SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 9| September 2026 - Peer Reviewed Journal

DTICS: AN OFFLINE-FIRST MULTILINGUAL DENTAL TRIAGE SYSTEM FOR NIGERIAN PRIMARY HEALTHCARE

Grace Ojochenemi Emmanuel-Anorue[1], Sadiya Mohammed Dala[2]
[1]Department of Computer Science, Federal College of Education (Technical), Gombe, Nigeria
[2]Management Information System (MIS) Director, Federal College of Education (Technical), Gombe, Nigeria

ABSTRACT

In many of the northern states, Nigeria has less than one dentist to serve 50,000 people and there are no dental officers of any cadre in Primary Healthcare Centres (PHCs) in rural areas. Community Health Extension Workers (CHEWs), who are the main first contact providers for rural patients, are poorly trained in formal dental care, and have to judge if a patient has an emergency that may result in life threatening complications like Ludwig's angina, a fast spreading infection of the airway which if not managed promptly has a reported mortality rate of 8-40%. In this paper, an Android application called DTICS (Dental Triage and Intelligent Classification System) is presented which categorizes the patient's reported dental complaints into Emergency, Urgent and Routine for English, Hausa and Nigerian Pidgin language. DTICS uses a three-layer safety architecture: deterministic, fully offline rules layer, which encodes 8 emergency rules based on evidences, and has absolute priority over AI inference; an LLM classification layer, which classifies free-text and multilingual symptom description that are not captured by the rules; and a human escalation layer, which routes emergency and low-confidence cases to a dental officer with the help of an auditable Confirm/Override interface. The Android client is built using Clean Architecture, Hilt dependency injection, Room persistence (encrypted using SQLCipher) and WorkManager-based background synchronisation, and has been informally evaluated by practicing dentists for usability. The system architecture is presented, as well as the rationale behind its conservative and rule-first safety design and a shadow-mode evaluation protocol measuring the agreement with dental officer triage decisions (Cohen's κ) and emergency sensitivity. DTICS proves that a clinically conservative, offline-resilient triage tool can be created in the connectivity and language landscape of Nigeria, and provides a replicable framework for AI-based triage in resource-limited environments.

KEYWORDS— *Dental Triage; Clinical Decision Support; Offline-First Mobile Health; Multilingual NLP; Healthcare AI Safety; Nigerian Mhealth*

I. INTRODUCTION

Dental disease is a common but under-resourced type of non-communicable disease in sub-Saharan Africa. The World Health Organization (WHO) and the Nigerian Dental Association (NDA) estimate the nation's dentist-to-population ratio to be about 1:50,000 in urban areas, and in northern Nigeria, beyond 1:80,000 in Yobe, Zamfara, Borno and northern Adamawa states [1,2]. At tertiary hospital like the Lagos University Teaching Hospital (LUTH) and Aminu Kano Teaching Hospital (AKTH), the duration of symptoms prior to seeking care was uniformly 3 - 7 days for all patients, with periapical abscess, spreading cellulitis and sometimes Ludwig's angina, a rapidly lethal bilateral submandibular space infection requiring emergency airway management, a common complication of reversible pulpitis [3, 4].

This is not a failure of treatment; Nigerian dental officers can handle drainage, extraction, antibiotic management, etc., but a failure of information: neither the patient nor the CHEWs have the clinical framework to recognise dental emergencies at first contact. About 60% of the population in Nigeria resides in rural/peri-urban communities where a dental officer of any cadre is not easily accessible [5] and CHEWs resort to only two indicators: visible swelling and self-reported pain, which is inadequate to differentiate between pericoronitis and a progressive fascial space infection.

In 2023, Nigeria's total mobile penetration rate stood at 89.4 million unique subscribers[6] and mHealth is a key focus of the country's national digital-economy policy for underserved populations [7]. While other systems show excellent triage performance, they are overwhelmingly cloud based, English only, and tested only on non-African populations, which precluding them from being effective in PHCs in the rural parts of Nigeria where connectivity is unstable and the language spoken is Hausa.

In this paper, we introduce DTICS, a combination of a deterministic offline emergency-detection layer, an LLM-based multilingual classification layer and a human-in-the-loop escalation workflow, integrated together into an Android application. Related work is reviewed in Section II, the system architecture is described in Section III, multilingual and offline data handling is discussed in Section IV, security and ethics is addressed in Section V, the status of implementation is presented in Section VI and the testing plan is presented in Section VII and Section VIII discuss and conclude.

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 9| September 2026 - Peer Reviewed Journal

II. RELATED WORK

Since BERT [10] has made significant progress for clinical text processing, domain-adapted transformer models have also made significant strides in biomedical and clinical concept extraction (BioBERT[11] and ClinicalBERT[12] respectively), but only on English-language data from cloud-connected American EPs, conditions which are not applicable to Nigerian PHCs. The feasibility of a future on-device migration of DTICS's classification layer on to a smaller device, like a mobile phone, is offered by compact architectures like MobileBERT[13] and DistilBERT[14] that achieve similar accuracy with a much smaller number of parameters, but which do not yet perform on-device inference (Section III-C).

Multilingual transformer models, such as mBERT[10], XLM-RoBERTa[15] and AfroXLM-R [16] have been applied to low-resource African languages, and AfroXLM-R has shown better performance on Hausa text classification than the general multilingual models. There is a vast amount of literature on code switching between English, Hausa and Pidgin in the same patient statement[17]. To our best knowledge, there has been no published work on the topic of multilingual clinical NLP for dental triage in a Nigerian language context previously.

This work follows the "cloud-privileged design bias"[18] that is commonly adopted in mainstream mHealth tools, even in settings with limited connectivity, and the offline-first architectural framework for primary-care decision support in Africa proposed by Wainaina et al.[19] that emphasises local inference and graceful degradation principles, and is followed in this work with Layer 1. On the safety side, both the WHO ethics guidance for AI in health[21] and Shneiderman's human-centred AI framework[20] state that bounded authority of AI and transparency are also necessary, and that consequential decisions must be made by humans.Both have been supported by the human-centred AI framework by Shneiderman[20] and WHO ethics guidance for AI in health[21] which state that bounded AI authority and transparency are essential, that human oversight is required for consequential decisions.

III. SYSTEM ARCHITECTURE

A) Overview

The DTICS Android client is done in Kotlin and follows Clean Architecture (MVVM presentation layer + Repository abstraction). Hilt offers dependency injection at compile-time between modules (Table I). Three roles (self-service patient, CHEW-assisted intake and dental-officer review) are enabled on 10 screens.

Table 1: Different model respond

Module	Responsibility	Key Tech.
:core:database	Local persistence, audit trail	Room, SQLCipher
:core:network	API client, sync scheduling	Retrofit, OkHttp, WorkManager
:core:rules	Deterministic emergency engine	Pure Kotlin, no deps.
:feature:chew_triage	CHEW-assisted clinical intake	Compose, vitals form
:feature:officer_inbox	Queue, review, Confirm/Override	Compose, RecyclerView

B. Layer 1: Deterministic Rule Engine

Layer 1 is a pure-Kotlin engine with zero network dependency and absolute priority over the AI layer: if any rule fires, the pipeline returns an 'Emergency' classification without invoking Layer 2. Eight rules (Table 2), derived from documented dental-emergency indicators, are evaluated synchronously against patient- or CHEW-entered symptoms, vitals, and pain data.

Table 2: Trigger conditions

Rule ID	Trigger condition
AIRWAY_COMPROMISE	Floor-of-mouth elevation + trismus <15mm, or bilateral swelling + dysphagia
FEVER_PLUS_SWELLING	Temp ≥38.5°C + facial swelling
UNCONTROLLED_BLEEDING	Bleeding with severe/prolonged presentation
AVULSION_WINDOW	Reported tooth avulsion (<60 min window)
SWALLOW_DIFFICULTY	Any reported dysphagia
SEVERE_TRISMUS	Inter-incisal opening <10mm
HIGH_FEVER	Temp ≥39.5°C with dental complaint
EXTREME_PAIN	Pain ≥9/10 for ≥3 days

C) Layer 2: AI Classification Layer

When a complaint does not fall under a Layer-1 rule, DTICS Generates a structured clinical prompt (symptoms, pain score and duration, patient selected language, and any voice transcription, if available, from CHEW), and sends this prompt to a hosted large language model that returns a structured JSON object (classification, confidence, reasoning, recommended action, and contributing

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 9| September 2026 - Peer Reviewed Journal

factors). The out is passed through per class confidence thresholds and if the class prediction is below threshold, a subsequent prediction is passed for that class to the next level of threshold (and is not accepted at face value, and maintains the asymmetric bias of the system towards over-triage).

The current prototype is based on the LLM approach, which is also utilized to implement the multilingual classification without the need of an annotated Nigerian dental-language training corpus that is not available yet, but can classify the input language within a second or two. Its restriction is a strict requirement for network connectivity at Layer 2, which is currently an on-device, quantised transformer (e.g., MobileBERT[13]) the target architecture for full offline operation at this layer, and therefore, is targeted for near-term future work (Section VIII) and is not a fully realized component.

D. Layer 3: Human Escalation

Each Emergency classification (Layer 1 or Layer 2) will be surfaced in the dental officer's case queue with the corresponding triggering rule flags or AI confidence with contributing factors for interpretability. Officers make a decision, and record this decision in an append-only audit log along with a time and officer identifier, which cannot be changed. Total outage SMS and push notification escalation channels are defined in the design and will be implemented in the next phase.

IV. MULTILINGUAL SUPPORT AND OFFLINE DATA HANDLING

Instead of a dedicated trained tokeniser, patients select one of the three options, English, Hausa or Nigerian Pidgin, at the beginning which defines the interface localisation, which is passed as an explicit language tag to the Layer-2 prompt (currently multilingual). Future work (Section VIII) refers to a domain-adapted Hausa/Pidgin subword vocabulary and an annotated code-switching corpus, which are not implemented components of the present system, but planned for on device migration.

All consultation records will captured into a SQLCipher encrypted (AES-256) Room database, but no attempt is made to send them over a network, so the CHEW flow and the patient facing flow continue to be usable even if it is not connected to the internet. There is a record level identifier generated by the client, there's a sync-state field and WorkManager opportunistically synchronises in the background when resources are available (when there is network/battery available, for instance). Architecture describes the remote synchronisation endpoint, which is a hub and spoke FastAPI service that will flatten the data from partner facilities but is not built yet, while the existing build is still queueing and persisting the records locally until the integration is complete.

V. SECURITY, ETHICS, AND GOVERNANCE

The field-level SQLCipher encryption protects data at rest, while the audit trail is append-only to ensure it cannot be tampered with. The LLM API credential is added to the client configuration when developing it during this prototype, but will not be sent in the client app during production, as is the case with other credentials. DTICS is designed to be consistent with the Nigeria Data Protection Act 2023 and the WHO ethical guidance for AI[21] which includes the principle of data minimisation, patients consent, capturing consent in a language of the patient's choice and a standing Clinical AI Governance Committee model for changes to rules and model versions.

The field-level SQLCipher encryption protects data at rest, while the audit trail is append-only to ensure it cannot be tampered with. The LLM API credential is added to the client configuration when developing it during this prototype, but will not be sent in the client app during production, as is the case with other credentials. DTICS is designed to be consistent with the Nigeria Data Protection Act 2023 and the WHO ethical guidance for AI[21] which includes the principle of data minimisation, patients consent, capturing consent in a language of the patient's choice and a standing Clinical AI Governance Committee model for changes to rules and model versions.

VI. IMPLEMENTATION STATUS AND EVALUATION PLAN

A. Current implementation status

All the features being implemented in this paper are available and running in the current codebase, such as the pattern of Clean Architecture module structure, the rule engine of Layer-1, the SQLCipher-encrypted local persistence and the WorkManager sync scaffold, as well as the patient /CHEW /officer Compose UI flows. The hosted LLM (Section III-C) enables the Layer-2 classification pathway. The remote sync backend, SMS/USSD escalation channel and on-device edge model are not implemented, but are design-completed and reported here as a future work product, and not as existing components.

B. Scheduled shadowing during the evaluation process. DTICS will be evaluated in a triage and deployment (shadow-mode) study with AKTH and in conjunction with any partnered PHC facilities, with sequential patients being triaged simultaneously by the DTICS and through standard human triage. Primary endpoint is Cohen's κ agreement with dental-officer classification, secondary endpoints include the sensitivity of confirmed emergencies and CHEW and officer System Usability Scale (SUS) scores. κ can be detected to be ≥ 0.80 when there are 280 triage encounters and the 95% CI half-width is 0.05 [22].

VII. DISCUSSION

The strong point about DTICs is the architectural: having an aggressive (or not so aggressive) AI classifier impose itself over an offline rule layer and an imperative human verification that ensures that if AI goes wrong, a patient does not come to any harm. This is

SJIF Impact Factor (2026): 8.863| ISI I.F. Value: 1.241| Journal DOI: 10.36713/epra2016 ISSN: 2455-7838(Online)
EPRA International Journal of Research and Development (IJRD)
Volume: 11 | Issue: 9| September 2026 - Peer Reviewed Journal

better than for cloud AI symptom checker solutions as Ada Health or Babylon (no offline back, no Nigerian language support or escalation automation possible at all or paper triage which gives no consistence and no escalation automation.

The big shortcoming of the present implementation is that offline-first is there fully on Layer 1, as well as offline first when data is being captured but not yet in using Layer 2 which requires a connection, due to its requirement for access to the hosted LLM. That's what it is; this prototype will not fix it, not a shortcoming inherent in the overall concept, LLMAI is as a plug and play component within the system, and the security framework within that can stay while on device replacement occurs, while keeping this concept

VIII. CONCLUSION AND FUTURE WORK

This article described DTICS, a 3 layer (fully-offline hard safety; natural language understanding through a LLM multi-language class; and an audit-trail human triage with back-routing) process that was tested informally against several working GPs/ dentists (with formal shadow mode study coming soon). Future work would be: (1) Moving Layer 2 to an on- device, quantised Transformer(e.g.,MobileBERTs) with adequate domain/cultural bias (using data fromNigerian dental corpus and corpus built in next 2 points); (2) Create a language tokeniser/corpus for Hausa, pidgin and codes-switching; (3)

Flesh out the FastAPI syncing back-end and the SMS/USSD backend. (4) Conduct the actualshadow mode clinical trial mentioned above

REFERENCES

1. *World Health Organization, "Global oral health status report: Towards universal health coverage for oral health by 2030," WHO Press, 2023.*
2. *Nigerian Dental Association, "Workforce and infrastructure gaps in Nigerian dental healthcare: Annual survey report 2022," NDA, 2022.*
3. *R. K. Bali, P. Sharma, S. Gaba, A. Kaur, and P. Ghanghas, "A review of complications of odontogenic infections," Natl. J. Maxillofac. Surg., vol. 14, no. 1, pp. 1–8, 2023.*
4. *B. F. Adeyemi, A. O. Olawale, and A. O. Lawal, "Ludwig's angina: Clinical presentation, management, and outcomes at a Nigerian tertiary centre," J. Oral Maxillofac. Surg., vol. 80, no. 3, pp. 512–519, 2022.*
5. *Federal Ministry of Health, "National Oral Health Policy for Nigeria," FMoH Nigeria, 2022.*
6. *GSMA, "The mobile economy: Sub-Saharan Africa 2024," GSMA Intelligence, 2024.*
7. *Federal Ministry of Communications, Innovation and Digital Economy, "National Digital Economy Policy and Strategy for Nigeria 2023–2027," FMCIDE, 2023.*
8. *C. Hong et al., "Clinical natural language processing in languages other than English: Opportunities and challenges," J. Biomed. Inform., vol. 138, 104108, 2023.*
9. *A. J. Thirunavukarasu et al., "Large language models in medicine," Nat. Med., vol. 29, no. 8, pp. 1930–1940, 2023.*
10. *J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in Proc. NAACL, 2019, pp. 4171–4186.*
11. *J. Lee et al., "BioBERT: A pre-trained biomedical language representation model for biomedical text mining," Bioinformatics, vol. 36, no. 4, pp. 1234–1240, 2020.*
12. *E. Alsentzer et al., "Publicly available clinical BERT embeddings," in Proc. 2nd Clinical NLP Workshop, 2019, pp. 72–78.*
13. *Z. Sun, H. Yu, X. Song, R. Liu, Y. Yang, and D. Zhou, "MobileBERT: A compact task-agnostic BERT for resource-limited devices," in Proc. ACL, 2020, pp. 2158–2170.*
14. *V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "DistilBERT, a distilled version of BERT," in Proc. NeurIPS EMC2 Workshop, 2019.*
15. *A. Conneau et al., "Unsupervised cross-lingual representation learning at scale," in Proc. ACL, 2020, pp. 8440–8451.*
16. *J. O. Alabi, J. Abbott, A. Krubel, and S. Ruder, "Adapting pre-trained language models to African languages via multilingual adaptive fine-tuning," in Proc. COLING, 2022, pp. 4336–4349.*
17. *G. I. Winata et al., "Language models are few-shot multilingual learners," in Proc. 1st Workshop on Multilingual Representation Learning, 2023, pp. 1–15.*
18. *M. Kluyts, A. Solomonides, and G. Marconi, "Offline-first clinical decision support in low-resource environments: A systematic framework review," J. Am. Med. Inform. Assoc., vol. 30, no. 4, pp. 748–759, 2023.*
19. *B. N. Wainaina, E. Ngassa, and A. Abubakar, "Towards offline-capable clinical decision support for African primary healthcare: An architectural review," Health Inform. J., vol. 30, no. 2, 2024.*
20. *B. Shneiderman, "Human-centered AI," Issues Sci. Technol., vol. 38, no. 4, pp. 56–61, 2022.*
21. *World Health Organization, "Ethics and governance of artificial intelligence for health: WHO guidance," WHO Press, 2023.*
22. *A. Donner and M. Eliasziw, "A goodness-of-fit approach to inference procedures for the kappa statistic," Stat. Med., vol. 11, no. 11, pp. 1511–1519, 1992.*