Multimodal Machine Learning for Stroke Risk Prediction in Nigerian Hospitals

Eze Chimdiya Chiemeka, **Lucky Nwabriaije Ahiakwo**, **Adline Chidinma Ekine-Pakaye**, **Johnson Chika Benjamin**, **Favour Ahiakwo**

Federal College of Education (Tech) Omoku, Rivers State, Nigeria

Corresponding author: Flobaba2222@gmail.com

ORCID: 0009-0009-8355-5207

Abstract

Non-communicable diseases are a growing cause of death in Nigeria, and ischaemic stroke, driven largely by hypertension and diabetes, is among the most damaging. Hospitals hold structured records and clinical notes that are incomplete and rarely used for risk prediction. This paper adapts a convolutional neural network-based multimodal risk prediction algorithm (CNN-MDRP), originally developed on hospital data from central China, to Nigerian conditions and tests the pipeline in a simulation. The framework combines latent factor imputation, expert-guided feature selection, convolutional feature learning from clinical notes and feature fusion, and is compared with four conventional baselines. Because no real Nigerian dataset was available, a synthetic cohort of 3,200 patients from three simulated hospitals was generated, with missing data and notes containing Nigerian Pidgin, and the models were evaluated by hold-out testing, bootstrap confidence intervals, calibration and leave-one-site-out validation. On these simulated data, CNN-MDRP reached an AUC of 0.90 (95% CI 0.87–0.93) against 0.78 for the best structured-data baseline, with similar results at each held-out site, while its gain over a text-only network was small. These figures show that the pipeline works, but they depend on assumptions built into the data and are not evidence of clinical performance. A multimodal, imputation-aware design therefore appears technically feasible for incomplete Nigerian hospital records, and validation on real, ethically approved data is the next step.

Keywords: disease risk prediction; machine learning; convolutional neural network; electronic health records; stroke; Nigeria

1. INTRODUCTION

Nigerian health planning has long centred on infectious disease, but the picture is changing. The World Health Organization estimated that non-communicable diseases (NCDs) caused about 29% of deaths in Nigeria in 2018, with cardiovascular disease the largest contributor at roughly 11% (World Health Organization [WHO], 2018), and the Federal Ministry of Health has responded with a national multisectoral action plan (Federal Ministry of Health, 2019). Hypertension is central: a meta-analysis put adult prevalence close to 30%, with wide variation between studies (Adeloye et al., 2015), and diabetes prevalence is roughly 6% (Uloko et al., 2018). The SIREN case-control study, run at 15 sites in Nigeria and Ghana, found hypertension to be the dominant modifiable risk factor for stroke (Owolabi et al., 2018), a conclusion echoed in wider African reviews (Akinyemi et al., 2021). Because many patients are undiagnosed until an acute event, tools that flag high-risk patients from data already collected in routine care have clear practical value.

Routine care does generate such data: demographics, vital signs, laboratory results and clinicians' notes. Yet electronic health record (EHR) adoption in sub-Saharan Africa is slow and uneven (Odekunle et al., 2017), so records are often incomplete and inconsistent between sites, and methods built on curated data from high-resource systems cannot simply be transferred. Chen et al. (2017) proposed CNN-MDRP for cerebral infarction, reconstructing missing values with a latent factor model and combining structured features with features learned from clinical text; they reported 94.8% accuracy on hospital data from central China. Combining both data types suits Nigeria, where laboratory data are sparse but narrative notes are comparatively rich, but a single-region result cannot be assumed to transfer.

This paper makes four contributions. It adapts CNN-MDRP to the tiered Nigerian health system; it builds consent, de-identification and Nigeria Data Protection Act compliance into the design; it treats incomplete records and locally varied clinical language as the normal case; and it specifies an evaluation protocol and tests the full pipeline in a simulation, with the limits of that exercise stated openly.

2. RELATED LITERATURE

Big data and health records. Pooling patient records can reveal patterns that single clinicians cannot see (Bates et al., 2014; Groves et al., 2013), but depends on data quality and on infrastructure for sharing and storage (Jensen et al., 2012; Chen et al., 2014; Wang et al., 2017). Where electricity and connectivity vary between facilities, practical systems must run on modest hardware and tolerate missing data. **Risk prediction.** Conventional models such as logistic regression, naive Bayes and decision trees usually rely on a few features chosen from clinical experience on modest datasets, and most published risk scores were derived in populations whose calibration on Nigerian patients is uncertain.

Clinical text. Convolutional neural networks apply learned filters over word vectors and keep the strongest local signals, and work well for sentence classification (Kim, 2014). Nigerian notes add complications: English is the working language, but notes may include abbreviations, local disease names and code-switching with Pidgin, Igbo, Yoruba or Hausa, and models trained on high-resource languages often perform poorly on African-language text (Adelani et al., 2021). **Missing data.** Deletion discards information and can bias results, whereas latent factor models, popularised in recommender systems (Koren et al., 2009), place patients and attributes in a shared low-dimensional space and reconstruct absent entries from their product. **Gaps.** Few multimodal risk studies target sub-Saharan Africa, text models are rarely tested on local clinical language, and performance is usually reported from one site without external validation.

3. PROPOSED FRAMEWORK

The framework first reconstructs missing structured values with a latent factor model. Priority diseases are identified from hospital statistics (here, ischaemic stroke), and structured features are chosen with clinicians. Features are learned automatically from clinical notes by a convolutional network; both feature sets are concatenated and passed to fully connected layers that output high or low risk. Figure 1 shows where the framework sits in the Nigerian health system and Figure 2 its architecture. Using both data types should improve discrimination (as Chen et al. reported), automatic text features reduce manual feature engineering, and imputation lets incomplete records contribute to training. Whether these advantages hold in Nigerian hospitals is an empirical question. Table 1 summarises the adaptation from the original study.

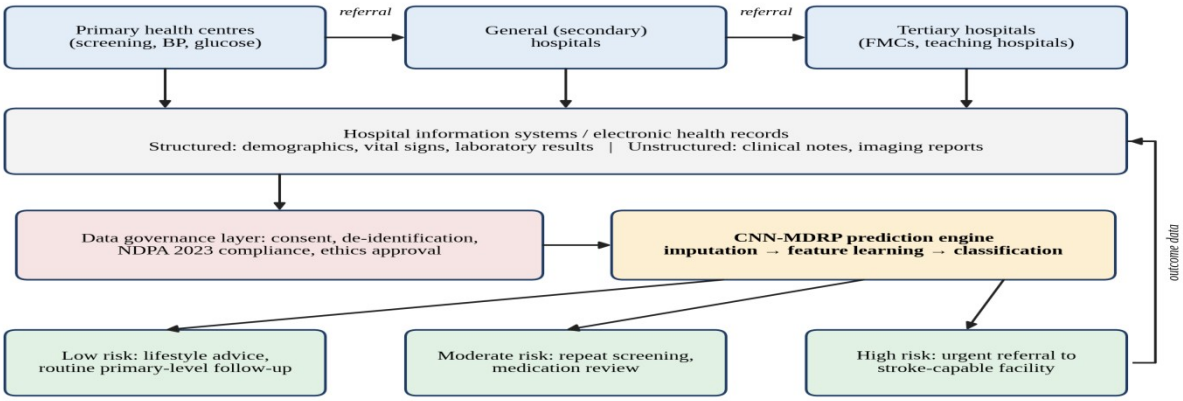

Figure 1. Data flow of the framework across the tiered Nigerian health system. NDPA, Nigeria Data Protection Act; FMC, federal medical centre.

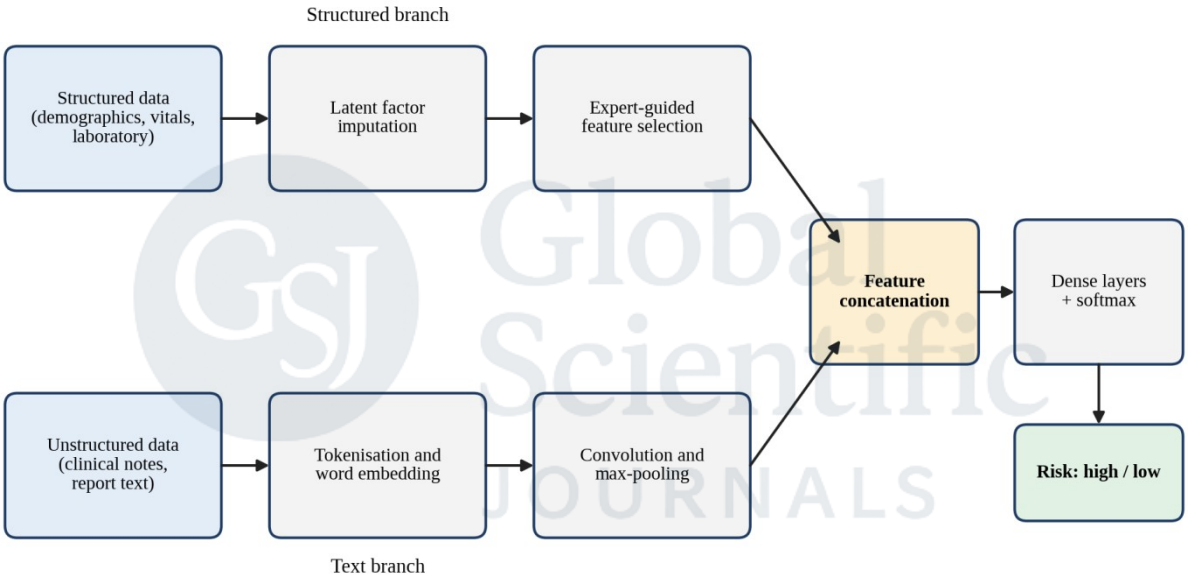

Figure 2. Architecture of the adapted CNN-MDRP model.

Table 1. *Adaptation of the original study design to the Nigerian setting.*

Aspect	Original study (Chen et al., 2017)	Adaptation proposed here
Setting	Hospital in central China, 2013–2015	Tertiary and secondary hospitals in southern Nigeria; primary health centres as referral sources
Target disease	Cerebral infarction	Ischaemic stroke, with hypertension and diabetes as main drivers (Owolabi et al., 2018)
Text data	Chinese-language clinical text	English notes with local terms, abbreviations and code-switching; validated on local notes
Governance	Not the focus of the report	Consent, de-identification, Nigeria Data Protection Act 2023, ethics review
Validation	Single-region data	Hold-out, leave-one-site-out, calibration and subgroup analysis

4. METHODOLOGY

4.1 Design, data and governance

The proposed study is a retrospective multi-site analysis of routine hospital data in southern Nigeria, covering urban and semi-urban catchments, followed by prospective validation. The outcome is a documented ischaemic stroke diagnosis (confirmed by neuroimaging where available), compared with patients having at least one major vascular risk factor and no stroke. Sample size will follow published guidance for prediction models (Riley et al., 2020) and reporting will follow TRIPOD (Collins et al., 2015). Structured data (Table 2) and the free-text parts of the record (complaints, history, examination, imaging reports) will be extracted. Approval will be sought from each institution's ethics committee under the National Code of Health Research Ethics (National Health Research Ethics Committee of Nigeria [NHREC], 2007). Processing will follow the Nigeria Data Protection Act (Federal Republic of Nigeria, 2023): records are de-identified at source, including names and addresses in free text, held on access-controlled institutional servers, and only aggregated results leave the institutions. The overall workflow is shown in Figure 3.

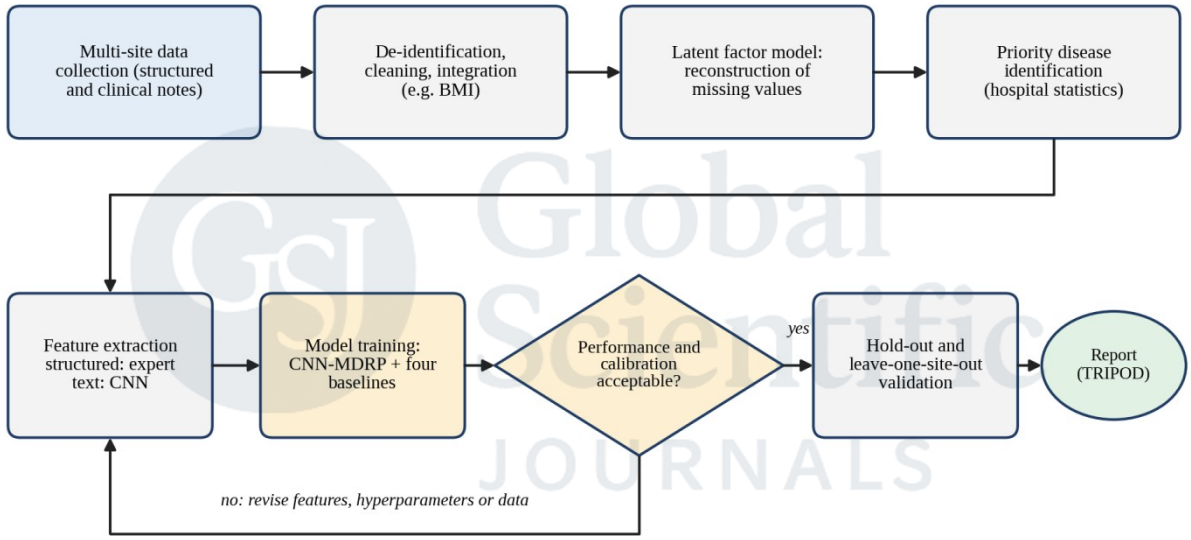

Figure 3. Methodological workflow from data collection to reporting.

Table 2. *Candidate structured predictors.*

Group	Examples	Relevance
Demographic	Age, sex, residence (urban or rural)	Risk varies by age and sex; residence proxies access to care
Vital signs	Blood pressure, heart rate, body mass index	Hypertension is the dominant modifiable risk factor
Laboratory	Fasting glucose, lipids, haemoglobin genotype	Diabetes and dyslipidaemia; sickle cell disease is common in Nigeria
History	Known hypertension or diabetes, adherence, smoking, alcohol, family history	Control of existing disease; modifiable behaviours
Cardiac	ECG findings, atrial fibrillation	Cardio-embolic pathway

4.2 Imputation and structured features

Records are checked for implausible values and duplicates and then integrated (height and weight become body mass index). Let R be the patient-by-attribute matrix with observed

entries in a set Ω. The latent factor model approximates R by the product of a patient-factor matrix P and an attribute-factor matrix Q, found by minimising

$$\Sigma_{(i,j)\in\Omega} (r_{ij} - p_i \cdot q_j)^2 + \lambda (\|p_i\|^2 + \|q_j\|^2)$$

and missing entries are filled from the product of P and Q. Rank and λ are tuned inside the training folds only, to avoid leakage. Because missingness may itself be informative (a test not ordered may signal limited access), the proportion missing is reported by site and variable, and missingness indicators are used as features. Clinicians propose features with established links to stroke, and statistical screening removes near-constant or near-duplicate variables.

4.3 Text branch and fusion

Notes are tokenised, standardised for common local abbreviations, padded to n words and mapped to d-dimensional embeddings. Filters of several widths slide across the embedding matrix to detect local patterns (for example, a phrase describing sudden weakness), max-over-time pooling keeps the strongest response of each filter, and the pooled values are concatenated and passed through dropout (Figure 4). The text vector is then concatenated with the structured vector and passed through dense layers to a softmax output, trained with class-weighted cross-entropy. A text-only version (CNN-UDRP, as in the original study) is trained for comparison.

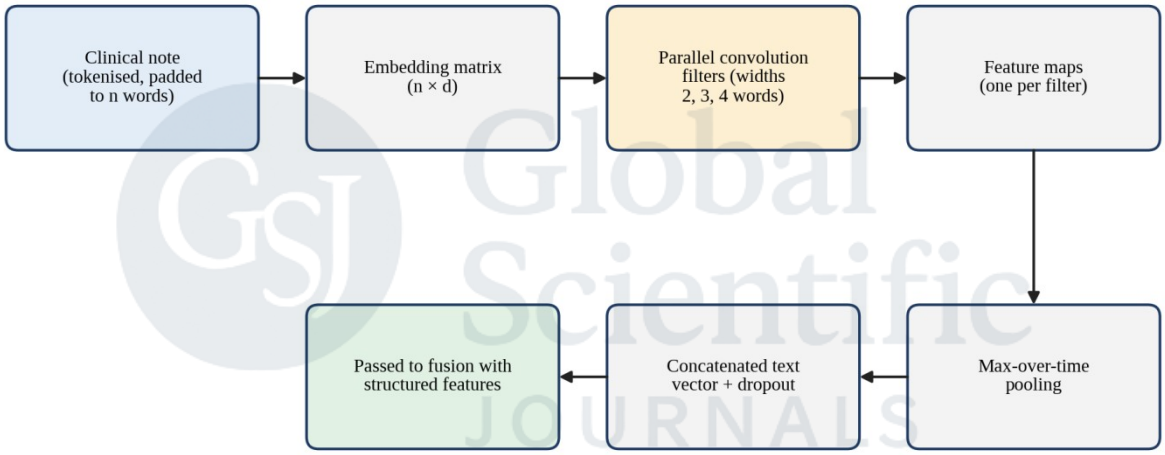

Figure 4. Convolutional text branch.

4.4 KNN baseline and evaluation

KNN is retained as a lightweight baseline because it needs no explicit training step, which suits limited computing resources (Cover & Hart, 1967). After min-max scaling, distances to all training records are computed, the k nearest are selected, and an inverse-distance-weighted vote gives a risk score (Figure 5). k is chosen by cross-validation; the value of about 10 found by Chen et al. (2017) is data-specific and is re-tuned. KNN is sensitive to feature scaling and slow on large datasets, and should not be confused with k-means clustering.

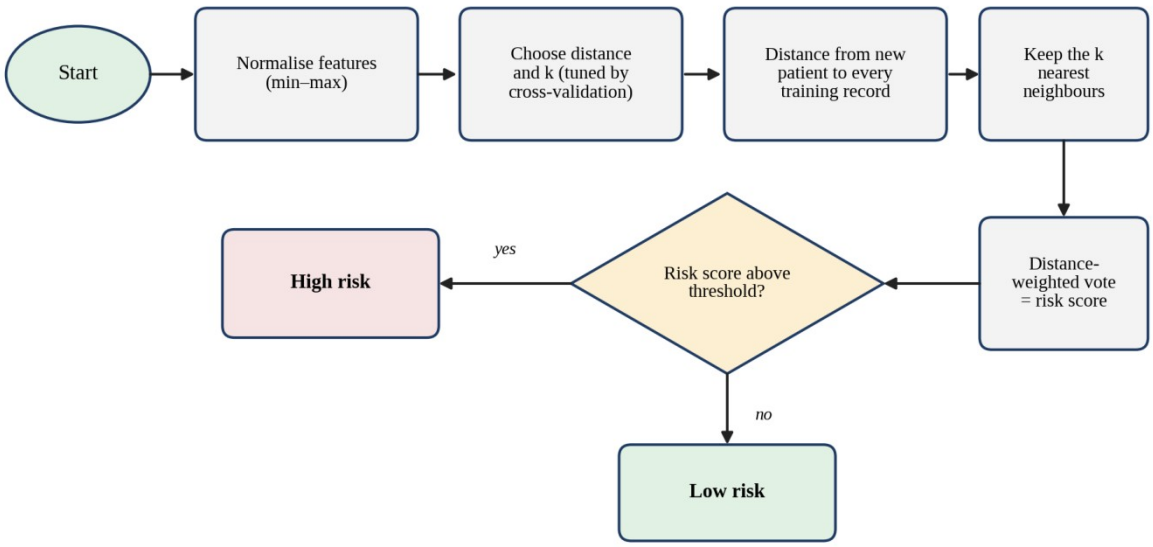

Figure 5. Flowchart of the KNN classification procedure.

Evaluation (Figure 6) uses an 80% development set with 10-fold stratified cross-validation for tuning, a 20% hold-out test set used once, and leave-one-site-out external validation. Measures are accuracy, sensitivity, specificity, precision, F1, area under the ROC curve (AUC), calibration and Brier score, each with 95% confidence intervals, plus subgroups by sex, age and catchment. Accuracy is not the principal measure because stroke cases are fewer than non-cases. Table 3 lists the models compared.

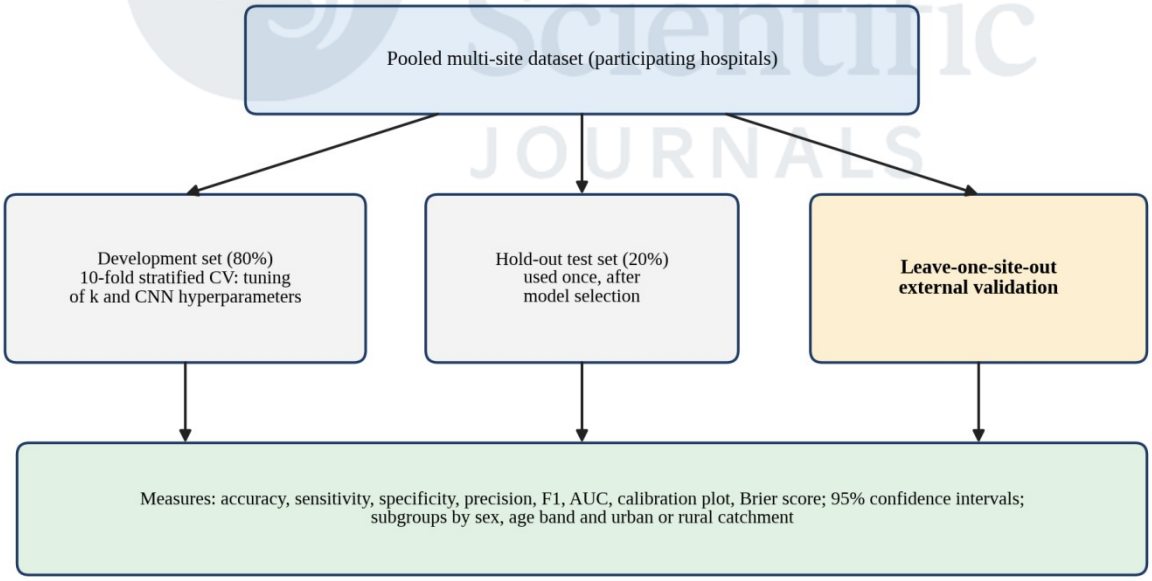

Figure 6. Pre-specified evaluation design.

Table 3. *Models compared.*

Model	Input	Purpose
Logistic regression, naive Bayes, decision tree	Structured	Conventional benchmarks (naive Bayes was used in the original study)
KNN	Structured	Lightweight baseline for low-resource settings
CNN-UDRP	Text only	Value of unstructured data alone
CNN-MDRP	Structured and text	Proposed multimodal model

5. SIMULATION EXPERIMENT AND RESULTS

5.1 Purpose and synthetic data

No real Nigerian hospital dataset was available, so to test the pipeline end to end a synthetic cohort was generated by the author. **All results below come from synthetic data; no real patient was involved.** The relationships between risk factors and stroke were set by the author, using directions of effect from the stroke literature (Owolabi et al., 2018), so the results show whether the design can recover assumed relationships and not how well it would predict stroke in real patients.

The cohort has 3,200 adults from three simulated hospitals: Site A, an urban teaching hospital (n = 1,300); Site B, a semi-urban federal medical centre (n = 1,100); and Site C, a general hospital (n = 800). Stroke prevalence was set at about 30%, as in a case-enriched hospital sample. Stroke status was drawn from a logistic function in which blood pressure, age and atrial fibrillation contributed most. Missing values rose from Site A to Site C and were higher for rural patients (Table 4). Each patient has a short note covering complaints, history and examination, written before diagnosis. Stroke symptoms appear in only a minority of stroke notes, non-stroke patients sometimes report overlapping symptoms such as headache, 10% to 40% of complaints (Site A to Site C) use Nigerian Pidgin (for example, “body dey weak”, “mouth don twist”), and spelling errors are added. The generator was changed once, before the final run, because the first version made the notes so informative that the text model reached an AUC of about 0.99. Data and code are available on request.

Table 4. *Synthetic cohort characteristics by simulated stroke status (n = 3,200).*

Characteristic	Stroke (n = 971)	No stroke (n = 2,229)	Missing, %
Age, years, mean (SD)	61.2 (12.4)	52.0 (11.9)	0.0
Male sex, %	52.1	49.9	0.0
Rural residence, %	33.1	29.9	0.0
Known hypertension, %	41.0	17.5	0.0
Poor treatment adherence, %	25.5	7.0	0.0
Systolic BP, mmHg, mean (SD)	135.4 (18.6)	120.6 (17.6)	2.9
Diastolic BP, mmHg, mean (SD)	83.2 (11.1)	78.0 (10.1)	2.9
Known diabetes, %	15.0	9.5	0.0
Fasting glucose, mmol/L, mean (SD)	5.8 (2.0)	5.6 (1.6)	26.4
LDL cholesterol, mmol/L, mean (SD)	3.3 (0.8)	3.0 (0.8)	41.5
Atrial fibrillation on ECG, % of those tested	5.3	1.3	43.3

5.2 Implementation

The latent factor model (rank 3, regularisation 1.0, alternating least squares) was fitted on training rows only, and test rows were imputed from their own observed values, so no test information entered training; missingness indicators were added for five variables. The text network (16-dimensional embeddings; 16 filters each of widths 2, 3 and 4; dropout 0.3; 32-unit hidden layer) was written in NumPy with gradients verified numerically, and trained with Adam, class-weighted cross-entropy and early stopping on a validation split. Notes were limited to 48 words and the vocabulary was 170 tokens. KNN k was chosen by five-fold cross-validation (25, the largest value tried, so a larger k might do better) and decision-tree depth likewise (3). The split was 80:20 (2,560 training, 640 test). Unlike the protocol in Section 4.4, the network was tuned on one validation split and not by 10-fold cross-validation.

5.3 Imputation

When extra observed values were hidden and reconstructed, the latent factor model had lower error than mean imputation at every masking rate: the root mean square error (standardised units) was 0.87 against 1.01 at 10% masking and 0.96 against 1.02 at 40%. The advantage was modest and did not improve classification: logistic regression reached an AUC of 0.780 with latent factor imputation and 0.781 with mean imputation.

5.4 Predictive performance

Table 5 and Figure 7 give hold-out results. The two neural models clearly outperformed the structured-data baselines: the AUC of CNN-MDRP was 0.90, against 0.78 for logistic regression and 0.72 to 0.74 for the others. The gap is large because, in this simulation, the notes carry presenting symptoms that the structured variables lack.

The difference between the multimodal and text-only networks is much smaller (AUC 0.90 against 0.88), with overlapping confidence intervals, so fusion cannot be said to improve discrimination. Accuracy was slightly higher for the text-only network (83.9% against 83.0%), but the multimodal model detected more strokes (sensitivity 77.3% against 67.0%) at lower specificity. Part of the structured information also appears in the notes, which limits what fusion can add. KNN's accuracy of 74.2% mainly reflects the larger non-stroke group; it found only 32.0% of strokes, which shows why accuracy alone misleads when classes are unequal.

Table 5. *Hold-out performance on synthetic data (n = 640; 194 stroke cases; threshold 0.5). AUC with 95% bootstrap confidence interval.*

Model	Acc. %	Sens. %	Spec. %	F1 %	AUC (95% CI)	Brier
Logistic regression	70.3	69.6	70.6	58.7	0.78 (0.74–0.82)	0.189
Naive Bayes	73.3	44.8	85.7	50.4	0.74 (0.70–0.78)	0.231
Decision tree	70.8	55.7	77.4	53.6	0.72 (0.68–0.76)	0.208
KNN	74.2	32.0	92.6	42.9	0.72 (0.68–0.77)	0.181
CNN-UDRP (text only)	83.9	67.0	91.3	71.6	0.88 (0.85–0.91)	0.115
CNN-MDRP (multimodal)	83.0	77.3	85.4	73.3	0.90 (0.87–0.93)	0.118

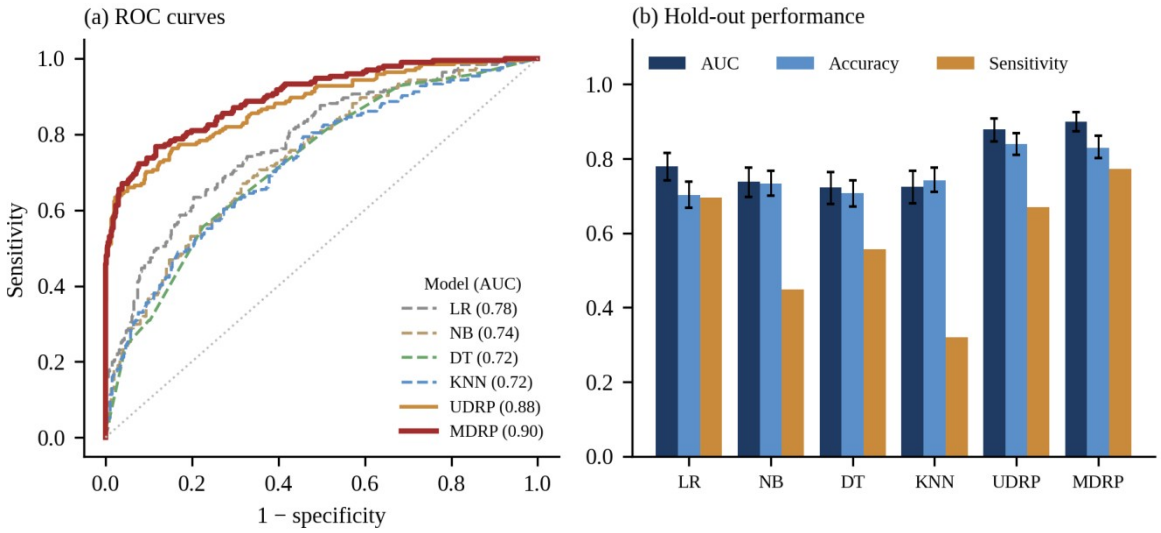

Figure 7. Hold-out results (simulated data). (a) ROC curves; (b) AUC, accuracy and sensitivity with 95% bootstrap intervals. LR, logistic regression; NB, naive Bayes; DT, decision tree; UDRP, CNN-UDRP; MDRP, CNN-MDRP.

5.5 Calibration and external validation

Calibration was weaker than discrimination (Figure 8a). Because the networks were trained with class weights, their probabilities run above observed proportions in the mid-range and should not be read as absolute risks; recalibration on a separate sample would be needed before showing probabilities to clinicians. Their Brier scores (0.118 and 0.115) were nonetheless well below those of the baselines (0.181 to 0.231). In leave-one-site-out testing (Figure 8b) the pattern held: CNN-MDRP reached AUCs of 0.88, 0.90 and 0.90 at Sites A, B and C, and baselines ranged from 0.70 to 0.79. The text models did not lose accuracy at Site C despite its heavier use of Pidgin, but this is weak evidence: the synthetic Pidgin comes from a small fixed list that the model can learn from other sites, which real clinical language will not respect.

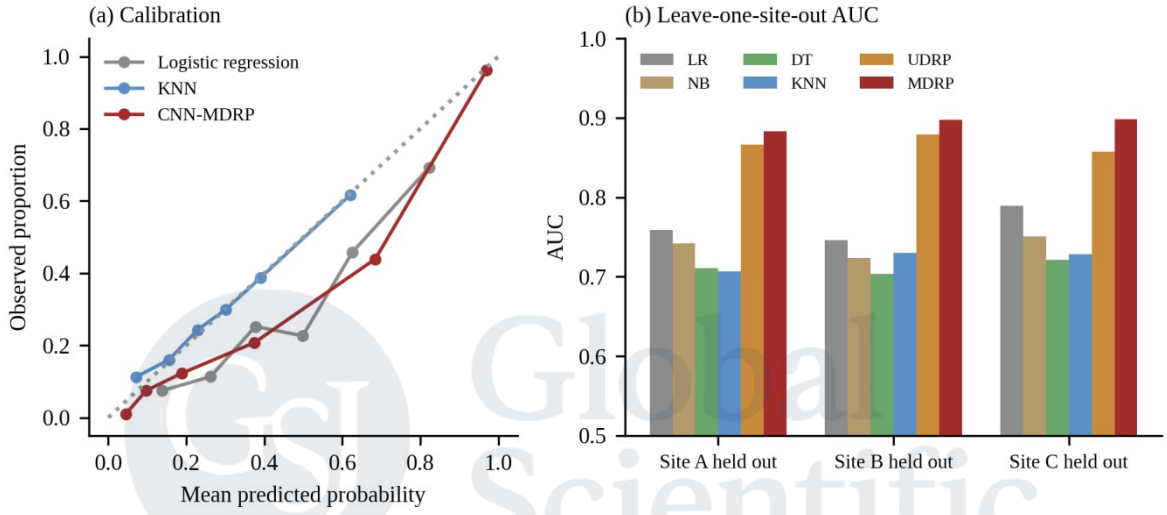

Figure 8. Calibration and external validation (simulated data). (a) Calibration on the hold-out set, six equal-sized probability groups; (b) AUC when each site is held out.

5.6 Relation to the original benchmark

The accuracy of 83.0% obtained here must not be compared with the 94.8% of Chen et al. (2017): datasets, outcomes, class balance and noise differ, and ours is synthetic. The useful comparison is qualitative: in both, adding text-derived features gave a clear gain over structured models, and the neural model was well ahead of conventional classifiers.

6. DISCUSSION

The simulation suggests the design works technically; the question is what would make it work in practice. At primary health centres, where blood pressure and glucose checks are most accessible, a lightweight structured-data model could help prioritise referral; at secondary and tertiary hospitals, where notes and laboratory data are richer, the multimodal model could support review of high-risk patients. Linking outputs to the national NCD action plan and to health insurance arrangements under the National Health Insurance Authority Act (Federal Republic of Nigeria, 2022) could in time support targeted screening, though that depends on policy choices beyond this paper.

Several challenges deserve open discussion. Latent factor imputation cannot recover information that was never collected, and in the simulation it improved reconstruction but not prediction. On structured data, well-tuned simple models are often hard to beat, and the benefit of text depends on note length and quality. Code-switching may degrade the text branch, so local validation is essential. Models needing graphics hardware or stable connectivity will not reach many facilities, which is why simple baselines are retained. If records come mainly from urban tertiary hospitals, the model may under-perform for rural patients who present late, which

the subgroup analyses are designed to expose. A risk score is decision support and not a diagnosis; patients should be informed in line with the Nigeria Data Protection Act, and deployment should follow prospective validation and ethics review.

7. LIMITATIONS AND FUTURE WORK

The central limitation is that the empirical results come from synthetic data. The strength of the text signal, the pattern of missing data and the risk-factor relationships were set by the author, so the figures describe the simulation and not real patients. Real notes are likely to be less informative, more varied and noisier, and the margin of the neural models over simple baselines may be smaller. The main comparison used one dataset and one split, and the network was tuned on a single validation split. Future work will run the protocol on real Nigerian hospital data, test other imputation methods and text encoders, assess whether Pidgin and local-language notes need dedicated models, add explainability for clinicians, and extend the framework to diabetes complications and sickle cell disease.

8. CONCLUSION

Rising chronic disease burden, large but imperfect hospital data and a tiered health system make risk prediction a promising but demanding application of machine learning in Nigeria. This paper adapted a multimodal design developed on Chinese hospital data to Nigerian conditions, combining imputation, expert-guided and learned features and a lightweight baseline within a governance structure that respects national law. A simulation on synthetic data showed that the pipeline runs end to end and that the multimodal model can outperform conventional baselines when notes carry useful information, but these results describe the simulation and not clinical performance. The next step is to run the protocol on real, ethically approved Nigerian hospital data.

REFERENCES

Adelani, D. I., Abbott, J., Neubig, G., D'souza, D., Kreutzer, J., Lignos, C., … Nakashole, N. (2021). MasakhaNER: Named entity recognition for African languages. *Transactions of the Association for Computational Linguistics, 9*, 1116–1131.

Adeloye, D., Basquill, C., Aderemi, A. V., Thompson, J. Y., & Obi, F. A. (2015). An estimate of the prevalence of hypertension in Nigeria: A systematic review and meta-analysis. *Journal of Hypertension, 33*(2), 230–242. https://doi.org/10.1097/HJH.0000000000000413

Akinyemi, R. O., Ovbiagele, B., Adeniji, O. A., Sarfo, F. S., Abd-Allah, F., Adoukonou, T., … Owolabi, M. O. (2021). Stroke in Africa: Profile, progress, prospects and priorities. *Nature Reviews Neurology, 17*(10), 634–656. https://doi.org/10.1038/s41582-021-00542-4

Bates, D. W., Saria, S., Ohno-Machado, L., Shah, A., & Escobar, G. (2014). Big data in health care: Using analytics to identify and manage high-risk and high-cost patients. *Health Affairs, 33*(7), 1123–1131. https://doi.org/10.1377/hlthaff.2014.0041

Chen, M., Hao, Y., Hwang, K., Wang, L., & Wang, L. (2017). Disease prediction by machine learning over big data from healthcare communities. *IEEE Access, 5*, 8869–8879. https://doi.org/10.1109/ACCESS.2017.2694446

Chen, M., Mao, S., & Liu, Y. (2014). Big data: A survey. *Mobile Networks and Applications, 19*(2), 171–209. https://doi.org/10.1007/s11036-013-0489-0

Collins, G. S., Reitsma, J. B., Altman, D. G., & Moons, K. G. M. (2015). Transparent reporting of a multivariable prediction model for individual prognosis or diagnosis (TRIPOD): The TRIPOD statement. *Annals of Internal Medicine, 162*(1), 55–63. https://doi.org/10.7326/M14-0697

Cover, T., & Hart, P. (1967). Nearest neighbor pattern classification. *IEEE Transactions on Information Theory, 13*(1), 21–27. https://doi.org/10.1109/TIT.1967.1053964

Federal Ministry of Health. (2019). *National multisectoral action plan for the prevention and control of non-communicable diseases (2019–2025)*. Federal Republic of Nigeria.

Federal Republic of Nigeria. (2022). *National Health Insurance Authority Act, 2022*. Official Gazette.

Federal Republic of Nigeria. (2023). *Nigeria Data Protection Act, 2023*. Official Gazette.

Groves, P., Kayyali, B., Knott, D., & Van Kuiken, S. (2013). *The "big data" revolution in healthcare: Accelerating value and innovation*. McKinsey & Company.

Jensen, P. B., Jensen, L. J., & Brunak, S. (2012). Mining electronic health records: Towards better research applications and clinical care. *Nature Reviews Genetics, 13*(6), 395–405. https://doi.org/10.1038/nrg3208

Kim, Y. (2014). Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing* (pp. 1746–1751). Association for Computational Linguistics. https://doi.org/10.3115/v1/D14-1181

Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer, 42*(8), 30–37. https://doi.org/10.1109/MC.2009.263

National Health Research Ethics Committee of Nigeria. (2007). *National code of health research ethics*. Federal Ministry of Health.

Odekunle, F. F., Odekunle, R. O., & Shankar, S. (2017). Why sub-Saharan Africa lags in electronic health record adoption and possible strategies to increase its adoption in this region. *International Journal of Health Sciences, 11*(4), 59–64.

Owolabi, M. O., Akpa, O. M., Agunloye, A. M., Akpalu, A., Sarfo, F., Ovbiagele, B., … SIREN Investigators. (2018). Dominant modifiable risk factors for stroke in Ghana and Nigeria (SIREN): A case-control study. *The Lancet Global Health, 6*(4), e436–e446. https://doi.org/10.1016/S2214-109X(18)30002-0

Riley, R. D., Ensor, J., Snell, K. I. E., Harrell, F. E., Martin, G. P., Reitsma, J. B., … Collins, G. S. (2020). Calculating the sample size required for developing a clinical prediction model. *BMJ, 368*, m441. https://doi.org/10.1136/bmj.m441

Uloko, A. E., Musa, B. M., Ramalan, M. A., Gezawa, I. D., Puepet, F. H., Uloko, A. T., … Borodo, M. M. (2018). Prevalence and risk factors for diabetes mellitus in Nigeria: A systematic review and meta-analysis. *Diabetes Therapy, 9*(3), 1307–1316. https://doi.org/10.1007/s13300-018-0441-1

Wang, J., Qiu, M., & Guo, B. (2017). Enabling real-time information service on telehealth system over cloud-based big data platform. *Journal of Systems Architecture, 72*, 69–79.

World Health Organization. (2018). *Noncommunicable diseases country profiles 2018*. World Health Organization.