

# Supplementary Information

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

This Supplementary Information provides mathematical and experimental details omitted from the main text for concision. It includes the RGB subgroup construction, angular assignment rules for the  $256^3$ ,  $128^3$ , and  $64^3$  RGB grids, region-of-point (RoP) and lookup-table (LUT) decoding, quantum circuit execution protocols, implementation pseudocode, extended qualitative visualizations, class-wise reconstruction tables, and shot-wise robustness results.

### Supplementary Note 1: RGB subgroup structure and quantized RGB grids

#### RGB subgroup decomposition

Let

$$\Sigma = \{(r, g, b) \mid r, g, b \in \{0, \dots, 255\}\} \quad (\text{S1})$$

denote the full 8-bit display-scale RGB color space. For an RGB triplet  $(r, g, b) \in \Sigma$ , define

$$M_x = \max\{r, g, b\}, \quad M_n = \min\{r, g, b\}. \quad (\text{S2})$$

For fixed integers  $0 \leq M_n \leq M_x \leq 255$ , we define the max–min subset

$$(M_x : M_n) = \{(r, g, b) \in \Sigma : M_x = \max(r, g, b), M_n = \min(r, g, b)\}. \quad (\text{S3})$$

We refer to these sets as subgroups for convenience, although they are not algebraic subgroups.

When  $M_x > M_n$ , the set  $(M_x : M_n)$  contains  $6(M_x - M_n)$  RGB values. With  $d = M_x - M_n$ , it can be written compactly as

$$(M_x : M_n) = \bigcup_{k=0}^{d-1} \{(M_x, M_n + k, M_n), (M_x - k, M_x, M_n), \\ (M_n, M_x, M_n + k), (M_n, M_x - k, M_x), \\ (M_n + k, M_n, M_x), (M_x, M_n, M_x - k)\}. \quad (\text{S4})$$

When  $M_x = M_n$ , the set reduces to the single grayscale element  $(M_x, M_x, M_x)$ .

For  $\ell = 0, 1, \dots, 510$ , we define the latitude-indexed RGB subset

$$\Sigma_\ell = \{(r, g, b) \in \Sigma : \max(r, g, b) + \min(r, g, b) = \ell\} = \bigcup_{M_x + M_n = \ell} (M_x : M_n). \quad (\text{S5})$$

The proposed quantum color space assigns RGB values with the same max–min descriptor  $(M_x, M_n)$  to a common latitude circle on the Bloch sphere. The value  $M_x + M_n$  determines the broad polar band, while the within-band max–min pair determines the precise polar angle.

#### Quantized RGB grids

For  $L \in \{64, 128, 256\}$ , the  $L$ -level RGB grid is

$$\Sigma_L = \{0, \dots, L - 1\}^3. \quad (\text{S6})$$

An 8-bit display-scale channel value  $v \in \{0, \dots, 255\}$  is mapped to an  $L$ -level channel value

$$v_L = \left\lfloor \frac{v}{s} \right\rfloor, \quad s = \frac{256}{L}. \quad (\text{S7})$$

After RoP/LUT decoding, the recovered value  $\hat{v}_L$  is mapped back to the display scale by

$$\hat{v} = s\hat{v}_L. \quad (\text{S8})$$

Thus, the  $128^3$  grid corresponds to display values  $\{0, 2, 4, \dots, 254\}^3$ , while the  $64^3$  grid corresponds to  $\{0, 4, 8, \dots, 252\}^3$ . These are quantized display-space RGB grids rather than raw intensity ranges  $\{0, \dots, 127\}^3$  or  $\{0, \dots, 63\}^3$ .

## Grayscale restriction

The grayscale restriction is the diagonal subset

$$\mathcal{G}_L = \{(x, x, x) \mid x \in \{0, \dots, L-1\}\}. \quad (\text{S9})$$

For grayscale values,  $M_x = M_n = x$ , and the azimuthal coordinate carries no chromatic information. In the primary map, we use the canonical representative  $\phi_x = 0$  for deterministic encoding and LUT construction. For the full 8-bit grayscale case, the polar angle is

$$\theta_x = \begin{cases} 0, & x = 0, \\ x \frac{\pi}{255} - \frac{\pi}{1020}, & x = 1, \dots, 127, \\ x \frac{\pi}{255} + \frac{\pi}{1020}, & x = 128, \dots, 254, \\ \pi, & x = 255, \end{cases} \quad \phi_x = 0. \quad (\text{S10})$$

The alternative anchor longitudes  $\{0, 2\pi/3, 4\pi/3\}$  may be interpreted as symmetric representatives of the same grayscale value, but the primary LUT uses  $\phi_x = 0$ .

## Supplementary Note 2: Bloch-sphere angular assignment

A point on the unit sphere is parameterized by

$$(x, y, z) = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta), \quad 0 \leq \theta \leq \pi, \quad 0 \leq \phi < 2\pi. \quad (\text{S11})$$

The proposed correspondence maps each RGB triplet to a spherical coordinate  $(\theta, \phi)$  and then to the corresponding point on the Bloch sphere.

### Azimuthal assignment

For  $M_x > M_n$ , let  $d = M_x - M_n$  and

$$\phi_k = \frac{k\pi}{3d}, \quad k = 0, 1, \dots, d-1. \quad (\text{S12})$$

The six cyclic RGB transition segments are assigned as

$$\phi(M_x, M_n + k, M_n) = \phi_k, \quad (\text{S13})$$

$$\phi(M_x - k, M_x, M_n) = \phi_k + \frac{\pi}{3}, \quad (\text{S14})$$

$$\phi(M_n, M_x, M_n + k) = \phi_k + \frac{2\pi}{3}, \quad (\text{S15})$$

$$\phi(M_n, M_x - k, M_x) = \phi_k + \pi, \quad (\text{S16})$$

$$\phi(M_n + k, M_n, M_x) = \phi_k + \frac{4\pi}{3}, \quad (\text{S17})$$

$$\phi(M_x, M_n, M_x - k) = \phi_k + \frac{5\pi}{3}. \quad (\text{S18})$$

This half-open convention avoids duplicate assignments at boundaries between adjacent transition segments. For grayscale values  $M_x = M_n$ ,  $\phi$  is redundant and the canonical representative is  $\phi = 0$ .

## Polar-angle assignment for the 256<sup>3</sup> RGB grid

For the full 256<sup>3</sup> grid, the channel range is  $\{0, \dots, 255\}$ . When  $M_x + M_n = 255$ , the polar angles are assigned around the equator as

$$\theta_{(255-i:i)} = \begin{cases} \frac{\pi}{2} - \frac{\pi}{1020} \frac{k}{64}, & i = 2k - 1, \quad k = 1, \dots, 64, \\ \frac{\pi}{2}, & i = 0, \\ \frac{\pi}{2} + \frac{\pi}{1020} \frac{k}{63}, & i = 2k, \quad k = 1, \dots, 63. \end{cases} \quad (\text{S19})$$

For  $\alpha = 1, \dots, 254$ , define

$$\hat{f}_{256}(\alpha) = \left\lfloor \frac{255 - \alpha}{2} \right\rfloor. \quad (\text{S20})$$

The upper-hemisphere family  $M_x + M_n = 255 - \alpha$  is assigned by

$$\theta_{(255-\alpha-i:i)} = \begin{cases} (255 - \alpha) \frac{\pi}{510} - \frac{\pi}{1020} + \frac{\pi}{510} \frac{i}{\hat{f}_{256}(\alpha) + 1}, & \alpha \text{ odd}, \\ (255 - \alpha) \frac{\pi}{510} + \frac{\pi}{1020} - \frac{\pi}{510} \frac{i + 1}{\hat{f}_{256}(\alpha) + 1}, & \alpha \text{ even}, \end{cases} \quad (\text{S21})$$

where  $i = 0, \dots, \hat{f}_{256}(\alpha)$ . The lower-hemisphere family  $M_x + M_n = 255 + \alpha$  is assigned by

$$\theta_{(255-i:\alpha+i)} = \begin{cases} (255 + \alpha) \frac{\pi}{510} + \frac{\pi}{1020} - \frac{\pi}{510} \frac{i}{\hat{f}_{256}(\alpha) + 1}, & \alpha \text{ odd}, \\ (255 + \alpha) \frac{\pi}{510} - \frac{\pi}{1020} + \frac{\pi}{510} \frac{i + 1}{\hat{f}_{256}(\alpha) + 1}, & \alpha \text{ even}, \end{cases} \quad (\text{S22})$$

where  $i = 0, \dots, \hat{f}_{256}(\alpha)$ .

## Polar-angle assignment for the 128<sup>3</sup> RGB grid

For the 128<sup>3</sup> grid, the channel range is  $\{0, \dots, 127\}$ . The equatorial family  $M_x + M_n = 127$  is assigned by

$$\theta_{(127-i:i)} = \begin{cases} \frac{\pi}{2} + \frac{\pi}{508} \frac{k}{31}, & i = 2k, \quad k = 1, \dots, 31, \\ \frac{\pi}{2}, & i = 0, \\ \frac{\pi}{2} - \frac{\pi}{508} \frac{k}{32}, & i = 2k - 1, \quad k = 1, \dots, 32. \end{cases} \quad (\text{S23})$$

For  $\alpha = 1, \dots, 126$ , define

$$\hat{f}_{128}(\alpha) = \left\lfloor \frac{127 - \alpha}{2} \right\rfloor. \quad (\text{S24})$$

Then

$$\theta_{(127-\alpha-i:i)} = \begin{cases} (127 - \alpha) \frac{\pi}{254} - \frac{\pi}{508} + \frac{\pi}{254} \frac{i}{\hat{f}_{128}(\alpha) + 1}, & \alpha \text{ odd}, \\ (127 - \alpha) \frac{\pi}{254} + \frac{\pi}{508} - \frac{\pi}{254} \frac{i + 1}{\hat{f}_{128}(\alpha) + 1}, & \alpha \text{ even}, \end{cases} \quad (\text{S25})$$

and

$$\theta_{(127-i:\alpha+i)} = \begin{cases} (127 + \alpha) \frac{\pi}{254} + \frac{\pi}{508} - \frac{\pi}{254} \frac{i}{\hat{f}_{128}(\alpha) + 1}, & \alpha \text{ odd}, \\ (127 + \alpha) \frac{\pi}{254} - \frac{\pi}{508} + \frac{\pi}{254} \frac{i + 1}{\hat{f}_{128}(\alpha) + 1}, & \alpha \text{ even}. \end{cases} \quad (\text{S26})$$

## Polar-angle assignment for the $64^3$ RGB grid

For the  $64^3$  grid, the channel range is  $\{0, \dots, 63\}$ . The equatorial family  $M_x + M_n = 63$  is assigned by

$$\theta_{(63-i:i)} = \begin{cases} \frac{\pi}{2} + \frac{\pi}{252} \frac{k}{15}, & i = 2k, \quad k = 1, \dots, 15, \\ \frac{\pi}{2}, & i = 0, \\ \frac{\pi}{2} - \frac{\pi}{252} \frac{k}{16}, & i = 2k - 1, \quad k = 1, \dots, 16. \end{cases} \quad (\text{S27})$$

For  $\alpha = 1, \dots, 62$ , define

$$\hat{f}_{64}(\alpha) = \left\lfloor \frac{63 - \alpha}{2} \right\rfloor. \quad (\text{S28})$$

Then

$$\theta_{(63-\alpha-i:i)} = \begin{cases} (63 - \alpha) \frac{\pi}{126} - \frac{\pi}{252} + \frac{\pi}{126} \frac{i}{\hat{f}_{64}(\alpha) + 1}, & \alpha \text{ odd}, \\ (63 - \alpha) \frac{\pi}{126} + \frac{\pi}{252} - \frac{\pi}{126} \frac{i + 1}{\hat{f}_{64}(\alpha) + 1}, & \alpha \text{ even}, \end{cases} \quad (\text{S29})$$

and

$$\theta_{(63-i:\alpha+i)} = \begin{cases} (63 + \alpha) \frac{\pi}{126} + \frac{\pi}{252} - \frac{\pi}{126} \frac{i}{\hat{f}_{64}(\alpha) + 1}, & \alpha \text{ odd}, \\ (63 + \alpha) \frac{\pi}{126} - \frac{\pi}{252} + \frac{\pi}{126} \frac{i + 1}{\hat{f}_{64}(\alpha) + 1}, & \alpha \text{ even}. \end{cases} \quad (\text{S30})$$

## Supplementary Note 3: Region-of-point and LUT construction

For each RGB value  $(r, g, b) \in \Sigma_L$ , let  $(\theta_{(r,g,b)}, \phi_{(r,g,b)})$  denote its encoded angle pair. The RoP decoder assigns a decision region

$$I_{(r,g,b)} = I_{(r,g,b)}^\theta \times I_{(r,g,b)}^\phi \quad (\text{S31})$$

to each encoded color point. Given an estimated angle pair  $(\hat{\theta}, \hat{\phi})$ , the decoded RGB value is

$$(\hat{r}_L, \hat{g}_L, \hat{b}_L) = \text{LUT}(\hat{\theta}, \hat{\phi}) \quad \text{s.t.} \quad (\hat{\theta}, \hat{\phi}) \in I_{(\hat{r}_L, \hat{g}_L, \hat{b}_L)}. \quad (\text{S32})$$

The recovered  $L$ -level values are then mapped back to the display scale using the rule in Supplementary Note 1.

### Azimuthal RoP interval

For grayscale values  $M_x = M_n$ , the azimuthal coordinate is ignored by the primary grayscale decoder. For non-grayscale values  $M_x > M_n$ , the azimuthal half-width is

$$\Delta\phi_{(r,g,b)} = \frac{\pi}{6(M_x - M_n)}. \quad (\text{S33})$$

The azimuthal RoP interval is

$$I_{(r,g,b)}^\phi = [\phi_{(r,g,b)} - \Delta\phi_{(r,g,b)}, \phi_{(r,g,b)} + \Delta\phi_{(r,g,b)}] \pmod{2\pi}. \quad (\text{S34})$$

The modulo operation handles points near the wrap-around boundary at 0 and  $2\pi$ .

### Polar RoP interval

For each RGB grid, the set of distinct polar angles is sorted in increasing order,

$$0 \leq \theta_1 < \theta_2 < \dots < \theta_m \leq \pi. \quad (\text{S35})$$

The polar interval around  $\theta_j$  is defined by neighboring midpoints:

$$I_j^\theta = \begin{cases} [0, \frac{\theta_1 + \theta_2}{2}), & j = 1, \\ [\frac{\theta_{j-1} + \theta_j}{2}, \frac{\theta_j + \theta_{j+1}}{2}), & 1 < j < m, \\ [\frac{\theta_{m-1} + \theta_m}{2}, \pi], & j = m. \end{cases} \quad (\text{S36})$$

The exact LUT used in the experiments was generated by combining this midpoint-based polar rule with the azimuthal interval rule above.

## Minimum RoP half-widths

The minimum angular half-widths for the three tested RGB grids are summarized in Table S1. Coarser grids increase the smallest polar and azimuthal decision margins, which improves tolerance to finite-shot angle-estimation error.

Table S1: Minimum RoP half-widths for the tested RGB grids.

RGB grid	$\Delta\theta^{\min}$	$\Delta\phi^{\min}$
$256^3$	$\pi/130560$	$\pi/1530$
$128^3$	$\pi/32512$	$\pi/762$
$64^3$	$\pi/8064$	$\pi/378$

## Supplementary Note 4: Quantum encoding protocols

### Two-qubit pixel-wise validation circuit

Figure S1 shows the two-qubit circuit used for pixel-wise validation. The circuit uses an angle qubit  $q_1$  and a flag qubit  $q_2$ , where the  $q_2 = 1$  branch encodes  $\theta$  and the  $q_2 = 0$  branch encodes  $\phi$ .




Figure S1: Two-qubit circuit used in the pixel-wise validation mode. The angle qubit  $q_1$  is measured to estimate the selected angle, while the flag qubit  $q_2$  selects the branch for  $\theta$  or  $\phi$ .

After measurement, let

$$p_{ij} = \Pr(|q_1 q_2\rangle = |ij\rangle). \quad (\text{S37})$$

The conditional probabilities are

$$p_{\theta,0} = \frac{p_{01}}{p_{01} + p_{11}}, \quad p_{\phi,0} = \frac{p_{00}}{p_{00} + p_{10}}. \quad (\text{S38})$$

The angle estimates are

$$\hat{\theta} = 2 \arccos(\sqrt{p_{\theta,0}}), \quad \hat{\phi} = 4 \arccos(\sqrt{p_{\phi,0}}). \quad (\text{S39})$$

The factor 4 in  $\hat{\phi}$  follows from the half-angle embedding in the  $\phi$  branch.

### Position-register full-image encoding

For an image of size  $2^n \times 2^n$ , the position-register full-image protocol uses  $2n$  position qubits to index pixel locations. The position register is equivalent to an address register in circuit implementation. For grayscale images, only one angle is required, so a  $32 \times 32$  grayscale image uses 10 position qubits and one angle qubit, giving 11 qubits in total. For RGB images, both  $\theta$  and  $\phi$  are required, and the position-register full-image RGB formulation uses  $2n + 2$  qubits.

In the RGB setting, the total shot budget is shared across all pixel positions and both angular branches. For an  $H \times W$  image, the effective number of conditional samples per pixel and branch scales approximately as

$$\frac{\text{shots}}{2HW}. \quad (\text{S40})$$

This explains why full-image RGB reconstruction is more sensitive to per-position sampling sparsity than pixel-wise validation.

## Supplementary Note 5: Algorithmic description

---

**Algorithm S1** LUT construction for an  $L$ -level quantum color space

---

**Require:** RGB level  $L \in \{64, 128, 256\}$

- 1: Initialize an empty table  $\mathcal{T}$
  - 2: **for** each  $(r, g, b) \in \{0, \dots, L-1\}^3$  **do**
  - 3:   Compute  $M_x = \max(r, g, b)$  and  $M_n = \min(r, g, b)$
  - 4:   Compute  $(\theta_{(r,g,b)}, \phi_{(r,g,b)})$  using the angular assignment rules
  - 5:   Compute  $I_{(r,g,b)} = I_{(r,g,b)}^\theta \times I_{(r,g,b)}^\phi$
  - 6:   Store  $I_{(r,g,b)} \mapsto (r, g, b)$  in  $\mathcal{T}$
  - 7: **end for**
  - 8: **return** LUT  $\mathcal{T}$
- 

---

**Algorithm S2** Measurement-driven RGB reconstruction

---

**Require:** Measurement counts and LUT  $\mathcal{T}$

- 1: Estimate branch-conditioned probabilities  $p_{\theta,0}$  and  $p_{\phi,0}$
  - 2: Compute  $\hat{\theta} = 2 \arccos(\sqrt{p_{\theta,0}})$
  - 3: Compute  $\hat{\phi} = 4 \arccos(\sqrt{p_{\phi,0}})$
  - 4: Find  $(\hat{r}_L, \hat{g}_L, \hat{b}_L) = \text{LUT}(\hat{\theta}, \hat{\phi})$
  - 5: Map to display scale:  $(\hat{r}, \hat{g}, \hat{b}) = s(\hat{r}_L, \hat{g}_L, \hat{b}_L)$
  - 6: **return** Reconstructed RGB value  $(\hat{r}, \hat{g}, \hat{b})$
- 

Algorithms S1 and S2 summarize the two main lookup-based procedures used in the proposed quantum color-space framework. Algorithm S1 constructs the region-of-point (RoP) lookup table by assigning each discrete RGB triplet to its corresponding angular interval on the Bloch sphere. Algorithm S2 describes the inverse process: measurement statistics are first converted into estimated angular values, and the nearest compatible RGB value is then recovered through the LUT. Together, these procedures provide a practical bridge between quantum measurement outcomes and classical color reconstruction.

## Supplementary Note 6: Extended reconstruction visualizations

### Pixel-wise validation results

Figure S2 provides extended pixel-wise reconstruction examples on CIFAR-10 under the  $64^3$ ,  $128^3$ , and  $256^3$  grids. These results complement the main-text visualizations by showing all three quantized grids in the pixel-wise validation mode.

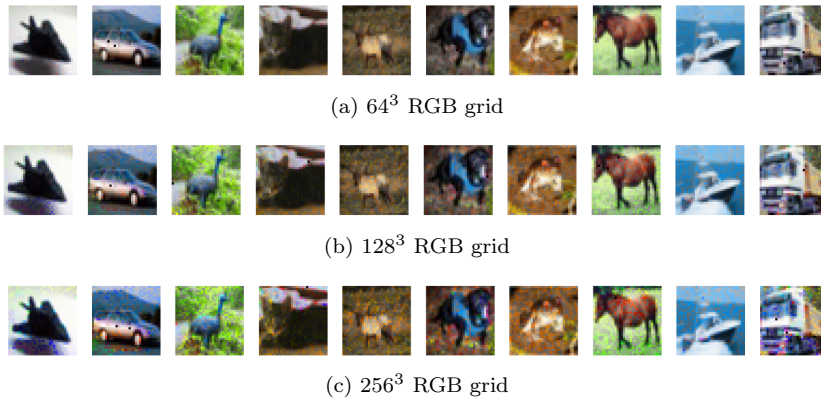


Figure S2: Extended pixel-wise CIFAR-10 RGB reconstruction results under the three quantized RGB grids.

## Position-register full-image results

Figure S3 shows reconstruction results obtained using the position-register full-image protocol. Compared with the pixel-wise validation protocol, all three grids exhibit degraded reconstruction quality because the total shot budget is shared across all pixel positions and both angular branches.



Figure S3: Extended position-register full-image CIFAR-10 reconstruction results under the three quantized RGB grids.

## Supplementary Note 7: Extended quantitative results

### MNIST grayscale reconstruction

Table S2: Digit-wise grayscale reconstruction performance on MNIST.

Metric	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

For grayscale images, the color-space recovery reduces to a one-dimensional angular estimation problem because the grayscale constraint fixes the color trajectory to a restricted subspace. This makes MNIST a useful diagnostic setting for verifying whether the proposed measurement-driven reconstruction preserves pixel-wise intensity structure before moving to the denser RGB color space.

Table S2 reports digit-wise grayscale reconstruction performance on MNIST using 100 samples per digit. The average MSE is 3.90 and the average PSNR is 42.44 dB, indicating that the reconstructed images preserve the main stroke structure with only minor intensity deviations. The lowest error is observed for digit '1', whose sparse foreground strokes reduce the number of pixels over which reconstruction errors can accumulate, whereas digit '0' shows the highest MSE due to its larger foreground coverage and curved structure. These results confirm that the proposed angular decoding and LUT-based recovery are reliable in the grayscale setting, where only a single dominant angular degree of freedom needs to be estimated.

### CIFAR-10 class-wise reconstruction results

Table S3: Class-wise CIFAR-10 reconstruction performance under the  $64^3$ ,  $128^3$ , and  $256^3$  RGB grids in the pixel-wise validation mode.

Grid	Metric	Airplane	Auto.	Bird	Cat	Deer	Dog	Frog	Horse	Ship	Truck	Avg.
$64^3$	MSE ↓	67.86	67.16	87.11	95.48	84.80	86.44	84.80	77.52	81.55	80.14	81.29
	PSNR ↑	30.58	30.26	29.68	29.59	29.44	29.66	29.70	29.82	29.71	29.98	29.84
	SSIM ↑	0.9128	0.9642	0.9036	0.9356	0.9071	0.9379	0.9326	0.9476	0.9146	0.9567	0.9313
$128^3$	MSE ↓	198.09	208.19	227.69	240.63	223.83	221.90	216.84	212.12	222.11	207.59	217.90
	PSNR ↑	25.63	25.26	25.15	24.93	24.91	25.01	25.13	25.12	24.98	25.28	25.14
	SSIM ↑	0.8284	0.9100	0.8145	0.8609	0.8110	0.8714	0.8557	0.8805	0.8309	0.9033	0.8567
$256^3$	MSE ↓	596.94	604.53	633.98	635.43	634.11	597.46	603.90	572.63	616.20	602.51	609.77
	PSNR ↑	20.77	20.52	20.45	20.42	20.40	20.60	20.62	20.76	20.46	20.53	20.55
	SSIM ↑	0.7121	0.8144	0.6887	0.7442	0.6762	0.7640	0.7341	0.7803	0.7162	0.8044	0.7435

Table S3 provides the full class-wise MSE, PSNR, and SSIM values for all three RGB grids in the pixel-wise validation mode. The results show a consistent trend across all CIFAR-10 classes: the  $64^3$  grid achieves the best reconstruction quality, with the lowest average MSE of 81.29, the highest average PSNR of 29.84 dB, and the highest average SSIM of 0.9313. Increasing the RGB grid resolution to  $128^3$  and  $256^3$  does not improve the reconstruction accuracy; instead, the average MSE increases to 217.90 and 609.77, while the average PSNR decreases to 25.14 dB and 20.55 dB, respectively.

This degradation is consistent with the finite-shot nature of the measurement process. A finer RGB grid creates a much denser angular codebook on the Bloch sphere, reducing the angular separation between neighboring color entries. Consequently, small measurement fluctuations in the estimated angles can lead to larger color-index mismatches in the LUT-based reconstruction. The effect is observed across all classes, although structured categories such as automobile and truck tend to retain relatively high SSIM, whereas visually more diverse or texture-rich classes such as bird, deer, and cat show larger degradation at higher grid resolutions. These class-wise results support the use of the  $64^3$  grid as the most stable operating point for the current finite-shot pixel-wise RGB reconstruction setting.

Table S4: Average CIFAR-10 reconstruction performance in the position-register full-image mode at  $2^{18}$  total shots.

RGB grid	$\theta$ -MAE	$\phi$ -MAE	MSE ↓	PSNR ↑	SSIM ↑
$64^3$	$\approx 0.072$ – $0.073$	$\approx 0.141$ – $0.147$	896.88	18.70	0.6960
$128^3$	$\approx 0.072$ – $0.073$	$\approx 0.141$ – $0.147$	880.82	18.78	0.6998
$256^3$	$\approx 0.072$ – $0.073$	$\approx 0.141$ – $0.147$	897.38	18.73	0.6999

Table S4 summarizes the average CIFAR-10 reconstruction performance in the position-register full-image mode at a fixed total shot budget of  $2^{18}$ . Unlike the pixel-wise validation results in Table S3, this experiment is intended to diagnose the measurement-allocation bottleneck rather than to compare class-dependent reconstruction behavior. Because the total shot budget is shared across all pixel positions and both angular branches, reconstruction quality remains uniformly low across the three RGB grids, indicating that per-position sampling sparsity dominates the full-image setting.

## Supplementary Note 8: Shot-wise robustness curves

### Shot-wise robustness analysis

Figure S4 summarizes the shot-wise reconstruction behavior for both the pixel-wise validation mode and the position-register full-image mode. Each panel reports MSE, PSNR, SSIM, and angular estimation error as the shot budget increases.

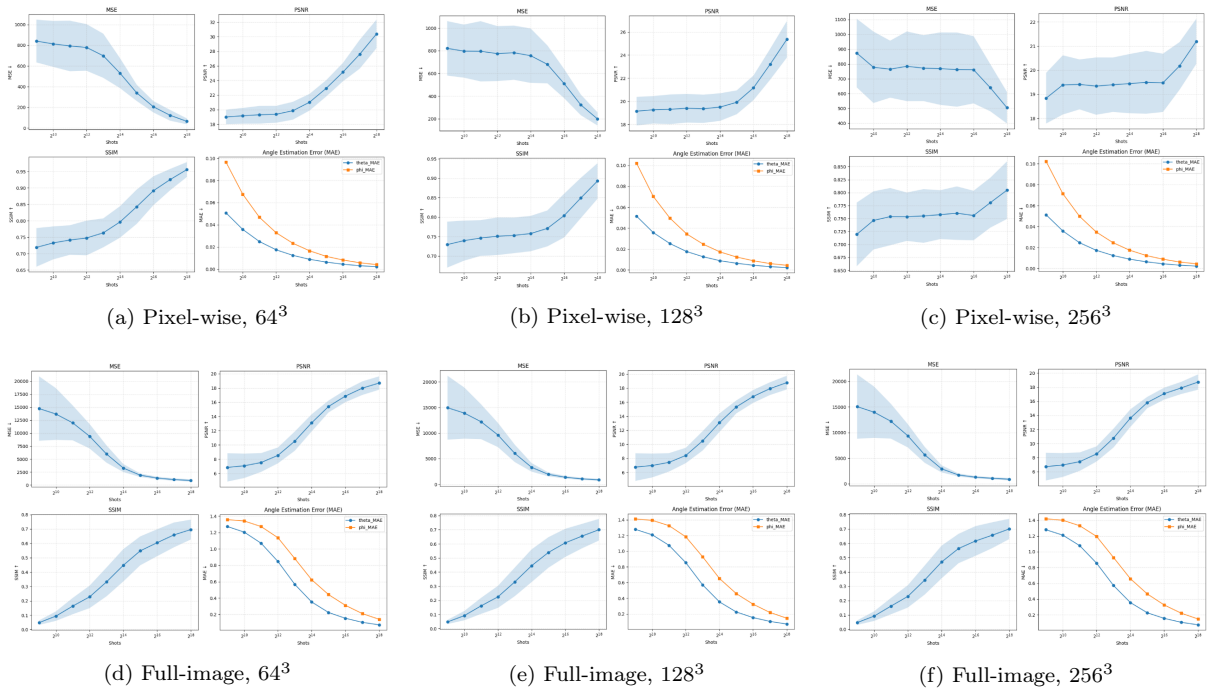


Figure S4: Shot-wise reconstruction behavior under the pixel-wise validation mode and the position-register full-image encoding mode. The top row shows pixel-wise validation results, where each pixel receives an independent shot budget. The bottom row shows position-register full-image results, where the total shot budget is shared across all pixel positions and both angular branches.