A Color Preserving Tone Mapping Framework in the Intrinsic Domain

S Melcarne[1], P Cyriac[2], J L Dugelay[1], A Artusi[3], F Banterle[4]

[1] Eurecom Research Center, Sophia Antipolis, France
[2] Huawei Nice Research Center, Mougins, France
[3] DeepCamera MRG, CYENS Centre of Excellence (CYENS CoE), Nicosia, Cyprus
[4] Visual Computing Lab, ISTI-CNR, Pisa, Italy

E-mail: simone.melcarne@eurecom.fr

Abstract. Tone mapping is an essential step in an acquisition or a rendering pipeline to map high dynamic range (HDR) content to a reference display range. The simplest tone mapping approach is to apply a function to the luminance channel of an HDR image and then to propagate the change to the red, green, and blue channels. However, this often causes color distortions since luminance and chrominance channels are interdependent, and modifying one affects the other. We propose a novel tone mapping approach that preliminarily decomposes the image into intrinsic components and leverages them to perform the actual operation. This strategy effectively mitigates color distortions, eliminating the need for post-processing color correction required by many state-of-the-art methods, and it also assists tone mapping operators (TMOs), improving the overall image quality. Our method was validated through quantitative metrics and a psychophysical experiment, both demonstrating its effectiveness.

1 Introduction

HDR imaging [1, 2] is a fundamental technology that allows users to have a digital experience that is one step closer to reality [3]. While HDR televisions and displays have entered the consumer markets, tone mapping—the process of adapting HDR content to a display with a smaller dynamic range—remains necessary due to the dynamic range mismatch between content and display. In other markets, this issue is even more pronounced. For instance, HDR head-mounted displays are largely in the prototype phase [4], making tone mapping an essential topic in this domain as well [5]. Although many TMOs have been introduced in the past [1, 2] and more advanced methods using deep learning have started to emerge [6], the majority of such TMOs typically do not handle color distortions between original and compressed content. Their focus is mostly on compressing luminance with either minimal color processing or no processing at all. However, it is important to consider that luminance compression of bright areas leads to an increase in color saturation, and luminance expansion of dark areas desaturates colors. In this work, we present a novel tone mapping framework that, through formulation of the problem in the intrinsic domain, inherently mitigates color distortion. Intrinsic Image Decomposition [7] is a technique that decomposes an image into its shading and reflectance layers. Our method employs this technique and can be applied in a plug-and-play manner to any existing TMO. We evaluate its effectiveness on HDR images spanning a wide luminance range, employing objective metrics and conducting a psychophysical experiment. Additional details and resources are available in the supplementary materials.

2 Related Work

2.1 Post-Processing Color Correction Algorithms

Tone mapping serves a dual purpose, requiring the manipulation of images to compress their absolute luminance range while also adjusting pixel relationships to enhance visible detail and alter overall contrast. Nevertheless, modifications to contrast and luminance frequently result in color shifts. This can negatively impact on saturation and hue of the images.

If the overarching purpose when performing tone mapping is to accurately preserve the authenticity and color representation of a scene, what is often done in the literature is to develop *a posteriori* color correction algorithms that reduce the distortion introduced in the tone mapped image. Schlick's approach [8] to color treatment in tone mapping involves introducing a saturation control parameter s into the equation:

$$C_{d'} = \left(\frac{C_w}{L_w}\right)^s \cdot L_d \qquad s \in (0,1], \tag{1}$$

with C representing each single color channel and L the luminance. The subscripts $_{d'}$, $_d$, and $_w$ are, respectively, corrected tone mapped values, tone mapped values before correction, and HDR values. This method reproduces exactly the traditional approach to tone mapping with the only novelty given by the management of the saturation parameter. However, this formulation has its drawbacks, as it can significantly impact the luminance, deviating from the desired outcome, especially for highly saturated pixels, by altering the chroma and lightness of colors. Mantiuk et al. [9] introduced a linear color correction formula:

$$C_{d'} = \left(\left(\frac{C_w}{L_w} - 1\right) s + 1\right) \cdot L_d \qquad s \in (0,1]. \tag{2}$$

Equation 2 mitigates lightness shifts, which happens in Equation 1, but it introduces a more pronounced hue shift. For both these two methods, the saturation parameter s is supposed to be manually picked by users, but it can also be automated when the TMO is global and differentiable [9].

Artusi et al. [10] proposed a fully automatic methodology that takes as input the original HDR and the tone mapped images, which are converted into the cylindrical color space *ICh* [11] of the *IPT* space [12]. Then, the hue of the tone mapped image is reset using the hue of the HDR image. Finally, the chroma component is scaled in such a way that it approximates the result if the tone mapping of the original HDR image is performed in the *ICh* space. This method can be extended to handle out-of-the-gamut colors [10]. As stated before, our problem is to minimize the color distortion between the input HDR image and the output tone mapped one. This is different from related topics such as Color Appearance Models (CAMs) and gamut mapping. Indeed, our problem does not take as input any viewing environment and display as CAMs and gamut-mapping methods [13].

2.2 Intrinsic Image Decomposition for Editing Tasks

Intrinsic Image Decomposition (IID) [7] is a core mid-level computer vision problem that aims to facilitate a deeper understanding of the image content by disentangling intrinsic scene properties from lighting variations. It separates an image I into two components: reflectance, R, and shading, S. Employing the widely recognized Lambertian-world assumption with a single-channel representation for shading (S_{gray}) and a color-embedded reflectance (R_{RGB}), this decomposition is formulated as a pixel-wise multiplication:

$$I_{RGB} = R_{RGB} \odot S_{gray}, \tag{3}$$

such that R_{RGB} is considered invariant to illumination changes and delivers information about the material's spectral properties and colors of the objects in a scene, while S_{gray} only accounts for the illumination effects over those objects. This decomposition provides a simplified yet insightful understanding of image formation, making it a valuable tool in various applications (*e.g.,* relighting, recoloring) [14] [15].

There is convincing evidence in the literature that numerous processes in the human visual system (HVS) compensate for the impact of light. For example, the HVS adjusts to illumination changes and approximately preserves an object's appearance by making its color stable over a wide range of lighting [16, 17]. If, in this sense, light appears of lesser significance, it is probable that alterations to the shading component will be less noticeable than changes to the reflectance. Limiting modifications to the shading component could be particularly beneficial in the context of tone mapping, as light is largely accountable for the extensive dynamic range in real-world settings. Indeed, the maximum dynamic range generated solely by reflectance is less than two orders of magnitude, whereas the dynamic range of shading surpasses four orders of magnitude [18]. This reasoning legitimizes the idea of applying tone mapping in the shading domain rather than the luminance one. Our work adopts IID as a new application, *i.e.,* applying it to color-preserving HDR tone mapping, a direction not previously explored.

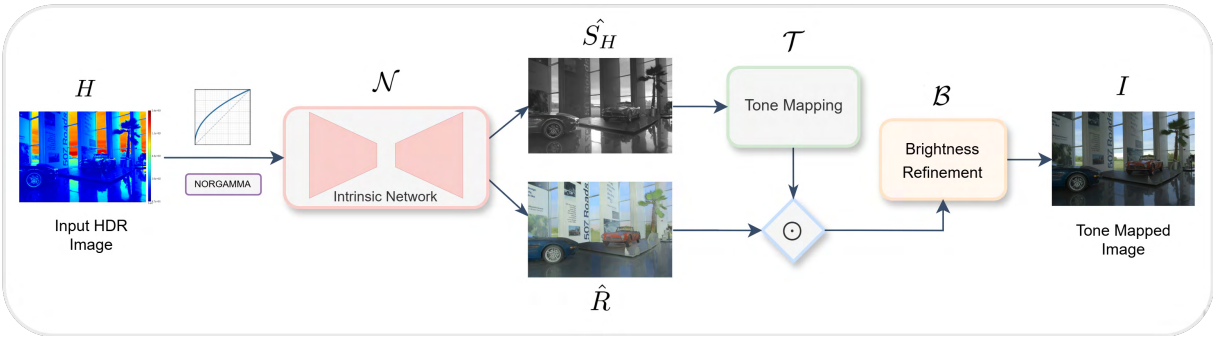

Figure 1: Pipeline of the proposed framework. The encoded HDR image is decomposed into shading and reflectance using an IID network [14]. Tone mapping is applied only to the shading, and the final image is obtained by recombining it with the reflectance, followed by brightness adjustment.

3 Methodology

In this section, we present the proposed framework. Figure 1 shows an overview of the pipeline: we first employ an off-the-shelf deep learning model to perform the intrinsic decomposition after which tone mapping is applied by leveraging the separation of the resulting intrinsic components.

3.1 Decomposition Network

To date, learning-based approaches have emerged as the preferred solution for addressing the IID task. Careaga and Aksoy [14] outlined an architecture that achieves state-of-the-art results in a wide range of possible scenarios (e.g., indoor, outdoor, with human faces, etc). This model was trained on several synthetic datasets [19], [20], [21] that include HDR renderings and corresponding intrinsic ground-truth. In the context of our framework, we conducted a comparative evaluation of two other IID models, CGIntrinsics [21] and PIE-Net [22], using HDR images sourced from the *HDR Photographic Survey* [23], and we performed task-specific fine-tuning on the *Hypersim* dataset [19]. However, Careaga and Aksoy's pretrained model consistently outperformed the others in terms of decomposition quality and computational efficiency. This confirmed its strong generalization ability, making it the most suitable choice for our tone mapping pipeline.

3.2 Input Encoding and Network Inference

As commonly done when applying deep learning models to HDR content, we encoded the images to ensure a suitable input representation before feeding them into the network. To identify the most effective encoding for generating reflectance and shading, we tested three different approaches: HLG [24], PU21 [25], and a robust normalization method (99.5-percentile) followed by 2.2 gamma correction, which we denote as NORGAMMA. Our experiments showed that NORGAMMA preserves more detail and reduces over-exposed regions in the final rendering of a tone mapped image (see Figure 2).

Considering an HDR image H, its NORGAMMA-encoded version $\hat{H}$ is passed through the decomposition network (denoted as $\mathcal{N}$) which computes:

$$\hat{H} = \hat{R} \odot \hat{S_H}, \tag{4}$$

where $\hat{R}$ naturally falls within the [0, 1] range, as it does not convey HDR information, while $\hat{S_H}$ retains the HDR content and remains unbounded.

3.3 Tone Mapping in the Intrinsic Domain

To compress the dynamic range while preserving colors, we formulate the tone mapping in the intrinsic domain. Denoting any TMO as $\mathcal{T}$, the outcome of the tone mapping process step can be expressed as:

$$\mathcal{T}^{IID}(\hat{R}; \hat{S_H}) = \hat{R} \odot \mathcal{T}(\hat{S_H}), \tag{5}$$

where $\mathcal{T}^{IID}(\hat{R}; \hat{S_H})$ is a tone mapped image obtained by first carrying out the tone mapping operation only over the shading, $\mathcal{T}(\hat{S_H})$, and then multiplying the result with the untouched reflectance $\hat{R}$.

(a) HLG [24] (b) PU21 [25] (c) NORGAMMA

Figure 2: An example of the application of different encodings for the IID network when applying our framework to the Drago et al. TMO [26] for the *Jesse Brown's Cabin* [23]. NORGAMMA allows to have a brighter image and to maintain details of trees and the sky.

3.4 Brightness Refinement

Despite preserving the overall structure and colors, we observed that images produced by Equation 5 generally exhibit a low brightness issue, appearing darker than those that would be obtained following the traditional methodology, as a tone mapped shading is typically darker than a tone mapped luminance. We realized that this difference is due to a global scalar multiplier that nearly matches the 99.5-th luminance percentile of the processed image. To address this, we introduced a step called *brightness refinement* ($\mathcal{B}$), where the output image is normalized by this scalar. A more sophisticated solution could be employed (*e.g.,* a refinement neural network to predict this value). To summarize, the complete pipeline to produce the final tone mapped image I can be expressed as:

$$I = \mathcal{B}(\mathcal{T}^{IID}(\mathcal{N}(\hat{H})). \tag{6}$$

4 Results

In this section, we present both quantitative and qualitative results obtained. To generate the test images, we used three traditional TMOs—Reinhard et al. [27], Reinhard and Devlin [28] and Drago et al. [26]—alongside a deep learning-based method by Vinker et al. [29], for a total of four TMOs. We juxtaposed our outcomes with those from Artusi et al.'s method [10] and Mantiuk et al.'s [9] and Schlick's [8] algorithms, with the saturation parameter estimated from the tone curve as described in [9]. For all the global TMOs, we employed the HDR Toolbox [30] implementation (please refer to the supplementary materials for the specifications of the parameter choices). Regarding Vinker et al.'s TMO [29], we employed the original implementation available on github [1], with automatic parameters estimation.

4.1 Objective Evaluation

We identified two evaluation metrics: *Hue Distance* (Δh) [10] and *ColorVideoVDP* (CVVDP) [31]. The Δh quantifies hue distortion between the original HDR image and its tone mapped counterpart, operating in the *ICh* color space [11], which is known for effectively decorrelating intensity and chromatic components. This metric has already been successfully applied to this type of evaluation in [10] and we computed the results using the HDR Toolbox implementation [30]. CVVDP is a perceptual metric that models spatial and temporal aspects of vision for both luminance and color, working for static images as well. To compute this metric, we converted the tone mapped 8-bit images into physical values by applying the display settings of an SDR monitor with peak luminance of 100 cd/m^2. HDR images were also displayed using the calibration parameters of an EIZO CG3145 HDR reference monitor with peak luminance equal to 1000 cd/m^2. We conducted the comparisons considering all 106 HDR images from *HDR Photographic Survey* [23]. Table 1 reports Δh and CVVDP values computed for each proposed method and for every tested TMO. Our method achieves the best performance across nearly all combinations of TMOs, consistently reducing hue distortion and improving perceptual quality. While Artusi et al. [10] specifically target hue preservation, their corrections operate in a color space where luminance and chromaticity are not fully decoupled. This partial coupling can lead to residual hue distortions in highly saturated regions, which our intrinsic domain formulation helps mitigate. It is also important to note that small hue changes can still be perceptible, as this value refers to the entire image, and such distortions are often concentrated in vivid colors covering small areas.

[1] Accessed on the 20-th of April 2025, https://github.com/yael-vinker/unpaired_hdr_tmo

Table 1: Δh and CVVDP scores for the tested TMOs on the *HDR Photographic Survey*[23]. Best in **bold**.

	Reinhard et al. [27]		Reinhard and Devlin [28]		Drago et al. [26]		Vinker et al. [29]	
Method	$\Delta h\downarrow$	CVVDP↑	$\Delta h\downarrow$	CVVDP↑	$\Delta h\downarrow$	CVVDP↑	$\Delta h\downarrow$	CVVDP↑
Schlick [8]	0.0142	7.81	0.0328	5.00	0.0137	6.92	0.0264	6.57
Mantiuk [9]	0.0152	7.67	0.0350	4.78	0.0145	6.79	0.0279	6.50
Artusi [10]	**0.0124**	7.82	0.0129	5.35	0.0125	6.94	0.0181	6.60
Ours	0.0157	**8.70**	**0.0120**	**7.49**	**0.0122**	**7.92**	**0.0137**	**7.35**

Table 2: Ranking of methods; total votes in parentheses. Methods are grouped together if they are statistically the same, using the critical value $R = 85$ at significance level $\alpha = 0.01$ [32].

	Groups		
	3rd	2nd	1st
Overall	Schlick(211)	Mantiuk(315) Artusi(392)	Ours(582)

4.2 Psychophysical Experiment

To validate the effectiveness of the proposed approach, we also conducted a psychophysical experiment. The experiment involved evaluating 25 images randomly selected out of the 106 from the *HDR Photographic Survey* [23] (selected images are shown in the supplementary materials). The type of experiment is 2AFC (2-Alternative-Forced-Choice), meaning that each time participants were asked to choose between two options. For the criterion of choice, participants were asked: *'Among the proposed options, which tone mapped image is most similar to the reference HDR image?'*. To reproduce the content of HDR images, we used the EIZO CG3145 HDR monitor with a peak luminance equal to 1000 cd/m^2. Initially, the reference HDR image is displayed alone in full screen. Upon pressing a key, two SDR tone mapped images are presented side by side. When the SDR images are shown for pair comparison, the monitor is calibrated to reproduce the luminance characteristics of a standard SDR display, with brightness limited to a peak of 100 cd/m^2. To control absolute luminance on the HDR monitor, the SDR images are mapped to PQ encoding [33] while preserving SDR brightness levels. The experiment involved 10 participants, comprised of 70% males and 30% females, all of whom were experts in color assessment and aged between 28 and 45 years. There were no time constraints imposed for making the choices. For each participant, the experiment consisted of a single session lasting approximately 40 minutes on average. The 2AFC protocol requires a large number of repeated comparisons, which significantly increases session duration. Therefore, we opted for a limited but qualified group of participants. This design choice is in line with common practice in tone mapping studies, which often rely on a comparable number of participants [29] [10]. Please, refer to the supplementary materials for more details. In this experiment, we tested images tone mapped using Reinhard et al.'s TMO [27]. The reasoning behind this choice is that we wanted to assess the quality of an operator which had not the best results in terms of Δh, even though visual comparisons showed promising results. We analyzed the collected data using paired comparison analysis [32], as previously employed in the HDR literature [34]. We computed: i) the coefficient of consistency u; ii) the coefficient of agreement ζ. The coefficient of consistency, $u \in [0, 1]$, indicates the degree to which voters are consistent in their preferences, with 0 meaning no consistency and 1 meaning perfect consistency. The coefficient of agreement, $\zeta \in [-1, 1]$, measures the level of agreement among voters, with -1 indicating complete disagreement and 1 indicating full agreement. Our results show that participants were highly consistent ($\zeta = 0.918$) and demonstrated significant agreement ($u = 0.199$). These results were statistically significant, as indicated by $\mathcal{X}^2 = 303.26$ (critical value at significance level $\alpha = 0.01$ is 16.81). Furthermore, the difference in scores, $D_n = 294.53$, was also statistically significant, exceeding the critical value at $\alpha = 0.01$ (11.33). These two tests are equivalent to the analysis of variance and Tukey's test respectively [32]. Table 2 shows the results of ranking the overall methods based on votes; note that we group methods that that are not distinguishable according to the critical value $R = 85$ [32]. From this table, we can notice that our method got most votes and differs significantly from the other methods. Moreover, our experiment confirms the literature, demonstrating that Mantiuk et al. [9] and Artusi et al. [10] methods achieve superior performance compared to Schlick [8].

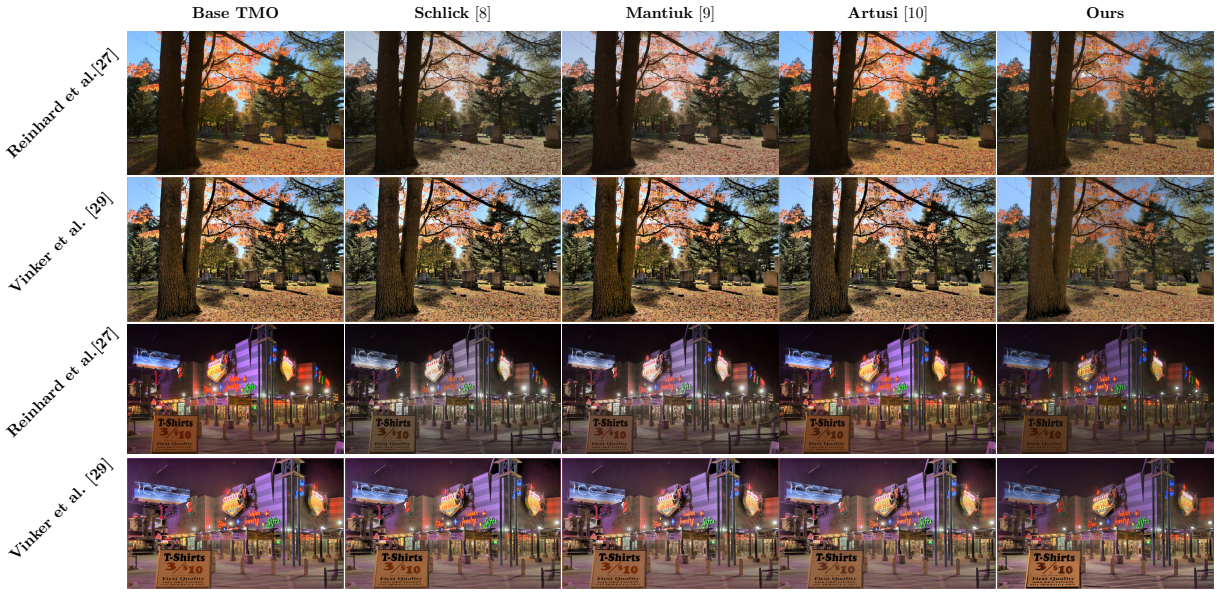

Figure 3: Visual comparison of our method against the state-of-the-art. The first column shows an HDR image tone mapped with a global TMO (Reinhard et al. [27]) and a deep-learning method (Vinker et al. [29]). Remaining columns report results for each specific method. From top to bottom: *Cemetery Tree* and *LasVegasStore* images from [23]. See supplementary materials for further comparisons.

4.3 Visual Comparison

Figure 3 compares images obtained by applying state-of-the-art color correction methods and following our preliminary decomposition approach, using [27, 29] as base TMOs. The strength of our framework lies in its ability to maintain accurate color reproduction. TMOs tend to disrupt the saturation of colors, resulting in less realistic visual outputs. Post-processing methods limit the problem but sometimes end up desaturating or oversaturating the images. In contrast, our approach prevents color distortion from the beginning and returns more balanced images (see the color of the leaves in *Cemetery Tree*). Furthermore, in challenging bright regions, finer details are more visible with our method, where others might fade or lose some information (see the neon signs in *LasVegasStore*), and it can also reduce halo artifacts in local TMOs, as observed in the same image tone mapped with Vinker et al. TMO [29], which is a local one.

4.4 Limitations

While our method performs well in most scenarios, we have observed that, in rare instances, it can produce images with lower brightness levels. This is particularly noticeable in predominantly dark images, such as night or in indoor scenes with few bright areas present (see *Peppermill* and *WaffleHouse* in Figure 5 of the supplementary materials) and it may occur because the reflectance-shading separation is not fully successful, given a case that probably the network cannot model. A possible solution could involve adapting the IID model to better handle these scenarios.

5 Conclusions

We presented a color-preserving tone mapping framework formulated in the intrinsic domain. Preferring this preliminary decomposition demonstrated advantages in minimizing color distortion and improving visual fidelity, outperforming post-processing state-of-the-art methods for color correction in most cases. The psychophysical experiment has confirmed this; *i.e.,* the analysis has shown participants consistently preferred images processed through our framework. Our work can be applied to any existing TMO in a plug-and-play fashion, requiring minimal effort for integration making it ideal for production pipelines.

While the current approach avoids additional training overhead by leveraging a pre-trained IID model, we acknowledge that jointly optimizing the IID and tone mapping modules by training the decomposition end-to-end with a learnable TMO could further enhance performance. We consider this direction as future research. Another possible extension involves incorporating temporal consistency constraints to handle HDR video tone mapping, ensuring coherent decomposition and appearance across frames.

References

[1] Reinhard E, Ward G, Pattanaik S N, Debevec P E, and Heidrich W, *High dynamic range imaging: acquisition, display, and image-based lighting*, Academic Press, 2 edition, 2010.

[2] Banterle F, Artusi A, Debattista K, and Chalmers A, *Advanced high dynamic range imaging*, A. K. Peters, USA, 2 edition, 2017.

[3] Zhong F, Jindal A, Yöntem A Ö, Hanji P, Watt S J, and Mantiuk R K, "Reproducing reality with a high-dynamic-range multi-focal stereo display," *ACM Trans. Graph.*, vol. 40, no. 6, pp. 241:1–241:14, 2021.

[4] Matsuda N, Zhao Y, Chapiro A, Smith C, and Lanman D, "Hdr vr," in *ACM SIGGRAPH 2022 Emerging Technologies*, New York, NY, USA, 2022, ACM.

[5] Tariq T, Matsuda N, Penner E, Jia J, Lanman D, Ninan A, and Chapiro A, "Perceptually adaptive real-time tone mapping," in *SIGGRAPH Asia 2023 Conference Papers*, New York, NY, USA, 2023, ACM.

[6] Wang L and Yoon K-J, "Deep learning for hdr imaging: state-of-the-art and future trends," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 12, pp. 8874–8895, 2022.

[7] Garces E, Rodríguez-Pardo C, Casas D, and Lopez-Moreno J, "A survey on intrinsic images: delving deep into lambert and beyond," *Int. J. Comput. Vis.*, vol. 130, no. 3, pp. 836–868, 2022.

[8] Schlick C, "Quantization techniques for visualization of high dynamic range pictures," in *Photorealistic rendering techniques*, Sakas G, Müller S, and Shirley P, Eds., pp. 7–20. Springer, Berlin, Heidelberg, 1995.

[9] Mantiuk R, Mantiuk R, Tomaszewska A M, and Heidrich W, "Color correction for tone mapping," *Comput. Graph. Forum*, vol. 28, no. 2, pp. 193–202, 2009.

[10] Artusi A, Pouli T, Banterle F, and Akyuz A O, "Automatic saturation correction for dynamic range management algorithms," *Signal Process. Image Commun.*, vol. 63, pp. 110–120, April 2018.

[11] Shen S, *Color difference formula and uniform color space modeling and evaluation*, Ph.D. thesis, Rochester Institute of Technology, 2008.

[12] Ebner F and Fairchild M D, "Development and testing of a color space (ipt) with improved hue uniformity," in *6th Color and Imaging Conference, CIC 1998*, Scottsdale, Arizona, USA, 1998, Society for Imaging Science and Technology, pp. 8–13.

[13] Touraj T and Hamid M, "Computationally-efficient hue-preserving gamut mapping in RGB and YUV," in *Electronic Imaging*, 2025, vol. 37, pp. 226–1–226–7.

[14] Careaga C and Aksoy Y, "Intrinsic image decomposition via ordinal shading," *ACM Trans. Graph.*, vol. 43, no. 1, November 2023.

[15] Bonneel N, Kovacs B, Paris S, and Bala K, "Intrinsic decompositions for image editing," *Computer Graphics Forum (Eurographics State of the Art Reports 2017)*, vol. 36, no. 2, 2017.

[16] D'Zmura M and Lennie P, "Mechanisms of color constancy," *J. Opt. Soc. Am. A*, vol. 3, no. 10, pp. 1662–1672, 1986.

[17] McCann J J, "Do humans discount the illuminant?," in *Proc. SPIE*. 2005, vol. 5666, pp. 5666–9, SPIE.

[18] Mantiuk R, Myszkowski K, and Seidel H-P, "High dynamic range imaging," *Wiley Online J.*, 2015, Published on April 18, 2016.

[19] Roberts M, Ramapuram J, Ranjan A, Kumar A, Bautista M A, Paczan N, Webb R, and Susskind J M, "Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*. 2021, pp. 10892–10902, IEEE.

[20] Li Z, Yu T-W, Sang S, Wang S, Song M, Liu Y, Yeh Y-Y, Zhu R, Gundavarapu N, Shi J, Bi S, Yu H-X, Xu Z, Sunkavalli K, Hašan M, Ramamoorthi R, and Chandraker M, "Openrooms: An open framework for photorealistic indoor scene datasets," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. 2021, pp. 7186–7195, IEEE.

[21] Li Z and Snavely N, "Cgintrinsics: Better intrinsic image decomposition through physically-based rendering," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Ferrari V, Hebert M, Sminchisescu C, and Weiss Y, Eds., Cham, 2018, pp. 381–399, Springer.

[22] Das P, Karaoglu S, and Gevers T, "Pie-net: Photometric invariant edge guided network for intrinsic image decomposition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. 2022, pp. 19758–19767, IEEE.

[23] Fairchild M, "Hdr photographic survey," 2008.

[24] ITU-R, "ITU-R 2100-2: Image parameter values for high dynamic range television for use in production and international programme exchange," Tech. Rep., International Telecommunication Union, 2018.

[25] Mantiuk R K and Azimi M, "Pu21: A novel perceptually uniform encoding for adapting existing quality metrics for hdr," in *Proc. Picture Coding Symp. (PCS)*, 2021, pp. 1–5.

[26] Drago F, Myszkowski K, Annen T, and Chiba N, "Adaptive logarithmic mapping for displaying high contrast scenes," in *Proc. Eurographics*, 2003.

[27] Reinhard E, Stark M, Shirley P, and Ferwerda J, "Photographic tone reproduction for digital images," *ACM Trans. Graph.*, vol. 21, no. 3, pp. 267–276, July 2002.

[28] Reinhard E and Devlin K, "Dynamic range reduction inspired by photoreceptor physiology," *IEEE Trans. Vis. Comput. Graph.*, vol. 11, no. 1, pp. 13–24, 2005.

[29] Vinker Y, Huberman-Spiegelglas I, and Fattal R, "Unpaired learning for high dynamic range image tone mapping," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Montreal, QC, Canada, 2021, pp. 14637–14646, IEEE.

[30] Banterle F, Artusi A, Debattista K, and Chalmers A, *Advanced High Dynamic Range Imaging (2nd Edition)*, AK Peters (CRC Press), Natick, MA, USA, July 2017.

[31] Mantiuk R K, Hanji P, Ashraf M, Asano Y, and Chapiro A, "Colorvideovdp: A visual difference predictor for image, video and display distortions," *ACM Trans. Graph.*, vol. 43, no. 4, 2024.

[32] David H A, *The method of paired comparisons*, Oxford University Press, 2 edition, 1988.

[33] Miller S, Nezamabadi M, and Daly S, "Perceptual signal coding for more efficient usage of bit codes," *SMPTE Motion Imaging Journal*, vol. 122, no. 4, pp. 52–59, 2013.

[34] Ledda P, Chalmers A, Troscianko T, and Seetzen H, "Evaluation of tone mapping operators using a high dynamic range display," *ACM Trans. Graph.*, pp. 640–648, 2005.

Supplementary Materials: A Color Preserving Tone Mapping Framework in the Intrinsic Domain

S Melcarne[1], P Cyriac[2], J L Dugelay[1], A Artusi[3], F Banterle[4]

[1] Eurecom Research Center, Sophia Antipolis, France
[2] Huawei Nice Research Center, Mougins, France
[3] DeepCamera MRG, CYENS Centre of Excellence (CYENS CoE), Nicosia, Cyprus
[4] Visual Computing Lab, ISTI-CNR, Pisa, Italy

E-mail: simone.melcarne@eurecom.fr

1 Global TMOs parameters choice

In this section, we present the parameters used for the three global tone mapping operators (TMOs) tested in our study. For Reinhard et al.'s [1] TMO parameters, α and L_{white} were estimated using Reinhard's estimation [2]. Regarding Reinhard and Devlin's TMO [3], we employed the default parameters suggested in the paper [3]: $m = 0.3 + 0.7k^{1.4}$ (where k is the key of the image), $f' = 0$, $c = 0$, and $a = 1$. Finally, for Drago et al.'s TMO [4], we employed the default parameters values as in the paper [4]: $L_{dmax} = 100\ cd/m^2$ and $b = 0.85$.

2 Experimental procedure

In Figure 1 we summarize the experimental setup. In an environment deliberately darkened to minimize all light sources, participants first viewed the HDR image, which was displayed alone in full-screen mode. When the participant pressed a key, a grayscale image was projected to allow for eye adaptation. After this adaptation period, participants were asked to express their preference between two SDR images shown side by side. Since there were 4 methods to compare in total, each HDR image required 6 pairwise comparisons of tone mapped SDR images. Participants were allowed to view the HDR image for as long as they wanted before proceeding to compare the SDR image pairs. Additionally, participants could bring up the HDR image again at any time before making their final choice. The frequent switching between HDR and SDR modes was facilitated using custom software. Participants viewed the images at a fixed distance of 80 *cm* from the monitor, with all images displayed in full-HD resolution. Figure 2 shows the 25 "baseline" images selected from the *HDR Photographic Survey* [5], which were used to generate all the variations for the 4 methods analyzed in this study. This resulted in a total of 100 (25×4) test images as experimental stimuli. However, due to space constraints, only the 25 baseline images are shown.

3 Additional results

In Figures 3, 4 and 5, we provide more visual comparisons between the proposed method and the others, showing the results for all tested TMOs. Note that these images represent additional examples of visual comparison from the 106 tested images of the *HDR Photographic Survey* [5] but not all of them were randomly selected for the experiment, except for *LabBooth*, *WaffleHouse* and *Zentrum*, which also appeared in the experiment.

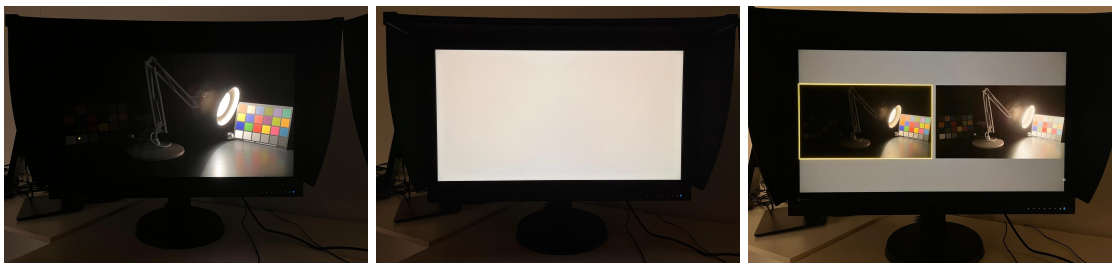

(a) HDR Scene Viewing (b) Visual Adaptation (c) Preference Selection

Figure 1: Experimental setup. The subject looks at the HDR image first. Then, a gray scale image is projected for an adaptation time for the eyes. Finally, the subject is asked to express a preference between 2 tone mapped options and for 6 times for each test image.

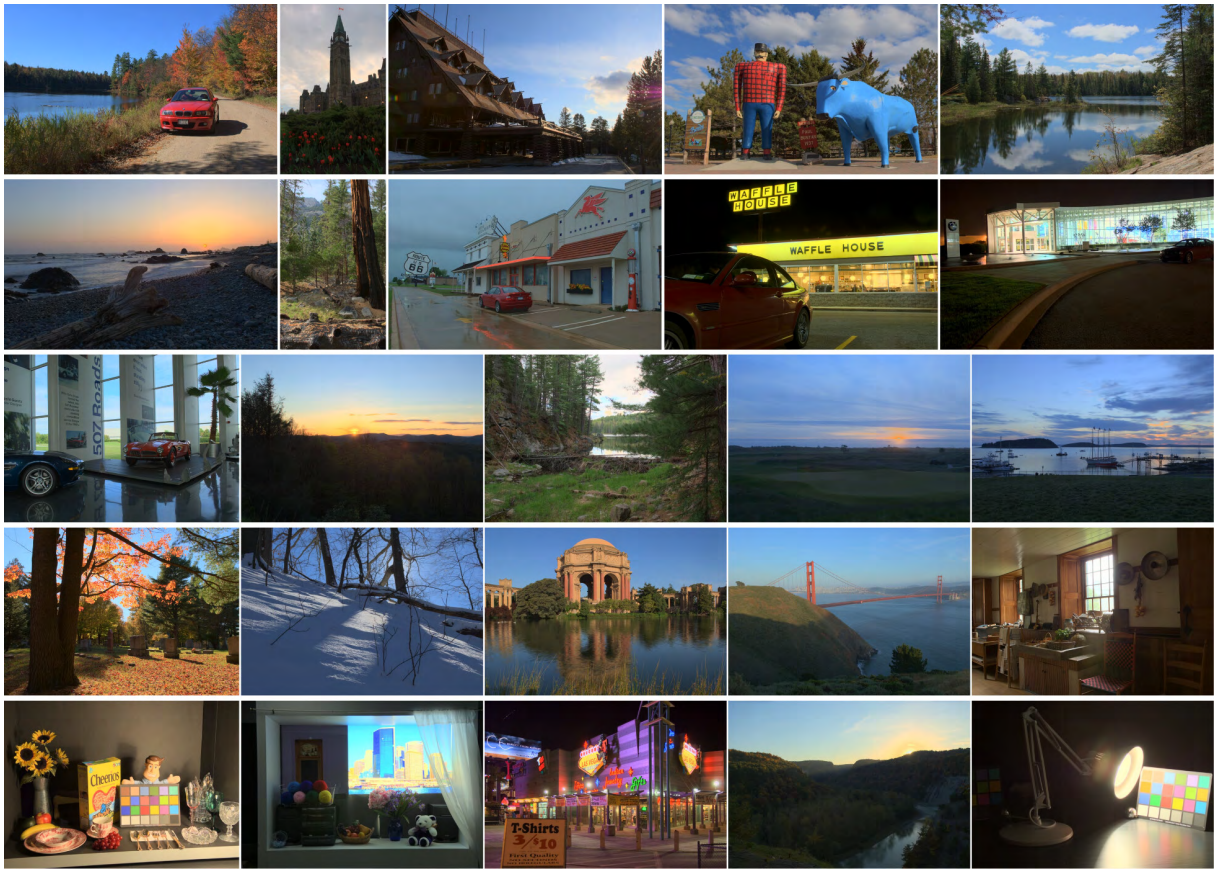

Figure 2: The 25 randomly selected images from the *HDR Photographic Survey* [5] for the experiment. Note that we show the "baseline" images tone mapped for visualization purposes using Reinhard et al.'s TMO [1], and not the versions obtained following each of the tested methods.

Figure 3: Visual comparison of our method against the state-of-the-art. The first column represents an HDR image respectively tone mapped using all the tested TMOs. Remaining columns represent the result obtained when each specific method is employed. In this figure, *LabBooth* and *MasonLake(2)* images from the *HDR Photographic Survey* [5] have been used from top to bottom respectively.

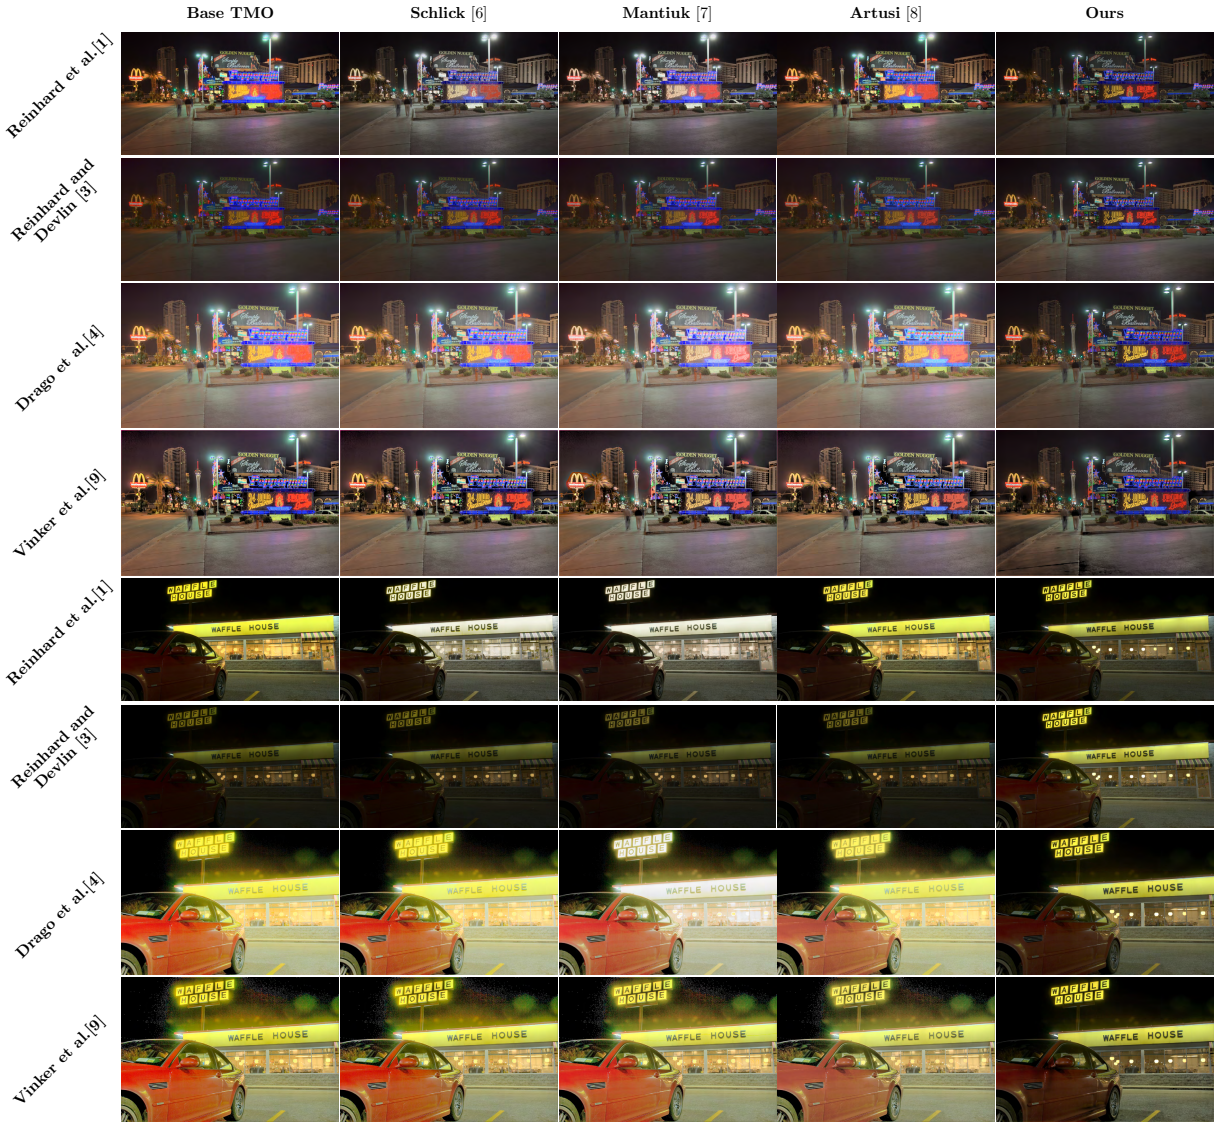

Figure 4: Visual comparison of our method against the state-of-the-art. The first column represents an HDR image respectively tone mapped using all the tested TMOs. Remaining columns represent the result obtained when each specific method is employed. In this figure, *Peppermill* and *Waffle House* images from the *HDR Photographic Survey* [5] have been used from top to bottom respectively.

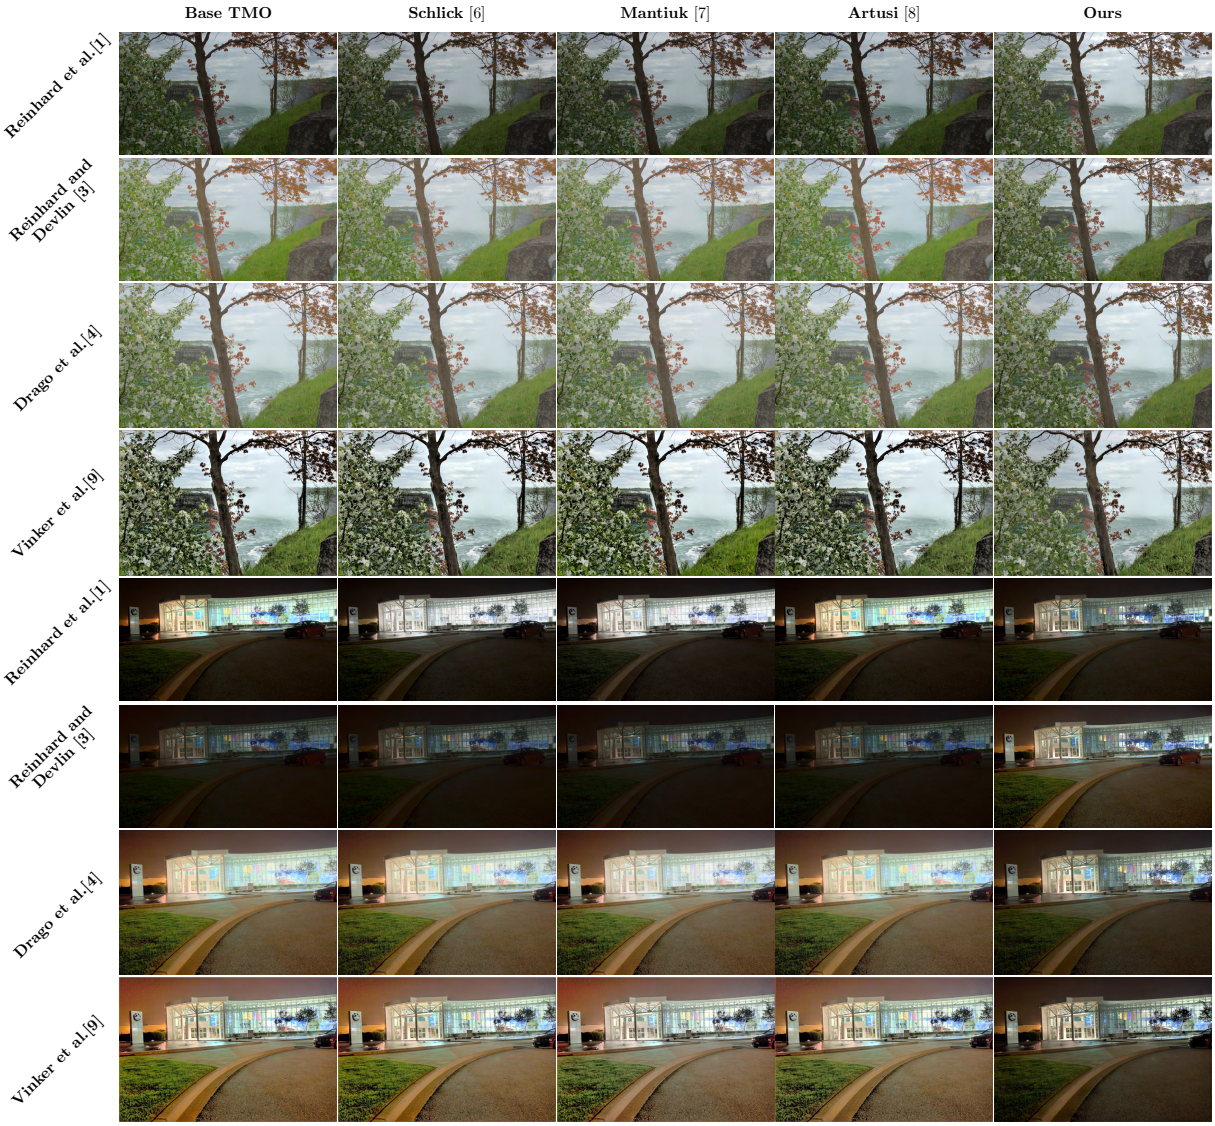

Figure 5: Visual comparison of our method against the state-of-the-art. The first column represents an HDR image respectively tone mapped using all the tested TMOs. Remaining columns represent the results obtained when each specific method is employed. In this figure, *Canadian Falls* and *Zentrum* images from the *HDR Photographic Survey* [5] have been used from top to bottom respectively.

References

[1] Reinhard E, Stark M, Shirley P, and Ferwerda J, "Photographic tone reproduction for digital images," *ACM Trans. Graph.*, vol. 21, no. 3, pp. 267–276, July 2002.

[2] Reinhard E, "Parameter estimation for photographic tone reproduction," *J. Graphics, GPU, & Game Tools*, vol. 7, no. 1, pp. 45–51, 2002.

[3] Reinhard E and Devlin K, "Dynamic range reduction inspired by photoreceptor physiology," *IEEE Trans. Vis. Comput. Graph.*, vol. 11, no. 1, pp. 13–24, 2005.

[4] Drago F, Myszkowski K, Annen T, and Chiba N, "Adaptive logarithmic mapping for displaying high contrast scenes," in *Proc. Eurographics*, 2003.

[5] Fairchild M, "Hdr photographic survey," 2008.

[6] Schlick C, "Quantization techniques for visualization of high dynamic range pictures," in *Photorealistic rendering techniques*, Sakas G, Müller S, and Shirley P, Eds., pp. 7–20. Springer, Berlin, Heidelberg, 1995.

[7] Mantiuk R, Mantiuk R, Tomaszewska A M, and Heidrich W, "Color correction for tone mapping," *Comput. Graph. Forum*, vol. 28, no. 2, pp. 193–202, 2009.

[8] Artusi A, Pouli T, Banterle F, and Akyuz A O, "Automatic saturation correction for dynamic range management algorithms," *Signal Process. Image Commun.*, vol. 63, pp. 110–120, April 2018.

[9] Vinker Y, Huberman-Spiegelglas I, and Fattal R, "Unpaired learning for high dynamic range image tone mapping," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Montreal, QC, Canada, 2021, pp. 14637–14646, IEEE.