Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A quantum color space for finite-shot RGB reconstruction in quantum image processing

Gang-Joon Yoon
National Institute for Mathematical Sciences
SeonMin Kim
Kookmin University
Byung-Gook Lee
Dongseo University
Jinjoo Song
Kookmin University
Sang Min Yoon
smyoon@kookmin.ac.kr

Kookmin University

Article

Keywords:

Posted Date: June 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9965533/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

A quantum color space for finite-shot RGB reconstruction in quantum image processing

Gang-Joon Yoon[1,+], Seonmin Kim[2,+], Byung-Gook Lee[3], Jinjoo Song[2], and Sang Min Yoon[2,*]

[1]National Institute for Mathematical Sciences, Daejeon, 34047, Korea
[2]Kookmin University, College of Computer Science, Seoul, 02707, Korea
[3]Dongseo University, Department of Computer Engineering, Busan, 47011, Korea
[*]Correspondence to: smyoon@kookmin.ac.kr
[+]These authors contributed equally to this work

ABSTRACT

Quantum image processing requires color representations that are compact in quantum registers and recoverable from finite-shot measurement statistics. Existing quantum image representations often emphasize storage efficiency, whereas the geometric organization of color information and its effect on reconstruction remain less explored. Here we propose a Bloch-sphere-based quantum color space that maps discrete RGB values to single-qubit state parameters (θ, ϕ). Colors are organized by maximum–minimum channel subgroups: each subgroup is assigned to a latitude band, and channel-permutation trajectories are encoded by the azimuthal angle. The construction also induces a canonical grayscale subspace. Finite-shot measurement counts are converted into angle estimates and decoded through a deterministic region-of-point/lookup-table (RoP/LUT) rule. In Qiskit Aer simulations, MNIST grayscale reconstruction achieves an average mean squared error (MSE) of 3.90 and peak signal-to-noise ratio (PSNR) of 42.44 dB. On CIFAR-10, the 64^3 RGB grid achieves the best reconstruction performance, with MSE 81.29, PSNR 29.84 dB, and structural similarity (SSIM) 0.9313. Shot-wise experiments show that pixel-wise validation benefits from increased measurement budgets, whereas position-register full-image encoding is limited by sparse per-position measurement counts.

Introduction

Quantum image processing (QIP) aims to represent, manipulate, and recover visual information using quantum states. Since images contain structured spatial and chromatic information, quantum image representation (QIR) plays a central role in determining how visual data can be encoded into quantum registers and how reliably it can be recovered from quantum measurements. While many QIR schemes have been designed to reduce qubit requirements or support elementary image operations, color representation remains a particularly challenging problem. A color image is not merely a collection of independent channel values; its red–green–blue (RGB) coordinates form a structured color space in which intensity, chromatic variation, and channel relationships are jointly organized. A quantum representation of color should therefore be compact, but it should also preserve an interpretable relationship with classical color geometry and support stable inverse reconstruction from finite-shot measurement statistics.

Early QIR models established several important encoding principles. Flexible representation of quantum images (FRQI) encodes grayscale intensity through rotation angles while using position qubits to index image coordinates, providing a compact angle-based representation[1]. This idea was extended to color imagery in multi-channel quantum image representation (MCQI), where RGB channels are represented by additional channel information and channel-wise controlled rotations[2]. These angle-based schemes are qubit-efficient and conceptually close to quantum state preparation, but image recovery is inherently statistical because pixel values must be inferred from measurement probabilities. As a result, reconstruction quality depends on the number of shots allocated to each encoded value, and small probability fluctuations can lead to visible reconstruction errors, especially when the encoded color space becomes dense.

A different line of work represents pixel values explicitly in computational-basis registers. Novel enhanced quantum representation (NEQR) stores grayscale values in basis states, enabling more direct value readout in principle, but requiring additional value qubits whose number increases with image precision[3]. For color images, optimized quantum representation for color digital images (OCQR) integrates color value, channel index, and pixel position into a compact register structure[4], while quantum representation model of color digital images (QRCI) encodes color information using qubit sequences and supports several channel- and bit-plane-related operations[5]. Hybrid qudit quantum representation (HQDQR) further broadens the design space by using qudits and hybrid qubit–qutrit structures for RGB images[6]. Recent studies have also explored qutrit-

or color-model-based representations, such as the Ternary Novel Colored Quantum Representation (TNCQR) and the adjacency Fourier quantum image representation of HSI (AFQIRHSI)[7,8].

These representations improve storage, indexing, retrieval, or operation design, but they primarily treat color as a value to be stored or manipulated in a register structure. They provide limited guidance on how the geometry of classical RGB coordinates should be organized in quantum state space, or how finite-shot perturbations affect color-space decision boundaries. This distinction is important because an encoded quantum state is not directly observable; it must be estimated from repeated measurements. For color-aware QIP, a representation should therefore make measurement dependence explicit by defining how RGB values are arranged in quantum state space, how measurement-derived estimates are decoded back to classical colors, and how the density of encoded color points controls reconstruction stability.

Here we propose a Bloch-sphere-based quantum color space for measurement-decodable RGB representation. The proposed framework defines a structure-preserving correspondence between the discrete RGB cube Σ and single-qubit state parameters on the Bloch sphere. Each RGB triplet is mapped to spherical coordinates (θ, ϕ), where the polar angle θ organizes colors according to maximum–minimum channel subgroups and the azimuthal angle ϕ parameterizes channel-permutation trajectories within each latitude band. In this construction, colors sharing the same maximum and minimum channel values lie on a common latitude circle, while chromatic transitions are represented as ordered trajectories along the azimuthal direction. The grayscale diagonal $\{(x,x,x)\}$ is naturally inherited as a one-dimensional subspace, in which intensity is encoded only by the polar angle and the azimuthal coordinate becomes redundant. The proposed color space is paired with a measurement-driven inverse reconstruction procedure. Instead of assuming exact analytic inversion from ideal quantum states, finite-shot measurement counts are converted into angle estimates $(\hat{\theta}, \hat{\phi})$, and these estimates are decoded into RGB values through a deterministic region-of-point (RoP) and lookup-table (LUT) rule. This explicitly connects color-space design with measurement stability. In particular, the density of RGB points on the Bloch sphere determines the angular decision margins available for inverse decoding. A fine RGB grid provides high nominal color resolution but creates narrow angular regions that are sensitive to finite-shot angle-estimation error. Conversely, coarser quantized RGB grids enlarge inter-color angular separability and can improve reconstruction robustness. We validate the proposed quantum color space using Qiskit-based simulations on grayscale and color image datasets. MNIST is used to examine the induced grayscale subspace, while CIFAR-10 is used to evaluate RGB reconstruction under quantized color spaces. By comparing 64^3, 128^3, and 256^3 RGB grids, we show that reconstruction quality is governed by a quantization–separability trade-off: coarser grids provide larger angular margins and more stable RoP/LUT decoding, whereas the full 256^3 RGB space is more vulnerable to boundary crossing under finite-shot perturbations. We further compare the pixel-wise validation mode and the position-register full-image mode to analyze how measurement allocation affects reconstruction. These results demonstrate that quantum color representations should be designed not only for compact storage but also for stable and interpretable recovery from finite-shot measurements.

Results

Reconstruction performance on grayscale images

We first evaluate the induced grayscale subspace using MNIST. In the proposed quantum color space, grayscale values lie on the diagonal set (x,x,x), where the azimuthal coordinate is redundant and only the polar angle θ varies. This reduces grayscale reconstruction to a one-dimensional region-of-point (RoP) decoding problem. We therefore use a position-register full-image circuit for MNIST, in which position qubits index all pixel locations and each pixel intensity is recovered by conditioning measurement statistics on the corresponding position state.

Figure 1 shows representative MNIST reconstruction results. The reconstructed digits preserve the overall stroke topology, digit shape, and thin foreground structures, with only minor intensity deviations. Table 1 reports digit-wise reconstruction performance over 100 samples per digit. The average mean squared error (MSE) is 3.90, and the average peak signal-to-noise ratio (PSNR) is 42.44 dB. The lowest MSE is observed for digit "1" (2.25), whereas digit "0" gives the highest MSE (4.80), which is consistent with its larger foreground coverage and the accumulation of small intensity deviations over more pixels.



Figure 1. Representative grayscale reconstruction results on MNIST. The proposed quantum color space reduces grayscale recovery to a one-dimensional polar-angle decoding problem.

These results indicate that the grayscale diagonal is a stable one-dimensional subspace of the proposed RGB construction. They also motivate the RGB reconstruction experiments below, where both θ and ϕ must be estimated for each pixel and the

Table 1. Grayscale reconstruction performance on MNIST.

Digits	0	1	2	3	4	5	6	7	8	9	Avg
MSE(↓)	4.80	2.25	4.32	4.37	3.56	3.90	3.96	3.50	4.56	3.78	3.90
PSNR(↑)	41.47	44.83	41.90	41.82	42.72	42.40	42.29	42.82	41.68	42.50	42.44

density of color points on the Bloch sphere becomes a more critical factor.

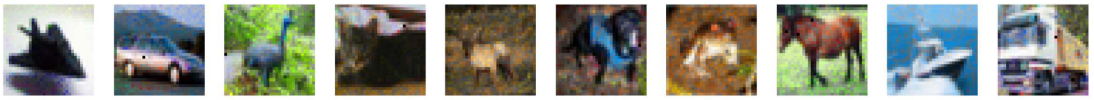
RGB reconstruction under quantized color spaces

We then evaluate full RGB reconstruction on CIFAR-10. Unlike grayscale reconstruction, RGB recovery requires estimating both the polar angle θ and the azimuthal angle ϕ for each pixel. This makes RGB reconstruction more sensitive to finite-shot angle-estimation error and to boundary crossing between neighboring region-of-point (RoP) decision regions. To isolate this intrinsic decoding behavior, we use pixel-wise execution for the main CIFAR-10 reconstruction experiments.

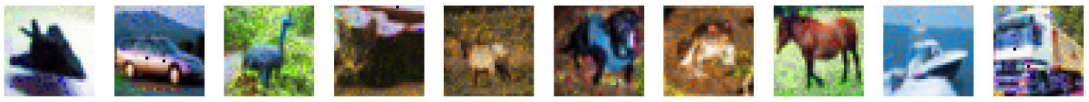
We compare three RGB grids: 64^3, 128^3, and 256^3. Figure 2 shows representative reconstruction results. The 128^3 grid preserves coarse object layout and class-dependent silhouettes, but visible hue and tone drift remain because neighboring color points are relatively close on the Bloch sphere.



(a) 64^3 RGB grid



(b) 128^3 RGB grid



(c) 256^3 RGB grid

Figure 2. Representative CIFAR-10 RGB reconstruction results under the three quantized RGB grids.

The 64^3 grid shows smaller chromatic deviations and more consistent reconstruction of large uniform regions than the denser grids. In contrast, the full 256^3 grid shows the largest degradation under finite-shot perturbations. Although it provides the highest nominal color resolution, its angular RoP regions are narrow, so small residual errors in $(\hat{\theta}, \hat{\phi})$ can cross decision boundaries and lead to color mis-decoding. In a few cases, isolated black null pixels appear due to rare RoP/LUT decoding exceptions under finite-shot perturbation and discretization; all results are reported without post-processing.

Class-wise structural similarity (SSIM) further shows how image content affects reconstruction quality. As shown in Table 2, classes with smoother object regions and more rigid structures, such as Automobile and Truck, tend to achieve higher SSIM values, whereas texture-rich classes with cluttered backgrounds, such as Bird and Deer, are more sensitive to chromatic and local reconstruction errors. This tendency is consistent across the three RGB grids, but the overall robustness improves substantially in the 64^3 setting. For example, Automobile and Truck achieve SSIM values of 0.9642 and 0.9567 in the 64^3 grid, compared with 0.9100 and 0.9033 in the 128^3 grid. In the 256^3 grid, all classes degrade because the RoP margins become too narrow for stable finite-shot decoding.

Table 2. Class-wise SSIM on CIFAR-10 under different RGB grids.

Grid	Airplane	Auto.	Bird	Cat	Deer	Dog	Frog	Horse	Ship	Truck	Avg.
64^3	0.9128	0.9642	0.9036	0.9356	0.9071	0.9379	0.9326	0.9476	0.9146	0.9567	0.9313
128^3	0.8284	0.9100	0.8145	0.8609	0.8110	0.8714	0.8557	0.8805	0.8309	0.9033	0.8567
256^3	0.7121	0.8144	0.6887	0.7442	0.6762	0.7640	0.7341	0.7803	0.7162	0.8044	0.7435

The average reconstruction performance is summarized in Table 3. The 64^3 grid achieves the best performance, with an average MSE of 81.29, PSNR of 29.84 dB, and SSIM of 0.9313. Compared with the 128^3 grid, the average MSE decreases from 217.90 to 81.29, corresponding to a 62.7% reduction, while PSNR increases from 25.14 dB to 29.84 dB and SSIM improves from 0.8567 to 0.9313. The full 256^3 grid yields the lowest reconstruction quality, with an average MSE of 609.77, PSNR of 20.55 dB, and SSIM of 0.7435.

Table 3. Average CIFAR-10 reconstruction performance under different RGB grids.

RGB grid	MSE ↓	PSNR ↑	SSIM ↑
64^3	81.29	29.84	0.9313
128^3	217.90	25.14	0.8567
256^3	609.77	20.55	0.7435

These results support the quantization–separability interpretation. Higher classical color resolution does not automatically lead to better quantum reconstruction, because denser RGB grids induce narrower angular decision regions on the Bloch sphere. Instead, reconstruction quality is governed by the balance between nominal color resolution and finite-shot angular separability. Full class-wise MSE, PSNR, and SSIM results are provided in Supplementary Table S3.

Shot-wise robustness in pixel-wise execution

To analyze finite-shot effects more directly, we vary the shot budget from 2^9 to 2^{18} and evaluate CIFAR-10 reconstruction in the pixel-wise validation mode. In this setting, each pixel circuit receives an independent shot budget, so increasing the number of shots directly improves the accuracy of the estimated angles. We use the 128^3 RGB grid as a representative intermediate-resolution setting in the main text, while the corresponding shot-wise results for the 64^3 and 256^3 grids are provided in the Supplementary Information.

Figure 3 shows the shot-wise behavior for the 128^3 RGB grid. As the shot budget increases, the angular estimation errors decrease and the image-level reconstruction metrics improve. This confirms that finite-shot sampling is a major source of reconstruction error in the proposed measurement-driven decoding pipeline. However, the improvement is not determined by shot count alone: the estimated angles must also remain within the correct region-of-point (RoP) decision region to recover the intended RGB value.

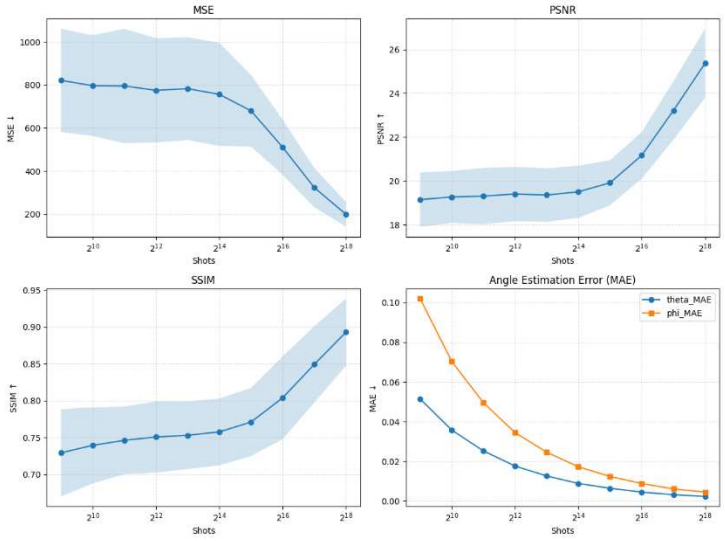


Figure 3. Shot-wise pixel-wise reconstruction under the 128^3 RGB grid. The four panels show MSE, PSNR, SSIM, and angular estimation error as the shot budget increases from 2^9 to 2^{18}.

The comparison with the other RGB grids further supports the quantization–separability interpretation. At 2^{18} shots, the three grids have comparable residual angular errors: θ-MAE is approximately 0.00224, 0.00220, and 0.00220 for 64^3, 128^3, and 256^3, respectively, while ϕ-MAE is approximately 0.00411, 0.00437, and 0.00439. Nevertheless, their image-level metrics remain substantially different. At the same shot budget, MSE increases from 64.96 to 198.69 and 504.29, PSNR decreases

from 30.38 dB to 25.37 dB and 21.19 dB, and SSIM decreases from 0.9560 to 0.8930 and 0.8047 for 64^3, 128^3, and 256^3, respectively.

These results show that the magnitude of the angle-estimation error is not the only determining factor. A residual angular perturbation that remains inside the correct RoP region in the 64^3 grid can cross a decision boundary in the denser 128^3 or 256^3 grids. Thus, increasing the shot budget improves finite-shot angle estimation, but robust RGB reconstruction also requires sufficiently wide angular decision regions. Full shot-wise curves are provided in Supplementary Fig. S4.

Measurement-allocation bottleneck in position-register full-image encoding

We finally examine reconstruction in the position-register full-image encoding mode, where all pixels are encoded in a single quantum circuit using position qubits and the total shot budget is shared across all pixel locations and both angle branches. This mode is more compact at the representation level than the pixel-wise validation mode, but it substantially changes the measurement-allocation pattern. Instead of assigning the full shot budget to each pixel independently, the circuit distributes the total number of shots across all position states and angle branches.

Figure 4 visualizes representative reconstruction results for the position-register full-image encoding under the 256^3 RGB grid. The recovered images show substantial degradation, including weakened structural preservation, chromatic distortion, and reduced local detail. This qualitative degradation illustrates the practical limitation of position-register full-image encoding under finite-shot measurements: although the representation is compact, the per-position measurement counts available for each pixel are too sparse to support reliable recovery of both θ and ϕ.

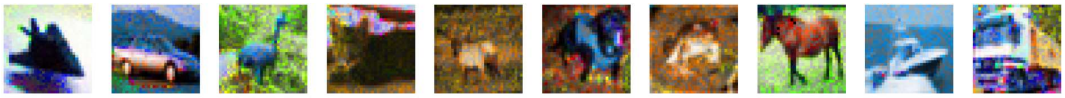


Figure 4. Representative CIFAR-10 reconstruction results in the position-register full-image encoding under the 256^3 RGB grid. The shared shot budget leads to sparse per-position measurement counts, resulting in visibly degraded reconstruction quality.

The quantitative results at 2^{18} total shots are summarized in Table 4. Reconstruction quality remains uniformly low across the three RGB grids: the MSE values are 896.88, 880.82, and 897.38 for 64^3, 128^3, and 256^3, respectively; the PSNR values are 18.70 dB, 18.78 dB, and 18.73 dB; and the SSIM values are 0.6960, 0.6998, and 0.6999. The corresponding angle-estimation errors are also much larger than in the pixel-wise mode, with θ-MAE around 0.072–0.073 and ϕ-MAE around 0.141–0.147. The full angle-error summary is provided in Supplementary Table S4.

Table 4. Average CIFAR-10 reconstruction performance in the position-register full-image encoding at 2^{18} total shots.

RGB grid	MSE ↓	PSNR ↑	SSIM ↑
64^3	896.88	18.70	0.6960
128^3	880.82	18.78	0.6998
256^3	897.38	18.73	0.6999

This behavior is markedly different from the pixel-wise setting. For a 32×32 image, the full-image circuit distributes the total shot budget across 1024 pixel positions and two angle branches. Therefore, the effective number of conditional samples per pixel and per branch scales approximately as shots$/(2HW)$. Even at 2^{18} total shots, this corresponds to only about 128 samples per pixel and branch on average. Consequently, the recovered angle estimates remain much noisier than in the pixel-wise setting, where each pixel receives an independent shot budget.

Taken together, the shot-wise experiments reveal two complementary constraints for quantum color reconstruction. The pixel-wise validation results show that increasing the shot budget can make finite-shot angle estimates sufficiently accurate, but robust recovery still requires sufficiently wide RoP decision regions. The position-register full-image results show that compact image-level encoding introduces a measurement-allocation bottleneck, because many pixel estimates must share a finite number of shots. Therefore, the proposed quantum color space should be interpreted not only in terms of qubit compactness, but also in terms of measurement-stable recovery under realistic shot budgets.

Discussion

This study introduced a Bloch-sphere-based quantum color space for measurement-driven RGB representation and reconstruction in quantum image processing. Rather than treating color as a value stored in a register, the proposed framework organizes

RGB triplets as geometrically interpretable points on the Bloch sphere. Colors sharing the same maximum–minimum subgroup descriptor (M_x, M_n) are assigned to common latitude bands, while channel-permutation trajectories are represented through the azimuthal angle ϕ. This construction links the structure of the classical RGB cube to single-qubit state parameters (θ, ϕ) and provides a deterministic route for inverse reconstruction through region-of-point (RoP) and lookup-table (LUT) decoding.

The experimental results show that reconstruction robustness is governed by the angular separability of the encoded color points. The induced grayscale subspace yields stable reconstruction on MNIST because the chromatic degree of freedom is removed and recovery reduces to a one-dimensional θ-partition. For RGB reconstruction on CIFAR-10, the quantized 64^3 color space achieves the strongest performance among the tested settings, whereas the denser 128^3 and 256^3 grids show progressively larger chromatic deviations. This trend indicates that higher nominal color resolution does not necessarily improve quantum reconstruction when finite-shot angle estimates can cross narrow RoP boundaries.

The shot-wise analysis further highlights the importance of measurement allocation. In the pixel-wise validation mode, increasing the shot budget reduces θ- and ϕ-estimation errors and improves RoP/LUT decoding. However, for a fixed residual angular error, coarser color grids remain more robust because their angular decision regions are wider. In contrast, the position-register full-image encoding exposes a per-position sampling bottleneck: the total shot budget is shared across all pixel locations and angle branches, leaving each pixel with sparse conditional measurements. This result suggests a practical trade-off between compact image-level encoding and reliable per-pixel color recovery.

The present study is limited by its controlled simulation setting and deterministic decoding strategy. All experiments are based on Qiskit Aer simulations without an explicit hardware noise model, and therefore do not fully capture gate noise, readout error, crosstalk, or calibration drift. In addition, the pixel-wise protocol is useful for analyzing intrinsic recoverability but requires many circuit executions. Future work will investigate likelihood-based decoding, confidence-aware RoP assignment, adaptive color-space partitioning, perceptually motivated color mappings, and validation on noisy simulators and real quantum devices. More broadly, the results suggest that quantum image representations should be designed not only for compact encoding, but also for stable and interpretable recovery from finite-shot measurements.

Methods

Bloch-sphere parameterization and quantum color-space construction

The proposed quantum color space is defined using the Bloch-sphere parameterization of a single-qubit pure state. Given the computational basis $|0\rangle, |1\rangle$, a qubit state can be written as

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\phi}\sin\left(\frac{\theta}{2}\right)|1\rangle, \tag{1}$$

where $0 \le \theta \le \pi$ is the polar angle and $0 \le \phi < 2\pi$ is the azimuthal angle. The corresponding point on the unit sphere is

$$(x, y, z) = (\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta). \tag{2}$$

We use this spherical coordinate system as the geometric domain for representing RGB colors. Rather than treating color values as independent labels or basis-register contents, the proposed construction organizes RGB values as finite points on the Bloch sphere so that finite-shot measurement estimates can be decoded through angular decision regions.

Let the full 8-bit RGB color space be

$$\Sigma = \{(r, g, b) \mid r, g, b \in \{0, \ldots, 255\}\}. \tag{3}$$

To analyze how color-space density affects reconstruction robustness, we consider three quantized RGB grids,

$$\Sigma_L = \{0, \ldots, L-1\}^3, \qquad L \in \{64, 128, 256\}. \tag{4}$$

For a display-scale channel value $v \in \{0, \ldots, 255\}$, the corresponding L-level channel value is

$$v_L = \left\lfloor \frac{v}{s} \right\rfloor, \qquad s = \frac{256}{L}. \tag{5}$$

After reconstruction, the decoded L-level value $\hat{v}_L$ is mapped back to the display scale by

$$\hat{v} = s\hat{v}_L. \tag{6}$$

Thus, the 128^3 grid corresponds to display values $\{0, 2, 4, \ldots, 254\}^3$, while the 64^3 grid corresponds to $\{0, 4, 8, \ldots, 252\}^3$. These grids are quantized display-space RGB grids, not raw intensity ranges $\{0, \ldots, 127\}^3$ or $\{0, \ldots, 63\}^3$.

For each RGB triplet $(r,g,b) \in \Sigma_L$, we define

$$M_x = \max\{r,g,b\}, \qquad M_n = \min\{r,g,b\}. \tag{7}$$

The proposed mapping

$$\Phi_L : \Sigma_L \to S^2 \tag{8}$$

assigns each RGB triplet to a Bloch-sphere coordinate (θ,ϕ) according to its max–min structure. The polar angle θ is determined by the pair (M_x,M_n), so that all colors with the same maximum and minimum channel values are assigned to a common latitude band. Within each latitude band, the azimuthal angle ϕ orders channel-permutation trajectories. For non-grayscale colors with $M_x > M_n$, one channel is fixed at M_x, one channel is fixed at M_n, and the remaining channel varies between them, yielding six cyclic transition segments around the latitude circle.

The polar-angle assignment is designed to distribute RGB points over the sphere while maintaining symmetry around the equator. Colors with lower max–min sums are placed in the upper hemisphere, colors with higher max–min sums are placed in the lower hemisphere, and paired families are arranged symmetrically. The complete case-wise angular assignments for the 256^3, 128^3, and 64^3 grids are provided in the Supplementary Information.

The grayscale subspace is the diagonal set

$$\mathscr{G}_L = \{(x,x,x) \mid x \in \{0,\dots,L-1\}\}. \tag{9}$$

For grayscale values, $M_x = M_n = x$, and there is no chromatic degree of freedom. We therefore use the canonical representative $\phi = 0$ for deterministic encoding and decoding. Under the full 8-bit grayscale case, black is mapped to the north pole, white is mapped to the south pole, and intermediate grayscale intensities are ordered monotonically along the polar direction. This restriction reduces grayscale reconstruction to a one-dimensional decision problem in θ.

Region-of-point decoding

Finite-shot measurements yield perturbed angle estimates rather than exact angle values. We therefore decode colors using a region-of-point (RoP) and lookup-table (LUT) rule. For each RGB value $(r,g,b) \in \Sigma_L$, we precompute a decision region $I_{(r,g,b)}$ around its encoded angle pair. Given a measurement-derived estimate $(\hat{\theta},\hat{\phi})$, the reconstructed RGB value is obtained by

$$(\hat{r}_L,\hat{g}_L,\hat{b}_L) = \mathrm{LUT}(\hat{\theta},\hat{\phi}) \quad \text{s.t.} \quad (\hat{\theta},\hat{\phi}) \in I_{(\hat{r}_L,\hat{g}_L,\hat{b}_L)}. \tag{10}$$

The recovered L-level values are then mapped back to display-scale RGB values using the quantization rule described above.

When evaluating RoP membership in the azimuthal coordinate, the 2π-periodicity of ϕ is explicitly handled so that estimates near the wrap-around boundary are assigned correctly. For grayscale inputs, the RoP partition reduces to a one-dimensional interval partition over $\theta \in [0,\pi]$, and the decoder ignores azimuthal variation. The complete RoP boundary definitions for all tested RGB grids are given in the Supplementary Information.

Quantum encoding protocols and simulation setup

We evaluate two quantum encoding protocols. The first is a pixel-wise validation mode, in which a two-qubit angle-estimation circuit is executed independently for each pixel. The circuit consists of an angle qubit q_1 and a flag qubit q_2. The flag qubit selects which angle is encoded: the $q_2 = 1$ branch encodes θ, and the $q_2 = 0$ branch encodes ϕ. After preparing the flag qubit in superposition, controlled R_y rotations are applied to the angle qubit, and the complementary control branch is implemented using X gates.

After measuring both qubits in the computational basis, let

$$p_{ij} = \Pr(|q_1 q_2\rangle = |ij\rangle) \tag{11}$$

denote the empirical outcome probabilities. The conditional probabilities for the two angle branches are

$$p_{\theta,0} = \Pr(q_1 = 0 \mid q_2 = 1) = \frac{p_{01}}{p_{01}+p_{11}}, \tag{12}$$

and

$$p_{\phi,0} = \Pr(q_1 = 0 \mid q_2 = 0) = \frac{p_{00}}{p_{00}+p_{10}}. \tag{13}$$

The corresponding finite-shot angle estimates are

$$\hat{\theta} = 2\arccos\left(\sqrt{p_{\theta,0}}\right), \qquad \hat{\phi} = 4\arccos\left(\sqrt{p_{\phi,0}}\right). \tag{14}$$

The factor 4 in the ϕ estimator follows from the half-angle embedding used in the ϕ branch, where the circuit applies $R_y(\phi/2)$ and therefore $p_{\phi,0} = \cos^2(\phi/4)$.

In the pixel-wise validation mode, this two-qubit circuit is executed independently for each pixel. For an image of size $H \times W$, the number of circuit executions is HW. This mode is used for the main CIFAR-10 RGB reconstruction experiments because it isolates the intrinsic finite-shot recoverability of the proposed RGB-to-sphere correspondence without position-conditioned sampling sparsity.

The second protocol is a position-register full-image encoding mode, in which all pixels of a $2^n \times 2^n$ image are encoded within a single quantum circuit using a position register. The position register, equivalently an address register in circuit implementation, consists of $2n$ qubits and indexes the two-dimensional pixel coordinates. Angle information is recovered by conditioning measurement counts on the observed position bitstring.

For grayscale images, only the polar angle θ is required; therefore, a 32×32 grayscale image can be represented using 10 position qubits and one angle qubit, resulting in an 11-qubit circuit. For RGB images, both θ and ϕ must be estimated for each pixel. The position-register full-image RGB formulation therefore uses the $2n$ position qubits together with the angle and flag qubits, resulting in $2n+2$ qubits for a $2^n \times 2^n$ image.

The position-register full-image mode is compact in terms of qubit count, but the measurement budget is shared across all pixel positions and angle branches. For an $H \times W$ RGB image, the effective number of conditional samples per pixel and per branch scales approximately as $shots/2HW$.

We use this mode to analyze the measurement-allocation bottleneck that arises when many pixel estimates share a finite number of shots.

All experiments are performed using Qiskit Aer simulations of the encode–measure–decode pipeline. Unless otherwise stated, we use the default `AerSimulator` backend without an additional hardware noise model. The simulations therefore isolate finite-shot statistical effects rather than gate noise, readout error, crosstalk, or device calibration drift. MNIST is used to evaluate reconstruction on the induced grayscale subspace, and CIFAR-10 is used to evaluate RGB reconstruction under the 64^3, 128^3, and 256^3 grids. For the main class-wise reconstruction experiments, we use 100 randomly sampled images per class. For shot-wise robustness analysis, the shot budget is varied over $2^9, 2^{10}, \ldots, 2^{18}$. In the pixel-wise validation mode, the reported shot count is allocated independently to each pixel circuit. In the position-register full-image mode, the reported shot count is the total number of shots allocated to the entire image circuit.

Evaluation metrics

Reconstruction quality is evaluated using mean squared error (MSE), peak signal-to-noise ratio (PSNR), and structural similarity (SSIM). MSE and PSNR are reported for both grayscale and RGB experiments, while SSIM is reported for RGB reconstruction. For shot-wise analysis, we additionally compute the mean absolute error of the recovered angles. The polar-angle error is measured as $|\hat{\theta} - \theta|$, and the azimuthal error is measured using the circular angular distance

$$d_\phi(\hat{\phi}, \phi) = \min\left(|\hat{\phi} - \phi|, 2\pi - |\hat{\phi} - \phi|\right). \tag{15}$$

The Supplementary Information provides the complete angular assignment rules, RoP boundary definitions, implementation pseudocode, additional reconstruction examples, class-wise CIFAR-10 tables, and detailed shot-wise numerical values.

Data availability

MNIST and CIFAR-10 are publicly available benchmark datasets. The data generated during the current study, including processed reconstruction results and lookup-table outputs, are available from the corresponding author upon reasonable request.

References

1. P. Q. Le, A flexible representation of quantum images for polynomial preparation, image compression, and processing operations, *Quantum Information Processing* **10**, 63–84 (2011).
2. B. Sun, A. Iliyasu, F. Yan, F. Dong, and K. Hirota, An RGB multi-channel representation for images on quantum computers, *Journal of Advanced Computational Intelligence and Intelligent Informatics* **17**(3), 404–417 (2013).
3. Y. Zhang, NEQR: A novel enhanced quantum representation of digital images, *Quantum Information Processing* **12**, 2833–2860 (2013).

4. K. Liu, Y. Zhang, K. Lu, X. Wang, and X. Wang, An optimized quantum representation for color digital images, *International Journal of Theoretical Physics* **57**, 2938–2948 (2018).
5. L. Wang, Q.-W. Ran, J. Ma, S. Yu, and L. Tan, QRCI: A new quantum representation model of color digital images, *Optics Communications* **438**, 147–158 (2019).
6. S. Das and F. Caruso, A hybrid-qudit representation of digital RGB images, *Scientific Reports* **13**, 13671 (2023).
7. M. Rofail, R. Montaser, and A. Younes, A novel qutrit representation for RGB digital images, *Scientific Reports* **15**, 43256 (2025).
8. N. A. Alwan, S. J. Obaiys, N. M. G. Al-Saidi, N. F. B. Mohd Noor, and Y. Karaca, Multilayered quantum computing and simulation system for enhanced image representation of HSI based Fourier transform and adjacency matrix, *Scientific Reports* **15**, 40286 (2025).

Funding

G.-J. Yoon was supported by the National Institute for Mathematical Sciences (NIMS) grant funded by the Korean government (No. NIMS-B26820000) and by the Ministry of Trade, Industry and Energy (MOTIE), Korea, under the Virtual Engineering Service Platform program (P0022335) supervised by the Korea Institute for Advancement of Technology (KIAT). S. M. Yoon was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (No. 2020-0-00457).

Author contributions statement

G.-J.Y. and S.K. contributed equally to this work. G.-J.Y. and S.K. conceived the study and developed the quantum color-space formulation. S.K. implemented the simulations and analyzed the experimental results. B.-G.L. and J.S. contributed to the interpretation of quantum image representation and reconstruction results. S.M.Y. supervised the project. All authors reviewed and approved the manuscript.

Additional information

Competing interests: The authors declare no competing interests.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryInformation.pdf