Phase-Dominant Spectral Encoding for Geometry-Preserving Compression in Scalable Visual Data Pipelines

Gehan Sathsara Vithanage, Dasuni Wimalachandra, Amila Dhanushka Karunanayaka, Ruchire Eranga Wijesinghe
CIET/DEEE, Sri Lanka Institute of Information Technology (SLIIT)
Malabe, Western Province, Sri Lanka
vithanage.gs@ieee.org,dasuniwimalachandra@gmail.com,karunanayakaamila197@gmail.com,eranga.w@sliit.lk

Abstract

Modern data engineering pipelines require compact representations that preserve structurally meaningful information rather than pixel-level fidelity. Conventional codecs such as JPEG prioritize transform magnitudes, implicitly under-utilizing phase despite its dominant role in spatial organization. This paper presents a phase-prioritized Fourier-domain compression framework for structure-aware data reduction in large-scale visual datasets. Low-frequency spectral components are preserved, while high-frequency magnitudes are replaced with compact priors and phase is uniformly quantized under Hermitian constraints for efficient reconstruction. Under extreme low-bitrate conditions, the method maintains structural coherence at substantially reduced storage cost. For a retinal OCT dataset, it achieves 0.051 bpp with SSIM = 1.000, whereas JPEG requires 0.144 bpp and yields SSIM = 0.578. In conventional photographic images, although PSNR is lower (15.8 dB vs. 25.9 dB), edge continuity is better preserved at nearly half the bitrate (0.094 bpp vs. 0.176 bpp). These results position phase-dominant encoding as a lightweight and computation-friendly alternative for storage-efficient vision datasets and structure-sensitive analytics in data engineering systems.

Keywords

Image compression, Fourier transform, phase quantization, magnitude approximation, structural similarity, and machine learning features.

1 Introduction

Image compression plays a critical role in modern data engineering ecosystems, where large-scale visual data must be efficiently stored, transmitted, and processed across distributed systems. From cloud-based content delivery networks to biomedical repositories and edge analytics platforms, scalable management of visual datasets is essential. Traditional compression techniques such as JPEG, wavelet-based coding, vector quantization, and fractal compression have been widely deployed in these environments [1]. These methods primarily reduce spatial redundancy and perceptual irrelevance, often guided by models of the human visual system (HVS). In particular, transform-based approaches built on the discrete cosine transform (DCT) and discrete wavelet transform (DWT) underpin dominant standards, offering a balance between compression efficiency and perceptual reconstruction quality.

Despite their success, prevailing transform-based codecs are fundamentally magnitude-centric, allocating most coding resources to coefficient magnitudes while implicitly treating phase. However, from a data representation perspective, phase encodes critical structural relationships, geometric alignment, and spatial organization, properties that are increasingly important in downstream analytics, feature extraction, and machine learning pipelines. Although standards such as JPEG and JPEG2000 rely on DCT or DWT frameworks, systematic exploration of phase-dominant coding remains limited. This gap suggests an opportunity to rethink compression as a structure-aware data reduction problem rather than purely a perceptual optimization task. Investigating phase-prioritized representations opens the possibility of designing lightweight codecs that preserve geometry-critical information while enabling efficient storage and computation within modern data engineering infrastructures.

1.1 Review of Related Articles

The application of frequency-domain methods for image compression is demonstrated in [2], where a fast fourier transform (FFT)-based codec is introduced. The approach leverages spectral domain representation to selectively retain dominant coefficients, enabling efficient reconstruction with acceptable quality for grayscale images. Similarly, [3] employs a 2D discrete fourier transform (DFT) coupled with a matrix minimization algorithm, achieving compression by reducing redundant coefficients while maintaining recognizable structural details in images.

Transform-based alternatives extend to the anamorphic stretch transform (AST). In [4], AST is proposed to map input images into a warped frequency domain where redundancies become more separable, yielding higher compression ratios compared to conventional transforms. Further, quaternion-based approaches have been introduced, as in [5], where reduced bit-plane quaternion singular value decomposition is applied to color images, enhancing compression while preserving perceptual quality across channels.

Recent innovations also explore phase-domain methods. In [6], an optimal phase coding strategy based on the root-mean-square (RMS) duration principle is proposed. This method emphasizes phase information as the primary carrier of structural features, enabling efficient spectral image compression with promising fidelity. Machine learning perspectives are reviewed in [7], which highlights clustering-based and neural approaches for adaptively modeling image redundancies, while [8] designs an unsupervised deep model capable of learning compression mappings without labeled data, demonstrating competitive performance with traditional codecs.

ADScAI Conference, University of Moratuwa, Sri Lanka
2026.

Finally, broader surveys [1], [9], [10] systematically classify both classical and learning-based methods. These works outline the progression from transform coding (e.g., JPEG, DWT, fractal) toward deep learning-driven end-to-end codecs, noting the increasing adoption of convolutional and generative models for image compression.

1.2 Limitations of Related Works

Certain limitations in the work thus far prove to be disadvantageous. These can be summarized as follows:

(1) FFT-based methods [1, 2] produce noticeable artifacts at low bitrates and often lag behind JPEG in quality.
(2) Transforms such as AST and quaternion SVD [3, 4] involve high mathematical complexity, making real-time use difficult.
(3) Machine learning approaches [6, 7, 10] demand large datasets and heavy computation, limiting lightweight applications.
(4) Survey works [8, 9] summarize existing techniques but do not introduce new paradigms.

Therefore, despite substantial progress in transform-based, and learning-driven image compression, the reviewed approaches reveal persistent drawbacks in balancing compression efficiency, reconstruction quality, and implementation feasibility.

1.3 Motivation and Contribution

Conventional image compression techniques have been designed primarily around magnitude-centric transforms, such as DCT and DWT, where the bulk of coding resources are allocated to representing coefficient magnitudes. This design philosophy has led to codecs that achieve high PSNR and SSIM values, but it underrepresents the structural role of phase. Decades of work in Fourier analysis demonstrate that image structure, edges, and geometry are predominantly encoded in phase, while magnitude conveys contrast and energy distribution. Yet, despite this well-known property, very few compression schemes place phase at the center of their design.

This work is motivated by the need to explore phase-dominant coding as a paradigm in its own right. Instead of competing with established standards such as JPEG or deep learning–based methods, the proposed codec serves as a demonstration of possibility: that phase-prioritized allocation, combined with lightweight magnitude priors, can achieve structurally faithful reconstructions at extremely low bitrates. The approach highlights how unconventional allocation strategies may be suitable for task-driven applications where structural preservation outweighs perceptual fidelity, such as machine learning feature extraction or lightweight computer vision pipelines.

The proposed codec visually represented in Fig. 1 operates in the Fourier domain. First, the image is mean-normalized and transformed using a 2D Fourier transform. Low-frequency magnitude and phase components are preserved to anchor overall brightness and global structure. High-frequency magnitude is not retained directly; instead, it is replaced by a compact prior, either in the form of a radial profile (averaged over frequency radius) or a low-resolution magnitude map obtained through downsampling. Phase values outside the low-frequency region are uniformly quantized with a fixed bit depth, while conjugate symmetry is enforced to ensure real-valued reconstruction. The inverse Fourier transform then produces the compressed image, with a final variance adjustment applied to restore contrast.

The key contributions can be summarized as follows:

(1) Introduction of a Fourier-domain codec that prioritizes phase representation, with high-frequency magnitudes replaced by either a radial profile or a low-resolution prior.
(2) Demonstration of structurally consistent reconstructions at low bitrates, despite lower PSNR and SSIM compared to magnitude-driven codecs.
(3) Positioning of phase-prioritized coding as a complementary paradigm, showcasing its potential in task-oriented compression rather than as a competitor to established image standards.

1.4 Paper Organization

The remainder of this paper is organized as follows. Section II outlines the design of the proposed codec for compression and decompression. Section III presents the experimental results obtained using the phase-prioritized paradigm, along with comparisons against baseline methods. Finally, Section IV concludes the work with final remarks.

2 Methodology

The proposed codec operates in the Fourier domain, where an input image is transformed, selectively represented, and reconstructed with compact phase-prioritized coding. The design consists of four main stages: pre-processing, frequency-domain decomposition, magnitude approximation, and phase quantization, followed by reconstruction.

2.1 Preprocessing and Centering

Given an image $I(x, y)$ of size $H \times W$, the mean intensity is first removed:

$$I_c(x, y) = I(x, y) - \mu, \quad \mu = \frac{1}{HW} \sum_{x=1}^{H} \sum_{y=1}^{W} I(x, y). \tag{1}$$

2.2 Fourier Transform and Decomposition

The pre-processed image is transformed into the frequency domain using the 2D Fourier transform:

$$F(u, v) = \mathcal{F}\{I_c(x, y)\}, \tag{2}$$

which is decomposed into magnitude and phase as,

$$F(u, v) = M(u, v)\, e^{j\Phi(u,v)}, \tag{3}$$

where $M(u, v) = |F(u, v)|$ is the magnitude component and, $\Phi(u, v) = \arg(F(u, v))$ is the phase component.

2.3 Low-Frequency Preservation

In $F(u, v)$, frequencies are organized radially from the spectrum center (u_0, v_0). This is expressed as,

$$R(u, v) = \sqrt{(u - u_0)^2 + (v - v_0)^2}, \quad R_{\max} = \max R(u, v), \tag{4}$$

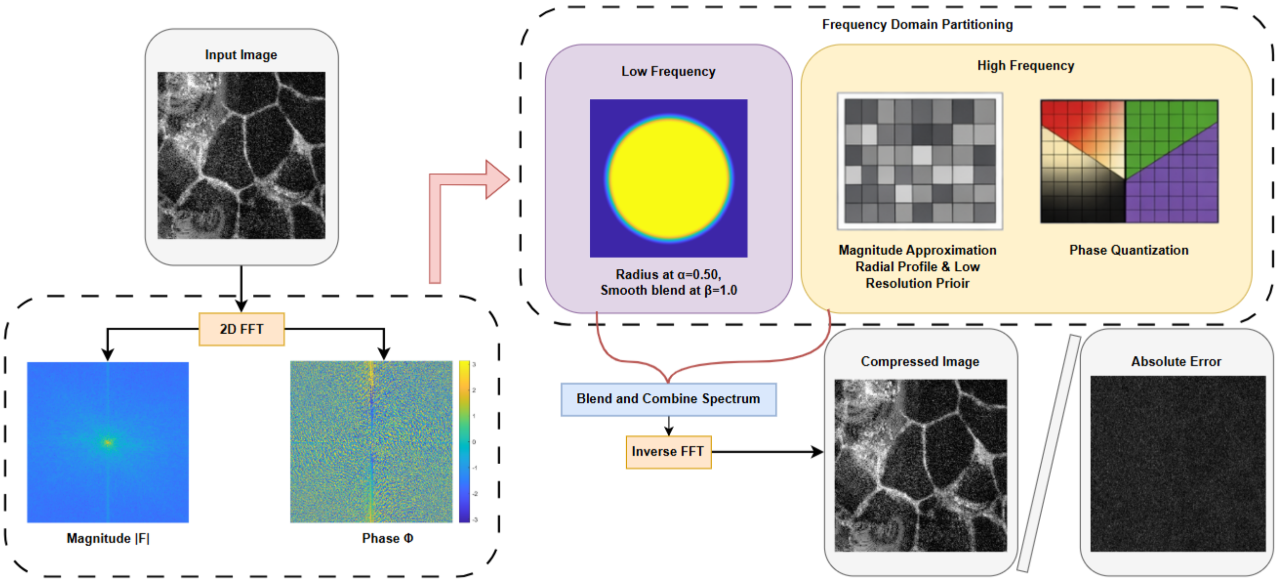

Figure 1: Proposed Fourier magnitude approximated image compression illustrated in a flow diagram.

where $R(u,v)$ is simply the radial distance of a frequency coefficient (u,v) from the center of the Fourier spectrum (u_0, v_0).

A low-frequency region of radius $R_k = \alpha R_{max}$ is preserved, with smooth blending over a width $\Delta = \beta R_{max}$. Here, $\alpha \in [0,1]$ defines the fraction of the maximum radius retained as the low-frequency region, while $\beta \in [0,1]$ specifies the fractional width of the smooth transition band applied around the cutoff. The blending mask is defined as,

$$w(r) = \begin{cases} 1, & r \le R_k, \\ \frac{1}{2}\big(1 + \cos\big(\pi(r - R_k)/\Delta\big)\big), & R_k < r < R_k + \Delta, \\ 0, & r \ge R_k + \Delta, \end{cases} \tag{5}$$

where $r \triangleq R(u,v)$.

2.4 Magnitude Approximation

Outside the low-frequency region, the true magnitude is replaced by a compact prior:

(1) Radial profile prior:

$$M_p(u,v) = \bar{M}(R(u,v)), \tag{6}$$

where $\bar{M}(r)$ is defined as,

$$\bar{M}(r) = \frac{1}{|\{(u,v) : R(u,v) = r\}|} \sum_{R(u,v)=r} M(u,v). \tag{7}$$

(2) Low-resolution prior:

$$M_p(u,v) = \text{Upsample}(M_d, H \times W), \tag{8}$$

where $M_d = \text{Downsample}(M_s, N{\times}N)$ and $M_s = G_\sigma * M(u,v)$. G_σ is a Gaussian smoothing filter.

The final blended magnitude is expressed as,

$$M'(u,v) = w(u,v)\, M(u,v) + (1 - w(u,v))\, M_p(u,v). \tag{9}$$

2.5 Phase Quantization and Symmetry

Outside the low-frequency region, phase is quantized using Q bits:

$$\Phi'(u,v) = \frac{2\pi}{L}\left(\left\lfloor \frac{\Phi(u,v) + \pi}{2\pi} L \right\rfloor + \frac{1}{2}\right) - \pi, \quad L = 2^Q. \tag{10}$$

Furthermore, Hermitian symmetry is enforced to ensure real-valued reconstruction as follows,

$$\Phi'(-u,-v) = -\Phi'(u,v). \tag{11}$$

2.6 Spectrum Reconstruction and Inverse Transform

For decoding the compressed image, the compressed spectrum is reconstructed as,

$$F'(u,v) = M'(u,v)\, e^{j\Phi'(u,v)}. \tag{12}$$

Finally, the decompressed image is obtained by inverse Fourier transform, with mean added back:

$$\hat{I}(x,y) = \mathcal{F}^{-1}\{F'(u,v)\} + \mu. \tag{13}$$

To restore global contrast, variance normalization is applied. This is defined as,

$$\hat{I}(x,y) \leftarrow \big(\hat{I}(x,y) - \bar{\hat{I}}\big) \cdot \frac{\sigma_I}{\sigma_{\hat{I}}} + \mu. \tag{14}$$

Here, $\bar{\hat{I}}$ denotes the mean of the reconstructed image, $\sigma_{\hat{I}}$ its standard deviation, σ_I the standard deviation of the original image, and μ the mean intensity of the original image.

3 Experimental Results

This section will rigorously evaluate the proposed phase-prioritized codec against JPEG under extreme low-bitrate conditions, emphasizing structural fidelity rather than conventional pixel-wise accuracy metrics. This analysis is conducted to demonstrate that, while

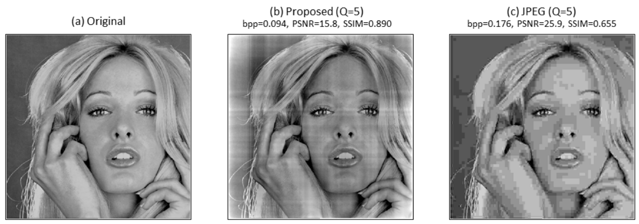

Figure 2: Image quality assessment of an image (woman) compressed with proposed adaptive Vs JPEG.

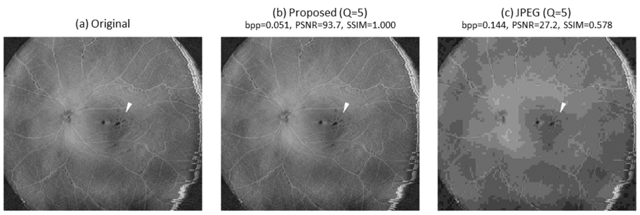

Figure 3: Image quality assessment of an OCT image (retina) compressed with proposed adaptive Vs JPEG

magnitude-centric techniques may optimize PSNR at moderate rates, the proposed paradigm exhibits superior structural robustness; particularly in preserving edge and geometric features under aggressive compression constraints.

3.1 Comparison of IQA Metrics for Proposed Codec with JPEG under Aggressive Compression

To evaluate structural robustness in the extreme compression regime, the proposed phase-prioritized codec was compared against baseline JPEG at very low quality settings $\alpha = 0.50$ and $\beta = 1.0$. Fig. 2. and Fig. 3. illustrate representative results for two images: a woman and an enface optical coherence tomography (OCT) scan of the retina [11]. The bitrate achieved by the proposed method is significantly lower than that of JPEG in both cases. Despite operating at nearly one-fifth of the bitrate, the reconstructed images preserve the global geometric structure and dominant edges.

For the woman image in Fig. 2, JPEG achieves higher PSNR (25.9 dB) compared to the proposed method (15.8 dB), reflecting superior pixel-wise fidelity. However, the JPEG reconstruction exhibits strong blocking artifacts and localized smoothing due to aggressive DCT coefficient quantization. However, the phase-driven representation preserves structural continuity, preventing the abrupt discontinuities typical of block-based compression at extreme quantization levels.

The structural advantage becomes even more evident in the retinal image in Fig. 3. The proposed codec achieves an SSIM of 1.000 at bpp of 0.051, indicating near-perfect preservation of structural relationships, while JPEG at over four times the bitrate (0.144 bpp) yields SSIM of 0.578. The vascular patterns and fine anatomical structures remain visually consistent in the phase-prioritized reconstruction, whereas JPEG introduces noticeable contrast distortion and detail loss. This behavior confirms that the preservation of Fourier phase information effectively maintains spatial organization, even when magnitude information is heavily approximated.

These results highlight a key distinction between magnitude-driven and phase-driven coding paradigms. JPEG optimizes energy compaction and perceptual smoothness, resulting in high PSNR and competitive SSIM at moderate bitrates. However, its performance degrades rapidly as quantization becomes aggressive, producing blocking and structural inconsistencies. The proposed phase-priority codec, while not optimizing pixel-level fidelity, demonstrates graceful degradation under severe bitrate constraints. Structural coherence and edge integrity remain largely intact even below 0.1 bpp.

To quantitatively evaluate the behavior of the proposed phase-prioritized codec, experiments were conducted on two independent image sets: natural baboon imagery and optical coherence microscopy (OCM) grape specimen scans as depicted in Fig.4. For each dataset, 10 grayscale images were processed using identical codec parameters. JPEG quality was selected through a constrained search procedure to operate in a comparable low-bitrate regime. Average PSNR, SSIM, bpp were computed across each dataset, and normalized bar plots were generated to enable direct metric-wise comparison.

For the baboon dataset, JPEG demonstrates superior PSNR performance, reflecting its magnitude-centric design optimized for minimizing mean squared error in natural scenes. The DCT-based energy compaction inherent to JPEG effectively preserves pixel-level fidelity in textured content. However, despite lower PSNR values, the proposed framework achieves competitive structural similarity (SSIM) at reduced bitrates. This behavior stems from the preservation of Fourier phase information, which maintains global edge alignment and geometric consistency even when magnitude precision is approximated.

The structural advantage of the proposed method becomes more pronounced in the OCM grape dataset. Biomedical imagery is characterized by fine-scale morphological patterns and continuous microstructural boundaries that are particularly sensitive to block-based quantization artifacts. In this context, the phase-prioritized codec yields higher normalized SSIM values while operating at a lower average bitrate than JPEG. The global spectral representation avoids localized discontinuities, thereby preserving morphological continuity that is degraded under aggressive JPEG compression.

These results demonstrate a consistent trade-off between pixel-wise fidelity and structural coherence. While JPEG excels in PSNR due to magnitude-driven optimization, the proposed codec emphasizes geometric preservation through phase-dominant allocation. The averaged results across both datasets confirm that phase-prioritized coding offers a complementary compression paradigm, particularly advantageous in structure-sensitive applications such as biomedical imaging and feature extraction pipelines.

3.2 Edge and Gradient Preservation at Ultra-Low Bitrates

In contrast, although the proposed method yields lower PSNR, it retains coherent edge contours and overall object geometry. To further investigate structural fidelity beyond conventional full-reference metrics, an edge-centric evaluation was conducted under extreme compression conditions. The proposed phase-prioritized codec was

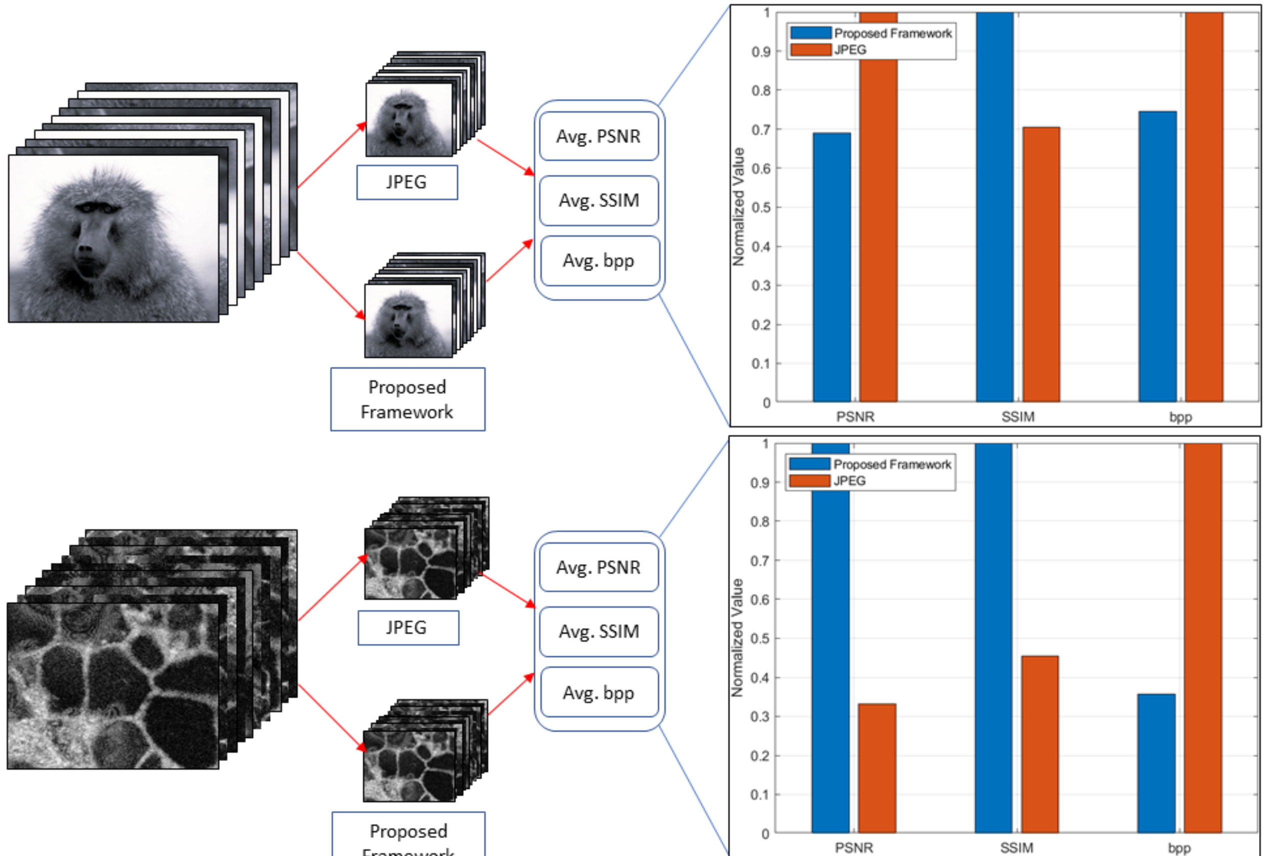

Figure 4: Comparative image quality assessment of two image datasets compressed using the proposed framework and conventional JPEG. The normalized average PSNR, SSIM, and bit-per-pixel (bpp) metrics are reported to illustrate the trade-off between compression efficiency and structural preservation across both texture-dominant and smooth-region images.

compared against JPEG at both quality level 5, resulting in bitrates of 0.094 bpp and 0.176 bpp, respectively. Rather than focusing solely on PSNR, this experiment evaluates structural preservation through Canny edge detection and gradient magnitude analysis, which are more directly related to geometric integrity and feature continuity.

Visual inspection of the reconstructed images depicted in Fig. 5 reveals that, although both methods operate in a severely constrained bitrate regime, their degradation characteristics differ fundamentally. JPEG introduces spatial blocking artifacts and local intensity discontinuities due to aggressive quantization of DCT coefficients within fixed-size blocks. These artifacts propagate into the edge maps, producing fragmented and irregular edge responses. In contrast, the proposed Fourier-domain codec preserves global phase relationships across the spectrum. Since phase encodes spatial localization and edge alignment, maintaining phase coherence even under coarse quantization results in structurally consistent edge contours.

The Canny edge maps confirm this observation. The proposed method retains a higher proportion of meaningful edge structures

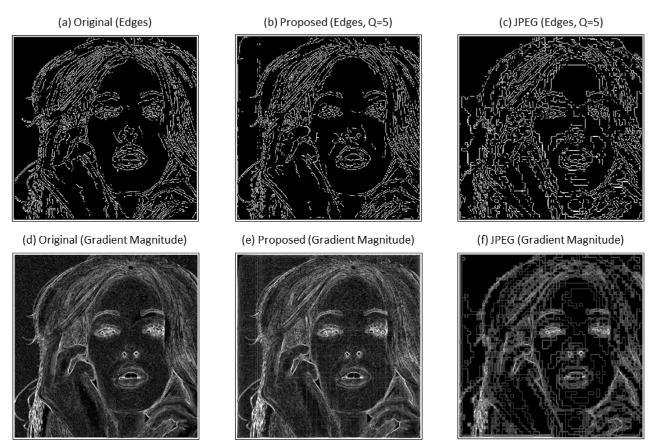

Figure 5: Comparison of edge and gradient preservation between the proposed framework and JPEG, showing improved structural continuity under phase-dominant coding.

Table 1: Comparison of Sizes of Image Batches Before and After Compression with Proposed Framework and JPEG.

Batch	Orig. Batch Size	Proposed Framework	JPEG
1	2.60 MB	28.25 KB	114.3 KB
2	2.82 MB	10.24 KB	15.94 KB

corresponding to facial boundaries, hair contours, and feature outlines, whereas JPEG exhibits spurious edges and partial edge suppression in high-frequency regions. Quantitatively, this behavior is reflected in the edge overlap metric, defined as the normalized intersection between reconstructed and reference edge maps. The higher overlap ratio achieved by the proposed codec indicates superior preservation of geometrically relevant structures relative to JPEG, despite operating at a lower bitrate.

Gradient magnitude analysis further supports this conclusion. The gradient maps of the proposed reconstruction exhibit spatial distributions closely aligned with the reference image, preserving edge strength localization. Conversely, the JPEG gradient map shows disrupted intensity transitions and localized artifacts corresponding to block boundaries. This demonstrates that magnitude-driven block quantization alters spatial derivatives more aggressively than phase quantization in the global Fourier domain.

These findings substantiate the central premise of the proposed paradigm: structural information is predominantly encoded in Fourier phase. By allocating coding resources toward preserving phase relationships while approximating high-frequency magnitude, the codec achieves graceful structural degradation even at ultra-low bitrates. The results highlight the suitability of the proposed phase-prioritized compression for applications where edge integrity and geometric consistency are critical, such as feature extraction, computer vision preprocessing, and morphology-sensitive imaging tasks.

3.3 Compression Efficiency and Dataset-Level Storage Reduction

Table 1 presents a quantitative comparison of dataset-level storage requirements before and after compression using the proposed phase-prioritized framework and conventional JPEG encoding. Two distinct image domains were evaluated: (i) Optical Coherence Microscopy (OCM) images representing high-frequency, structure-dominant biomedical data (Batch 1), and (ii) natural texture-rich baboon images (Batch 2). All input images were originally stored in lossless PNG format. The results demonstrate a substantial reduction in storage footprint using the proposed framework, achieving compression from 2.60 MB to 28.25 KB for OCM data and from 2.82 MB to 10.24 KB for natural images. In both cases, the proposed method yields markedly lower storage requirements than JPEG.

Large-scale imaging datasets particularly in biomedical and computer vision pipelines pose substantial challenges in storage, transmission bandwidth, and memory constraints during model training and inference. The proposed compression strategy enables aggressive dimensionality reduction at the dataset level while preserving structurally salient information, which is often more critical for downstream analytical tasks than pixel-wise fidelity. In OCM data, where morphological continuity and boundary information drive quantitative analysis, the structural bias of the phase-prioritized codec aligns well with task-oriented compression principles. Similarly, for natural images used in computer vision benchmarking, the method offers a favorable trade-off between compact representation and structural integrity.

4 Conclusion

This work introduced a phase-prioritized image compression paradigm that strategically reallocates coding resources toward preserving Fourier phase while approximating high-frequency magnitude through compact spectral priors. Across natural and biomedical datasets, the proposed framework demonstrated that, although it does not explicitly optimize conventional pixel-wise distortion metrics such as PSNR, it maintains strong structural robustness under extreme bitrate constraints, consistently preserving edge continuity, geometric consistency, and gradient coherence more effectively than magnitude-dominant JPEG compression at comparable or lower bitrates. By emphasizing structurally informative spectral components, the method provides a compact yet semantically meaningful representation of image data, offering a favorable balance between compression efficiency and structural fidelity. These results indicate that phase-dominant coding constitutes a viable alternative to traditional transform-based approaches, particularly in scenarios where preserving morphological and geometric structure is more critical than achieving high pixel-level reconstruction accuracy.

References

[1] S. Dhawan, “A Review of Image Compression and Comparison of its Algorithms,” *International Journal of Electronics & Communication Technology*, vol. 2, no. 1, pp. 22–26, Mar. 2011.

[2] S. Karmakar et al., "Image Compression and Decompression based on FFT algorithm," 2023 7th International Conference on Electronics, Materials Engineering & Nano-Technology (IEMENTech), Kolkata, India, 2023, pp. 1-6.

[3] M. H. Rasheed, O. M. Salih, M. M. Siddeq, and M. A. Rodrigues, “Image compression based on 2D Discrete Fourier Transform and matrix minimization algorithm,” Array, vol. 6, p. 100024, Jul. 2020.

[4] M. H. Asghari and B. Jalali, "Image compression using the Anamorphic Stretch Transform," IEEE International Symposium on Signal Processing and Information Technology, Athens, Greece, 2013, pp. 000233-000236.

[5] Manju Dabass, S. Vashisth, and R. Vig, “Lossy Color Image Compression Technique using Reduced Bit Plane-Quaternion SVD,” 2022 12th International Conference on Cloud Computing, Data Science & Engineering (Confluence), pp. 504–509, Jan. 2019.

[6] Ayman Alfalou, Marwa El Bouz, A. Mansour, and Gilels Keryer, “New spectral image compression method based on an optimal phase coding and the RMS duration principle,” Journal of Optics, vol. 12, no. 11, pp. 115403–115403, Sep. 2010.

[7] B. F. A. B H and P. R, "Overview on Machine Learning in Image Compression Techniques," 2021 Innovations in Power and Advanced Computing Technologies (i-PACT), Kuala Lumpur, Malaysia, 2021, pp. 1-8.

[8] S. Han, B. Mo, J. Zhao, J. Xu, S. Sun, and B. Jin, “High-Quality Image Compression Algorithm Design Based on Unsupervised Learning,” Sensors, vol. 24, no. 20, p. 6503, Oct. 2024.

[9] R. Kaur and P. Choudhary, “A Review of Image Compression Techniques,” International Journal of Computer Applications, vol. 142, no. 1, pp. 8–11, May 2016.

[10] D. Mishra, S. K. Singh, and R. K. Singh, “Deep Architectures for Image Compression: A Critical Review,” Signal Processing, vol. 191, p. 108346, Feb. 2022.

[11] M. Akada and H. Tabuchi, “Time course of eccentric macular hole formation after pars plana vitrectomy for epiretinal membrane detected by optical coherence tomography,” American Journal of Ophthalmology Case Reports, vol. 32, p. 101914, Dec. 2023.