

# Improving the color accuracy of lighting estimation models

Zitian Zhang<sup>†</sup>, Joshua Urban Davis<sup>‡</sup>, Jeanne Phuong Anh Vu<sup>‡</sup>, Jiangtao Kuang<sup>‡</sup>, and Jean-François Lalonde<sup>†</sup>.

<sup>†</sup> Université Laval, Canada <sup>‡</sup> Meta, United States

## Abstract

Advances in high dynamic range (HDR) lighting estimation from a single image have opened new possibilities for augmented reality (AR) applications. Predicting complex lighting environments from a single input image allows for the realistic rendering and compositing of virtual objects. In this work, we investigate the color robustness of such methods—an often overlooked yet critical factor for achieving visual realism. While most evaluations conflate color with other lighting attributes (e.g., intensity, direction), we isolate color as the primary variable of interest. Rather than introducing a new lighting estimation algorithm, we explore whether simple adaptation techniques can enhance the color accuracy of existing models. Using a novel HDR dataset featuring diverse lighting colors, we systematically evaluate several adaptation strategies. Our results show that preprocessing the input image with a pre-trained white balance network improves color robustness, outperforming other strategies across all tested scenarios. Notably, this approach requires no retraining of the lighting estimation model. We further validate the generality of this finding by applying the technique to three state-of-the-art lighting estimation methods from recent literature. Our project webpage is available at: <https://lusn.github.io/coloraccuracy>.

## Introduction

Approaches for automatically estimating lighting from images have emerged in the past several years and they have enabled new augmented reality (AR) applications, such as the insertion of realistic virtual content in images. This task aims at recovering a representation of the lighting environment in a scene in high dynamic range (HDR) from a single input image.

Gardner et al. [8] have proposed the first deep learning approach for solving this challenging task. Their method has since been improved along several axes: the accuracy of the predicted HDR [30], the generation of highly detailed environment textures [17], the speed of inference to obtain real-time performance [9, 21], the editability of the estimated lighting to allow for artistic control [27, 24, 4], and predicting spatially-varying lighting representations [14, 23, 31].

In this literature, not much attention has been given to the color accuracy of such models. Indeed, while methods evaluate themselves on standard benchmarks [8] and report a variety of quantitative metrics, color accuracy is not specifically targeted. However, accurate color reproduction is a very important factor in realism when inserting virtual objects in images [18]. Consider the examples in fig. 1, where virtual objects are relit by the prediction of an automatic lighting estimation algorithm. While the overall lighting direction is accurately rendered, the wrong color (fig. 1, middle) makes the object seem out of place. Achieving accurate color rendering by adjusting the lighting estimation model with one of the strategies we explore in this paper (fig. 1, right) results



**Figure 1.** Getting the “color right” is critical for realistic virtual object insertion. Left: input image. Middle: virtual object (sofa) relit with estimated lighting from a baseline model. Right: a lighting estimation model with higher color accuracy results in a more realistic composite.

in a much more believable composition.

In this paper, we aim to systematically study how lighting estimation models produce color-consistent results for a given input image. Specifically, we focus on the case of indoor scenes. To this end, we capture a new test dataset of different indoor environments, where each scene is captured under several different white balance conditions. This simulates the presence of different illuminants in the scenes and, in each case, the corresponding ground truth HDR lighting map is available. We then leverage this dataset, used only at test time, to quantify the impact of various color adaptation strategies on the robustness of the lighting estimation model, especially for the case of virtual object insertion, where an object is relit and composited in the image.

Our goal is to identify strategies that are agnostic to the specific lighting estimation approach, with the aim of providing practitioners a set of “best color practices” they can employ with their own approach. Specifically, we explore the use of color-specific loss function, chromatic data augmentation, and leveraging pre-trained white balance networks at both training and test time. Our experiments show that correcting for the overall color of the input image, with a pre-trained white balance network, achieves the best performance across all tested strategies. This is particularly useful since it requires no re-training of the lighting estimation network itself: it can be used as is. We demonstrate the practical impact of this solution by evaluating the same strategy on two other lighting estimation approaches from the recent literature, and show similar improvements in color accuracy. We hope this paper will both bring attention to the problem of color accuracy in lighting estimation, and provide practitioners with the tools to improve this accuracy in their approaches.

## Related work

Estimating lighting from images has been explored from a variety of angles in recent years. Here, we focus on works who aim to estimate lighting from a single image captured indoors. We refer the reader to [6] for a broader overview of deep learning models for lighting estimation. Gardner et al. [8] pioneered the use of deep learning for lighting estimation in indoor scenes, and proposed a network that can predict an environment map (360° lighting map) in HDR from a single input image. This technique was later improved in several ways. Legendre et al. [13] train a single model for indoor and outdoor images, Song and Funkhouser [22] take into account the object insertion position. Garon et al. [9] focus on speed and infer fast spatially-varying light estimates using spherical harmonics. Somanath and Kurz [21] train a UNet and focus on real-time on-device estimation. A variety of other lighting parametrizations have been used. For instance, parametric lights [7] and spherical gaussians [30, 29] typically provide more controllability on the output since they possess few, intuitive parameters. This has explicitly been exploited by [27, 4] who combine a parametric lighting with powerful GANs for editable lighting estimation. StyleLight [24] employ a GAN-inversion approach to achieve some level of control over the output. It has also been proposed to learn volumetric lighting representations from images. Of note, Lighthouse [23] learns multi-scale volumetric lighting from a stereo pair, [14, 31] predicts a dense 2D grid of spherical gaussians which is further extended into a 3D volumetric representation by Wang et al. [26]. More recently, diffusion models have also been used [17] for environment map generation. Most lighting estimation methods are evaluated on the Laval Indoor HDR Dataset [8]. To this end, the typical evaluation methodology employed is to render a simple virtual object (e.g., one or many spheres with different material properties [30, 27, 24]) with both the ground truth and estimated lighting, and compute image quality metrics to establish the degree of similarity. While color metrics are commonly used (e.g., RGB angular error [13, 24, 4]), no paper has, to our knowledge, specifically evaluated the color robustness of lighting estimation models.


## Strategies for improved color accuracy

We wish to identify methods for improving the color accuracy of lighting estimation models, but that are independent of the specific approach used. In that sense, we avoid approach-specific mechanisms, for example: a specific type of convolutions can only be applicable to CNN-based approaches, and would not work on transformer-based architectures.

Therefore, we seek to adjust the behavior of a lighting estimation model  $\mathcal{L}(I) = L$ , which produces an estimated lighting  $L$  from a single input image  $I$ , without modifying  $\mathcal{L}$  itself. We also assume the estimated lighting  $L$  and its ground truth target  $L^*$  can both be expressed as an image (for example, in spherical coordinates through equirectangular projection), although the following could easily be adapted to other representations (e.g., spherical harmonics or spherical gaussians can readily be converted to spherical images).

### Adding a color-specific loss

The first strategy is to add a color-specific loss, to enforce that the predicted lighting  $L$  yields a color as close as possible to



**Figure 2.** Through chromatic adaptation, we augment the training data of the lighting estimation network  $\mathcal{L}$  to present it with a wider variety of illuminant during training. The illuminant name is shown below each image.

the ground truth target  $L^*$ . We denote

$$\ell_{\text{ang}} = \frac{1}{4\pi} \sum_{i=1}^N (1 - \cos(\tilde{L}_i, \tilde{L}_i^*)) d\omega_i, \quad (1)$$

where  $N$  is the total number of pixels in  $L$ ,  $i$  is a pixel index in  $L$ , and  $d\omega_i$  the solid angle subtended by that pixel, each  $\tilde{L}_i = \frac{(R_i, G_i, B_i)}{R_i + G_i + B_i}$  is the chromaticity-normalized RGB vector of pixel  $i$  (similarly for  $\tilde{L}_i^*$ ). This loss penalizes large angular differences between the predicted lighting  $L$  and the target  $L^*$ . We name this strategy ANGLLOSS.

### Chromatic data augmentation

Another strategy is to rely on data augmentation and present the network with a wider variety of illuminant color during training. For this, we apply chromatic adaptation to map the target lighting  $L^*$  to a different standard illuminant, which are chosen uniformly at random from: (no change), A, B, C, D50, D55, D65, D75, E, F2, F7 and F11. Fig. 2 shows an example of such data augmentation strategy. Note that, while typical white balance algorithms in real-world ISPs may never create such extreme illuminants, we employ this strategy strictly as a way to provide a greater variability in the training data [20]. We name this strategy AUGMENT.

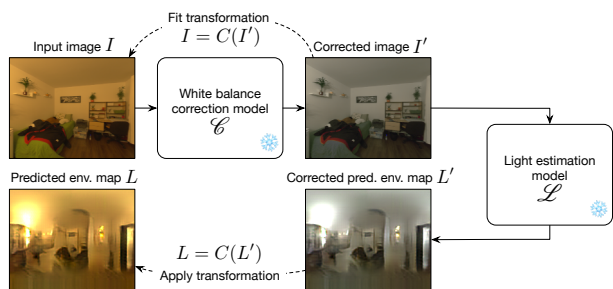
### Adjusting for white balance

A last strategy is to leverage a white balance editing network [1]  $\mathcal{C}$ , which has been trained to convert an input image  $I$  to a neutral color version  $I'$ , i.e.,  $I' = \mathcal{C}(I)$ . We explore its use first at test time only, then during training as well.

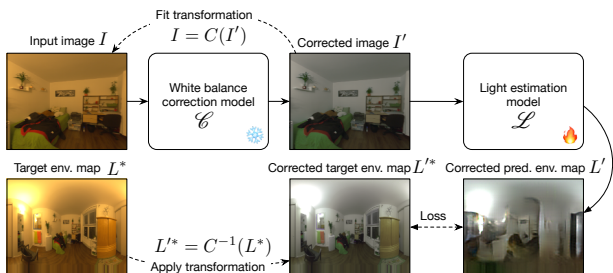
**At test time** First, the input image  $I$  is color-corrected to obtain  $I'$ , which is then fed to the lighting estimation model  $\mathcal{L}$ . The predicted lighting  $L' = \mathcal{L}(I')$  must be corrected back to the original white balance setting of the input by applying the inverse transformation  $\mathcal{C}^{-1}$ . In practice,  $\mathcal{C}$  is a feed-forward neural network, so we cannot assume it to be invertible. We obtain an approximation  $C \approx \mathcal{C}^{-1}$  by fitting a simple linear operator (e.g., a  $3 \times 3$  color transformation matrix) such that  $C(I') = I$ , via a least-squares fit. The resulting  $C$  can then be applied to  $L'$  to obtain  $L = C(L')$ . Compactly, this strategy amounts to

$$L = C(L') = C(\mathcal{L}(\mathcal{C}(I))), \quad (2)$$

We dub this strategy WBTEST and fig. 3-(a) illustrates this process graphically.



(a) Exploiting a white balance network  $\mathcal{C}$  at test time.



(b) Exploiting a white balance network  $\mathcal{C}$  during training.

**Figure 3.** Using a white balance network  $\mathcal{C}$  for improving color robustness. We experiment with employing it at (a) test time only, where all networks ( $\mathcal{C}$  and  $\mathcal{L}$  have frozen weights); and (b) during the training of the lighting estimation network  $\mathcal{L}$  as well (only  $\mathcal{C}$  is frozen). We posit that removing the burden of reasoning about color from the lighting estimation network  $\mathcal{L}$  will make its task easier and thus achieve greater performance.

**During training** A similar strategy can also be applied during training. Here, the idea is to train the lighting estimator model on white-balanced images  $I' = \mathcal{C}(I)$ , with the corresponding ground truth (target) lighting also color-corrected,  $\mathcal{C}(L^*)$ . To ensure both the input image and target lighting are color-corrected using the same transformation  $T$ , the same strategy as in eq. (2) is applied to obtain  $T$ . Then, the corrected target environment map is obtained by  $L'^* = T^{-1}(L^*)$ . This is illustrated graphically in fig. 3. At inference time, eq. (2) is also applied. This focuses training solely on the estimating lighting instead of color. Intuitively, this should result in an easier training since the lighting estimation network  $\mathcal{L}$  no longer needs to adapt to the color of the input image. We name this strategy WBTRAIN.

## Dataset

While there exists a variety of HDR lighting datasets [8, 16, 12], none are particularly appropriate to evaluate the color accuracy of lighting estimation methods as they either focus on dynamic range [8], lighting directions [16], or do not contain full 360° lighting [12]. Therefore, we captured a novel dataset of HDR 360° panoramas in a variety of indoor scenes along with in-camera JPG images under a variety of different white balance settings. This dataset, specifically targeted to the task of evaluating color robustness, accurately replicates a real-world scenario that would be hard to reproduce using data augmentation because we capture images with the hardware ISP on-board the camera.

## Capture protocol

For a given scene, we employ the following capture protocol. First, a Ricoh Theta Z1 camera is placed on a tripod in (approx-



**Figure 4.** Capturing scenes under different white balance settings. For two different indoor scenes captured in HDR (top, tone-mapped with [19] for display), we also capture a series of JPG images under different white balance settings (bottom), to acquire a rich dataset of realistic color renderings. A color checker is placed in each scene for accurate color calibration.

imately) the center of the scene. A Macbeth color chart is also installed on a small tripod and placed approximately 1.5m in front of the camera, in an area of mostly uniform lighting.

As in [2], a set of 11 exposures are captured in exposure bracketing mode, with the camera set in RAW. The resulting files are subsequently merged into a single EXR file using [5]. After HDR capture, a photo is captured with the Theta Z1 camera in all white balance settings supported by the camera, namely: auto, daylight, shade, cloudy, incandescent 1, incandescent 2, daylight fluorescent, neutral white fluorescent, cool white fluorescent, and warm white fluorescent. This is, of course, not equivalent to physically changing the properties of the illuminants, as was done in the LSMI dataset [12] by turning on/off individual light sources. Unfortunately, the HDR environment maps were not captured so full lighting is not available. In our case, we approximate the change of illuminant in a global fashion through the white balance setting onboard the camera ISP. The color checker (present in all frames) is used to adjust the white balance of the HDR panorama to each of the white balance settings.

## Dataset overview

Example images for two scenes from our dataset are shown in fig. 4. The variety of white balance settings simulates vastly different illuminants. Fig. 5 shows a preview of additional scenes in our dataset. As can be observed in the figure, a wide variety of indoor lighting conditions are present in the dataset, either coming from different scenes (e.g., hallways, classrooms, meeting rooms, kitchens, corridors, basements, parking lots, etc.), or from the same scene but with different lighting conditions.

In all, a total of 2,233 images over 78 different scenes are captured. For each panorama, we extract three perspective images at 90° field of view, at 120° intervals, to be used as input images to the lighting estimation networks, at test time only. The dataset is available publicly at the project webpage.



Figure 5. Example scenes (tonemapped for display) in the test dataset.

## Evaluation

We now proceed to evaluate the proposed strategies on our new test dataset. Note that the dataset is used only at test time: none of its images are used during training.

### Baseline lighting estimation method

We define a simple BASELINE lighting estimation method on which to test potential color adaptation strategies.

**Network architecture** We draw inspiration from [21, 27] and adopt the conditional GAN architecture of pix2pixHD [25]. This network is composed of a CNN-based generator and discriminator. The generator follows an encoder-decoder setup in which the input image is downsampled using a series of strided convolutional layers followed by a bottleneck of residual blocks, the output of which is then upsampled to the original resolution using a series of transpose convolutions. The multi-scale discriminator [25] contains two separate networks: the first operates at the original image resolution and receives a concatenation of the input image and the generator prediction as input, whereas the second operates at half the scale and takes a (average-pooled) downsampled version of the aforementioned concatenation as input.

The network takes as input a single (square) LDR image  $I \in \mathbb{R}^{256 \times 256 \times 3}$ , and directly outputs a full 360° HDR panorama  $L$  (equirectangular projection) of the same spatial resolution. To help the network deal with high dynamic range targets, the network outputs are passed through a (differentiable) inverse tonemapping function  $f_{iTM}$  prior to computing the losses,  $f_{iTM}(x) = x^{2.2}$ .

**Training details** The Laval Indoor HDR Dataset [8] is used for training the BASELINE model. To generate the input image from the panorama, we extract 10 rectified crops from each HDR panoramas in the training set. As in [27], the resulting images are converted to LDR by re-exposing to make the median intensity




Figure 6. Distribution of errors on our test dataset for the (a)  $\Delta E$  and (b) RGB angular error metrics for the different color adaptation strategies.

equal to 0.45, clipping to 1, and applying a  $\gamma = 1/2.2$  tonemapping. The same exposure factor is subsequently applied to the target HDR panorama to ensure consistency.

In addition, a subset of the Zillow Indoor Dataset [3] is used for training, which contains a total of 67 448 LDR indoor panoramas of 1575 unfurnished residences. This dataset contains LDR data, so no high dynamic range lighting values are available. However, this dataset is useful to provide additional visual examples of indoor scenes so the network can learn to generate more realistic outputs. We use a subset of the same size as the Laval dataset.

**Loss functions** A combination of GAN, feature matching, perceptual losses are employed [25]. The same default weights as in [25] are used in training. Since the Zillow dataset does not contain HDR, we ensure we do not penalize the network for generating HDR values by clipping its output  $L$  to the  $[0, 1]$  interval before computing the losses. That way, any produced HDR value will not be considered in the losses. Since  $L$  is of equirectangular projection, the feature matching loss is spatially weighted by the solid angle at each pixel. The discriminator employs minibatch standard deviation [11] to minimize mode collapse.

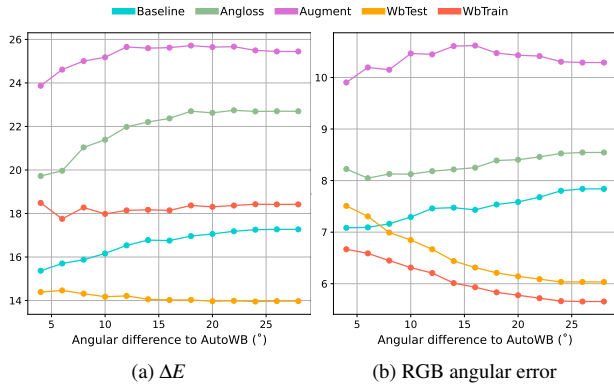
Since we are interested in studying the impact of color on the task of virtual object insertion for augmented reality, we also employ a render loss. To compute this loss, we render a virtual scene with the predicted lighting and compare it with the same virtual scene rendered with the ground truth. Since the scene is static, we pre-compute a light transport matrix  $T$ :

$$\ell_{\text{render}} = \|TL - TL^*\|_1. \quad (3)$$


Here,  $T$  is the transport matrix for a lambertian scene (without interreflections) composed of a flat ground plane with 9 spheres arranged in a grid fashion, observed from above as in [27].

### Evaluation methodology

We retrain several versions of the baseline network by employing the proposed strategies, and evaluate the trained networks on all the images of our test dataset. We reiterate that the images in the test dataset do not overlap the training set: both were captured with different cameras. To measure performance and since we are interested in the use of lighting estimation for virtual object insertion, we first render a virtual scene with the estimated lighting  $L$ , and compare it with a rendering of the same virtual scene with the ground truth lighting  $L^*$ . The virtual scene follows a setup similar to [27], featuring 9 grid-aligned diffuse spheres above a plane shadow catcher, viewed from a bottom-facing camera. Then, the  $\Delta E$  [28] and RGB angular error [10] metrics are computed as they are color-specific.



**Figure 7.** Error metric as a function of the angular distance to the reference automatic white balance (“AutoWB”) setting onboard the camera in the test dataset.



**Figure 8.** Distribution of errors on our test dataset for Weber et al. [27], StyleLight [24], and DiffusionLight [17]. The figure shows the (a)  $\Delta E$  and (b) RGB angular error metrics and for the WBTEST strategy.

## Experimental results

We begin by presenting the distribution of errors obtained by applying each of the strategies on our test dataset. The results are shown in fig. 6. Compared to BASELINE, we first observe that several strategies actually make things *worse* in terms of color. Indeed, ANGLOSS and AUGMENT both result in a worsening of the overall color metrics. In contrast, WBTEST and WBTRAIN improve the color accuracy of lighting estimation. We observe that the best results across all metrics are obtained with WBTEST, which indicates that including a white balance correction step at test time only works better than doing so both at training and test time. This strategy is promising: it is both simpler (as it requires no training) and achieves an improvement in performance that is greater than any of the other strategies tried. We also observe that the worse-case results (shown by the top ticks in the graphs) is significantly improved using WBTEST.

In addition to the statistics computed over the entire test set reported in fig. 6, we also wish to evaluate how well each strategy work on images when their white balance setting differs from the default auto white balance setting (AWB). For this, recall that for each scene in the test set, we captured images under different white balance settings, including AWB. In fig. 7, we thus plot each metric as a function of the RGB angular difference (in degrees) between each test scene and its corresponding AWB image. Here, observe how all metrics progressively become worse as the angular difference increases for BASELINE, ANGLOSS and AUGMENT. In contrast, more stable performance is observed for both WBTEST and WBTRAIN, demonstrating once more that explicitly including

white balance correction indeed makes light estimation more robust to color errors. As before, WBTEST, which applies the correction only at test time, is the most stable of all strategies. Qualitative examples are shown in fig. 9 to visually demonstrate the impact of each strategy.


## Other lighting estimation baselines

It was established that the WBTEST strategy resulted in improved color accuracy when estimating lighting, when evaluated on our BASELINE model. Here, we evaluate whether these findings can be corroborated with other lighting estimation models. Specifically, we experiment with Weber et al. [27], StyleLight [24] and DiffusionLight [17]. As observed in fig. 8, we observe that WBTEST improves color robustness quite significantly, especially visible in the angular error metric. This is corroborated visually in fig. 10, which illustrates qualitative examples applying WBTEST to both techniques. While the difference is less noticeable for StyleLight and DiffusionLight, the color balance can still be improved.

## Discussion

This paper examines the color robustness of lighting estimation algorithms. Rather than proposing a new method, our focus is on developing simple yet effective strategies to enhance the color accuracy of *existing* algorithms. To this end, we introduce a new dataset of indoor HDR lighting environments and simulate varying color conditions by applying different on-camera white balance presets. We establish a baseline lighting estimation model and use it to evaluate several color adaptation strategies: modifying the loss function (ANGLOSS), applying data augmentation (AUGMENT), and incorporating a white balance network either at test time (WBTEST) or during training (WBTRAIN). Our experiments demonstrate that WBTEST offers a particularly robust way to improve a pre-trained lighting estimation network’s resilience to variations in illuminant color. Importantly, this approach requires no retraining of the original model. We validate its generality by applying it to two additional lighting estimation methods from the recent literature.

While the paper presents practical strategies and a valuable dataset for exploring color robustness, we also acknowledge some limitations that can inform future work. First, simulating different illuminants through camera white balance settings is not equivalent to capturing scenes illuminated by physically different light sources (e.g., as in [12]). In some cases, the chosen white balance settings significantly deviate from the scene’s actual illumination, resulting in unrealistic color shifts that would not typically occur with automatic white balance in real-world scenarios. Although our dataset includes such cases intentionally as a way to stress-test models under extreme conditions, future work could improve realism by capturing scenes under varied, physically accurate illuminants. Second, our baseline model relies on a GAN-based architecture, which introduces its own set of challenges. GANs are well known for training instability and issues such as mode collapse [15], phenomena we also observed during our experiments. While we employed techniques to mitigate these issues (e.g., [11]), these problems are specific to GANs and might be avoided altogether by exploring alternative architectures—such as diffusion models—which offer greater stability and are increasingly popular in generative modeling tasks.



**Figure 9.** Qualitative results comparing the estimated HDR environment maps with the various strategies tested. We show the input image (left), the predicted environment maps for each strategy (first row), and a virtual scene rendered with each predicted environment map (bottom row). "GT" (right) is the ground truth.



**Figure 10.** Qualitative results comparing the estimated HDR environment maps when applying the WBTEST strategy to the techniques from Weber et al. [27], StyleLight [24], and DiffusionLight [17]. Original (top) and color-corrected (bottom) predictions are shown.

## References

- [1] M. Afifi and M. S. Brown. Deep white-balance editing. In *Conf. Comput. Vis. Pattern Recog.*, 2020.
- [2] C. Bolduc, J. Giroux, M. Hébert, C. Demers, and J.-F. Lalonde. Beyond the pixel: a photometrically calibrated HDR dataset for luminance and color temperature prediction. In *Int. Conf. Comput. Vis.*, 2023.
- [3] S. Cruz, W. Hutchcroft, Y. Li, N. Khosravan, I. Boyadzhiev, and S. B. Kang. Zillow indoor dataset: Annotated floor plans with 360° panoramas and 3D room layouts. In *Conf. Comput. Vis. Pattern Recog.*, 2021.
- [4] M. R. K. Dastjerdi, J. Eisenmann, Y. Hold-Geoffroy, and J.-F. Lalonde. EverLight: Indoor-outdoor editable HDR lighting estimation. In *Int. Conf. Comput. Vis.*, 2023.
- [5] P. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *ACM SIGGRAPH*. 1997.
- [6] F. Einabadi, J.-Y. Guillemot, and A. Hilton. Deep neural models for illumination estimation and relighting: A survey. *Comput. Graph. Forum*, 40(6):315–331, 2021.
- [7] M.-A. Gardner, Y. Hold-Geoffroy, K. Sunkavalli, C. Gagné, and J.-F. Lalonde. Deep parametric indoor lighting estimation. In *Int. Conf. Comput. Vis.*, 2019.
- [8] M.-A. Gardner, K. Sunkavalli, E. Yumer, X. Shen, E. Gambaretto, C. Gagné, and J.-F. Lalonde. Learning to predict indoor illumination from a single image. *ACM Trans. Graph.*, 9(4), 2017.
- [9] M. Garon, K. Sunkavalli, S. Hadap, N. Carr, and J.-F. Lalonde. Fast spatially-varying indoor lighting estimation. In *Conf. Comput. Vis. Pattern Recog.*, 2019.
- [10] A. Gijssenij, T. Gevers, and M. P. Lucassen. Perceptual analysis of distance measures for color constancy algorithms. *J. Opt. Soc. Am. A*, 26(10):2243–2256, Oct 2009.
- [11] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *Int. Conf. Learn. Represent.*, 2017.
- [12] D. Kim, J. Kim, S. Nam, D. Lee, Y. Lee, N. Kang, H.-E. Lee, B. Yoo, J.-J. Han, and S. J. Kim. Large scale multi-illuminant (LSMI) dataset for developing white balance algorithm under mixed illumination. In *Int. Conf. Comput. Vis.*, 2021.
- [13] C. LeGendre, W.-C. Ma, G. Fyffe, J. Flynn, L. Charbonnel, J. Busch, and P. Debevec. DeepLight: Learning illumination for unconstrained mobile mixed reality. In *Conf. Comput. Vis. Pattern Recog.*, 2019.
- [14] Z. Li, M. Shafiei, R. Ramamoorthi, K. Sunkavalli, and M. Chandraker. Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and svbrdf from a single image. In *Conf. Comput. Vis. Pattern Recog.*, 2020.
- [15] M. Lucic, K. Kurach, M. Michalski, S. Gelly, and O. Bousquet. Are GANs created equal? a large-scale study. In *Adv. Neural Inform. Process. Syst.*, 2018.
- [16] L. Murmann, M. Gharbi, M. Aittala, and F. Durand. A dataset of multi-illumination images in the wild. In *Int. Conf. Comput. Vis.*, 2019.
- [17] P. Phongthawee, W. Chinchuthakun, N. Sinsunthithet, V. Jampani, A. Raj, P. Khungurn, and S. Suwajanakorn. Diffusionlight: Light probes for free by painting a chrome ball. In *Conf. Comput. Vis. Pattern Recog.*, 2024.
- [18] E. Reinhard, M. Adhikhmin, B. Gooch, and P. Shirley. Color transfer between images. *IEEE Comp. graph. and appl.*, 21(5):34–41, 2001.
- [19] E. Reinhard and K. Devlin. Dynamic range reduction inspired by photoreceptor physiology. *IEEE Trans. Vis. Comput. Graph.*, 11(1):13–24, 2005.
- [20] C. Shorten and T. M. Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of big data*, 6(1):1–48, 2019.
- [21] G. Somanath and D. Kurz. HDR environment map estimation for real-time augmented reality. In *Conf. Comput. Vis. Pattern Recog.*, 2021.
- [22] S. Song and T. Funkhouser. Neural illumination: Lighting prediction for indoor environments. In *Conf. Comput. Vis. Pattern Recog.*, 2019.
- [23] P. P. Srinivasan, B. Mildenhall, M. Tancik, J. T. Barron, R. Tucker, and N. Snavely. Lighthouse: Predicting lighting volumes for spatially-coherent illumination. In *Conf. Comput. Vis. Pattern Recog.*, 2020.
- [24] G. Wang, Y. Yang, C. C. Loy, and Z. Liu. StyleLight: HDR panorama generation for lighting estimation and editing. In *Eur. Conf. Comput. Vis.*, 2022.
- [25] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Conf. Comput. Vis. Pattern Recog.*, 2018.
- [26] Z. Wang, J. Phillion, S. Fidler, and J. Kautz. Learning indoor inverse rendering with 3D spatially-varying lighting. In *Int. Conf. Comput. Vis.*, 2021.
- [27] H. Weber, M. Garon, and J.-F. Lalonde. Editable indoor lighting estimation. In *Eur. Conf. Comput. Vis.*, 2022.
- [28] G. Wyszecki and W. S. Stiles. *Color science: concepts and methods, quantitative data and formulae*, volume 40. John Wiley & Sons, 2000.
- [29] F. Zhan, Y. Yu, R. Wu, C. Zhang, S. Lu, L. Shao, F. Ma, and X. Xie. GMLight: Lighting estimation via geometric distribution approximation. *IEEE Trans. Image Process.*, 2022.
- [30] F. Zhan, C. Zhang, Y. Yu, Y. Chang, S. Lu, F. Ma, and X. Xie. EM-Light: Lighting estimation via spherical distribution approximation. In *Assoc. Adv. of Art. Int.*, 2021.
- [31] R. Zhu, Z. Li, J. Matai, F. Porikli, and M. Chandraker. Irisformer: Dense vision transformers for single-image inverse rendering in indoor scenes. In *Conf. Comput. Vis. Pattern Recog.*, 2022.