HarmoQ: Harmonized Post-Training Quantization for High-Fidelity Image Super-Resolution

Hongjun Wang[1,2], Jiyuan Chen[2,3], Xuan Song[4*], Yinqiang Zheng[1*]

[1]The University of Tokyo
[2]Southern University of Science and Technology
[3]The Hong Kong Polytechnic University
[4]Jilin University

Abstract

Post-training quantization offers an efficient pathway to deploy super-resolution models, yet existing methods treat weight and activation quantization independently, missing their critical interplay. Through controlled experiments on SwinIR, we uncover a striking asymmetry: weight quantization primarily degrades structural similarity, while activation quantization disproportionately affects pixel-level accuracy. This stems from their distinct roles—weights encode learned restoration priors for textures and edges, whereas activations carry input-specific intensity information. Building on this insight, we propose HarmoQ, a unified framework that harmonizes quantization across components through three synergistic steps: structural residual calibration proactively adjusts weights to compensate for activation-induced detail loss, harmonized scale optimization analytically balances quantization difficulty via closed-form solutions, and adaptive boundary refinement iteratively maintains this balance during optimization. Experiments show HarmoQ achieves substantial gains under aggressive compression, outperforming prior art by 0.46 dB on Set5 at 2-bit while delivering 3.2× speedup and 4× memory reduction on A100 GPUs. This work provides the first systematic analysis of weight-activation coupling in super-resolution quantization and establishes a principled solution for efficient high-quality image restoration.

Code — https://github.com/Dreamzz5/HarmoQ

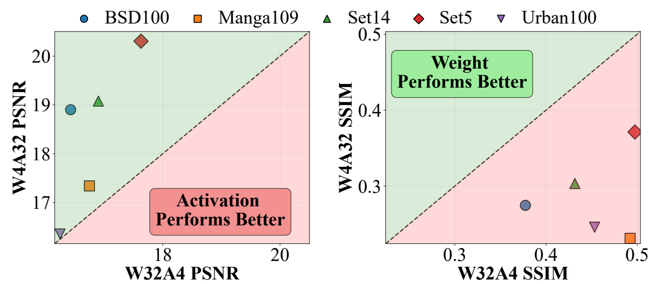

(a) PSNR and SSIM under weight and activation quantization

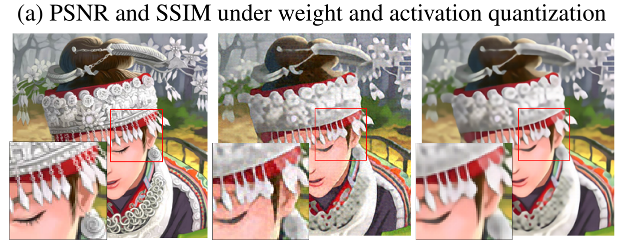

(b) Ground Truth (c) W32A4 result (d) W4 A32 result

Figure 1: Weight vs. activation quantization analysis. (a) Performance comparison showing that W32A4 (activation quantization) primarily degrades PSNR while W4A32 (weight quantization) shows stronger SSIM degradation, indicating complementary effects on pixel-level accuracy versus structural similarity. (b, c) Visual comparison of W32A4 and W4A32 quantization results on anime image reconstruction.

Introduction

Single-image super-resolution (SISR) has witnessed remarkable progress driven by deep neural networks, especially Transformer-based architectures such as SwinIR (Liang et al. 2021) and Restormer (Zamir et al. 2022). These models achieve impressive reconstruction accuracy but come at a steep cost: their high parameter counts and computational demands pose major barriers to deployment in real-world scenarios, particularly in edge or mobile environments.

Post-training quantization (PTQ) offers an appealing solution to compress pretrained models without retraining. However, applying PTQ to super-resolution remains an open challenge. Prior works diverge on how to balance quantization: some target both weights and activations (Liu et al. 2024), while others focus exclusively on activations (Wang et al. 2025). Yet, the relative impact of these two components on visual fidelity remains poorly understood.

To systematically investigate this trade-off, we conduct a controlled comparison between two complementary 4-bit quantization regimes—W4A32 (quantized weights, full-precision activations) and W32A4 (full-precision weights, quantized activations)—on SwinIR across standard benchmarks. As shown in Figure 1, a clear pattern emerges: **weight quantization (W4A32)** has a stronger effect on SSIM (drops to 0.85 on Manga109), indicating corruption of fine-grained textures and edge structures, while **ac-**

[*]Corresponding Author

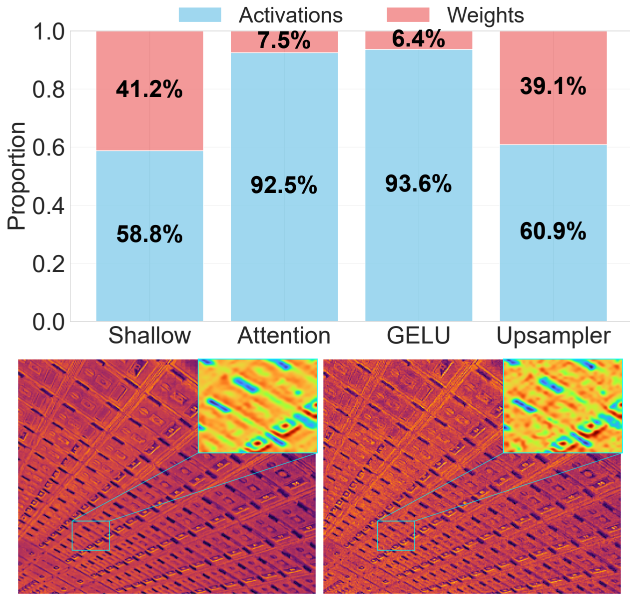

(a) W32A4 Attention Map (b) W4A32 Attention Map

Figure 2: Layer-wise quantization sensitivity analysis in SwinIR. Stacked bar chart shows the proportion of sensitivity to weight versus activation quantization across four layer types. Attention and GELU layers are highly sensitive to activation quantization (92.5% and 93.6%), while Shallow layers show balanced sensitivity. Attention maps compare (a) W32A4 and (b) W4A32 configurations, illustrating differential quantization impacts in attention layer.

tivation quantization (W32A4) disproportionately affects PSNR (drops to 32.5 dB on Urban100), reflecting degradation of pixel-level accuracy and global attributes such as brightness and contrast.

> **Key Observation:** Activation quantization contributes disproportionately more to overall quantization losses than weight quantization, especially in attention and nonlinear layers.

This asymmetry arises from the distinct functional roles of weights and activations. Weights encode distributed, global restoration priors learned during training; their quantization introduces systematic errors that primarily corrupt representations of fine textures and edges. In contrast, activations carry local, input-specific intensity information; their quantization introduces stochastic noise that disproportionately affects pixel-level accuracy, degrading global attributes like brightness uniformity and overall contrast.

Despite these findings, most existing quantization methods either treat weight and activation quantization independently, as in DBDC (Tu et al. 2023) and 2DQuant (Liu et al. 2024), or focus on a single component, as in dynamic activation-only quantization (Wang et al. 2025). These methods overlook the intrinsic coupling between weights and activations. Classical PTQ techniques such as MinMax (Jacob et al. 2018) or Percentile (Li et al. 2019) further exacerbate this issue by applying uniform clipping or scaling heuristics that ignore signal content related to structure and texture, leading to suboptimal quantization behavior.

To address these challenges, we introduce HarmoQ, a harmonized post-training quantization framework designed specifically for super-resolution models. HarmoQ jointly minimizes the compound error caused by weight-activation interaction by incorporating structure-aware mechanisms. First, it introduces a structural residual calibration scheme that adjusts weights to proactively compensate for structural degradation arising from activation quantization. Second, it derives a closed-form optimal scaling factor that analytically balances the quantization difficulty across components. Third, it applies adaptive boundary refinement to maintain this balance consistently throughout optimization, even under aggressive quantization regimes.

This work contributes the first systematic analysis of the asymmetric effects of weight versus activation quantization in super-resolution. It also proposes a closed-form structural calibration technique to mitigate the loss of fine-grained details, and an analytically grounded scaling strategy for harmonizing quantization across components. Together, these advances lead to strong empirical gains in low-bit regimes, offering a robust and theoretically motivated solution for efficient super-resolution deployment.

Related Work

Image Super-resolution. Image super-resolution aims to reconstruct high-resolution images from low-resolution inputs, with applications in medical imaging (Greenspan 2008; Isaac and Kulkarni 2015) and surveillance (Zhang et al. 2010). Early CNN approaches like SRCNN (Dong et al. 2014) pioneered deep learning for super-resolution. EDSR (Lim et al. 2017) achieved state-of-the-art results using enhanced residual networks, while SRGAN (Ledig et al. 2017) introduced adversarial training for photo-realistic results. RCAN (Zhang et al. 2018a) and RDN (Zhang et al. 2018b) further improved performance through attention mechanisms and dense connections. Recent transformer-based methods have shown promising results. SwinIR (Liang et al. 2021) adapted Swin Transformer (Liu et al. 2021) for image restoration, while Restormer (Zamir et al. 2022) proposed efficient transformer architectures for high-resolution restoration. CAT (Chen et al. 2022) and DAT (Chen et al. 2023b) explore cross-aggregation and dual-aggregation mechanisms for enhanced feature learning.

Model Quantization. Model quantization reduces computational complexity by representing weights and activations with lower precision, enabling efficient deployment on resource-constrained devices. Early quantization-aware training methods include BinaryNets (Courbariaux et al. 2016), which constrained weights to binary values, and DoReFa-Net (Zhou et al. 2016), which extended low-bitwidth quantization. PACT (Choi et al. 2018) introduced parameterized clipping for quantized networks. Post-training quantization offers practical advantages by quantizing pre-trained models without retraining. BRECQ (Li and et al. 2021) achieved significant improvements through block-wise reconstruction, while recent work (Ding et al. 2022) addresses PTQ challenges for vision transformers. For super-resolution specifically, PAMS (Li et al. 2020) proposed parameterized max scale quantization, DAQ (Hong et al. 2022b) developed channel-wise distribution-aware methods, and CADYQ (Hong et al. 2022a) introduced content-aware dynamic quantization. Recent advances (Tu et al. 2023) have made progress toward accurate post-training quantization for super-resolution models.

Problem Formulation

Consider a linear transformation in a super-resolution network: $y = Wx + b$, where $W \in \mathbb{R}^{m \times n}$ and $x \in \mathbb{R}^n$ denote the weights and input activation, respectively. Let $\tilde{W}$ and $\tilde{x}$ be their quantized counterparts, with corresponding quantization errors $\delta_W = \tilde{W} - W$ and $\delta_x = \tilde{x} - x$. The quantized forward pass expands to:

$$\tilde{y} = Wx + W\delta_x + \delta_W x + \delta_W \delta_x + b. \tag{1}$$

Our analysis in Figure 1 reveals that weight and activation quantization exhibit complementary degradation patterns: weight quantization primarily impacts structural similarity (degrading SSIM), while activation quantization corrupts pixel-level accuracy (degrading PSNR). This metric differentiation makes independent optimization suboptimal, as existing methods (Jacob et al. 2018; Tu et al. 2023) fail to minimize the critical compound error term $W\delta_x + \delta_W x$.

We propose a harmonized quantization framework that jointly optimizes quantization parameters to preserve critical structural and pixel-level information. Our approach formulates the problem as:

$$\begin{aligned} \min_{s} \quad & \mathcal{L}_{\text{total}}(s) = \mathbb{E}\left[\|W\delta_x(s) + \delta_W(s)x\|_2^2\right], \\ \text{s.t.} \quad & |\text{MSE}_x(s) - \text{MSE}_w(s)| \leq \epsilon, \end{aligned} \tag{2}$$

where s is a scaling parameter shared across quantization modules, and ϵ enforces distortion symmetry. The optimal s^* satisfies $\text{MSE}_x(s^*) \approx \text{MSE}_w(s^*)$, yielding minimal total error under balanced quantization.

Methodology

As illustrated in Figure 4, traditional approaches like 2DQuant and DOBI (Figure 4a) calibrate weight and activation quantization parameters separately, resulting in suboptimal boundary choices that neglect their interplay. In contrast, our HarmoQ framework (Figure 4b) introduces a unified three-stage optimization strategy: (1) *calibration* to set initial parameter ranges, (2) *harmonized scaling* to balance quantization difficulty via an optimal scale factor s^*, and (3) *adaptive boundary refinement* to jointly adjust clipping ranges, minimizing the compound error. This coordinated process preserves critical image fidelity components and achieves more robust super-resolution performance.

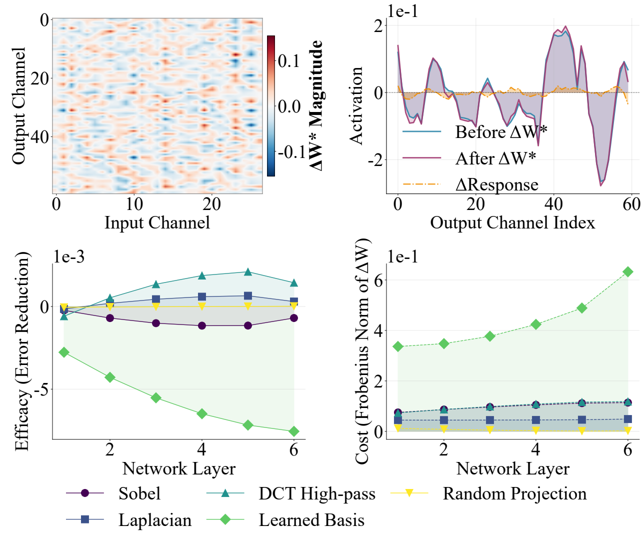

Figure 3: Structural Residual Calibration analysis in SwinIR. The top-left panel presents the optimal weight calibration ΔW^* heatmap using Laplacian filter, revealing structured channel relationships. The top-right panel shows response comparison before and after ΔW^* calibration, demonstrating structure-aware modulation. The bottom panel illustrates the efficacy across layers for different projection matrices. Laplacian filter achieves superior performance (see Table 3 for quantitative comparison), while random projection fails due to lack of structure-aware design. Cost analysis via Frobenius norm of ΔW shows consistent magnitude across projection types.

Step 1: Structural Residual Calibration

In quantized super-resolution networks, the compound error term $W\delta_x + \delta_W x$ exhibits metric-specific degradation patterns. The core challenge stems from weight quantization errors δ_W, which systematically corrupt the learned structural priors (e.g., for edges and textures) in W. This degradation is particularly severe in SR tasks where preserving edge sharpness and texture details is critical. Unlike weight quantization that introduces systematic biases affecting global image properties uniformly, activation quantization errors vary both spatially and temporally, making them particularly detrimental to the preservation of fine-grained visual details that are critical for super-resolution quality.

We propose a Structural Residual Calibration (SRC) mechanism that adjusts weight parameters to compensate for structural degradation introduced by weight quantization δ_W itself. This approach leverages the insight that weight parameters, with their distributed representation capacity, can be strategically modified to encode compensatory information that mitigates activation-induced loss of fine-grained detail before they manifest in the output.

Subspace projection operator. Let $H \in \mathbb{R}^{k \times d}$ be a linear projection matrix that extracts structural features like edges and textures from d-dimensional output features. This matrix can be constructed using fixed filters such as DCT high-

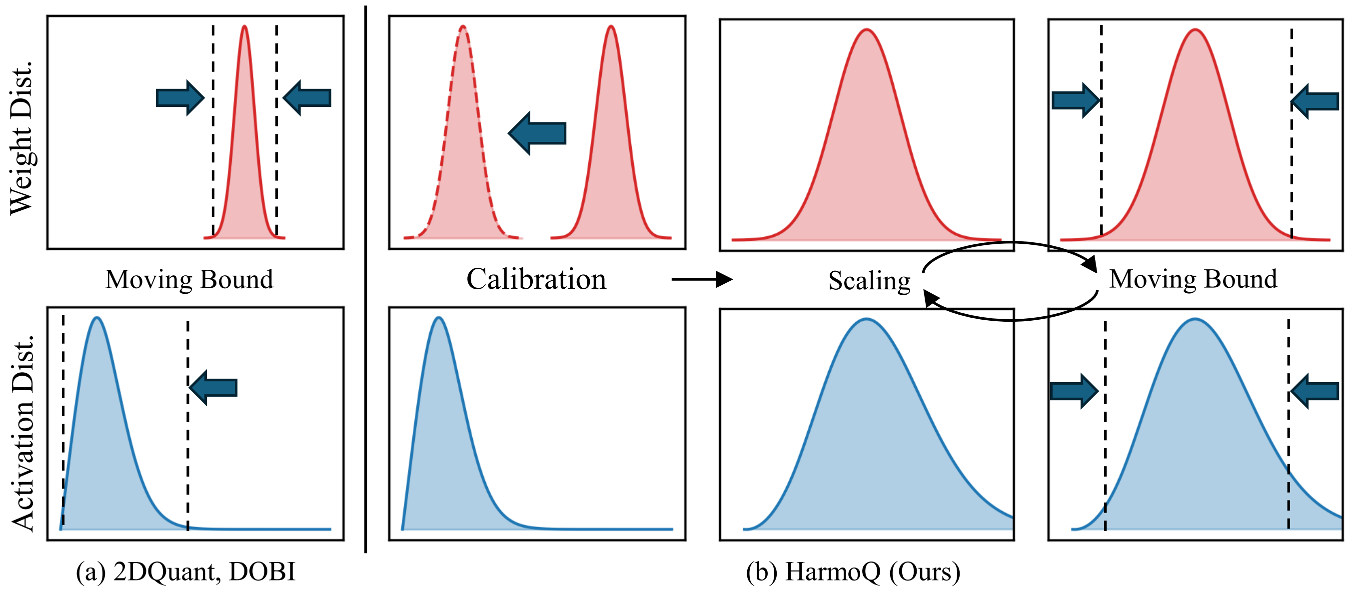

Figure 4: Comparison of quantization optimization strategies. (a) Existing methods (2DQuant, DOBI) independently optimize weight and activation quantization boundaries through separate calibration processes, leading to suboptimal parameter selection. (b) Our HarmoQ framework employs a unified three-step optimization: calibration for initial range estimation, harmonized scaling using optimal factor s^* to balance quantization difficulty, and iterative moving bounds refinement to jointly minimize compound quantization errors. The arrows indicate the iterative optimization flow between scaling and boundary adjustment steps.

pass masks or Laplacian edge filters bases derived from calibration data.

Structure-aware calibration objective. We seek a calibration term δ_W that suppresses the projected residual energy while maintaining stability through regularization:

$$\min_{\delta_W} \mathbb{E}\left[\|H(W\delta_x + \delta_W x)\|_2^2\right] + \lambda\|\delta_W\|_F^2, \tag{3}$$

where $\lambda > 0$ balances calibration strength and smoothness.

Closed-form solution. By expanding the objective and taking its derivative with respect to δ_W, we obtain the following optimality condition:

$$\delta_W \left(H\mathbb{E}[xx^T]H^T + \lambda I\right) = -W\mathbb{E}[\delta_x x^T]H^T. \tag{4}$$

Solving for δ_W yields the closed-form structure-aware calibration:

$$\delta_W^* = -W\mathbb{E}[\delta_x x^T]H^T \left(H\mathbb{E}[xx^T]H^T + \lambda I\right)^{-1} \tag{5}$$

The complete derivation of this result is provided in *Appendix*.

Projection analysis. To validate our structural residual calibration, we analyze the optimal weight calibration ΔW^* across different projection matrices. The top-left panel of Figure 3 shows the heatmap of ΔW^* derived via a Laplacian filter, revealing structured, non-random patterns that target specific channel relationships, confirming the structure-specific corrective capability of our method. The top-right panel illustrates how ΔW^* modulates activation responses while preserving key signal characteristics. The smooth transitions in the corrected profile support our structure-aware calibration principle and demonstrate reduced structural distortion. The bottom panel further shows consistent error reduction across layers for various structural projection filters. DCT high-pass and learned basis filters yield superior results, especially in deeper layers where quantization errors accumulate, while random projections perform poorly, underscoring the importance of a structure-aware design.

Step 2: Harmonized Scale Optimization

Following the structural calibration, we address the fundamental challenge of *non-uniform quantization difficulty* across weights and activations due to their distinct statistical properties. We seek a harmonizing scale factor s that balances quantization difficulty between components.

Scale harmonization objective. Given clipping boundaries $\theta = \{\alpha_x, \beta_x, \alpha_w, \beta_w\}$, we seek a scale factor s that ensures equal quantization difficulty: $\text{MSE}_x(s) = \text{MSE}_w(s)$. In SR networks, the activation range is further modulated by an input scaling factor s, which affects MSE differently for activations and weights:

$$\text{MSE}_x(s) = \frac{(\beta_x - \alpha_x)^2}{12s^2(2^{b_x} - 1)^2}, \ \text{MSE}_w(s) = \frac{(\beta_w - \alpha_w)^2 s^2}{12(2^{b_w} - 1)^2}.$$

Setting $\text{MSE}_x(s) = \text{MSE}_w(s)$ and solving for s gives the closed-form optimal scale:

$$s^* = \sqrt{\frac{(\beta_x - \alpha_x)(2^{b_w} - 1)}{(\beta_w - \alpha_w)(2^{b_x} - 1)}}. \tag{6}$$

The detailed derivation of the optimal scale factor is given in *Appendix*. This balances the difficulty of quantizing weights and activations, particularly in structure-sensitive regimes, ensuring that neither component dominates the quantization error.

Step 3: Adaptive Boundary Refinement

With the harmonized scale s^* established, we optimize the quantization boundaries θ to minimize the total reconstruction error while maintaining the balanced quantization difficulty achieved in Step 2.

Boundary optimization objective. We minimize the total quantization error by updating the quantization boundaries θ:

$$\mathcal{L}_{\text{total}}(s^*, \theta) = \mathbb{E}\left[\|W \cdot \delta_x(s^*, \theta) + \delta_W(s^*, \theta) \cdot x\|_2^2\right].$$

The gradient with respect to any boundary parameter $\theta_i \in \theta$ is:

$$\frac{\partial \mathcal{L}_{\text{total}}}{\partial \theta_i} = \mathbb{E}\left[2\left(W\delta_x + \delta_W x\right)^T \cdot \left(W\frac{\partial \delta_x}{\partial \theta_i} + \frac{\partial \delta_W}{\partial \theta_i}x\right)\right].$$

Gradient computation for clipping boundaries. Since δ_x and δ_W are functions of the quantizer's clipping range, their derivatives w.r.t. α and β depend on the quantization function. For a symmetric uniform quantizer $Q(z) = \text{clip}(z; \alpha, \beta)$ with step size $\Delta = \frac{\beta-\alpha}{2^b-1}$, we can compute:

$$\frac{\partial \delta}{\partial \alpha} = \frac{\partial}{\partial \alpha}(Q(z) - z), \qquad \frac{\partial \delta}{\partial \beta} = \frac{\partial}{\partial \beta}(Q(z) - z),$$

which are differentiable almost everywhere except at clipping points. The complete gradient derivation for boundary optimization is presented in *Appendix*.

Boundary update rule. The boundaries are updated using projected gradient descent:

$$\theta^{(t+1)} = \text{Proj}_{\Omega}\left(\theta^{(t)} - \eta\nabla_\theta \mathcal{L}_{\text{total}}(s^*, \theta^{(t)})\right), \quad (7)$$

where Ω is a feasible set ensuring $\alpha < \beta$ and η is the learning rate.

Iterative Refinement Process

The unified framework alternates between three coordinated operations:

1. **Step 1 - Structural Residual Calibration:** Applies closed-form structural calibration $\delta W_\ell^* = -W_\ell \Sigma_{\delta x} H^T (H\Sigma_{xx}H^T + \lambda I)^{-1}$ to minimize structural distortion.
2. **Step 2 - Harmonized Scale Optimization:** Computes optimal scale factor $s^* = \sqrt{\frac{(\beta_x - \alpha_x)(2^{b_w}-1)}{(\beta_w - \alpha_w)(2^{b_x}-1)}}$ to maintain equal quantization difficulty.
3. **Step 3 - Adaptive Boundary Refinement:** Updates clipping boundaries θ via gradient descent while enforcing the constraint $|\text{MSE}_x(s^*, \theta) - \text{MSE}_w(s^*, \theta)| \leq \epsilon$.

Experiments

Baseline Methods. We compare HarmoQ against several categories of quantization methods. Traditional approaches include MinMax (Jacob et al. 2018) and Percentile (Li et al. 2019) quantization, which use activation extrema and percentiles respectively to determine quantization boundaries. General post-training quantization methods include BRECQ (Li and et al. 2021), which pushes quantization limits through block-wise reconstruction, and DOBI, a data-free method that independently optimizes weight and activation quantization. For super-resolution specific methods, we evaluate against PAMS (Li et al. 2020) (parameterized max scale quantization), DBDC+Pac (Tu et al. 2023) (accurate post-training quantization for image super-resolution), 2DQuant (Liu et al. 2024) (low-bit post-training quantization), and Dynamic Granularity (Wang et al. 2025) (activation-only quantization strategy).

Experimental Setup. We evaluate HarmoQ on standard super-resolution benchmarks using SwinIR-light (Liang et al. 2021) and HAT-S (Chen et al. 2023a) as backbone networks. Following common practice, we train on DF2K dataset (combination of DIV2K (Timofte et al. 2017) and Flickr2K (Lim et al. 2017)) and test on five widely-used benchmarks: Set5 (Bevilacqua et al. 2012), Set14 (Zeyde, Elad, and Protter 2010), B100 (Martin et al. 2001), Urban100 (Huang, Singh, and Ahuja 2015), and Manga109 (Matsui et al. 2017). We conduct experiments with ×2 and ×4 upscaling factors under 2-bit and 3-bit quantization settings. Performance is measured using PSNR and SSIM computed on the Y channel of YCbCr color space.

Implementation Details. We use Adam optimizer with learning rate $1{\times}10^{-2}$, batch size 32, and 3000 total iterations. The structural projection matrix H is constructed using Laplacian filters (see Table 3 for ablation), which empirically outperform DCT high-pass and learned bases by 0.26 dB on Set5. The regularization parameter is set to $\lambda = 0.01$. Convergence is achieved when $|\mathcal{L}_{\text{total}}^{(t+1)} - \mathcal{L}_{\text{total}}^{(t)}| < 10^{-4}$, typically within 15-20 iterations.

Quantitative Evaluation. Table 1 presents comprehensive comparisons of HarmoQ against existing quantization methods on standard super-resolution benchmarks. Our method achieves consistent improvements under 2-bit and 3-bit quantization, particularly on structure-rich datasets like Urban100. This validates our core hypothesis that weight quantization affects structural similarity (SSIM) while activation quantization corrupts pixel-level accuracy (PSNR). Through unified optimization of structural residual calibration, harmonized scale balancing, and adaptive boundary refinement, HarmoQ transforms compound quantization errors from independent linear superposition to coordinated control. Cross-architecture experiments from SwinIR-light to HAT-S confirm the method's generalizability, while performance on challenging 4× super-resolution demonstrates robustness under extreme compression. The quantitative results validate both HarmoQ's technical superiority and the effectiveness of coupling analysis, providing a practical pathway for efficient super-resolution model deployment.

Method	Bit	Set5 (×2) PSNR/SSIM	Set14 (×2) PSNR/SSIM	BSD100 (×2) PSNR/SSIM	Urban100 (×2) PSNR/SSIM	Set5 (×4) PSNR/SSIM	Set14 (×4) PSNR/SSIM	BSD100 (×4) PSNR/SSIM	Urban100 (×4) PSNR/SSIM
SwinIR-light	32	38.15/0.9611	33.86/0.9206	32.31/0.9012	32.76/0.9340	32.45/0.8976	28.77/0.7858	27.69/0.7406	26.48/0.7980
Bicubic	32	32.25/0.9118	29.25/0.8406	28.68/0.8104	25.96/0.8088	27.56/0.7896	25.51/0.6820	25.54/0.6466	22.68/0.6352
MinMax	2	33.88/0.9185	30.81/0.8748	29.99/0.8535	27.48/0.8501	23.96/0.4950	22.92/0.4407	22.70/0.3943	21.16/0.4053
Percentile	2	30.82/0.8016	28.80/0.7616	27.95/0.7232	26.30/0.7378	23.03/0.4772	22.12/0.4059	21.83/0.3816	20.45/0.3951
DBDC+Pac	2	34.55/0.9386	31.12/0.8912	30.27/0.8706	27.63/0.8649	25.01/0.5554	23.82/0.4995	23.64/0.4544	21.84/0.4631
DOBI	2	35.25/0.9361	31.72/0.8917	30.62/0.8699	28.52/0.8727	28.82/0.7699	26.46/0.6804	25.97/0.6319	23.67/0.6407
Granular-DQ	2	35.85/0.9485	31.85/0.8995	30.80/0.8795	28.45/0.8800	29.35/0.8355	26.75/0.7305	26.35/0.6910	23.70/0.6895
2DQuant	2	36.00/0.9497	31.98/0.9012	30.91/0.8810	28.62/0.8819	29.53/0.8372	26.86/0.7322	26.46/0.6927	23.84/0.6912
HarmoQ	2	**36.46/0.9515**	**32.24/0.9037**	**31.12/0.8836**	**29.18/0.8893**	**30.23/0.8543**	**27.29/0.7446**	**26.70/0.7031**	**24.27/0.7123**
MinMax	3	28.19/0.6961	26.40/0.6478	25.83/0.6225	25.19/0.6773	19.41/0.3385	18.35/0.2549	18.79/0.2434	17.88/0.2825
Percentile	3	34.37/0.9170	31.04/0.8646	29.82/0.8339	28.25/0.8417	27.55/0.7270	25.15/0.6043	24.45/0.5333	22.80/0.5833
DBDC+Pac	3	35.07/0.9350	31.52/0.8873	30.47/0.8665	28.44/0.8709	27.91/0.7250	25.86/0.6451	25.65/0.6239	23.45/0.6249
DOBI	3	36.37/0.9496	32.33/0.9041	31.12/0.8836	29.65/0.8967	29.59/0.8237	26.87/0.7156	26.24/0.6735	24.17/0.6880
Granular-DQ	3	37.20/0.9555	32.75/0.9095	31.50/0.8900	30.30/0.9070	30.75/0.8690	27.65/0.7560	26.90/0.7115	24.70/0.7340
2DQuant	3	37.32/0.9567	32.85/0.9106	31.60/0.8911	30.45/0.9086	30.90/0.8704	27.75/0.7571	26.99/0.7126	24.85/0.7355
HarmoQ	3	**37.98/0.9602**	**34.04/0.9220**	**32.23/0.9004**	**33.01/0.9368**	**31.09/0.8740**	**27.91/0.7618**	**27.11/0.7186**	**25.31/0.7548**
HAT-S	32	38.58/0.9628	34.70/0.9261	32.59/0.9050	34.31/0.9459	32.92/0.9047	29.15/0.7958	27.97/0.7505	27.87/0.8346
Bicubic	32	32.25/0.9118	29.25/0.8406	28.68/0.8104	25.96/0.8088	27.56/0.7896	25.51/0.6820	25.54/0.6466	22.68/0.6352
MinMax	2	34.63/0.9263	31.23/0.8798	30.59/0.8595	30.95/0.8777	25.95/0.5746	24.16/0.4888	23.78/0.4483	22.37/0.4734
Percentile	2	32.06/0.8756	29.89/0.8118	28.97/0.7842	28.68/0.8288	24.18/0.5398	22.89/0.4507	22.34/0.4162	21.01/0.4346
DBDC+Pac	2	35.89/0.9457	32.29/0.9001	31.28/0.8777	30.61/0.8934	26.97/0.6285	25.16/0.5493	24.56/0.5044	23.18/0.5127
DOBI	2	36.63/0.9413	32.93/0.9008	31.69/0.8769	31.68/0.9004	31.04/0.8297	27.51/0.7104	26.89/0.6859	25.13/0.7089
Granular-DQ	2	36.70/0.9525	32.55/0.9070	31.30/0.8870	30.05/0.9045	30.30/0.8495	27.30/0.7405	26.80/0.7005	25.00/0.7410
2DQuant	2	36.81/0.9538	32.66/0.9085	31.42/0.8883	30.21/0.9061	30.48/0.8513	27.42/0.7423	26.89/0.7019	25.14/0.7425
HarmoQ	2	**38.34/0.9617**	**34.35/0.9246**	**32.14/0.8961**	**33.87/0.9434**	**31.18/0.8749**	**27.95/0.7626**	**27.12/0.7191**	**25.33/0.7560**
MinMax	3	30.86/0.7707	27.61/0.6957	26.07/0.6607	27.44/0.7567	21.77/0.4146	20.02/0.3287	19.77/0.3061	19.96/0.3764
Percentile	3	35.66/0.9283	32.17/0.8789	30.42/0.8453	31.37/0.8938	29.36/0.7855	26.13/0.6542	25.19/0.5998	25.16/0.7056
DBDC+Pac	3	36.41/0.9436	32.63/0.8977	30.99/0.8719	31.70/0.9018	29.82/0.7951	26.55/0.6769	25.71/0.6264	25.47/0.7178
DOBI	3	37.77/0.9558	33.49/0.9156	31.70/0.8901	33.09/0.9274	31.38/0.8586	27.83/0.7466	26.94/0.7085	26.89/0.8002
Granular-DQ	3	37.00/0.9520	33.15/0.9115	31.40/0.8860	32.65/0.9225	30.90/0.8520	27.45/0.7405	26.60/0.7030	26.40/0.7935
2DQuant	3	37.11/0.9532	33.26/0.9128	31.52/0.8872	32.78/0.9239	31.02/0.8534	27.57/0.7421	26.71/0.7043	26.54/0.7948
HarmoQ	3	**38.35/0.9618**	**34.42/0.9250**	**32.36/0.9021**	**34.21/0.9441**	**31.90/0.8821**	**28.34/0.7700**	**27.41/0.7268**	**26.71/0.8106**

Table 1: Quantitative comparison of HarmoQ against SOTA quantization methods on standard SR benchmarks.

Components				Set5 (×2)		Urban100 (×4)	
SRC	HSO	ABR	Config	PSNR	SSIM	PSNR	SSIM
✗	✗	✗	SwinIR-Light	35.12	0.9387	23.12	0.6234
✓	✗	✗	SRC	36.24	0.9468	23.89	0.6567
✗	✓	✗	HSO	36.45	0.9482	24.15	0.6634
✗	✗	✓	ABR	35.89	0.9434	23.56	0.6423
✓	✓	✗	SRC+HSO	37.12	0.9531	24.67	0.7012
✓	✗	✓	SRC+ABR	36.78	0.9503	24.23	0.6789
✗	✓	✓	HSO+ABR	37.01	0.9521	24.45	0.6934
✓	✓	✓	**Full HarmoQ**	**37.98**	**0.9602**	**25.31**	**0.7548**

Table 2: Component-wise ablation analysis.

Projection Matrix H	Set5 (×2)		Urban100 (×4)	
	PSNR	SSIM	PSNR	SSIM
w/o SRC	37.01	0.9521	24.45	0.6934
Laplacian Filter	**37.98**	**0.9602**	**25.31**	**0.7548**
Sobel Edge Filter	37.85	0.9589	25.18	0.7521
DCT High-pass	37.72	0.9575	25.02	0.7495
Learned Basis	37.91	0.9594	25.24	0.7535
Random Projection	37.34	0.9548	24.78	0.7312

Table 3: Structural feature projection matrix analysis.

Component-wise ablation. Our ablation study in Table 2 demonstrates that HarmoQ's performance gains arise from complex nonlinear synergies rather than linear superposition. The baseline method, which processes weight and activation quantization in isolation while ignoring their coupling, embodies the dilemma of traditional super-resolution quantization. In contrast, our components show significant advantages. The SRC mechanism markedly improves per-

Figure 5: Visual comparison of different quantization methods on benchmark dataset.

formance by compensating for loss of structural information from activation quantization. HSO outperforms SRC when applied alone, highlighting the critical role of balancing quantization difficulty to enhance system robustness. ABR shows limited standalone effectiveness, as its performance depends on an established quantization balance. Pairwise combinations reveal strong synergy between SRC and HSO, where calibration is amplified by the quantization balance. Without HSO as a bridge, however, SRC and ABR struggle to synergize. Ultimately, the full HarmoQ framework outperforms all two-component combinations, demonstrating that its unified optimization enables a qualitative leap from local to global coordination. Consistent results on the Urban100 dataset validate this universality in scenes with rich textures.

Structural feature projection analysis. Our ablation study in Table 2 explores the quest for optimal structure-aware calibration. Laplacian filters emerge victorious (37.98/0.9602), their second-order nature perfectly capturing edge discontinuities that activation quantization disrupts. Sobel filters (37.85/0.9589) and learned basis (37.91/0.9594) engage in a close contest, proving both classical gradient wisdom and data-driven intelligence can decode quantization artifacts. DCT high-pass filters (37.72/0.9575) show promise yet fall short, suggesting spatial gradients trump pure frequency decomposition. Random projection's failure (37.34/0.9548) delivers the final verdict: effective error compensation demands explicit structural semantics, not blind dimensionality reduction.

Visual comparison. Figure 5 demonstrates HarmoQ's superiority under aggressive 3-bit quantization. MinMax produces severe artifacts with blurred giraffe fur and loss of building edge sharpness. DBDC and 2DQuant show improved structure but still exhibit texture degradation in fine details. Dynamic Granularity maintains better overall quality yet fails to preserve crisp window frames and fur texture boundaries. The ablation comparison reveals SRC's critical contribution: HarmoQ (w.o. SRC) preserves the giraffe's body contour but loses individual fur strand definition and produces softened building corner edges. Full HarmoQ recovers these fine-grained details—sharp fur textures, clear window boundaries, and preserved architectural edges—demonstrating that structural residual calibration effectively compensates for activation quantization artifacts in texture-sensitive regions.

Conclusion

This work addresses the deployment challenge of high-performance image SR models under aggressive compression. We proposed HarmoQ, a harmonized post-training quantization framework that unifies three strategies: structural residual calibration, harmonized scale optimization, and adaptive boundary refinement. HarmoQ enables practical deployment of high-quality SR models on resource-constrained devices, representing a significant advance in model compression for image restoration tasks.

Acknowledgments

This work was partially supported by JST-Mirai Program JPMJMI23G1, JSPS KAKENHI Grant Numbers 24KK0209, 24K22318, and 22H00529, the SPRING GX project of the University of Tokyo (grant number JPMJSP2108), the Research Institute of Trustworthy Autonomous Systems at Southern University of Science and Technology, the Jilin Provincial International Cooperation Key Laboratory for Super Smart City, and the Jilin Provincial Key Laboratory of Intelligent Policing.

References

Bevilacqua, M.; Roumy, A.; Guillemot, C.; and Alberi-Morel, M. L. 2012. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. In *BMVC*.

Chen, X.; Wang, X.; Zhou, J.; Qiao, Y.; and Dong, C. 2023a. Activating more pixels in image super-resolution transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 22367–22377.

Chen, Z.; Zhang, Y.; Gu, J.; Kong, L.; Yang, X.; and Yu, F. 2023b. Dual aggregation transformer for image super-resolution. In *CVPR*.

Chen, Z.; Zhang, Y.; Gu, J.; Zhang, Y.; Kong, L.; and Yuan, X. 2022. Cross aggregation transformer for image restoration. In *NeurIPS*.

Choi, J.; Wang, Z.; Venkataramani, S.; Chuang, P. I.-J.; Srinivasan, V.; and Gopalakrishnan, K. 2018. Pact: Parameterized clipping activation for quantized neural networks. *arXiv preprint arXiv:1805.06085*.

Courbariaux, M.; Hubara, I.; Soudry, D.; El-Yaniv, R.; and Bengio, Y. 2016. Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. *arXiv preprint arXiv:1602.02830*.

Ding, Y.; Qin, H.; Yan, Q.; Chai, Z.; Liu, J.; Wei, X.; and Liu, X. 2022. Towards accurate post-training quantization for vision transformer. In *ACM MM*.

Dong, C.; Loy, C. C.; He, K.; and Tang, X. 2014. Learning a deep convolutional network for image super-resolution. In *ECCV*.

Greenspan, H. 2008. Super-resolution in medical imaging. *The Computer Journal*.

Hong, C.; Baik, S.; Kim, H.; Nah, S.; and Lee, K. M. 2022a. Cadyq: Content-aware dynamic quantization for image super-resolution. In *ECCV*.

Hong, C.; Kim, H.; Baik, S.; Oh, J.; and Lee, K. M. 2022b. Daq: Channel-wise distribution-aware quantization for deep image super-resolution networks. In *WACV*.

Huang, J.-B.; Singh, A.; and Ahuja, N. 2015. Single image super-resolution from transformed self-exemplars. In *CVPR*.

Isaac, J. S.; and Kulkarni, R. 2015. Super resolution techniques for medical image processing. In *ICTSD*.

Jacob, B.; Kligys, S.; Chen, B.; Zhu, M.; Tang, M.; Howard, A.; Adam, H.; and Kalenichenko, D. 2018. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *CVPR*.

Ledig, C.; Theis, L.; Huszár, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A.; Tejani, A.; Totz, J.; Wang, Z.; and Shi, W. 2017. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*.

Li, H.; Yan, C.; Lin, S.; Zheng, X.; Zhang, B.; Yang, F.; and Ji, R. 2020. Pams: Quantized super-resolution via parameterized max scale. In *ECCV*.

Li, R.; Wang, Y.; Liang, F.; Qin, H.; Yan, J.; and Fan, R. 2019. Fully quantized network for object detection. In *CVPR*.

Li, Z.; and et al. 2021. Brecq: Pushing the limit of post-training quantization with bias rounding and equalization. In *ICML*.

Liang, J.; Cao, J.; Sun, G.; Zhang, K.; Gool, L. V.; and Timofte, R. 2021. Swinir: Image restoration using swin transformer. In *ICCVW*.

Lim, B.; Son, S.; Kim, H.; Nah, S.; and Lee, K. M. 2017. Enhanced deep residual networks for single image super-resolution. In *CVPRW*.

Liu, K.; Qin, H.; Guo, Y.; Yuan, X.; Kong, L.; Chen, G.; and Zhang, Y. 2024. 2DQuant: Low-bit post-training quantization for image super-resolution. *Advances in Neural Information Processing Systems*, 37: 71068–71084.

Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*.

Martin, D.; Fowlkes, C.; Tal, D.; and Malik, J. 2001. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV*.

Matsui, Y.; Ito, K.; Aramaki, Y.; Fujimoto, A.; Ogawa, T.; Yamasaki, T.; and Aizawa, K. 2017. Sketch-based manga retrieval using manga109 dataset. *Multimedia Tools and Applications*.

Timofte, R.; Agustsson, E.; Gool, L. V.; Yang, M.-H.; Zhang, L.; Lim, B.; Son, S.; Kim, H.; Nah, S.; Lee, K. M.; et al. 2017. Ntire 2017 challenge on single image super-resolution: Methods and results. In *CVPRW*.

Tu, Z.; Hu, J.; Chen, H.; and Wang, Y. 2023. Toward accurate post-training quantization for image super resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5856–5865.

Wang, M.; Zhang, Z.; Li, F.; Xu, K.; Miao, K.; and Wang, M. 2025. Thinking in Granularity: Dynamic Quantization for Image Super-Resolution by Intriguing Multi-Granularity Clues. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 7826–7834.

Zamir, S. W.; Arora, A.; Khan, S.; Hayat, M.; Khan, F. S.; and Yang, M.-H. 2022. Restormer: Efficient transformer for high-resolution image restoration. In *CVPR*.

Zeyde, R.; Elad, M.; and Protter, M. 2010. On single image scale-up using sparse-representations. In *Proc. 7th Int. Conf. Curves Surf.*

Zhang, L.; Zhang, H.; Shen, H.; and Li, P. 2010. A super-resolution reconstruction algorithm for surveillance images. *Elsevier Signal Processing*.

Zhang, Y.; Li, K.; Li, K.; Wang, L.; Zhong, B.; and Fu, Y. 2018a. Image super-resolution using very deep residual channel attention networks. In *ECCV*.

Zhang, Y.; Tian, Y.; Kong, Y.; Zhong, B.; and Fu, Y. 2018b. Residual dense network for image super-resolution. In *CVPR*.

Zhou, S.; Wu, Y.; Ni, Z.; Zhou, X.; Wen, H.; and Zou, Y. 2016. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*.