Journal of Computer and Electrical Sciences (JCES)- Volume-1 Issue-4, April 2026
ISSN: 3049-2602

DEEP LEARNING-BASED CLINICAL DECISION SUPPORT SYSTEM FOR DIABETES AND BLOOD PRESSURE RISK PREDICTION

B. Venkata Sudeesh Kumar[1], Mr. B. Rajasekhar Reddy[2], P. Vivek[3], G. Yaswanth Reddy[4]
[1,3,4]Student, Department of Computer Science and Engineering (Data Science)
Andhra Loyola Institute of Engineering and Technology, Vijayawada, Andhra Pradesh, India
[2]*Assistant Professor, Department of CSE (Data Science)*
Andhra Loyola Institute of Engineering and Technology, Vijayawada, Andhra Pradesh, India
Email: sudeeshborra924@gmail.com, bomma689@gmail.com, vivekpallagani@gmail.com, gudayash06@gmail.com

Abstract: *Chronic diseases like diabetes mellitus and hypertension silently affect over a billion people worldwide, yet early intervention remains inaccessible for millions in low-resource settings. This paper presents a Deep Learning-Based Clinical Decision Support System (CDSS) that simultaneously predicts the risk of diabetes and blood pressure abnormalities using patient health parameters. Four predictive models were trained and compared — Logistic Regression, Support Vector Machine (SVM), Random Forest, and a Deep Neural Network (DNN) — on the PIMA Indians Diabetes Dataset (768 records) and Cardiovascular Disease Dataset (70,000 records). The DNN achieves the highest accuracy of 91%, precision of 0.89, recall of 0.88, F1-Score of 0.88, and AUC-ROC of 0.95. A rigorous preprocessing pipeline including SMOTE oversampling, median imputation, and StandardScaler normalization ensures data quality. The system is deployed as an accessible Flask web application with real-time predictions and personalized health recommendations, validated through unit testing, integration testing, and 10-fold cross-validation (mean accuracy 90.4%). This work demonstrates that an affordable, multi-disease AI screening tool can meaningfully contribute to preventive healthcare goals.*

Keywords: Deep Neural Network, Clinical Decision Support System, Diabetes Prediction, Hypertension, Flask, SMOTE, PIMA Dataset, Healthcare AI.

1. INTRODUCTION

Managing chronic non-communicable diseases has become one of the most pressing challenges facing global healthcare systems today. Diabetes mellitus and hypertension — individually serious conditions — are increasingly presenting together in what clinicians call cardiometabolic syndrome, dramatically compounding the risk of cardiovascular events, kidney damage, and early mortality. According to the World Health Organization (WHO), roughly 537 million adults live with diabetes globally and over 1.28 billion are hypertensive, yet a significant fraction of cases in developing countries like India remain undiagnosed until complications arise.

The fundamental obstacle is structural: traditional screening depends on laboratory tests, specialist consultations, and physical examinations — resources that are expensive, time-consuming, and unevenly distributed. Rural and semi-urban communities, where nearly 65% of India's population lives, often lack even basic diagnostic infrastructure. This gap makes a compelling case for intelligent, software-based tools that can perform preliminary health risk screening using commonly available parameters.

Machine learning and deep learning have demonstrated strong predictive capability on structured medical data. However, most existing systems either target a single disease or remain confined to academic settings without practical deployment. This paper addresses both gaps by proposing a unified CDSS that integrates deep learning-based prediction for two related diseases — diabetes

https://doi.org/10.5281/zenodo.19627813

and hypertension — within a single, freely accessible web application. The system is built, evaluated, and deployed following a complete software engineering process including UML design, testing, and performance analysis.

2. LITERATURE SURVEY

Computational approaches to diabetes prediction date back to Smith et al. (1988), who applied the ADAP learning algorithm to the PIMA Indians Dataset and achieved a benchmark accuracy of 76%. Kavakiotis et al. (2017) reviewed 85 publications over two decades and concluded that SVMs and neural networks consistently outperform rule-based methods on clinical classification tasks. LeCun, Bengio, and Hinton (2015) established the theoretical foundation for multi-layer deep learning, enabling feature learning directly from raw data — a capability particularly valuable in healthcare where interactions between clinical variables are complex and non-linear.

Zou et al. (2018) compared Naive Bayes, Decision Trees, and Random Forest on the PIMA dataset, with Random Forest achieving 77.21% accuracy and identifying glucose concentration as the dominant predictor. Hasan et al. (2020) applied ensemble learning combining KNN, Naive Bayes, and SVM classifiers, reaching 82.35%. Weng et al. (2017) demonstrated, using 378,256 patient records, that all four ML models tested outperformed the traditional ACC cardiovascular risk scoring tool — a pivotal result establishing ML's clinical value over conventional methods.

Despite strong individual performances, a recurring weakness across reviewed studies is the absence of deployed, user-accessible systems. Rajkomar et al. (2019) highlighted that prediction alone is insufficient; clinical utility requires interpretable, accessible interfaces integrated into real-world workflows. The proposed CDSS specifically addresses this deployment gap.

3. PROPOSED SYSTEM

The proposed CDSS is architected as a three-tier web application: a frontend Presentation Layer (HTML5, CSS3, JavaScript), an Application Layer (Flask backend), and a Model Layer (trained DNN artifacts). This separation ensures modularity and allows the predictive models to be updated independently of the user interface.

Two trained DNN models are deployed: one for diabetes prediction trained on the PIMA Indians Diabetes Dataset, and one for cardiovascular/BP risk trained on the 70,000-record Cardiovascular Disease Dataset. Each prediction request flows from the browser to Flask, through a preprocessing pipeline (StandardScaler normalization), through the DNN model, and back to the user as a risk level (High/Low), probability percentage, and personalized health recommendations — all within two seconds.

System Architecture:

- Presentation Layer: Responsive HTML forms for health parameter input, JavaScript validation, dynamic result rendering.
- Application Layer: Flask routes (/predict/diabetes, /predict/bp), input sanitization, JSON API responses, rate limiting (100 req/hr per IP).
- Model Layer: DNN model files (.h5) and preprocessing scalers (.pkl) loaded at startup; inference in 4.5ms per prediction.

Protection mechanisms include input range validation on both client and server sides, HTTPS encryption, and a prominent non-dismissible medical disclaimer on all pages confirming that predictions are for screening purposes only and do not replace clinical diagnosis.

4. METHODOLOGY

The development follows a five-phase data-driven pipeline:

Phase 1 – Data Collection and Exploration:

Both datasets are sourced from public repositories (UCI ML Repository, Kaggle). Exploratory Data Analysis using Pandas, Matplotlib, and Seaborn reveals feature distributions, outliers, class imbalance, and inter-feature correlations. The PIMA dataset shows a 65:35 class ratio (non-diabetic : diabetic); the cardiovascular dataset is approximately balanced.

Phase 2 – Preprocessing:

Zero values in clinical features (Glucose, BMI, Blood Pressure, Skin Thickness, Insulin) are biologically implausible and are replaced with feature-specific medians. SMOTE oversampling is applied to the PIMA dataset to correct class imbalance, creating synthetic minority samples through nearest-neighbor interpolation. All features are then standardized using StandardScaler. An 80:20 stratified train-test split is applied.

Phase 3 – Model Training:

Four models are trained: Logistic Regression (max_iter=1000), SVM (RBF kernel, probability=True), Random Forest (100 estimators), and the proposed DNN. The DNN is built using TensorFlow/Keras with four layers: Input (8 or 11 features) → Dense(128, ReLU, BatchNorm, Dropout 0.3) → Dense(64, ReLU, BatchNorm, Dropout 0.3) → Dense(32, ReLU) → Dense(1, Sigmoid). The Adam optimizer with learning rate 0.001 and binary cross-entropy loss are used. EarlyStopping (patience=15) prevents overfitting.

Phase 4 – Evaluation:

All models are evaluated on the held-out test set using Accuracy, Precision, Recall, F1-Score, and AUC-ROC. 10-fold cross-validation confirms generalizability. McNemar's test verifies that performance differences are statistically significant ($p < 0.05$).

Phase 5 – Deployment:

The best DNN models are serialized and integrated into the Flask application. The full system is tested through unit, integration, system, and user acceptance testing before deployment.

5. RESULTS AND DISCUSSION

Table 1 compares all four models on five evaluation metrics across the PIMA diabetes test set (20% split, 154 samples). All models used identical preprocessing pipelines to ensure fairness.

Table 1: Model Performance Comparison – Diabetes Prediction (PIMA Dataset)

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
Logistic Regression	78%	0.74	0.71	0.72	0.82
SVM	82%	0.79	0.76	0.77	0.87
Random Forest	86%	0.84	0.81	0.82	0.91
DNN (Proposed)	91%	0.89	0.88	0.88	0.95

The DNN achieves a 13 percentage-point accuracy gain over Logistic Regression, 9 points over SVM, and 5 points over Random Forest. More clinically significant is the recall improvement: the DNN correctly identifies 88% of all actual diabetic cases (vs 81% for Random Forest), meaning 7 fewer missed diagnoses per 100 at-risk patients — a meaningful reduction in potentially dangerous

false negatives. The 10-fold cross-validation mean accuracy of 90.4% (SD = 0.77%) confirms the model generalizes well without overfitting.

On the cardiovascular dataset, the DNN achieves 87.6% accuracy and 0.92 AUC-ROC, again outperforming all baseline models, though slightly lower than diabetes performance owing to higher complexity in the 70,000-record dataset. McNemar's test confirms all pairwise model differences are statistically significant (DNN vs Random Forest: $p = 0.022$; DNN vs Logistic Regression: $p < 0.001$).

Feature importance analysis using permutation importance and SHAP values consistently ranks Glucose (SHAP = 0.243), BMI (0.198), and Age (0.156) as the top three predictors — findings that align precisely with established clinical evidence. The model's learned feature hierarchy mirrors clinical diagnostic logic, supporting its trustworthiness in practice.

Table 2: System Test Case Results

Test Case	Input Parameters	Expected Output	Actual Output	Status
TC-01	Glucose:90, BMI:22.5, Age:30	Low Risk	Low Risk – 18%	PASS
TC-02	Glucose:180, BMI:35, Age:52, Insulin:200	High Risk	High Risk – 87%	PASS
TC-03	Systolic BP:165, Diastolic:105	High BP Risk	High Risk – 79%	PASS
TC-04	BMI:42, Age:45, Glucose:130	High Risk	High Risk – 76%	PASS

User Acceptance Testing with 15 participants yielded an average satisfaction score of 4.5/5, with response speed rated 4.7/5 (94%) and ease of use 4.6/5 (92%). Trust in predictions scored 3.6/5 — an appropriate and expected finding confirming users correctly perceive the system as a screening aid rather than a clinical diagnosis tool, validating the effectiveness of the medical disclaimer design.

6. CONCLUSION

This paper has presented a fully deployed, multi-disease Deep Learning-Based CDSS that addresses a real and urgent gap in preventive healthcare access. By integrating rigorous preprocessing, a high-accuracy DNN model (91% accuracy, 0.95 AUC-ROC), and a user-friendly web interface into a single cohesive system, this work moves beyond the typical research prototype to deliver a practical, testable application with demonstrated clinical relevance.

The DNN consistently outperforms all three baseline models across every evaluation metric, and these improvements are statistically verified. Feature importance analysis confirms the model's predictions align with clinical expectations, and comprehensive testing validates its reliability. The system's no-PII-storage policy, HTTPS encryption, rate limiting, and prominent medical disclaimer ensure responsible deployment.

Future work will integrate Explainable AI (SHAP-based per-prediction explanations), federated learning for privacy-preserving model improvement across institutions, wearable device data integration for continuous monitoring, and multilingual support for India's linguistically diverse

population. The system as presented demonstrates that affordable, AI-powered health screening tools can be built, validated, and deployed with the rigor required for real-world healthcare impact.

REFERENCES

1. *Smith, J. W., Everhart, J. E., Dickson, W. C., Knowler, W. C., & Johannes, R. S. (1988). Using the ADAP learning algorithm to forecast the onset of diabetes mellitus. Proceedings of the Annual Symposium on Computer Application in Medical Care, 261–265.*
2. *Kavakiotis, I., Tsave, O., Salifoglou, A., Maglaveras, N., Vlahavas, I., & Chouvarda, I. (2017). Machine learning and data mining methods in diabetes research. Computational and Structural Biotechnology Journal, 15, 104–116.*
3. *LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444.*
4. *Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. New England Journal of Medicine, 380(14), 1347–1358.*
5. *Weng, S. F., Reps, J., Kai, J., Garibaldi, J. M., & Qureshi, N. (2017). Can machine-learning improve cardiovascular risk prediction using routine clinical data? PLOS ONE, 12(4), e0174944.*
6. *Zou, Q., Qu, K., Luo, Y., Yin, D., Ju, Y., & Tang, H. (2018). Predicting diabetes mellitus with machine learning techniques. Frontiers in Genetics, 9, 515.*
7. *Hasan, M. K., Alam, M. A., Das, D., Hossain, E., & Hasan, M. (2020). Diabetes prediction using ensembling of different machine learning classifiers. IEEE Access, 8, 76516–76531.*
8. *Sisodia, D., & Sisodia, D. S. (2018). Prediction of diabetes using classification algorithms. Procedia Computer Science, 132, 1578–1585.*
9. *Mohan, V., et al. (2019). A novel machine learning based hybrid classification model for prediction of heart disease. International Journal of Advanced Computer Science and Applications, 10(5), 234–242.*
10. *Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. Journal of Artificial Intelligence Research, 16, 321–357.*
11. *Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30.*
12. *World Health Organization. (2021). Global report on diabetes. WHO Press.*