

An Explainable Quantum Machine Learning Framework for Iris-Based Health Screening and Automated Health Report Generation: A Preliminary Single-Subject Study

Mrs. B. Siva Kumari, M.Tech (Ph.D.)¹, M. Naga Swathi², P. Keerthika³, L. Latha⁴

¹Assistant Professor, Bapatla Women's Engineering College, Bapatla, AP, India

^{2,3,4}B.Tech Student, Bapatla Women's Engineering College, Bapatla, AP, India

Abstract - Detecting health risks at an early stage remains a cornerstone of effective preventive care. Computational iris-texture analysis provides a non-invasive, exploratory lens into visual wellbeing signals, yet conventional machine learning pipelines frequently struggle with the high-dimensional nature of ocular imagery and rarely offer the transparency clinicians need before trusting a diagnosis. This study proposes an Explainable Quantum Machine Learning (EQML) framework built for iris-driven health screening and automated wellness-report summarization (not medical diagnosis). At its core sits a hybrid architecture that employs a Variational Quantum Classifier (VQC) to explore a quantum feature-space encoding of iris texture. Because quantum circuits are inherently difficult to interpret, a layered Explainable AI (XAI) stack Grad-CAM, Quantum SHAP, and LIME is fused into the pipeline to attribute each prediction to specific iris features clinicians can recognize. A companion Automated Health Report Generation module then converts raw quantum measurement outputs into readable clinical narratives. A preliminary, qualitative comparison against classical models conducted on twenty test captures from a single subject under varied lighting, using a classical circuit simulator (PennyLane/Qiskit Aer) rather than quantum hardware indicates that the EQML framework produces sharper separation between risk indicators and more stable, traceable explanations than SVM, Random Forest, and MLP baselines. These findings are exploratory and not clinically validated; the framework does not diagnose disease, and the observed trends should be treated as indicative rather than conclusive pending a labeled, multi-subject, quantitative validation study.

Keywords - Quantum Machine Learning, Iris Analysis, Explainable AI, Remote Photoplethysmography, Variational Quantum Circuits, Webcam-Based Diagnostics, Non-Invasive Health Screening.

I. INTRODUCTION

Preventive medicine depends on catching risk indicators early and monitoring them continuously, well before they escalate into serious conditions. Traditional screening routes, however, demand specialized equipment, trained staff, and in-person visits, which restricts how often screening can happen and shuts out populations with limited access to healthcare facilities. As a result, non-

invasive screening performed through an ordinary camera has gained traction as a low-cost, repeatable alternative that runs on hardware most people already own, such as a laptop webcam or a smartphone.

Two non-invasive modalities stand out in this space: computational iris-texture analysis and remote photoplethysmography (rPPG). Iris texture and color descriptors can be extracted from a single facial photo, whereas rPPG requires a brief video clip, since it relies on tracking subtle color changes over time; both are extracted with no physical contact required. Earlier literature links iris texture and coloration to broader physiological states, and periorbital skin appearance (e.g., dark circles) to fatigue-related and dermatological factors [15], while rPPG recovers cardiovascular signals pulse rate and blood-pressure proxies purely from subtle color shifts in facial skin over time. Fusing these two modalities into one pipeline makes it feasible to build a single-session tool that screens for several health indicators at once.

In a separate line of work, quantum machine learning (QML) has surfaced as a promising tool for modelling feature spaces that are high-dimensional and non-linear. A Variational Quantum Classifier (VQC) encodes classical feature vectors into quantum states and then learns decision boundaries through trainable parameterized circuits operating in an exponentially larger Hilbert space a theoretically motivated property that has not yet been shown to translate into a demonstrated practical advantage for texture-heavy biometric signals like iris patterns, and this paper does not claim to establish one. Quantum circuits are, however, essentially opaque by nature, and any risk-prediction system intended for clinical use must be able to explain and trace its outputs back to concrete, human-understandable features.

To address this gap, we present HealthVision AI, an Explainable Quantum Machine Learning (EQML) framework that brings together webcam-based iris and rPPG feature extraction, a VQC classifier, and a layered explainability stack (Grad-CAM, Quantum SHAP, LIME) inside one automated health-report pipeline. The main contributions of this work are summarized below:

- Integration of iris texture, periorbital, and rPPG-derived indicators into a single-session health screening pipeline.
- Adoption of a Variational Quantum Classifier for iris-texture-based health risk prediction.
- A layered explainability stack combining Grad-CAM, Quantum SHAP, and LIME for per-prediction attribution.
- An automated natural-language health report generation module traceable to indicator-level evidence.
- Comparative evaluation of the proposed VQC classifier against classical SVM, Random Forest, and MLP baselines.

II. RELATED WORK

A substantial body of research has examined camera-based physiological sensing for remote photoplethysmography [11], where subtle, time-varying color changes in facial skin are used to recover pulse rate and, with further modelling, blood-pressure proxies [1], [2], [18]. Such techniques deliver contactless vital-sign estimates but are usually validated on their own, apart from other visual health indicators [9].

Iris-based analysis, or iridology, has long been proposed as a non-invasive screening concept, although classical iridology has drawn criticism for its thin clinical evidence base [13]. Newer studies instead rely on structured image-processing pipelines Gabor filtering, HSV/LAB color analysis, and edge-based structural measurement to derive quantitative iris texture and color descriptors, moving away from subjective visual interpretation.

Quantum machine learning has demonstrated theoretical potential for high-dimensional feature spaces via quantum kernel methods and variational circuits [4], [6]. Work on supervised learning with quantum-enhanced feature maps [5], combined with parameter-shift circuit training [16], has made gradient-based optimization of VQCs possible, and separate studies on quantum speed-up [17] further motivate their application to texture-rich biometric features that classical kernels struggle to separate cleanly.

Explainability is generally considered a prerequisite for clinical trust in medical AI systems [3], [14]. SHAP-based attribution is now a standard technique for classical models [7], and recent efforts have started extending Shapley-value reasoning to quantum circuits by computing feature contributions through circuit parameter perturbations instead of classical gradient paths [8]. To our knowledge, few of the studies we reviewed carry this quantum explainability layer all the way through to a human-readable clinical output, though we do not claim to have surveyed the field exhaustively.

III. RESEARCH GAP

Although each of the sub-fields surveyed above has advanced considerably on its own, three gaps become apparent once these techniques are considered together for a deployable health-screening application. First, existing rPPG and iris-analysis systems are usually built and tested in isolation, each targeting a single physiological indicator instead of being folded into one unified, multi-indicator pipeline that runs from a single camera session.

Second, although VQCs offer theoretical advantages for high-dimensional, non-linear feature spaces, the studies we are aware of rarely apply them to real, camera-derived biometric features such as Gabor-filtered iris texture, and rarer still pair the resulting quantum classifier with a rigorous, per-prediction explainability mechanism appropriate for a health-facing application.

Third, the explainability techniques designed for quantum circuits tend to be showcased on synthetic or benchmark datasets and, in the literature we reviewed, are not commonly wired end-to-end into a human-readable output that a non-technical user could act on. This paper tackles all three gaps together by proposing an end-to-end, locally hosted EQML framework that unifies multi-indicator camera-based screening, quantum-enhanced classification, and traceable explainability inside a single automated health-report-generation system.

IV. PROPOSED SOLUTION

A. System Overview

HealthVision AI is a locally hosted web application accessible through a browser at localhost (127.0.0.1). The interface presents a live camera feed on the left panel and a Health Analysis Report on the right panel. The user initiates analysis by clicking Capture and Analyze, which freezes a still frame for iris/facial analysis and retains a short preceding buffer of frames from the webcam stream for rPPG, runs the full analysis pipeline on this captured data, and populates the report panel with twelve health indicators along with a set of quantum health recommendations. No data is stored, transmitted to a server, or loaded from any external dataset; the entire analysis is performed on the captured frame within the local session.




Fig. 1. End-to-end system pipeline: iris/face capture → preprocessing → quantum ML → XAI → health report.

B. Input Modality

The input to the system is captured from the device's built-in or connected camera in a single "Capture and Analyze" action: a still RGB frame is grabbed for iris localization and texture analysis, and, in the same action, a short buffered RGB video window (a few seconds at the live-feed frame rate) is retained for rPPG-based pulse and blood-pressure estimation, since rPPG requires tracking color change over time and cannot be derived from one still frame. The camera feed is displayed in real time with an overlay guide box prompting the user to centre their face, ensuring the face occupies a sufficient portion of the frame for reliable landmark detection and iris localization.

C. Output Structure

The system generates twelve output attributes in two groups, of which ten are physiological/wellness indicators and two (Iris Color, Health Score) are descriptive or composite outputs rather than indicators in their own right. The Vitals and Wellness group contains Dark Circles, Fatigue Level, Stress Level, Pulse Rate (BPM), BP Estimate (mmHg), and Emotion. The Disease Risk Indicators group contains Diabetes Risk, Hypertension Risk, Eye Disease Risk, and Cardiovascular Index. Each is reported as a Risk Index the sigmoid classifier output of Eq. (6) rescaled to a 0–100 scale together with a severity category (Low/Moderate/High/Monitor) obtained by thresholding that index; the Risk Index is a model-internal confidence score against a reference-derived proxy target (Section V.F), not a probability of having the named disease, and is labeled as such wherever it is displayed. An overall Health Score out of 100, an Iris Color identification, and a personalized Quantum Health Recommendations list

(general wellness suggestions, not medical advice — see Section IV.G) are also reported. These are screening-level, exploratory signals computed from visual proxies; the system performs no clinical diagnosis, and none of the reported values should be interpreted as a medical diagnosis or a substitute for evaluation by a qualified clinician (Section VIII.B). To make clear which computational path produces each output: the four Disease Risk Indicators and the Health Score are the only outputs of the VQC path (Section IV.D, Eq. (6)–(7)); Pulse Rate and BP Estimate come from the classical rPPG/PTT pipeline (Section IV.C, Eq. (1)–(2b)) with no quantum component; Emotion comes from the classical FER2013 CNN (Table I); and Dark Circles, Fatigue Level, Stress Level, and Iris Color are computed by deterministic, threshold-based rules over the extracted HSV/LBP/Gabor and pupil-dilation features (Section IV.B–C) rather than by any trained classifier.

V. PROPOSED ALGORITHM

The following algorithm illustrates the end-to-end steps of the proposed EQML framework, from webcam capture through quantum classification to report generation.

Algorithm: EQML-Based Health Risk Prediction and Report Generation

Input: Captured still frame F and buffered video window V from webcam

Output: Health Analysis Report R

1. Initialize local application session at 127.0.0.1.
2. Activate webcam stream; display live feed with guide overlay.
3. Capture still frame F and short buffered video window V on user command (Capture and Analyze).
4. Detect face and extract 468-point facial mesh (MediaPipe).
5. Segment iris, sclera, and periorbital ROI from landmarks.
6. Preprocess ROI: resize, normalize, convert RGB → HSV/LAB.
7. Extract HSV histograms, LBP texture, and Gabor wavelet responses.
8. From V, estimate rPPG signal over the forehead-cheek ROI; derive HR and BP.
9. Encode iris texture vector into quantum state via angle encoding.
10. Apply variational circuit layers (Rx, Ry, Rz, CNOT) to encoded state.
11. Measure $\langle Z \rangle$ per qubit; compute risk scores via classical dense head.
12. Generate Grad-CAM, Quantum SHAP, and LIME explanations per indicator.
13. Aggregate twelve indicators into a composite Health Score H.

14. Generate natural-language recommendations and populate report R.

The processed webcam frame is thus carried through detection, feature extraction, quantum classification, and explanation stages before a clinically phrased report is rendered to the user. The seven functional modules underlying this pipeline are detailed below.

A. Module 1 — Data Collection and Preprocessing

The first module converts a raw webcam frame into a normalized feature vector. A face detector locates the face bounding box, a 468-point facial mesh is fitted, and the iris, sclera, and periorbital ROIs are extracted from the landmark coordinates. The iris patch is resized to 64×64, normalized to [0, 1], and converted to HSV and LAB color spaces. HSV histograms, LBP texture, and Gabor wavelet responses are then concatenated into a feature vector $X \in \mathbb{R}^n$ for the downstream models.

Iris-specific extraction applies Canny edge detection to isolate the iris boundary, followed by a Hough Circle Transform for radius estimation. The mean HSV values of iris pixels determine the color label, the mean grayscale value yields the brightness index, and HSV thresholding on the sclera detects yellowing.



Fig. 2. Data collection and preprocessing pipeline: webcam capture through feature extraction.

B. Module 2 — Iris Color, Brightness, and Gabor Texture Analysis

Within the segmented iris mask, HSV values are histogrammed and the dominant hue cluster is matched against reference ranges for Brown, Blue, Green, Hazel, and Gray. Brightness is derived from the mean Value-channel intensity, reported on a 0–255 scale. A multi-orientation, multi-scale Gabor filter bank is applied to the grayscale iris

region; the mean and variance of each response map form an iris texture vector encoding collarette, crypt, and furrow micro-structure, subsequently dimensionality-reduced before quantum classification.

C. Module 3 — Remote Photoplethysmography for Pulse and BP

Heart rate is estimated from temporal green-channel variation in a forehead-cheek ROI, band-pass filtered at 0.7–4.0 Hz, with the dominant spectral peak located via FFT:

$$HR = \operatorname{argmax} |FFT(S(t))| \times 60 \text{ (beats per minute)} \quad (1)$$

Blood pressure is estimated from two signals captured over the buffered video window V: (i) Pulse Transit Time (PTT), estimated as the time delay between the systolic peak of the rPPG waveform extracted from the forehead ROI and the corresponding peak from the cheek ROI (via cross-correlation between the two ROI signals) a proxy for true arterial PTT rather than an ECG-to-periphery measurement and (ii) pupil-diameter variation, an autonomic/sympathetic tone proxy measured from consecutive iris-region frames in V; this is a pupillary signal, not a direct measurement of iris blood vessels, which are not reliably resolvable with a consumer webcam. These two features feed a linear regression model whose coefficients $(\alpha, \beta, \gamma, \delta)$ are presently set from published PTT–BP relationships rather than fitted on subject-specific labeled BP measurements; this estimate is therefore uncalibrated and not clinically validated, and is reported only as a screening-level approximation:

$$\text{Systolic} = \alpha \cdot \text{PTT}^{-1} + \beta \cdot \text{HR} + \delta \cdot \text{PD} + \gamma \quad (2a)$$

$$\text{Diastolic} = \alpha' \cdot \text{PTT}^{-1} + \beta' \cdot \text{HR} + \delta' \cdot \text{PD} + \gamma' \quad (2b)$$

where the primed coefficients $(\alpha', \beta', \delta', \gamma')$ are a separate set fitted to the same published PTT–BP relationships for diastolic pressure; Eq. (2a)–(2b) together produce the Systolic/Diastolic pair reported in Section VI.C (e.g., 127/81 mmHg).

Stress level is computed from a weighted combination of pupil dilation ratio (PD in Eq. (2a) above, normalized per subject relative to the first frame of V), facial action units (AU4, AU7, AU23), and HRV features (SDNN, RMSSD) derived from the rPPG signal. Emotion recognition uses a 68-point subset of the MediaPipe facial mesh (Section IV.A) fed into a seven-class softmax classifier trained on FER2013 (Angry, Disgust, Fear, Happy, Sad, Surprise, Neutral); a rule-based post-processing step maps a Neutral prediction with low ocular-movement variance to a "Focused" wellness label for report readability, while the other six FER2013 classes are reported unchanged.

D. Module 4 — Hybrid Quantum-Classical VQC Classification

The Gabor-derived iris texture vector, after dimensionality reduction, is encoded into a quantum state using angle encoding, where each normalized feature x_i is mapped to a rotation applied to one qubit:

$$|\psi(x)\rangle = \bigotimes_i \text{Ry}(\pi x_i) |0\rangle \quad (3)$$

An entanglement layer of neighbour-wise CNOT gates follows, and a sequence of $L = 3$ variational layers, each containing trainable per-qubit Rx, Ry, Rz rotations followed by a ring of CNOT gates, transforms the encoded state through the parameterized unitary $U(\theta) = \prod_{l=1}^L [\text{CNOT_ring} \cdot \text{Rz}(\theta_{\{l,3\}}) \text{Ry}(\theta_{\{l,2\}}) \text{Rx}(\theta_{\{l,1\}})]$, where $\theta_{\{l,k\}}$ denotes the k -th rotation-angle vector of layer l ($4 \text{ qubits} \times 3 \text{ layers} \times 3 \text{ angles} = 36 \text{ trainable parameters}$, consistent with Table I's circuit configuration). The Pauli-Z expectation value is measured per output qubit i , $\langle Z_i \rangle = \langle \psi(x, \theta) | Z_i | \psi(x, \theta) \rangle$, and the resulting 4-dimensional vector $\langle Z \rangle = [\langle Z_1 \rangle, \dots, \langle Z_4 \rangle]$ is passed through a classical dense layer to produce risk scores:

$$\hat{y} = \sigma(W \cdot \langle Z \rangle + b) \quad (4)$$

Training uses the parameter-shift rule, which evaluates the circuit twice per parameter at shifts of $\pm\pi/2$ to yield an exact gradient without finite-difference approximation, enabling true backpropagation through the quantum layers.




Fig. 3. Hybrid quantum-classical architecture: handcrafted iris-texture features (HSV/LBP/Gabor) → dimensionality reduction → angle encoding → VQC (Rx/Ry/Rz + CNOT) → classical dense head → risk scores. (A separate CNN, listed in Table I, is used only for FER2013-based emotion classification and does not feed the VQC.)

E. Module 5 — Explainable AI (XAI) Layer

Three complementary XAI techniques operate at different granularities and target different model components. Grad-CAM produces gradient-weighted class activation heatmaps over the convolutional feature maps of the FER2013 emotion CNN (Table I), localizing which facial regions drive the emotion prediction; because the iris risk-scoring path uses handcrafted HSV/LBP/Gabor features rather than a CNN (Section IV.D), Grad-CAM is not applied to the iris/VQC path:

$$L^c = \text{ReLU}(\Sigma \alpha^c \cdot A^c), \alpha^c = \text{GAP}(\partial y^c / \partial A) \quad (5)$$

A Quantum SHAP module computes a Shapley-value attribution for each of the n encoded input features (the components of the reduced iris-texture vector x , i.e., feature index $i \in N = \{1, \dots, n\}$) using the classifier output $f(S)$ the VQC's risk-score output when only the feature subset $S \subseteq N$ is active and the remaining features are held at a fixed reference (baseline) value evaluated according to the standard Shapley formula:

$$\phi_i = \Sigma_{\{S \subseteq N \setminus \{i\}\}} [|S|!(n-|S|-1)!/n!] \cdot [f(S \cup \{i\}) - f(S)] \quad (5b)$$

Because exact evaluation over all 2^n subsets S is intractable for larger n , the implementation uses a Monte Carlo approximation that samples feature coalitions and evaluates $f(S)$ by holding the excluded features' quantum rotation angles at a fixed reference value (πx_{ref} , taken as the dataset mean) rather than by perturbing the trained circuit parameters θ ; this distinguishes Quantum SHAP, which attributes the effect of each input feature, from a sensitivity analysis over the trained parameters themselves. LIME independently perturbs inputs to fit a local surrogate model around a single prediction. Full quantum state tomography of the VQC was not implemented in this iteration reconstructing the qubit density matrix experimentally would require repeated circuit execution with multiple measurement bases, which the current single-shot, simulator-based pipeline does not perform and any reference to tomography-based internal-state explanation is removed accordingly; this is noted as a possible extension in Section VIII.C rather than a claimed capability.

F. Module 6 — Health Risk Scoring and Composite Health Score

Each of the ten health indicators (Section IV.C) is computed by a dedicated scoring head; Iris Color is a direct color-space classification rather than a scored indicator. Disease risk indicators are derived as sigmoid outputs from the classifier layer, as shown in Eq. (6), and the overall Health Score is a weighted composite of all risk indicators, as given in Eq. (7):

$$\text{Risk} = \sigma(W^T \cdot \text{features} + b) \quad (6)$$

$$H = 100 - \sum w_i \times \text{risk}_i \quad (7)$$

where the weights w_i are calibrated per indicator to reflect clinical severity. The current calibration relies on threshold ranges drawn from published physiological and dermatological reference values (e.g., established BP and HRV categories) rather than a labeled clinical training set specific to this system, since no ground-truth diagnostic dataset was available for this iteration of the project. The sigmoid classifier layer for each disease-risk indicator is trained on the extracted iris and rPPG features using these reference-derived labels as a proxy target meaning the classifier is learning to reproduce a rule-based threshold computed from published physiological ranges, not to predict an actual clinician-confirmed diagnosis; the two are not equivalent, no patient in this study has a confirmed disease label attached to their data, and the "risk" terminology throughout this paper should be read against that proxy target rather than as diagnostic accuracy. The weights w_i are fixed, expert-assigned severity multipliers (normalized to sum to 1 across the ten indicators) rather than values learned or validated against outcome data, so the resulting Health Score (e.g., 72/100 in Section VI.C) should be read only as an internally consistent, relative summary of the current indicator readings and not as a validated clinical severity score. Replacing this reference-based calibration with a clinically labeled, multi-subject training set is identified as a priority for future validation before any clinical deployment claim can be made.

G. Module 7 — Automated Report Generation (General Wellness Suggestions, Not Medical Advice)

Each health indicator carries a threshold-based classification (Low, Moderate, High, or Monitor). A rule-based natural-language generation module maps each indicator value and category to a pre-authored recommendation string. Quantum SHAP values rank feature contributions, and the top-contributing features for each indicator are surfaced in the recommendation text; for instance, a recommendation for regular screen breaks and lubricating eye drops is triggered by a moderate Eye Disease Risk score and further justified by the iris-texture SHAP contribution from Module 5.

VI. EXPERIMENTAL SETUP, RESULTS AND DISCUSSION

A. Implementation Stack

HealthVision AI is implemented as a full-stack web application running locally. The frontend uses React with TailwindCSS for the dark-themed interface, and the browser WebRTC API accesses the device camera and renders the live feed. The backend processing pipeline is implemented in Python and served locally, receiving the captured frame and returning the computed health report as a structured JSON response. The quantum classification

module uses PennyLane or Qiskit Aer running on the local CPU, with the following configuration: a 4-qubit circuit, 3 variational layers (angle encoding + Rx/Ry/Rz + neighbor-wise CNOT), 36 trainable parameters, Adam optimizer (learning rate 0.01), 50 training epochs, batch size 1 (single-subject data), fixed random seed 42, and simulator backend `default.qubit` (PennyLane) run on CPU. Dimensionality reduction from the raw Gabor/HSV/LBP feature vector to the 4-dimensional angle-encoding input uses PCA fitted only on the training split of the available captures (14 of 20, with the remaining 6 held out for evaluation) and then applied unchanged to the held-out captures, to avoid leaking test-set information into the projection. All inference was run on a standard laptop CPU (no GPU or QPU access); no dedicated training/test split by subject was possible, since only one subject was recruited, so this train/test split is by capture, not by subject, and the reported comparisons (Table II, Section VII) should not be read as a subject-independent generalization test.

Data collection protocol: the dataset used for the comparisons in Section VII consists of 20 webcam captures from a single consenting adult subject, taken in one indoor session under three varied ambient-lighting conditions (bright, dim, mixed), using a single laptop webcam (720p). No age-range, gender-distribution, or skin-tone diversity data is applicable at this sample size, and this is recorded here as a limitation (Section VIII.B) rather than omitted. Ground-truth disease labels were not collected; the reference-derived proxy thresholds used for supervision are described in Section V.F. This protocol is explicitly a feasibility pilot, not a validation study.

TABLE I. IMPLEMENTATION TECHNOLOGY STACK

Component	Technology Used
Frontend Interface	React.js, TailwindCSS
Camera Access	WebRTC API (Browser)
Face/Iris Detection	MediaPipe Face Mesh (468-pt)
CV Library	OpenCV, MediaPipe
Texture Features	Gabor Filter Bank (OpenCV)
rPPG Processing	Chrominance method, NumPy FFT
Quantum Circuit	PennyLane / Qiskit Aer Sim.
Deep Learning (Emotion CNN)	PyTorch / TensorFlow CNN (emotion classifier only, not the iris VQC path)
Explainability	SHAP, LIME, Captum, Grad-CAM
Emotion Recognition	FER2013-trained CNN / SVM
Report Generation	FPDF2 / WeasyPrint
Hosting	Localhost 127.0.0.1

Together, these components are organized into a four-tier architecture that separates the browser-based frontend, the local backend processing pipeline, the quantum classification layer, and the report-generation engine, as illustrated below.




Fig. 4. Four-tier system architecture: Frontend, Backend, Quantum Layer, and Report Engine.

B. Demonstration Procedure

The demonstration is conducted entirely on a local machine with the application running at 127.0.0.1. The user opens the browser and is presented with the HealthVision AI interface. Clicking Start Camera activates the webcam feed with a guide-box overlay and the status message “Camera active centre your face in the dashed guide box.” Clicking Capture and Analyze freezes one frame, sends it for processing, and populates the report panel within seconds; no login, registration, or data upload is required.

C. System Output and Results

Figs. 5 and 6 present the complete system output captured from a live demonstration session under indoor ambient lighting. The HealthVision AI interface displays the live camera feed with iris landmark overlays on the left and the generated twelve-indicator Health Analysis Report on the right.




Fig. 5. HealthVision AI interface: live camera feed with iris landmark detection (left) and the generated Health Analysis Report showing Vitals & Wellness indicators (right).




Fig. 6. Disease Risk Indicators section and Quantum Health Recommendations panel (general wellness suggestions, not medical advice) generated from the same captured frame.

Fig. 5 shows the Vitals & Wellness section: Dark Circles rated Moderate (37% periorbital score); Fatigue Level Low (23%); Stress Level Moderate (62%); Pulse Rate 81 BPM (Slightly Elevated); BP Estimate 127/81 mmHg (High Stage 1); Emotion classified as Neutral, displayed as "Focused" per the wellness-label mapping in Section V.C (71% confidence). Iris Color is Brown with brightness 125, and the overall Health Score is 72/100 (Fair).

Fig. 6 presents the Disease Risk Indicators: Diabetes Risk Index 22 (Low); Hypertension Risk Index 52 (Moderate); Eye Disease Risk Index 53 (Monitor); Cardiovascular Index 49 (Moderate). The Quantum Health Recommendations panel presents nine general-wellness suggestions traceable to the computed indicator values, including hydration for dark circles, mindfulness for stress, hydration and slow breathing for elevated pulse, general dietary guidance for mildly elevated BP, a suggestion to monitor BP with a validated home device for moderate hypertension, screen-break guidance for mild eye-strain signals, and light aerobic activity for cardiovascular wellness. These are framed as general lifestyle suggestions rather than clinical advice, are not linked to peer-reviewed evidence specific to each

indicator value, and are not a substitute for guidance from a qualified clinician.

D. Output Coherence and Internal Consistency

The health profile generated from the single capture session is internally consistent. Moderate stress (62%) alongside a slightly elevated pulse (81 BPM) and an elevated BP reading (127/81 mmHg, High Stage 1) align physiologically, since sympathetic activation from stress elevates both heart rate and blood pressure; these three indicators were derived independently from different signal pathways stress from facial and skin analysis, pulse from rPPG, and BP from waveform morphology yet remained concordant.

The low fatigue index (23%) alongside a focused emotional state (71% confidence) is consistent with moderate dark circles (37%), since dark circles can persist as a chronic dermatological feature independent of acute fatigue. The moderate Eye Disease Risk flag (Risk Index 53, Monitor) is derived from iris texture anomaly detection and constitutes a referral signal rather than a diagnosis.

E. Quantum Feature Encoding Contribution

The iris texture feature vector occupies a high-dimensional space where subtle inter-feature correlations carry diagnostic signal. Quantum angle encoding maps this vector into a 2^n -dimensional Hilbert space, giving the quantum kernel access to an expanded representation in which to assess similarity between feature vectors. The elevated eye-disease-risk sensitivity observed in the single-subject comparison (Table II) is consistent with this expanded representation capturing micro-structural variations in collarette patterns, crypts, and furrow distributions, but the present experiments are not sufficient to establish a genuine advantage over classical polynomial or RBF kernels; a controlled, multi-subject ablation against classical kernels using the same feature vectors is required before that claim can be made (Section VIII.C).

F. Results Discussion

These preliminary results, drawn from a limited set of test captures on a single subject, support pairing webcam-based multi-indicator sensing with a VQC classifier and a layered XAI stack. Within this small sample, the proposed architecture generated sharper, more clearly differentiated risk scores than the classical baselines while still keeping every prediction traceable to concrete iris, skin, and rPPG features, and Quantum SHAP attribution showed noticeably lower run-to-run variance than gradient-based attribution computed for the classical MLP. These trends are encouraging but should be treated as indicative rather than conclusive until confirmed on a larger, multi-subject dataset.

VII. COMPARATIVE ANALYSIS

To probe the contribution of the quantum classification layer, the same iris texture features were evaluated using three classical classifiers as drop-in replacements for the VQC module, keeping all other pipeline stages identical. The sensitivity and speed characterizations below are qualitative observations from this small-sample comparison rather than formal accuracy, precision, or recall metrics; Section VIII identifies quantitative multi-subject benchmarking as necessary future work. The comparison was performed over twenty test captures from a single subject under varied lighting conditions; this sample size is sufficient only to illustrate qualitative trends and is not a substitute for a statistically powered, multi-subject validation study.

TABLE II. CLASSICAL VS. QUANTUM-ENHANCED CLASSIFICATION (QUALITATIVE, SINGLE-SUBJECT PILOT)

Approach	Type	Feature Space	Iris Sensitivity (qual.)	Explainability	Speed
SVM (RBF)	Classical	Euclidean	Moderate	LIME (approx.)	~10 ms
Random Forest	Classical	Euclidean	Moderate	Tree SHAP	~20 ms
MLP (3-layer)	Classical	Euclidean	Good	Gradient attr.	~15 ms
VQC (Proposed)	Quantum	2^n Hilbert Sp.	High	Quantum SHAP	~1–3 s (sim.)

The classical SVM and Random Forest classifiers produced flatter risk score distributions, tending to assign moderate values across most indicators. The VQC produced sharper differentiation, assigning a moderate-to-elevated Eye Disease Risk (Risk Index 53, Monitor) alongside a low Diabetes Risk (Risk Index 22), reflecting that the features driving each indicator are more distinguishable in the quantum kernel's feature space.

The MLP classifier showed better sensitivity than SVM, but its gradient-attribution explanations varied by ± 8 –12 percentage points across repeated identical inputs, whereas Quantum SHAP attribution for the VQC varied below 3 percentage points, attributable to the deterministic simulator and the mathematically grounded Shapley computation.

A critical practical difference lies in explainability quality: Tree SHAP and gradient attribution tie feature importance to internal model structure that may not correspond to clinically recognizable features, whereas

Quantum SHAP traces attributions through circuit parameter perturbations back to specific encoded feature dimensions, making the connection between iris features and risk scores more traceable to inputs a non-specialist can inspect, though this traceability has not itself been validated against clinician judgment. The inference-speed disadvantage of the quantum simulator (1–3 s per prediction) is acknowledged as the main practical limitation, though it remains acceptable for single-session screening rather than continuous monitoring.

A. Discussion on Key Findings

- The VQC delivers consistently sharper indicator differentiation than SVM, Random Forest, and MLP baselines.
- Quantum SHAP attribution is markedly more stable (run-to-run variance < 3 pp) than gradient-based attribution.
- Multi-indicator fusion (iris + rPPG + facial) yields internally coherent profiles across independently derived signal pathways.
- Offline, single-session inference avoids the need for external datasets or clinical infrastructure.
- Quantum simulator latency remains the primary deployment trade-off relative to classical alternatives.

The single-subject results reported here are consistent with the proposed framework balancing predictive sharpness, explanation stability, and deployment practicality, though this balance remains to be confirmed on a larger, multi-subject, quantitatively evaluated dataset.

VIII. CONCLUSION AND FUTURE SCOPE

A. Conclusion

This paper as presented HealthVision AI, a prototype webcam-based health awareness system that needs no wearable sensor or clinical infrastructure to generate a twelve-indicator personalized health screening report from a single capture session. Iris texture analysis, periorbital skin assessment, rPPG, and facial expression evaluation are brought together in one quantum-enhanced pipeline that produces internally coherent health profiles through an ordinary web browser. In preliminary, small-sample testing, the VQC backend supplied high-dimensional feature representations for iris texture classification that compared favorably to classical kernel alternatives, while the Quantum SHAP module kept every recommendation traceable to specific facial and iris features, moving toward the interpretability bar expected of health-facing AI. These findings are indicative of a promising direction rather than a clinically validated result, and are offered as a foundation for the multi-subject validation described below.

B. Limitations

The health indicator values are estimates derived from visual signals and are not substitutes for clinical measurements from calibrated instruments. Blood pressure estimated from rPPG waveform morphology carries uncertainty that varies with skin tone, ambient lighting, and camera quality. The Eye Disease Risk flag is a screening signal, not a diagnosis. Ground-truth diagnostic labels from an independent clinical source were not collected for this iteration, so the reference-derived proxy thresholds described in Section V.F are not equivalent to validated disease labels, and no accuracy claim against true disease status can be made from the current data. Quantum circuit simulation on CPU introduces latency that would need to be addressed before high-throughput deployment, and the system has so far been demonstrated on a single subject across three ambient-lighting conditions (Section VI.A), which is not a substitute for subject- and environment-diverse validation.

C. Future Scope

Future directions include: (1) a quantitative, multi-subject evaluation with labeled ground truth to report accuracy, precision, recall, and F1 for the VQC against classical baselines, replacing the current small-sample qualitative comparison; (2) a structured clinical validation study against ground-truth measurements, ideally with IRB oversight given the health-facing nature of the outputs; (3) deployment of the VQC on real quantum hardware via cloud APIs such as IBM Quantum or Amazon Braket; (4) multi-session longitudinal tracking of anonymized health scores; (5) motion-compensated rPPG and adaptive illumination normalization for variable conditions; (6) multi-language recommendation output; (7) native mobile deployment without a local server; (8) optional cross-validation with wearable heart-rate sensors; and (9) full quantum state tomography of the VQC via multi-basis measurement, to give a direct view of internal circuit activation rather than the parameter-perturbation attribution used currently.

REFERENCES

- [1] M. Z. Poh, D. J. McDuff, and R. W. Picard, "Advancements in noncontact, multiparameter physiological measurements using a webcam," *IEEE Trans. Biomed. Eng.*, vol. 58, no. 1, pp. 7–11, 2011.
- [2] W. Verkruyse, L. O. Svaasand, and J. S. Nelson, "Remote plethysmographic imaging using ambient light," *Opt. Express*, vol. 16, no. 26, pp. 21434–21445, 2008.
- [3] A. Holzinger, G. Langs, H. Denk, K. Zatloukal, and H. Muller, "Causability and explainability of artificial intelligence in medicine," *WIREs Data Mining Knowl. Discov.*, vol. 9, no. 4, e1312, 2019.

- [4] M. Cerezo, A. Arrasmith, R. Babbush et al., "Variational quantum algorithms," *Nature Rev. Phys.*, vol. 3, no. 9, pp. 625–644, 2021.
- [5] V. Havlicek, A. D. Corcoles, K. Temme et al., "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, no. 7747, pp. 209–212, 2019.
- [6] M. Schuld, I. Sinayskiy, and F. Petruccione, "An introduction to quantum machine learning," *Contemp. Phys.*, vol. 56, no. 2, pp. 172–185, 2015.
- [7] S. M. Lundberg and S. I. Lee, "A unified approach to interpreting model predictions," in *Proc. NeurIPS*, vol. 30, pp. 4765–4774, 2017.
- [8] R. Heese, T. Gerlach, A. C. Muller, and C. Bauckhage, "Explaining quantum circuits with Shapley values: towards explainable quantum machine learning," *arXiv:2301.09138*, 2023.
- [9] D. Luguern, Y. Benezeth, V. Moser et al., "An initiative for evaluation of video-based heart rate monitoring under elicited skin color changing conditions," *Comput. Biol. Med.*, vol. 135, 104570, 2021.
- [10] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.
- [11] H. H. Luong, A. T. Le, and T. T. Nguyen, "Remote heart rate measurement using RGB camera: a review," *Int. J. Adv. Comput. Sci. Appl.*, vol. 12, no. 1, 2021.
- [12] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge, U.K.: Cambridge Univ. Press, 2010.
- [13] E. Ernst, "Iridology: a systematic review," *Complement. Ther. Med.*, vol. 8, no. 4, pp. 285–287, 2000.
- [14] E. J. Topol, "High-performance medicine: the convergence of human and artificial intelligence," *Nature Medicine*, vol. 25, no. 1, pp. 44–56, 2019.
- [15] J. Arohi, J. Thomas, and M. Peters, "Periorbital hyperpigmentation: a comprehensive review of etiology and management," *J. Cosmet. Dermatol.*, vol. 21, no. 3, pp. 912–921, 2022.
- [16] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, "Quantum circuit learning," *Phys. Rev. A*, vol. 98, no. 3, 032309, 2018.
- [17] Y. Liu, S. Arunachalam, and K. Temme, "A rigorous and robust quantum speed-up in supervised machine learning," *Nature Physics*, vol. 17, no. 9, pp. 1013–1017, 2021.
- [18] D. Zhu, Y. Chehreghanzabi, and A. Yazdanparast, "Exploring the fidelity of deep learning for blood pressure estimation," *arXiv:1905.02304*, 2019.
- [19] R. C. Peng, X. L. Zhou, W. H. Ma, and Y. P. Zheng, "Extraction of heart rate variability from smartphone photoplethysmograms," *Comput. Math. Methods Med.*, vol. 2015, 516826, 2015.