PSG: PRESERVING ACTIVATION STRUCTURE FOR LOW-BIT IMAGE RESTORATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Post-training quantization (PTQ) is a practical route to deploying image restoration models, yet restoration fidelity degrades sharply at ultra-low bit widths. A key challenge is channel-wise outliers, which are particularly pronounced in Mamba-based models and substantially inflate the activation range. Suppressing these extremes can tighten the range, but is not a reliable remedy, as even limited removal can severely degrade reconstruction quality. We show that these outliers are not random spikes, but exhibit strong persistence: their dominant direction remains stable across inputs, and even within each input, the response along this direction is effectively shared across tokens. We term this shared directional component the *pedestal*. Crucially, the pedestal not only inflates the range, but also defines a direction highly sensitive to quantization error, making it a structure that should be preserved rather than suppressed. Building on this, we propose PSG, a PTQ framework that combines Pedestal Shifting and GRILL. Pedestal Shifting keeps the pedestal outside the quantizer while reducing the activation range, and GRILL calibrates quantization bounds for the residual activations and weights through closed-form updates. Across diverse image restoration tasks and architectures, PSG consistently outperforms existing methods, delivering up to 3.73 dB higher PSNR at 2 bits, $4.33\times$ smaller models, and $6.13\times$ fewer bit-operations, while deploying as standard TensorRT engines with no additional inference overhead. Our code is available at: `https://anonymous.4open.science/r/PSG-AD60`.

1 INTRODUCTION

Image restoration (IR) has advanced from convolutional architectures (Dong et al., 2014; Zhang et al., 2017; Lim et al., 2017) to Transformer-based models (Chen et al., 2021; Liang et al., 2021; Zamir et al., 2022) and, more recently, state-space models (SSMs) such as Mamba (Gu & Dao, 2024; Guo et al., 2024; 2025; Li et al., 2025). Transformers have substantially improved restoration quality by effectively capturing global context, while SSM-based restorers efficiently leverage linear-complexity sequence modeling and have further improved restoration performance. Despite these advances, modern IR backbones still incur substantial computational and memory costs, limiting their deployment on resource-constrained devices (Liu et al., 2024; Chen et al., 2026).

Low-bit quantization maps floating-point weights and activations to low-bit representations, reducing model size and inference latency (Jacob et al., 2018). While quantization-aware training (QAT) can effectively recover quantization-induced accuracy loss, it typically requires access to the full training data and costly retraining of the model and quantizer parameters (Esser et al., 2020). Post-training quantization (PTQ), in contrast, operates on a pretrained model using only a small set of unlabeled calibration samples, making it attractive for practical deployment (Nagel et al., 2020; Li et al., 2021). However, applying PTQ to image restoration models often leads to notable degradation in fidelity, especially at ultra-low bit widths (< 4 bits) (Liu et al., 2024).

A key source of this degradation is the trade-off between rounding and clipping errors in uniform quantization (Banner et al., 2019). This trade-off is particularly severe in image restoration, where activations often exhibit strongly asymmetric and long-tailed distributions (Tu et al., 2023; Liu et al., 2024). Mamba-based models further exhibit pronounced channel-wise outliers that persist across inputs (Xu et al., 2025; Cho et al., 2025). Under per-tensor quantization, a few extreme channels

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

stretch the pooled range, forcing a fixed number of quantization levels to cover the tails and leaving fewer levels to represent the high-density region of activations (Xiao et al., 2023; Dettmers et al., 2022). Simply truncating the extremes to tighten the range, however, is not a reliable solution: recent work on PTQ for image super-resolution has shown that outlier removal can significantly degrade reconstruction quality (Wang et al., 2025). Our analysis further shows that a $1.5\times$ range reduction clips only 0.14% of activations yet degrades restoration fidelity by 2.86 dB (Appendix A).

Rather than treating these extremes as random spikes, we examine their underlying structure. We find that the recurring extreme-channel pattern consistently aligns with a dominant direction across inputs. This input-level persistence extends to the token level: within each input, a single value can represent the token responses along this direction with limited effect on restoration quality. The input-wise persistence and token-wise functional constancy allow us to isolate a shared directional offset, which we term the *pedestal*. The pedestal is more than a range-inflating structure: distortions along this direction have a disproportionately large impact on restoration quality. These properties suggest that the pedestal should be preserved rather than suppressed.

To preserve this structure while removing its adverse effect on quantization, we introduce Pedestal Shifting, which keeps the pedestal outside the quantizer. Channel-wise shifting is an established tool for handling activation asymmetry, where the shift is an unconstrained per-channel vector set from range statistics (Wei et al., 2023) or learned by block-wise reconstruction (Shao et al., 2024). Pedestal Shifting instead constrains the shift to a single structured component of the activation, the pedestal, and restores it in full precision through offline folding. This removes the sensitive and range-inflating activation structure from the quantizer, reducing both the quantization error along the pedestal direction and the per-tensor activation range. We complement this with GRILL (GRId-constrained Lloyd–Max), which calibrates uniform quantization bounds for the residual activations and weights through closed-form updates under tensor-wise and input-induced objectives. Pedestal Shifting and GRILL form PSG, our post-training quantization framework for image restoration.

Across multiple tasks, PSG achieves state-of-the-art performance on diverse Mamba-based restorers under ultra-low precision, while evaluation on a Transformer model demonstrates its generality across architectures. We further evaluate deployment efficiency, from model size and bit-operations to latency and peak memory on TensorRT engines. Our main contributions are:

- We uncover the structure underlying channel-wise activation outliers in image restoration: a single dominant direction that persists across inputs, is functionally constant across tokens, and along which quantization error disproportionately degrades fidelity. We term its shared component the *pedestal* and show that it should be preserved rather than suppressed.
- Building on this finding, we propose **PSG**, a PTQ framework that optimizes only the quantization bounds. **Pedestal Shifting** removes the pedestal before quantization and restores it in full precision through offline folding. **GRILL** then fits the bounds of the residual activations and weights on the deployable uniform grid through closed-form updates.
- Extensive experiments across IR tasks and architectures validate PSG: up to 3.73 dB higher PSNR than existing methods at 2 bits, $4.33\times$ smaller models, $6.13\times$ fewer bit-operations, and deployment as standard TensorRT engines with no additional inference overhead.

2 Related Work

2.1 Image Restoration

Deep image restoration has progressed from convolutional architectures to Transformers and, more recently, to state-space models. Early CNNs established learned restoration, from SRCNN (Dong et al., 2014) for super-resolution and DnCNN (Zhang et al., 2017) for denoising to the deeper residual design of EDSR (Lim et al., 2017). Transformers subsequently improved long-range modeling: IPT (Chen et al., 2021) pre-trained a Transformer for restoration, SwinIR (Liang et al., 2021) adopted shifted-window attention for efficient spatial modeling, and Restormer (Zamir et al., 2022) reformulated attention across channels for high-resolution restoration. More recently, state-space models built on Mamba (Gu & Dao, 2024) have emerged as efficient alternatives through linear-complexity global context modeling. MambaIR (Guo et al., 2024) augments selective state-space modeling with local enhancement and channel attention, while MambaIRv2 (Guo et al., 2025) alle-

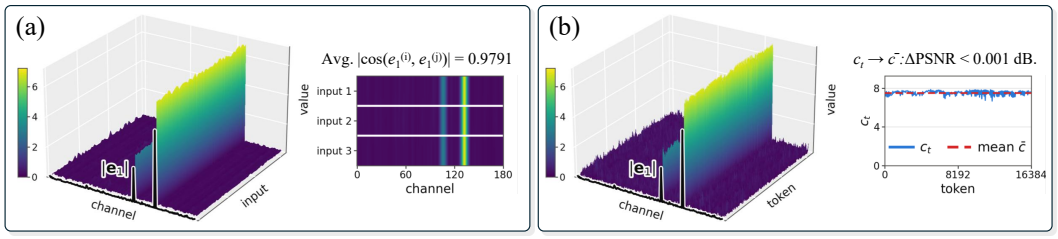


Figure 1: (a) Extreme-channel patterns recur across inputs and align with e_1. (b) The pattern persists across tokens, while e_1 responses remain functionally constant under $c_t \to \bar{c}$.

viates causal scanning limitations through attentive state-space modeling and provides a lightweight variant. MaIR (Li et al., 2025) further preserves image locality and spatial continuity through nested scanning. Despite these advances, deploying strong restoration backbones on resource-constrained devices remains challenging, as many still incur substantial computation and memory demands.

2.2 Quantization for Image Restoration

Low-bit quantization offers an effective way to reduce model size and inference latency in restoration models. QAT can recover low-bit accuracy but typically requires full training data and costly retraining, whereas PTQ calibrates a pretrained model with unlabeled calibration data, making it practical for deployment (Tu et al., 2023). Most restoration-specific PTQ has focused on super-resolution, with task-specific exceptions such as image demoiréing (Chen et al., 2025). DBDC+PaC (Tu et al., 2023) and 2DQuant (Liu et al., 2024) tailor quantization ranges to activation distributions and refine them with full-precision guidance, targeting CNN- and Transformer-based SR models, respectively. Recent methods address complementary error sources through separate quantization of dense and outlier regions (Wang et al., 2025), joint treatment of weight and activation quantization (Wang et al., 2026), or quantization-aware low-rank rectification in SPR^{2}Q (Xin et al., 2026). Mamba-specific PTQ has identified pronounced channel-wise outliers, heavy-tailed distributions, and persistent channel structure (Cho et al., 2025; Xu et al., 2025; Chiang et al., 2025), and suppresses them through smoothing, rotation, or channel reordering, but primarily targets generic vision or language models. Although SPR^{2}Q evaluates MambaIRv2-light, it is limited to super-resolution, while Q-MambaIR (Chen et al., 2026) studies Mamba-based image restoration across multiple tasks but requires quantization-aware retraining. Thus, PTQ for broader image restoration tasks remains comparatively underexplored, particularly for Mamba-based models.

3 Persistence and Sensitivity Along the Dominant Direction

We analyze the recurring extreme-channel pattern in Mamba-based restoration models through its dominant direction and identify two notable properties. First, this direction remains highly consistent across inputs, and its token-wise responses exhibit constancy within each input: a single value can represent the responses across all tokens. Second, quantization-induced distortion along this direction disproportionately degrades restoration quality.

Let $X = [x_1, \ldots, x_T]^\top \in \mathbb{R}^{T \times C}$ denote an activation matrix, where $x_t \in \mathbb{R}^C$ is the activation vector of token t. We formalize the dominant channel-space direction e_1 as follows:

$$M = \frac{1}{T} X^\top X, \qquad e_1 = \arg\max_{\|e\|_2 = 1} e^\top M e. \tag{1}$$

e_1 maximizes the average squared projection of token activations over all unit directions in channel space. We use the uncentered second moment to retain channel-wise offset structure.

3.1 Persistence Across Inputs and Tokens

Input-wise persistence. Prior work has shown that activation outliers in Mamba-based models recur in specific channels across inputs (Cho et al., 2025; Chiang et al., 2025). We observe this

extreme-channel pattern in Mamba-based restoration models. In Fig. 1(a), large activation magnitudes recur in similar channels and this pattern aligns with e_1. Dominant directions estimated independently from individual inputs are also highly aligned, with an average pairwise absolute cosine similarity of 0.979, indicating that the channel-space direction remains consistent across inputs.

Token-wise functional constancy. This persistence further extends across tokens within each input. As shown in Fig. 1(b), the extreme channels remain prominent throughout the token dimension. The right panel shows $c_t = x_t^\top e_1$, the projection coefficient of each token onto the dominant direction. We decompose each activation into its component along e_1 and the remaining orthogonal component, and replace the token-dependent coefficient c_t with a single constant $\bar{c}$:

$$\begin{aligned} c_t &= x_t^\top e_1, & x_t &= c_t e_1 + x_t^\perp, \\ \bar{c} &= \frac{1}{T}\sum_{t=1}^{T} c_t, & \tilde{x}_t &= \bar{c} e_1 + x_t^\perp. \end{aligned} \tag{2}$$

Replacing x_t with $\tilde{x}_t$ has only a limited effect on restoration quality, even under joint replacement. This indicates that, despite token-wise variation in c_t, the response along e_1 is functionally constant across tokens and can be represented by a single value while largely preserving restoration quality. Detailed analyses of coefficient variation and joint replacement are provided in Appendix B.

3.2 Sensitivity to Directional Distortion

Directional perturbation. While Sec. 3.1 shows that token-wise deviations from $\bar{c}$ can be removed while preserving restoration quality, perturbing the shared offset along e_1 produces a markedly different effect. We inject equal-norm perturbations along e_1 or random channel-space directions. Perturbations along e_1 cause a 6.76 dB PSNR drop on average across Mamba-based restorers, compared with only 0.56 dB for random directions, a $12.1\times$ larger functional impact. The same tendency persists across perturbation magnitudes, indicating that the persistent e_1-aligned component is substantially more sensitive to distortion than unstructured channel-space offsets. Perturbing e_2 to e_5 also degrades PSNR far more than random directions, with the most harmful of them reaching 5.47 dB on average. Their responses, however, vary strongly across tokens (Appendix B), so e_1 is the only sensitive direction whose response can be represented by a single value.

Actual quantization error. The same directional asymmetry arises in actual low-bit quantization. Let $\epsilon_t = Q(x_t) - x_t$ denote the activation quantization error, which we decompose into its component along e_1 and the remaining component orthogonal to e_1:

$$\epsilon_t = \underbrace{(\epsilon_t^\top e_1)e_1}_{\epsilon_{t,\parallel}} + \underbrace{\left[\epsilon_t - (\epsilon_t^\top e_1)e_1\right]}_{\epsilon_{t,\perp}}. \tag{3}$$

For each error component, we measure its contribution to total activation distortion and the recovery in restoration fidelity after its removal. For MambaIR, $\epsilon_{t,\parallel}$ accounts for only 2.67% of the distortion, yet its removal recovers 53.3% of the lost fidelity. In contrast, removing $\epsilon_{t,\perp}$ eliminates 97.3% of the distortion but recovers only 58.0%. This disparity between distortion magnitude and restoration impact indicates that e_1-aligned errors disproportionately affect fidelity.

4 Method

Fig. 2 provides an overview of PSG, which consists of Pedestal Shifting and GRILL. Pedestal Shifting extracts a fixed e_1-aligned pedestal, removes it before quantization, and restores its contribution through offline folding into adjacent static operators. GRILL then calibrates the bounds of the residual activations and weights on a uniform affine grid using closed-form updates. For weights, the same update structure is extended to an input-induced objective that better reflects output distortion.

Uniform affine quantization. We adopt uniform affine quantization for both weights and activations (Jacob et al., 2018; Nagel et al., 2021). Given a tensor Z, bit-width b, and quantization bounds (l, u), let $K = 2^b - 1$ and $s = (u - l)/K$. We define the quantize–dequantize operation as:

$$Q(Z; l, u) = l + s \left\lfloor \frac{\mathrm{clip}(Z, l, u) - l}{s} \right\rceil. \tag{4}$$

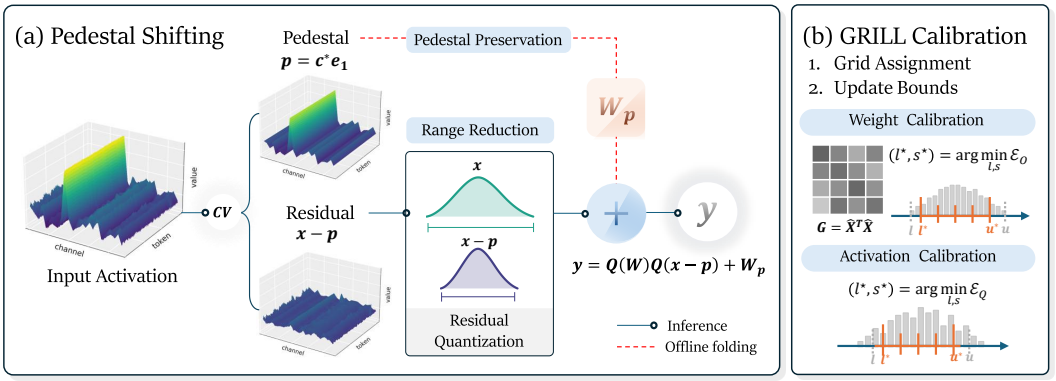


Figure 2: Overview of PSG. (a) Pedestal Shifting preserves the e_1-aligned component while reducing the activation range. (b) GRILL calibrates weight and activation bounds on a uniform grid.

4.1 Pedestal Shifting

We apply Pedestal Shifting only where the pedestal shift and its compensation can be absorbed into adjacent static operators through offline reparameterization, incurring no additional inference overhead (Wei et al., 2023). We refer to each such input as a *site*.

Pedestal estimation. For each site, we estimate the dominant direction e_1 from the calibration set by accumulating the uncentered second moment M in Eq. 1. Building on the functional constancy in Sec. 3.1, we represent $\{c_t\}$ by a single token-independent value and define the pedestal as:

$$p = c^\star e_1. \tag{5}$$

After removing the pedestal, the activation retains the residual $(c_t - c^\star)e_1$ and orthogonal component $x_t^\perp$. Since the pedestal is intended to capture the persistent structure along e_1, rather than to serve as a free offset for optimizing the entire activation distribution, we determine its magnitude $c^\star$ solely from $\{c_t\}$ and focus on the range of the e_1-aligned residual. Although Sec. 3.1 shows that their token-wise variation has limited functional importance, the residual still contributes to the range for quantization. We thus choose $c^\star$ to minimize the range occupied by the e_1-aligned residual.

For each channel, subtracting the pedestal does not change the width of the residual interval; it only changes its position. For the e_1-aligned residual, the shared bounds must cover the union of these channel-wise intervals, so the pooled range depends on their overlap. To make this union as compact as possible, we place $c^\star$ halfway between the minimum and maximum projections onto e_1:

$$c^\star = \arg\min_c \max_t |c_t - c| = \frac{\max_t c_t + \min_t c_t}{2}. \tag{6}$$

This choice centers every channel-wise residual interval at zero, making the intervals nested. The pooled range therefore coincides with the widest individual interval and attains its minimum possible width. The midrange thus serves as a representative of the persistent e_1-aligned response while minimizing the pooled range left by this structured component. Unlike channel-wise shifting, which estimates a separate offset for each activation channel, our pedestal is estimated as a single structured component along the persistent dominant direction. Even when estimated from small calibration sets, it remains substantially more consistent in both direction and overall scale (Table 1(a)).

Constancy-gated shifting. To safeguard the token-constancy assumption, we apply the shift only when the projection coefficients show limited token-wise variation. Specifically, we measure their coefficient of variation (CV) over the calibration set and apply the shift when:

$$\mathrm{CV}(c) = \frac{\sigma_c}{|\bar{c}| + \delta} \le \tau, \tag{7}$$

where $\bar{c}$ and σ_c denote the mean and standard deviation of $\{c_t\}$, respectively, and δ is a small constant for numerical stability. Sites exceeding the threshold are left untouched, as a single constant can be

Table 1: (a) Pedestal vs. per-channel midrange estimators over five calibration sets, other calibration settings fixed (Cos., pairwise cosine similarity; Dist., relative ℓ_2 distance). (b) τ sensitivity at 2 bits.

(a) Pedestal robustness

Estimator	Cos. ↑	Dist. (%) ↓	PSNR ↑
$c^\star e_1$	**0.9974**	**7.6**	**35.42±0.07**
$[c_1^\star, \ldots, c_C^\star]^\top$	0.8764	47.0	31.09±1.56

(b) Constancy-gate sensitivity

τ	0.00	0.05	**0.10**	0.15	∞
PSNR	35.563	36.131	**36.197**	36.128	35.402

insufficient to represent strongly token-dependent responses. A single threshold τ is used across all settings, with its sensitivity evaluated in Table 1(b). At a gated site, the quantized layer becomes:

$$y = Q(W)\, Q(x - p) + Wp. \tag{8}$$

Here, Wp is precomputed from the full-precision weight, allowing the pedestal to bypass both weight and activation quantization. The persistence and functional constancy established in Sec. 3.1 allow the pedestal to be represented by a fixed site-wise vector p. Since Pedestal Shifting is applied only at the foldable sites defined above, both the shift and its compensation are absorbed offline and introduce no additional inference overhead (Appendix C). The bias term is omitted for clarity.

4.2 GRILL Calibration

Pedestal Shifting reduces the activation range, but residual activations can remain asymmetric, requiring lower and upper quantization bounds that are not symmetrically coupled. Existing SR PTQ methods initialize these bounds from statistics or range search and refine them with full-precision guidance (Tu et al., 2023; Liu et al., 2024). Recent methods keep this framework while adding weight-side corrections (Wang et al., 2026; Xin et al., 2026). We retain the two-stage structure but leave the weights fixed, replacing the initialization with GRILL (GRId-constrained Lloyd–Max), which fits the bounds in closed form. Motivated by Lloyd–Max (Lloyd, 1982; Max, 1960) and MSE-optimal clipping (Sakr et al., 2022), GRILL adopts an alternating assignment–update structure but constrains the reconstruction levels to a uniform affine grid, so only the bounds are updated.

Grid-constrained Lloyd–Max calibration. Under the quantizer in Eq. 4, the deployable quantization levels are $\{l + sk\}_{k=0}^{K}$. For the entries $\{v_i\}$ of a tensor, GRILL alternates two steps:

$$\begin{aligned}
&\textbf{Step 1: Grid assignment} \quad k_i = \mathrm{clip}\left(\left\lfloor \frac{v_i - l}{s} \right\rceil, 0, K\right), \qquad \mathcal{E}_Q(l, s; \mathbf{k}) = \sum_i (v_i - l - sk_i)^2, \\
&\textbf{Step 2: Bound update} \quad (l^\star, s^\star) = \arg\min_{l,s} \mathcal{E}_Q(l, s; \mathbf{k}), \\
&\qquad\qquad\qquad\qquad s^\star = \frac{\sum_i (k_i - \bar{k})(v_i - \bar{v})}{\sum_i (k_i - \bar{k})^2}, \qquad l^\star = \bar{v} - s^\star \bar{k}, \qquad u^\star = l^\star + K s^\star.
\end{aligned} \tag{9}$$

For fixed assignments $\mathbf{k}$, Step 2 gives the exact minimizer of the squared quantization error $\mathcal{E}_Q$ in closed form. Each Step 2 update therefore directly solves for the optimal grid parameters under the current assignments. Alternating these two steps reduces $\mathcal{E}_Q$ while keeping all quantization levels on a single uniform affine grid. Since the grid is not constrained to symmetric bounds, it can naturally adapt to asymmetric residual distributions. For weight quantization, however, the tensor-wise error $\mathcal{E}_Q$ does not account for how weight distortion affects the layer output. We therefore preserve the same closed-form update structure while replacing the Step 2 objective with an input-induced quadratic error. Let $\hat{X}$ denote the shifted and quantized input to a weight layer, and let $\mathbf{w}_{\mathrm{FP}}$ and $\mathbf{w}_{\mathrm{Q}} = l\mathbf{1} + s\mathbf{k}$ denote its full-precision and quantized weight row, respectively.

$$\begin{aligned}
&\textbf{Step 2: Bound update} \quad \mathcal{E}_O(l, s; \mathbf{k}) = (\mathbf{w}_{\mathrm{Q}} - \mathbf{w}_{\mathrm{FP}})^\top G (\mathbf{w}_{\mathrm{Q}} - \mathbf{w}_{\mathrm{FP}}), \qquad G = \hat{X}^\top \hat{X}, \\
&\qquad\qquad\qquad\qquad (l^\star, s^\star) = \arg\min_{l,s} \mathcal{E}_O(l, s; \mathbf{k}), \\
&\qquad\qquad\qquad\qquad \begin{bmatrix} \mathbf{1}^\top G \mathbf{1} & \mathbf{1}^\top G \mathbf{k} \\ \mathbf{k}^\top G \mathbf{1} & \mathbf{k}^\top G \mathbf{k} \end{bmatrix} \begin{bmatrix} l^\star \\ s^\star \end{bmatrix} = \begin{bmatrix} \mathbf{1}^\top G \mathbf{w}_{\mathrm{FP}} \\ \mathbf{k}^\top G \mathbf{w}_{\mathrm{FP}} \end{bmatrix}.
\end{aligned} \tag{10}$$

$\mathcal{E}_O$ is the layer-wise reconstruction objective that underlies weight-optimizing PTQ (Nagel et al., 2020; Frantar et al., 2023). GRILL keeps round-to-nearest weights and moves only the two bounds

Table 2: Quantitative comparison on MambaIRv2-light for lightweight SR $\times 3$.

Method	Bit	Set5		Set14		B100		Urban100		Manga109	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
MambaIRv2-light	32/32	34.707	0.9298	30.680	0.8483	29.262	0.8098	29.012	0.8689	34.412	0.9497
OS+	4/4	32.872	0.9056	29.343	0.8190	28.328	0.7824	26.310	0.8004	30.712	0.9112
PTQ4VM	4/4	32.226	0.8944	28.825	0.8062	27.940	0.7698	25.895	0.7863	29.972	0.8954
2DQuant	4/4	33.418	0.9148	29.715	0.8285	28.601	0.7909	27.044	0.8243	31.839	0.9272
SPR^{2}Q	4/4	33.536	0.9179	29.821	0.8316	28.685	0.7941	27.146	0.8288	32.032	0.9309
HarmoQ	4/4	33.519	0.9154	29.767	0.8292	28.612	0.7906	27.051	0.8233	31.979	0.9279
PSG (Ours)	4/4	**33.854**	**0.9210**	**29.984**	**0.8353**	**28.818**	**0.7976**	**27.596**	**0.8402**	**32.511**	**0.9363**
OS+	3/3	31.805	0.8781	28.670	0.7949	27.857	0.7600	25.570	0.7655	29.347	0.8750
PTQ4VM	3/3	31.303	0.8697	28.229	0.7769	27.598	0.7462	25.225	0.7443	28.468	0.8476
2DQuant	3/3	31.317	0.8853	28.339	0.8013	27.727	0.7662	25.086	0.7632	27.997	0.8759
SPR^{2}Q	3/3	31.113	0.8827	28.209	0.7990	27.639	0.7639	24.949	0.7596	27.774	0.8737
HarmoQ	3/3	32.326	0.8999	28.981	0.8128	28.114	0.7770	25.803	0.7863	29.596	0.8983
PSG (Ours)	3/3	**33.018**	**0.9124**	**29.476**	**0.8249**	**28.466**	**0.7879**	**26.604**	**0.8145**	**31.132**	**0.9216**
OS+	2/2	30.469	0.8358	27.867	0.7599	27.307	0.7224	24.809	0.7158	27.556	0.8157
PTQ4VM	2/2	27.962	0.6824	25.960	0.6046	25.862	0.5954	23.702	0.5798	25.872	0.6414
2DQuant	2/2	30.909	0.8799	28.056	0.7955	27.539	0.7603	24.817	0.7554	27.544	0.8709
SPR^{2}Q	2/2	30.906	0.8796	28.054	0.7953	27.537	0.7600	24.816	0.7551	27.541	0.8706
HarmoQ	2/2	30.863	0.8710	28.057	0.7897	27.539	0.7561	24.847	0.7504	27.591	0.8612
PSG (Ours)	2/2	**32.270**	**0.9008**	**28.956**	**0.8146**	**28.123**	**0.7786**	**25.755**	**0.7878**	**29.514**	**0.9001**

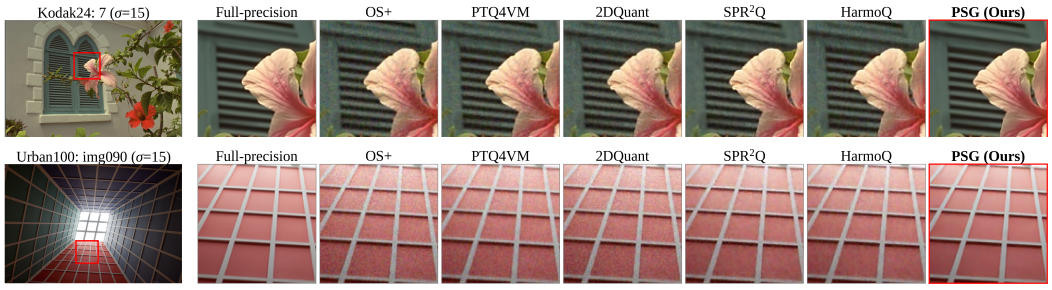


Figure 3: Qualitative comparison on MaIR for CDN ($\sigma = 15$) under 3-bit (top) and 2-bit (bottom).

in closed form. We evaluate $\mathcal{E}_O$ on the shifted and quantized input $\hat{X}$ of the partially quantized model, so that the bounds of each layer are calibrated against the error accumulated from preceding layers. Following existing IR PTQ methods (Tu et al., 2023; Liu et al., 2024), we further refine only the GRILL-calibrated bounds through distillation from the full-precision model. The network weights remain frozen, and PSG introduces no learnable parameters beyond the bounds.

5 Experiments

5.1 Experimental Setup

Dataset and evaluation. Following the evaluation setup of MambaIR (Guo et al., 2024), we evaluate image super-resolution (classic, lightweight, and real-world SR), Gaussian color denoising, and JPEG compression artifact reduction (JPEG CAR). Our experiments cover MambaIR (Guo et al., 2024), MambaIRv2-light (Guo et al., 2025), and MaIR (Li et al., 2025), with SwinIR (Liang et al., 2021) further included to assess generalization beyond Mamba architectures. For SR, we use Set5 (Bevilacqua et al., 2012), Set14 (Zeyde et al., 2012), B100 (Martin et al., 2001), Urban100 (Huang et al., 2015), and Manga109 (Matsui et al., 2017). For Gaussian color denoising, we use BSD68 (Martin et al., 2001), Kodak24 (Franzen, 1999), McMaster (Zhang et al., 2011), and Urban100 (Huang et al., 2015); for JPEG CAR, Classic5 (Foi et al., 2007) and LIVE1 (Sheikh et al., 2006). We report PSNR and SSIM for quantitative evaluation, computed on the Y channel of YCbCr for SR and JPEG CAR, and on RGB images for denoising.

Table 3: Quantitative comparison on MambaIR for Gaussian color denoising and JPEG CAR.

(a) Color denoising ($\sigma = 25$)

Method	Bit	BSD68		Kodak24		McMaster		Urban100	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
MambaIR	32/32	31.860	0.895	32.992	0.894	33.433	0.907	33.210	0.933
OS+	3/3	28.699	0.779	29.192	0.759	29.430	0.760	27.937	0.797
PTQ4VM	3/3	28.389	0.754	28.695	0.726	28.800	0.724	27.324	0.761
2DQuant	3/3	30.042	0.838	30.623	0.826	30.454	0.824	29.315	0.855
SPR^{2}Q	3/3	29.892	0.832	30.403	0.819	30.060	0.811	28.986	0.847
HarmoQ	3/3	28.996	0.802	29.645	0.792	29.754	0.793	28.248	0.824
PSG (Ours)	3/3	**31.455**	**0.886**	**32.436**	**0.883**	**32.685**	**0.890**	**32.035**	**0.918**
OS+	2/2	27.857	0.739	28.224	0.713	28.482	0.711	26.706	0.747
PTQ4VM	2/2	27.521	0.723	27.924	0.699	28.146	0.698	26.204	0.730
2DQuant	2/2	27.812	0.734	28.190	0.708	28.149	0.697	26.483	0.740
SPR^{2}Q	2/2	28.393	0.752	28.717	0.725	28.564	0.714	27.339	0.761
HarmoQ	2/2	27.499	0.713	27.874	0.687	27.994	0.682	26.513	0.727
PSG (Ours)	2/2	**31.116**	**0.878**	**31.947**	**0.872**	**32.027**	**0.875**	**31.064**	**0.904**

(b) JPEG CAR ($q = 10$)

Method	Bit	Classic5		LIVE1	
		PSNR	SSIM	PSNR	SSIM
MambaIR	32/32	30.302	0.826	29.926	0.830
OS+	3/3	28.476	0.769	28.348	0.784
PTQ4VM	3/3	28.536	0.769	28.430	0.786
2DQuant	3/3	28.723	0.780	28.661	0.796
SPR^{2}Q	3/3	28.870	0.784	28.820	0.798
HarmoQ	3/3	28.536	0.773	28.374	0.788
PSG (Ours)	3/3	**29.953**	**0.817**	**29.641**	**0.823**
OS+	2/2	28.361	0.771	28.234	0.788
PTQ4VM	2/2	28.311	0.764	28.192	0.781
2DQuant	2/2	28.513	0.772	28.371	0.786
SPR^{2}Q	2/2	28.659	0.777	28.597	0.792
HarmoQ	2/2	28.352	0.771	28.220	0.786
PSG (Ours)	2/2	**29.621**	**0.809**	**29.381**	**0.817**

Table 4: Deployment efficiency on the $\times 2$ SR backbones with a 512×384 input. (a) Size and MACs of the exported models, with compression and bit-operation (BOPs) reduction relative to FP32 and the selective-scan recurrence kept at full precision. (b) TensorRT engines on an NVIDIA B200.

Backbone	(a) Model size and compute								(b) TensorRT engines (B200)			
	Size	Compression			MACs	BOPs reduction			Latency (ms)		Peak memory (MB)	
	(MB)	W4A4	W3A3	W2A2	(G)	W4A4	W3A3	W2A2	FP32	W4A16	FP32	W4A16
MambaIRv2-light	3.4	1.84×	1.90×	1.97×	233.1	5.78×	5.98×	6.13×	68.5	51.0	3236	1581
MambaIR	82.0	3.53×	3.89×	4.33×	5266.3	2.67×	2.70×	2.72×	1884.9	1774.7	16526	15739
MaIR	78.1	3.33×	3.64×	4.02×	5036.9	2.56×	2.59×	2.61×	1968.4	1810.5	17167	10408

Quantization settings. We quantize weights and activations to the same bit width and evaluate W4A4, W3A3, and W2A2. For fair comparison, we align all methods to a single quantization scope: every linear, convolutional, and attention operator inside the blocks and the operands of the selective scan (Chiang et al., 2025) are quantized, while the trunk convolutions, normalizations, and stored SSM constants remain in full precision, following prior SR PTQ (Liu et al., 2024; Xin et al., 2026). PSG uses eight calibration images, with the effect of calibration-set size analyzed in Table 5(b). For the other baselines, we follow their official implementations and recommended settings. Further implementation and optimization details are provided in Appendix C.

5.2 Results

We compare PSG with OS+ (Wei et al., 2023), 2DQuant (Liu et al., 2024), SPR^{2}Q (Xin et al., 2026), and HarmoQ (Wang et al., 2026), and additionally include PTQ4VM (Cho et al., 2025) for Mamba architectures. OS+ also serves as a baseline for conventional channel-wise shifting.

Main results. Tables 2 and 3 show consistent improvements across restoration tasks and bit widths. PSG outperforms methods that additionally correct the weights (Wang et al., 2026; Xin et al., 2026) while optimizing only the quantization bounds. For lightweight SR $\times 3$, PSG achieves the best PSNR and SSIM across all five benchmarks and bit widths, with the margin widening toward lower precision: at 2 bits, the PSNR gain over the strongest baseline reaches 1.92 dB. The advantage extends beyond SR on MambaIR, where PSG improves by up to 2.72 dB at 3 bits and 3.73 dB at 2 bits for color denoising. For JPEG CAR, PSG remains close to full precision at 3 bits, with PSNR gaps of only 0.35 and 0.29 dB on Classic5 and LIVE1, respectively. Qualitative results on MaIR in Fig. 3 show that PSG better suppresses residual noise while preserving fine structures, producing restorations close to the full-precision outputs. The results on MambaIRv2-light, MambaIR, and MaIR demonstrate consistent effectiveness across Mamba-based restoration backbones. Additional task and backbone results are provided in Appendix D, including SwinIR results confirming Transformer generality. The best and second-best results are marked in red and underlined.

Table 5: Ablation studies of PSG components and calibration-set size on MambaIR classic SR $\times 2$.

(a) Component analysis

Pedestal	GRILL	Set5	B100	Manga109
		36.180	31.198	35.715
✓		36.572	31.300	36.096
	✓	37.013	31.634	36.941
✓	✓	37.423	31.796	37.613
+ Distillation		37.869	32.159	38.410

(b) Calibration-set size

N	2	4	**8**	16	32	64
PSNR	36.442	36.487	36.491	36.482	36.496	36.504
Std.	0.120	0.060	0.040	0.046	0.052	0.031
Cos.	0.982	0.990	0.996	0.997	0.998	0.999

Deployment efficiency. Table 4(a) reports the size and compute of the exported models. At 2 bits, PSG reduces the model size by up to $4.33\times$ and BOPs by up to $6.13\times$. Table 4(b) then deploys the models with a standard toolchain: each model is exported to ONNX, compiled by TensorRT into an engine, and run on an NVIDIA B200. The pedestal shift and its compensation are folded into constant terms before export, so the engine incurs no additional inference overhead. TensorRT provides no INT4 activation kernels, so the 4-bit models run as W4A16. Even in this weight-only regime, the engines reduce peak memory by up to $2.05\times$ and latency by up to $1.34\times$ over FP32.

5.3 Ablation Study

PSG components. Table 5(a) isolates the contribution of each component, starting from the calibration baseline adopted in prior PTQ methods (Liu et al., 2024). We conduct this study under 3-bit precision rather than 2-bit, because the baseline collapses at 2-bit and no longer serves as a meaningful reference for measuring the contribution of each component. Pedestal Shifting and GRILL each improve over the baseline, their gains are additive, and the combination performs best. PSG already achieves strong performance with gradient-free calibration alone, while the final distillation step provides an additional improvement.

Calibration set size. Table 5(b) examines calibration-set size N over five randomly sampled sets. Performance largely saturates at $N = 2$, while its variation remains low for $N \geq 8$, indicating robust pedestal estimation with few calibration samples. This robustness is also reflected in restoration quality: the pedestal estimate yields higher and more stable performance than channel-wise shifting under otherwise identical settings (Table 1(a)). We therefore use eight images for calibration.

Effect of Pedestal Shifting. Pedestal Shifting substantially contracts the activation range by removing the persistent directional offset before quantization. Under 2-bit MambaIR quantization, it reduces the pooled activation range by a median of $2.70\times$ (up to $6.13\times$) and the e_1-aligned quantization error by a median of $39.3\times$. These results verify both the range-reduction and sensitive-direction preservation effects of Pedestal Shifting, with similar behavior observed across backbones.

6 Conclusion

We showed that channel-wise outliers in image restoration share a persistent, quantization-sensitive component, the pedestal. PSG preserves this pedestal outside the quantizer with no additional inference overhead and fits the remaining bounds in closed form, which yields state-of-the-art low-bit restoration across tasks and architectures. Two limitations remain. Pedestal Shifting is applied only where one token-independent vector suffices, so token-varying sites rely on GRILL alone, and a token-adaptive pedestal with bounded runtime cost is a natural extension. The measured deployment is limited to weight-only INT4 engines, since standard toolchains lack INT4 activation and low-bit selective-scan kernels, so the compute reduction at 2 and 3 bits awaits such kernels.

AI use statement

In this work, we used generative AI tools to assist in the writing of proofs, namely the derivations of the closed-form updates, to provide feedback on the design and presentation of experiments, to assist with translation, and to assist in implementing methods. We did not use generative AI tools

to develop the conceptual framework or the proposed method, to formulate mathematical claims or provide critical ingredients for their proofs, to propose or refine hypotheses, or to help interpret results. Generating synthetic data, cleaning or reformatting datasets, and qualitative or thematic data analysis are not applicable to this work. Additionally, we used AI tools to edit parts of the paper for readability, to summarize and search relevant literature, to format references, and to edit software code. All AI-assisted derivations were checked by the authors, all AI-assisted code was tested by the authors, and all AI-assisted text was revised by the authors. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

Ethics statement

This work studies post-training quantization of image restoration models and uses only publicly available benchmark datasets. It involves no human subjects, personal data, or newly collected data. We do not foresee negative societal impacts of this work, and we adhere to the ICLR Code of Ethics.

Reproducibility statement

The quantization scope shared by all methods is specified in Sec. 5.1, and Appendix C gives the calibration data, the optimization settings, and the sites of Pedestal Shifting. The analyses of Sec. 3 are detailed in Appendices A and B, and the results for the remaining tasks, backbones, and bit widths are given in Appendix D. All baselines are run with their official implementations aligned with identical quantization scope. The inference code is provided in the anonymized repository linked in the abstract; the calibration code and quantized models will be released upon acceptance.

References

Ron Banner, Yury Nahshan, and Daniel Soudry. Post training 4-bit quantization of convolutional networks for rapid-deployment. In *Advances in Neural Information Processing Systems*, volume 32, pp. 7948–7956, 2019.

Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie-Line Alberi Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. In *Proceedings of the British Machine Vision Conference*, pp. 135.1–135.10, 2012.

Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. Pre-trained image processing transformer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12299–12310, 2021.

Yujie Chen, Haotong Qin, Zhang Zhang, Michele Magno, Luca Benini, and Yawei Li. Q-MambaIR: Accurate quantized Mamba for efficient image restoration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Findings*, pp. 2524–2533, 2026.

Zheng Chen, Kewei Zhang, Xiaoyang Liu, Weihang Zhang, Mengfan Wang, Yifan Fu, and Yulun Zhang. QuantDemoire: Quantization with outlier aware for image demoiréing. *arXiv preprint arXiv:2510.04066*, 2025.

Hung-Yueh Chiang, Chi-Chih Chang, Natalia Frumkin, Kai-Chiang Wu, Mohamed S. Abdelfattah, and Diana Marculescu. Quamba2: A robust and scalable post-training quantization framework for selective state space models. In *International Conference on Machine Learning*, pp. 10411–10427, 2025.

Younghyun Cho, Changhun Lee, Seonggon Kim, and Eunhyeok Park. PTQ4VM: Post-training quantization for visual Mamba. In *IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 1176–1185, 2025.

Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. GPT3.int8(): 8-bit matrix multiplication for transformers at scale. In *Advances in Neural Information Processing Systems*, volume 35, pp. 30318–30332, 2022.

Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *European Conference on Computer Vision*, pp. 184–199, 2014.

Steven K. Esser, Jeffrey L. McKinstry, Deepika Bablani, Rathinakumar Appuswamy, and Dharmendra S. Modha. Learned step size quantization. In *International Conference on Learning Representations*, 2020.

Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Pointwise shape-adaptive dct for high-quality denoising and deblocking of grayscale and color images. *IEEE Transactions on Image Processing*, 16(5):1395–1411, 2007.

Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. GPTQ: Accurate post-training quantization for generative pre-trained transformers. In *International Conference on Learning Representations*, 2023.

Rich Franzen. Kodak lossless true color image suite. `http://r0k.us/graphics/kodak/`, 1999.

Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. In *First Conference on Language Modeling*, 2024.

Hang Guo, Jinmin Li, Tao Dai, Zhihao Ouyang, Xudong Ren, and Shu-Tao Xia. MambaIR: A simple baseline for image restoration with state-space model. In *European Conference on Computer Vision*, pp. 222–241, 2024.

Hang Guo, Yong Guo, Yaohua Zha, Yulun Zhang, Wenbo Li, Tao Dai, Shu-Tao Xia, and Yawei Li. MambaIRv2: Attentive state space restoration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 28124–28133, 2025.

Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5197–5206, 2015.

Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2704–2713, 2018.

Boyun Li, Haiyu Zhao, Wenxin Wang, Peng Hu, Yuanbiao Gou, and Xi Peng. MaIR: A locality- and continuity-preserving Mamba for image restoration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7491–7501, 2025.

Yuhang Li, Ruihao Gong, Xu Tan, Yang Yang, Peng Hu, Qi Zhang, Fengwei Yu, Wei Wang, and Shi Gu. BRECQ: Pushing the limit of post-training quantization by block reconstruction. In *International Conference on Learning Representations*, 2021.

Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. SwinIR: Image restoration using Swin transformer. In *IEEE/CVF International Conference on Computer Vision Workshops*, pp. 1833–1844, 2021.

Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pp. 136–144, 2017.

Kai Liu, Haotong Qin, Yong Guo, Xin Yuan, Linghe Kong, Guihai Chen, and Yulun Zhang. 2DQuant: Low-bit post-training quantization for image super-resolution. In *Advances in Neural Information Processing Systems*, volume 37, pp. 71068–71084, 2024.

Stuart P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28 (2):129–137, 1982.

David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings of the IEEE International Conference on Computer Vision*, volume 2, pp. 416–423, 2001.

Yusuke Matsui, Kota Ito, Yuji Aramaki, Azuma Fujimoto, Toru Ogawa, Toshihiko Yamasaki, and Kiyoharu Aizawa. Sketch-based manga retrieval using Manga109 dataset. *Multimedia Tools and Applications*, 76(20):21811–21838, 2017.

Joel Max. Quantizing for minimum distortion. *IRE Transactions on Information Theory*, 6(1):7–12, 1960.

Markus Nagel, Rana Ali Amjad, Mart van Baalen, Christos Louizos, and Tijmen Blankevoort. Up or down? adaptive rounding for post-training quantization. In *International Conference on Machine Learning*, pp. 7197–7206, 2020.

Markus Nagel, Marios Fournarakis, Rana Ali Amjad, Yelysei Bondarenko, Mart van Baalen, and Tijmen Blankevoort. A white paper on neural network quantization. *arXiv preprint arXiv:2106.08295*, 2021.

Charbel Sakr, Steve Dai, Rangha Venkatesan, Brian Zimmer, William Dally, and Brucek Khailany. Optimal clipping and magnitude-aware differentiation for improved quantization-aware training. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162, pp. 19123–19138, 2022.

Wenqi Shao, Mengzhao Chen, Zhaoyang Zhang, Peng Xu, Lirui Zhao, Zhiqian Li, Kaipeng Zhang, Peng Gao, Yu Qiao, and Ping Luo. OmniQuant: Omnidirectionally calibrated quantization for large language models. In *International Conference on Learning Representations*, 2024.

Hamid R. Sheikh, Muhammad F. Sabir, and Alan C. Bovik. A statistical evaluation of recent full reference image quality assessment algorithms. *IEEE Transactions on Image Processing*, 15(11): 3440–3451, 2006.

Zhijun Tu, Jie Hu, Hanting Chen, and Yunhe Wang. Toward accurate post-training quantization for image super resolution. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5856–5865, 2023.

Hailing Wang, Jianglin Lu, Yitian Zhang, and Yun Fu. Outlier-aware post-training quantization for image super-resolution. In *IEEE/CVF International Conference on Computer Vision*, pp. 16175–16184, 2025.

Hongjun Wang, Jiyuan Chen, Xuan Song, and Yinqiang Zheng. HarmoQ: Harmonized post-training quantization for high-fidelity image super-resolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 9802–9810, 2026.

Xiuying Wei, Yunchen Zhang, Yuhang Li, Xiangguo Zhang, Ruihao Gong, Jinyang Guo, and Xianglong Liu. Outlier suppression+: Accurate quantization of large language models by equivalent and effective shifting and scaling. In *Conference on Empirical Methods in Natural Language Processing*, pp. 1648–1665, 2023.

Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. SmoothQuant: Accurate and efficient post-training quantization for large language models. In *International Conference on Machine Learning*, pp. 38087–38099, 2023.

Jingwei Xin, Wenhao Li, Nannan Wang, Jie Li, and Xinbo Gao. SPR^{2}Q: Static priority-based rectifier routing quantization for image super-resolution. In *International Conference on Learning Representations*, 2026.

Zukang Xu, Yuxuan Yue, Xing Hu, Dawei Yang, Zhihang Yuan, Zixu Jiang, Zhixuan Chen, Jiangyong Yu, Chen Xu, and Sifan Zhou. MambaQuant: Quantizing the Mamba family with variance aligned rotation methods. In *International Conference on Learning Representations*, 2025.

Zongsheng Yue, Jianyi Wang, and Chen Change Loy. ResShift: Efficient diffusion model for image super-resolution by residual shifting. In *Advances in Neural Information Processing Systems*, volume 36, pp. 13294–13307, 2023.

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5728–5739, 2022.

Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *Curves and Surfaces*, pp. 711–730. Springer, 2012.

Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017.

Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. In *IEEE/CVF International Conference on Computer Vision*, pp. 4791–4800, 2021.

Lei Zhang, Xiaolin Wu, Antoni Buades, and Xin Li. Color demosaicking by local directional interpolation and nonlocal adaptive thresholding. *Journal of Electronic Imaging*, 20(2):023016, 2011.

A Range-Targeted Clipping of Extreme Channels

Setup. We further examine range reduction by clipping extreme values across all activation sites subject to quantization in the full-precision MambaIR super-resolution $\times 2$ model. For each site, asymmetric thresholds are chosen to achieve a target range reduction while clipping as few values as possible. We report the total clipped fraction, median clipped-value concentration across activation sites, and the Set5 Y-channel PSNR change from the unclipped model.

Results. Range-targeted clipping removes values primarily from a few extreme channels and rapidly degrades restoration fidelity. At a $1.5\times$ range reduction, only 0.14% of all activation values are clipped, yet the most affected channel accounts for a median 48% of the removed values and the top five channels for 95%, while PSNR drops by 2.86 dB. Increasing the range reduction to $2\times$ clips only 0.45% of values but causes a 6.77 dB drop. These results show that directly truncating extreme channels is not a reliable way to reduce the per-tensor activation range.

Table 6: Range-targeted clipping on MambaIR SR $\times 2$.

Range reduction	Clipped	Clipped-value concentration		ΔPSNR
		Top channel	Top 5 channels	
$1.25\times$	0.04%	57%	100%	-0.60
$1.5\times$	0.14%	48%	95%	-2.86
$2\times$	0.45%	37%	85%	-6.77

B Token-Wise Functional Constancy

Site-wise and joint replacement. We further evaluate the token-wise functional constancy in Sec. 3.1 by replacing c_t with its token mean $\bar{c}$, as in Eq. 2, on the full-precision MambaIR $\times 2$ model. We consider both site-wise replacement and joint replacement across all sites passing the gate for Pedestal Shifting. Site-wise replacement changes PSNR by at most 0.003 dB, with a median absolute change of < 0.001 dB across sites, consistent with Fig. 1(b). When the same replacement is applied jointly to all gated sites, PSNR changes by only -0.014 dB, so the token-dependent variation along e_1 has little functional effect even when accumulated across the sites used by our method. Extending the replacement to all 108 candidate sites, including those rejected by the constancy gate, increases the degradation to -0.180 dB, supporting gated application based on token-wise variation.

Rank-k extension. We examine whether the pedestal can extend beyond e_1. Let $\{e_r\}_{r=1}^{k}$ denote the top-k eigenvectors of M, with $c_{t,r} = x_t^\top e_r$. We replace each $c_{t,r}$ with its token mean $\bar{c}_r$, extending Eq. 2 to the top-k directions. As shown in Table 7(b), only e_1 exhibits strong token-wise constancy: higher-order directions have median CVs of 64.7–74.6, compared with 0.06 for e_1. Accordingly, joint replacement increasingly degrades restoration quality as k increases. We therefore restrict the pedestal to the single dominant direction e_1.

Table 7: Token-wise functional constancy (Set5). (a) Replacement scope. (b) Rank-k extension.

(a) Replacement scope

Scope	ΔPSNR (dB)
Site-wise (median $\|\Delta\|$)	< 0.001
Site-wise (max $\|\Delta\|$)	0.003
Joint (gated sites)	-0.014
Joint (all candidate sites)	-0.180

(b) Rank-k extension

k	CV_k (median)	ΔPSNR (dB)
1	0.06	-0.014
2	68.0	-0.174
3	64.7	-0.301
4	74.6	-0.403

C Implementation Details

Calibration and optimization data. Calibration images and distillation inputs are drawn from the training data of each task, without ground truth: DF2K for super-resolution on all four backbones, DF2K with synthetic Gaussian noise for color denoising, and DFWB for JPEG CAR. For real-world SR, they are drawn from the 45 images of RealSet65 (Yue et al., 2023) that do not overlap with the test set RealSRSet (Zhang et al., 2021). Distillation uses only degraded inputs, since the full-precision model serves as the teacher, and no ground truth is used at any stage. The eight calibration images are sampled once per task and reused for all bit widths for PSG.

Optimization. Only the quantization bounds are optimized, with all network weights frozen. We use Adam with a learning rate of 10^{-2}, 4,000 iterations, an effective batch size of 16, and 64×64 degraded input patches. The loss is the L_1 distance between the outputs of the quantized and full-precision models plus a feature-matching term on the normalized outputs of each residual group, following 2DQuant. The same setting is used for every backbone, task, and bit width.

Pedestal Shifting sites. A site is the input of a block-internal linear or convolutional layer whose producer is a static operator, so that the shift $-p$ folds into the producer and the compensation Wp into the bias of the consumer. The constancy gate with $\tau = 0.1$ passes 50 of the 108 candidate sites on MambaIR, 44 of 160 on MambaIRv2-light, and 56 of 108 on MaIR. For a linear consumer the fold is exact. For a zero-padded $k \times k$ convolution the compensation deviates from Wp on the $(k-1)/2$-pixel border, and we fold the interior value only. We verified the fold on the full-precision MaIR $\times 2$ model over the 328 images of the five SR benchmarks: the folded network changes PSNR by less than 0.001 dB, with similar behavior observed across backbones. An exact fold is also available at negligible overhead: adding the eight border constants per channel at the convolutional sites reproduces the full-precision output to 6×10^{-6} with an extra constant addition on the border, about 3×10^{-7} of the network's MACs. We nevertheless keep the parameter-only fold, which adds no operation at inference.

D Additional Results

This section complements the main results with the remaining scales, bit widths, and backbones for SR and Gaussian color denoising, extends the comparison to a Transformer backbone, and adds qualitative comparisons. All experiments follow the setup of the main text, and the best and second-best results are marked in red and underlined.

D.1 Super-Resolution

Table 8 reports classic SR on MambaIR at $\times 2$, $\times 3$, and $\times 4$, and Table 9 reports classic SR on MaIR at $\times 2$ together with lightweight SR on MambaIRv2-light at $\times 2$ and $\times 4$, complementing the $\times 3$ results of Table 2, all at W4A4, W3A3, and W2A2. On MambaIR and MaIR, PSG achieves the best PSNR and SSIM on every benchmark, scale, and bit width. At W2A2 on MambaIR, plain channel-wise shifting (OS+) becomes the strongest baseline, and PSG still improves over it on every benchmark, which is consistent with the estimator comparison in Table 1(a). On MambaIRv2-light, PSG leads on every benchmark and bit width with a single exception. At $\times 2$ W4A4, SPR^{2}Q is ahead on Set5 and Manga109, and the two methods are otherwise closely matched.

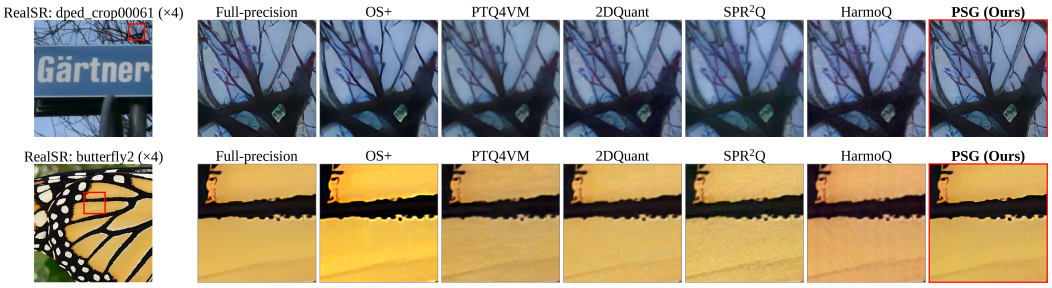


Figure 4: Qualitative comparison on MambaIR for real-world SR under 3-bit (top) and 2-bit (bottom). No ground truth is available, so the full-precision output serves as the reference.

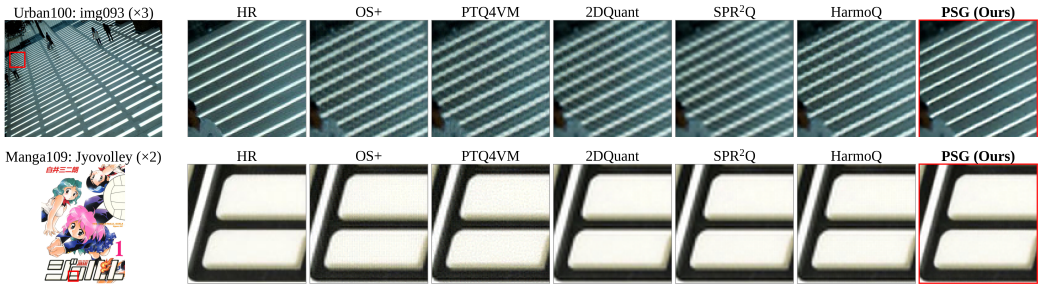


Figure 5: Qualitative comparison on MambaIRv2-light for lightweight SR under 3-bit at $\times 3$ (top) and 2-bit at $\times 2$ (bottom).

D.2 Gaussian Color Denoising

Table 10 reports denoising on MaIR at $\sigma = 15$ for all bit widths and adds the W4A4 results on MambaIR at $\sigma = 25$ to Table 3. PSG achieves the best PSNR and SSIM on every benchmark and bit width, and on MaIR its W2A2 results exceed the W4A4 results of every baseline.

D.3 Generalization to Transformers

Tables 11 and 12 evaluate PSG on SwinIR at W2A2 for classic SR, Gaussian color denoising, and JPEG CAR. HarmoQ is the strongest baseline for SR and SPR^{2}Q for denoising and JPEG CAR, and PSG achieves the best PSNR and SSIM on every benchmark. These results confirm that the pedestal structure and its treatment are not specific to Mamba backbones.

D.4 Qualitative Results

Figure 4 shows real-world SR on MambaIR at W3A3 and W2A2, where no ground truth exists and the full-precision output serves as the reference. Figure 5 shows lightweight SR on MambaIRv2-light at $\times 3$ W3A3 and $\times 2$ W2A2, corresponding to Tables 2 and 9. In Fig. 4, PSG preserves the thin branches that the baselines blur or smear (top) and keeps both the edges and the colors of the wing pattern, where several baselines introduce visible color shifts (bottom). In Fig. 5, the baselines produce jagged stripes on the crosswalk, whereas PSG reconstructs them cleanly (top), and PSG yields cleaner edges on the line art (bottom). In both figures, the outputs of PSG remain closest to the reference.

Table 8: Quantitative comparison on MambaIR for classic SR $\times 2$, $\times 3$, and $\times 4$.

Method	Bit	Set5		Set14		B100		Urban100		Manga109	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
MambaIR ($\times 2$)	32/32	38.574	0.9627	34.672	0.9261	32.583	0.9048	34.151	0.9446	40.277	0.9806
OS+	4/4	37.483	0.9573	33.198	0.9144	31.843	0.8946	31.463	0.9219	37.641	0.9735
PTQ4VM	4/4	37.781	0.9582	33.438	0.9157	32.049	0.8969	31.953	0.9260	38.230	0.9749
2DQuant	4/4	37.166	0.9553	33.006	0.9110	31.665	0.8913	30.739	0.9120	37.366	0.9717
SPR^{2}Q	4/4	36.794	0.9543	32.515	0.9068	31.317	0.8870	29.458	0.8940	35.942	0.9681
HarmoQ	4/4	37.742	0.9589	33.466	0.9175	32.048	0.8974	31.912	0.9266	38.265	0.9756
PSG (Ours)	4/4	**38.202**	**0.9607**	**34.104**	**0.9225**	**32.370**	**0.9018**	**33.225**	**0.9388**	**39.191**	**0.9782**
OS+	3/3	36.408	0.9520	32.360	0.9055	31.209	0.8861	29.720	0.9005	35.626	0.9651
PTQ4VM	3/3	36.267	0.9507	32.220	0.9043	31.104	0.8839	29.109	0.8901	35.204	0.9628
2DQuant	3/3	36.794	0.9539	32.551	0.9069	31.355	0.8871	29.657	0.8974	36.235	0.9684
SPR^{2}Q	3/3	36.961	0.9551	32.659	0.9087	31.444	0.8886	30.015	0.9025	36.657	0.9706
HarmoQ	3/3	36.881	0.9545	32.663	0.9091	31.446	0.8892	30.197	0.9064	36.659	0.9700
PSG (Ours)	3/3	**37.869**	**0.9595**	**33.665**	**0.9190**	**32.159**	**0.8989**	**32.269**	**0.9314**	**38.410**	**0.9765**
OS+	2/2	35.270	0.9450	31.622	0.8977	30.624	0.8775	28.277	0.8771	33.525	0.9549
PTQ4VM	2/2	34.812	0.9317	31.368	0.8864	30.430	0.8677	28.072	0.8623	33.059	0.9360
2DQuant	2/2	34.670	0.9405	31.189	0.8930	30.334	0.8725	27.697	0.8664	32.258	0.9480
SPR^{2}Q	2/2	34.820	0.9433	31.247	0.8948	30.374	0.8743	27.715	0.8684	32.332	0.9510
HarmoQ	2/2	34.977	0.9427	31.387	0.8956	30.477	0.8753	27.903	0.8705	32.665	0.9502
PSG (Ours)	2/2	**37.333**	**0.9574**	**32.925**	**0.9120**	**31.674**	**0.8925**	**30.585**	**0.9131**	**37.304**	**0.9735**
MambaIR ($\times 3$)	32/32	35.084	0.9323	30.988	0.8536	29.508	0.8157	29.934	0.8841	35.430	0.9546
OS+	4/4	33.845	0.9198	30.004	0.8348	28.861	0.7971	27.895	0.8455	32.854	0.9366
PTQ4VM	4/4	34.224	0.9240	30.178	0.8377	28.968	0.8002	27.948	0.8473	32.934	0.9386
2DQuant	4/4	33.731	0.9185	29.886	0.8313	28.794	0.7940	27.623	0.8377	32.478	0.9331
SPR^{2}Q	4/4	33.167	0.9087	29.675	0.8292	28.706	0.7933	27.668	0.8402	32.201	0.9295
HarmoQ	4/4	33.800	0.9196	29.994	0.8346	28.841	0.7966	27.673	0.8410	32.617	0.9354
PSG (Ours)	4/4	**34.565**	**0.9276**	**30.510**	**0.8464**	**29.250**	**0.8086**	**28.749**	**0.8667**	**33.909**	**0.9465**
OS+	3/3	32.669	0.9042	29.219	0.8186	28.280	0.7815	26.399	0.8049	30.491	0.9109
PTQ4VM	3/3	32.715	0.9061	29.174	0.8177	28.270	0.7813	26.085	0.7957	30.157	0.9060
2DQuant	3/3	33.172	0.9111	29.511	0.8233	28.496	0.7856	26.728	0.8129	31.492	0.9219
SPR^{2}Q	3/3	32.887	0.9094	29.362	0.8208	28.371	0.7831	26.344	0.8027	30.868	0.9165
HarmoQ	3/3	33.133	0.9114	29.514	0.8246	28.492	0.7878	26.642	0.8136	31.372	0.9218
PSG (Ours)	3/3	**34.199**	**0.9237**	**30.211**	**0.8402**	**29.020**	**0.8025**	**28.030**	**0.8520**	**33.226**	**0.9413**
OS+	2/2	31.257	0.8783	28.350	0.7981	27.696	0.7626	25.155	0.7615	28.235	0.8725
PTQ4VM	2/2	31.136	0.8670	28.272	0.7887	27.644	0.7558	25.150	0.7511	28.190	0.8520
2DQuant	2/2	31.100	0.8834	28.181	0.8004	27.634	0.7657	24.920	0.7602	27.730	0.8746
SPR^{2}Q	2/2	31.059	0.8827	28.157	0.7999	27.617	0.7652	24.899	0.7596	27.691	0.8738
HarmoQ	2/2	31.029	0.8780	28.186	0.7972	27.603	0.7624	24.965	0.7593	27.846	0.8711
PSG (Ours)	2/2	**33.054**	**0.9124**	**29.431**	**0.8255**	**28.472**	**0.7878**	**26.544**	**0.8134**	**31.097**	**0.9221**
MambaIR ($\times 4$)	32/32	33.030	0.9046	29.202	0.7961	27.979	0.7503	27.681	0.8287	32.317	0.9272
OS+	4/4	31.547	0.8811	28.214	0.7703	27.309	0.7248	25.676	0.7692	29.617	0.8927
PTQ4VM	4/4	32.091	0.8917	28.475	0.7772	27.542	0.7333	26.006	0.7818	30.105	0.9017
2DQuant	4/4	31.616	0.8826	28.235	0.7693	27.296	0.7224	25.650	0.7659	29.487	0.8899
SPR^{2}Q	4/4	31.716	0.8855	28.270	0.7725	27.363	0.7267	25.828	0.7755	29.590	0.8923
HarmoQ	4/4	31.679	0.8856	28.279	0.7714	27.345	0.7261	25.709	0.7709	29.702	0.8957
PSG (Ours)	4/4	**32.436**	**0.8965**	**28.784**	**0.7863**	**27.736**	**0.7412**	**26.722**	**0.8062**	**30.929**	**0.9123**
OS+	3/3	30.167	0.8507	27.385	0.7464	26.786	0.7049	24.532	0.7224	27.446	0.8472
PTQ4VM	3/3	30.695	0.8643	27.529	0.7514	26.892	0.7096	24.494	0.7214	27.603	0.8529
2DQuant	3/3	30.746	0.8661	27.650	0.7549	26.954	0.7112	24.752	0.7312	28.128	0.8634
SPR^{2}Q	3/3	30.416	0.8596	27.451	0.7504	26.839	0.7074	24.514	0.7224	27.355	0.8506
HarmoQ	3/3	30.756	0.8685	27.740	0.7583	26.962	0.7128	24.763	0.7351	28.272	0.8702
PSG (Ours)	3/3	**31.591**	**0.8833**	**28.271**	**0.7732**	**27.340**	**0.7276**	**25.670**	**0.7726**	**29.607**	**0.8921**
OS+	2/2	29.332	0.8293	26.785	0.7301	26.382	0.6888	23.773	0.6874	26.018	0.8116
PTQ4VM	2/2	29.140	0.8110	26.630	0.7145	26.302	0.6797	23.682	0.6686	25.866	0.7819
2DQuant	2/2	28.988	0.8258	26.488	0.7259	26.267	0.6893	23.472	0.6794	25.439	0.8046
SPR^{2}Q	2/2	29.027	0.8266	26.513	0.7266	26.280	0.6898	23.496	0.6804	25.479	0.8054
HarmoQ	2/2	28.769	0.8042	26.397	0.7107	26.139	0.6705	23.442	0.6639	25.435	0.7861
PSG (Ours)	2/2	**30.384**	**0.8614**	**27.445**	**0.7540**	**26.840**	**0.7102**	**24.489**	**0.7263**	**27.311**	**0.8547**

Table 9: Quantitative comparison for lightweight SR on MambaIRv2-light and classic SR on MaIR.

Method	Bit	Set5		Set14		B100		Urban100		Manga109	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
MambaIRv2-light (×2)	32/32	38.257	0.9615	34.079	0.9221	32.363	0.9019	33.257	0.9378	39.355	0.9785
OS+	4/4	36.354	0.9455	32.325	0.8997	31.133	0.8802	29.537	0.8913	35.535	0.9586
PTQ4VM	4/4	35.711	0.9421	31.821	0.8959	30.723	0.8751	28.935	0.8813	34.848	0.9548
2DQuant	4/4	37.354	0.9564	33.075	0.9121	31.743	0.8931	31.181	0.9182	37.471	0.9725
SPR^{2}Q	4/4	**37.600**	**0.9586**	33.241	0.9149	31.888	**0.8957**	31.383	0.9214	**37.951**	**0.9749**
HarmoQ	4/4	37.397	0.9563	33.097	0.9121	31.779	0.8932	31.171	0.9173	37.599	0.9723
PSG (Ours)	4/4	37.507	0.9579	**33.281**	**0.9153**	**31.894**	0.8955	**31.485**	**0.9223**	37.500	0.9739
OS+	3/3	35.266	0.9341	31.687	0.8889	30.668	0.8692	28.500	0.8697	33.891	0.9435
PTQ4VM	3/3	34.459	0.9186	31.111	0.8703	30.063	0.8474	27.852	0.8341	32.855	0.9088
2DQuant	3/3	36.444	0.9513	32.323	0.9042	31.172	0.8843	29.224	0.8914	35.558	0.9642
SPR^{2}Q	3/3	36.896	0.9552	32.590	0.9082	31.423	0.8886	29.734	0.9000	36.247	0.9689
HarmoQ	3/3	36.337	0.9500	32.202	0.9027	31.082	0.8830	29.057	0.8883	35.354	0.9621
PSG (Ours)	3/3	**37.092**	**0.9558**	**32.743**	**0.9099**	**31.518**	**0.8903**	**30.077**	**0.9058**	**36.883**	**0.9708**
OS+	2/2	32.760	0.8692	30.297	0.8363	29.394	0.8027	27.272	0.8064	31.594	0.8865
PTQ4VM	2/2	32.480	0.8612	29.990	0.8214	29.172	0.7943	27.113	0.7975	31.140	0.8632
2DQuant	2/2	34.633	0.9413	31.146	0.8932	30.295	0.8723	27.640	0.8666	32.138	0.9489
SPR^{2}Q	2/2	34.661	0.9414	31.164	0.8934	30.311	0.8727	27.655	0.8667	32.171	0.9486
HarmoQ	2/2	34.638	0.9374	31.187	0.8911	30.352	0.8718	27.730	0.8656	32.287	0.9444
PSG (Ours)	2/2	**36.272**	**0.9514**	**32.155**	**0.9040**	**31.090**	**0.8845**	**28.887**	**0.8874**	**35.069**	**0.9631**
MambaIRv2-light (×4)	32/32	32.501	0.8992	28.844	0.7878	27.748	0.7426	26.818	0.8079	31.243	0.9182
OS+	4/4	30.761	0.8617	27.599	0.7488	26.832	0.7051	24.719	0.7279	28.211	0.8616
PTQ4VM	4/4	30.094	0.8515	27.054	0.7369	26.584	0.6989	24.108	0.7049	27.211	0.8448
2DQuant	4/4	31.162	0.8764	28.002	0.7649	27.153	0.7205	25.242	0.7553	28.868	0.8834
SPR^{2}Q	4/4	31.352	0.8805	28.113	0.7691	27.222	0.7247	25.352	0.7621	29.148	0.8887
HarmoQ	4/4	31.325	0.8776	28.010	0.7651	27.182	0.7207	25.335	0.7586	29.113	0.8864
PSG (Ours)	4/4	**31.506**	**0.8829**	**28.188**	**0.7709**	**27.266**	**0.7259**	**25.482**	**0.7671**	**29.279**	**0.8920**
OS+	3/3	29.760	0.8304	27.043	0.7287	26.532	0.6889	24.050	0.6935	26.616	0.8165
PTQ4VM	3/3	29.279	0.8117	26.559	0.6960	26.240	0.6651	23.653	0.6491	26.069	0.7792
2DQuant	3/3	29.373	0.8329	26.740	0.7301	26.403	0.6917	23.708	0.6873	25.800	0.8104
SPR^{2}Q	3/3	28.742	0.8189	26.337	0.7178	26.148	0.6809	23.352	0.6715	25.230	0.7975
HarmoQ	3/3	28.903	0.8113	26.479	0.7163	26.240	0.6799	23.515	0.6711	25.507	0.7932
PSG (Ours)	3/3	**30.757**	**0.8683**	**27.701**	**0.7576**	**26.969**	**0.7142**	**24.806**	**0.7383**	**28.028**	**0.8670**
OS+	2/2	28.400	0.7681	26.136	0.6726	25.931	0.6385	23.320	0.6235	25.191	0.7238
PTQ4VM	2/2	28.146	0.7419	25.951	0.6507	25.760	0.6184	23.191	0.6040	24.996	0.6977
2DQuant	2/2	28.785	0.8204	26.358	0.7190	26.167	0.6823	23.369	0.6728	25.254	0.7990
SPR^{2}Q	2/2	28.734	0.8193	26.329	0.7178	26.143	0.6809	23.343	0.6715	25.219	0.7982
HarmoQ	2/2	28.628	0.7970	26.287	0.7051	26.104	0.6710	23.354	0.6596	25.265	0.7768
PSG (Ours)	2/2	**29.888**	**0.8477**	**27.068**	**0.7417**	**26.622**	**0.7020**	**24.033**	**0.7043**	**26.494**	**0.8305**
MaIR (×2)	32/32	38.526	0.9626	34.740	0.9267	32.577	0.9048	34.153	0.9448	40.272	0.9807
OS+	4/4	37.031	0.9560	32.789	0.9097	31.464	0.8899	30.472	0.9100	36.537	0.9711
PTQ4VM	4/4	37.919	0.9597	33.537	0.9174	32.040	0.8972	31.779	0.9241	38.596	0.9765
2DQuant	4/4	37.141	0.9558	32.771	0.9090	31.517	0.8897	30.151	0.9045	36.908	0.9713
SPR^{2}Q	4/4	36.997	0.9557	32.711	0.9089	31.456	0.8892	29.787	0.8996	36.526	0.9703
HarmoQ	4/4	37.708	0.9591	33.452	0.9169	31.982	0.8969	31.810	0.9251	38.343	0.9759
PSG (Ours)	4/4	**38.246**	**0.9612**	**34.203**	**0.9241**	**32.356**	**0.9021**	**33.153**	**0.9390**	**39.304**	**0.9787**
OS+	3/3	35.066	0.9362	31.510	0.8843	30.397	0.8618	28.539	0.8688	33.961	0.9488
PTQ4VM	3/3	36.974	0.9549	32.573	0.9076	31.369	0.8878	29.771	0.8994	36.176	0.9681
2DQuant	3/3	36.204	0.9505	32.143	0.9020	31.021	0.8819	28.850	0.8848	34.768	0.9619
SPR^{2}Q	3/3	36.583	0.9532	32.408	0.9059	31.242	0.8861	29.265	0.8922	35.756	0.9667
HarmoQ	3/3	36.402	0.9522	32.264	0.9050	31.120	0.8847	29.199	0.8913	35.522	0.9662
PSG (Ours)	3/3	**38.004**	**0.9603**	**33.864**	**0.9216**	**32.211**	**0.9002**	**32.486**	**0.9336**	**38.897**	**0.9774**
OS+	2/2	33.927	0.9144	30.942	0.8717	30.008	0.8502	27.849	0.8528	32.761	0.9273
PTQ4VM	2/2	35.273	0.9386	31.626	0.8922	30.630	0.8749	28.366	0.8737	33.773	0.9481
2DQuant	2/2	34.738	0.9431	31.195	0.8946	30.337	0.8740	27.672	0.8684	32.192	0.9508
SPR^{2}Q	2/2	34.860	0.9437	31.262	0.8953	30.386	0.8746	27.726	0.8689	32.331	0.9514
HarmoQ	2/2	34.931	0.9421	31.393	0.8946	30.463	0.8744	27.958	0.8711	32.845	0.9501
PSG (Ours)	2/2	**37.634**	**0.9588**	**33.196**	**0.9151**	**31.916**	**0.8963**	**31.211**	**0.9209**	**37.878**	**0.9751**

Table 10: Quantitative comparison for Gaussian color denoising on MaIR and MambaIR.

Method	Bit	BSD68		Kodak24		McMaster		Urban100	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
MaIR ($\sigma = 15$)	32/32	34.479	0.9361	35.533	0.9310	35.706	0.9356	35.348	0.9532
OS+	4/4	31.250	0.8575	31.705	0.8361	31.790	0.8394	30.530	0.8672
PTQ4VM	4/4	32.988	0.9067	33.491	0.8933	33.299	0.8880	32.467	0.9154
2DQuant	4/4	32.753	0.9023	33.437	0.8912	33.072	0.8851	32.300	0.9118
SPR^{2}Q	4/4	33.334	0.9160	33.922	0.9048	33.526	0.8972	32.788	0.9231
HarmoQ	4/4	33.179	0.9122	33.869	0.9018	33.578	0.8968	32.895	0.9236
PSG (Ours)	4/4	**34.214**	**0.9325**	**35.190**	**0.9267**	**35.224**	**0.9279**	**34.667**	**0.9477**
OS+	3/3	30.436	0.8216	30.783	0.7966	30.826	0.7924	29.402	0.8268
PTQ4VM	3/3	31.942	0.8755	32.208	0.8501	32.076	0.8452	30.943	0.8794
2DQuant	3/3	31.020	0.8514	31.402	0.8287	31.335	0.8221	29.624	0.8521
SPR^{2}Q	3/3	33.046	0.9115	33.515	0.8979	32.978	0.8873	32.148	0.9149
HarmoQ	3/3	31.412	0.8608	31.904	0.8421	31.736	0.8377	30.383	0.8661
PSG (Ours)	3/3	**34.074**	**0.9307**	**34.991**	**0.9241**	**34.927**	**0.9230**	**34.259**	**0.9445**
OS+	2/2	30.021	0.8001	30.219	0.7666	30.356	0.7662	29.054	0.8068
PTQ4VM	2/2	30.091	0.8328	30.533	0.8128	31.038	0.8168	28.547	0.8352
2DQuant	2/2	30.491	0.8245	30.774	0.7969	30.642	0.7851	29.095	0.8234
SPR^{2}Q	2/2	32.218	0.8894	32.661	0.8742	32.020	0.8594	30.927	0.8911
HarmoQ	2/2	31.404	0.8682	31.988	0.8555	31.909	0.8523	30.241	0.8758
PSG (Ours)	2/2	**33.760**	**0.9262**	**34.543**	**0.9181**	**34.326**	**0.9132**	**33.493**	**0.9368**
MambaIR ($\sigma = 25$)	32/32	31.860	0.8951	32.992	0.8942	33.433	0.9074	33.210	0.9330
OS+	4/4	29.817	0.8258	30.414	0.8113	30.422	0.8062	29.691	0.8511
PTQ4VM	4/4	30.057	0.8357	30.642	0.8236	30.603	0.8213	29.394	0.8456
2DQuant	4/4	31.428	0.8843	32.399	0.8810	32.655	0.8895	32.107	0.9175
SPR^{2}Q	4/4	30.702	0.8640	31.403	0.8550	31.404	0.8580	30.391	0.8860
HarmoQ	4/4	30.386	0.8496	31.057	0.8388	31.022	0.8427	30.444	0.8822
PSG (Ours)	4/4	**31.638**	**0.8908**	**32.691**	**0.8887**	**33.026**	**0.8983**	**32.560**	**0.9252**

Table 11: Quantitative comparison on SwinIR for classic SR at W2A2.

Method	Bit	Set5		Set14		B100		Urban100		Manga109	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
SwinIR ($\times 2$)	32/32	38.423	0.9623	34.462	0.9250	32.530	0.9041	33.831	0.9427	39.908	0.9796
OS+	2/2	33.747	0.9253	30.405	0.8754	30.012	0.8624	27.162	0.8476	29.757	0.8996
2DQuant	2/2	36.092	0.9509	32.119	0.9034	30.997	0.8832	28.990	0.8889	34.824	0.9624
SPR^{2}Q	2/2	36.433	0.9528	32.380	0.9062	31.221	0.8860	29.487	0.8965	35.631	0.9661
HarmoQ	2/2	37.008	0.9555	32.736	0.9100	31.518	0.8903	30.318	0.9079	36.560	0.9690
PSG (Ours)	2/2	**37.688**	**0.9591**	**33.330**	**0.9166**	**31.985**	**0.8973**	**31.467**	**0.9236**	**37.768**	**0.9753**
SwinIR ($\times 3$)	32/32	34.953	0.9318	30.929	0.8534	29.452	0.8145	29.751	0.8826	35.121	0.9536
OS+	2/2	26.335	0.6992	24.875	0.6634	26.335	0.6888	23.085	0.6507	21.826	0.6513
2DQuant	2/2	31.569	0.8879	28.550	0.8051	27.846	0.7703	25.300	0.7705	28.462	0.8774
SPR^{2}Q	2/2	31.872	0.8940	28.726	0.8095	27.969	0.7739	25.523	0.7787	28.946	0.8895
HarmoQ	2/2	32.492	0.9043	29.046	0.8155	28.216	0.7804	26.070	0.7970	29.836	0.8993
PSG (Ours)	2/2	**33.811**	**0.9215**	**30.000**	**0.8365**	**28.859**	**0.7988**	**27.483**	**0.8391**	**32.314**	**0.9352**
SwinIR ($\times 4$)	32/32	32.928	0.9043	29.086	0.7949	27.925	0.7490	27.455	0.8255	32.036	0.9260
OS+	2/2	23.349	0.6178	22.290	0.5634	24.316	0.5978	20.468	0.5215	19.000	0.5150
2DQuant	2/2	29.493	0.8355	26.825	0.7320	26.462	0.6941	23.785	0.6897	25.988	0.8119
SPR^{2}Q	2/2	29.471	0.8358	26.806	0.7321	26.458	0.6939	23.768	0.6896	25.968	0.8127
HarmoQ	2/2	30.590	0.8636	27.463	0.7529	26.899	0.7119	24.643	0.7327	27.527	0.8535
PSG (Ours)	2/2	**31.533**	**0.8842**	**28.196**	**0.7727**	**27.338**	**0.7278**	**25.530**	**0.7684**	**29.044**	**0.8884**

Table 12: Quantitative comparison on SwinIR for Gaussian color denoising ($\sigma = 25$, BSD68 to Urban100) and JPEG CAR ($q = 10$, Classic5 and LIVE1).

Method	Bit	BSD68		Kodak24		McMaster		Urban100		Classic5		LIVE1	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
SwinIR	32/32	31.778	0.8940	32.889	0.8927	33.312	0.9057	32.921	0.9301	30.267	0.8249	29.855	0.8287
OS+	2/2	21.543	0.6350	22.561	0.6596	21.797	0.6571	19.933	0.6150	26.929	0.7247	26.933	0.7354
2DQuant	2/2	29.242	0.8138	29.767	0.8021	29.225	0.7843	28.108	0.8257	29.080	0.7917	28.756	0.8004
SPR^{2}Q	2/2	30.333	0.8552	31.023	0.8468	30.744	0.8389	29.603	0.8746	29.342	0.8005	29.090	0.8087
HarmoQ	2/2	29.391	0.8316	29.939	0.8276	29.222	0.8029	27.998	0.8377	28.954	0.7905	28.664	0.8003
PSG (Ours)	2/2	**31.043**	**0.8759**	**31.897**	**0.8701**	**31.978**	**0.8740**	**30.895**	**0.9022**	**29.703**	**0.8116**	**29.416**	**0.8182**