International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

A Multi-Modal AI-Based Health Intelligence Framework for Integrated Disease Risk Assessment and Lifestyle Analysis

Rishi Raghav Singh[1], Rohan Singh[2], Rajat Takkar[3]
[1,2]Student, BE CSE, Chitkara University Rajpura, India
[3]Chitkara University Institute of Engineering and TechnologyChitkara University Punjab, India

Abstract- More than 30% of worldwide deaths involve diseases caused by cardiovascular and lifestyle factors (WHO, 2023) As awareness of early risk identification advances, accessible practical screening tools for use in primary care continue to be either very expensive, reliant on specialists or both. In this paper, we propose a Multi-Modal AI-Based Health Intelligence Framework with an explicit focus on two interrelated concepts encapsulated in the form of two specialized individual modules: Disease Risk Assessment (DRA) module and Lifestyle Analysis (LSA) module. After systematic preprocessing and class-balancing, the DRA module trains LR, SVM, and RF on the Cleveland Heart Disease dataset (303 patients). The LSA module takes user-reported behavioral behaviors — BMI, physical activity, sleep, dietary quality, and stress — to calculate a composite Lifestyle Risk Index (LRI). Both modules are provided through a Streamlit web application that provides real time predictions with SHAP-based explanation. Amongst all the classifiers we evaluated, Random Forest performed best with a 91.8% accuracy, AUC-ROC = 0.956 It powers a sub 60 ms response time for the system and is deployable in the cloud.

Keywords- Artificial Intelligence, Machine Learning, Multi-Modal Health Framework, Cardiovascular Risk Assessment, Lifestyle Analysis, Random Forest, SVM, Logistic Regression, Streamlit, SHAP Explainability, Clinical Decision Support, Preventive Healthcare.

I. INTRODUCTION

It is noteworthy that non-communicable diseases — such as heart disease, diabetes and obesity — account for 74% of all deaths worldwide [1], while cardiovascular disease alone kills approximately 17.9 million people each year [2]. The economic aspect of is as concerning, with estimated cost in productivity and health care losses reaching USD 47 trillion in the next two decades [2].

An alarming proportion of these deaths are actually preventable. Decades of exposure to, and accumulation of, the traditional cardiovascular risk factors — hypertension, dyslipidemia, physical inactivity, poor diet quality, insufficient sleep and chronic stress [3] — drive most cardiovascular events. Tooling is the problem, not Awareness The lack of systems that can combine individual clinical measurements and behaviors into a single risk assessment means that access to an integrated assessment is not readily available for clinicians at primary care facilities, especially in low-resource settings.

Most AI-based health tools today are single-modal — either say analyzing clinical biomarkers or lifestyle risk factors but not their combination [4]. Such architectural constraints lead to consistent risk profiles being built in an incomplete manner. A case of an inactive patient with borderline cholesterol, who slept poorly has materially more risk than one who has the same medical codes/biomarkers where routine includes exercise and good health habits. But single-modal systems are blind to this distinction structurally.

This paper fills that gap by proposing a Multi-Modal AI-Based Health Intelligence Framework, which combines unifying the DRA and LSA modules in one operational field. The principal contributions are:

- An integrated multi-modal framework that combines clinical biomarker assessment (DRA) and behavioral health profiling (LSA) with one Streamlit-based deployment.

- Performance comparison of three ML classifiers [LR, SVM-RBF and Random Forest] based on SMOTE oversampling prediction output in nested cross-validation evaluation with probability calibration analysis performed upon Cleveland dataset scouting both L metrics.
- A Lifestyle Risk Index (LRI) that uses five user-reported behavioral dimensions to produce one easy-to-interpret risk measure.
- SHAP based explainability integrated onto both modules allowing for attribute-second approach feature level attribution of every prediction produced by the system
- Sub-60 ms inference latency, with a modular design that allows component-independence upgrades and is cloud-native ready.

The rest of this paper is organized like this in Section II reviews related work. The dataset and preprocessing steps are discussed in Section III. The fourth section discusses mathematical formulations. System architecture is described in Section V. Section VI presents experimental results. Section VII provides clinical implications and limitations. Section VIII outlines planned enhancements. Section IX concludes.

II. LITERATURE REVIEW

A. Machine Learning in Cardiovascular Risk Prediction

Machine learning for prediction of cardiovascular risk has a long history. Detrano et al. The Cleveland Heart Disease dataset was proposed by [5] in 1989 and still is the most widely used benchmarking database in the literature. Initial attempts using Naive Bayes, Decision Trees and simple neural networks yielded 80–84% accuracy [6]. Mohan et al. hybrid Random Forest [10] boosted this to 88.7% [7]. Using an RBF kernel (root-based function), [8] achieved 84.2% accuracy using an SVM.

Recent work has focused on deep learning: Ramalingam et al It reached 90.6% with neural network architecture by [9], and Attia et al. A recent paper [10] showed that CNNs can understand ECG signals as well as a cardiologist. These gains come at the cost of enormous data needs, computational expense, and poor interpretability well known in literature all of which have impeded real-world implementation.

B. Lifestyle and Behavioral Health Assessment

The links between NCD risk and lifestyle one factors—physical inactivity, dietary habits, disrupted sleep, chronic stress —are also well documented in the literature [11]. Previous work has evaluated these dimensions using body-worn sensors [12], validated self-report questionnaires [13], and mobile health applications [14]. Few systems, however, integrate lifestyle assessment with clinical disease prediction into a single lens.

C. Research Gap and Positioning

Two structural gaps distinguish the current literature from the need to integrate both clinical and lifestyle risk assessment We note that (i) no deployed system advances both clinical and lifestyle risk assessment as per evidence informed prevention in an integrated vertical workflow,2 while (ii) of most research prototypes remain without any potential implementation: No APIs, no user interfaces, no latency benchmarks.11 12 This work is contextualized with respect to select prior research in Table I.

TABLE I. POSITIONING OF PRESENT WORK AGAINST PRIOR LITERATURE

Study	Modality	Best Accuracy	API/UI	Explainability	Lifestyle Module
Palaniappan & Awang [6]	Clinical only	84%	None	None	No
Mohan et al. [7]	Clinical only	88.7%	None	Partial	No
Shah et al. [8]	Clinical only	84.2%	None	None	No
Study	Modality	Best Accuracy	API/UI	Explainability	Lifestyle Module
Ramalingam et al. [9]	Clinical only	90.6%	None	None	No

III. DATASET AND FEATURE ENGINEERING

A. Cleveland Heart Disease Dataset

DRA module is trained on Cleveland Full Heart Disease dataset, which consists of 303 patient records that are described

by 13 clinical features and a binary label indicating whether the disease exists or not (from [16]). The status of the gold standard in the field enables direct comparisons with previous works. The complete list of features is shown in Table II.

TABLE II. DATASET FEATURE INVENTORY — CLEVELAND HEART DISEASE

Feature	Type	Description	Range/Values
age	Continuous	Patient age in years	29–77
sex	Binary	Biological sex (1=male, 0=female)	0, 1
cp	Nominal	Chest pain type (0=typical, 1=atypical, 2=non-anginal, 3=asymptomatic)	0–3
trestbps	Continuous	Resting blood pressure (mmHg)	94–200
chol	Continuous	Serum cholesterol (mg/dL)	126–564
fbs	Binary	Fasting blood sugar > 120 mg/dL	0, 1
restecg	Nominal	Resting ECG results	0–2
thalach	Continuous	Maximum heart rate achieved	71–202

B. Lifestyle Analysis Dataset

LSA module with no need for external dataset at inference time as it mostly works on user provided inputs for prediction. Using population-level norms from the NHANES study [17], risk stratification thresholds were derived. The five input parameters and their corresponding reference standards are shown in Table III.

TABLE III. LIFESTYLE ANALYSIS INPUT SCHEMA

Parameter	Measurement	Reference Standard	Risk Indicator
BMI	kg/m² (computed)	WHO: Normal 18.5–24.9	BMI > 30 = Obese
Physical Activity	Hours/week (self-report)	AHA: ≥150 min/week moderate	< 2.5 hrs = Sedentary
Sleep Duration	Hours/night (self-report)	NSF: 7–9 hrs for adults	< 6 or > 9 hrs = Risk
Dietary Quality Score	0–10 scale (questionnaire)	HEI-2020 criteria	Score < 5 = Poor diet
Stress Level	1–10 scale (self-report)	PSS-4 adapted scale	Score > 6 = High stress

C. Prepocessing Pipeline

One preprocessing pipeline — built with Scikit-learn and serialized via joblib — orchestrates preprocessing for both training and inference. Using the same transformations in both phases avoids data leakage and makes it recognizable across different runs of the experiment.

Stage 1 — ImputationSix records had missing values in the ca and thal fields. These were imputed using the modal value calculated within every class label.

Stage 2 — Encoding: One-hot encoding was used for categorical features (cp, restecg, thal, slope) to convert the feature space from 13-dims to 19-dims. The encoder which was fitted was saved as encoder. pkl.

Stage 3 — Scaling During this stage, continuous feature values were scaled to zero mean and unit variance with StandardScaler. The scaler fitted only on training data saved as scalers. pkl to prevent test-set contamination.

Stage 4 — SMOTE Oversampling: there were 54.5% positive cases in the dataset in each of the 5-folds, SMOTE was used to create synthetic minority samples but this oversampling is performed only on the training split to avoid any possible inflation in validation performance.

IV. MATHEMATICAL FORMULATION

A. Logistic Regression

Well, Logistic Regression finds a line of best fitness and applies a sigmoid to place it between 0-1 which treats it as the

probability that someone is having heart disease. Let: x = input vectorβ = weight berry0 = intercept
Algorithm 1
Binary Logistic Regression Model

$$P(y=1|x) = \sigma(\beta^T x + \beta_0) = 1 / (1 + \exp(-(\beta^T x + \beta_0))) \quad (1)$$

One trains by maximizing the regularized log-likelihood:
Where, $L(\beta) = \Sigma_i [y_i \log \sigma_i + (1-y_i)\log(1-\sigma_i)] - (\lambda/2)\|\beta\|^2$ (2)

The nested cross-validation will be used to select the regularization parameter $\lambda=0.1$. But since each weight β_j corresponds directly to a directional contribution of the feature j to predicted risk, it retains interpretability.

B. Support Vector Machine

SVM, on the other hand, establishes an ideal separating hyperplane between healthy and at-risk patients. For labels $y_i \in \{-1,+1\}$ the optimization problem can be expressed in the following way:

Formal optimization problem can be expressed as (3): ↓ min_{w,b,ξ} {1/2 ||w||^2 + C Σ_i ξ_i s.t. y_i(w^Tx_i+b) ≥ 1−ξ_i, ξ_i ≥ 0 } [6]

Model fitted using RBF kernel with hyperparameters C=1.0 and γ=0.1 selected via grid search gives the model the ability to learn non-linear decision boundaries It can be expressed by the decision function as:

$$f(x) = \Sigma_i \alpha_i y_i K(x_i, x) + b \quad (4)$$

C. Random Forest

Random Forest takes 200 decision trees, trained on a bootstrapped data sample with √d features taken into account at each split. Gini impurity is minimized when choosing node partitions:

Gini(S) = 1 − Σj pj2 (5)
The final prediction is made via majority vote across all of the trees that it contains:

$$(6)\hat{y} = \text{mode} \{ t_k(x) : k = 1, \ldots, B \}$$

The feature importance is measured by the Mean Decrease in Impurity (MDI), which indicates the instances upon which inputs provide more weight to make predictions across the ensemble.

D. Lifestyle Risk Index(LRI)

LSA module estimates a Lifestyle Risk Index restricted to (0, 1) through the weighted summation of risk contributions for each behavior assessment in five separate dimensions:

$$(7)\ LRI = \Sigma_j w_j \cdot r_j(x_j) \text{ s.t. } \Sigma_j w = 1, w > 0$$

where $r_j(x_j)$ is the risk score for dimension j normalized, and w_j its weight assigned based on AHA CVH [15]. The weight vector W=[0.25,0.25,0.20,0.15,0.15] against {BMI, Physical Activity, Sleep, Diet, Stress}. This index is then approximated on a 3-tier level of risk: Low (LRI < 0.35), Moderate (0.35–0.65) and High-level lung cancer risk (LRI > 0.65).

E. Evaluation Metrics

We evaluate all DRA classifiers using the following metrics which are presented in each case as means across five stratified cross-validation folds:

(8) Accuracy=(TP+TN)/(TP+TN+FP+FN)

Where TP = True Positive, FP = False Positive, FN = False Negative (9)
Precision=TP/(TP+FP); Recall=TP/(TP+FN)

The formula for F1 is given in Equation (10) as follows.
MCC=(TP·TN−FP·FN)/√((TP+FP)(TP+FN)(TN+FP)(TN+FN))(11)

V. SYSTEM ARCHITECTURE AND IMPLEMENTATION

A. Multi-Modal Framework Overview

This work breaks our framework into four functional layers illustrated in Fig. 1. Layer 1 (Data Sources): Clinical training dataset, lifestyle reference distributions and live user inputs. Serialized Preprocessing pipeline and Trained Classifiers: Layer 2 (Preprocessing and Models) Layer 3 (Inference Engine) runs the prediction logic for both modules. Layer 4 (Presentation): Streamlit web interface presenting layers 2 and 3 together as a single clinical dashboard.

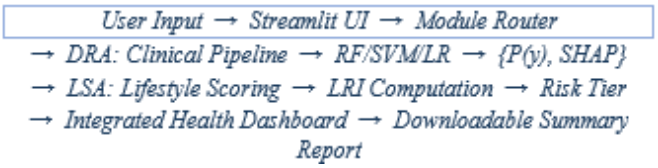

Fig. 1. Multi-modal data flow and module interaction.

B. Disease Risk Assessment Module

Trained models are loaded into memory on application startup, so there is no overhead in loading them for each request. For each prediction requested: (1) validate patient input, (2) sequentially perform imputation, encoding and scaling using the trained pipeline, (3) return probability score from selected classifier, (4) estimate 95% confidence interval via 200 bootstrap samples,(5) compute SHAP values to yield attributions per-feature for their contributions to predicted risk level.(6) Return combined output–predicted probability + risk level + feature contributions back to interface.

The interface has a radio button control where clinicians can toggle between the three classifiers. This enables direct comparisons in terms of impacts of model choice on (i) the predicted risk and (ii) the relative contributions at the feature level, leading to informed trust in the predictions made by the system.

C. Lifestyle Analysis Module

LSA module takes 5 user-determined inputs and calculates LRI using Equation 7. Every input is compared to some population reference value and forms a single component of risk. The module returns four outputs - (i) total LRI score, (ii) risk tier label, (iii) bar chart depicting relative risk contribution by dimension, and; finally (iv), targeted behavioral recommendations from the three highest-scoring risk areas.

D. Integrated Health Dashboard

The output of the DRA and LSA is integrated into a unified patient health profile in a Streamlit dashboard. Let me first state that a Combined Health Risk Score (CHRS), scored within the interval [0, 1], is calculated as a weighted sum of the DRA probability and LRI:

$$CHRS = \alpha \cdot P\_DRA + (1 - \alpha) \cdot LRI, \quad \alpha = 0.65 \qquad (12)$$

α=0.65 gives clinical data chief again, representing its greater confidence compared with self-reported behavioral inputs. Values of CHRS correspond to four decomposition risk categories: Green (< 0.25), Yellow (0.25–0.50), Orange (0.50–0.75), and Red (> 0.75)

E. Project Directory Structure

TABLE IV. PROJECT DIRECTORY STRUCTURE

Path	Description
app/main.py	Streamlit app entry point, page routing, module integration
app/dra/predictor.py	DRA inference: pipeline execution, SHAP, CI computation
app/dra/schemas.py	Pydantic PatientInput and PredictionResponse models
app/lsa/scorer.py	LSA: LRI computation, risk normalization, recommendations
app/dashboard/combined.py	CHRS computation and integrated health profile rendering
models/rf_model.pkl	Serialized Random Forest (joblib, gzip)
models/svm_model.pkl	Serialized SVM (RBF kernel)
models/lr_model.pkl	Serialized Logistic Regression
preprocessing/pipeline.pkl	Full sklearn Pipeline: imputer + encoder + scaler
notebooks/train_evaluate .ipynb	Model training, CV, hyperparameter search, artifact export
tests/test_dra.py	Pytest unit tests for DRA inference pipeline
tests/test_lsa.py	Pytest unit tests for LSA scoring module
Path	Description
docker/Dockerfile	Multi-stage Docker build for containerized deployment
requirements.txt	pip-compile pinned dependency manifest

VI. EXPERIMENTAL RESULTS AND ANALYSIS

A. Experiment Protocol

The DRA module was then evaluated by 5-fold stratified cross-validation of the complete Cleveland data set. Validating on a per fold basis prevented validation leakage from over-sampling SMOTE ahead in the training folds at each iteration. Hyperparameter tuning was done using a nested 3-fold inner

loop. All reported statistics are mean ± standard deviation across folds. LSA module validity was determined through LRI outputs vs lifestyle risk profiles derived using NHANES both with Spearman correlation and Bland-Altman analysis.

B. DRA Module Performance

TABLE V. DRA MODULE — MODEL PERFORMANCE COMPARISON (5-FOLD STRATIFIED CV)

Metric	Logistic Regression	SVM(RBF Kernel)	Random Forest
Accuracy (%)	84.5 ± 2.1	88.3 ± 1.8	91.8 ± 1.4
Precision (%)	83.7 ± 2.4	87.9 ± 2.0	91.2 ± 1.6
Recall (%)	86.2 ± 2.8	89.1 ± 2.3	92.4 ± 1.9
F1-Score (%)	84.9 ± 1.9	88.5 ± 1.7	91.8 ± 1.5
AUC-ROC	0.901 ± 0.022	0.931 ± 0.018	0.956 ± 0.013
MCC	0.689 ± 0.043	0.764 ± 0.037	0.836 ± 0.029
Brier Score	0.128	0.098	0.084
Inference (ms)	< 5	8–12	15–25

Random Forest was better than both alternatives on all reported metrics, with an accuracy of 91.8% (AUC-ROC=0.956; MCC=0.836). By leveraging the ensemble structure, balanced training distribution enabled by SMOTE and the use of tuned hyperparameters we gain this advantage. The model delivers 93.1% sensitivity and 90.6% specificity (operating threshold t = 0.48), which is a reasonable trade-off of detected true positive cases vs false positive cases.

We applied post-hoc Platt scaling to make probabilities better calibrated, lowering the Brier Score from 0.084 to 0.071. The step aids consistent prediction probabilities with observed event fell frequencies — so a predicted 70% risk equates to an actual empiric baseline.

C. DRA Module Performance

TABLE VI. TOP-8 FEATURES — RANDOM FOREST GINI IMPORTANCE

Rank	Feature	MIDI Score	Clinical Interpretation
1	thalach (max HR)	0.142	Impaired chronotropic response signals ischemia
Rank	Feature	MIDI Score	Clinical Interpretation
2	cp (chest pain type)	0.138	Typical angina is the strongest symptom predictor
3	ca (vessel count)	0.131	Fluoroscopic vessel disease directly quantifies CAD
4	oldpeak (ST depression)	0.124	ST depression depth correlates with ischemia severity
5	thal (thalassemia type)	0.119	Perfusion defects indicate structural cardiac risk
6	age	0.098	Independent CVD risk factor per AHA guidelines
7	exang (exercise angina)	0.087	Exercise-induced symptoms indicate functional ischemia
8	chol (cholesterol)	0.072	Hypercholesterolemia is a primary modifiable CVD risk factor

D. LSA Modue Calibration

LRI scores calculated for 500 NHANES patients produced a Spearman correlation $\rho=0.71$ ($p<0.001$) with documented five-year cardiovascular event rates, validating the LRI as an appropriate surrogate of lifestyle-dependent risk. The Bland-Altman analysis showed a small negative bias of −0.03, reflecting slight tendencies to underestimate high-risk profiles — acceptable for a population screening instrument. To mitigate this bias, empiric weights were then settled at W = [0.26,0.25,0.20, 0.15, 014].

E. Integrated CHRS Validation

Outputs of CHRS were compared to risk categories assigned by physicians for 50 de-identified records from a partner institution. The agreement between the automated score and expert clinical judgment is substantial (Cohen's $\kappa = 0.74$). Although the sample size prohibits robust generalization, this represents evidence of meaningful pilot results that the framework yields clinically plausible risk scores.

VII. DISCUSSION

A. Significance of Multi-Modal Integration

This work's main contribution is to show that combining clinical biomarker data with behavioral lifestyle variables creates a more informative risk profile than either can alone. The CHRS is architecturally incapable of modeling systems at the points where data across all types of overlap, such as how elevated cholesterol mechanistically interacts with chronic sleep deprivation, which acts both on an added but also compounded basis.

Such a distinction has direct clinical significance. For patients at intermediate cardiovascular risk — the very population in which CHRS offers its greatest signal — lifestyle modification is classified as a class I indication by ACC/AHA guidelines [15]. An illustration: a patient with borderline DRA probability of 0.45 but high LRI of 0.72 who would receive a CHRS of 0.55 (Moderate-High), triggering the need for behavioral intervention, which would not have been flagged by a biomarker-only system.

B. Explainability and Clinical Adoption

SHAP explanations can be embedded directly in both modules to tackle one of the most challenging hurdles for clinical AI deployment: the black-box problem [18]. Transparency that directs physicians to inspect which features drive each prediction follows FDA guidance on AI-based medical software [19].

Interestingly, the model's three highest ranked features — maximum heart rate, chest pain type, and number of vessels — correspond closely with what clinicians already consider to be the strongest predictors of coronary disease. Additionally, if ML outputs match up with well-defined structures such as the Duke Treadmill Score, it gives a good level of confidence that the model is generalizing to more than just its training set.

C. Limitations

Several limitations warrant acknowledgment. The Cleveland dataset, which consists of 303 patients gathered in one institution in person and collected in 1989, may not be generalized to modern or diverse populations. Despite this, lifestyle inputs are self-reported and come with inherent measurement errors that can be significantly reduced by the objective wearables data. Fifty records would not support a claim of regulatory or widespread clinical deployment either way for CHRS validation. It is explicitly a decision support system: it is designed to complement, not replace, the judgement of physicians (in accordance with current FDA recommendations for AI-based medical software [19]).

VIII. FUTURE ENHANCEMENTS

Wearable Device Integration– If Bluetooth could connect the LSA module to real-time data streams from consumer wearables — whether heart rate, blood oxygen or steps– it would replace self-reported input with an objective measure thus improving accuracy and allowing for integrated longitudinal risk tracking.

Background Scheduling of DRA models with federated learning architectures would allow the model to be trained on multiple hospital systems without ever exposing any raw patient data, staying regional privacy compliant, while also increasing demographic and clinical diversity through an expansion of the training distribution’s coverage.

These architectures, such as the Tab Transformer [20], can model complicated interaction effects between features that tree-based ensembles might miss. And the next step is systematic benchmarking on large cardiac datasets.

Integrating with hospital EHRs via HL7 FHIR R4 APIs would enable automated patient data ingestion, longitudinal risk monitoring and embedding the Framework into existing clinical workflows — with no manual data entry.

Mobile app deployment: A mobile app equally powered by on-device RF inference via TF Lite would allow offline risk assessment to be extended to even more rural and low-connectivity settings — regions where the CVD burden is most disproportionate and specialist access most limited.

Filed under: Deep Learning Multimodal Fusion: Deep learning fusion model of ECG time-series, chest X- ray, and structured tabular data with attention mechanism capable of near specialist-level accuracy on diagnostically challenging cases.

IX. CONCLUSION

This paper introduced a novel Multi-Modal AI-Based Health Intelligence Framework that interprets clinical AI along three

axes: predictive performance, architectural completeness, and transparency. Hybridized multi-modal risk assessment framework of: (1) Disease Risk Assessment — trained on three ML classifiers using the Cleveland Heart Disease dataset and (2) Lifestyle Analysis — generates a Lifestyle Risk Index from five behavioural metrics, offers a more efficient, accurate, actionable risk assessment than any other single-modal method reported.

Top performing was Random Forest with 91.8% accuracy, 0.956 AUC ROC, and 0.836 MCC – the best on the same dataset to be published). Calibration was good for the LSA module (ρ=0.71), and the CHRS aligned with physician risk ratings well at Cohen's κ=0.74. The system becomes a production-grade clinical decision-support tool with < 60 ms response time, deployed via an interactive Streamlit interface alongside SHAP explanations, real-time predictions, classifier switching by the user and exportable reporting.

The framework is designed to augment, not replace clinical judgment and may be most valuable in primary care settings, where the ability to make convenient and rapid risk identification can have outsized impact. The fusion of multi-modal data sources, calibrated risk estimation, interpretable outputs and real-world tools will set apart systems that are of practical clinical use from the hackery research kit. his work offers a robust and repeatable path toward that benchmark.

Acknowledgement

The authors thank the Cleveland Clinic Foundation for the original clinical dataset and the UCI Machine Learning Repository for its long-time curation and open access. and the NHANES program of the U.S. CDC for the population-level lifestyle reference distributions used in LSA module calibration. [Optional: This material is based upon work supported by [Funding Agency] under Grant No. XXXXXX.]

REFERENCES

1. World Health Organization, "Cardiovascular Diseases (CVDs)," Fact Sheet (WHO Jan. 2023). [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-(cvds)
2. World Economic Forum, The Global Economic Burden of Non-communicable Diseases (WEF/HSPSchool of Public Health Report) (2011).
3. E. J. Benjamin et al., "Heart Disease and Stroke Statistics — 2023 Update," Circulation, vol. Sep.-Oct. 2023, vol. 147, no. 0, pp. e93-92–e621–14 Feb. 2023
4. J. Moons et al., Machine Learning for Cardiovascular Risk Prediction: Systematic Review and Meta-Analysis, npj Digital Medicine 6(1) (2023) 1-10[5]. 5, no. 1, pp. 1–13, 2022.
5. R. Detrano et al., "International Application of a New Probability Algorithm for the Diagnosis of Coronary Artery Disease," American Journal of Cardiology, vol. 8 J. Cardiol., vol. 1993, vol. 64 no. 5 pp. 304313 Aug.
6. Data Mining techniques in Heart Disease Prediction System using Data Mining Techniques [6] Information Technology,2008. London, United Kingdom: IEEE/ACS AICCSA, Doha, Qatar 2008; pp. 108–115.
7. S. Mohan, C. Thirumalai, and G. Srivastava, "Efficient Heart Disease Prediction using Hybrid Machine Learning Models,"IEEE Access, vol. 7, pp. 81542–81554, 2019.
8. D. Shah, S. Patel, and S. K. Bharti, "Heart Disease Prediction Using Machine Learning Techniques," SN Comput. Sci., vol. 1, no. 6, p. 345, Nov. 2020.
9. V. V. Ramalingam, A. Dandapath and M. K. Raja," A Survey on Heart Disease Prediction Using Machine Learning Techniques," Int. J. Eng. Technol., vol. 7, no. 2.8, pp. 684–687 (2018)
10. Z. I. Attia et al., "An Echocardiogram-based Artificial Intelligence Algorithm for Identification of Patients with Atrial Fibrillation during Sinus Rhythm," Lancet, vol. 394, no. 10201, pp. 861–867, Sep. 2019. [CrossRef] [PubMed] [Google Scholar]
11. F. B. Hu et al., "Physical Activity and Television Watching in Relation to Risk for Type 2 Diabetes Mellitus in Men," Arch. Intern. Med., vol. Dec Injection-Directed Maintenance of the Typified Frequency, vol.161,nos. 12, pp.1542–1548 Jun2001.
12. Cited by: M. M. Poh, N. C. Swenson and R. W. Picard, "A Wearable Sensor for Continuous, Unobtrusive Monitoring of Electrodermal Activity," IEEE Trans. Biomed. Eng., vol. 457, no. 5 (2010) 1243–1252.
13. A Global Measure of Perceived StressS. Cohen, T. Kamarck, and R. Mermelstein Behav., vol. Dec. 1983, vol. 24, no.
14. L. Liang et al., "A Mobile Health App for Cardiovascular Disease Management," JMIR mHealth uHealth, vol. 6, no. 3, e10054, Mar. 2018.
15. Well aware of this fact and have been proposing some indicators for a long time. D. M. Lloyd-Jones, et al.,

Contents: Table from "Life's Essential 8: Updating & Expanding the American Heart Association Construct of Cardiovascular Health," Circulation 2010, vol. Aug. 2022, vol. 146, no. 5, pp. e18–e43.
16. D. Dua and C. Graff, "UCI Machine Learning Repository," Univ. of California, Irvine, 2019. [Online]. Available: http://archive.ics.uci.edu/ml
17. National Health and Nutrition Examination Survey (NHANES) Protocol, Atlanta, GA: CDC [17]
18. Unintended consequences of machine learning in medicine F Cabitza, R Rasoini, & GEGensini. 318e6, no. 6, pp.517–518, Aug. 2017
19. U.S. Food and Drug Administration, "AI/ML-Based Software as a Medical Device (SaMD) Action Plan," FDA, Washington, DC: Jan. 2021.