Development of a Multi-scale Hierarchical Representation of Quantum Images as a Basis for Applied Noise Reduction of B-SCAN OCT Images

Anonymous Author[1], **Anonymous Author**[2], and **Anonymous Author**[3],

ABSTRACT

This study develops a Multi-scale Hierarchical Representation of Quantum Images (MHRQI) and a heuristic denoising operation for B-scan OCT speckle suppression. Validation was performed in classical quantum simulation using a stratified sample of 32 B-scan OCT images from Kermany2018 (Normal, CNV, DME, Drusen). For encoder validation on a 128×128 synthetic image (25 qubits, 2,000,000 shots), reconstruction reached MSE = 0.0242 and PSNR = 64.29 dB. For denoiser validation on a synthetic gradient image (25 qubits, 2,000,000 shots), MSE improved from 1934.77 to 426.60 and PSNR from 15.26 to 21.83 dB. On the 32-image benchmark, classical baselines yielded stronger absolute speckle suppression (best SSI: NLM 0.545 ± 0.110; best NSF: SRAD 0.579 ± 0.180), while the proposed method preserved structure more consistently (EPF 0.924 ± 0.042, Rank 1). Wilcoxon signed-rank testing ($\alpha = 0.05$) showed EPF was significantly better than BM3D, NLM, and SRAD ($p = 9.90 \times 10^{-6}$, 0.011, 4.66×10^{-10}). These results position MHRQI as a structure-aware quantum image representation with demonstrated empirical behavior in simulation.

Keywords: Optical coherence tomography, speckle noise reduction, quantum image processing, hierarchical quantum representation, diabetic retinopathy, B-scan OCT

1. INTRODUCTION

In 2024, diabetes affected an estimated 800 million adults worldwide, primarily in low- and middle-income countries [1]. Among its complications, diabetic retinopathy (DR) stands out as a leading cause of preventable vision impairment.

DR is a microvascular complication where high blood sugar damages retinal blood vessels. Studies estimate that 25% to 27% of individuals with diabetes develop DR, with the disease burden projected to rise through 2045 [2, 3]. As such, the growing impact necessitates enhanced screening and early detection techniques; and among them is optical coherence tomography (OCT).

OCT is a non-invasive imaging modality that uses light waves to create 3D cross-sectional images of the retina [4]. Clinicians use 2D slices, known as Brightness Scans (B-SCANs), to inspect retinal layers [5].

However, B-SCAN OCT images are inherently corrupted by speckle noise, created by the constructive and destructive interference of scattered light waves. This noise acts multiplicatively: $y = x \cdot n$, where y is the observed noisy intensity, x is the true signal intensity, and n is the noise component. Because the noise is embedded within the observed signal, the true signal intensity and the noise component are unknowns [5]. Mathematically, this is an ill-posed inverse problem.

Researchers have tackled OCT speckle noise through various approaches. Classical digital image processing techniques effectively reduce noise but often require careful parameter tuning and risk over-smoothing critical anatomical features [6–8]. To address these limitations, machine learning approaches algorithmically and iteratively optimize parameters by learning from paired datasets of clean and noisy images. However, obtaining truly noise-free in vivo OCT reference images is virtually due to the nature of their acquisition process. Training datasets must therefore rely on synthetically produced references or averaged scans, that of which introduces errors and can compromise the accuracy of retinal structures despite improved overall denoising performance [5].

Quantum computing is also important because it is now being applied to concrete reliability and optimization problems, not only theoretical demonstrations. Recent ECTI-CIT studies include a surface-code error-correction analysis for quantum gate translation of a classical ALU under depolarizing noise and a D-Wave quantum annealing implementation for disaster-response resource allocation in Marikina City [9, 10].

Consequently, researchers have explored quantum-based denoising frameworks to model these pixel interactions as particles in a controlled quantum system [11–14]. Most of these approaches, however,

Manuscript received on December 15, 2025 ; revised on March 25, 2026.
Final manuscript received on March 25, 2026.
[1,2,3] Author information omitted for double-blind review.
Corresponding author information omitted for double-blind review.

simulate the physical image formation process rather than performing direct image domain processing. Upon observation of the quantum image processing domain, it is found that existing quantum image representations primarily encode pixel intensities, lacking the structural encoding required to preserve spatial relationships and multiscale anatomy. To our knowledge, speckle suppression relies on structural context, this limitation obscures how anatomical boundaries might be distinguished from noise natively within a quantum state. As a result, structure-aware quantum image processing for practical OCT denoising remains an open theoretical challenge.

This study introduces the Multi-scale Hierarchical Representation of Quantum Images (MHRQI) framework, functioning as a basis for a quantum denoising operation on B-scan OCT images. The objective is to formally define and analytically verify this quantum image representation, and to empirically test it through classical simulation whether introducing this method can offer a statistical advantage in noise suppression compared to established classical filtering methods.

2. RELATED WORKS

This section summarizes OCT denoising approaches from classical filtering, to machine learning, to quantum and quantum-inspired methods, then motivates the representational gap addressed by MHRQI.

2.1 Classical Denoising Methods

Classical OCT denoising began with spatial filtering and wavelet-domain methods, but these approaches can blur retinal boundaries or under-handle multiplicative speckle noise [15]. Nonlocal and transform-domain methods (e.g., NLM and BM4D) improve denoising by exploiting patch similarity [16, 17], while diffusion and regularization-based methods such as SRAD, TV, and WNNM better model OCT noise statistics and edge preservation [8, 18, 19]. Their main practical limitation is parameter sensitivity, where stronger smoothing can suppress diagnostically relevant structures.

2.2 Artificial Intelligence and Machine Learning

Deep learning methods, including CNN and GAN variants, often improve quantitative denoising metrics compared with classical filters [20, 21]. However, supervised OCT denoising is constrained by the absence of truly noise-free in vivo references, so training commonly depends on synthetic or averaged targets [5, 22]. This can reduce performance stability under distribution shift and introduces risk of anatomically misleading outputs when optimization prioritizes learned objectives over faithful structure retention [23].

2.3 Quantum and Quantum-Inspired Image Processing

Quantum-related OCT research follows two directions. Hardware-focused studies use nonclassical light and report benefits such as dispersion-related improvements [11, 14, 24]. Algorithmic quantum-inspired studies apply quantum formalisms to classically acquired images and report adaptive denoising behavior [25–27]. Both directions are valuable, but neither directly performs denoising on a natively encoded quantum image representation used for downstream image-domain operations.

2.4 Quantum Image Representations

Existing quantum image representations (amplitude or basis-state families) encode positions and intensities efficiently and reversibly [28–30]. Yet they primarily act as coordinate-addressing schemes and provide limited native multiscale spatial context. This limitation complicates local structural consistency checks needed for speckle suppression, motivating the hierarchical spatial register introduced by MHRQI.

3. METHODOLOGY

This section details the Multi-scale Hierarchical Representation of Quantum Images (MHRQI) algorithm and the corresponding quantum denoising operation developed for B-scan OCT images. The encoder translates classical image data into a structural quantum state, and the denoiser evaluates local neighborhood consistency to suppress speckle noise.

3.1 Quantum Image Representation

First, the classical mathematical formalism preprocesses positional and intensity data before executing the quantum encoding circuit.

3.1.1 Classical Preprocessing

3.1.1.a Pixel Intensity NormalizationLet a classical image be represented as a matrix of pixel intensity values. For a grayscale image with dimensions $N \times N$, the matrix is defined as:

$$I = \begin{bmatrix} I(0,0) & I(1,0) & \cdots & I(N-1,0) \\ I(0,1) & I(1,1) & \cdots & I(N-1,1) \\ \vdots & \vdots & \ddots & \vdots \\ I(0,N-1) & I(1,N-1) & \cdots & I(N-1,N-1) \end{bmatrix} \quad (1)$$

For a b-bit grayscale image, the number of possible intensity levels L is given by:

$$L = 2^b \quad (2)$$

Each pixel intensity satisfies the valid range:

$$0 \leq I(x, y) \leq L - 1 \tag{3}$$

Intensity values are then expressed in binary form using exactly b bits. For a pixel at position (x, y), the binary representation is:

$$I(x, y) = (b_{b-1} b_{b-2} \dots b_0)_2 \tag{4}$$

where each bit $b_k \in \{0, 1\}$. The equivalent bitwise expanded form is:

$$I(x, y) = \sum_{k=0}^{b-1} b_k 2^k \tag{5}$$

This converts each grayscale value into binary components stored in a 1D array using row-major ordering. For an image width N, each pixel (x, y) is assigned a linear index:

$$i = y \cdot N + x \tag{6}$$

The binary vector stored per pixel is:

$$B[i] = (b_{b-1}, b_{b-2}, \dots, b_0)_{(x,y)} \tag{7}$$

Producing the final linear array of binary intensity values:

$$B = [B[0], B[1], B[2], \dots, B[N^2 - 1]] \tag{8}$$

3.1.1.b Position Coordinate RemappingTo organize pixel positions hierarchically, the image is divided into progressively smaller subdivisions. At each level k, the size of each subdivision s_k is defined as:

$$s_k = \frac{N}{d^k} \tag{9}$$

where d is the subdivision factor. The maximum classical subdivision level L_{max} is computationally determined as:

$$L_{max} = \lfloor \log_d N \rfloor \tag{10}$$

At each subdivision level k, a pixel's hierarchical index in the x and y directions is computed individually to locate its sub-block:

$$\begin{cases} Q_k(x) = \min(\lfloor \frac{x \pmod{s_{k-1}} \cdot d}{s_{k-1}} \rfloor, d - 1) \\ Q_k(y) = \min(\lfloor \frac{y \pmod{s_{k-1}} \cdot d}{s_{k-1}} \rfloor, d - 1) \end{cases} \tag{11}$$

Concatenating these indices across all levels produces the multi-scale hierarchical coordinate vector $H(x, y)$:

$$\begin{aligned} H(x, y) = (&Q_1(x), Q_1(y), Q_2(x), Q_2(y), \\ &\dots, Q_{L_{max}}(x), Q_{L_{max}}(y)) \end{aligned} \tag{12}$$

The complete combination across the image forms the hierarchical matrix $\mathbf{H}$:

$$\mathbf{H} = \begin{bmatrix} H(0,0) & \cdots & H(N-1,0) \\ H(0,1) & \cdots & H(N-1,1) \\ \vdots & \ddots & \vdots \\ H(0,N-1) & \cdots & H(N-1,N-1) \end{bmatrix} \tag{13}$$

3.1.2 Encoder Quantum Formulae

The quantum system requires a position register of paired qudits per hierarchical level, $|q_k^x\rangle$ and $|q_k^y\rangle$, alongside an intensity register of b qubits. The full system is initialized to the zero state:

$$|\Psi_0\rangle = |0\rangle_{q_x}^{\otimes L} \otimes |0\rangle_{q_y}^{\otimes L} \otimes |0\rangle_I^{\otimes b} \tag{14}$$

The position register creates a uniform superposition over all hierarchical spatial coordinates, while the intensity register remains zero at this stage:

$$\begin{aligned} |\Psi_1\rangle = \frac{1}{\sqrt{d^{2L}}} \sum_{q_1^x=0}^{d-1} \sum_{q_1^y=0}^{d-1} \cdots \sum_{q_L^x=0}^{d-1} \sum_{q_L^y=0}^{d-1} \\ |q_1^x, q_1^y, \dots, q_L^x, q_L^y\rangle \otimes |0\rangle_I^{\otimes b} \end{aligned} \tag{15}$$

The computational basis state representing the intensity for pixel (x, y) is then explicitly encoded:

$$|I(x, y)\rangle = |b_{b-1}\rangle \otimes |b_{b-2}\rangle \otimes \cdots \otimes |b_0\rangle \tag{16}$$

A controlled intensity loading operator sets the intensity register to $|I(x, y)\rangle$ whenever the position register equals $|H(x, y)\rangle$:

$$\begin{aligned} W_{x,y} = &|H(x,y)\rangle\langle H(x,y)| \otimes X_{I(x,y)} \\ &+ (\hat{I} - |H(x,y)\rangle\langle H(x,y)|) \otimes \hat{I}_I \end{aligned} \tag{17}$$

The global sequence applies this across all pixels while properly conditioning on the position register:

$$W = \prod_{y=0}^{N-1} \prod_{x=0}^{N-1} W_{x,y} \tag{18}$$

Yielding the quantum state after intensity encoding:

$$|\Psi_2\rangle = W |\Psi_1\rangle \tag{19}$$

Creating the final MHRQI state, directly entangling hierarchical position with exact binary intensity:

$$|MHRQI\rangle = \frac{1}{\sqrt{d^{2L}}} \sum_{y=0}^{N-1} \sum_{x=0}^{N-1} |H(x, y)\rangle \otimes |I(x, y)\rangle \tag{20}$$

Algorithm 1 MHRQI Encoder Algorithm

Require: Grayscale image I of size $N \times N$, bit depth b, subdivision factor d
Ensure: Encoded quantum state $|MHRQI\rangle$
Calculate max subdivision levels $L_{\max} = \lfloor \log_d N \rfloor$
Initialize position register q with $2L_{\max}$ qudits to $|0\rangle$
Initialize intensity register i with b qubits to $|0\rangle$
Initialize workspace ancilla w to $|0\rangle$
Apply Hadamard gates to position register q to create uniform superposition
for each pixel (x, y) in image I **do**
Map (x, y) to hierarchical vector $H(x, y)$
Convert intensity $I(x, y)$ to binary string $(b_{b-1} \dots b_0)$
Apply Pauli-X to position register q where $H(x, y)$ bits are 0
Apply multi-controlled NOT from q to flip ancilla w to $|1\rangle$
for $k = 0$ to $b - 1$ **do**
if $b_k = 1$ **then**
Apply CNOT from ancilla w to intensity qubit i_k
end if
end for
Uncompute multi-controlled NOT to restore ancilla w to $|0\rangle$
Uncompute Pauli-X on position register q
end for

3.2 Quantum Denoising Operation

3.2.1 Denoiser Formulae

The transition from the initialized state to the final denoised state is governed by a composite unitary operator, $\hat{U}_{denoise}$. To ensure strict mathematical rigor, each sub-operation of the denoising circuit is defined as a tensor product across the entire joint Hilbert space. The identity operator ($\hat{I}$) is explicitly applied to all registers not actively manipulated by a given gate sequence.

The initialized state introduces three ancillae (parent-average $|0\rangle_p$, consistency $|0\rangle_c$, and outcome $|0\rangle_o$) directly to the encoded image:

$$|\Psi\rangle = \frac{1}{\sqrt{M}} \sum |H(x, y)\rangle \otimes |I(x, y)\rangle \otimes |0\rangle_p \otimes |0\rangle_c \otimes |0\rangle_o \tag{21}$$

Next, a sibling superposition operator ($\hat{U}_{H_{fine}}$) isolates the finest-level spatial subdivision by applying a Hadamard gate ($\hat{H}$) to both the x and y coordinate qubits at level $L-1$, while explicitly maintaining the state of all parent-level positions, the intensity register, and the ancillae.

$$\begin{aligned}\hat{U}_{H_{fine}} = \left(\bigotimes_{k=0}^{L-2} \hat{I}_{q_k^y} \otimes \hat{I}_{q_k^x} \right) \otimes \left(\hat{H}_{q_{L-1}^y} \otimes \hat{H}_{q_{L-1}^x} \right) \\ \otimes \hat{I}_{int} \otimes \hat{I}_{p,c,o}\end{aligned} \tag{22}$$

Subsequently, a parent-level proxy operator ($\hat{U}_{R_p}$) accumulates the parent-level intensity proxy within the work ancilla p. The rotation applied to p is strictly dependent on the computational basis state of the b-qubit intensity register.

$$\begin{aligned}\hat{U}_{R_p} = \left(\bigotimes_{k=0}^{L-1} \hat{I}_{q_k^y} \otimes \hat{I}_{q_k^x} \right) \otimes \hat{I}_c \otimes \hat{I}_o \\ \otimes \left(\sum_{i \in \{0,1\}^b} |i\rangle\langle i|_{int} \otimes \hat{R}_p(i)_p \right)\end{aligned} \tag{23}$$

Where $\hat{R}_p(i)_p$ represents the descending rotational weights controlled by the specific bits $(i_{b-1}, \dots, i_{b-4})$ of the intensity state $|i\rangle$:

$$\begin{aligned}\hat{R}_p(i)_p = \hat{R}_y \left(\frac{\pi}{16} \right)^{i_{b-1}} \hat{R}_y \left(\frac{\pi}{8} \right)^{i_{b-2}} \\ \cdot \hat{R}_y \left(\frac{\pi}{4} \right)^{i_{b-3}} \hat{R}_y \left(\frac{\pi}{2} \right)^{i_{b-4}}\end{aligned} \tag{24}$$

Then, a consistency logic operator ($\hat{U}_{XNOR}$) executes the reversible XNOR comparison. It evaluates the Most Significant Bit (MSB) of the intensity register (i_{b-1}) against the state of the parent-average ancilla (p), targeting the consistency ancilla (c).

$$\begin{aligned}\hat{U}_{XNOR} = \left(\bigotimes_{k=0}^{L-1} \hat{I}_{q_k^y} \otimes \hat{I}_{q_k^x} \right) \otimes \Bigg(\sum_{i \in \{0,1\}^b} \sum_{v \in \{0,1\}} \\ |i\rangle\langle i|_{int} \otimes |v\rangle\langle v|_p \otimes \hat{X}_c^{\delta(i_{b-1}, v)} \Bigg) \otimes \hat{I}_o\end{aligned} \tag{25}$$

The exponent uses the Kronecker delta function $\delta(i_{b-1}, v)$, which equals 1 if the MSB of the intensity state (i_{b-1}) matches the state of the parent proxy (v), applying the Pauli-X gate to c.

Following this, an outcome marking operator ($\hat{U}_{mark}$) records the denoising decision on the outcome qubit (o). If the consistency ancilla indicates a mismatch ($c = 0$), the outcome qubit is flipped.

$$\begin{aligned} \hat{U}_{mark} = & \left(\bigotimes_{k=0}^{L-1} \hat{I}_{q_k^y} \otimes \hat{I}_{q_k^x} \right) \otimes \hat{I}_{int} \otimes \hat{I}_p \\ & \otimes \Big(|0\rangle\langle 0|_c \otimes \hat{X}_o + |1\rangle\langle 1|_c \otimes \hat{I}_o \Big) \end{aligned} \tag{26}$$

The final composite unitary operator $\hat{U}_{denoise}$ is the forward sequence of these transformations followed by the uncomputation sequence required to cleanly release the working ancillae and restore the position register basis. Because the Hadamard and XNOR operations are their own inverses, their adjoints (†) reflect this precise uncomputation:

$$\begin{aligned} \hat{U}_{denoise} = & \hat{U}^{\dagger}_{H_{fine}} \cdot \hat{U}^{\dagger}_{R_p} \cdot \hat{U}^{\dagger}_{XNOR} \\ & \cdot \hat{U}_{mark} \cdot \hat{U}_{XNOR} \cdot \hat{U}_{R_p} \cdot \hat{U}_{H_{fine}} \end{aligned} \tag{27}$$

Applying this composite operator to the initialized state yields the final, mathematically reversible denoised state:

$$\begin{aligned} |\Psi_{denoise}\rangle &= \hat{U}_{denoise} |\Psi\rangle \\ &= \frac{1}{\sqrt{M}} \sum |H(x,y)\rangle \otimes |I(x,y)\rangle \\ & \quad \otimes |0\rangle_p \otimes |0\rangle_c \otimes |o(x,y)\rangle \end{aligned} \tag{28}$$

3.3 Proof Development

To validate core correctness properties before simulation, the framework was mechanically checked in Lean 4. The proof goals were: (i) bijective coordinate-to-hierarchy mapping for lossless positional encoding, (ii) reversibility of the denoising circuit through proper uncomputation, and (iii) correctness of the outcome-qubit consistency logic under the formal rule set.

3.4 Software Development

The implementation uses a Python/Qiskit workflow executed through modular scripts for circuit construction, preprocessing/reconstruction, baseline comparison, and statistical benchmarking. Large-batch runs were executed on high-performance computing resources using SLURM-managed jobs and isolated software environments, enabling reproducible dataset-scale experiments.

3.5 Experimental Setup

Experiments used a stratified sample of 32 B-scan OCT images from Kermany2018 (Normal, CNV, DME, Drusen). The proposed method was compared against BM3D, NLM, and SRAD under shared input conditions. Evaluation used speckle-focused metrics (SSI, SMPI, NSF, ENL, CNR) and structure-focused metrics (EPF, EPI, OMQDI), followed by paired significance testing using the Wilcoxon signed-rank test at $\alpha = 0.05$.

4. RESULTS AND DISCUSSION

This section presents the analytical validation and empirical results of the encoding algorithm, the quantum denoising operation, and the overall system evaluation applied to B-SCAN OCT images. The system was benchmarked against classical denoising methods BM3D, NLM, and SRAD.

4.1 Results of the Encoding Algorithm

Analytical proofs confirm that the encoding algorithm mathematically guarantees a lossless representation of the classical image data.

4.1.1 Analytical Validation

The structural mapping equations of the MHRQI encoder were mechanically verified using Lean 4. The formal proofs confirm that the spatial mapping from the 2D Cartesian plane of pixel coordinates into the multiscale hierarchical register is strictly bijective. This guarantees that the quantum image representation algorithm introduces zero data loss and zero index collisions, thereby losslessly encoding classical spatial data into a quantum state.

4.1.2 Empirical Results

The encoding circuit was executed on a 128×128 synthetic classical test image. The 25-qubit circuit (depth 130,197) was simulated using 2,000,000 measurement shots. The reconstructed image obtained a Mean Squared Error (MSE) of 0.0242 and a Peak Signal-to-Noise Ratio (PSNR) of 64.29 dB when compared directly to the original image. Minor spatial deviations were distributed randomly, consistent with expected statistical sampling variance rather than structural distortion. These results confirm the analytically proven lossless correlation.

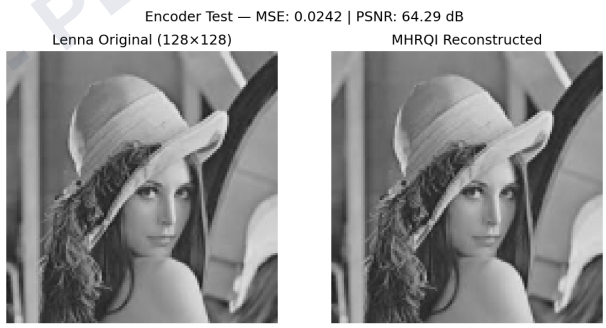

Fig.1: *Encoder Empirical Results using a Synthetic Test Image*

Algorithm 2 Quantum Denoising Consistency Operation

Require: Encoded $|MHRQI\rangle$ state, bit depth b
Ensure: Denoised consistency outcome state $|\Psi_{denoise}\rangle$
Initialize ancillae parent-average p, consistency c, outcome o to $|0\rangle$
Apply Hadamard to finest-level position qubits q^x_{L-1}, q^y_{L-1} (sibling superposition)
Apply controlled $RY(\pi/16)$ from I_{b-1} to ancilla p
Apply controlled $RY(\pi/8)$ from I_{b-2} to ancilla p
Apply controlled $RY(\pi/4)$ from I_{b-3} to ancilla p
Apply controlled $RY(\pi/2)$ from I_{b-4} to ancilla p
Apply Hadamard to finest-level position qubits q^x_{L-1}, q^y_{L-1}
Apply Toffoli-sequence (XNOR) between I_{b-1} and p, target c
Apply CNOT from c to outcome o conditioned on inconsistency
Uncompute Toffoli-sequence on c
Uncompute RY operations on p

4.2 Results of the Denoising Operation

4.2.1 Analytical Validation

The discrete logic of the quantum denoising circuit was similarly verified in Lean 4. The formal proofs confirm that the algorithmic sequence is theoretically reversible, verifying that temporary computational workspaces uncompute to their initial states to preserve overall quantum coherence. Additionally, the verification ensures that the outcome qubit evaluates to $|1\rangle$ if and only if the structural inconsistency condition is met, confirming the denoiser's logic operates strictly according to its mathematical definition.

4.2.2 Empirical Results

To isolate the consistency logic, the denoising circuit operated on a 128×128 synthetic image featuring a sharp vertical gradient and artificially injected pixel-level speckle noise. The 25-qubit simulated circuit (2,000,000 shots) recorded quantitative error reduction. The noisy input registered an MSE of 1934.77 (15.26 dB PSNR), which improved to an MSE of 426.60 (21.83 dB PSNR) after denoising. The outcome qubit flagged inconsistent pixels as defined, smoothing the speckle while maintaining the gradient boundary in the simulation environment.

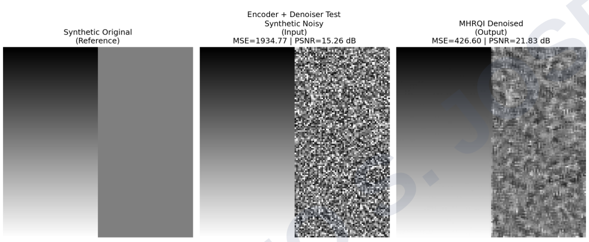

Fig.2: *Empirical Validation of the Denoiser on a Synthetic Image*

4.3 System Evaluation on B-SCAN OCT Images

The system was evaluated against baseline methods using a stratified sample of 32 B-SCAN OCT images (Normal, CNV, DME, Drusen).

4.3.1 Speckle Reduction Performance

Table 1 reports the mean quantitative results for Speckle Suppression Index (SSI), Speckle Suppression and Mean Preservation Index (SMPI), Noise Suppression Factor (NSF), Equivalent Number of Looks (ENL), and Contrast-to-Noise Ratio (CNR).

Classical baselines achieved stronger absolute variance suppression: NLM gave the best SSI (0.545 ± 0.110), SRAD gave the best NSF (0.579 ± 0.180), and BM3D gave the highest CNR (59.359 ± 155.675). The proposed method ranked fourth in SSI, SMPI, NSF, ENL, and CNR (SSI 0.933 ± 0.062, NSF 0.127 ± 0.075), consistent with its conservative, structure-preserving denoising rule.

4.3.2 Structural Similarity Performance

The structural consistency findings (Edge Preservation Factor (EPF), Edge Preservation Index (EPI), and Objective Measure of Quality of Denoised Images (OMQDI)) are summarized in Table 2.

For structural metrics, the proposed method achieved the highest EPF (0.924 ± 0.042, Rank 1), with lower EPF dispersion than BM3D (± 0.130), NLM (± 0.105), and SRAD (± 0.177). It ranked third in EPI (0.924 ± 0.043) and second in OMQDI (1.032 ± 0.033), indicating a favorable edge-preservation profile despite weaker absolute speckle suppression.

4.3.3 Statistical Significance

A Wilcoxon signed-rank test ($\alpha = 0.05$) compared differences between the proposed method and the baseline algorithms. Table 3 reports the p-values and hypothesis testing results.

Wilcoxon tests reject H_0 for SSI and ENL against all baselines ($p = 4.66 \times 10^{-10}$), confirming stronger absolute smoothing by classical methods. For EPF,

Table 1: *Mean Speckle Reduction Metrics Results*

Algorithm	SSI↓ (Rank)	SMPI↓ (Rank)	NSF↑ (Rank)	ENL↑ (Rank)	CNR↑ (Rank)
BM3D	0.707 ± 0.083 (#3)	0.976 ± 0.019 (#3)	0.212 ± 0.171 (#2)	160.139 ± 74.083 (#3)	**59.359 ± 155.675 (#1)**
Non-Local Means	**0.545 ± 0.110 (#1)**	0.975 ± 0.017 (#2)	0.137 ± 0.115 (#3)	**292.836 ± 171.602 (#1)**	42.189 ± 105.090 (#2)
SRAD	0.669 ± 0.077 (#2)	**0.950 ± 0.023 (#1)**	**0.579 ± 0.180 (#1)**	175.648 ± 77.196 (#2)	20.820 ± 45.559 (#3)
Proposed	0.933 ± 0.062 (#4)	0.977 ± 0.012 (#4)	0.127 ± 0.075 (#4)	87.427 ± 28.797 (#4)	9.994 ± 17.194 (#4)

Table 2: *Mean Structural Similarity Metrics Results*

Algorithm	EPF (Rank)	EPI (Rank)	OMQDI (Rank)
BM3D	0.820 ± 0.130 (3)	0.937 ± 0.039 (2)	0.982 ± 0.080 (4)
NLM	0.865 ± 0.105 (2)	0.944 ± 0.033 (1)	0.982 ± 0.064 (3)
SRAD	0.513 ± 0.177 (4)	0.823 ± 0.068 (4)	1.047 ± 0.091 (1)
Proposed	0.924 ± 0.042 (1)	0.924 ± 0.043 (3)	1.032 ± 0.033 (2)

Table 3: *Wilcoxon Signed-Rank Test Results*

Metric	vs. BM3D (p)	vs. NLM (p)	vs. SRAD (p)
SSI	4.66×10^{-10} (Reject H_0)	4.66×10^{-10} (Reject H_0)	4.66×10^{-10} (Reject H_0)
SMPI	0.019 (Reject H_0)	0.500 (Fail to reject H_0)	4.66×10^{-10} (Reject H_0)
NSF	0.012 (Reject H_0)	0.465 (Fail to reject H_0)	4.66×10^{-10} (Reject H_0)
ENL	4.66×10^{-10} (Reject H_0)	4.66×10^{-10} (Reject H_0)	4.66×10^{-10} (Reject H_0)
CNR	0.003 (Reject H_0)	2.31×10^{-6} (Reject H_0)	0.028 (Reject H_0)
EPF	9.90×10^{-6} (Reject H_0)	0.011 (Reject H_0)	4.66×10^{-10} (Reject H_0)
EPI	0.304 (Fail to reject H_0)	0.018 (Reject H_0)	3.26×10^{-9} (Reject H_0)
OMQDI	2.84×10^{-5} (Reject H_0)	6.91×10^{-7} (Reject H_0)	0.747 (Fail to reject H_0)

the proposed method is significantly different from BM3D, NLM, and SRAD ($p = 9.90 \times 10^{-6}$, 0.011, 4.66×10^{-10}), supporting the observed structural-preservation advantage.

4.3.4 Visual Results

Visual inspection indicates the proposed simulated algorithm smooths high-frequency speckle while retaining localized layer interfaces and pathological structures more clearly than the evaluated classical diffusion approaches, though at the expense of leaving residual granular noise.

5. CONCLUSION AND RECOMMENDATIONS

This study formalized MHRQI as a hierarchical quantum image representation with a reversible consistency-based denoising circuit, then validated both analytically and empirically. Encoder simulation on a 128×128 test image (25 qubits, 2,000,000 shots) achieved MSE 0.0242 and PSNR 64.29 dB. Denoiser simulation on synthetic noisy gradients improved MSE from 1934.77 to 426.60 and PSNR from 15.26 to 21.83 dB.

On 32 B-scan OCT images, baseline methods remained stronger for absolute speckle suppression (best SSI: NLM 0.545 ± 0.110; best NSF: SRAD 0.579 ± 0.180), while MHRQI achieved the strongest EPF (0.924 ± 0.042, Rank 1). Wilcoxon testing confirmed this tradeoff: absolute-noise metrics favored classical methods, whereas EPF showed significant differences versus all baselines. The results support MHRQI as a structure-aware representation in simulation, while highlighting two constraints for future work: stronger variance reduction and lower circuit depth for practical hardware execution.

Recommended next steps are:

- Improve denoising strength while preserving structural metrics.
- Reduce circuit complexity (including qudit-oriented variants).
- Validate compressed variants on available quantum hardware.
- Extend the hierarchy to additional operations (e.g., edge/region analysis) and higher-dimensional representations.

ACKNOWLEDGMENT

Acknowledgment omitted for double-blind review.

References

[1] UN News. Global diabetes epidemic reaches critical levels with 800 million cases | UN News, November 2024.

[2] Temesgen Gebeyehu Wondmeneh and Jemal Abdu Mohammed. Prevalence of diabetic retinopathy and its associated risk factors among adults in Ethiopia: a systematic review and meta-analysis. *Scientific Reports*, 14(1):28266, November 2024.

[3] Bin Zhou, Archie W Rayner, Edward W Gregg, et al. Worldwide trends in diabetes prevalence and treatment from 1990 to 2022: a pooled analysis of 1108 population-representative studies with 141 million participants. *The Lancet*, 404(10467):2077–2093, November 2024.

[4] Ahlem Aziz, Necmi Serkan Tezel, Seydi Kaçmaz, and Youcef Attallah. Early Diabetic Retinopathy Detection from OCT Images Using Multifractal Analysis and Multi-Layer Perceptron Classification. *Diagnostics*, 15(13):1616, June 2025.

[5] Yongqiang Huang, Zexin Lu, Zhimin Shao, Maosong Ran, Jiliu Zhou, Leyuan Fang, and Yi Zhang. Simultaneous denoising and super-resolution of optical coherence tomography images based on generative adversarial network. *Optics Express*, 27(9):12289, April 2019.

[6] Ashkan Abbasi and Amirhassan Monadjemi. Optical coherence tomography retinal image reconstruction via nonlocal weighted sparse representation. *Journal of Biomedical Optics*, 23(03):1, March 2018.

[7] Jing Mao, Lianming Sun, Jie Chen, and Shunyuan Yu. Overview of Research on Digital Image

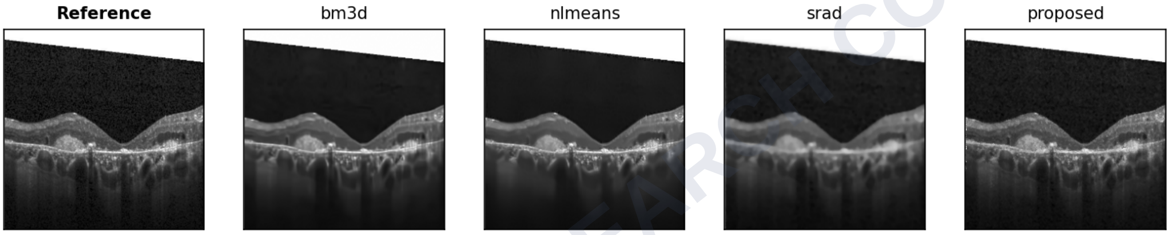

Fig.3: *Visual Denoising Comparison for a Representative CNV Sample*

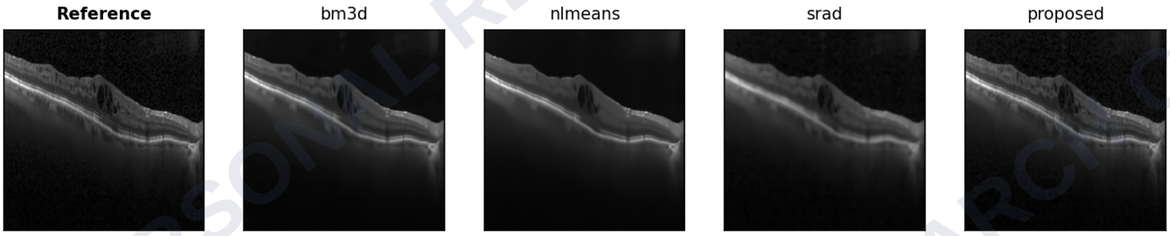

Fig.4: *Visual Denoising Comparison for a Representative DME Sample*

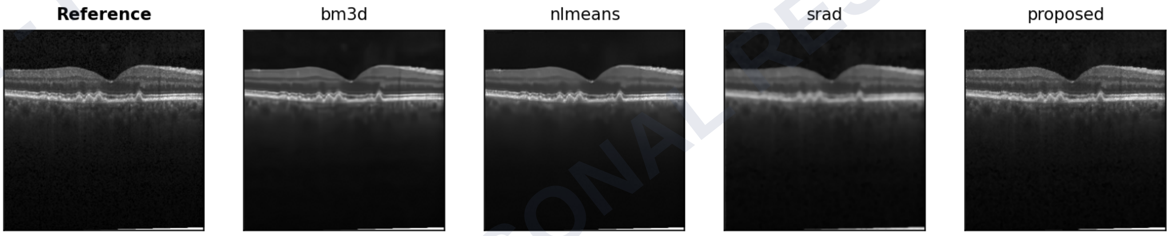

Fig.5: *Visual Denoising Comparison for a Representative Drusen Sample*

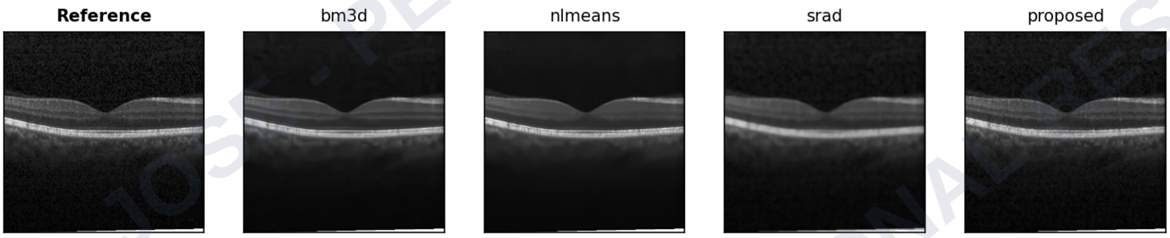

Fig.6: *Visual Denoising Comparison for a Representative Normal Sample*

Denoising Methods. *Sensors*, 25(8):2615, April 2025.

[8] Yongjian Yu and S.T. Acton. Speckle reducing anisotropic diffusion. *IEEE Transactions on Image Processing*, 11(11):1260–1270, November 2002.

[9] Shani Pineda, Kenneth Cabobos, Brient Rebulanan, Joshua Rodriguez, and Orland Tubola. Analyzing surface code as error correction method for quantum gate translation of a classical arithmetic logic unit (alu) in the presence of depolarizing noise. *ECTI Transactions on Computer and Information Technology (ECTI-CIT)*, 19:97–107, January 2025.

[10] Allyza Cadelina, Mori Cuevas, Marselean Kallos, Morris Montemayor, Ronald Uy, Joshua Rodriguez, and Orland Tubola. D-wave implementation of quantum annealing for optimal resource allocation in disaster response operation of marikina city. *ECTI Transactions on Computer and Information Technology (ECTI-CIT)*, 18:24–33, January 2024.

[11] Sylwia M. Kolenderska, Frédérique Vanholsbeeck, and Piotr Kolenderski. Fourier domain quantum optical coherence tomography. *Optics Express*, 28(20):29576, September 2020.

[12] Jonas Landman, Natansh Mathur, Yun Yvonna Li, Martin Strahm, Skander Kazdaghli, Anupam Prakash, and Iordanis Kerenidis. Quantum Methods for Neural Networks and Application to

Medical Image Classification. *Quantum*, 6:881, December 2022.

[13] Krzysztof A. Maliszewski, Magdalena A. Urbańska, Piotr Kolenderski, Varvara Vetrova, and Sylwia M. Kolenderska. Extracting Group Velocity Dispersion values using quantum-mimic Optical Coherence Tomography and Machine Learning. *Scientific Reports*, 13(1):6596, April 2023.

[14] Magued B. Nasr, Bahaa E. A. Saleh, Alexander V. Sergienko, and Malvin C. Teich. Demonstration of Dispersion-Canceled Quantum-Optical Coherence Tomography. *Physical Review Letters*, 91(8):083601, August 2003.

[15] Fangfang Li, Qizhou Wu, Bei Jia, and Zhicheng Yang. Speckle Noise Removal in OCT Images via Wavelet Transform and DnCNN. *Applied Sciences*, 15(12):6557, June 2025.

[16] Huaiguang Chen and Jing Gao. Non-Local Mean Denoising Algorithm Based on Fractional Compact Finite Difference Scheme Effectively Reduces Speckle Noise in Optical Coherence Tomography Images. *Micromachines*, 13(12):2039, November 2022.

[17] M. Maggioni, V. Katkovnik, K. Egiazarian, and A. Foi. Nonlocal Transform-Domain Filter for Volumetric Data Denoising and Reconstruction. *IEEE Transactions on Image Processing*, 22(1):119–133, January 2013.

[18] Parisa Ghaderi Daneshmand and Hossein Rabbani. Total variation regularized tensor ring decomposition for OCT image denoising and super-resolution. *Computers in Biology and Medicine*, 177:108591, July 2024.

[19] Shuhang Gu, Lei Zhang, Wangmeng Zuo, and Xiangchu Feng. Weighted Nuclear Norm Minimization with Application to Image Denoising. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2862–2869, Columbus, OH, USA, June 2014. IEEE.

[20] Kerry J. Halupka, Bhavna J. Antony, Matthew H. Lee, Katie A. Lucy, Ravneet S. Rai, Hiroshi Ishikawa, Gadi Wollstein, Joel S. Schuman, and Rahil Garnavi. Retinal optical coherence tomography image enhancement via deep learning. *Biomedical Optics Express*, 9(12):6205, December 2018.

[21] Nilesh A. Kande, Rupali Dakhane, Ambedkar Dukkipati, and Phaneendra Kumar Yalavarthy. SiameseGAN: A Generative Model for Denoising of Spectral Domain Optical Coherence Tomography Images. *IEEE Transactions on Medical Imaging*, 40(1):180–192, January 2021.

[22] Furkan Kirik, Farid Iskandarov, Kamile Melis Erturk, and Hakan Ozdemir. Quantitative analysis of deep learning-based denoising model efficacy on optical coherence tomography images with different noise levels. *Photodiagnosis and Photodynamic Therapy*, 45:103891, February 2024.

[23] Weijia Mao, Hao Chen, Zhenheng Yang, and Mike Zheng Shou. The Image as Its Own Reward: Reinforcement Learning with Adversarial Reward for Image Generation, November 2025. arXiv:2511.20256 [cs].

[24] Kyohei Hayama, Bo Cao, Ryo Okamoto, Shun Suezawa, Masayuki Okano, and Shigeki Takeuchi. High-depth-resolution imaging of dispersive samples using quantum optical coherence tomography. *Optics Letters*, 47(19):4949, October 2022.

[25] Sayantan Dutta, Adrian Basarab, Bertrand Georgeot, and Denis Kouamé. Quantum mechanics-based signal and image representation: application to denoising. 2020. Version Number: 3.

[26] Sayantan Dutta, Adrian Basarab, Bertrand Georgeot, and Denis Kouamé. A Novel Image Denoising Algorithm Using Concepts of Quantum Many-Body Theory. *Signal Processing*, 201:108690, December 2022.

[27] Amirreza Hashemi, Sayantan Dutta, Bertrand Georgeot, Denis Kouamé, and Hamid Sabet. Quantum inspired approach for denoising with application to medical imaging. *Research Square*, pages rs.3.rs–4600863, July 2024.

[28] Phuc Q. Le, Fangyan Dong, and Kaoru Hirota. A flexible representation of quantum images for polynomial preparation, image compression, and processing operations. *Quantum Information Processing*, 10(1):63–84, February 2011.

[29] Yi Zhang, Kai Lu, Yinghui Gao, and Mo Wang. NEQR: a novel enhanced quantum representation of digital images. *Quantum Information Processing*, 12(8):2833–2860, August 2013.

[30] Yingying Hu, Dayong Lu, Qianqian Zhang, and Meiyu Xu. Quantum image representations based on density matrices in open quantum systems. *EPJ Quantum Technology*, 11(1):28, December 2024.