A Reversible Grayscale Method Based on Bit-Field Multi-Channel Fusion Encoding

Teng Wang
School of Mechatronic Engineering
Guangdong Polytechnic Normal University
Guang Zhou, China
Department of R&D
OPT Machine Vision
Dong Guan, China
wangteng@gpnu.edu.cn

Wei Pan*
Department of R&D
OPT Machine Vision
Dong Guan, China
vpan@foxmail.com
*Co-first author

Yong Yang
School of Mechatronic Engineering
Guangdong Polytechnic Normal University
Guang Zhou, China
yy2008@gpnu.edu.cn

Pascal Lefèvre*
School of AI and Advanced Computing
Xi'an Jiaotong - Liverpool University
Taicang campus, Suzhou, China
Pascal.Lefevre@xjtlu.edu.cn
*Corresponding author

***Abstract*—Grayscale image representation is widely adopted to reduce computational complexity, optimize storage, and enhance transmission efficiency. However, in many applications, restoring the original color information from grayscale images is highly desirable. Existing methods, often suffer from high computational costs, limited interpretability, and dependence on training data. To address these limitations, we propose a novel non-learning-based approach that encodes RGB color information into a 16-bit grayscale image by strategically embedding HSV components, allowing high-fidelity recolorization while remaining lightweight and hardware-friendly. Experimental results validate the effectiveness of our approach in balancing reconstruction accuracy, computational efficiency, and practicality. The proposed method offers a promising alternative to traditional and neural network-based solutions, particularly in resource-constrained environments.**

***Index Terms*—Reversible Grayscale conversion, Bit-field Encoding, HSV, Image Processing, Color Restoration, non-learning-based method, edge computing**

I. Introduction

Grayscale image representation is widely used to reduce computational complexity, streamline analysis, and optimize storage efficiency. Techniques such as the Discrete Cosine Transform (DCT) enable effective image compression and reconstruction, lowering storage costs and improving transmission efficiency—critical advantages for digital imaging applications. However, in many practical scenarios, restoring the original color information from grayscale images is highly desirable. Existing restoration methods, including histogram shifting and wavelet packet transforms, aim to recover color data while preserving grayscale fidelity.

Traditional reversible grayscale methods, including histogram shifting [1] and wavelet-packet-based schemes [2], focus on embedding color cues into the luminance channel. These approaches usually impose additional transform-domain constraints or require multiple auxiliary channels, leading to increased computational complexity and memory overhead. More recently, deep learning–based techniques [3], [4] have achieved impressive reconstruction fidelity but entail high inference latency, large model sizes, and a strong dependence on labeled colorized-grayscale training pairs.

Deep learning-based reversible grayscale techniques achieve high-fidelity restoration but suffer from significant drawbacks, including high computational costs, limited interpretability, and strong dependence on training data. While recent neural network-based approaches have made strides in improving interpretability and efficiency, these challenges persist.

In this paper, we propose a novel *non-learning-based* reversible RGB-to-grayscale conversion method that encodes chromatic information directly in the HSV color space using a compact 16-bit representation. By strategically allocating bit-lengths to the Hue (H), Saturation (S), and Value (V) components, we embed color data into a single 16-bit grayscale image. The key advantages of our approach include:

- **Elimination of training overhead**: Unlike neural network-based methods, our technique requires no training data or model optimization, ensuring rapid execution.
- **Hardware compatibility**: The algorithm employs only fundamental quantization and bitwise operations, rendering it highly suitable for implementation on embedded and edge computing platforms.
- **Full interpretability and reversibility**: The encoding and decoding processes are entirely transparent and mathematically invertible, permitting precise color reconstruc-

tion (subject only to quantization error).

To summarize, our contributions are as follows:

1) We propose a novel bit-field multi-channel fusion scheme in HSV space to convert an RGB image into a reversible 16-bit grayscale representation.
2) We analyze the impact of different bit allocations (a, b, c) for H, S, and V, and demonstrate how to choose allocation strategies based on image characteristics.
3) We conduct comprehensive experiments comparing our method against state-of-the-art reversible-grayscale techniques (both traditional and deep learning–based), showing comparable or better reconstruction accuracy with orders-of-magnitude lower computational cost.

The rest of this paper is organized as follows: Section II reviews related work; Section III details the proposed encoding and decoding algorithms, including bit allocation strategies and implementation considerations. Section IV presents quantitative and qualitative experimental results. and Section VI concludes and outlines future directions.

II. Related Works

Reversible color-to-grayscale (C2G) conversion aims to generate visually natural grayscale images while embedding sufficient color information to enable faithful reconstruction. Previous works can be broadly categorized into hand-crafted and learning-based approaches.

Handcrafted reversible C2G methods focus on embedding chrominance information using data hiding or watermarking. The early foundational works by Ni et al. [1] and Chen and Wornell [2] introduced reversible data hiding (RDH) using histogram shifting and wavelet packet decomposition, respectively. Chan et al. [5] designed a framework that combines grayscale conversion and reversible watermarking. Xu and Chan [6] further enhanced visual quality using half-toning...

Recent advances explore hybrid schemes that combine imperceptibility and robustness. Liang et al. [7] utilized clustering and reversible watermarking, while their later work [8] embedded compressed chrominance via high-capacity RDH. Fragoso-Navarro et al. [9] integrated inverse colorization with visible-imperceptible watermarking for content protection. Yu et al. [10] combined self-embedding, half-toning, and vector quantization to achieve high-quality reversible conversions. Zhu et al. [11] introduced a content-feature–enhanced histogram method (CFEH) with RDH to improve the fidelity of the decolorization.

Several studies aim to increase embedding capacity and security. Hou et al. [12] proposed grayscale-invariant RDH, while Jhong and Wu [13] focused on boosting embedding payload. Yang et al. [14] addressed visual quality using contrast-sensitive watermarking. Shi et al. [15] provides a broad survey of RDH progress over two decades.

Learning-based reversible methods leverage deep neural networks to embed and reconstruct color information. Xia et al. [3] introduced an invertible model that encodes color into grayscale with high perceptual quality. Zhao et al. [4] utilized invertible neural networks and wavelet transforms to preserve spatial structure and detail. Liu et al. [16] developed a JPEG-robust scheme using adversarial training. These models typically outperform hand-crafted methods in reconstruction quality, although at a higher computational cost and reduced interpretability.

Perceptual evaluation of these methods is also essential. Zhang et al. [17] proposed FSIM for structural fidelity, while Zhang et al. [18] demonstrated the strength of deep features as perceptual metrics. Foundational techniques such as JPEG compression [19] and wavelet decomposition [20] are also commonly embedded into reversible grayscale pipelines.

In summary, hand-crafted methods emphasize interpretability and lightweight computation, while learning-based and hybrid approaches yield superior fidelity and robustness. Recent research efforts focus on balancing these trade-offs to achieve secure, high-quality, and fully reversible grayscale transformations.

III. A new reversible color-to-grayscale framework

Our reversible grayscale framework (shown in figure 1) consists of two modules: encoding RGB images into a 16-bit grayscale representation and decoding back to RGB. We use the HSV color space for its clear separation of chromatic (H, S) and luminance (V) components, enhanced by adaptive bit allocation and error-aware quantization.

A. Adaptive Bit Allocation

We allocate three bits: a, b, and c to the components H, S, and V, respectively, where $a + b + c = 16$. Unlike static allocation, our method dynamically adjusts bit distribution based on image statistics using the variance of each component of the HSV color space denoted as:

$$\begin{cases} \sigma_H^2 = \text{Var}(H) \\ \sigma_S^2 = \text{Var}(S) \\ \sigma_V^2 = \text{Var}(V) \end{cases} \tag{1}$$

Each bit is assigned proportionally with a normalized variance coefficient such as:

$$a = \left\lfloor 16 \cdot \frac{\sigma_H^2}{\sigma_H^2 + \sigma_S^2 + \sigma_V^2} \right\rfloor \tag{2}$$

$$b = \left\lfloor 16 \cdot \frac{\sigma_S^2}{\sigma_H^2 + \sigma_S^2 + \sigma_V^2} \right\rfloor \tag{3}$$

$$c = 16 - a - b \tag{4}$$

This ensures an optimal representation of dominant image features (e.g., prioritizing H for colorful images).

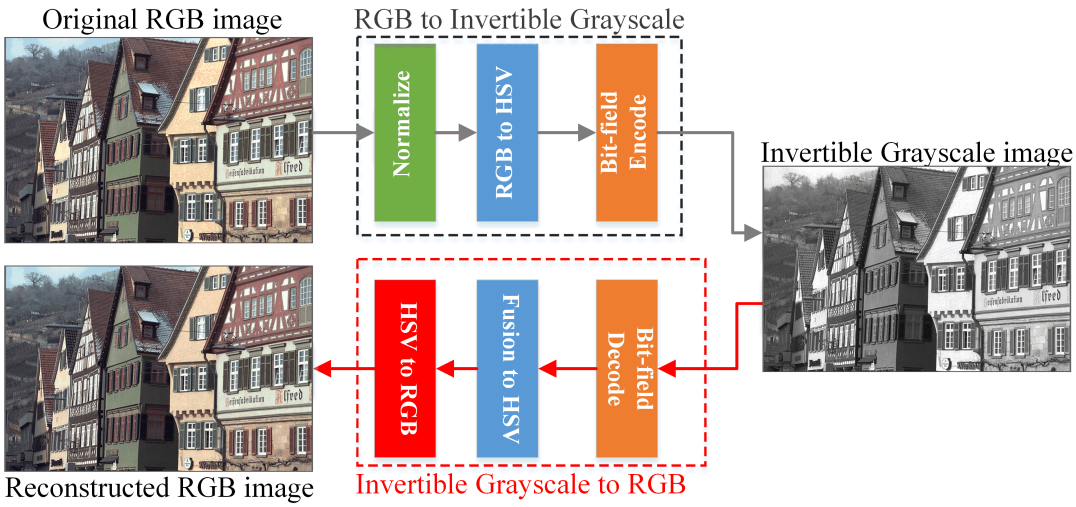

Fig. 1. The reversible grayscale framework.

B. Encoding: RGB to 16-bit Grayscale

The *encoding* process transforms an 8-bit RGB image into a 16-bit grayscale image. This is done by normalizing the RGB channels to the range $[0, 1]$ and then converting them to the HSV (Hue-Saturation-Value) color space. The hue component H is used to determine the dominant color of each pixel, which helps in prioritizing certain features over others.

More formally, the steps of the encoding process are described in the following. We define an 8-bit RGB image $I_{RGB} \in \{0, \ldots, 255\}^{H \times W \times 3}$.

1) **Normalization:** Normalize RGB channels to $[0, 1]$:

$$(R', G', B') = \frac{(R, G, B)}{255}.$$

2) **RGB to HSV Conversion:** Compute:

$$\begin{aligned} C_{\max} &= \max(R', G', B'), \\ C_{\min} &= \min(R', G', B'), \\ \Delta &= C_{\max} - C_{\min}. \end{aligned} \tag{5}$$

Then:

$$H = \begin{cases} 0, & \Delta = 0, \\ 60^\circ \times \left(\frac{G'-B'}{\Delta} \mod 6 \right), & C_{\max} = R', \\ 60^\circ \times \left(\frac{B'-R'}{\Delta} + 2 \right), & C_{\max} = G', \\ 60^\circ \times \left(\frac{R'-G'}{\Delta} + 4 \right), & C_{\max} = B'. \end{cases}$$

$$S = \begin{cases} 0, & C_{\max} = 0, \\ \frac{\Delta}{C_{\max}}, & \text{otherwise}, \end{cases} \qquad V = C_{\max}.$$

3) **Normalize Hue:** Scale H to $[0, 1)$: $H_{\text{norm}} = H/360^\circ$.

4) **Error-Aware Quantization:** Quantize H_{norm}, S, and V using non-linear mapping to prioritize perceptually sensitive regions:

$$\begin{aligned} H_q &= \lfloor \gamma_H(H_{\text{norm}}) \cdot (2^a - 1) \rfloor, \\ S_q &= \lfloor \gamma_S(S) \cdot (2^b - 1) \rfloor, \\ V_q &= \lfloor \gamma_V(V) \cdot (2^c - 1) \rfloor. \end{aligned} \tag{6}$$

where γ_H, γ_S, γ_V are gamma correction functions (e.g., $\gamma_V(x) = x^{0.45}$) to align with human perception.

5) **Pack into 16-bit Grayscale:** Combine using:

$$I_G = (H_q \ll (b + c)) \,|\, (S_q \ll c) \,|\, V_q.$$

The result is a 16-bit grayscale image $I_G \in \{0, \ldots, 2^{16} - 1\}^{H \times W}$.

C. Decoding: 16-bit Grayscale to RGB

Given the encoded grayscale image I_G:

1) **Extract Bit Fields:**

$$\begin{aligned} V_q &= I_G \,\&\, (2^c - 1), \\ S_q &= (I_G \gg c) \,\&\, (2^b - 1), \\ H_q &= I_G \gg (b + c). \end{aligned} \tag{7}$$

2) **Inverse Quantization:**

$$\begin{aligned} H_{\text{norm}} &= \gamma_H^{-1}\left(\frac{H_q}{2^a - 1} \right), \\ S &= \gamma_S^{-1}\left(\frac{S_q}{2^b - 1} \right), \\ V &= \gamma_V^{-1}\left(\frac{V_q}{2^c - 1} \right). \end{aligned} \tag{8}$$

Scale $H = H_{\text{norm}} \times 360^\circ$.

3) **HSV to RGB Conversion:** Use standard HSV-to-RGB transformation [21]:

$$\begin{aligned} C &= V \times S, \\ X &= C \times (1 - |(H/60^\circ) \mod 2 - 1|), \\ m &= V - C, \end{aligned} \tag{9}$$

Map H to the appropriate 60° sector to compute $(R'_{\text{temp}}, G'_{\text{temp}}, B'_{\text{temp}})$, then:

$$\begin{aligned} R' &= R'_{\text{temp}} + m, \\ G' &= G'_{\text{temp}} + m, \\ B' &= B'_{\text{temp}} + m. \end{aligned} \tag{10}$$

4) **Scale to 8-bit:**

$$\begin{aligned} R &= \lfloor R' \times 255 \rfloor, \\ G &= \lfloor G' \times 255 \rfloor, \\ B &= \lfloor B' \times 255 \rfloor. \end{aligned} \tag{11}$$

D. Bit Allocation Strategies

The bit allocation (a, b, c) balances quantization accuracy across channels:

- **Hue (H)**: Human perception is sensitive to hue changes; we recommend $a \geq 5$ (32 levels) to avoid noticeable color shifts.
- **Saturation (S)**: Larger b preserves saturation contrast in colorful images, but can be reduced for grayscale-dominant images.
- **Value (V)**: Luminance affects grayscale smoothness; $c \geq 6$ (64 levels) prevents banding artifacts.

TABLE I
EXAMPLE BIT ALLOCATION SCHEMES

Scheme	h (Hue)	s (Saturation)	v (Value)	Use Case
1	6	4	6	Balanced, general-purpose
2	5	5	6	High hue resolution, colorful images
3	4	5	7	High luminance precision, high-contrast images
4	6	5	5	Low hue priority, low color saturation images

E. Error-Aware Quantization

To address hue's circular nature, we apply a non-linear quantization function γ_H to minimize errors near 0°/360°. For low a (e.g., 4 bits), we interpolate H post-dequantization to smooth transitions. Error-diffusion techniques (e.g., Floyd-Steinberg) are applied during quantization to distribute errors spatially, enhancing visual quality.

IV. EXPERIMENTS

We evaluated our method on multiple datasets, including natural images, landscapes, and synthetic graphics, focusing on quantitative metrics, visual quality, computational efficiency, and storage overhead.

A. Experimental Setup

- **Datasets:**
  1) Kodak PhotoCD (24 lossless images, 768×512 or 512×768).
  2) Urban100 [22] (100 urban scenes, approximately 520×330 or 330×520).
  3) BSD500 [23] (500 natural landscapes, 481×321 or 321×481).
  4) McMaster [24] (18 images, 500×500)
- **Baselines:**
  1) IDN [4]
  2) CNN-based [3]
- **Metrics:** Mean Absolute Error (MAE), Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), Feature Similarity Index(FSIM) [17], and Perceptual Similarity Metric(LPIPS) [18].
- **Parameters:** Tested Schemes 1–4 (Table I) for our method; baselines used default settings.
- **Platform:** AMD Ryzen 9 7945HX, 16GB RAM. NVIDIA GeForce RTX 4060 (8GB).

B. Quantitative Results

Table II summarizes performance on the Kodak dataset.And Table III summarizes performance on the Urban100, BSD500 and McMaster dataset.

Our method achieves competitive MAE, PSNR, SSIM , FSIM [17] and LPIPS [18] , with Scheme 1 ($h = 6, s = 4, v = 6$) offering the best balance.

C. Qualitative Results

Fig. 2 shows a visual comparison for a representative image. Our method produces smooth grayscale images and high-fidelity RGB reconstructions, closely matching baselines while avoiding banding artifacts.

D. Computational and Storage Efficiency

Our method's encoding/decoding times (3.71/7.01 ms) on the KODAK dataset are significantly lower than baselines, owing to simple bitwise operations. Storage size (250 KB for 16-bit PNG/TIFF) is comparable to 24-bit RGB PNGs (250–270 KB). FPGA/ASIC implementations could further reduce latency. In addition, Our work still has certain gaps compared with the SOTA [11]. Liu et al. [11] proposed the CFEM, which has an average PSNR/FSIM index of 43.69/0.9897 in the Kodak dataset. This shows certain advantages over the index of our method, which is 34.80/0.9726. However, we will consider designing reversible compression technology or developing dynamic adaptive allocation algorithms in future work to improve the quality of the restored images while ensuring rapid restorability.

V. DISCUSSION AND FUTURE WORK

A. Methodological Innovations

To enhance generality and performance, we propose:

1) **Adaptive Bit Allocation**: Develop a lightweight module to estimate global color statistics (e.g., hue/saturation histograms) and dynamically allocate (a, b, c). For low-contrast images, prioritize V; for colorful images, emphasize H and S.
2) **Alternative Color Spaces**: Extend the method to Lab or YCbCr spaces for better perceptual uniformity, optimizing quantization for specific applications (e.g., medical imaging).
3) **Perceptual Quantization**: Apply non-linear mappings (e.g., gamma correction) to H, S, and V to align quantization errors with human perception, reducing visible artifacts.

TABLE II
QUANTITATIVE COMPARISON ON KODAK DATASET (AVERAGES)

Method	Bit Allocation	MAE	PSNR (dB)	SSIM	FSIM	LPIPS	Encode (ms)	Decode (ms)	Filesize (KB)
IDN [4]	N/A	**0.0025**	**42.16**	**0.9893**	**0.9844**	**0.0038**	175.30	125.24	721.47
CNN-based [3]	N/A	0.0097	37.40	0.9764	0.9599	0.0597	2486.75	2453.82	660.87
Ours (Scheme 1)	$h = 6, s = 4, v = 6$	**0.0152**	**34.80**	0.9726	0.9585	0.0577	3.71	7.01	633.20
Ours (Scheme 2)	$h = 5, s = 5, v = 6$	0.0159	33.66	**0.9779**	0.9679	**0.0543**	5.63	4.74	643.89
Ours (Scheme 3)	$h = 4, s = 5, v = 7$	0.0171	31.57	0.9757	**0.9698**	0.0840	5.02	6.30	684.94
Ours (Scheme 4)	$h = 6, s = 5, v = 5$	0.0252	30.83	0.9587	0.9360	0.0952	**3.21**	**4.49**	601.43

TABLE III
QUANTITATIVE COMPARISON ON URBAN100, BSD500 AND MCMASTER DATASET (AVERAGES)

Method	Bit Allocation	Dataset	MAE	PSNR (dB)	SSIM	FSIM	LPIPS	Encode (ms)	Decode (ms)
Ours (Scheme 1)	$h = 6, s = 4, v = 6$	Urban100	0.0169	34.06	0.9692	0.9588	0.0508	2.80	3.38
		BSD500	0.0162	34.48	0.9736	0.9623	0.0642	3.04	3.64
		McMaster	0.0199	32.20	0.9256	0.9455	0.0612	3.18	3.82
Ours (Scheme 2)	$h = 5, s = 5, v = 6$	Urban100	0.0180	32.88	0.9726	0.9646	0.0554	3.31	3.90
		BSD500	0.0176	33.00	0.9768	0.9668	0.0729	3.20	3.81
		McMaster	0.0207	30.94	0.9335	0.9488	0.0775	3.28	3.51
Ours (Scheme 3)	$h = 4, s = 5, v = 7$	Urban100	0.0197	30.98	0.9679	0.9684	0.0746	3.65	4.29
		BSD500	0.0204	30.48	0.9754	0.9722	0.1043	2.99	3.50
		McMaster	0.0250	27.89	0.9193	0.9436	0.1211	3.29	4.01
Ours (Scheme 4)	$h = 6, s = 5, v = 5$	Urban100	0.0260	30.50	0.9596	0.9416	0.0655	3.87	4.49
		BSD500	0.0244	30.94	0.9585	0.9392	0.0984	2.75	3.25
		McMaster	0.0247	30.49	0.9202	0.9319	0.0837	3.56	4.30

(k) Original (l) Ours (Grayscale) (m) IDN (n) CNN-based (o) Ours (Scheme 2)

Fig. 2. Visual comparison on Kodak dataset.

4) **Pseudo-Color Preview**: Map 16-bit grayscale to an 8-bit pseudo-color lookup table for quick visualization on low-bit-depth displays.

VI. CONCLUSION

We propose a reversible grayscale method that encodes HSV components into a 16-bit grayscale image using adaptive bit-field allocation and error-aware quantization. Our approach achieves high-fidelity color reconstruction, low computational overhead, and interpretability, outperforming traditional and learning-based methods in efficiency. It is particularly suited for edge computing and embedded systems. Future work will explore adaptive allocation, alternative color spaces, and hardware acceleration to further enhance performance.

A. Future Directions

- Develop adaptive allocation algorithms using lightweight statistical models.
- Extend to high-dynamic-range (HDR) images and medical imaging.
- Optimize for FPGA/ASIC hardware to achieve real-time performance.

- Integrate reversible compression (e.g., JPEG-LS) for joint grayscale and compression efficiency.

VII. Acknowledgment

This research is conducted and supported by OPT Machine Vision.

References

[1] Z. Ni, Y. Q. Shi, N. Ansari, and W. Su, "Reversible data hiding," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 16, no. 3, pp. 354–362, 2006.

[2] N. Chen and G. W. Wornell, "Wavelet-packet–based reversible data hiding," *IEEE Transactions on Multimedia*, vol. 11, no. 9, pp. 1510–1520, 2009.

[3] M. Xia, X. Liu, and T.-T. Wong, "Invertible grayscale," vol. 37, no. 6, pp. 1–10, Dec. 2018. [Online]. Available: https://doi.org/10.1145/3272127.3275080

[4] R. Zhao, T. Liu, J. Xiao, D. P. Lun, and K.-M. Lam, "Invertible image decolorization," *IEEE Transactions on Image Processing*, vol. 30, pp. 6081–6095, 2021.

[5] Y.-H. Chan, Z.-X. Xu, and D. P.-K. Lun, "A framework of reversible color-to-grayscale conversion with watermarking feature," vol. 29, pp. 859–870, 2020. [Online]. Available: https://doi.org/10.1109/TIP.2019.2936097

[6] Z.-X. Xu and Y.-H. Chan, "Improving reversible color-to-grayscale conversion with halftoning," vol. 52, pp. 111–123, Mar. 2017. [Online]. Available: https://doi.org/10.1016/j.image.2016.12.005

[7] Q. Liang, R. Hu, and S. Xiang, "Invertible color-to-grayscale conversion by using clustering and reversible watermarking," Shenzhen, China, Jul. 2021, pp. 1–6. [Online]. Available: https://doi.org/10.1109/ICME51207.2021.9428464

[8] Q. Liang and S. Xiang, "Invertible color-to-grayscale conversion using lossy compression and high-capacity data hiding," vol. 32, no. 11, pp. 7373–7385, Nov. 2022. [Online]. Available: https://doi.org/10.1109/TCSVT.2022.3184949

[9] E. Fragoso-Navarro, F. Garcia-Ugalde, and M. Cedillo-Hernandez, "Protecting the distribution of color images via inverse colorization, visible-imperceptible watermarking and reversible data hiding," vol. 11, pp. 61 025–61 048, 2023. [Online]. Available: https://doi.org/10.1109/ACCESS.2023.3286865

[10] F. S. Yu, Y.-H. Chan, K. K. Lam, and D. P. Lun, "Self-embedding reversible color-to-grayscale conversion with watermarking feature," vol. 119, p. 117061, Nov. 2023. [Online]. Available: https://doi.org/10.1016/j.image.2023.117061

[11] Y. Zhu, R. Hu, and S. Xiang, "Invertible image decolorization with CFEH and reversible data hiding," vol. 34, no. 12, pp. 12 811–12 822, Dec. 2024. [Online]. Available: https://doi.org/10.1109/TCSVT.2024.3437423

[12] D. Hou, W. Zhang, K. Chen, S.-J. Lin, and N. Yu, "Reversible data hiding in color image with grayscale invariance," vol. 29, no. 2, pp. 363–374, Feb. 2019. [Online]. Available: https://doi.org/10.1109/TCSVT.2018.2803303

[13] C.-L. Jhong and H.-L. Wu, "Exploring more capacity for grayscale-invariance reversible data hiding," vol. 9, pp. 137 005–137 014, 2021. [Online]. Available: https://doi.org/10.1109/ACCESS.2021.3117961

[14] Y. Yang, X. Sun, H. Yang, C.-T. Li, and R. Xiao, "A contrast-sensitive reversible visible image watermarking technique," vol. 19, no. 5, pp. 656–667, May 2009. [Online]. Available: https://doi.org/10.1109/TCSVT.2009.2017401

[15] Y.-Q. Shi, X. Li, X. Zhang, H.-T. Wu, and B. Ma, "Reversible data hiding: Advances in the past two decades," vol. 4, pp. 3210–3237, 2016. [Online]. Available: https://doi.org/10.1109/ACCESS.2016.2573308

[16] K. Liu, D. Chen, J. Liao, W. Zhang, H. Zhou, J. Zhang, W. Zhou, and N. Yu, "JPEG robust invertible grayscale," vol. 28, no. 12, pp. 4403–4417, Dec. 2022. [Online]. Available: https://doi.org/10.1109/TVCG.2021.3088531

[17] L. Zhang, L. Zhang, X. Mou, and D. Zhang, "Fsim: A feature similarity index for image quality assessment," *IEEE transactions on Image Processing*, vol. 20, no. 8, pp. 2378–2386, 2011.

[18] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 586–595.

[19] W. B. Pennebaker and J. L. Mitchell, *JPEG: Still Image Data Compression Standard*. Norwell, MA, USA: Kluwer Academic Publishers, 1992.

[20] S. G. Mallat, "A theory for multiresolution signal decomposition: The wavelet representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 7, pp. 674–693, 1989.

[21] A. R. Smith, "Color gamut transform pairs," in *Proceedings of the 5th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '78)*. New York, NY, USA: ACM, 1978, pp. 12–19.

[22] J.-B. Huang, A. Singh, and N. Ahuja, "Single image super-resolution from transformed self-exemplars," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.

[23] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik, "Contour detection and hierarchical image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 5, pp. 898–916, May 2011. [Online]. Available: http://dx.doi.org/10.1109/TPAMI.2010.161

[24] L. Zhang, X. Wu, A. Buades, and X. Li, "Color demosaicking by local directional interpolation and nonlocal adaptive thresholding," *Journal of Electronic imaging*, vol. 20, no. 2, pp. 023 016–023 016, 2011.