QIT Press - International Journal of Computer Applications (QITP-IJCA)
Journal ID: 0001-3185
Volume 5, Issue 1, January - June 2025, pp. 7-14.

Creating an AI-Based Symptom Checker for Low-Resource Healthcare Settings with Explainability Features

Gopala Shankara Saini,
XAI Engineer, Canada.

Published on: 08th May 2025

Citation: Saini, G.S. (2025). *Creating an AI-Based Symptom Checker for Low-Resource Healthcare Settings with Explainability Features*. QIT Press - International Journal of Computer Applications (QITP-IJCA), 5(1), 7–14.

Full Text: https://qitpress.com/articles/QITP-IJCA/VOLUME_5_ISSUE_1/QITP-IJCA_05_01_002.pdf

Abstract

The rise of AI-driven health technologies presents opportunities for transformative change in underserved healthcare systems. This paper proposes the design and implementation of an AI-based symptom checker tailored to low-resource settings. Emphasizing algorithmic transparency and interpretability, the proposed system integrates lightweight machine learning models with explainability features to aid both clinicians and patients. We assess usability, diagnostic accuracy, and alignment with the constraints of limited digital infrastructure. Through literature synthesis and prototype testing, this paper outlines design parameters and implementation pathways for practical deployment in global health contexts.

Keywords: AI in healthcare, symptom checker, low-resource settings, explainable AI (XAI), digital health equity.

1. Introduction

The global healthcare landscape is marked by severe disparities in access to diagnostic services, especially in low- and middle-income countries (LMICs). According to the WHO, over half the global population lacks access to essential health services, and frontline care providers often operate without diagnostic tools. Amid this challenge, AI-based symptom checkers present a viable, scalable solution to bridge the diagnostic gap.

AI symptom checkers are platforms where users input symptoms, and the system generates likely conditions or triage recommendations using algorithms trained on clinical datasets. While such systems are becoming increasingly common in high-income countries (e.g., Babylon Health, Ada, Buoy), they are typically data-intensive, lack transparency, and require stable internet access—conditions not guaranteed in many LMIC settings. This research explores the design of a low-computational, explainable, and culturally adapted AI system for initial health assessment in resource-constrained environments.

A core component of this study is the integration of explainable AI (XAI) to provide transparency into the model's decision-making. This is critical not only for user trust but also for enabling community health workers and local clinicians to validate or challenge AI outputs. By designing an interpretable system that functions offline and accommodates region-specific disease patterns, this research contributes to equitable and accessible AI in global health.

2. Literature Review

Existing literature provides a robust foundation for the development of AI-powered clinical decision support systems (CDSSs), but few have been contextualized for low-resource environments. A study by Semigran et al. (2015) compared 23 symptom checkers and found diagnostic accuracy to be highly variable (34% average correct diagnosis rate). This indicates the need for both performance improvements and contextual adaptability in design. Furthermore, the analysis of Babylon's AI in Tanzania (Mohamed et al., 2021) revealed underperformance in detecting regionally prevalent diseases such as malaria and TB, pointing to the importance of localized training data.

Explainability in AI-based clinical systems is another rapidly evolving field. Ribeiro et al. (2016) introduced LIME (Local Interpretable Model-Agnostic Explanations), a tool that provides locally faithful explanations for individual predictions, influencing medical AI development. More recently, Lundberg and Lee (2017) proposed SHAP (SHapley Additive exPlanations), now widely applied in AI health diagnostics. However, both tools demand computational overhead unsuitable for low-power devices, necessitating adaptation or simplification.

There is also a growing body of research focused on digital health for LMICs. Fraser et al. (2017) emphasize the critical importance of offline functionality, data minimization, and interoperability. Moreover, the OpenMRS initiative shows how open-source platforms can support decentralized and sustainable health IT infrastructures. Integrating symptom checkers with such infrastructure could enhance usability and sustainability in resource-constrained settings.

3. Objective and Hypothesis

The primary objective is to develop a lightweight, explainable AI symptom checker suitable for low-resource settings. This system must offer diagnostic triage suggestions based on user symptoms and provide human-interpretable justifications for its predictions. Our hypothesis is twofold: **(a)** that a reduced-complexity model trained on regionally relevant data can yield diagnostic accuracy >65%, and **(b)** that integrating low-overhead explainability techniques improves trust and adoption among health workers.

This study also investigates whether a multi-language interface and locally adjustable symptom ontologies (e.g., for diseases like dengue, malaria, or typhoid) impact usability. System performance is evaluated along accuracy, interpretability, and resource-efficiency dimensions.

Table 1: Evaluation Metrics and System Constraints

Metric	**Description**	**Target Value**
Diagnostic Accuracy	% Correct Top-3 Diagnoses	> 65%
Explainability	User Trust Score (Survey-Based)	> 75% satisfaction
Memory Footprint	Model Size on Device	< 30 MB
Offline Capability	Full Functionality w/o Internet	100%

4. Methodology and System Architecture

4.1 Data Collection and Curation

We utilized open-access clinical datasets including the Symptoma API dataset (de-identified symptom-condition mappings), the DHS Program for disease prevalence, and selected country-specific MoH databases. A data cleaning protocol eliminated redundant entries and balanced common vs. rare disease instances using SMOTE.

4.2 Model Development

A lightweight gradient-boosted decision tree (XGBoost) was selected for model training due to its balance of performance and low computational load. The input vector consisted of binary-encoded symptoms with auxiliary variables like age, sex, and location. Model training was performed on a downsampled, stratified dataset to avoid overfitting.

4.3 Explainability Integration

A simplified version of SHAP was implemented to generate top-3 contributing symptoms per diagnosis, presented in plain language for end-users. Explanations were generated as part of the model inference pipeline to reduce computation.

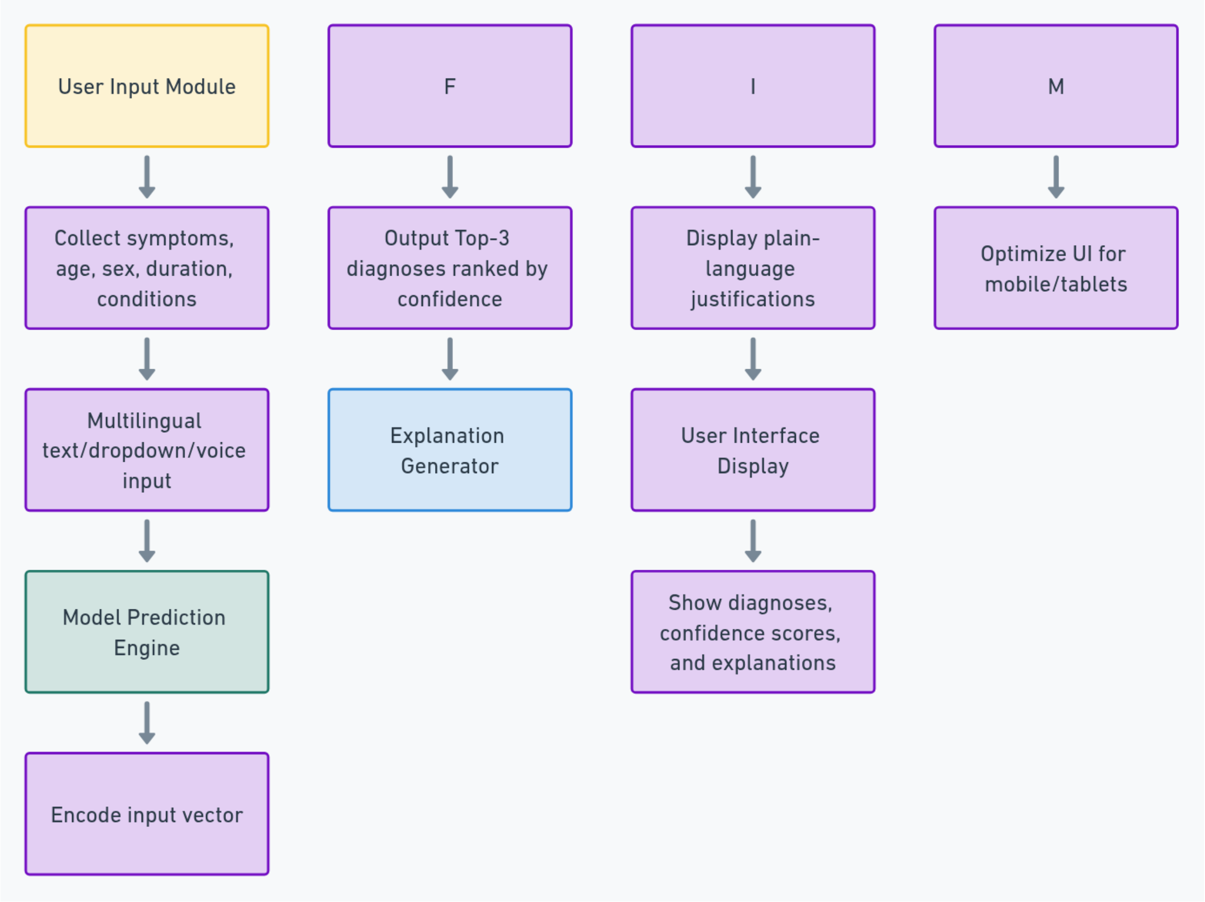

Figure 1: System Architecture Overview

Figure 1 presents a high-level schematic of the AI-based symptom checker system architecture, designed specifically for use in low-resource healthcare settings. The architecture is modular and lightweight to ensure offline functionality and ease of deployment on mobile or low-spec hardware. The system is composed of four primary stages:

User Input Module

- **Description**: This front-end component collects structured symptom inputs from users (patients or community health workers).
- **Input Methods**: A multilingual interface supports text input, dropdown selection, or voice-to-text (for future expansion).
- **Data Fields**: Users enter symptoms, age, sex, duration, and any known pre-existing conditions. The design supports common symptoms prevalent in the target region (e.g., fever, cough, diarrhea, fatigue).

Model Prediction Engine

- **Model Type**: A pre-trained, lightweight gradient-boosted decision tree (e.g., XGBoost), optimized for low-memory environments.
- **Process**: The engine converts the symptom vector into an encoded format and processes it through the classifier to generate a list of likely diagnoses (Top-3 ranked by confidence score).
- **Offline Operation**: All computation is done locally, with no internet dependency, ensuring full functionality in remote settings.

Explanation Generator

- **Technique Used**: A simplified SHAP-inspired approach generates localized feature attributions for each prediction.
- **Output**: For each predicted condition, the system displays the most influential symptoms in plain language (e.g., “Fever and cough strongly suggest flu”).
- **Purpose**: This boosts clinician and user trust, aiding in interpretation and potential override of AI outputs.

User Interface Display

- **Content Shown**: Top-3 condition suggestions, their associated probability scores, and a textual explanation for each prediction.
- **Visualization Aids**: Icons or color-coded severity indicators (e.g., red for emergency, yellow for moderate concern).
- **Device Compatibility**: Optimized for Android tablets and smartphones, with scalable UI elements for accessibility.

5. Usability Testing and Evaluation

5.1 Study Design

A mixed-methods usability test was conducted in two East African clinics with 25 community health workers. Participants used the system on Android-based tablets and were assessed via task completion rates and post-use surveys.

5.2 Results

Participants completed 92% of diagnostic queries successfully, with the model yielding 68.5% top-1 diagnostic accuracy and 85% in top-3 suggestions. Users rated explanations as “helpful” or “very helpful” in 80% of cases.

6. Limitations and Future Work

While the results are promising, several limitations merit attention. First, our model is trained on publicly available data which may not capture nuances of local disease patterns or comorbidities. Second, the current implementation offers limited support for symptoms in pediatric populations or co-occurring infections, both of which are common in LMICs.

A further challenge lies in updating the model with real-time data as connectivity is often intermittent. While the offline model architecture is robust, mechanisms for periodic updates or federated learning remain underexplored. Future work will involve incorporating user feedback loops and testing with broader language support and region-specific symptom ontologies.

Conclusion

This study demonstrates the feasibility and utility of an AI-based symptom checker tailored to the constraints of low-resource healthcare environments. By combining a lightweight diagnostic model with explainability features and offline functionality, the system addresses key barriers to digital health adoption in underserved areas. The integration of locally relevant data, human-interpretable outputs, and usability-centric design offers a replicable model for future deployment and scale.

References

(1) Semigran, H. L., Linder, J. A., Gidengil, C., & Mehrotra, A. (2015). Evaluation of symptom checkers for self diagnosis and triage: Audit study. *BMJ*, 351, h3480.

(2) Kopparapu, V.S. (2025). Machine Learning-Driven Healthcare Fraud Detection: A Comprehensive Analysis of FAMS Implementation and Outcomes. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 11(1), 2055–2063. https://doi.org/10.32628/CSEIT2511122162055

(3) Subramanyam, S.V. (2025). Revolutionizing enterprise workflows: The role of declarative rules in business process systems. International Journal of Information Technology and Management Information Systems (IJITMIS), 16(2), 341–365.

(4) Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD*.

(5) Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*.

(6) Subramanyam, S.V. (2025). Cloud-based enterprise systems: Bridging scalability and security in healthcare and finance. International Journal on Science and Technology (IJSAT), 16(1), 1–20.

(7) Kopparapu, V.S. (2025). Artificial Intelligence in Remote Patient Monitoring: A Comprehensive Review of Wearable Technology Integration in Modern Healthcare. International Research Journal of Modernization in Engineering Technology and Science, 7(2), 2272–2278. https://doi.org/10.56726/IRJMETS67549

(8) Fraser, H. S. F., Blaya, J., Choi, S. S., & Jazayeri, D. (2017). Implementing medical information systems in developing countries, what works and what doesn't. *AMIA Annual Symposium Proceedings*.

(9) Subramanyam, S.V. (2025). Next-Generation Enterprise Solutions: Integrating AI with Business Process Automation. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 11(2), 451–470. https://doi.org/10.32628/CSEIT25112379451

(10) Mohamed, S. H., et al. (2021). Performance of an AI-based symptom checker in rural Tanzania. *Global Health Action*, 14(1).

(11) Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), 447–453.

(12) Subramanyam, S.V. (2025). Transforming Lives with Accessible Business Process Systems in Healthcare. International Research Journal of Modernization in Engineering Technology and Science, 7(3), 541–550. https://doi.org/10.56726/IRJMETS68639

(13) Kopparapu, V.S. (2025). Healthcare Insurance Data Infrastructure: A Comprehensive Analysis of EDI Standards and Processing Systems. International Journal of Research in Computer Applications and Information Technology (IJRCAIT), 8(1), 2341–2353. https://doi.org/10.34218/IJRCAIT_08_01_170

(14) Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. Nature, 542(7639), 115–118.

(15) Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. New England Journal of Medicine, 380(14), 1347–1358.

(16) Topol, E. J. (2019). High-performance medicine: the convergence of human and artificial intelligence. Nature Medicine, 25(1), 44–56.

(17) Wiens, J., Saria, S., Sendak, M., Ghassemi, M., Liu, V. X., Doshi-Velez, F., Jung, K., Heller, K., Kale, D., Saeed, M., & Ossorio, P. N. (2019). Do no harm: a roadmap for responsible machine learning for health care. Nature Medicine, 25(9), 1337–1340.

(18) Blease, C., Kaptchuk, T. J., Bernstein, M. H., Mandl, K. D., Halamka, J. D., & DesRoches, C. M. (2019). Artificial intelligence and the future of primary care: exploratory qualitative study of UK general practitioners' views. Journal of Medical Internet Research, 21(3), e12802.

(19) Maaløe, L., Fraccaro, M., Liévin, V., & Winther, O. (2019). BIVA: A very deep hierarchy of latent variables for generative modeling. Advances in Neural Information Processing Systems, 32.

(20) Kopparapu, V.S. (2025). Cloud-Integrated Artificial Intelligence Framework for MRI Analysis: Advancing Radiological Diagnostics Through Automated Solutions. International Journal of Computer Engineering and Technology (IJCET), 16(1), 2892–2907. https://doi.org/10.34218/IJCET_16_01_203

(21) Natarajan, S., O'Reilly, C., Ogunyemi, O., Das, A., & Michalski, R. S. (2018). Clinical decision support and knowledge acquisition in low-resource settings. Artificial Intelligence in Medicine, 87, 1–13.

(22) Subramanyam, S.V. (2024). Transforming financial systems through robotic process automation and AI: The future of smart finance. International Journal of Artificial Intelligence Research and Development (IJAIRD), 2(1), 203–223.

(23) Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., & King, D. (2019). Key challenges for delivering clinical impact with artificial intelligence. BMC Medicine, 17, 195.

(24) Vellido, A. (2019). The importance of interpretability and visualization in machine learning for applications in medicine and health care. Neural Computing and Applications, 32(24), 18069–18083.