloom's level is fixed by the input tag rather than left to the model to infer, the four rows in Table 1 remain cognitively distinct even though they share a single topic, directly addressing the tendency of unconstrained LLMs to default to lower-order phrasing (Elkins et al., 2024).

Note on implementation status: Listings 1–2 and the architectural description of the system constitute a reference implementation only. None of the components above have been executed end-to-end there has been no LoRA training run, no live Hausa ASR/TTS test, and no integration test of the full ClassroomInteractionPipeline. The code is provided to make the design concrete and reproducible, not as evidence of a working system; see the Limitations subsection in Section 7 for further update.

5. Evaluation Plan and Expected Outcomes

The expected outcomes of this research include:

- A conceptual framework for AI-supported Integrated Science teacher education.
- Improved access to pedagogically relevant and curriculum-aligned instructional support for pre-service Integrated Science teachers.
- Enhanced capacity of teachers to integrate concepts across science disciplines and adopt inquiry-based teaching strategies.
- Increased inclusivity through multilingual and voice-based pedagogical support.
- A foundation for future empirical studies and scalable deployment of AI tools in Integrated Science teacher education programmes.

For the speech and translation components, word-error rate (WER) will be reported for the fine-tuned XLSR-Wav2Vec2 model on held-out Hausa classroom speech, benchmarked against the sub-15% range reported for fine-tuned Hausa ASR in comparable recent work (DeRenzi et al., 2025), with BLEU/chrF scores tracked for the Hausa-English translation stage before and after in-domain adaptation (Adelani et al., 2022). For pedagogical alignment, independent raters will score a sample of AI-generated prompts against their intended Bloom's-Taxonomy level, following the inter-rater protocol used to validate Bloom-conditioned question generation (Elkins et al., 2024). For overall usability and adoption, a structured survey and think-aloud sessions with pre-service teachers across partner colleges of education will assess

perceived usefulness, ease of use, and trust, consistent with technology-acceptance approaches used elsewhere in Sub-Saharan African higher education. These components together constitute a mixed-methods pilot to be conducted following the present design-and-implementation phase, rather than a claim of completed empirical validation.

6. Ethical and Privacy Considerations

Ethical AI principles are embedded throughout the framework. Regular evaluations are conducted to identify and mitigate bias, cultural insensitivity, or pedagogical misalignment in generated content. Privacy safeguards are implemented through anonymized data handling, secure storage, and strict access controls to protect institutional and user confidentiality. These measures ensure that the framework supports Integrated Science teacher education responsibly while maintaining teacher trust.

7. Limitations and Conclusion

Before the conclusion, three limitations are worth stating plainly. First, this framework has not yet been implemented or evaluated end-to-end: the architecture, code listings, and the Table 1 walkthrough describe an intended design, not a deployed or tested system. Second, the reference implementation depends on components that do not yet exist for this domain, a Bloom's-tagged Integrated Science corpus, a curated Hausa speech corpus, and an in-domain Hausa-English parallel corpus, and Hausa to English translation part explicitly uses a general-purpose multilingual MT model as a placeholder pending that data. Third, this framework's architecture closely parallels a related AI-powered pedagogical support framework for general teacher education developed by overlapping authors, readers should be aware of this relationship, and any submission of this paper should cite that related work explicitly and summarize rather than re-describe the shared architecture, to avoid the appearance of duplicate or overlapping publication.

This paper has presented the design, architecture, and Python reference implementation of an AI-powered, multilingual pedagogical support framework for Integrated Science teacher education in resource-constrained African contexts, structuring classroom interaction as an input-processing-output pipeline that combines curriculum curation and Bloom's-Taxonomy tagging, LoRA-fine-tuned LLM generation paired with fine-tuned Hausa speech recognition and translation, and bilingual text-and-speech delivery. The framework operationalizes the

three objectives set out for this research: strengthening lesson planning and conceptual integration, embedding pedagogical principles into AI-generated content, and enabling inclusive, voice-based, multilingual interaction. Future work will execute the evaluation plan outlined in Section 5 across selected colleges of education, extend the speech and translation components to additional Nigerian languages, and explore end-to-end translanguaging.

Acknowledgments

This work was supported by Isa Kaita College of Education, Dutsin-Ma, Katsina, through the TETFund Institution Based Research (IBR) intervention programme. The research grant number is TETF/DR&D/CE/COE/KATSINA/IBR/2026/VOL.I.

References

- Adelani, D. I., Ji, Y., Dossou, B. F. P., Orife, I., Akinode, V., Amuok, K., ... Nekoto, W. (2022). A few thousand translations go a long way: Leveraging pre-trained models for African languages. In Findings of the Association for Computational Linguistics: ACL 2022 (pp. 305–316). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.findings-acl.25>
- Baevski, A., Zhou, Y., Mohamed, A., & Auli, M. (2020). wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems*, 33, 12449–12460.
- Birhane, A. (2021). Algorithmic injustice: A relational ethics approach. *Patterns*, 2(2), Article 100205. <https://doi.org/10.1016/j.patter.2021.100205>
- Bonner, E., Lege, R., & Frazier, E. (2023). Large language model-based artificial intelligence in the language classroom: Practical ideas for teaching. *Teaching English with Technology*, 23(1), 23–41.
- Buzaaba, H., Wettig, A., Adelani, D. I., & Fellbaum, C. (2025). Luga-Llama: Adapting large language models for African languages. *arXiv*. <https://doi.org/10.48550/arXiv.2504.06536>
- Conneau, A., Baevski, A., Collobert, R., Mohamed, A., & Auli, M. (2021). Unsupervised cross-lingual representation learning for speech recognition. In *Proceedings of Interspeech 2021* (pp. 2426–2430). <https://doi.org/10.21437/Interspeech.2021-329>
- Cukurova, M., Luckin, R., & Kent, C. (2020). Artificial intelligence and big data for education: Ethical implications. *Computers and Education: Artificial Intelligence*, 1, Article 100001. <https://doi.org/10.1016/j.caeai.2020.100001>
- DeRenzi, B., Dixon, A., Farhi, M. A., & Resch, C. (2025). Synthetic voice data for automatic speech recognition in African languages. In *Proceedings of the First Workshop on Advancing NLP for Low-Resource Languages* (pp. 152–186). Association for Computational Linguistics. <https://aclanthology.org/2025.lowresnlp-1.16/>
- Elkins, S., Kochmar, E., Cheung, J. C. K., & Serban, I. (2024). How teachers can use large language models and Bloom's Taxonomy to create educational quizzes. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(21), 23084–23091. <https://doi.org/10.1609/aaai.v38i21.30353>
- García, O., & Wei, L. (2014). *Translanguaging: Language, bilingualism and education*. Palgrave Macmillan.
- Holmes, W., Porayska-Pomsta, K., & Holstein, K. (2024). AI in the classroom: Understanding teacher–AI partnerships. *Computers & Education*, 205, Article 104852. <https://doi.org/10.1016/j.compedu.2023.104852>
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., & Chen, W. (2022). LoRA: Low-rank adaptation of large language models. In *Proceedings of the International Conference on Learning Representations (ICLR 2022)*.
- Luckin, R., Moss, E., & Holmes, W. (2022). Developing AI literacy for teachers: A framework for professional growth. *International Journal of Artificial Intelligence in Education*, 32(4), 745–765. <https://doi.org/10.1007/s40593-022-00280-1>
- Makalela, L. (2018). Moving out of linguistic boxes: The effects of translanguaging strategies for multilingual classrooms. In *Language, Learning and Epistemic Access in Multilingual Education* (pp. 24–41). Routledge. <https://doi.org/10.4324/9780429453113>
- Ojo, J., Ogundepo, O., Oladipo, A., Ogueji, K., Lin, J., Stenertorp, P., & Adelani, D. I. (2024). AfroBench: How good are large language

- models on African languages? arXiv. <https://doi.org/10.48550/arXiv.2311.07978>
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., ... Scialom, T. (2023). Llama 2: Open foundation and fine-tuned chat models. arXiv. <https://doi.org/10.48550/arXiv.2307.09288>
- UNESCO. (2023). Education and digital equity in Africa: Challenges and opportunities. UNESCO Publishing.
- Yakubu, Y. (2024). Application of artificial intelligence in technology-enhanced learning for sustainable growth and inclusive development in Nigeria. *Zaria Journal of Educational Studies*, 24(1), 64–75.
- Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education. *International Journal of Educational Technology in Higher Education*, 16(1), Article 39. <https://doi.org/10.1186/s41239-019-0171-0>
- Zhai, X., He, J., Liu, G., & Kuo, Y.-L. (2021). A review of artificial intelligence in education from 2010 to 2020. *Complexity*, 2021, Article 8812542. <https://doi.org/10.1155/2021/8812542>