Dawn: A Robust Tone Mapping Operator for Multi-Illuminant and Low-Light Scenarios

Furkan Kınlı[1,a], Barış Özcan[2,b] and Furkan Kıraç[1,c]
[1]*Department of Computer Science, Özyeğin University, İstanbul, Turkey*
[2]*Department of Computer Engineering, Bahçeşehir University, İstanbul, Turkey*
{furkan.kinli, furkan.kirac}@ozyegin.edu.tr, baris.ozcan@bau.edu.tr

Keywords: High Dynamic Range, Image Statistics, Low-Light Imaging, Naka-Rushton Equation, Tone Mapping Operator, Retinex.

Abstract: We introduce Dawn, a novel Tone Mapping Operator (TMO) designed to address the limitations of state-of-the-art TMOs such as Flash and Storm, particularly in challenging lighting conditions. While existing methods perform well in stable, well-lit, single-illuminant environments, they struggle with multi-illuminant and low-light scenarios, often leading to artifacts, amplified noise, and color shifts due to the additional step to adjust overall scene brightness. Dawn solves these issues by adaptively inferring the scaling parameter for the Naka-Rushton Equation through a weighted combination of luminance mean and variance. This dynamic approach allows Dawn to handle varying illuminant conditions, reducing artifacts and improving image quality without requiring additional adjustments to scene brightness. Our experiments show that Dawn matches the performance of current state-of-the-art TMOs on HDR datasets and outperforms them in low-light conditions, providing superior visual results. The source code for Dawn will be available at https://github.com/birdortyedi/dawn-tmo/.

1 INTRODUCTION

Tone Mapping Operators (TMOs) are essential in High Dynamic Range (HDR) imaging, enabling the compression of HDR content into a Standard Dynamic Range (SDR) format while preserving essential visual details. TMOs typically operate on the luminance channel, often calculated from the Y channel in the YUV color space (Koschan and Abidi, 2008), with alternative representations available in other color spaces such as HSV, HSL and Lab (Banić and Lončarić, 2014; Nguyen and Brown, 2017). Two major categories of TMOs exist: global and local. Global TMOs apply the same transformation to all pixels, offering faster processing and making them more suitable for real-time applications (Tumblin and Rushmeier, 1993; Larson et al., 1997). On the other hand, local TMOs process intensities based on spatial location by providing higher quality results at the cost of increased computational complexity (Durand and Dorsey, 2002; Reinhard et al., 2023; Mantiuk et al., 2006; Mantiuk et al., 2008).

Recent advances in tone mapping literature (Banić and Lončarić, 2016; Banic and Loncaric, 2018) have focused on achieving a balance between computational efficiency and image quality. These operators, commonly known as Flash and Storm, use the Naka-Rushton equation to model the human visual response to the luminance channel and offer a per-pixel complexity of $O(1)$, which makes them highly practical for real-time applications. However, their effectiveness is limited to well-lit, single-illuminant conditions. In more complex scenarios, such as low-light or multi-illuminant scenes, these TMOs often introduce artifacts such as noise amplification and color distortions, shown in Figure 1. Although deep learning-based TMOs have shown promising results, they demand significant computational resources, making them unsuitable for real-time applications. Evaluating deep learning methods would also require additional performance metrics, such as training time and memory usage, which would shift the focus away from our primary objective of developing an adaptive, non-learning-based solution optimized for real-time tone mapping. For these reasons, deep learning approaches are considered beyond the scope of this study.

[a] https://orcid.org/0000-0002-9192-6583
[b] https://orcid.org/0000-0001-8598-1239
[c] https://orcid.org/0000-0001-9177-0489

Kınlı, F., Özcan, B. and Kıraç, F.
Dawn: A Robust Tone Mapping Operator for Multi-Illuminant and Low-Light Scenarios.
DOI: 10.5220/0013134600003912

In *Proceedings of the 20th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2025) - Volume 3: VISAPP*, pages 62-68
ISBN: 978-989-758-728-3; ISSN: 2184-4321

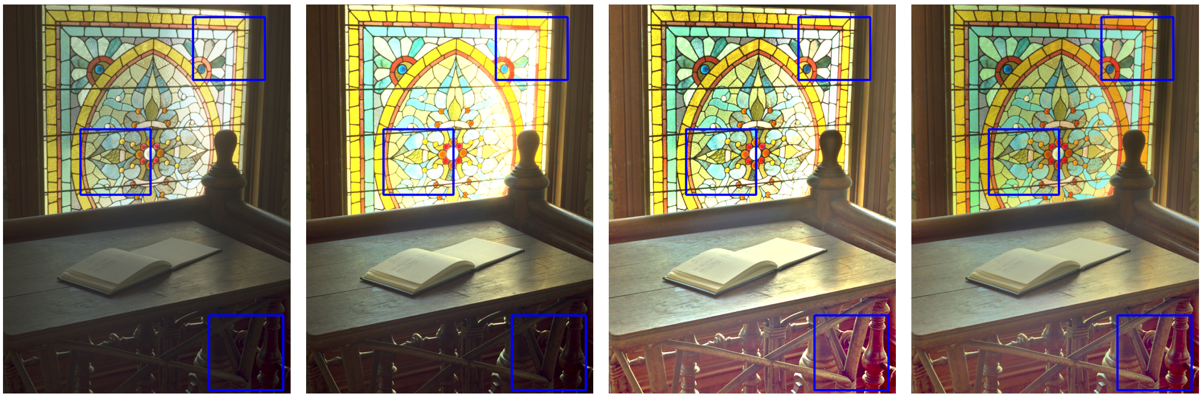

(a) Reinhard (b) $\text{Flash}_{10}\text{+L}_{110}$ (c) Storm_{20} (d) Dawn (**ours**)

Figure 1: Comparison of tone mapping results in complex lighting conditions. Flash and Storm (Banic and Loncaric, 2018) produce noise amplification and color distortions, as exemplified in regions marked with blue boxes, while *Dawn* effectively handles challenging scenarios by preserving image quality.

In response to the limitations of current TMOs in low-light and multi-illuminant scenarios, we propose *Dawn*, a novel TMO designed to dynamically adapt to varying lighting conditions in scenes. *Dawn* leverages image statistics—specifically, the luminance mean and variance—to infer the scaling parameter for the Naka-Rushton Equation, allowing it to reduce artifacts and preserve color distribution and image quality. Our method provides significant improvements in challenging scenarios where existing TMOs struggle, offering robust performance while maintaining computational efficiency.

Our contributions can be summarized as follows.

- **Adaptation to Complex Lighting.** *Dawn* introduces a robust method for handling multi-illuminant and low-light conditions, significantly reducing noise and color distortions.
- **Dynamic Scaling Mechanism.** By using image luminance statistics to infer the Naka-Rushton scaling parameter adaptively, *Dawn* outperforms the current state-of-the-art TMOs using a static or user-defined scaling parameter in challenging scenarios.
- **Overall Robust Performance.** Experimental results demonstrate that *Dawn* matches the performance of existing TMOs in HDR dataset (Ward, 2015) and outperforms them in scenes containing low-light and multi-illuminant cases.

2 METHODOLOGY

In this section, we describe the methodology behind *Dawn*, a novel Tone Mapping Operator (TMO) that builds upon the foundation of Flash, introducing a more robust and adaptive scaling mechanism that eliminates the need for static or user-defined scaling, and avoids additional brightness adjustments.

2.1 Flash: The Foundation

Flash, introduced by (Banic and Loncaric, 2018), utilizes a global tone mapping approach that compresses the luminance values of HDR images using the Naka-Rushton equation. This method applies a global transformation uniformly to all pixels, which makes the mapping operator computationally efficient. The core equation is formulated as follows.

$$L' = \frac{L}{L + a \cdot L_w} \tag{1}$$

where L represents pixel luminance, L_w is the geometric mean luminance (*i.e.*, image key), and a is a static or user-defined scaling parameter. This static scaling parameter, while efficient, is unable to adapt to varying lighting conditions within an image, resulting in suboptimal performance in complex scenarios containing very bright or low-lit areas in the scene.

To mitigate suboptimal results, Flash employs an additional step, called Leap, which adjusts the overall brightness of the tone-mapped image. Leap is an optional post-processing step intended to correct global brightness by normalizing the mean luminance of the tone-mapped image to a predefined target mean value. Specifically, it is used to ensure that the final LDR output maintains a consistent brightness, particularly when the static scaling parameter a in Flash does not account for the image's varying luminance distribution. The equation for Leap is applied as follows

$$L'_{\text{Leap}} = L' \cdot \frac{M_{\text{target}}}{M_{\text{output}}} \tag{2}$$

where M_{target} is the target mean of L (*i.e.*, predefined or user-defined) and M_{output} is the mean luminance of the tone-mapped image.

While Leap helps maintain consistent brightness in Flash, its reliance on a user-defined target mean luminance introduces complexity and reduces flexibility. This dependency may not be optimal for all images or lighting conditions, particularly in challenging scenarios such as varying illuminant sources or dynamic lighting environments. In these cases, manual adjustment of the target mean can exacerbate inconsistencies, leading to suboptimal tone mapping results.

2.2 Dawn: Robust and Adaptive Scaling for Flash

To address the limitations of Flash and Storm in complex lighting conditions, we propose *Dawn*, a novel Tone Mapping Operator (TMO) that introduces an adaptive scaling mechanism into the core equation. Unlike the static or user-defined scaling parameters used by Flash, *Dawn* dynamically adjusts the scaling parameter a based on the luminance statistics of the image. This adaptive approach eliminates the need for Leap and inherently handles brightness normalization throughout the tone mapping process.

The adaptive scaling parameter a for *Dawn* is computed using the following equation

$$a = k_1 \cdot \mu_L + k_2 \cdot \sigma_L + k_3 \tag{3}$$

where μ_L and σ_L represent the mean and variance of the luminance values, respectively. The constants k_1, k_2, and k_3 play a crucial role in this computation, as they control the influence of brightness and contrast on the scaling mechanism. Specifically, k_1 adjusts the contribution of the mean luminance μ_L, affecting the overall brightness response, while k_2 determines how much variance σ_L affects contrast adaptation. The constant k_3 serves as a base value, ensuring stability in different luminance ranges. By fine-tuning these constants, *Dawn* can be tailored to provide optimal tone mapping in a wide range of lighting conditions.

By leveraging image statistics and sweeping post-processing corrections away, *Dawn* continuously adapts the tone mapping process to each image's luminance distribution, which ensures smooth transitions in sudden brightness changes, minimized artifacts, and optimal brightness. This dynamic scaling mechanism allows *Dawn* to handle varying brightness levels and contrasts more effectively than static parameters. This makes this approach more robust in delivering higher-quality outputs in both low-light and multi-illuminant scenes.

2.3 Why Adaptive Scaling Improves Quality

The adaptive scaling mechanism in *Dawn* offers significant advantages over static parameters by dynamic adjustment with respect to the luminance statistics of each image. In low-light conditions, the Leap operation proposed in (Banic and Loncaric, 2018) frequently amplifies noise as it tries to globally adjust brightness and enhance contrast. In contrast, *Dawn* adapts to luminance variance locally, selectively increasing contrast and recovering details by injecting less amount of noise. Next, in multi-illuminant scenes, where static scaling often causes color shifts or haloing, *Dawn* leads to adjusting to brightness variations in different regions, which tailors the tone mapping to specific lighting conditions and minimizing these artifacts. Moreover, *Dawn* ensures consistent tone mapping across regions with varying brightness, such as shadows, midtones, and highlights, maintaining balanced exposure throughout the scene. This adaptability, which does not require manual adjustments or predefined parameters, enables *Dawn* to handle a wide range of lighting scenarios, from high-contrast daylight to complex, low-light environments, with ease and reliability.

2.4 Nonlinear Scaling for Complex Scenarios

In more extreme lighting environments, *Dawn* can employ an optional nonlinear scaling variant to further enhance performance. The scaling parameter in this case is computed as

$$a = \exp(k_1 \cdot \mu_L) + k_2 \cdot \log(1 + \sigma_L) \tag{4}$$

where μ_L represents the mean luminance of the image, and σ_L is the variance of the luminance values, which captures the contrast within the image. The constants k_1 and k_2 control the contribution of the mean and variance to the scaling process, respectively. Specifically, k_1 governs the degree to which the mean luminance influences the exponential adjustment, while k_2 determines the impact of the variance on the logarithmic correction. The addition of 1 to the logarithmic function ensures numerical stability when handling low contrast values.

This nonlinear approach emphasizes the dynamic response to rapid changes in luminance, offering greater flexibility in complex scenarios. By applying exponential and logarithmic transformations, *Dawn* can adapt more aggressively to scenes with large variations in brightness or contrast, ensuring better preservation of detail and consistency of tone.

(a) Drago (b) Reinhard (c) Flash$_{10}$+L$_{110}$ (d) Storm$_{20\text{-}(..,\frac{1}{16})}$+L$_{110}$ (e) Dawn (**ours**)

Figure 2: Comparison of Tone Mapping Operators (TMOs) across different scenes.

2.5 Implementation Details

The implementation of *Dawn* retains the computational efficiency of Flash and Storm, maintaining a per-pixel complexity of $O(1)$, which is crucial for real-time applications. However, unlike Flash, *Dawn* dynamically adjusts the scaling parameter a based on the luminance statistics of the image and eliminates the need for additional brightness correction steps such as Leap.

The processing starts by calculating the maximum luminance value, $v = \max(R, G, B)$, where R, G, and B are the red, green, and blue channels of the pixel. Using either the adaptive or nonlinear scaling strategy, the scaling parameter a is calculated based on the luminance mean (μ_L) and variance (σ_L). The final luminance value with the nonlinear approach for each pixel is calculated using the following equation

$$L' = \frac{L}{L + (\exp(k_1 \cdot \mu_L) + k_2 \cdot \log(1 + \sigma_L)) \cdot L_w} \quad (5)$$

By relying on this dynamic scaling mechanism, *Dawn* ensures a more natural tone-mapped image, eliminating the need for post-processing brightness adjustments, such as Leap. This simplifies the tone mapping pipeline and reduces the reliance on hyperparameters, making the method more flexible and robust across different lighting conditions. Finally, a gamma correction is applied to adjust the brightness and contrast of the tone-mapped image. After gamma correction, the pixel values are clipped to the valid dynamic range (*i.e.*, [0, 1]).

In our implementation, the constants are set to $k_1 = 0.5$, $k_2 = 0.5$, and $k_3 = 1.0$, striking a balance between the luminance mean and variance, allowing *Dawn* to maintain consistent brightness and contrast across a range of lighting conditions. For the nonlinear version, the constants are $k_1 = 0.5$ and $k_2 = 0.2$, chosen to provide more aggressive scaling in challenging lighting environments. The source code for *Dawn* will be available at https://github.com/birdortyedi/dawn-tmo/.

Looking ahead, local kernel-wise improvements, such as those used in Storm, could be readily integrated into *Dawn* to further enhance performance in scenarios involving significant local brightness variations. These adjustments would allow for more localized control of tone mapping, enhancing its ability to handle highly complex lighting environments.

3 RESULTS AND DISCUSSION

The qualitative comparison in Figure 2 highlights the performance of various TMOs in different chal-

Table 1: Quantitative comparison on HDR dataset (Ward, 2015). Metrics used in this comparison: TMQI and FSITM^G_TMQI. Cumulative computation times are also provided, which includes the metric computation.

TMO	**TMQI**	**FSITM^G_TMQI**	**t(s)**
Ashikhmin (Debevec and Gibson, 2002)	0.6620	0.7338	225.23
Drago (Drago et al., 2003)	0.7719	0.8158	30.69
Durand (Durand and Dorsey, 2002)	0.8354	0.8405	225.14
Fattal (Fattal et al., 2023)	0.7198	0.7810	64.78
Mantiuk (Mantiuk et al., 2006)	0.8225	0.8266	88.03
Mantiuk (Mantiuk et al., 2008)	0.8443	0.8494	36.20
Pattanaik (Pattanaik et al., 2000)	0.6813	0.7635	46.91
Reinhard (Reinhard et al., 2023)	0.8695	0.8581	33.41
Reinhard (Reinhard and Devlin, 2005)	0.6968	0.7679	30.01
Flash$_{10}$ (Banic and Loncaric, 2018)	0.8072	0.8315	**21.19**
Flash$_{10}$+Leap$_{110}$ (Banic and Loncaric, 2018)	**0.8755**	0.8625	**21.26**
Storm$_{20\text{ - }(1, \frac{1}{4}, \frac{1}{16})}$ (Banic and Loncaric, 2018)	0.7675	0.8004	24.35
Storm$_{20\text{ - }(1, \frac{1}{4}, \frac{1}{16})}$+Leap$_{110}$ (Banic and Loncaric, 2018)	**0.8782**	0.8551	24.59
Dawn -linear (ours)	**0.8654**	**0.8827**	**21.22**
Dawn -nonlinear (ours)	0.8590	**0.8795**	**21.11**

lenging scenes. The performance of *Dawn* was assessed using the HDR dataset provided in (Ward, 2015), which contains 33 HDR images. Drago and Reinhard preserve midtones but struggle with high-luminance regions, particularly in scenes like the cathedral, where significant blooming and loss of detail occur in bright areas. Flash improves highlight handling but introduces artifacts due to its reliance on static scaling and Leap, which flattens details in bright regions, such as the desk lamp scene. Storm mitigates some of these issues, but still suffers from loss of highlight detail and local inconsistencies.

Dawn, however, consistently outperforms the other methods, handling both low-light and high-luminance regions effectively. Zoomed-in regions are highlighted in Figure 3. For example, in the cathedral scene, *Dawn* preserves detail in the bright windows while maintaining contrast in shadow areas. The desk lamp scene also shows balanced highlights without the flattening seen in other operators. Using adaptive scaling and nonlinear adjustments, *Dawn* delivers more natural, artifact-free images under various lighting conditions. Overall, *Dawn* demonstrates superior robustness and consistency, particularly in challenging lighting environments, where other TMOs introduce artifacts or lose critical details.

Table 1 shows the quantitative comparison of TMOs based on TMQI (Yeganeh and Wang, 2012), FSITM^G_TMQI (Nafchi et al., 2014). *Dawn*, for both linear and nonlinear methods, delivers competitive performance, with its image quality metrics closely matching or surpassing other leading methods such as Flash and Storm. In particular, the linear variant of *Dawn* achieves one of the highest FSITM^G_TMQI scores, indicating superior visual quality.

(a) Flash$_{10}$+L$_{110}$ (b) Storm$_{20}$+L$_{110}$ (c) Dawn (**ours**)

Figure 3: Zoomed-in regions from Figure 2, highlighting the performance differences between Flash$_{10}$+L$_{110}$, Storm$_{20\text{-}(\ldots, \frac{1}{16})}$+L$_{110}$, and *Dawn*.

In terms of efficiency, *Dawn* maintains low execution times comparable to the fastest TMOs such as Flash. This balance of high-quality results and real-time efficiency makes *Dawn* a strong candidate for practical applications, particularly in environments

(a) Default (b) $\sigma\uparrow$ (c) $\mu\uparrow\sigma\downarrow$ (d) $\mu\downarrow\sigma\downarrow$

Figure 4: Influence of adjusting mean (μ) and variance (σ) on the output of *Dawn* with nonlinear scaling.

(a) Linear (b) Nonlinear

Figure 5: Comparison of tone mapping results using linear and nonlinear scaling methods of *Dawn* in more challenging lighting conditions.

that demand both performance and speed.

Figure 4 illustrates the impact of adjusting the luminance mean (μ) and variance (σ) on *Dawn*, showcasing how our optimized parameter choices affect image quality. In (a), the default image serves as the baseline. Increasing variance in (b) enhances local contrast and detail visibility in high-dynamic regions, such as stained glass. However, this comes at the cost of slight noise amplification in shadow regions. In (c), increasing the mean and decreasing the variance produce a brightening of the overall image, similar to the default in terms of exposure, but at the cost of flattening some details in shadow regions. In (d), decreasing both mean and variance results in an overexposed image. These variations confirm that our selected parameters strike a balance, enhancing brightness and contrast without sacrificing detail, demonstrating the adaptability of *Dawn* to different lighting conditions.

Figure 5 introduces an example of linear and nonlinear scaling methods of *Dawn* under more challenging lighting conditions. This example mainly highlights the effectiveness of the nonlinear scaling approach, particularly in scenarios with rapid variations in brightness. The nonlinear approach demonstrates improved handling of high-contrast areas, such as stained glass in the scene, where it preserves more details in the brightly lit regions. The linear method, while effective, tends to flatten the contrast slightly, resulting in less detail retention in these high-luminance areas.

The different TMOs exhibit distinct performance characteristics in extremely low-light scenarios, as shown in Figure 6. Flash significantly brightens the image, but introduces heavy noise, oversaturation, and color shifts, particularly around the lamps, which results in blown-out highlights and lost details. While enhancing brightness, Storm introduces substantial noise and a strong red-yellow tint across the image, overexposing the lamps and distorting the colors. Next, *Dawn* with linear scaling offers a more balanced approach compared to Flash and Storm, retaining better color accuracy and reducing noise, though some areas still appear overexposed and detail in the shadows is limited. Finally, *Dawn* with nonlinear scaling provides the best overall result, preserving natural colors, controlling noise, and maintaining both highlight and shadow details without overexposure.

4 CONCLUSIONS

In this paper, we introduced *Dawn*, a novel tone mapping operator that builds upon the foundation of Flash and Storm by incorporating an adaptive scaling mechanism based on image luminance statistics. *Dawn* eliminates the need for post-processing corrections like Leap, reducing artifacts and improving image quality in low-light and multi-illuminant conditions. Our results demonstrate that *Dawn* consistently outperforms existing methods in terms of both image

(a) $\text{Flash}_{10}+\text{L}_{110}$ (b) $\text{Storm}_{20\text{-}(..,\ \frac{1}{16})}+\text{L}_{110}$ (c) Dawn -linear (**ours**) (d) Dawn -nonlinear (**ours**)

Figure 6: Comparison of Tone Mapping Operators (TMOs) under an extremely low-light scenario.

quality and computational efficiency, making it a robust solution for real-time tone mapping in diverse lighting environments.

REFERENCES

Banić, N. and Lončarić, S. (2014). Improving the tone mapping operators by using a redefined version of the luminance channel. In *Image and Signal Processing: 6th International Conference, ICISP 2014, Cherbourg, France, June 30–July 2, 2014. Proceedings 6*, pages 392–399. Springer.

Banić, N. and Lončarić, S. (2016). Puma: A high-quality retinex-based tone mapping operator. In *2016 24th European Signal Processing Conference (EUSIPCO)*, pages 943–947. IEEE.

Banic, N. and Loncaric, S. (2018). Flash and storm: Fast and highly practical tone mapping based on naka-rushton equation. In *VISIGRAPP (4: VISAPP)*, pages 47–53.

Debevec, P. and Gibson, S. (2002). A tone mapping algorithm for high contrast images. In *13th eurographics workshop on rendering: Pisa, Italy. Citeseer*, volume 2.

Drago, F., Myszkowski, K., Annen, T., and Chiba, N. (2003). Adaptive logarithmic mapping for displaying high contrast scenes. In *Computer graphics forum*, volume 22, pages 419–426. Wiley Online Library.

Durand, F. and Dorsey, J. (2002). Fast bilateral filtering for the display of high-dynamic-range images. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 257–266.

Fattal, R., Lischinski, D., and Werman, M. (2023). Gradient domain high dynamic range compression. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 671–678.

Koschan, A. and Abidi, M. (2008). *Digital color image processing*. John Wiley & Sons.

Larson, G. W., Rushmeier, H., and Piatko, C. (1997). A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):291–306.

Mantiuk, R., Daly, S., and Kerofsky, L. (2008). Display adaptive tone mapping. In *ACM SIGGRAPH 2008 papers*, pages 1–10.

Mantiuk, R., Myszkowski, K., and Seidel, H.-P. (2006). A perceptual framework for contrast processing of high dynamic range images. *ACM Transactions on Applied Perception (TAP)*, 3(3):286–308.

Nafchi, H. Z., Shahkolaei, A., Moghaddam, R. F., and Cheriet, M. (2014). Fsitm: A feature similarity index for tone-mapped images. *IEEE Signal Processing Letters*, 22(8):1026–1029.

Nguyen, R. M. and Brown, M. S. (2017). Why you should forget luminance conversion and do something better. In *Proceedings of the ieee conference on computer vision and pattern recognition*, pages 6750–6758.

Pattanaik, S. N., Tumblin, J., Yee, H., and Greenberg, D. P. (2000). Time-dependent visual adaptation for fast realistic image display. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 47–54.

Reinhard, E. and Devlin, K. (2005). Dynamic range reduction inspired by photoreceptor physiology. *IEEE transactions on visualization and computer graphics*, 11(1):13–24.

Reinhard, E., Stark, M., Shirley, P., and Ferwerda, J. (2023). Photographic tone reproduction for digital images. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 661–670.

Tumblin, J. and Rushmeier, H. (1993). Tone reproduction for realistic images. *IEEE Computer graphics and Applications*, 13(6):42–48.

Ward, G. (2015). High dynamic range image examples.

Yeganeh, H. and Wang, Z. (2012). Objective quality assessment of tone-mapped images. *IEEE Transactions on Image processing*, 22(2):657–667.