

Article

Perceptually Optimal Tone Mapping of HDR Images Through Two-Stage Bayesian Optimization

Naif Alasmari 

Department of Computer Science, Applied College, King Khalid University, Muhayil 61913, Saudi Arabia; nasmary@kku.edu.sa

Abstract

Critical details in both bright and dark regions are frequently lost in high dynamic range (HDR) images when they are displayed on low dynamic range (LDR) devices. To mitigate this issue, tone mapping operators (TMOs) have been developed to convert HDR images into LDR representations while maintaining perceptual quality. However, it is challenging to effectively balance various key visual attributes, such as naturalness and structural fidelity. To overcome this limitation, a two-stage Bayesian optimization approach was proposed in this work to enhance the perceptual quality of tone-mapped images across multiple evaluation metrics. The first stage adaptively optimizes TMQI parameters to capture image-specific perceptual characteristics, while the second stage refines the tone mapping function to further improve detail preservation and visual realism. Extensive experiments using three distinct HDR benchmark datasets were conducted, indicating that the proposed method generally performs better than the existing tone mapping techniques across most evaluated metrics, including TMQI, Naturalness, and Structural Fidelity. Our adaptive approach offers a robust and effective solution for optimizing HDR image conversion, resulting in a significantly improved perceptual quality compared to traditional methods.

Keywords: tone mapping; HDR images; tone mapping quality indicator; optimization



Academic Editor: Luca Mesin

Received: 1 September 2025

Revised: 14 October 2025

Accepted: 15 October 2025

Published: 17 October 2025

Citation: Alasmari, N. Perceptually Optimal Tone Mapping of HDR Images Through Two-Stage Bayesian Optimization. *Electronics* **2025**, *14*, 4080. <https://doi.org/10.3390/electronics14204080>

Copyright: © 2025 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The rapid development of imaging technologies has reshaped our ability to capture and reproduce visual scenes with unprecedented fidelity. Among these advancements, high dynamic range (HDR) imaging has emerged as a pivotal technology for capturing and representing the full range of luminance levels present in real-world scenes, thereby exceeding the inherent limitations of traditional low dynamic range (LDR) systems [1]. The capability to represent scenes with extreme contrast ratios not only introduces new opportunities for creative expression but also improves the precision of data interpretation across a wide range of applications. HDR imaging has been adopted in various domains, including cinematography, computer graphics, augmented and virtual reality, and scientific visualization [2,3]. Accurately rendering both bright and dark regions is crucial for ensuring both perceptual realism and analytical accuracy.

Despite the significant advancements in HDR imaging, a major practical limitation persists: most consumer displays, web platforms, and image processing workflows are still restricted to low dynamic range (LDR) output. As a result, the HDR content must be converted to LDR-compatible formats before it can be displayed or processed. This conversion

process, known as tone mapping, involves compressing the dynamic range while maintaining key perceptual attributes, such as contrast, texture, and color fidelity [4]. However, achieving this balance is a non-trivial task. The conventionally applied naive compression methods often lead to visual artifacts, loss of critical scene details, or degradation of overall image quality.

The tone mapping quality index (TMQI) is a widely adopted metric for evaluating tone-mapped images, combining structural fidelity (S) that preserves fine details, and naturalness (N), which reflects perceived realism [5].

$$\text{TMQI} = a \cdot S^\alpha + (1 - a) \cdot N^\beta,$$

where

- S denotes structural fidelity;
- N represents statistical naturalness;
- a balances structural fidelity, naturalness, and $\in [0, 1]$;
- β adjusts the emphasis on naturalness;
- α controls the sensitivity to structural similarity.

However, this formulation employs fixed parameters that may not be generalized across diverse image types or application contexts, highlighting the need for adaptive and content-aware evaluation methods.

Existing TMOs face several challenges. Many rely on globally fixed parameters that cannot adapt to the diverse luminance distributions, textures, and semantic structures found in real-world scenes. Consequently, they often produce inconsistent perceptual outcomes, preserving structure in some cases but sacrificing naturalness or color fidelity in others. Moreover, while TMQI provides a valuable quantitative framework, its static weighting between structural fidelity and naturalness fails to account for scene-dependent perceptual variations. These limitations motivate the development of an adaptive, perceptually guided optimization approach.

To address the aforementioned issue, a novel global tone mapping approach that integrates perceptual modeling within an adaptive two-stage optimization process was introduced in this work. In the first stage, the TMQI parameters (a , α , and β) are dynamically adjusted to align with the specific perceptual characteristics of the input HDR image. This iterative refinement ensures that the metric accurately reflects the human visual perception within the context of the given scene. In the second stage, the above-mentioned parameters are maintained. In addition, the tone mapping function is optimized using the established TMQI configuration, with an emphasis on maximizing structural detail and enhancing overall perceptual realism. Specifically, two key values were adjusted: the histogram mixing weight (w_1) and the luminance compression exponent (s) to improve the structural fidelity that conveys important details and enhances naturalness in the final image. This adaptive strategy facilitates the content-aware HDR-to-LDR conversion and demonstrates robustness across a wide range of image types and scenes, mitigating limitations inherent in traditional tone mapping approaches.

To assess the effectiveness of the proposed approach, comprehensive experiments were carried out on two widely used HDR benchmark datasets, with results compared against several established tone mapping operators. The evaluation incorporated multiple perceptual quality measures, further supported by statistical analyses to confirm both the consistency and the significance of the observed improvements.

The main contributions of this paper can be summarized as follows:

- We introduce an adaptive two-stage Bayesian optimization approach for tone mapping, which dynamically adjusts both TMQI parameters and tone-curve functions according to image content.
- The proposed approach integrates perceptual modeling to jointly optimize structural fidelity and naturalness, addressing the limitations of fixed-parameter TMOs.
- A comprehensive experimental evaluation is conducted using multiple HDR datasets and perceptual quality metrics, demonstrating statistically significant improvements over conventional methods.

The remainder of this article is structured as follows: Section 2 reviews related work, while Section 3 details the proposed methodology. The implementation aspects are outlined in Section 4. Section 5 provides an in-depth analysis of the experimental results, and Section 6 concludes the study with key findings and potential directions for future research.

2. Related Work

Tone mapping is a technique that is used to render HDR content on devices with limited LDR capabilities. Its primary objective is to reduce the wide luminance range in HDR images while preserving visual fidelity, structural integrity, and perceptual realism on LDR displays. Several approaches utilize global tone mapping operators for their simplicity and deterministic behavior. For instance, Reinhard et al. [6] introduced a tone mapping method adapted from the zone system developed by Ansel Adams. This approach consists of a global operator that first applies a linear scaling step to the image luminance, and, if required, a local operator that automatically performs dodging-and-burning to compress the dynamic range while preserving local detail. This approach offered intuitive controls, proved robust across scenes, and laid a perceptual foundation for many subsequent tone mapping operators. Drago et al. [7] proposed an adaptive tone mapping algorithm based on the principles of human visual perception. Logarithmic functions with different bases were employed depending on the luminance range: base two is used to enhance contrast and visibility in dark regions, while base ten is applied to the brightest regions to achieve stronger contrast compression. For intermediate luminance values, a bias function is used to enable artifact-free image reproduction. This approach demonstrated that combining perceptual anchoring with logarithmic compression can lead to visually satisfying results in a computationally efficient manner. Khan et al. [8] developed an adaptive tone technique (ATT). In particular, a global histogram-based method based on the threshold-versus-intensity (TVI) curve was introduced, and just-noticeable difference (JND) modeling was performed. By incorporating human visual system (HVS) sensitivity into histogram construction, the algorithm allocated display levels to perceptually significant regions while avoiding excessive contrast enhancement in dominant bins. This perceptually driven binning approach reduced artifacts and preserved details across both bright and dark areas. A tone mapping method was subsequently introduced that utilized a perceptual quantizer to transform the real-world luminance values of HDR images into a perceptual domain. The implementation of this transformation, which closely emulates the response characteristics of the human visual system, yielded an enhancement in the contrast of low-luminance regions while the intensities in excessively bright areas were compressed. Subsequently, a luminance histogram is generated, from which a tone mapping curve can be derived to improve visual quality by revealing detail in darker regions and mitigating information loss in high-intensity zones [9]. In parallel, Quraishi et al. [10] proposed a learning-based Application-Specific Tone Mapping Operator (ASTMO) aimed at improving the OCR accuracy under diverse lighting conditions. This method is built on the ATT by iteratively adjusting histogram bin locations using the OCR accuracy as the optimization objective, thereby enhancing the text contrast in tone-mapped HDR images.

While the global TMOs attain speed and robustness, they often lack the necessary local adaptivity to maintain the fine spatial details. Cao et al. [11] addressed this issue by developing a perceptually optimized and self-calibrated tone mapping operator (PS-TMO) approach based on a two-stage deep learning framework for generating high-quality LDR images. In the first stage, HDR images are decomposed into a normalized Laplacian pyramid, and a convolutional neural network is used to estimate the corresponding LDR pyramid. In the second stage, self-calibration on the HDR image is performed to adaptively produce the final tone-mapped result. Mantiuk et al. [12] emphasized the importance of display-adaptive tone mapping and proposed an optimization framework that minimizes contrast distortions relative to display characteristics and ambient conditions. This method combined a human visual system model with quadratic programming to dynamically adapt tone curves across viewing scenarios, laying the groundwork for perceptual optimization in display-aware pipelines.

Hybrid methods have also been developed to merge the strengths of global and local operators. For instance, Zhang et al. [13] introduced a tone mapping that uses both global and local operators, guided by the mean intensity of the HDR image. The local operator emphasizes fine detail preservation, while the global operator ensures the overall visual consistency. Then, the Curvelet transform fused the results to yield LDR images with enhanced detail and natural visual quality. Furthermore, ref. [14] devised an adaptive tone mapping algorithm implemented in an FPGA, where global and local operators are combined through a histogram-based function. This separation of global illumination and local detail is conducive to suppressing halos and reducing noise while enabling efficient real-time processing. Cheng et al. [15] proposed the TWT algorithm, an enhancement of the ATT method that integrates the weighted least squares (WLS) decomposition with human visual system-based histogram adjustment. In this approach, the HDR image is decomposed into base and detail layers. The former layer undergoes dynamic range compression, while the latter is selectively enhanced, followed by a color balance correction step to preserve natural colors. Both global brightness and local contrast are maintained during these operations.

Deep learning has also emerged as a powerful tool in tone mapping. Ref. [16] introduced a generative adversarial network-based tone mapping operator to eliminate manual parameter tuning and generate high-quality LDR images, although at the cost of higher computational complexity. In [17], a convolutional neural network was proposed for HDR image fusion and tone mapping in the context of autonomous driving, which demonstrated robustness against motion artifacts but was mainly tailored to traffic scenes. In addition, ref. [18] presented a lightweight method based on convolutional neural networks that converts HDR images into the standard dynamic range with improved efficiency. This approach is suitable for real-time applications, though with limited generalization outside its target scenarios.

In addition to these developments, several other learning-based approaches have been developed to further enhance adaptability and perceptual quality. Rana et al. [19] proposed DeepTMO, a parameter-free operator based on conditional generative adversarial networks that produces high-resolution outputs and mitigates common artifacts such as saturation and blurring through a multi-scale generator–discriminator design. Guo and Jiang [20] extended this direction by introducing a semi-supervised CNN framework in which image quality assessment metrics were embedded into the training objective, enabling the operator to generate visually consistent results across diverse HDR scenes. In parallel, Rana et al. [21] explored the role of tone mapping in computer vision tasks, presenting a learning-based operator optimized for image matching under challenging illumination conditions through a guidance model trained with support vector regression. More re-

cently, Todorov et al. [22] proposed TGTM, a TinyML-based global operator designed for HDR sensors in automotive contexts, achieving competitive visual performance while significantly reducing computational cost by leveraging histogram-driven neural networks. Although these methods highlight the versatility of data-driven solutions, they fall outside the scope of our work, which focuses on traditional tone mapping operators.

Despite these advances, the existing approaches still suffer from detail loss, high complexity, or limited interpretability. To overcome these issues, a two-stage optimization approach that adapts both the evaluation metric and the tone mapping function was proposed, yielding an interpretable, data-adaptive, and robust solution across diverse HDR scenes.

3. Methodology

The proposed approach presents a perceptually guided two-stage Bayesian optimization method for tone mapping, aimed at jointly refining both the perceptual evaluation metric and the tone mapping parameters. The complete workflow is shown in Figure 1. In the first stage, the algorithm adaptively adjusts the parameters of the TMQI to align with image-specific perceptual characteristics. In the second stage, the tone-curve function is further optimized to improve visual fidelity while maintaining the achieved TMQI performance. This integration of perceptual modeling and optimization forms the central contribution of this work, delivering a content-adaptive and reliable HDR to LDR conversion process.




Figure 1. High-level diagram of the proposed two-stage optimization strategy for perceptual tone mapping. The ATT operator generates the initial LDR image, while TMQI provides perceptual feedback that directly guides parameter optimization across both stages of the Bayesian process.

The detailed implementation of this two-stage optimization framework is illustrated in Figure 1. As shown, the process begins with a global tone mapping step based on the adaptive tone mapping technique (ATT) [8], which serves as the baseline global operator in the proposed approach. ATT converts the input HDR image into an initial LDR image by modeling the HVS adaptation to luminance. It estimates global scene statistics, such as mean and maximum luminance, to determine the required degree of compression, and incorporates local adaptation through JND modeling derived from the TVI function. A perceptually guided histogram is then constructed in which each bin corresponds to a JND

in luminance, ensuring finer quantization for perceptually significant regions. This histogram is transformed into a nonlinear tone curve that compresses the dynamic range while preserving local detail and contrast. Two adaptive parameters, w_1 (global–local mixing weight) and s (tone–curve exponent), jointly control the balance between global luminance compression and local contrast preservation. Although the original ATT formulation includes a temporal adaptation mechanism for HDR video sequences, this component is intentionally excluded in this work, as the proposed approach is designed for static HDR images. Applying the resulting tone curve to the HDR image produces the initial LDR representation, which is then used as input for the subsequent optimization stages.

In this approach, the ATT algorithm acts as the baseline tone mapping operator that generates the initial LDR image, while the TMQI metric serves as the perceptual guide for optimization. Rather than applying TMQI only for post-evaluation, it is integrated within the optimization loop to iteratively refine the tone mapping parameters. In Stage 1, TMQI's weighting parameters (a , α , and β) are optimized to align with image-specific perceptual responses. In Stage 2, the ATT parameters (w_1 , s) are re-optimized using the calibrated TMQI configuration, ensuring that perceptual feedback directly drives the refinement of tone mapping behavior. This integration enables a closed-loop optimization process that jointly balances perceptual fidelity and naturalness.

As depicted in Figure 1, the first stage of the approach focuses on calibrating the perceptual quality metric. TMQI is employed as the primary objective function, combining structural fidelity and naturalness components. Although TMQI requires both HDR and LDR inputs for evaluation, in this approach, the LDR image is generated internally during the optimization process. Therefore, only the HDR image is provided as input, while LDR candidates are iteratively produced and evaluated within each optimization stage. In this stage, the TMQI parameters (a , α , and β) are optimized on a per-image basis using Bayesian optimization [23]. These parameters govern the relative contributions of structural fidelity and naturalness in TMQI computation. The optimization proceeds until a predefined TMQI threshold is met, ensuring a minimum perceptual quality standard while reducing unnecessary computation. A fixed TMQI threshold is used to establish a consistent convergence baseline across diverse HDR images. This fixed target does not assume identical perceptual characteristics for all content but ensures that Stage 1 reaches a stable level of perceptual quality before Stage 2 adaptively refines tone mapping parameters according to image-specific characteristics.

If the optimal TMQI score in Stage 1 satisfies the quality criterion, the process continues to Stage 2. This stage aims to further refine the tone mapping curve by adjusting two core parameters: the histogram mixing weight w_1 and the luminance compression exponent s . These parameters influence how global luminance and local contrast are balanced in the final tone-mapped image. The second stage is accepted only if it is maintained or improved upon the TMQI score from Stage 1 while also enhancing the combined structural fidelity and naturalness scores. The approach includes an optional early stopping parameter, *stage2_earlyStop*, which halts Stage 2 once convergence is reached. It was disabled (*False*) in all reported experiments to enable full optimization within the 150-iteration limit, though it remains useful in extended or less constrained runs.

Upon completion of both stages, the approach generates the final LDR image, the corresponding TMQI score, S , N , and the optimized parameter values. For reproducibility, all optimized parameters are logged, and the best-performing configurations are retained for future initialization.

By integrating perceptual adaptation with parameterized refinement of tone mapping, the proposed approach ensures robust and content-aware rendering. The conditional structure of the two-stage optimization allows the system to tailor the tone mapping strategy

to the unique characteristics of each HDR image, promoting consistency in perceptual quality across a wide range of scenes. Bayesian optimization, implemented with Optuna's tree-structured Parzen estimator (TPE) sampler, is employed to efficiently explore the parameter space. In Stage 1, the TMQI weighting parameters (a, α, β) are optimized for each image. In Stage 2, the TMO parameters (w_1, s) are tuned while preserving the TMQI value determined in Stage 1. The term optimize in Algorithm 1 refers to the Bayesian optimization process implemented using Optuna's tree-structured Parzen estimator (TPE). In Stage 1, the objective function maximizes TMQI by searching for the best combination of (a, α, β) , while in Stage 2, the goal is to maximize the combined perceptual score $(S + N)$ while maintaining or improving the TMQI score obtained in Stage 1. Early stopping is applied when either the TMQI target or the $S + N$ threshold is reached.

The two-stage optimization approach is summarized in Algorithm 1.

Algorithm 1 Two-stage Bayesian optimization procedure using the tree-structured Parzen estimator (TPE) to iteratively refine TMQI and tone mapping parameters

Require: HDR image I_{HDR} , initial parameters $\theta_0 = \{a, \alpha, \beta, w_1, s\}$

Ensure: Optimized LDR image I_{LDR} , best parameters $\theta^* = \{a^*, \alpha^*, \beta^*, w_1^*, s^*\}$

- 1: **Compute luminance** from I_{HDR} and construct perceptual histogram using TVI model
 - 2: **Generate tone curve** using JND binning and initial parameters w_1, s
 - 3: Apply tone mapping to obtain initial I_{LDR}
 - 4: **Stage 1: TMQI Parameter Optimization**
 - 5: Fix tone curve (w_1, s)
 - 6: Optimize $\text{TMQI}(I_{HDR}, I_{LDR})$ by tuning (a, α, β)
 - 7: **if** $\text{TMQI} \geq \text{threshold}$ **then**
 - 8: Store optimal (a^*, α^*, β^*)
 - 9: **else**
 - 10: Use best result so far
 - 11: **end if**
 - 12: **Stage 2: Tone Mapping Refinement**
 - 13: Fix (a^*, α^*, β^*)
 - 14: Optimize $(S + N)$ by tuning (w_1, s) subject to:
 - 15: $\text{TMQI}(I_{HDR}, I_{LDR}) \geq \text{TMQI from Stage 1}$
 - 16: Store (w_1^*, s^*) if improved
 - 17: Generate final tone curve using (w_1^*, s^*)
 - 18: Apply tone mapping to obtain final I_{LDR}
 - 19: **return** $I_{LDR}, \theta^* = \{a^*, \alpha^*, \beta^*, w_1^*, s^*\}$
-

4. Implementation Details

Experiments were conducted on 238 HDR images from three benchmark datasets: the Funt dataset [24,25], a publicly available dataset by Greg Ward [26], and the UltraFusion benchmark proposed by Chen et al. [27]. All computations were performed on a desktop system equipped with an Intel i7 processor (2.1 GHz) and 64 GB of RAM.

5. Results and Discussion

To evaluate the performance of the proposed two-stage tone mapping approach, experiments on three benchmark HDR datasets were conducted: the Funt dataset (105 images), the HDR collection from Greg Ward's site (33 images), and the UltraFusion benchmark (100 images). Four widely used tone mapping methods were implemented in Python as baselines for comparison: ATT [8], a simplified Mantiuk operator [12], the global photographic tone mapping operator of Reinhard [6], and a simplified Drago operator [7].

The evaluation employed both reference-based and no-reference image quality metrics to ensure a comprehensive assessment of performance. Specifically, the analysis included

TMQI, S, and N, together with two additional perceptual quality metrics, BRISQUE and NIQE, which evaluate image naturalness and perceptual quality without requiring reference HDR images. To enhance analytical robustness, each metric was reported using the mean, median, standard deviation, and 95% confidence intervals (CI). Statistical significance was assessed using the Friedman and Wilcoxon tests.

5.1. Performance on the Funt Dataset

This section presents the evaluation results on the Funt dataset, demonstrating the effectiveness of the proposed approach compared with existing methods. As shown in Table 1, the proposed two-stage method consistently achieves the highest performance across all quantitative and perceptual metrics.

For TMQI, the proposed method achieved a mean score of 0.9101 ± 0.0107 and a median of 0.9274 ± 0.0107 , exceeding ATT (mean 0.8688 ± 0.0137), Mantiuk (0.8058 ± 0.0169), Reinhard (0.7369 ± 0.0103), and Drago (0.6035 ± 0.0245). The Friedman test, applied to ranked metric values, confirmed statistically significant differences among all methods ($\chi^2 = 361.92$, $p = 4.68 \times 10^{-77}$), and the post-hoc Wilcoxon analysis indicated that the two-stage method performed significantly better than each baseline ($p < 0.05$). The narrow confidence intervals demonstrate high consistency and stability across the dataset.

For structural fidelity, the two-stage method achieved a mean of 0.8638 ± 0.0140 and a median of 0.8665 ± 0.0140 , marginally higher than ATT (mean 0.8618 ± 0.0146 , median 0.8670 ± 0.0146). Although the difference between these two methods was modest, both exhibited higher fidelity compared to Mantiuk (0.7993 ± 0.0267), Reinhard (0.7417 ± 0.0284), and Drago (0.4289 ± 0.0439). The Friedman test ($\chi^2 = 259.57$, $p = 5.63 \times 10^{-55}$) indicated statistically significant differences across all methods, supporting the consistency of the structural fidelity improvements.

For the naturalness metric, the proposed method achieved a mean of 0.5098 ± 0.0605 and a median of 0.5685 ± 0.0605 , exceeding ATT (mean 0.4343 ± 0.0572 , median 0.4504 ± 0.0572), Mantiuk (0.2301 ± 0.0506), Reinhard (0.0176 ± 0.0072), and Drago (0.0216 ± 0.0125). The Friedman test ($\chi^2 = 356.08$, $p = 8.52 \times 10^{-76}$) confirmed statistically significant differences, demonstrating the perceptual robustness of the proposed approach in generating natural-looking tone-mapped images.

The no-reference metrics BRISQUE and NIQE further supported the perceptual consistency of the proposed approach. It achieved the lowest BRISQUE (mean 7.36 ± 0.49 , median 7.58 ± 0.49) and NIQE (mean 4.01 ± 0.32 , median 3.43 ± 0.32) scores among all compared methods, indicating minimal perceptual distortion and high image naturalness. In contrast, ATT produced BRISQUE and NIQE means of 25.30 ± 2.42 and 4.09 ± 0.33 , respectively, while Drago exhibited the weakest perceptual quality (BRISQUE 55.92 ± 5.67 , NIQE 6.78 ± 0.72). Both BRISQUE ($\chi^2 = 255.71$, $p = 3.83 \times 10^{-54}$) and NIQE ($\chi^2 = 130.10$, $p = 3.72 \times 10^{-27}$) differences were statistically significant.

Overall, the proposed two-stage tone mapping method ranked first across all quantitative and perceptual measures, with consistently high mean and median values, narrow confidence intervals, and low variability. These results confirm the method's robustness, stability, and high perceptual fidelity in tone mapping HDR images from the Funt dataset.

Table 1. Statistical performance results for the Funt dataset (105 images) across five tone mapping methods. Values include mean, median, standard deviation (Std), and 95% confidence intervals (CI). Lower values indicate better performance for BRISQUE and NIQE. Bold values indicate the best performance among the methods.

Metric	Measure	Two-Stage	ATT	Drago	Mantiuk	Reinhard
TMQI	Mean	0.9101	0.8688	0.6035	0.8058	0.7369
	Median	0.9274	0.8841	0.6184	0.7957	0.7464
	Std	0.0552	0.0710	0.1267	0.0874	0.0530
	95% CI	± 0.0107	± 0.0137	± 0.0245	± 0.0169	± 0.0103
Structural Fidelity	Mean	0.8638	0.8618	0.4289	0.7993	0.7417
	Median	0.8665	0.8670	0.4272	0.8325	0.7546
	Std	0.0722	0.0755	0.2268	0.1380	0.1468
	95% CI	± 0.0140	± 0.0146	± 0.0439	± 0.0267	± 0.0284
Naturalness	Mean	0.5098	0.4343	0.0216	0.2301	0.0176
	Median	0.5685	0.4504	0.0000	0.0988	0.0056
	Std	0.3126	0.2956	0.0647	0.2613	0.0372
	95% CI	± 0.0605	± 0.0572	± 0.0125	± 0.0506	± 0.0072
BRISQUE	Mean	7.3551	25.3025	55.9216	27.5389	36.0330
	Median	7.5760	23.8820	57.2629	27.1398	33.0637
	Std	2.5481	12.5085	29.3048	11.4936	16.6543
	95% CI	± 0.4931	± 2.4207	± 5.6712	± 2.2243	± 3.2230
NIQE	Mean	4.0070	4.0877	6.7807	4.6825	4.3791
	Median	3.4273	3.4679	5.6102	3.7025	4.1612
	Std	1.6690	1.7018	3.7423	2.3005	1.2703
	95% CI	± 0.3230	± 0.3293	± 0.7242	± 0.4452	± 0.2458

To complement the quantitative evaluations, visual comparisons were provided to highlight perceptual differences among the five evaluated tone mapping methods. Figures 2–4 present representative images from the first dataset, illustrating each method’s ability to preserve detail, maintain contrast, reproduce color fidelity, and enhance overall perceptual quality.

Figure 2 focuses on the TMQI results for a texture-rich terrain, linking perceptual observations with the quantitative scores shown below each image. The proposed two-stage method attains the highest TMQI value and preserves surface texture with balanced brightness and contrast. The ATT output is visually close but slightly flatter in shadow regions. The Mantiuk result appears over-sharpened, while the Drago output is overexposed and loses midtone detail. In contrast, the Reinhard method produces a darker rendering that obscures fine features.

Figure 3 illustrates the structural-fidelity-based assessment of an open landscape. The proposed method achieves a natural balance between sky luminance and foreground texture. The ATT method preserves this balance reasonably well, whereas the Mantiuk method significantly increases brightness, washing out the midtones. In contrast, the Drago method overexposes large regions, while the Reinhard method yields a dull and overly dark appearance.

Figure 4 evaluates perceptual naturalness using a forested scene with strong luminance gradients. The two-stage method successfully retains highlight and shadow details while preserving texture. The ATT method achieves comparable luminance control but with slightly muted contrast. The Drago method produces a washed-out result due to overexposure, whereas the Reinhard method renders the scene overly dark. The Mantiuk method increases local contrast but introduces unnatural color shifts, particularly in the vegetation.



Figure 2. TMQI-based comparison on the Funt dataset across five tone mapping methods.



Figure 3. Structural fidelity-based comparison on the Funt dataset across five tone mapping methods.




Figure 4. Naturalness-based comparison on the Funt dataset across five tone mapping methods.

5.2. Performance on the Greg Ward HDR Dataset

To further evaluate the robustness and generalizability of the proposed two-stage tone mapping approach, additional experiments were conducted on the Greg Ward HDR dataset, which includes 33 images covering a wide range of luminance levels and scene content. The results, summarized in Table 2, are consistent with those obtained from the Funt dataset and reinforce the effectiveness of the proposed optimization framework.

The two-stage approach achieved the highest overall performance across all five objective metrics. For TMQI, it achieved a mean score of 0.9331 and a median of 0.9457, exceeding ATT (mean 0.9079, median 0.9350), Drago (0.8239, 0.8312), Mantiuk (0.7851, 0.7858), and Reinhard (0.7597, 0.7669). The relatively narrow 95% confidence interval (± 0.0167) and low standard deviation (0.047) indicate that the method delivers consistent high-quality results. The Friedman test confirmed statistically significant differences among all evaluated methods ($\chi^2 = 109.31$, $p = 1.022 \times 10^{-22}$), and the post-hoc Wilcoxon analysis verified that the two-stage approach performed significantly better than ATT ($p < 0.05$).

In terms of structural fidelity, the two-stage approach achieved a mean of 0.8505 and a median of 0.8670, slightly higher than ATT (0.8499 and 0.8677). Although the numerical difference is small, the two-stage method exhibited greater stability, as reflected by its lower variability (Std = 0.0748) and narrower confidence interval (± 0.0265). The Friedman test confirmed significant differences among the methods ($\chi^2 = 33.33$, $p = 1.021 \times 10^{-6}$), supporting the consistency of the structural fidelity improvements.

For the naturalness metric, the proposed approach achieved a mean of 0.7142 and a median of 0.8589, exceeding ATT (0.6717, 0.8302), Drago (0.3705, 0.3806), Mantiuk (0.1799, 0.0678), and Reinhard (0.0542, 0.0382). These results indicate improved perceptual realism and tonal balance while effectively reducing unnatural contrast artifacts. The Friedman test confirmed that these differences were statistically significant ($\chi^2 = 100.29$, $p = 8.529 \times 10^{-21}$).

The no-reference perceptual quality metrics BRISQUE and NIQE further supported the perceptual consistency of the proposed approach. It achieved the lowest BRISQUE (mean 20.45, median 18.38) and NIQE (mean 3.71, median 3.73) scores among all methods, indicating reduced visual distortion and enhanced naturalness. The Friedman tests for BRISQUE ($\chi^2 = 36.80$, $p = 1.981 \times 10^{-7}$) and NIQE ($\chi^2 = 51.68$, $p = 1.605 \times 10^{-10}$) confirmed that these differences were statistically significant.

Overall, the results on the Greg Ward dataset demonstrate that the proposed two-stage tone mapping approach consistently achieves high objective and perceptual performance. The combination of elevated TMQI, stable Structural Fidelity, and favorable no-reference scores (BRISQUE and NIQE) highlights its robustness, stability, and adaptability across diverse HDR imaging conditions.

Table 2. Statistical performance results for the Greg Ward HDR dataset (33 images) across five tone mapping methods. Values include mean, median, standard deviation (Std), and 95% confidence intervals (CI). Lower values indicate better performance for BRISQUE and NIQE. Bold values indicate the best performance among the methods.

Metric	Measure	Two-Stage	ATT	Drago	Mantiuk	Reinhard
TMQI	Mean	0.9331	0.9079	0.8239	0.7851	0.7597
	Median	0.9457	0.9350	0.8312	0.7858	0.7669
	Std	0.0470	0.0634	0.1003	0.0689	0.0446
	95% CI	± 0.0167	± 0.0225	± 0.0356	± 0.0244	± 0.0158
Structural Fidelity	Mean	0.8505	0.8499	0.7628	0.7599	0.7685
	Median	0.8670	0.8677	0.8007	0.8041	0.7929
	Std	0.0748	0.0771	0.1462	0.1283	0.1197
	95% CI	± 0.0265	± 0.0273	± 0.0519	± 0.0455	± 0.0425
Naturalness	Mean	0.7142	0.6717	0.3705	0.1799	0.0542
	Median	0.8589	0.8302	0.3806	0.0678	0.0382
	Std	0.2934	0.2972	0.2997	0.2055	0.0668
	95% CI	± 0.1040	± 0.1054	± 0.1063	± 0.0729	± 0.0237
BRISQUE	Mean	20.4456	20.8070	27.3883	24.4149	28.2521
	Median	18.3825	18.6369	24.0256	21.0905	23.8581
	Std	14.9931	14.8915	17.9727	15.7582	18.1919
	95% CI	± 5.3163	± 5.2803	± 6.3728	± 5.5876	± 6.4506
NIQE	Mean	3.7059	3.7166	4.3470	4.3274	4.4188
	Median	3.7251	3.5157	4.5504	4.2880	4.2440
	Std	1.0738	1.0822	1.3921	1.1659	1.4329
	95% CI	± 0.3807	± 0.3837	± 0.4936	± 0.4134	± 0.5081

The visual comparisons of the second dataset are shown in Figures 5–7, providing perceptual evidence alongside the quantitative scores displayed below each image. These examples allow an independent assessment of how each tone mapping method preserves structural detail, contrast, and overall visual realism.

Figure 5 illustrates a high-contrast indoor lighting scene. The proposed two-stage approach achieves the highest TMQI value, preserving fine reflections on metallic surfaces while maintaining color fidelity and balanced illumination. The ATT result is visually comparable but slightly less consistent across bright highlights. Mantiuk introduces local over-enhancement, whereas Drago and Reinhard either overexpose or darken extensive regions, thereby reducing perceptual realism.

Figure 6 presents the structural fidelity assessment for a complex natural scene. The two-stage method produces the most coherent textures across both shadows and illuminated areas, corresponding to the highest S value. ATT performs similarly but loses minor detail near the bright trunk. Mantiuk and Reinhard generate darker mid-tones, while Drago exhibits severe highlight clipping.

Finally, Figure 7 compares perceptual naturalness for an indoor scene. The proposed method produces a well-balanced rendering with realistic brightness and contrast, reflected in the highest naturalness score. ATT maintains perceptual balance with slight contrast compression, whereas Mantiuk and Reinhard yield underexposed results, and Drago produces excessive brightness that reduces depth perception.



Figure 5. TMQI-based comparison on HDR images from Greg Ward's site across five tone map ping methods.



Figure 6. TMQI-based comparison on HDR images from Greg Ward's site across five tone map ping methods.




Figure 7. N-based comparison on HDR images from Greg Ward's site across five tone map ping methods.

5.3. Performance on the UltraFusion Benchmark Dataset

To further assess scalability and cross-dataset generalization, the proposed two-stage approach was evaluated on the UltraFusion benchmark, a large and diverse HDR dataset encompassing substantial variations in lighting conditions, dynamic range, and scene composition. The statistical results, summarized in Table 3, provide a comprehensive comparison across the five tone mapping algorithms.

The proposed approach achieved the highest overall performance in TMQI, with a mean of 0.9687 and a median of 0.9760, substantially surpassing ATT (0.9319, 0.9435), Drago (0.8217, 0.8487), Mantiuk (0.6811, 0.6955), and Reinhard (0.7698, 0.7792). The narrow confidence interval (± 0.0046) and low standard deviation (0.0231) indicate exceptional stability across the dataset. The Friedman test confirmed statistically significant differences among the compared methods ($\chi^2 = 357.59$, $p = 4.02 \times 10^{-76}$), and post-hoc analysis identified the proposed approach as the top performer ($p < 0.05$).

For structural fidelity, the proposed method maintained strong performance, achieving a mean of 0.9356 and a median of 0.9387, outperforming ATT (0.9023, 0.9112) and all remaining baselines. These results highlight its ability to preserve fine structural details while avoiding over-enhancement artifacts. The Friedman test further confirmed statistically significant differences among the five algorithms ($\chi^2 = 300.71$, $p = 7.61 \times 10^{-64}$).

The naturalness metric further validates the perceptual advantage of the proposed approach. It achieved a mean of 0.8404 and a median of 0.9090, exceeding ATT (0.7244, 0.8038) and substantially outperforming Drago (0.3409, 0.2976), Mantiuk (0.0004, 0.0001), and Reinhard (0.0626, 0.0263). The extremely low variability (Std = 0.1915) and narrow confidence interval (± 0.0380) emphasize both perceptual realism and consistency. These differences were statistically significant ($\chi^2 = 340.38$, $p = 2.09 \times 10^{-72}$).

In the no-reference evaluation, BRISQUE and NIQE results further corroborate the perceptual strength of the proposed method. It achieved a low BRISQUE mean of 19.17 and median of 18.38, second only to Reinhard's unusually low values (mean 11.07, median 8.78), which correspond to perceptually underexposed outputs. The Friedman test ($\chi^2 = 66.30$, $p = 1.37 \times 10^{-13}$) confirmed significant variation across methods. For NIQE, the proposed approach achieved the best results (mean 2.86, median 2.88), indicating high perceptual naturalness and minimal structural distortion. The NIQE differences were also statistically significant ($\chi^2 = 208.26$, $p = 6.30 \times 10^{-44}$).

Overall, the UltraFusion benchmark results further substantiate the effectiveness of the proposed two-stage tone mapping approach. The consistent top-ranked performance across all full-reference and perceptual metrics demonstrates strong scalability, robustness, and adaptability to diverse HDR imaging conditions. These findings confirm that the proposed approach generalizes effectively across datasets, maintaining high perceptual fidelity even under complex luminance variations.

Table 3. Statistical performance results for the UltraFusion benchmark dataset across five tone mapping methods. Values include mean, median, standard deviation (Std), and 95% confidence intervals (CI). Lower values indicate better performance for BRISQUE and NIQE. Bold values indicate the best performance among the methods.

Metric	Measure	Two-Stage	ATT	Drago	Mantiuk	Reinhard
TMQI	Mean	0.9687	0.9319	0.8217	0.6811	0.7698
	Median	0.9760	0.9435	0.8487	0.6955	0.7792
	Std	0.0231	0.0431	0.1068	0.0645	0.0497
	95% CI	± 0.0046	± 0.0086	± 0.0212	± 0.0128	± 0.0099
Structural Fidelity	Mean	0.9356	0.9023	0.7820	0.6034	0.7991
	Median	0.9387	0.9112	0.8357	0.6279	0.8129
	Std	0.0307	0.0482	0.1695	0.1647	0.1192
	95% CI	± 0.0061	± 0.0096	± 0.0336	± 0.0327	± 0.0236
Naturalness	Mean	0.8404	0.7244	0.3409	0.0004	0.0626
	Median	0.9090	0.8038	0.2976	0.0001	0.0263
	Std	0.1915	0.2490	0.2884	0.0015	0.0960
	95% CI	± 0.0380	± 0.0494	± 0.0572	± 0.0003	± 0.0190
BRISQUE	Mean	19.1667	17.2910	19.9697	20.7830	11.0663
	Median	18.3843	16.8283	14.8681	19.9074	8.7787
	Std	9.6638	15.4986	23.7815	16.5519	12.6130
	95% CI	± 1.9175	± 3.0753	± 4.7188	± 3.2843	± 2.5027
NIQE	Mean	2.8616	3.3168	3.7641	4.3276	3.4411
	Median	2.8819	3.2462	3.2486	4.1981	3.3970
	Std	0.5206	0.7003	1.6407	1.0729	0.6949
	95% CI	± 0.1033	± 0.1390	± 0.3256	± 0.2129	± 0.1379

5.4. Ablation Study

To evaluate the specific contribution of each stage in the proposed two-stage optimization approach, an ablation study was conducted using the three HDR datasets. Five objective quality metrics were employed to quantify the proportion of images that exhibited improvement after applying the complete two-stage process. For each image, an improvement was recorded when TMQI, S, or N increased in Stage 2 compared to Stage 1,

or when BRISQUE and NIQE decreased. The resulting percentages, summarized in Table 4, represent the proportion of improved images for each dataset and metric.

Table 4. Percentage of images showing improvement after applying the full two-stage process across the three datasets.

Dataset	TMQI	S	N	BRISQUE	NIQE
Funt dataset	67%	44%	61%	35%	40%
Greg Ward HDR dataset	97%	55%	91%	67%	61%
UltraFusion benchmark	98%	56%	95%	45%	52%

As shown in Table 4, the second optimization stage provides consistent improvements across all datasets. Significant gains were observed in TMQI and Naturalness, reflecting enhanced perceptual quality and improved tone adaptation. Moderate improvements in Structural Fidelity indicate that the refinement stage helps preserve fine structural details while maintaining overall visual balance. Moreover, a higher proportion of images with reduced BRISQUE and NIQE values suggests lower perceptual distortion and increased visual realism. Collectively, these findings confirm that incorporating both optimization stages results in a more stable, perceptually consistent, and robust tone-mapping performance compared to using only the initial stage.

5.5. Overall Analysis

Across all three datasets, the proposed approach consistently ranked first or second across all evaluation metrics. The most notable gains were observed in TMQI and naturalness, aligning with the approach's design emphasis on perceptual enhancement and adaptive tone mapping.

While the ATT method achieved slightly higher values in Structural Fidelity, these differences were marginal and not consistently statistically significant. By contrast, the proposed method delivered balanced improvements across both perceptual and structural criteria, indicating stronger generalization and robustness.

These findings support the effectiveness of the proposed approach, particularly in preserving image realism while maintaining detail and contrast. The method's adaptive optimization strategy, based on the perceptual feedback and early stopping, enables efficient computation without compromising quality. The performance remained consistent across varied scenes and lighting conditions, reinforcing the method's applicability to a wide range of HDR imaging scenarios.

Furthermore, these visual comparisons uphold the quantitative analysis. The proposed two-stage method consistently outperformed others by balancing contrast, structure, and perceptual realism across varied scenarios. Common artifacts, such as highlight clipping (Drago method), oversaturation (Mantiuk method), and excessive darkening (Reinhard method), were mitigated, while surpassing the ATT method in global tonal balance and local detail fidelity. By sequentially optimizing perceptual metrics and refining luminance and contrast, the robustness of the proposed approach across diverse lighting conditions, scene complexities, and texture distributions was proven.

6. Conclusions and Future Work

A two-stage optimization approach for HDR tone mapping that jointly refines the TMQI evaluation metric and the tone mapping operator was presented. By adapting perceptual parameters in the first stage and optimizing tone curve parameters in the second, the proposed method achieved consistent improvements in both quantitative metrics and visual quality. Evaluations on benchmark datasets demonstrated superior

performance over several established tone mapping techniques, validating the method's ability to produce perceptually faithful LDR images.

Future work will explore extending this approach to HDR video, addressing temporal coherence and real-time processing needs. Additionally, the integration of learned perceptual models or semantic scene analysis may further enhance adaptation to diverse visual content. These directions could expand the method's applicability in practical domains, such as immersive media, mobile imaging, and interactive rendering systems.

Funding: This work was funded by the Deanship of Research and Graduate Studies at King Khalid University, KSA, through the Small Research Group under grant number (219/46).

Data Availability Statement: All datasets employed in this study are publicly accessible benchmark HDR datasets. Specifically, the study utilized data from the Funt dataset [24,25], the publicly available HDR dataset by Greg Ward [26], and the UltraFusion benchmark introduced by Chen et al. [27].

Acknowledgments: The author extend their appreciation to the Deanship of Research and Graduate Studies at King Khalid University, KSA, for funding this work through Small Research Group under grant number (219/46).

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

Abbreviation	Definition
HDR	High Dynamic Range
LDR	Low Dynamic Range
TMO	Tone Mapping Operator
TMQI	Tone Mapping Quality Index
S	Structural Fidelity
N	Naturalness
BRISQUE	Blind/Referenceless Image Spatial Quality Evaluator
NIQE	Natural Image Quality Evaluator
ATT	Adaptive Tone mapping Technique
TVI	Threshold Versus Intensity
JND	Just Noticeable Difference
HVS	Human Visual System

References

1. Kunkel, T.; Wanat, R. Achieving and maintaining perceptual intent with hdr imaging. *Inf. Disp.* **2019**, *35*, 9–26. [[CrossRef](#)]
2. See, Z.S.; Goodman, L.; Hight, C.; Sunar, M.S.; Dey, A.; Ng, Y.K.; Billingham, M. Creating high fidelity 360° virtual reality with high dynamic range spherical panorama images. *Virtual Creat.* **2019**, *9*, 73–109. [[CrossRef](#)]
3. Guindy, M.; Kara, P.A.; Balogh, T.; Simon, A. Analysis of high dynamic range light field images in practical utilization contexts. In Proceedings of the Novel Optical Systems, Methods, and Applications XXV, San Diego, CA, USA, 21–26 August 2022; SPIE: Bellingham, WA, USA, 2022; Volume 12216, pp. 144–152.
4. Salih, Y.; Md-Esa, W.b.; Malik, A.S.; Saad, N. Tone mapping of HDR images: A review. In Proceedings of the 2012 4th International Conference on Intelligent and Advanced Systems (ICIAS2012), Kuala Lumpur, Malaysia, 12–14 June 2012; IEEE: New York, NY, USA, 2012; Volume 1, pp. 368–373.
5. Yeganeh, H.; Wang, Z. Objective Quality Assessment of Tone-Mapped Images. *IEEE Trans. Image Process.* **2013**, *22*, 657–667. [[CrossRef](#)] [[PubMed](#)]
6. Reinhard, E.; Stark, M.; Shirley, P.; Ferwerda, J. Photographic tone reproduction for digital images. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*; Association for Computing Machinery: New York, NY, USA, 2023; pp. 661–670.
7. Drago, F.; Myszkowski, K.; Annen, T.; Chiba, N. Adaptive logarithmic mapping for displaying high contrast scenes. In *Computer Graphics Forum*; Wiley Online Library: Hoboken, NJ, USA, 2003; Volume 22, pp. 419–426.

8. Khan, I.R.; Rahardja, S.; Khan, M.M.; Movania, M.M.; Abed, F. A tone-mapping technique based on histogram using a sensitivity model of the human visual system. *IEEE Trans. Ind. Electron.* **2017**, *65*, 3469–3479. [\[CrossRef\]](#)
9. Khan, I.R.; Aziz, W.; Shim, S.O. Tone-mapping using perceptual-quantizer and image histogram. *IEEE Access* **2020**, *8*, 31350–31358. [\[CrossRef\]](#)
10. Quraishi, A.R.; Nadeem, S.; Khan, N.A.; Khan, I.R. Optimizing Parameters of Tone-Mapping Operation for Maximizing an Objective Function. In Proceedings of the 2022 IEEE 24th International Workshop on Multimedia Signal Processing (MMSP), Shanghai, China, 26–28 September 2022; IEEE: New York, NY, USA, 2022; pp. 1–6.
11. Cao, P.; Le, C.; Fang, Y.; Ma, K. A Perceptually Optimized and Self-Calibrated Tone Mapping Operator. *IEEE Trans. Vis. Comput. Graph.* **2025**, *31*, 8268–8282. [\[CrossRef\]](#) [\[PubMed\]](#)
12. Mantiuk, R.; Daly, S.; Kerofsky, L. Display adaptive tone mapping. In *ACM SIGGRAPH 2008 Papers*; Association for Computing Machinery: New York, NY, USA, 2008; pp. 1–10.
13. Zhang, F.; Liu, Y.; Wu, S.; Liu, W. Fusion of global and local Tone Mapping for high-dynamic range images. In Proceedings of the 2015 International Conference on Information and Communications Technologies (ICT 2015), Xi'an, China, 24–26 April 2015; IET: Stevenage, UK, 2015; pp. 1–5.
14. Ambalathankandy, P.; Ikebe, M.; Yoshida, T.; Shimada, T.; Takamaeda, S.; Motomura, M.; Asai, T. An Adaptive Global and Local Tone Mapping Algorithm Implemented on FPGA. *IEEE Trans. Circuits Syst. Video Technol.* **2020**, *30*, 3015–3028. [\[CrossRef\]](#)
15. Cheng, H.; Wei, Y.; Zhang, J.; Xu, Z. A detail-maintained tone mapping algorithm based on weighted least squares filter. In Proceedings of the 2019 IEEE 2nd International Conference on Automation, Electronics and Electrical Engineering (AUTEEE), Shenyang, China, 22–24 November 2019; IEEE: New York, NY, USA, 2019; pp. 269–276.
16. Panetta, K.; Kezebou, L.; Oludare, V.; Agaian, S.; Xia, Z. Tmo-net: A parameter-free tone mapping operator using generative adversarial network, and performance benchmarking on large scale hdr dataset. *IEEE Access* **2021**, *9*, 39500–39517. [\[CrossRef\]](#)
17. Shopovska, I.; Aelterman, J.; Van Hamme, D.; Philips, W. Low-Complexity Deep HDR Fusion And Tone Mapping for Urban Traffic Scenes. In Proceedings of the 2023 IEEE Intelligent Vehicles Symposium (IV), Anchorage, AK, USA, 4–7 June 2023; IEEE: New York, NY, USA, 2023; pp. 1–6.
18. Shopovska, I.; Stojkovic, A.; Aelterman, J.; Van Hamme, D.; Philips, W. High-dynamic-range tone mapping in intelligent automotive systems. *Sensors* **2023**, *23*, 5767. [\[CrossRef\]](#) [\[PubMed\]](#)
19. Rana, A.; Singh, P.; Valenzise, G.; Dufaux, F.; Komodakis, N.; Smolic, A. Deep tone mapping operator for high dynamic range images. *IEEE Trans. Image Process.* **2019**, *29*, 1285–1298. [\[CrossRef\]](#)
20. Guo, C.; Jiang, X. Deep tone-mapping operator using image quality assessment inspired semi-supervised learning. *IEEE Access* **2021**, *9*, 73873–73889. [\[CrossRef\]](#)
21. Rana, A.; Valenzise, G.; Dufaux, F. Learning-based tone mapping operator for efficient image matching. *IEEE Trans. Multimed.* **2018**, *21*, 256–268. [\[CrossRef\]](#)
22. Todorov, P.; Hartig, J.; Meyer-Siemon, J.; Fiedler, M.; Schewior, G. TGTM: TinyML-based Global Tone Mapping for HDR Sensors. *arXiv* **2024**, arXiv:2405.05016.
23. Frazier, P.I. A tutorial on Bayesian optimization. *arXiv* **2018**, arXiv:1807.02811. [\[CrossRef\]](#)
24. Funt, B.; Shi, L. The rehabilitation of MaxRGB. In Proceedings of the Eighteenth Color Imaging Conference, San Antonio, TX, USA, 8–12 November 2010; Simon Fraser University: Burnaby, BC, Canada, 2010.
25. Funt, B.; Shi, L. The effect of exposure on MaxRGB color constancy. In Proceedings of the Human Vision and Electronic Imaging XV, San Jose, CA, USA, 17–21 January 2010; SPIE: Bellingham, WA, USA, 2010; Volume 7527, pp. 282–288.
26. Ward, G. High Dynamic Range Image Examples. 2025. Available online: <http://www.anyhere.com/gward/hdrenc/pages/originals.html> (accessed on 1 June 2025).
27. Chen, Z.; Wang, Y.; Cai, X.; You, Z.; Lu, Z.; Zhang, F.; Guo, S.; Xue, T. UltraFusion: Ultra High Dynamic Imaging using Exposure Fusion. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Nashville, TN, USA, 10–17 June 2025; pp. 16111–16121. Available online: <https://openimaginglab.github.io/UltraFusion/> (accessed on 25 September 2025).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.