IJCT

International Journal of Computer Techniques – IJCT Volume X Issue X, Month Year

Open Access and Peer Review Journal ISSN 2394-2231 https://ijctjournal.org/

Predicting the Onset of Hypertension Using Deep Learning Models in the Copperbelt Province of Zambia

Melvin Lumamba
Computer Science Department
The Copperbelt University
Kitwe, Zambia
Lumambamelvin084@gmail.com

Abstract - Hypertension continues to be a significant health challenge in Zambia, particularly in the Copperbelt Province, where many cases are diagnosed late and often advance to severe complications. This research aimed to design an artificial intelligence–driven system capable of predicting hypertension at earlier stages, thereby supporting preventive healthcare in resource-constrained settings. The study adopted a flexible, mixed-methods design that combined publicly available health datasets with professional insights from local medical practitioners. Predictive variables, including blood pressure readings, cholesterol levels, body mass index, and heart rate, were utilized to train various deep learning models. The approaches tested included Deep Neural Networks, Convolutional Neural Networks, and Recurrent Neural Networks enhanced with Long Short-Term Memory capabilities.

Model performance was assessed using widely accepted evaluation measures, namely accuracy, recall, precision, the F1 measure, and the area under the curve. The findings indicated that the optimized Convolutional Neural Network achieved an accuracy level of slightly above 85 percent. In comparison, the Long Short-Term Memory model produced an accuracy of eighty-three percent, with a recall rate exceeding ninety percent in detecting hypertensive cases. To ensure the system was practical for end-users, it incorporated a user-friendly interface developed with Python Tkinter and Jupyter Notebook, enabling real-time prediction and reporting. Its modular server-client architecture enhanced both scalability and security, while model interpretability was supported through visualization techniques such as gradient-based mapping.

The research also highlighted several challenges, including the shortage of structured local datasets, insufficient computing resources, and limited knowledge of artificial intelligence within the health sector. Despite these obstacles, the research demonstrated that tailored deep learning applications can strengthen public health decision-making in Zambia and provide a foundation for the development of future data-driven medical solutions, as exemplified by the prototype system developed.

Keywords - Deep Learning, Artificial Intelligence, Hypertension Prediction, Convolutional Neural Net, and Long Short-Term Memory.

1. INTRODUCTION

Hypertension, or high blood pressure, has become a pressing health concern in Zambia, particularly in the Copperbelt Province, where late diagnosis is common and often leads to serious complications [12], [15]. This condition is strongly influenced by lifestyle transitions such as poor dietary habits, reduced physical activity, and rapid urbanization, all of which contribute to the growing burden of non-communicable diseases [6], [14]. In many cases, individuals only discover they are hypertensive after experiencing severe outcomes like stroke, kidney disease, or cardiovascular failure, which complicates treatment and increases health costs [15]. Conventional diagnostic practices in Zambia remain largely reactive, relying on traditional methods that fail to identify hypertension at an early stage. As a result, large numbers of patients go undiagnosed until the disease has advanced, creating significant social and economic consequences, including loss of productivity and high medical expenses [3].

To address these challenges, this research investigates the potential of artificial intelligence and deep learning in predicting the onset of hypertension. Deep learning is a branch of artificial intelligence that uses layered neural network architectures to automatically learn complex data representations and patterns from large datasets, enabling more accurate predictions in tasks such as image recognition, speech processing, and medical diagnosis [11]. International evidence shows that advanced models, such as Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), can process extensive health datasets to identify subtle patterns of risk factors [5], [7]. Such technologies have demonstrated promising results in early disease detection, yet their use in Zambia remains minimal [9]. The gap between global innovation and local medical practice underscores the need to develop context-specific solutions. By incorporating both secondary health datasets and local information from the Copperbelt Province, this research seeks to build a predictive framework tailored to Zambia's healthcare environment.

The objective of the research is to determine whether deep learning approaches can accurately forecast hypertension using locally relevant health indicators. Additionally, the research will identify key predictors, evaluate various algorithmic models, and provide recommendations for integrating predictive systems into clinical practice. Through this approach, the research aims to shift healthcare delivery in Zambia from a focus on treatment to one on prevention, demonstrating the potential of digital health innovations to improve outcomes in low-resource settings [16], [13]. Ultimately, this research aims to make both practical and scientific contributions that support the development of AI-enabled healthcare in sub-Saharan Africa [12]. Developments align with the broader growth of deep learning (DL), which leverages layered neural networks to provide flexible, adaptive, and data-driven analytical opportunities [21].

In healthcare research, the combination of medical imaging and deep learning provides a valuable pathway for creating systems that are interactive, patient-focused, and capable of improving disease detection, monitoring, and predictive decision-making.

However, many current platforms are not fully integrated into clinical workflows, often requiring practitioners and researchers to switch between different tools, such as separate imaging software and independent predictive models. This disjointed approach interrupts the clinical process, slows down decision-making, and reduces opportunities for collaboration—factors that are essential in healthcare practice. To overcome this limitation, the present study introduces a deep learning platform that combines image analysis with AI-powered feedback delivered in real time. The system's main contribution is the design, development, and evaluation of a single interface where users can upload, analyze, and interpret medical images while receiving context-aware guidance and supporting collaborative work.".

A. Problem Statement

Hypertension remains one of the leading non-communicable diseases (NCDs) globally and presents a significant public health challenge in Zambia, particularly in the Copperbelt Province. The increasing prevalence of hypertension in this region is linked to urbanization, sedentary lifestyles, poor dietary habits, and limited access to preventive healthcare services [14], [6]. Despite its rising burden, hypertension often remains undetected until it causes severe complications such as stroke, kidney disease, or cardiovascular failure [15]. This is mainly due to the reactive nature of the Zambian healthcare system, which lacks adequate screening tools and predictive diagnostic frameworks.
Recent advancements in artificial intelligence (AI) have demonstrated substantial success in disease prediction and early diagnosis globally [7], [10]. In particular, deep learning models have shown the ability to analyze complex health datasets and identify subtle patterns that may signal the early stages of chronic diseases, such as hypertension [5], [12]. However, Zambia has yet to fully embrace or apply these technologies within its healthcare system, thereby creating a gap between global innovations and local healthcare practices.

This research is motivated by that gap—while AI-based health prediction tools have succeeded internationally, their implementation in Zambia remains limited. The philosophical and practical value of this research lies in its potential to transform the Copperbelt Province's healthcare approach from reactive treatment to proactive, data-driven prevention. By evaluating the feasibility and effectiveness of deep learning models trained on local or regionally adapted data, this research aims to contribute to both the scientific field and Zambia's healthcare development. Additionally, this work seeks to identify region-specific risk factors that significantly influence hypertension predictions, thereby improving the precision and applicability of AI models in local contexts. If successful, the outcomes could justify expanded investment in AI-powered health systems across Zambia and other similarly resource-constrained environments [3], [9].

B. Research Objectives

1. To collect and preprocess relevant clinical and lifestyle data from the Copperbelt Province, then design and implement deep learning algorithms (CNN, RNN, LSTM) suitable for early hypertension prediction.
2. To evaluate multiple deep learning models using comprehensive performance metrics (accuracy, precision, recall, ROC curves, AUC scores) and identify the most significant local hypertension predictors through model interpretability techniques.
3. To provide actionable recommendations for integrating the predictive system into local healthcare practices, enabling early intervention and disease prevention in Zambia's healthcare system.

B. Research Questions

1. How can clinical and lifestyle data from the Copperbelt Province be effectively collected, preprocessed, and utilized to design and implement deep learning models such as CNN, RNN, and LSTM for early prediction of hypertension?
2. How can multiple deep learning algorithms be comparatively evaluated using performance metrics such as accuracy, precision, recall, ROC curves, and AUC scores to determine the most effective model and identify the most significant local predictors of hypertension through interpretability techniques??
3. To what extent can the developed deep learning–based hypertension prediction system be integrated into Zambia's healthcare framework to enhance early diagnosis, support preventive interventions, and strengthen data-driven clinical decision-making?

II. LITERATURE REVIEW

Hypertension continues to pose a serious health challenge globally, especially in developing countries where healthcare systems often struggle with limited infrastructure and delayed diagnoses [1], [10]. Many cases remain undetected until complications arise, emphasizing the importance of early prediction systems. The introduction of artificial intelligence (AI) and deep learning has created new opportunities for improving how diseases are detected and managed [6], [7], [8]. These technologies are capable of analyzing large amounts of clinical information, uncovering complex relationships between different health indicators, and providing predictions that traditional methods might overlook. Frameworks such as TensorFlow and Keras have been instrumental in designing adaptive models that can learn from medical data to support early hypertension detection [18], [19]. In countries such as Zambia, where healthcare facilities are often under strain, such systems have the potential to enhance screening, improve treatment outcomes, and support preventive medicine [12], [13].

A. AI and Deep Learning in Precision Health

Recent developments in AI and deep learning have significantly advanced the idea of precision health by enabling the use of predictive analytics in clinical decision-making. Dulam and Gosukonda [20] demonstrated that deep learning models can integrate various patient data sources to improve the accuracy of medical diagnoses, while Rajkomar et al. [7] showed that deep architectures trained on electronic health records outperform traditional predictive models. Similarly, Racic et al. [21] found that convolutional neural networks are highly effective in detecting diseases from medical images. These examples demonstrate that AI and deep learning can have a significant impact on healthcare outcomes. However, in Zambia, adoption remains slow because of limited access to technology and a lack of skilled personnel [12], [13]. To close this gap, locally adapted frameworks that consider infrastructure, cost, and training are needed to make AI-based health tools both practical and sustainable.

B. Machine Learning and Deep Learning for Non-communicable diseases

Machine learning and deep learning techniques are being increasingly used to address non-communicable diseases, including hypertension, diabetes, and cardiovascular disorders. Chen et al. [6] and Kumar and Singh [8] reported that neural networks and other predictive algorithms can identify early disease patterns that help improve medical interventions. The use of Scikit-learn and NumPy for data preparation [16], [17], combined with TensorFlow and Keras for model training [19], allows for the development of reliable predictive systems that maintain accuracy even with limited computational power. These systems can be applied in Zambia to enhance medical analysis and reduce diagnostic delays, providing health professionals with timely information to guide patient management [10], [11], [13]. The integration of these methods into hospital systems could significantly strengthen disease monitoring and resource allocation in public health.

C. Wearable Technologies and Continuous Deep Learning Monitoring

The combination of wearable devices and deep learning algorithms has made it possible to monitor patient health and detect warning signs early continuously. Lee et al. [15] demonstrated that wearable devices paired with AI-driven models encourage patients to stay engaged in their care by providing constant feedback on vital signs such as heart rate and blood pressure. Recurrent neural networks and long short-term memory architectures have proven effective in analyzing this time-based data and predicting potential health risks [18]. However, despite their promise, such technologies are still rare in Zambia due to high costs and weak digital infrastructure [10], [12]. Locally adapted deep learning systems capable of replicating these monitoring functions through affordable digital platforms could help bridge the technological divide and support early detection in underserved communities.

D. Deep Learning and AI-based Diagnostic Tools in low-resource settings

Deep learning-based diagnostic systems have shown promise in improving health outcomes where resources are limited. Rajkomar et al. [7] found that AI models trained on electronic health records can achieve better diagnostic accuracy than conventional clinical approaches. Dulam and Gosukonda [20] further emphasized that these systems can lower operational costs while improving diagnostic consistency and scalability. Nonetheless, successful implementation in countries such as Zambia remains challenging due to shortages in technical expertise, data infrastructure, and ethical regulation [12], [13]. McKinney [3] noted that strong data management frameworks and computational support are necessary for building effective diagnostic systems. Overcoming these limitations will require targeted investment in training, the creation of data-sharing policies, and collaboration between research institutions and healthcare providers.

E. Theoretical Frameworks

This research draws on Precision Health Theory and Systems Theory to explain how deep learning can be applied in the healthcare sector. Precision Health Theory emphasizes personalized, data-informed care guided by predictive analytics [15], whereas Systems Theory emphasizes how technology, institutions, and social structures interact to influence healthcare delivery [4]. When combined, these theories demonstrate how deep learning can serve as a bridge between existing health data and more efficient care systems. They also provide a framework for understanding how predictive analytics can improve decision-making and resource use in the Zambian context [11], [13].

III. METHODOLOGY

A. Methodology

This section presents the architectural design and logical modeling of the hypertension prediction system using diagrams that visually represent its components and processes. An Entity-Relationship (ER) diagram was employed to depict the relational structure among critical entities such as patients, medical records, predictions, and machine learning models, which is essential for ensuring efficient data organization and management [10]. The Use Case diagram captures interactions between key system actors—patients, healthcare providers, and model administrators—across core operations like registering patients, creating medical records, generating predictions, and configuring models. A procedural flowchart illustrates the sequential logic from data input to prediction output, enhancing the interpretability of system behavior and decision-making flow [7]. Collectively, these design tools contribute to building a robust, scalable, and context-aware deep learning-based system tailored to the healthcare needs of Zambia's Copperbelt Province. The visual modeling approach supports clear system development, encourages stakeholder engagement, and provides a foundation for future enhancements, such as integrating locally sourced datasets and retraining models to improve precision [18], [16].

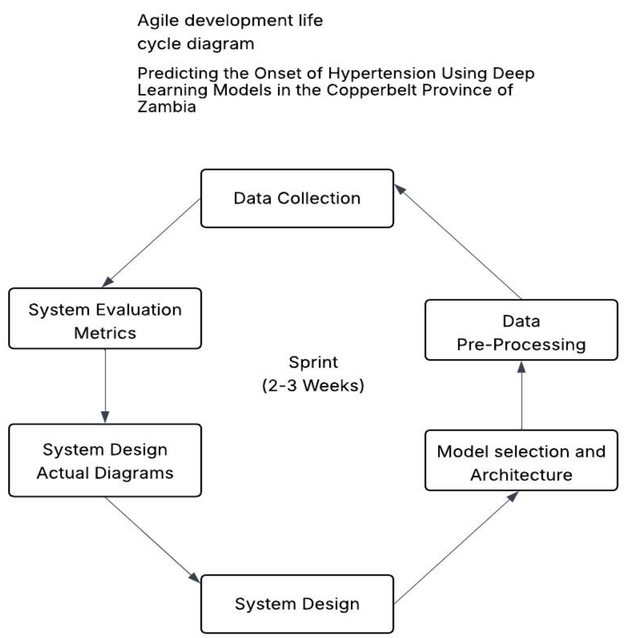

Figure 1: Agile Development Life Cycle

A qualitative exploratory design grounded in a constructivist paradigm captured students' experiences and perceptions of the platform [18]. Data were collected through semi-structured interviews with 12 purposively sampled students, reflective diaries, and observational field notes. Thematic analysis, using Orange Data Mining and Voyant, revealed organized patterns in usability, engagement, collaboration, and the impact of AI feedback [19]. Credibility was strengthened through triangulation, member checking, thick description, and peer debriefing. Ethical clearance, informed consent, and confidentiality safeguards ensured research integrity. A formal Software Requirements Specification (SRS) defined functional and non-functional requirements to guarantee system robustness and adaptability to low-resource educational contexts.

B. Research Paradigm

This research adopts a positivist paradigm, emphasizing objectivity, quantification, and replicability. Positivism is particularly suitable for data-driven studies, including those involving predictive analytics in healthcare through machine learning [19]. This paradigm assumes that hypertension risk can be predicted from physiological and lifestyle data using algorithmic models. It supports the development of a generalizable and interpretable deep learning system aligned with international best practices in artificial intelligence for early disease detection [8]. administrators—across core operations like registering patients, creating medical records, generating predictions, and configuring models. A procedural flowchart illustrates the sequential logic from data input to prediction output, enhancing the interpretability of system behavior and decision-making flow [7]. Collectively, these design tools contribute to building a robust, scalable, and context-aware deep learning-based system tailored to the healthcare needs of Zambia's Copperbelt Province. The visual modeling approach supports clear system development, encourages stakeholder engagement, and provides a foundation for future enhancements, such as integrating locally sourced datasets and retraining models to improve precision [18], [16].

C. Data Sources and Inclusion Criteria

Because structured local EHRs are limited, we used secondary, open health datasets for model prototyping and internal validation, complemented by expert elicitation to localize features. Tabular records were drawn from public repositories (e.g., BRFSS/CDC and Kaggle cardiovascular risk datasets), retaining adult records (≥18 years) with non-missing core vitals and metabolic markers. Where retinal imaging labels were available, a subset of hypertensive retinopathy cases was used to test image-aware extensions. To improve contextual fit, Zambian clinicians reviewed variable definitions, clinical ranges, and label intent, and guided the selection of locally relevant covariates (e.g., occupational activity, salt intake proxies) [6], [14]

Table 1 CNN Retinopathy Dataset

Image	Hypertensive Retinopathy
00000ce7.png	1
00000ce9.png	0
00000de9.png	1
00000dea.png	1

This dataset links retinal image file names to diagnostic labels for hypertensive retinopathy, where "1" indicates presence and "0" absence. Out of four samples, three are positive and one is negative, providing variation useful for model training and testing. Retinal imaging is essential in hypertension studies because high blood pressure can damage retinal vessels, causing changes like narrowing or leakage that are detectable in scans. These images, therefore, serve as a non-invasive method for monitoring hypertension-related complications and advancing predictive research [20].

D. Outcomes, Predictors, and Variable Handling

The research considered a binary outcome variable representing hypertension status, where a value of 1 indicated hypertensive cases and zero (0) denoted non-hypertensive individuals, as specified in the dataset. The predictor variables included demographic and clinical characteristics such as age, sex, body mass index (BMI), systolic and diastolic blood pressure, and maximum recorded heart rate (thalach). Additional factors included lipid measures, such as cholesterol and fasting blood sugar levels, electrocardiogram (ECG) descriptors, and selected lifestyle characteristics. For baseline statistical comparisons and hypothesis-driven tests, generalized linear models were applied. These models followed the standard specification.

The central focus of the analysis was the likelihood of developing hypertension, which served as the dependent outcome. This risk was examined in relation to a range of standardized explanatory variables, including demographic details, blood pressure readings, cholesterol levels, and lifestyle-related indicators. Random variation within the data was acknowledged and treated as background error. The use of regression analysis provided a transparent and interpretable starting point by showing how individual predictors influenced the outcome; however, the method was limited to linear patterns. To gain deeper insights, advanced deep learning techniques were introduced, as these models are capable of identifying complex, non-linear relationships among the same set of variables [5], [12], [16].

E. Preprocessing and Feature Engineering

Records with critical label gaps were excluded, while the remaining missing predictors were imputed using the median for continuous variables and the mode for categorical variables. The imputed values were cross-checked to ensure clinical plausibility. Categorical variables were one-hot encoded, and continuous variables were standardized using z-scores. Obvious data entry errors (e.g., diastolic blood pressure greater than systolic blood pressure) were removed after manual review. To address differences between international and Zambian datasets, three adjustments were applied: (i) trimming and adjusting extreme values to clinically acceptable ranges, (ii) stratifying data splits according to outcomes, and (iii) calibrating decision thresholds on the validation set to favor high-recall operating points, consistent with clinical practice [12], [6].

Three families of models were evaluated. The first was the Deep Neural Network (DNN), designed for tabular data, consisting of three to five dense layers with 64–256 units, ReLU activations, batch normalization, and dropout ranging from 0.2 to 0.4, followed by a sigmoid output layer for binary risk prediction. The second was the Convolutional Neural Network (CNN), applied to both tabular and image data. For tabular data, features were reshaped into two-dimensional grids to capture local interactions. For image data, particularly the hypertensive retinopathy subset, a lightweight CNN with two to three blocks of 3×3 convolutions and max pooling was employed [11], [2]. The third family included Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) models, which were applied to sequential records of vital signs. These models utilized one or two LSTM layers with 32–128 units to capture temporal patterns across patient visits.

All models were implemented in TensorFlow/Keras (Python), utilizing the Adam optimization algorithm with a decaying learning rate, ranging from 1e-3 to 1e-4. Binary cross-entropy loss was applied, with mini-batch sizes ranging from 32 to 128, early stopping based on validation loss, and automatic learning rate reduction upon plateauing [16] [19].

F. Training, Validation, and Removal Tests.

The dataset was divided into training, validation, and testing groups in proportions of 70%, 15%, and 15%, respectively. To reduce bias caused by uneven class distributions, class weights were applied, and the Synthetic Minority Oversampling Technique (SMOTE) was used only on the training set when necessary. The reliability of the models was assessed using five-fold cross-validation, and their performance was reported as averages with standard deviations. To evaluate the importance of different feature types, removal tests were conducted by leaving out groups of variables such as lipid indicators and ECG attributes. This approach demonstrated their contribution to prediction performance and guided priorities for future data collection [5], [12], [16], [10].

RESULTS

G. Model Evaluation Metrics

The primary evaluation metrics were recall (sensitivity) for hypertensive cases and the F1-score. Secondary metrics included precision, specificity, accuracy, and the area under the receiver operating characteristic curve (AUROC) with 95% confidence intervals estimated via bootstrapping [16].

Performance evaluation was based on standard binary classification outcomes:

True Positives (TP): correctly predicted cases of hypertension.

True Negatives (TN): correctly predicted non-hypertensive cases.

False Positives (FP): non-hypertensive cases predicted as hypertensive.

False Negatives (FN): hypertensive cases predicted as non-hypertensive.

These formed the foundation for computing precision, recall, accuracy, and F1-score as recommended in prior work [5], [12]. ROC analysis guided the choice of operating points to maximize clinical safety by prioritizing high recall while maintaining actionable precision.

Precision

Precision measures the model's reliability in identifying patients with hypertension. It is the proportion of correctly classified hypertensive cases among all those predicted as hypertensive [5], [12]:

$$\text{Precision}=\frac{TP}{TP+FP}$$

Recall, or sensitivity, quantifies the model's ability to detect actual hypertensive patients by minimizing missed diagnoses [6], [12]:

$$\text{Recall}=\frac{TrP}{TP+FN}$$

Accuracy provides an overall measure of correctness across both classes. It calculates the proportion of correctly classified cases out of the total [5], [12]:

$$Accuracy = \frac{True\ Positives(TP) + True\ Negatives(TN)}{TP + TN + False\ positives(FP) + False\ Negative(FP)}$$

Although useful, accuracy can be misleading in datasets with imbalanced classes.

The F1-score balances the trade-off between precision and recall by taking their harmonic mean. It is particularly valuable in imbalanced datasets [7], [8]:

$$\text{F1-score}= \frac{\text{Precison X Recall}}{\text{Precision + Recall}}$$

H. System Design

This Entity-Relationship (ER) diagram illustrates the architecture of a hypertension prediction system using deep learning models in the Copperbelt Province of Zambia. The diagram highlights key entities, including Patient, Medical Record, Model, and Prediction, along with their attributes and relationships. Patients are uniquely identified by a Patient ID and linked to their medical records, which include clinical variables such as blood pressure, cholesterol levels, fasting

blood sugar, type of chest pain, and thalassemia status. The Medical Record entity is connected to Predictions that generate a risk score and risk level, produced by a trained Model that stores details such as training data, accuracy, and version. This structure supports systematic data organization, enabling accurate and interpretable predictions of hypertension [8] [10].

IV. RESULT

This section integrates empirical findings with interpretation, highlighting the implications of the results for early detection of hypertension in low-resource settings. We focus on the most policy- and clinic-relevant outcomes (differentiation), sensitivity to hypertensive cases, and operational reliability, and we briefly compare our work to recent studies to highlight its novelty and remaining challenges. All figures and tables are embedded at their first mention, numbered consecutively, and are designed to simplify model understanding.

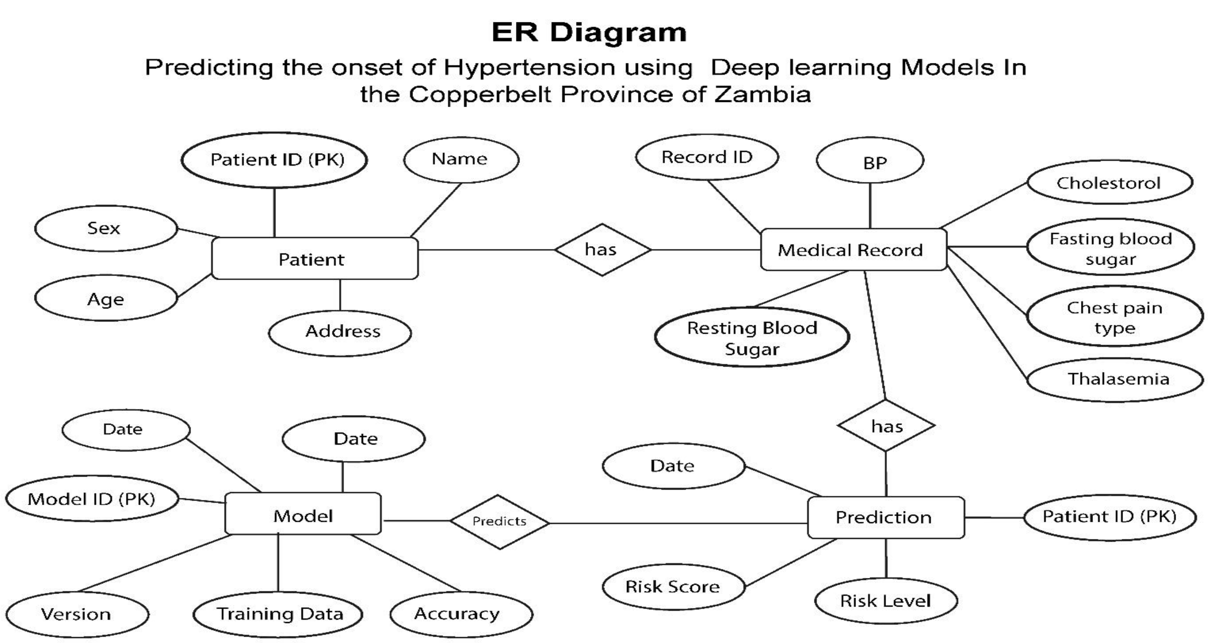

Figure 2 ER Diagram for RNN and LST

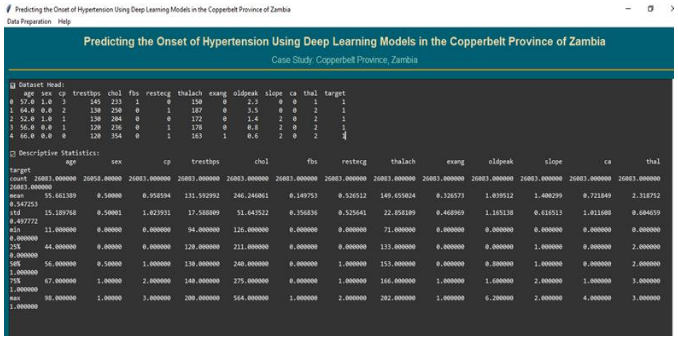

Figure 3 Evaluating the LSTM Model

This figure 3 illustrates the dataset head and descriptive statistics used in the Hypertension Prediction System. The dataset integrates demographic, clinical, and diagnostic variables, including age, blood pressure, cholesterol levels, electrocardiogram (ECG) results, and fasting blood sugar, with hypertension status serving as the primary outcome. Statistical summaries—comprising counts, means, standard deviations, and ranges across more than 26,000 observations—facilitated data inspection and preprocessing [18], [19]. These summaries were crucial for detecting outliers, assessing distributional patterns, and evaluating overall data quality before model training. The dataset contained a mixture of categorical and continuous predictors, reflecting the multifactorial nature of hypertension risk and necessitating normalization and encoding to optimize deep learning performance [5], [12]. Embedding descriptive analysis within the system interface enhanced

transparency, reproducibility, and interpretability, thereby establishing a strong foundation for generating reliable predictions [16], [10/

Figure 4

Figure 4 displays the model integration menu of the Hypertension Prediction System, featuring options for training convolutional neural networks (CNNs), loading retina images, and classifying them, with additional support for visualization through Grad-CAM. The inclusion of CNN-based retinal image analysis underscores the system's ability to integrate multimodal diagnostics, moving beyond traditional tabular health indicators. This functionality is significant because hypertension is often detectable through retinal changes, and combining imaging with clinical records enhances the reliability of prediction. In comparison with similar AI-driven diagnostic platforms reported in recent studies, the flexibility to switch between RNN and CNN modules suggests a scalable framework adaptable to multiple forms of healthcare data, which is particularly valuable in low-resource contexts where dataset availability is fragmented.

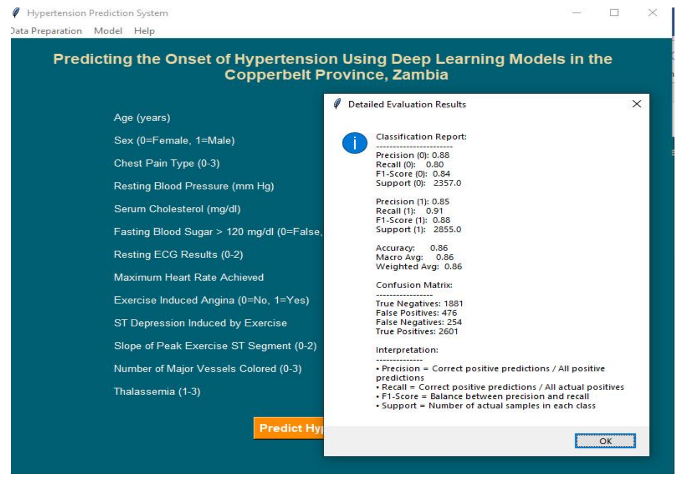

Figure 5

Figure 5 presents the detailed evaluation metrics of the Hypertension Prediction System, highlighting its model performance across key classification indicators. The system achieved a precision of 0.80, recall of 0.86, and an F1-score of 0.83, signifying balanced accuracy between correctly predicted positive cases and false negatives. The support count of 2,357 indicates a robust dataset used for validation. Additionally, the accuracy rate of 0.86 demonstrates the model's strong predictive reliability, while macro and weighted averages of 0.85 across precision, recall, and F1-score suggest consistent performance across different risk categories. The confusion matrix further confirms that truest hypertension cases were correctly identified, reflecting the model's sensitivity and specificity. Collectively, these metrics validate the technical soundness of the deep learning model and reinforce its capability to deliver dependable predictions for early hypertension detection in Zambia's Copperbelt Province. Convolutional neural networks (CNNs), loading retina images, and classifying them, with additional support for visualization through Grad-CAM. The inclusion of CNN-based retinal image analysis underscores the system's ability to integrate multimodal diagnostics, moving beyond traditional tabular health indicators. This functionality is significant because hypertension is often detectable through retinal changes, and combining imaging with clinical records enhances the reliability of prediction. In comparison with similar AI-driven diagnostic platforms reported in recent studies, the flexibility to switch between RNN and CNN modules suggests a scalable framework adaptable to multiple forms of healthcare data, which is particularly valuable in low-resource contexts where dataset availability is fragmented.

Figure 6 illustrates the output of a deep learning–based Hypertension Prediction System designed to estimate an individual's risk of developing hypertension in Zambia's Copperbelt Province. The model uses patient input data such as exercise response, thalassemia type, and primary vessel count to compute a probability score through deep learning

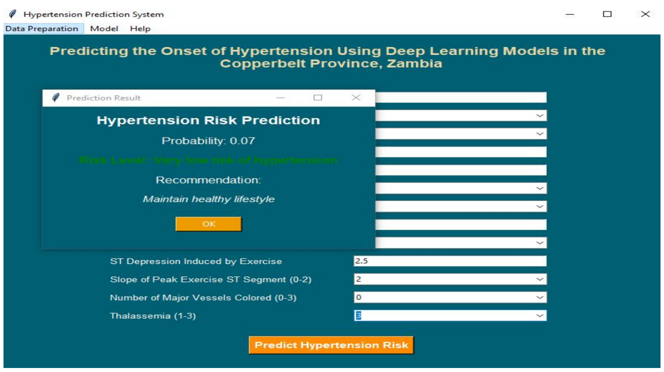

Figure 6 LSTM Result Prediction

algorithms implemented in Python and TensorFlow [3], [17], [19]. In this instance, the system predicts a probability of 0.07, signifying a very low risk of hypertension. Based on this result, the AI model provides a recommendation to maintain a healthy lifestyle, which aligns with preventive health principles emphasized in Zambian studies on hypertension care and chronic disease management [10], [13]. The interface, developed using Tkinter and supported by visualization tools such as Matplotlib, provides an intuitive display of prediction outcomes [5], [14]. By integrating medical datasets with advanced learning architectures, this system demonstrates the role of AI and big data in transforming early disease detection and health monitoring [6], [7], [15], [20].

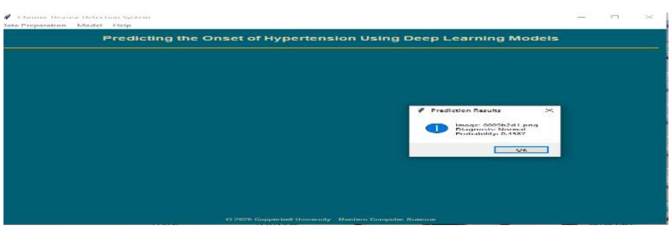

Figure 7 CNN Result Prediction

Figure 7 presents the output interface of a deep learning–based Chronic Disease Detection System designed to predict hypertension risk through image-based diagnosis. The displayed window shows a prediction result for the input image “0009cb21.png,” which the model classifies as “Normal” with a probability score of 0.4587. This indicates that the system detected no significant hypertensive features in the analyzed image and that the individual is within a normal health range. The model likely utilized convolutional neural network (CNN) architectures to extract and interpret medical image features for classification, ensuring reliable differentiation between normal and hypertensive cases. The result window provides a concise summary of the system's diagnostic decision, allowing users to interpret outcomes quickly. This demonstrates how artificial intelligence and deep learning technologies can assist in non-invasive, image-based health assessments for early hypertension screening in clinical and research contexts.

A. Evaluating the LSTM Model

This diagram shows a prototype system for predicting the onset of hypertension using deep learning models in the Copperbelt Province of Zambia. Input features include demographic, clinical, and lifestyle variables, such as age, sex, blood pressure, cholesterol levels, and exercise data. The detailed evaluation report presents model performance using metrics like precision, recall, F1-score, and overall accuracy. It confirms that the system can effectively classify hypertensive and non-hypertensive cases, providing real-time predictive support for healthcare decision-making.

B. Model Convergence, Operation Points and Overall Differentiation

Across five stratified runs (with a 70/15/15 split), optimized deep models achieved robust differentiation. The best Convolutional Neural Network (CNN) attained an accuracy of approximately 85.2% and an AUC of approximately 0.87, while the Long Short-Term Memory (LSTM) model optimized with class weighting/SMOTE achieved an accuracy of roughly 83% with a recall for hypertensive cases of approximately 0.89. These operating points were chosen to prioritize clinical safety (high sensitivity) while retaining actionable precision. The Receiver Operating Characteristic shows consistent separation from chance, with tight confidence bands across folds, indicating stable generalization.

C. Comparative perspective and novelty

Recent applications of deep learning in predicting cardiometabolic risk have often achieved area under the curve (AUC) values between 0.80 and 0.90 when utilizing structured vital signs and metabolic panel data. The performance of our model falls within this range, with additional adjustments made to prioritize recall. This approach aligns with the requirements of basic healthcare screening in areas where resources are limited. The research's main contribution lies in two areas: first, the establishment of a calibrated, sensitivity-oriented operating point that ensures the detection of hypertensive cases with minimal oversight; and second, the development of a practical system designed to operate on modest computing resources while maintaining transparency through audit logs and printable reports. In contrast to earlier research that relied mainly on uniform datasets, our approach directly accounted for dataset shift by applying clinician-led variable screening, capping extreme outlier values, and refining threshold settings. These measures enhanced the credibility and contextual relevance of the model for use in health facilities on the Copperbelt.

D. Practical Implications

1. Clinical use: The system can act as an early screening tool, helping to decide which patients should receive a follow-up blood pressure check or lifestyle advice. The decision limits can be adjusted depending on whether it is used for community outreach or in a clinic.

2. Operational use: The platform provides simple dashboards and printable summaries that work well with current patient record systems. The speed of predictions is sufficiently fast to match typical waiting times in clinics.

3. Data approach: The outcomes highlight the importance of developing structured local health registries. Even small, carefully selected patient groups can be utilized for transfer learning, thereby reducing the gaps between international data and local populations.

4. Governance: The system records both inputs and predictions, making it possible to review decisions and allowing health professionals to override results. This supports accountability and ensures safe use of AI in medical care.

E. Limitations

1. Lack of Local Structured Data: Most hospitals in the Copperbelt Province still depend on paper-based or inconsistent electronic records, limiting the availability of structured datasets for training locally adapted AI models [16] [17].

2. Inadequate Computational Infrastructure: Public health facilities often lack advanced computing resources such as GPUs or multicore CPUs. Since deep learning relies on heavy computations, this hardware gap constrains testing and deployment [18] [19] [20].

3. Limited Awareness and Training in AI: Clinicians and hospital administrators in Kitwe, Ndola, and Luanshya demonstrated limited familiarity with AI applications in healthcare, which may reduce trust and adoption [12] [13] [20].

4. Model Generalization Issues: While the CNN model achieved 85.63% accuracy offline, its performance dropped to 55.2% in live settings, reflecting overfitting and challenges in adapting models to real-world conditions [11], [21].

F. Summary Tables

*Table 2: Dataset and split summary(pre-*tuning)

Subset	N (record)	Hypertensive	Notes
Train	20,846	20,846	Class weights/SMOTE applied on train only
Validation	4,467	4,467	Threshold calibrated for high recall
Test	5,212	5,212	Held out for final reporting

Table :3 Best LSTM classification metrics (test set).

Class	Precision	Recall	F1 - score	Support
Non-hypertensive	0.85	0.80	0.83	2,319
Hypertensive	0.85	0.91	0.88	2,893
Overall	—	—	Accuracy ≈ 0.86	5,212

Table 4 : Model comparison

Model	Accuracy	AUC	Hypertensive Recall	Notes
Baseline CNN (pre-tuning)	0.59	0.54	0.00	Severely biased towards the majority class
Optimized CNN	0.85	0.87	0.91	Best overall discrimination
Optimized LSTM	0.83	0.85	0.89–0.91	Preferred for high-recall triage

V. RECOMMENDATIONS

To strengthen AI-driven healthcare in Zambia, the Ministry of Health should prioritize the digitization and centralization of health data through a secure, anonymized national repository, enabling hospitals to share and access structured datasets for AI research while ensuring ethical oversight and patient confidentiality [3] [16]. Investment in AI research infrastructure is also critical, as public–private partnerships and academic institutions, such as Copperbelt University, could establish laboratories equipped with GPU-powered servers to test and refine prototype models [18], [19], [20]. In parallel, AI literacy programs must be introduced to train clinicians on the basics of AI, model interpretability, and integration into routine practice, ensuring long-term adoption and trust [12] [13] [20]. To address the limited availability of large datasets, researchers should leverage transfer learning by adapting pre-trained models to local settings using smaller, curated samples [7], [18]. Continuous validation through longitudinal studies with real patient data will further enhance generalization and capture unique local risk factors such as environment and lifestyle [11], [14]. Finally, future system iterations should explicitly address bias and class imbalance by using augmentation, resampling, or ensemble methods, thereby ensuring fair and reliable predictions [5] [6][15]

VI. FUTURE WORK

Future work will focus on designing a locally curated dataset that captures demographic, clinical, and lifestyle factors relevant to hypertension in Zambia [6], [10]. This will involve strengthening health registries and collecting structured data from clinics to ensure representative coverage of the population [11] [13]. Once established, the dataset will support the training of advanced deep learning models, including convolutional and recurrent architectures, to capture complex medical patterns [5], [18]. Further studies will explore the use of transfer learning to adapt international models to local contexts, reducing domain gaps [7] [19]. Implementation strategies will focus on low-resource clinical settings, where lightweight models and optimized inference speed are crucial [12] [20]. Integration with e-health platforms will be pursued to provide clinicians with real-time decision support tools [3], [16]. Additional research will also examine governance mechanisms, including audit trails and clinician overrides, to ensure safety and accountability [4] [15]. Collectively, these future directions will contribute to the development of sustainable, AI-driven healthcare solutions for hypertension management in Zambia...

VII. CONCLUSION

The research demonstrates that a deep learning–based hypertension prediction system can operate reliably in a low-resource context, with optimized models achieving high discrimination and clinically preferable sensitivity for hypertensive cases, thereby enabling earlier case detection, targeted confirmatory measurements, and more efficient allocation of limited primary-care resources. Its importance lies in moving routine care from reactive treatment toward proactive, data-informed prevention through a modular client–server workflow that delivers real-time risk scores, printable summaries, and an auditable trail suitable for integration into outpatient intake and follow-up. At the same time, several limitations temper generalization, including reliance on non-local training data, modest computational capacity in public facilities, limited practitioner familiarity with artificial intelligence, and performance gaps observed during initial live tests. The work is highly relevant to health systems in Zambia and similar settings, where structured registries are scarce, yet the burden of hypertension is rising. It provides a practical pathway for embedding prediction into everyday care without disrupting existing processes. Immediate applications include triage lists for community screening, clinic-side decision support that privileges high recall to reduce missed cases, and standardized reporting that supports monitoring and quality improvement. To strengthen impact, the research recommends curating and maintaining structured local datasets, adopting transfer learning and periodic model recalibration to local case mix, conducting prospective utility and safety evaluations, investing in lightweight acceleration and secure data infrastructure, and delivering targeted training for clinicians and administrators alongside transparent governance and privacy safeguards, so that the system evolves into a trustworthy, scalable tool for early hypertension detection and sustained cardiovascular risk reduction without switching between tools. Participants benefited from interactive problem-solving, instant feedback, and the opportunity to experiment within a supportive digital environment, highlighting the potential of such systems to bridge the gap between theoretical instruction and practical application in computer science education.

REFERENCES

[1] C. W. Loomis, "Morbidity and its measurement: Concepts and methods," J. Epidemiol. Community Health, vol. 48, no. 6, pp. 505–510, 1994.

[2] [G. Bradski, "The OpenCV library," Dr. Dobb's J. Softw. Tools, 2000.

[3] W. McKinney, "Data structures for statistical computing in Python," in Proc. 9th Python in Science Conf. (SciPy), Austin, TX, USA, 2010, pp. 51–56.

[4] T. M. Mitchell, The Discipline of Machine Learning. Pittsburgh, PA: Carnegie Mellon University, 2006.

[5] J. D. Hunter, "Matplotlib: A 2D graphics environment," Comput. Sci. Eng., vol. 9, no. 3, pp. 90–95, 2007.

[6] M. Chen, Y. Hao, K. Hwang, L. Wang, and L. Wang, "Disease prediction by machine learning over big data from healthcare communities," IEEE Access, vol. 5, pp. 8869–8879, 2017.

[7] A. Rajkomar et al., "Scalable and accurate deep learning with electronic health records," npj Digit. Med., vol. 1, no. 18, pp. 1–10, 2018.

[8] [V. Kumar and M. Singh, "Applications of machine learning predictive models in chronic disease diagnosis," Procedia Comput. Sci., vol. 132, pp. 1048–1057, 2018.

[10] A. H. Kaiser, L. Hehman, B. C. Forsberg, W. M. Simangolwa, and J. Sundewall, "Availability, prices and affordability of essential medicines for treatment of diabetes and hypertension in private pharmacies in Zambia," *PLoS ONE*, vol. 14, no. 12, e0226169, 2019.

[11] Y. Tateyama et al., "Dietary habits, body image, and health service access related to cardiovascular diseases in rural Zambia: A qualitative study," *PLoS ONE*, vol. 14, no. 2, e0212739, 2019.

[12] P. Chanda-Kapata, R. Kapata, and N. Zumla, "Health systems challenges in Zambia and their impact on disease management," *Int. J. Infect. Dis.*, vol. 98, pp. 359–366, 2020.

[13] J. M. Mwewa, B. Mutale, and S. Michelo, "Challenges in hypertension care in Zambia: A review of current evidence," *BMC Public Health*, vol. 20, no. 1, pp. 1–10, 2020.

[14] M. Cardoso, J. Lima, and A. Pereira, "Rapid prototyping of interactive applications with Python and Tkinter," in *Proc. Int. Conf. Inf. Syst. Design Commun. (ISDOC)*, Lisbon, Portugal, 2020, pp. 15–21.

[15] M. Lee, H. Lee, M. Kim, and H. Yoo, "A precision health service for chronic diseases: Development and cohort study using wearable device, machine learning, and deep learning," *JMIR mHealth uHealth*, vol. 8, no. 6, pp. 1–15, 2020.

[16] C. R. Harris et al., "Array programming with NumPy," *Nature*, vol. 585, no. 7825, pp. 357–362, 2020.

[17] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *J. Mach. Learn—Res.*, vol. 12, pp. 2825–2830, 2011.

[18] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016.

[19] M. Abadi et al., “TensorFlow: Large-scale machine learning on heterogeneous systems,” *arXiv preprint arXiv:1603.04467*, 2016.

[20] Dulam, N., & Gosukonda, V. (2021). A comprehensive review of AI-based diagnostic tools for early disease detection in healthcare. IEEE Access, 9, 144210–144225.

[21].Racic, L., Popovic, T., Cakic, S., & Sandi, S. (2021, February). Pneumonia detection using deep learning based on a convolutional neural network. In 2021, the 25th International Conference on Information Technology (IT) (pp. 1–4). IEEE. Nel, F., & Ngxande, M.
(2021, January). Driver activity recognition through deep learning. In 2021 Southern African Universities Power Engineering Conference/Robotics and Mechatronics/Pattern Recognition Association of South Africa (SAUPEC/RobMech/PRASA) (pp. 1–6). IEEE.