Exploring the JPEG Algorithm: Impacts of Modifications in the Quantization and Transform Stages

Luís F. Krause de Castro[1] and Tanilson D. dos Santos[1]

[1] *Federal University of Tocantins, Computer Science Department, Tocantins, Brazil*

Reception date of the manuscript: 17/11/2025
Acceptance date of the manuscript: 14/02/2026
Publication date: 14/02/2026

Abstract— The JPEG algorithm is one of the most widely used methods for image compression, achieving substantial file size reduction by removing perceptually less significant information. Among its stages, quantization and frequency-domain transformation play a crucial role in determining the balance between visual quality and compression efficiency. This study explores structural modifications in both stages of the JPEG process: first, by testing alternative quantization matrices with varying levels of aggressiveness, and second, by replacing the standard DCT (Discrete Cosine Transform) with other mathematical transforms, such as Fourier, Laplace, and Wavelet. The results highlight how adjustments in quantization and transform selection can significantly influence the trade-off between efficiency and reconstructed image quality.

Keywords— JPEG algorithm, Transform, Image compression, Quantization, Frequency analysis.

I. Introduction

Image compression plays a fundamental role in modern computing, enabling the efficient storage and transmission of visual information while minimizing perceptual degradaFtion. With the exponential increase in multimedia data volume, driven by high-resolution imaging systems and global content sharing, compression techniques are indispensable for reducing bandwidth and storage requirements in both online and offline environments [1, 2].

Among the various compression techniques, the JPEG (Joint Photographic Experts Group) standard remains one of the most widely adopted lossy compression methods worldwide. Officially standardized as ISO/IEC 10918-1 in 1994 [3], JPEG relies on transforming the image from the spatial to the frequency domain using the DCT (Discrete Cosine Transform) and subsequently quantizing the coefficients based on perceptual properties of the human visual system [4, 5]. The quantization step, which discards high-frequency information less perceptible to human vision, is primarily responsible for the achieved compression ratio while maintaining acceptable visual quality [6, 7].

However, despite its proven efficiency and simplicity, JPEG applies a fixed quantization table and a single transform type (DCT) to all images, regardless of their specific spatial-frequency content. This limitation raises critical questions: how would the algorithm behave under structural modifications in its core stages? What are the perceptual and statistical impacts of adjusting quantization aggressiveness or replacing the DCT with alternative transforms such as the Fourier, Laplacian, or Wavelet? Prior research has explored some of these directions — evaluating alternative quantization strategies [8] and transform-based variations such as the DWT (Discrete Wavelet Transform) used in JPEG2000 [9]. Yet, there remains a lack of systematic comparative studies quantifying the impact of such modifications on image quality and compression ratio under controlled experimental conditions.

This work aims to fill this gap by investigating how modifications in the quantization and transform stages influence the balance between compression efficiency and reconstructed image quality in the JPEG pipeline. Specifically, three quantization matrices — standard, moderate, and aggressive — were evaluated alongside four different transforms: DCT, DFT (Discrete Fourier Transform), Laplacian, and DWT.

Despite the emergence of advanced compression standards and neural codecs, JPEG remains the most widely adopted image format worldwide due to its simplicity, hardware support, and backward compatibility [2]. Investigating how subtle modifications in its core stages—transform and quantization—impact compression efficiency and perceptual quality remains highly relevant, particularly for embedded or bandwidth-limited applications where modern codecs are computationally prohibitive.

Contact data: Luís F. Krause de Castro, lfelipe.krause@gmail.com

II. Related Work

Several studies have explored alternative transforms and quantization strategies within the JPEG framework. Watson [10] proposed perceptual weighting of DCT coefficients based on the sensitivity characteristics of the human visual system, leading to visually optimized quantization matrices. More recently, Liu et al. [11] investigated adaptive quantization guided by machine-learning-based saliency models, demonstrating the potential of perceptual adaptivity in traditional compression schemes.

Other works have examined alternative transforms to the DCT. Parmar [12] showed that wavelet-based approaches can achieve higher compression ratios and superior performance at low bitrates, albeit with increased computational cost. The JPEG2000 standard later formalized the use of the Discrete Wavelet Transform, enabling both lossy and lossless compression within the same framework [9]. Despite these advances, the DCT remains predominant in practice due to its simplicity, compatibility, and efficient hardware implementation.

Alternative quantization methods have also been extensively studied. Delp and Mitchell [8] demonstrated that customized quantization matrices can preserve perceptual quality even under higher compression levels. More recent works by Alakuijala et al. [13] and Ma et al. [14] employed perceptual and learning-based models to optimize quantization adaptively, showing that quantization is inherently tied to the characteristics of the human visual system [7, 6]. These studies emphasize that low-frequency components contribute most to perceived quality, supporting the selective coefficient suppression strategy used in JPEG. Recent neural and hybrid models revisit this principle, learning to reproduce similar perceptual trade-offs through data-driven optimization [14].

In summary, although modern codecs and neural compression models continue to evolve rapidly, the classical JPEG algorithm remains an essential baseline for both academic research and practical applications, owing to its transparency, interpretability, and computational efficiency.

III. Methodological Procedures

The experiments were designed to isolate the influence of each variable — transform type and quantization aggressiveness — while keeping the remaining stages consistent with the standard JPEG implementation.

The general workflow of the research methodology — from algorithm implementation to experimental evaluation — is summarized in Fig.1, which outlines the main development and analysis stages.

a. Implementation and Tools

All experiments were implemented in *Python*, leveraging open-source scientific and analytical libraries to ensure transparency and reproducibility. The development environment was designed to emulate the full JPEG compression pipeline while enabling flexible modifications in the quantization and transform stages. The following tools and resources were employed:

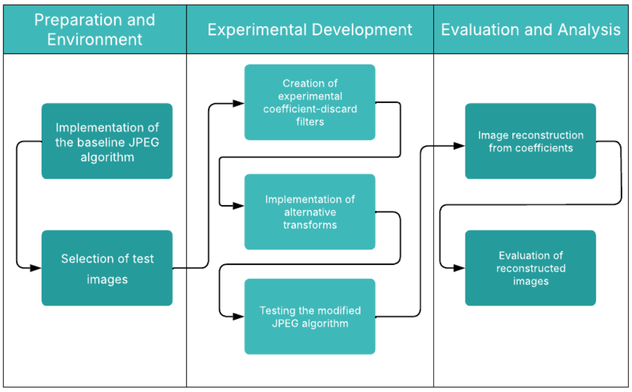

Fig. 1: Flowchart summarizing the main stages of the research process, including preparation, experimental development, and evaluation.

- **OpenCV (cv2)**: for loading, displaying, and performing basic image manipulation, including color space conversion and block segmentation;
- **NumPy (numpy)**: for matrix operations and numerical computation, including splitting and merging of 8×8 image blocks;
- **SciPy (scipy.fftpack)**: for computing the two-dimensional DCT and its inverse (IDCT);
- **NumPy FFT (numpy.fft)**: for computing the two-dimensional DFT and its inverse (IDFT);
- **PyWavelets (pywt)**: for implementing wavelet-based compression using the Haar basis [15];
- **scikit-image (skimage.metrics)**: for computing objective image quality metrics such as PSNR (Peak Signal-to-Noise Ratio) and SSIM (Structural Similarity Index) [16, 7];
- **Collections (Counter, namedtuple)** and **Heapq**: for implementing Huffman coding, including frequency counting, tree construction, and code generation;
- **Matplotlib (matplotlib.pyplot)** and **Seaborn**: for visualizing images and generating statistical plots for result interpretation;
- **OS**: for file and directory handling and retrieving file sizes for compression ratio computation;
- **Pandas**: for organizing and tabulating experimental results, including metric aggregation and compression statistics.

The full source code and dataset used in the experiments are publicly available in a dedicated GitHub repository (https://github.com/LuisFelipeKrause/JPEG-Algorithm-TCC) to promote reproducibility and further research.

b. Dataset and Experimental Setup

The experiments were conducted using a balanced dataset composed of 420 color and grayscale images of 512×512 pixels, equally drawn from two widely recognized databases:

 ISSN: 2675-3588

- **USC-SIPI Image Database** [17]: includes natural, aerial, and synthetic scenes with varied textural complexity and contrast levels;
- **ImageNet Sample Subset** [18]: provides a diverse range of natural and structured images, ensuring generalization across different spatial frequency distributions.

The compression process was then executed under two independent experimental conditions:

1. **Quantization Variation:** three quantization matrices were tested: the standard JPEG table, a *moderate* version (proposed by the author) with slightly increased coefficient suppression, and an *aggressive* version designed to enhance compression at the expense of detail loss.
2. **Transform Variation:** four frequency-domain transforms were compared: the standard DCT, the Discrete Fourier Transform (DFT), the discrete Laplacian operator, and the Discrete Wavelet Transform.

c. Evaluation Metrics

To ensure a consistent and objective comparison across experiments, three primary quantitative metrics were used:

- **PSNR (Peak Signal-to-Noise Ratio):** measures the fidelity between the original and reconstructed images in decibels (dB). Higher PSNR values indicate smaller mean-squared error and better visual preservation. Typical thresholds classify images above 35 dB as high quality [16].
- **SSIM (Structural Similarity Index):** quantifies structural and luminance similarity between the original and compressed images, yielding values from 0 to 1, where values closer to 1 denote greater perceptual similarity [7].
- **Compression Ratio (CR):** measures storage efficiency, computed as the ratio between the uncompressed and compressed file sizes.

IV. Results and Discussion

To analyze the practical behavior of the JPEG algorithm, a series of experiments were designed to test how specific stages of the compression pipeline influence the balance between image quality and compression efficiency. Two complementary sets of experiments were carried out: one focused on modifying the quantization parameters, and another on replacing the mathematical transform used in the frequency-domain representation. Together, these analyses allow a deeper understanding of how each stage contributes to the final compression outcome and how structural adjustments can optimize performance for different image types.

a. Impact of Quantization Matrices

Three quantization tables were used:

- **Standard**: the quantization table used in the conventional JPEG algorithm;

(a) Standard JPEG quantization mask

16	11	10	16	24	40	51	61
12	12	14	19	26	58	60	55
14	13	16	24	40	57	69	56
14	17	22	29	51	87	80	62
18	22	37	56	68	109	103	77
24	35	55	64	81	104	113	92
49	64	78	87	103	121	120	101
72	92	95	98	112	100	103	99

- **Moderate**: a quantization table proposed by the author with slightly higher discarding compared to the standard table;

(b) Moderate quantization mask

40	43	45	50	60	70	90	100
43	45	50	60	70	90	100	110
45	50	60	70	90	100	110	128
50	60	70	90	100	110	128	128
60	70	90	100	110	128	128	128
70	90	100	110	128	128	128	128
90	100	110	128	128	128	128	128
100	110	128	128	128	128	128	128

- **Aggressive**: a quantization table proposed by the author with significantly higher discarding.

(c) Aggressive quantization mask

80	85	90	100	120	140	180	200
85	90	100	120	140	180	200	220
90	100	120	140	180	200	220	255
100	120	140	180	200	220	255	255
120	140	180	200	220	255	255	255
140	180	200	220	255	255	255	255
180	200	220	255	255	255	255	255
200	220	255	255	255	255	255	255

The standard table follows the JPEG specification [2], while the moderate and aggressive versions were designed by the author to incrementally increase coefficient suppression in the high-frequency region.

Fig.2 illustrates the compression ratio distributions for the three quantization settings. As expected, more aggressive quantization yielded higher compression ratios, with the aggressive table achieving an average ratio above 14 : 1, compared to approximately 11 : 1 for the standard configuration. The moderate matrix achieved intermediate results, maintaining lower variance, which indicates more predictable compression behavior across diverse image types.

The relationship between compression and quality follows the expected inverse correlation. Fig.3 and Fig.4 show that the standard quantization matrix achieved the best perceptual balance, with median PSNR values around 34 dB and SSIM near 0.93—values typically considered visually lossless for photographic content [16]. The moderate quantization introduced only a small degradation (PSNR $\approx$ 30 dB; SSIM $\approx$ 0.88) while improving compression efficiency by

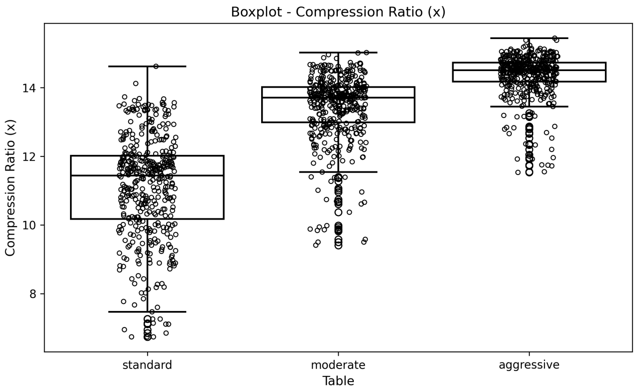

Fig. 2: Distribution of compression ratios for different quantization matrices.

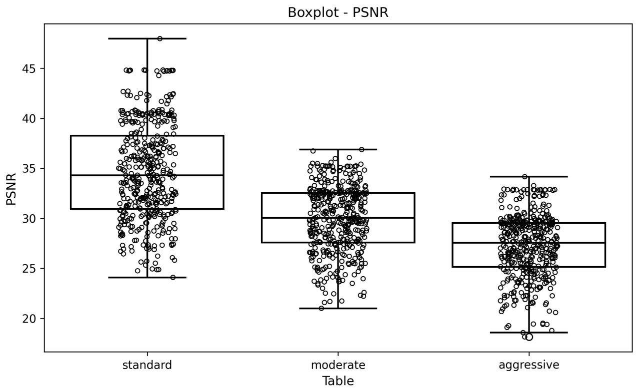

Fig. 3: PSNR distribution across quantization matrices.

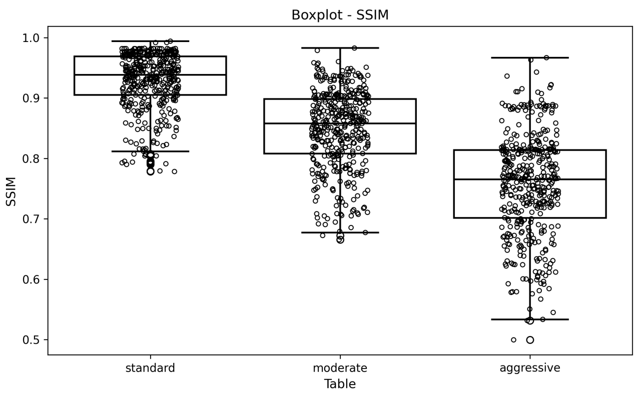

Fig. 4: SSIM distribution across quantization matrices.

roughly 20%. In contrast, the aggressive matrix resulted in substantial degradation, especially around edges and high-frequency textures, dropping PSNR to $\approx$ 27 dB and SSIM below 0.8.

The scatter plots in Fig.5 and Fig.6 provide a clearer view of the trade-off between visual quality and compression ratio. The results reveal a strong negative correlation between PSNR/SSIM and compression ratio ($r \approx -0.85$), consistent with established compression literature [7]. The standard quantization occupies the upper-left region, corresponding to higher quality at lower compression, while the aggressive setup shifts toward higher compression but at a significant perceptual cost.

From a perceptual perspective, visual inspection confirmed that the moderate quantization preserved overall

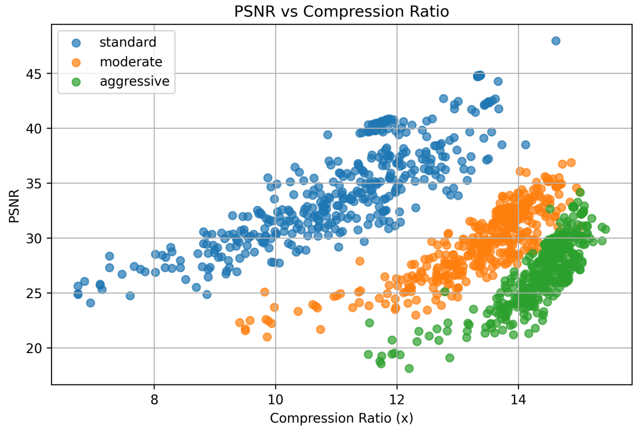

Fig. 5: Correlation between PSNR and compression ratio for quantization matrices.

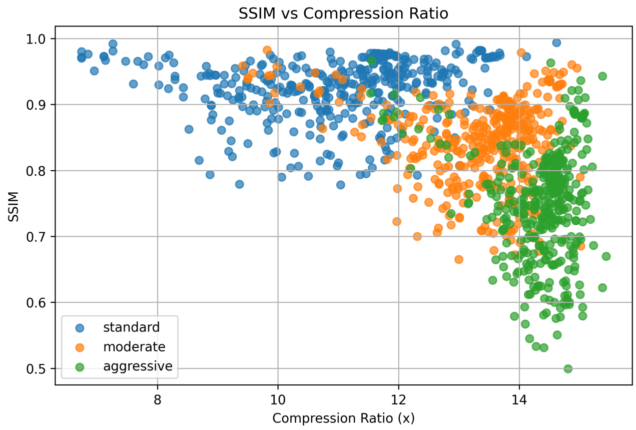

Fig. 6: Correlation between SSIM and compression ratio for quantization matrices.

image structure and fine details in most scenarios. Edges remained stable, and color transitions exhibited limited blocking artifacts.

To illustrate the perceptual impact of quantization strength, Fig.7, Fig.8b, Fig.8a and Fig.8c shows a comparison between the original image and the reconstructions obtained using moderate, standard, and aggressive quantization tables. While the standard configuration preserves most structural details, the moderate quantization (Fig.8b) achieves a superior balance, effectively preserving overall image structure and fine details in most scenarios. In contrast, the aggressive table leads to noticeable loss of fine detail and pronounced blocking artifacts, particularly in regions with high spatial frequency.

The visual comparison clearly illustrates that increasing quantization aggressiveness reduces the preservation of fine spatial detail. The standard quantization table maintains edges and texture transitions with minimal distortion, whereas the aggressive table produces noticeable blocking artifacts and smoothing of high-frequency regions. This behavior aligns with the fact that quantization discards smaller DCT coefficients first, which correspond to fine detail and texture, causing visually perceptible degradation when the quantization intervals become too coarse.

Fig. 7: Original Image

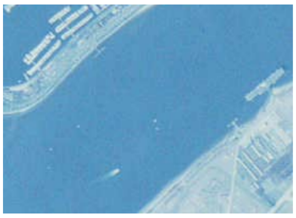

(a) Standard Quantization

(b) Moderate Quantization

(c) Aggressive Quantization

Fig. 8: Comparison of zoomed-in image regions using different quantization levels.

b. Impact of Transform Substitution

Four two-dimensional mathematical transforms were implemented and compared, each representing distinct mathematical and perceptual characteristics:

- **2D Discrete Cosine Transform (2D-DCT, Type-II)**: employed with orthogonal normalization, identical to the transform used in the original JPEG standard, serving as the baseline for comparison;
- **2D Discrete Fourier Transform (2D-DFT)**: computed using the Fast Fourier Transform (FFT) algorithm provided by NumPy, enabling efficient frequency-domain analysis with complex-valued coefficients;
- **Discrete Laplacian Operator (2D Laplace)**: applied as a local differential operator on each 8×8 block to capture rapid intensity variations, emphasizing edges and spatial discontinuities;
- **2D Discrete Wavelet Transform (2D-DWT)**: implemented using the Haar wavelet basis via the PyWavelets library, providing a multiresolution representation through hierarchical decomposition of approximation and detail coefficients [15].

These transforms were selected to represent distinct mathematical domains and properties: frequency-based (DCT and DFT), spatial-differential (Laplacian), and multiscale

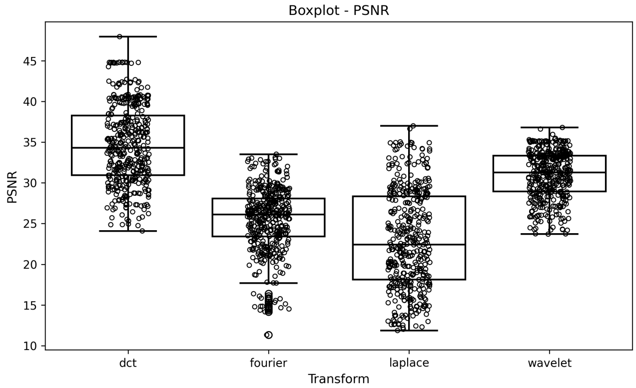

Fig. 9: PSNR distribution for different transform domains.

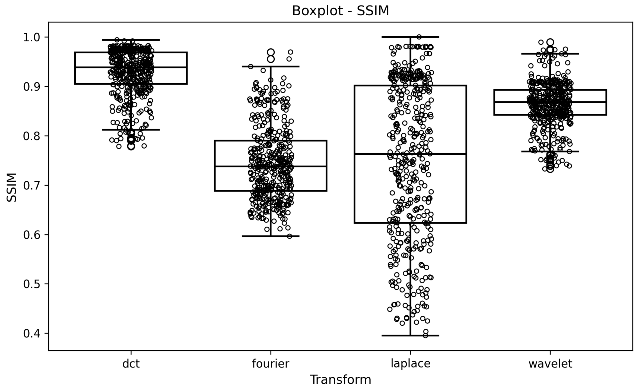

Fig. 10: SSIM distribution for different transform domains.

(Wavelet). Their inclusion enables a comparative analysis of how different frequency and spatial representations affect compression efficiency and reconstructed image quality.

To ensure a controlled evaluation, the quantization matrix was kept fixed (standard JPEG table) across all tests, allowing the analysis to focus exclusively on the impact of the transform domain on objective and perceptual quality metrics.

Fig.9 and Fig.10 summarize the statistical distributions of PSNR and SSIM across all images for each transform. The DCT and DWT outperformed the others, with the DWT exhibiting slightly higher perceptual scores and lower variance, suggesting a more stable reconstruction quality across texture types. The DFT suffered from energy dispersion across the frequency spectrum, while the Laplacian operator—being spatially local but not energy-compacting—produced the weakest compression and visible artifacts in smooth regions.

Fig.11 and Fig.12 reinforce these trends by showing the correlation between compression ratio and perceptual quality. The DWT maintained consistent PSNR and SSIM even as compression increased, reflecting its superior multi-resolution energy compaction. The DCT maintained a close second position, confirming its efficiency and computational advantage. Both the Laplacian and DFT, despite their theoretical relevance, demonstrated poor rate-distortion trade-offs, reaffirming that their energy distributions are less compatible with scalar quantization schemes [9].

In summary, the experiments confirm that while the DCT remains the most computationally efficient and robust choice for JPEG compression, wavelet-based transforms such as

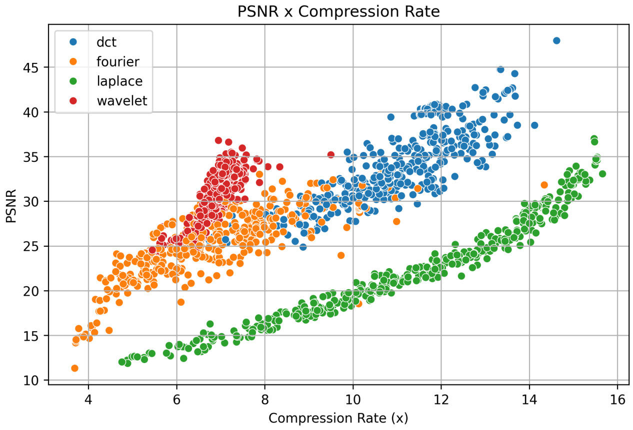

Fig. 11: Relationship between PSNR and compression ratio across transforms.

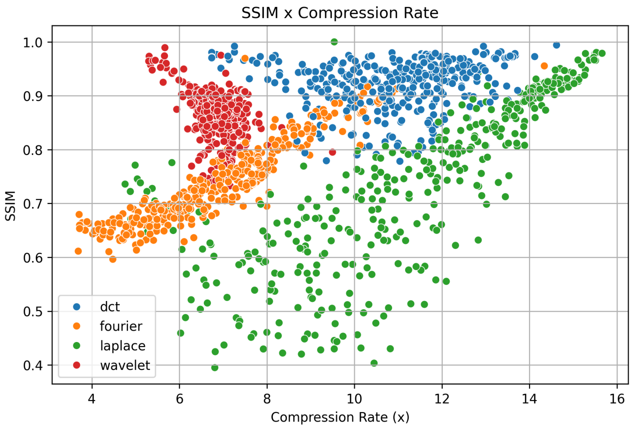

Fig. 12: Relationship between SSIM and compression ratio across transforms.

Fig. 13: Original Image

DWT can achieve better perceptual results at equivalent bitrates. This aligns with the motivation behind the development of the JPEG2000 standard [9], which replaced block-based DCT with DWT to eliminate blocking artifacts and improve scalability.

Fig.14a, Fig.14b, Fig.14c and Fig.14d compares the reconstructed images obtained after applying different transform domains in the JPEG compression pipeline, while keeping the same quantization table. The DCT reconstruction closely resembles the original image, confirming its strong energy compaction properties. The DFT reconstruction preserves global structure but introduces ringing artifacts due to spectral spreading. The Laplacian operator emphasizes edges excessively, creating over-sharpened regions with visual noise. The DWT (Haar) reconstruction, in contrast, provides smoother gradients and reduced blocking artifacts, benefiting from multiresolution spatial-frequency localization.

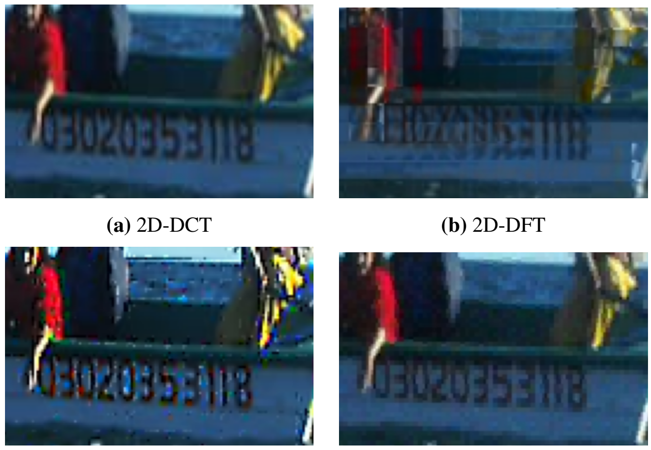

(a) 2D-DCT **(b)** 2D-DFT

(c) Laplacian Transform **(d)** 2D-DWT (Haar) Transform

Fig. 14: Comparison of zoomed-in image regions obtained using different frequency and spatial domain transforms: DCT, DFT, Laplacian, and DWT (Haar).

The transform comparison highlights how each mathematical basis influences the distribution and reconstruction of image information. The DCT provides compact energy representation and yields reconstructions that closely resemble the original image, justifying its adoption in the JPEG standard. The DFT, by spreading energy globally, introduces ringing artifacts near edges. The Laplacian operator overemphasizes local intensity variations, resulting in oversharpened and noisy regions. In contrast, the DWT preserves structural continuity and smooth gradients due to its multiresolution nature, reducing blocking artifacts and improving perceptual quality at similar compression rates.

c. Discussion and Interpretation

The experimental findings reveal consistent patterns across both the quantization and transform analyses, emphasizing that small structural modifications in the JPEG pipeline can lead to measurable variations in compression efficiency and perceptual fidelity. The trade-off between these two dimensions—data reduction and visual preservation—remains at the core of lossy image compression, yet the results obtained here provide a more granular understanding of how each stage contributes to this balance.

From the quantization experiments, the data confirm that perceptual quality degradation follows a nonlinear curve with respect to quantization aggressiveness. The moderate quantization table demonstrated a near-optimal point on this curve, achieving approximately 20–25% better compression than the standard matrix with negligible perceptual loss. This suggests that the standard JPEG quantization matrix, originally hand-tuned in the early 1990s [2], may still be improved through statistically guided parameter adjustment. Similar findings have been reported in perceptual compression studies that adapt quantization weights based on human visual sensitivity models [7], as well as in modern machine learning–based codecs [14].

Regarding the transform stage, the results indicate that the DWT provides the most effective energy compaction among the tested alternatives, producing higher PSNR and SSIM scores. This superiority is linked to the DWT's ability to represent spatially localized frequency components and maintain sharp edge continuity [15, 9]. Notably, the DCT produces a considerably smoother visual output in textured areas, which may be preferable in specific aesthetic contexts, even though it is slightly less efficient in energy compaction than the DWT. Despite this smoothness and the DCT's computational practicality for hardware, the DWT remains qualitatively superior in preserving complex structural details. By contrast, the Laplacian and DFT-based schemes showed weaker performance due to overemphasis of high-frequency components and lack of spatial localization, respectively [19].

A key observation emerging from both experimental fronts is the strong correlation between perceptual metrics (SSIM) and human visual tolerance to high-frequency suppression. Consequently, moderate quantization and wavelet-based transforms succeed because they selectively preserve low-frequency content while distributing quantization error across perceptually less relevant frequencies.

Methodologically, these results demonstrate that perceptual improvements can be achieved without altering the bitstream structure, ensuring backward compatibility and facilitating the integration of adaptive quantization or hybrid DCT–DWT into existing JPEG encoders. Consequently, these findings reinforce the continuing relevance of the classic transform–quantization framework, even in the era of deep learning-based compression. By systematically quantifying these modifications, this study provides empirical evidence for the development of perceptually optimized JPEG variants tailored for edge devices and low-latency applications

V. Concluding Remarks

This work presented a systematic experimental evaluation of the influence of frequency-domain transforms and quantization strategies on the performance of the JPEG compression pipeline. By maintaining the original JPEG encoding structure and independently varying the transform and quantization components, it was possible to isolate their respective contributions to perceptual quality and compression efficiency.

The results demonstrated that the **Discrete Cosine Transform** remains a robust and computationally efficient choice for image compression, offering an optimal balance between energy compaction and simplicity of implementation. However, the **Discrete Wavelet Transform**—particularly when used with the Haar basis—showed superior perceptual performance (higher PSNR and SSIM) at comparable compression rates. This confirms the theoretical advantage of wavelet representations in localizing both spatial and frequency information [15, 9].

In terms of quantization, the results suggest that **moderately aggressive quantization tables** can provide substantial reductions in bit rate (up to 25%) while maintaining high perceptual fidelity. This finding indicates that the standard JPEG quantization matrix, originally designed through empirical tuning [2], can be improved via data-driven optimization or perceptual weighting models [7]. The analysis also reinforces that excessive quantization, although improving compression ratio, quickly deteriorates structural similarity and perceived image quality.

Overall, the experiments highlight the enduring relevance of transform-based compression models and underscore that meaningful improvements to the JPEG standard are still achievable without altering its bitstream structure or entropy coding stage. Such backward-compatible modifications are of great practical importance, particularly for embedded systems, low-bandwidth communications, and real-time applications, where deep learning–based codecs may be computationally prohibitive.

Future work includes exploring perceptually guided quantization based on saliency or attention maps, as well as hybrid transform approaches that combine DCT and DWT domains for improved rate–distortion trade-offs. Machine-learning-based optimization of quantization tables also represents a promising direction for enhancing compression adaptivity.

References

[1] D. R.-J. G.-J. Rydning, J. Reinsel, and J. Gantz, "The digitization of the world from edge to core," *Framingham: International Data Corporation*, vol. 16, pp. 1–28, 2018.

[2] G. K. Wallace, "The jpeg still picture compression standard," *Communications of the ACM*, vol. 34, no. 4, pp. 30–44, 1991.

[3] "Information technology – digital compression and coding of continuous-tone still images – requirements and guidelines," International Organization for Standardization, Geneva, Switzerland, Tech. Rep. ISO/IEC 10918-1:1994, 1994.

[4] W. B. Pennebaker and J. L. Mitchell, *JPEG: Still image data compression standard*. Springer Science & Business Media, 1992.

[5] N. Ahmed, T. Natarajan, and K. R. Rao, "Discrete cosine transform," *IEEE transactions on Computers*, vol. 100, no. 1, 2006.

[6] R. C. Gonzalez and R. E. Woods, *Processamento Digital de Imagens*, 3rd ed. São Paulo: Pearson Prentice Hall, 2010.

[7] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.

[8] E. Delp and O. Mitchell, "Image compression using block truncation coding," *IEEE transactions on Communications*, vol. 27, no. 9, pp. 1335–1342, 1979.

[9] D. S. Taubman, M. W. Marcellin, and M. Rabbani, "Jpeg2000: Image compression fundamentals, standards and practice," *Journal of Electronic Imaging*, vol. 11, no. 2, pp. 286–287, 2002.

[10] A. B. Watson, "Dct quantization matrices visually optimized for individual images," in *Human vision, visual processing, and digital display IV*, vol. 1913. SPIE, 1993, pp. 202–216.

[11] J. Chiu, "Modeling image quantization tradeoffs for optimal compression," *arXiv preprint arXiv:2112.07207*, 2021.

[12] H. M. Parmar, "Comparison of dct and wavelet based image compression techniques," *International journal of engineering development and research*, vol. 2, no. 1, pp. 664–669, 2014.

[13] J. Alakuijala, R. Obryk, O. Stoliarchuk, Z. Szabadka, L. Vandevenne, and J. Wassenberg, "Guetzli: Perceptually guided jpeg encoder," *arXiv preprint arXiv:1703.04421*, 2017.

[14] S. Ma, X. Zhang, C. Jia, Z. Zhao, S. Wang, and S. Wang, "Image and video compression with neural networks: A review," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 6, pp. 1683–1698, 2019.

[15] I. Daubechies, *Ten lectures on wavelets*. SIAM, 1992.

[16] A. Hore and D. Ziou, "Image quality metrics: Psnr vs. ssim," *2010 20th International Conference on Pattern Recognition*, pp. 2366–2369, 2010.

[17] USC Signal and Image Processing Institute - Department of Electrical and Computer Engineering, "The usc-sipi image database," 2025, available at: https://sipi.usc.edu/database/database.php. Accessed on: September 5, 2025.

[18] Schwartz, Pullen, "Imagenet sample images," 2025, available at: https://github.com/EliSchwartz/imagenet-sample-images. Accessed on: September 5, 2025.

[19] A. V. Oppenheim, A. S. Willsky, and S. H. Nawab, *Signals & systems*. Pearson Educación, 1997.

ISSN: 2675-3588