

Relightful Harmonization: Lighting-aware Portrait Background Replacement

Mengwei Ren^{12*} Wei Xiong² Jae Shin Yoon² Zhixin Shu²
 Jianming Zhang² HyunJoon Jung² Guido Gerig¹ He Zhang²
¹New York University ²Adobe



Figure 1. Relightful Harmonization on four real-world images. Each set shows a direct composition (*upper left*) of the foreground subject onto a new background (*lower left*), and our harmonized result (*right*) that accounts for both **lighting** and **color**.

Abstract

Portrait harmonization aims to composite a subject into a new background, adjusting its lighting and color to ensure harmony with the background scene. Existing harmonization techniques often only focus on adjusting the global color and brightness of the foreground and ignore crucial illumination cues from the background such as apparent lighting direction, leading to unrealistic compositions. We introduce Relightful Harmonization, a lighting-aware diffusion model designed to seamlessly harmonize sophisticated lighting effect for the foreground portrait using any background image. Our approach unfolds in three stages. First, we introduce a lighting representation module that allows our diffusion model to encode lighting information from target image background. Second, we introduce an alignment network that aligns lighting features learned from image background with lighting features learned from panorama

environment maps, which is a complete representation for scene illumination. Last, to further boost the photorealism of the proposed method, we introduce a novel data simulation pipeline that generates synthetic training pairs from a diverse range of natural images, which are used to refine the model. Our method outperforms existing benchmarks in visual fidelity and lighting coherence, showing superior generalization in real-world testing scenarios, highlighting its versatility and practicality.

1. Introduction

Portrait harmonization [60, 64] stands as a crucial element in both photography and creative image editing, seeking to seamlessly composite a subject into a new background while maintaining realism and aesthetic uniformity in terms of lighting and color. The process initiates with the segmentation of the subject from its original image, followed by the composition into a new background. To enhance visual consistency, the subsequent step entails meticulous ad-

*Work done during an internship at Adobe.

justments to the foreground, aligning it with the new background—considering factors such as color, brightness, saturation, and lighting conditions. The manual effort could be labor-intensive, particularly when dealing with intricate lighting scenarios in portraiture.

There are two principal sets of methods that could automatically adjust the foreground to match the background for human portraits: (1) image harmonization techniques, and (2) portrait relighting methods. Harmonization-based methods [5, 7, 8, 14–16, 23, 27, 30, 31, 57, 62, 68, 69] aim to match the color statistics of the foreground with those of the background, by adjusting the foreground color tone, contrast, and illumination. Yet, they often overlook the lighting characteristics and leave the foreground illumination effects unchanged, such as the lighting direction and shadows, potentially resulting in an unnatural appearance when the background has distinct lighting conditions. For instance, compositing a person photographed under a top-down light into a sunset scene might make the composite non-realistic to the human eye. On the other hand, recent work on portrait relighting [34, 39, 43, 56, 65, 70, 73, 77, 80] are designed for matching the lighting of the subject towards a new environment by using the paired training data acquired with the light stage system [10]. Nevertheless, current relighting methods often require dynamic range (HDR) panorama environment maps [43] during training and inference, which are not always feasible to acquire, especially in casual photography settings. In this work, we explore the possibility of generating realistic and lighting-aware composition images in a straightforward harmonization set up. Given a foreground image (with its corresponding alpha mask) and an arbitrary background image, we propose a unified and end-to-end framework that encompasses both *color* and *lighting* harmonization. We approach the task through a conditional generative framework, leveraging a pretrained diffusion model [3, 8], and develop a three-stage training pipeline.

In the first phase, we conduct **Lighting-aware Diffusion Training** to integrate explicit lighting conditioning into a pretrained diffusion model. This involves a lighting representation learning module that derives lighting conditions from a selected background image. The resulting lighting information is then integrated into the diffusion UNet backbone to guide the generative process. The training is performed on a pairwise relighting-specialized light stage dataset to effectively capture the lighting dynamics.

Given the challenge of accurately inferring environmental lighting from a single background image, which is inherently an ill-posed problem, we employ paired environment maps to augment the physical plausibility of our background-derived lighting representation. This is achieved through a second stage of **Lighting Representation Alignment**, designed to align the lighting representa-

tion extracted from background images with that learned from their corresponding panorama environment maps.

Finally, we perform **Finetuning for Photorealism** on an expanded dataset to improve the photorealism of the harmonization. We propose a novel data synthesis pipeline using our initially trained model to create additional data from natural images. Notably, once trained, our pipeline *does not* rely on any external environment maps, which greatly empowers the proposed framework for flexible background replacement and portrait harmonization.

Our contributions are threefold. (1) We enable the lighting effects to be encoded in a pretrained image-conditioned latent diffusion model by incorporating a spatial lighting feature extraction and conditioning module to the diffusion backbone. The background-extracted lighting representation is further aligned with the feature extracted from panorama environment map to ensure better physical plausibility. (2) We use our model as a data augmentor and propose a novel data simulation pipeline to synthesize training pairs from natural images. The model is then refined with the enlarged dataset to further boost the photorealism of the results. (3) Compared with existing harmonization and relighting methods, our pipeline demonstrates improvements of the harmonized results in both lighting coherence visual fidelity, providing a versatile solution for real-world portrait harmonization in a variety of settings.

2. Related Works

Image Harmonization aims to rectify color, contrast, and style differences between foreground and background to ensure natural and consistent composition. In deep learning, this task is approached as an end-to-end image-to-image translation problem [5, 7, 8, 14–16, 23, 31, 57, 58, 62, 68, 81], where the network is trained to predict a harmonized image from the input composite. Pixel-aligned datasets are created by altering foreground color in real images with pre-designed [7] or learned [42] augmentations. Yet, existing harmonization methods primarily focus on global color adjustment, overlooking the subtle but important discrepancies between foreground and background lighting, i.e., direction, intensity, and shadow. This can lead to a harmonized image that, despite matching colors, still appears unnatural due to mismatched lighting conditions. Therefore, we postulate that enhancing the lighting-awareness of harmonization models is a vital yet underexplored area for natural and realistic composition.

Portrait Relighting Recent advancements in portrait relighting has been driven by deep learning methods [34, 35, 39, 43, 45, 56, 65, 70, 73, 77, 80]. They leverage supervised training with the paired training data acquired with the light stage system [10]. These methods require a target HDR environment map as the external input source, and

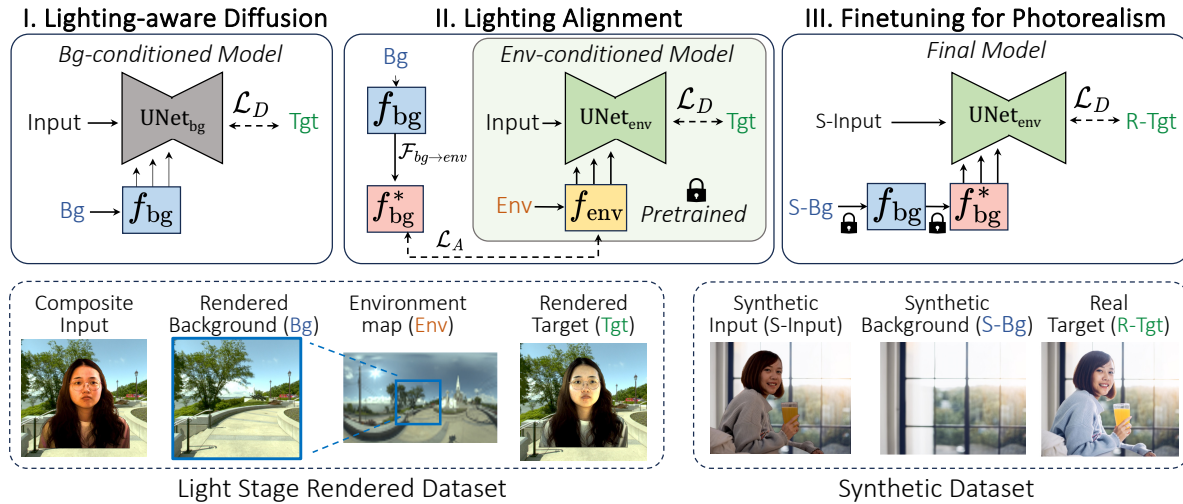


Figure 2. The Pipeline of Relightful Harmonization. In Stage I, a lighting representation module is integrated into the diffusion model, conditioning the generation on lighting information encoded from the background image, trained with a light stage dataset for relighting (*lower left*). Stage II aligns lighting features derived from the background with the environment map for enhanced physical accuracy. Finally, Stage III refines the model on a real image dataset (*lower right*) obtained via a novel data simulation pipeline.

typically involve the intermediate prediction of surface normals, albedo, and/or a set of diffuse and specular maps with ground truth supervision. However, reliance on HDR maps for background replacement and harmonization tasks poses significant limitations on their applicability in everyday scenarios where HDR maps cannot be easily captured alongside [43]. In our framework, we do not require any external input sources except a single target background image. Furthermore, many current relighting systems employ multistage frameworks and/or heavily rely on external packages [12]. The accuracy and performance of these systems are consequently contingent on the precision of each individual stage, making the overall process complicated and prone to errors propagated through these intermediate steps. Additionally, the datasets commonly employed for training are rendered from limited light stage illumination acquisition, which means the target images utilized during the training phase are not captured in real-world conditions but are rendered composites. Therefore, its generalization on unseen images, as well as the photorealism when applied to arbitrary background replacement tasks remain unclear.

Diffusion Models [11, 20, 52, 53] have significantly advanced the image and video synthesis quality [19, 21, 25, 26, 46, 48]. Image-conditioned diffusion models [49, 63, 74] typically take an image as additional input to perform an image-conditioned generation such as image enhancement [24, 26, 29, 47], harmonization [6, 32, 33] and translation [28, 41, 51, 63, 74, 79], typically trained on task-specific pairwise data. Recently, the application of pre-trained text-to-image diffusion models [40, 48, 50] has been extended to image editing tasks [1–4, 9, 13, 17, 18, 22, 36–38, 54, 55, 59, 61, 66, 78]. These models leverage text-image correlations to perform context modifications in image editing, such as ‘adding a sunset’ with Instruct-

Pix2Pix [3], which is loosely connected to our lighting-aware set up. However, text-based editing does not incorporate spatial information, thus lacking finegrained control to the model. Instead, we propose to use a spatial lighting representation as the new ‘instruction’ that guide the diffusion model to perform lighting-aware editing.

3. Method

We aim to develop a conditional diffusion model that processes a composite image (along with its alpha mask) as the input, conditioned on the target background, and produce *color* and *lighting* harmonized output. To do so, we develop a three-stage training strategy presented in Fig. 2.

Stage I: Lighting-aware Diffusion Training: We build our model on a pretrained diffusion model [3] and enable its lighting awareness by attaching a lighting representation learning branch to encode lighting information from the background image, which is then injected into the UNet backbone as illustrated in Fig. 2-I. The training is conducted with relighting-specialized light stage rendered dataset, as shown on the bottom left of Fig. 2-I.

Stage II: Lighting Representation Alignment: As one of our goals is to enable lighting-aware portrait harmonization *without* relying on environment maps during inference, we propose a representation alignment step (Fig. 2-II) to adapt the lighting representation extracted from a background image towards the learned representation from its environment map. We assume the aligned representation is more robust and physically plausible.

Stage III: Finetuning for Photorealism: In the third stage, we finetune our model using high-quality pixel-aligned training pairs from natural images, where these paired datasets are generated via a novel data synthesis pipeline (Fig. 2-III) using the stage 2 model as a data augementer.

3.1. Lighting-aware Diffusion Training

Our lighting-aware diffusion model learns to generate a harmonized image given the composite input, conditioned on a lighting representation extracted from the target background image. Adhering to established practice for training relighting models [43, 56, 76], we assume access to the light stage rendered training dataset. A training tuple includes the input image (alongside its alpha mask), the target background, the target environment map, as well as the target image. An example of training data is shown on the bottom left of Fig. 2.

Formally, we represent a rendered image sample as x_i^a , indicating a portrait image of subject i illuminated under the lighting condition a . The corresponding environment map is denoted as z_i^a , and the background image y_i^a for x_i^a is generated by projecting the HDR map with a specified field of view and resolution. Subject masks m_i are obtained using methods described in [71, 72].

Lighting Conditioning: We further modify the diffusion backbone to incorporate explicit lighting conditions. As depicted in Fig. 2-I, a lighting conditioning branch is integrated atop the UNet backbone, injecting a lighting representation f encoded from the target background image by a CNN \mathcal{F} , at multiple scales within the UNet. The conditioning mechanism is designed in a similar fashion as [74] where conditional feature maps are added to the UNet features at respective resolution within the encoder.

With the light stage dataset, we postulate that the lighting representation f can be learned from the pairwise training. Specifically, a training tuple from the same subject i is sampled as (x^a, m, y^b, x^b) . Noise is progressively added to the target image x^b until time step t , resulting in a noisy image x_t^b . The UNet, denoted as \mathcal{U}_{bg} , is conditioned on the background-extracted lighting feature $F_{bg}(y^b)$, and is trained to predict the noise ϵ , with the following objective:

$$\mathcal{L}_D = \mathbb{E}_{x^a, y^b, x_0^b, t, \epsilon} \left[\|\epsilon - \mathcal{U}_{bg}(x_t^b, t, x^a, F_{bg}(y^b))\|_2^2 \right] \quad (1)$$

where $\epsilon \sim \mathcal{N}(0, 1)$. At this stage, we initialize the weights of the UNet backbone from [3], and jointly train *both* UNet and the conditioning branch.

3.2. Lighting Representation Alignment

Given that a background image is a partial projection of the environment map which encapsulates panoramic lighting information (see example on bottom left of Fig. 2), we suspect that the lighting cues learned from an environment map will be inherently more comprehensive than those from f_{bg} , implying that an environment-conditioned model could potentially offer more physically plausible relighting under the same training scenario. This is further empirically verified in our ablation detailed in Sec. 4.5. However, in real-world photography, environment maps are usually not co-acquired which poses practical challenges, limiting the applicability

of environment map dependent models. To circumvent this, we align the lighting representation extracted from a background image with features derived from its ground truth environment map, ultimately enabling effective portrait harmonization with just a single background image.

As shown in Fig. 2 II., we first pretrain an environment map conditioned harmonization model (in green box) to generate a ground truth environment map-derived lighting representation. The model architecture is identical to the background-based model in Fig. 2 I, while substituting the Stage I input condition from the background image to its corresponding environment map. It is trained with the same light stage dataset, under a denoising loss analogous to Eq. 1, where the background lighting feature f_{bg} is replaced with an environment map-derived feature $\mathcal{F}_{env}(z^b)$.

Then, we freeze the environment-conditioned model and introduce an alignment network $\mathcal{F}_{bg \rightarrow env}$ that calibrates the background-extracted lighting representation to align with its environment map equivalent. We formulate such a process as an inverse problem that can be learned with a network $\mathcal{F}_{bg \rightarrow env}$, under a supervised loss. For a training tuple (x^a, m, y^b, z^b, x^b) , the alignment network takes $f_{bg} = \mathcal{F}_{bg}(y^b)$ as input, and maps it with the alignment network to f_{bg}^* . The environment extracted feature $f_{env} = \mathcal{F}_{env}(z^b)$ is utilized as the ground truth, and we use a L_1 objective:

$$f^* = F_{bg \rightarrow env}(F_{bg}(y^b)),$$

$$\mathcal{L}_A = \mathbb{E}_{y^b, z^b} \left[\|F_{env}(z^b) - f_{bg}^*\|_1 \right]. \quad (2)$$

During this phase, we update only $\mathcal{F}_{bg \rightarrow env}$, while freezing the other networks. We assume that this alignment enhances the background-derived lighting representation to more accurately encode the environmental lighting, which is empirically verified where the aligned feature maps reflect more global illumination information (Fig. 7). Once trained, we integrate the aligned feature extraction and conditioning into \mathcal{U}_{env} , formulating our final model in Fig. 2-III.

3.3. Finetuning for Photorealism

The light stage dataset serves as a valuable resource for learning lighting representations, providing physically constrained relighting pairs. However, it is essential to recognize that this dataset is *compositional* by nature. It uses backgrounds projected from environment maps, combined with relit foreground subjects, to create composites that serve as ground truth for the diffusion model. However, these composites differ from real photographs, leading to potential concerns about the photorealism of the model's outputs. Moreover, due to the cost of light stage data acquisition, it restricts the number and diversity of the subjects that can be collected. The diversity of the background images is also bounded by the scale of available environment maps during rendering (a few thousands). These limitations

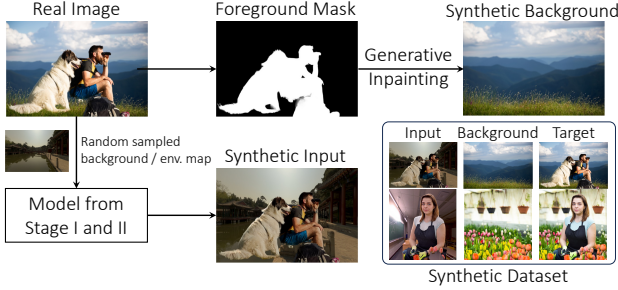


Figure 3. Data synthesis pipeline. Given a real image, and inpaint the subject region, we obtain a synthetic background. The foreground lighting is then altered with our model trained in Stage I/II, to create an input image with distinct lighting. Two example pairs are shown on the lower right.

could affect the model’s generalization ability to real images, and the capacity to produce realistic and varied lighting effects. Therefore, we propose a third stage (Fig. 2 III) that finetunes our final model for improved photorealism.

We introduce a novel data synthesis pipeline that creates pairwise training pairs from natural images, to ensure that the ground truth for finetuning the diffusion model remains *real* images. As depicted in Fig. 3, the process starts with a portrait photograph, from which we extract the foreground mask [71, 72]. We then inpaint the foreground region [67] with text guidance ‘clear background’ to create a clean background image for the real image, which can serve as the condition input for the training. Next, the lighting of the foreground subject(s) is altered by running our trained model from stage I/II with a randomly chosen background image or environment map as the condition. This produces a synthetic input image with distinct foreground lighting and color compared to the target image. Two sets of generated training tuples are displayed at bottom right of Fig. 3. Once we obtain a sufficient number of synthetic data, we combine the original light stage dataset with the synthetic data to refine our model. During this stage, we freeze the lighting representation extraction and conditioning branch and only finetune the UNet backbone to refine the synthesis quality while maintaining the learned lighting plausibility. Once trained, our final model in Fig. 2 III is used to perform portrait harmonization given arbitrary background images, eliminating the need for environment maps.

4. Experiments

4.1. Setup and Metrics

Three testing scenarios are created for evaluation: (1) 500 Light stage rendered test pairs to evaluate the lighting accuracy; (2) 200 natural image test pairs, synthetically created from real images using our data synthesis pipeline, to assess the lighting realism; and most importantly, (3) Real-world portraits combined with arbitrary backgrounds, examining

the model’s generalizability and adaptability in real-life scenarios. For (1) and (2), we also quantify the results with common metrics MSE, SSIM, PSNR and LPIPS [75]. To benchmark, we compare Relightful Harmonization with both established harmonization methods INR [5], PCT [14], Harmonizer [27] and PIH [62], and relighting method TR [43]. We also construct a relighting baseline with a transformer architecture and trained it with light stage data. More details are provided in the appendix. Note that relighting methods are applied only on the light stage test set as they are not applicable without HDR maps.

4.2. Implementation Details

Our model is implemented in PyTorch [44] using 8×80 GB A100 at 512×512 resolutions, with 96 batch size. In stage I, we initialize UNet from the pretrained weights of the InstructPix2Pix [3] checkpoint. In the first and second stage, we use in total 400k training image pairs, rendered from a arbitrary combination of 100 unique light stage subjects, and 2908 HDR environment maps. We also randomly rotate the HDR maps and use various FoVs to increase the diversity of the background. In the third stage, we train the network with additional 200k pairs of images synthesized from natural images. We set learning rate to $5e-5$. More details are provided in the appendix.

4.3. Benchmark Results

In Figs. 4a and 4b, we present visual comparisons of our method with selected benchmarks across on both test sets. Our approach performs better in adjusting both the foreground color and lighting, aligning more closely with the ground truth on the right. Quantitative results in Table 1 further demonstrate the benefit of our methods. Fig. 4c showcases test results on natural images. While harmonization methods adjust for color, the composition still lacks fidelity due to counterintuitive lighting on the composed image. For example, the strong cast shadow on the foreground in the first row, and opposite lighting directions between foreground/background in the second and third rows. Full visual results with all methods are provided in appendix.


We further conducted a user study to verify the visual plausibility of our method on the real world testing set. Given pairwise comparison between our method and each baseline, we asked Amazon Mechanical Turk raters to select the better harmonized image from a given pair of results sample from 70 image. The results on 1750 ratings are collected, and we report the fraction of times that raters preferred our results over the baseline method in Table. 2.

4.4. Arbitrary Portrait Background Replacement


We further test our model on various in-the-wild portraits by compositing them onto arbitrary natural backgrounds.



(a) Comparison on the light stage test set with [43] and [27].



(b) Comparison on the natural image synthetic test set with [27, 62].



(c) Comparison with [62] and [27] on real test images. Our method more effectively harmonizes incoherent foreground lighting and shadow.

Figure 4. Visual comparisons with benchmark methods.

Table 1. Quantitative results on both light stage test set and the natural image test set. The relighting methods including TR [43] and the transformer baseline require HDR maps during the inference and are thus non-applicable on the natural image test set.

Method	Light stage test set				Natural image test set			
	MSE ↓	PSNR ↑	SSIM ↑	LPIPS ↓	MSE ↓	PSNR ↑	SSIM ↑	LPIPS ↓
TR [43]	0.044 (0.057)	15.889 (4.318)	0.757 (0.087)	0.354 (0.092)	N/A	N/A	N/A	N/A
Transformer*	0.026 (0.021)	17.259 (3.715)	0.742 (0.096)	0.337 (0.095)	N/A	N/A	N/A	N/A
INR [5]	0.016 (0.014)	19.147 (3.353)	0.823 (0.081)	0.327 (0.083)	0.009 (0.005)	21.566 (2.943)	0.904 (0.038)	0.113 (0.031)
Harmonizer [27]	0.015 (0.011)	19.304 (2.980)	0.822 (0.077)	0.338 (0.087)	0.010 (0.007)	21.419 (3.506)	0.905 (0.039)	0.108 (0.032)
PCT [14]	0.020 (0.016)	18.339 (3.454)	0.808 (0.093)	0.408 (0.082)	0.014 (0.010)	19.647 (3.279)	0.898 (0.038)	0.147 (0.039)
PIH [62]	0.018 (0.015)	18.865 (3.087)	0.807 (0.087)	0.330 (0.089)	0.010 (0.007)	21.147 (3.097)	0.901 (0.038)	0.112 (0.033)
Ours	0.012 (0.010)	20.527 (3.136)	0.848 (0.076)	0.159 (0.058)	0.005 (0.004)	23.562 (2.830)	0.913 (0.034)	0.097 (0.044)

Table 2. User preference. Each value represents the fraction of times that raters preferred our results than the baseline method.

Method	PIH [62]	INR [5]	PCT [14]	Harmonizer [27]
Preference Rate	0.713	0.702	0.845	0.639


Lighting Plausibility: We start by evaluating the lighting plausibility and visual fidelity by replacing the backgrounds with strong lighting indications, like sunlight. Example results are shown in Fig. 5a. As can be seen, the lighting tone and direction in the harmonized foreground matched the background effectively. We further vertically flip the backgrounds on the right side, and observe a corresponding change in the lighting effects as expected.

Shadow Plausibility: We then examine how our model handles shadows. Fig. 5b depicts the test cases on input subjects with prominent shadows. When compositing and harmonizing them onto background images with more ambient lighting, our model is able to remove the strong shadow and estimate a shadow-free output, while adapting to the background light. Conversely, in Fig. 5c, under backgrounds in daytime scenes with potential overhead sunlight, our model

generates visually plausible self-occlusion shadows.

Creative Background Replacement: Our method allows for creative background replacement for portrait images. In Fig. 5d (right), we create a sequence of background crops from one panorama image with a sliding window, so that the major light source consistently changes from left to right. We observe visually consistent lighting changes on the foreground. Similarly, Fig. 5d (left) shows the harmonization results by placing a subject into different timelapse video frames at different timepoints, resulting in a series of portrait images that effectively mimicked a timelapse effect.

Reference-based Harmonization: Our pipeline also extends to reference-based harmonization, where users can integrate their portrait photos with scenes from a chosen reference portrait image. This is achieved by first inpainting the reference portrait to generate a background image, onto which the desired foreground is composited. Our method ensures that the final output matches the tone and lighting of the original scene. Fig. 5e showcases examples where our results effectively align with the tone and lighting of the reference image, displayed on the top left.




(a) In cases where the target backgrounds offer clear lighting cues, our method generates visually convincing **lighting effects**. Additionally, upon flipping the background, we note consistent and appropriate adjustments to the lighting direction in the output.



(b) Our method effectively **neutralizes pronounced shadows** in the input while accommodating the ambient lighting of the background.



(c) When applied to backgrounds with intense lighting conditions, e.g., with overhead sunlight, our method casts plausible **self-occlusion shadows**.



(d) Our method consistently adjusts lighting when applied to moving backgrounds with **temporally**(left) and **spatially**(right) changing lighting directions.



(e) Our approach allows for reference-based harmonization tasks. This involves removing the subject from the reference image (*upper left*) to create a background (*lower left*) for composition. The harmonized results (*right*) achieve lighting effects closely resembling those in the reference.

Figure 5. Real-world testing results under different scenarios to examine the lighting and shadow effects. For each pair of results in row (a)-(c), we display the composite image (*left*) and the harmonized image (*right*). In row (d), we omit the composition for better visibility. Full visualization is provided in the Appendix.

Table 3. Ablation to verify the effects of lighting conditioning (‘Cond’), alignment (‘Align’), and finetuning (‘Finetune’).

Model	Cond	Align	Finetune	Light stage test set				Natural image test set			
				MSE ↓	PSNR ↑	SSIM ↑	LPIPS ↓	MSE ↓	PSNR ↑	SSIM ↑	LPIPS ↓
0	-	-	-	0.018 (0.015)	18.768 (3.531)	0.815 (0.085)	0.188 (0.070)	0.012 (0.008)	20.158 (3.024)	0.866 (0.045)	0.103 (0.028)
1	Bg	-	-	0.014 (0.012)	19.748 (3.161)	0.835 (0.084)	0.168 (0.063)	0.013 (0.008)	19.787 (2.684)	0.864 (0.042)	0.106 (0.028)
2	Env	-	-	0.009 (0.009)	21.626 (3.162)	0.866 (0.077)	0.148 (0.056)	0.013 (0.008)	19.824 (2.634)	0.863 (0.045)	0.105 (0.028)
3	Bg	✓	-	0.012 (0.009)	20.439 (2.987)	0.842 (0.077)	0.163 (0.059)	0.012 (0.007)	20.006 (2.586)	0.866 (0.044)	0.106 (0.028)
4	Bg	✓	✓	0.012 (0.010)	20.527 (3.136)	0.848 (0.076)	0.159 (0.058)	0.005 (0.004)	23.562 (2.830)	0.913 (0.035)	0.097 (0.044)



Figure 6. Example testing results from our ablation. Model 1 to Model 4 correspond to the configurations in Table 3.




Figure 7. The L_2 norm of learned lighting representations. The aligned background-derived feature on the right matches the panorama much closer, indicating a better lighting representation.

4.5. Ablation

As our proposed method involves multiple stages, we conducted an ablation study to isolate individual and collective effects. We define the base Model#0 as a baseline diffusion model without lighting conditioning. Model#1 and Model#2 incorporate lighting conditioning via background and environment map, respectively. Model#3 further introduces the alignment module and Model#4 is our final model that includes the finetuning. Table 3 illustrates the quantitative performance of each model configuration on both light stage and natural image test sets. We include more visual comparisons in the appendix.

Lighting-conditioning: Upon integrating lighting conditioning into Model#0, we observe notable improvements across all metrics on light stage data with both background conditioned Model#1 and environmental conditioned Model#2. Model#2 further outperforms Model#1, verifying our assumption that an environmental map facilitates a more accurate encoding of lighting.

Lighting Representation Alignment: In Model#3, we introduce the embedding alignment. The improvements from Model#1 to Model#3 validated the benefits of utilizing an additional adaptation step to extract robust lighting representation from a single background image. As shown in Fig. 7, the right column representing the aligned feature norms, shows a visual convergence towards the features extracted from the environment map.

Finetuning for Photorealism: Model#4 is a finetuned version of Model#3 with the newly synthesized data. We observe significant improvements in particular on the natural image test set, verifying our assumption that the finetuning boosts the photorealism on natural images. As can be also seen in examples in Fig. 6, while Model#2 and Model#3 estimate plausible lighting direction, the lighting effects are much less realistic than the finetuned Model#4.

5. Conclusion

We present Relightful Harmonization, a novel lighting-aware diffusion model to blend advanced lighting effects into foreground portraits when compositing onto diverse background images. Limitations exist including a resolution cap of 512x512, potentially affecting facial detail, especially in smaller faces. Also, subtle variations may occur in the subject’s clothing and skin tones. Detailed analyses and additional failure cases are included in the appendix.

6. Acknowledgement

We are grateful for Yannick Hold-Geoffroy who provides the panorama environment maps and set up the rendering scripts. We thank Chaowei Company for the support of light stage dataset. We thank David Futschik for running the testing on Total Relighting. We thank Scott Cohen for the inspiration for our project name.

References

- [1] Yuval Alaluf, Daniel Garibi, Or Patashnik, Hadar Averbuch-Elor, and Daniel Cohen-Or. Cross-image attention for zero-shot appearance transfer. *arXiv preprint arXiv:2311.03335*, 2023. [3](#)
- [2] Manuel Brack, Felix Friedrich, Dominik Hintersdorf, Lukas Struppek, Patrick Schramowski, and Kristian Kersting. Sega: Instructing diffusion using semantic dimensions. *arXiv preprint arXiv:2301.12247*, 2023.
- [3] Tim Brooks, Aleksander Holynski, and Alexei A. Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18392–18402, 2023. [2](#), [3](#), [4](#), [5](#), [1](#)
- [4] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaohu Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. *arXiv preprint arXiv:2304.08465*, 2023. [3](#)
- [5] Jianqi Chen, Yilan Zhang, Zhengxia Zou, Keyan Chen, and Zhenwei Shi. Dense pixel-to-pixel harmonization via continuous image representation. *IEEE Transactions on Circuits and Systems for Video Technology*, pages 1–1, 2023. [2](#), [5](#), [6](#)
- [6] Jianqi Chen, Zhengxia Zou, Yilan Zhang, Keyan Chen, and Zhenwei Shi. Zero-shot image harmonization with generative model prior, 2023. [3](#)
- [7] Wenyan Cong, Jianfu Zhang, Li Niu, Liu Liu, Zhixin Ling, Weiyuan Li, and Liqing Zhang. Dovenet: Deep image harmonization via domain verification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8394–8403, 2020. [2](#)
- [8] Wenyan Cong, Xinhao Tao, Li Niu, Jing Liang, Xuesong Gao, Qihao Sun, and Liqing Zhang. High-resolution image harmonization via collaborative dual transformations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18470–18479, 2022. [2](#)
- [9] Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. Diffedit: Diffusion-based semantic image editing with mask guidance. *arXiv preprint arXiv:2210.11427*, 2022. [3](#)
- [10] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 145–156, 2000. [2](#)
- [11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. [3](#)
- [12] Yao Feng, Haiwen Feng, Michael J Black, and Timo Bolkart. Learning an animatable detailed 3d face model from in-the-wild images. *ACM Transactions on Graphics (ToG)*, 40(4): 1–13, 2021. [3](#)
- [13] Tsu-Jui Fu, Wenze Hu, Xianzhi Du, William Yang Wang, Yinfei Yang, and Zhe Gan. Guiding instruction-based image editing via multimodal large language models. *arXiv preprint arXiv:2309.17102*, 2023. [3](#)
- [14] Julian Jorge Andrade Guerreiro, Mitsuru Nakazawa, and Björn Stenger. Pct-net: Full resolution image harmonization using pixel-wise color transformations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5917–5926, 2023. [2](#), [5](#), [6](#)
- [15] Zonghui Guo, Dongsheng Guo, Haiyong Zheng, Zhaorui Gu, Bing Zheng, and Junyu Dong. Image harmonization with transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 14870–14879, 2021.
- [16] Zonghui Guo, Haiyong Zheng, Yufeng Jiang, Zhaorui Gu, and Bing Zheng. Intrinsic image harmonization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16367–16376, 2021. [2](#)
- [17] Ligong Han, Song Wen, Qi Chen, Zhixing Zhang, Kunpeng Song, Mengwei Ren, Ruijiang Gao, Anastasis Stathopoulos, Xiaoxiao He, Yuxiao Chen, et al. Proxedit: Improving tuning-free real image editing with proximal guidance. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 4291–4301, 2024. [3](#)
- [18] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. [3](#)
- [19] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. [3](#)
- [20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. [3](#)
- [21] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *arXiv:2204.03458*, 2022. [3](#)
- [22] Inbar Huberman-Spiegelglas, Vladimir Kulikov, and Tomer Michaeli. An edit friendly ddpm noise space: Inversion and manipulations. *arXiv preprint arXiv:2304.06140*, 2023. [3](#)
- [23] Yifan Jiang, He Zhang, Jianming Zhang, Yilin Wang, Zhe Lin, Kalyan Sunkavalli, Simon Chen, Sohrab Amirghodsi, Sarah Kong, and Zhangyang Wang. Ssh: A self-supervised framework for image harmonization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4832–4841, 2021. [2](#)
- [24] Bahjat Kavar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. In *Advances in Neural Information Processing Systems*, 2022. [3](#)
- [25] Bahjat Kavar, Roy Ganz, and Michael Elad. Enhancing diffusion-based image synthesis with robust classifier guidance. *arXiv preprint arXiv:2208.08664*, 2022. [3](#)
- [26] Bahjat Kavar, Jiaming Song, Stefano Ermon, and Michael Elad. Jpeg artifact correction using denoising diffusion restoration models. *arXiv preprint arXiv:2209.11888*, 2022. [3](#)
- [27] Zhanghan Ke, Chunyi Sun, Lei Zhu, Ke Xu, and Rynson W.H. Lau. Harmonizer: Learning to perform white-box image and video harmonization. In *European Conference on Computer Vision (ECCV)*, 2022. [2](#), [5](#), [6](#)
- [28] Gihyun Kwon and Jong Chul Ye. Diffusion-based image translation using disentangled style and content representation. *arXiv preprint arXiv:2209.15264*, 2022. [3](#)

- [29] Haoying Li, Yifan Yang, Meng Chang, Shiqi Chen, Huajun Feng, Zhihai Xu, Qi Li, and Yueting Chen. Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing*, 479:47–59, 2022. 3
- [30] Jingtang Liang, Xiaodong Cun, and Chi-Man Pun. Spatial-separated curve rendering network for efficient and high-resolution image harmonization. *arXiv preprint arXiv:2109.05750*, 2021. 2
- [31] Lingxiao Lu, Jiangtong Li, Junyan Cao, Li Niu, and Liqing Zhang. Painterly image harmonization using diffusion model. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 233–241, 2023. 2
- [32] Lingxiao Lu, Jiangtong Li, Junyan Cao, Li Niu, and Liqing Zhang. Painterly image harmonization using diffusion model. In *Proceedings of the 31st ACM International Conference on Multimedia*. ACM, 2023. 3
- [33] Shilin Lu, Yanzhu Liu, and Adams Wai-Kin Kong. Tf-icon: Diffusion-based training-free cross-domain image composition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2294–2305, 2023. 3
- [34] Yiqun Mei, He Zhang, Xuaner Zhang, Jianming Zhang, Zhixin Shu, Yilin Wang, Zijun Wei, Shi Yan, HyunJoon Jung, and Vishal M. Patel. Lightpainter: Interactive portrait relighting with freehand scribble. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 195–205, 2023. 2
- [35] Yiqun Mei, Yu Zeng, He Zhang, Zhixin Shu, Xuaner Zhang, Sai Bi, Jianming Zhang, HyunJoon Jung, and Vishal M. Patel. Holo-relighting: Controllable volumetric portrait relighting from a single image. *arXiv preprint arXiv:2403.09632*, 2024. 2
- [36] Chenlin Meng, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. 3
- [37] Daiki Miyake, Akihiro Iohara, Yu Saito, and Toshiyuki Tanaka. Negative-prompt inversion: Fast image inversion for editing with text-guided diffusion models. *arXiv preprint arXiv:2305.16807*, 2023.
- [38] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. *arXiv preprint arXiv:2211.09794*, 2022. 3
- [39] Thomas Nestmeyer, Jean-François Lalonde, Iain Matthews, and Andreas Lehrmann. Learning physics-guided face relighting under directional light. In *CVPR*, pages 5124–5133, 2020. 2
- [40] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. 3
- [41] Shen Nie, Hanzhong Allan Guo, Cheng Lu, Yuhao Zhou, Chenyu Zheng, and Chongxuan Li. The blessing of randomness: Sde beats ode in general diffusion-based image editing. *arXiv preprint arXiv:2311.01410*, 2023. 3
- [42] Li Niu, Junyan Cao, Wenyan Cong, and Liqing Zhang. Deep image harmonization with learnable augmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7482–7491, 2023. 2
- [43] Rohit Pandey, Sergio Orts Escolano, Chloe Legendre, Christian Haene, Sofien Bouaziz, Christoph Rhemann, Paul Debevec, and Sean Fanello. Total relighting: learning to relight portraits for background replacement. *ACM Transactions on Graphics (TOG)*, 40(4):1–21, 2021. 2, 3, 4, 5, 6
- [44] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 5
- [45] Puntawat Ponglertnapakorn, Nontawat Tritrong, and Supasorn Suwajanakorn. Difareli: Diffusion face relighting. 2023. 2
- [46] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 3
- [47] Mengwei Ren, Mauricio Delbracio, Hossein Talebi, Guido Gerig, and Peyman Milanfar. Multiscale structure guided diffusion for image deblurring. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10721–10733, 2023. 3
- [48] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022. 3
- [49] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022. 3
- [50] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. 3
- [51] Hiroshi Sasaki, Chris G Willcocks, and Toby P Breckon. Unit-ddpm: Unpaired image translation with denoising diffusion probabilistic models. *arXiv preprint arXiv:2104.05358*, 2021. 3
- [52] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. 3
- [53] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. 3
- [54] Yizhi Song, Zhifei Zhang, Zhe Lin, Scott Cohen, Brian Price, Jianming Zhang, Soo Ye Kim, and Daniel Aliaga. Object-stitch: Object compositing with diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18310–18319, 2023. 3

- [55] Yizhi Song, Zhifei Zhang, Zhe Lin, Scott Cohen, Brian Price, Jianming Zhang, Soo Ye Kim, He Zhang, Wei Xiong, and Daniel Aliaga. Imprint: Generative object compositing by learning identity-preserving representation. *arXiv preprint arXiv:2403.10701*, 2024. 3
- [56] Tiancheng Sun, Jonathan T Barron, Yun-Ta Tsai, Zexiang Xu, Xueming Yu, Graham Fyffe, Christoph Rhemann, Jay Busch, Paul Debevec, and Ravi Ramamoorthi. Single image portrait relighting. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019. 2, 4, 1
- [57] Linfeng Tan, Jiangtong Li, Li Niu, and Liqing Zhang. Deep image harmonization in dual color spaces. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 2159–2167, 2023. 2
- [58] Yi-Hsuan Tsai, Xiaohui Shen, Zhe Lin, Kalyan Sunkavalli, Xin Lu, and Ming-Hsuan Yang. Deep image harmonization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3789–3797, 2017. 2
- [59] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1921–1930, 2023. 3
- [60] Jeya Maria Jose Valanarasu, He Zhang, Jianming Zhang, Yilin Wang, Zhe Lin, Jose Echevarria, Yinglan Ma, Zijun Wei, Kalyan Sunkavalli, and Vishal M Patel. Interactive portrait harmonization. *arXiv preprint arXiv:2203.08216*, 2022. 1
- [61] Bram Wallace, Akash Gokul, and Nikhil Naik. Edict: Exact diffusion inversion via coupled transformations. *arXiv preprint arXiv:2211.12446*, 2022. 3
- [62] Ke Wang, Michaël Gharbi, He Zhang, Zhihao Xia, and Eli Shechtman. Semi-supervised parametric real-world image harmonization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5927–5936, 2023. 2, 5, 6
- [63] Tengfei Wang, Ting Zhang, Bo Zhang, Hao Ouyang, Dong Chen, Qifeng Chen, and Fang Wen. Pretraining is all you need for image-to-image translation. *arXiv preprint arXiv:2205.12952*, 2022. 3
- [64] Yijiang Wang, Yuqi Li, Chong Wang, and Xulun Ye. Harmonized portrait-background image composition. In *Computer Graphics Forum*, page e14921. Wiley Online Library, 2023. 1
- [65] Zhibo Wang, Xin Yu, Ming Lu, Quan Wang, Chen Qian, and Feng Xu. Single image portrait relighting via explicit multiple reflectance channel modeling. *ACM TOG*, 39(6):1–13, 2020. 2
- [66] Chen Henry Wu and Fernando De la Torre. A latent space of stochastic diffusion models for zero-shot image editing and guidance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7378–7387, 2023. 3
- [67] Shaoan Xie, Zhifei Zhang, Zhe Lin, Tobias Hinz, and Kun Zhang. Smartbrush: Text and shape guided object inpainting with diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22428–22437, 2023. 5
- [68] Yazhou Xing, Yu Li, Xintao Wang, Ye Zhu, and Qifeng Chen. Composite photograph harmonization with complete background cues. In *Proceedings of the 30th ACM international conference on multimedia*, pages 2296–2304, 2022. 2
- [69] Su Xue, Aseem Agarwala, Julie Dorsey, and Holly Rushmeier. Understanding and improving the realism of image composites. *ACM Transactions on graphics (TOG)*, 31(4):1–10, 2012. 2
- [70] Yu-Ying Yeh, Koki Nagano, Sameh Khamis, Jan Kautz, Ming-Yu Liu, and Ting-Chun Wang. Learning to relight portrait images via a virtual light stage and synthetic-to-real adaptation. *ACM TOG*, 2022. 2
- [71] Qihang Yu, Jianming Zhang, He Zhang, Yilin Wang, Zhe Lin, Ning Xu, Yutong Bai, and Alan Yuille. Mask guided matting via progressive refinement network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1154–1163, 2021. 4, 5
- [72] He Zhang, Jianming Zhang, Federico Perazzi, Zhe Lin, and Vishal M. Patel. Deep image compositing. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 365–374, 2021. 4, 5
- [73] Longwen Zhang, Qixuan Zhang, Minye Wu, Jingyi Yu, and Lan Xu. Neural video portrait relighting in real-time via consistency modeling. In *ICCV*, pages 802–812, 2021. 2
- [74] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 3, 4, 1
- [75] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 5
- [76] Xuaner Zhang, Jonathan T Barron, Yun-Ta Tsai, Rohit Pandey, Xiuming Zhang, Ren Ng, and David E Jacobs. Portrait shadow manipulation. *ACM Transactions on Graphics (TOG)*, 39(4):78–1, 2020. 4
- [77] Xiuming Zhang, Sean Fanello, Yun-Ta Tsai, Tiancheng Sun, Tianfan Xue, Rohit Pandey, Sergio Orts-Escolano, Philip Davidson, Christoph Rhemann, Paul Debevec, et al. Neural light transport for relighting and view synthesis. *ACM TOG*, 40(1):1–17, 2021. 2
- [78] Zhixing Zhang, Ligong Han, Arnab Ghosh, Dimitris N Metaxas, and Jian Ren. Sine: Single image editing with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6027–6037, 2023. 3
- [79] Min Zhao, Fan Bao, Chongxuan Li, and Jun Zhu. Egsde: Unpaired image-to-image translation via energy-guided stochastic differential equations. *Advances in Neural Information Processing Systems*, 35:3609–3623, 2022. 3
- [80] Hao Zhou, Sunil Hadap, Kalyan Sunkavalli, and David W Jacobs. Deep single-image portrait relighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7194–7202, 2019. 2
- [81] Jun-Yan Zhu, Philipp Krahenbuhl, Eli Shechtman, and Alexei A Efros. Learning a discriminative model for the

perception of realism in composite images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3943–3951, 2015. [2](#)

Relightful Harmonization: Lighting-aware Portrait Background Replacement

Supplementary Material

7. Additional Results

We note that the all testing input portrait images shown in our paper are sampled from Unsplash or Adobe Stock.

Comparison with benchmarks To supplement Fig. 4c, we present additional visual comparison with benchmarks on the real world data in Fig. 8, 9 and 10. To supplement Fig. 4a and Fig. 4b, we show full benchmark comparison on the light stage test dataset in Fig. 11, and on the natural image test set in Fig. 12.

Ablation We present additional visual comparison among our ablation models on the natural image test and light stage test set in Fig. 14 and Fig. 15 respectively. The ablation comparison on the real test set is shown in Fig. 13 as a supplement to Fig. 6. In Fig. 16, we include additional visualization of the feature norms to illustrate the affects of the alignment module to supplement Fig. 7.

Real world testing results Fig. 17 shows reference based harmonization example as in Fig. 5e. Fig. 18 shows harmonization results when we flip the background image as in Fig 5a. Fig. 19 shows the results under spatially and temporally changing lighting as in Fig. 5d.

8. Additional Implementation Details

8.1. Network Architecture

Lighting-conditioned diffusion is built on the Instruct-Pix2Pix [3] backbone. The core rationale for selecting the pretrained InstructPix2Pix model [3] as the foundation, rather than the stable diffusion model, is on its capability to incorporate an additional input image channel. Therefore, at the beginning of our training, the input and output image will be identical (i.e., no editing on the input image), and it will gradually incorporate the lighting conditioning from the extra lighting representation. We use a dummy editing prompt ‘portrait’ during our training and inference.

The lighting conditioning branch architecture follows ControlNet [74], where an encoder structure identical to the diffusion UNet backbone is applied and the intermediate feature maps are added to the UNet encoder at respective resolutions. The lighting representation is extracted from a 4-layer CNN. We train our model with the input resolution of 512×512 (for both input image and the background image), and the lighting representation is a tensor with shape $64 \times 64 \times 320$. We empirically found that training with a higher resolution (e.g., 768×768) led to better identity preservation, but performed worse in terms of the relighting. We speculate that this is related to the stable diffusion pretraining, which is on 512×512 resolution.

The Alignment Network is an encoder-decoder architecture built with Residual blocks. The encoder is composed of three sequential residual blocks. Each of these blocks is coupled with a subsequent downsampling layer. The decoder is symmetrical to the encoder, with three residual blocks, and each of them followed by an upsampling layer. The input and output dimensions of the alignment network are consistent, maintaining a shape of $64 \times 64 \times 320$.

Ablation Models Specifics Model#0 is a baseline diffusion model without lighting conditioning and its implementation follows InstructPix2Pix [3] with the text prompt fixed as ‘Portrait’. Model#1 takes the background image as the conditional input, which is resized to 512×512 . Model#2 shares the same architecture as Model#1 but replaces the conditional input to the LDR environment map. Model#3 introduces the alignment module after the conditional branch from Model#1. The Unet backbone from Model#2 is used as diagramed in Fig. 2. Model#4 fine-tunes on Model#3 with the synthetic data.

8.2. Transformer relighting model

To train a relighting baseline on our light stage dataset, we built a transformer based encoder-decoder network. The network input is a concatenated input image, foreground mask, and the parsing mask, which is divided into patches of 4×4 . A hierarchical Transformer encoder is applied to obtain multi-level features at $\{\frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \frac{1}{32}\}$ of the original resolution. A decoder with transpose convolution is then followed to get the final result with the same resolution as the input. The target LDR environment map is concatenated at the bottleneck latent space in a similar manner as [56].

9. Failure case and analysis

We illustrate several example failure cases in Fig. 20. In our training approach, since we do not impose constraints on the subject’s identity, there are instances where the model struggles to retain identity-specific details. For instance, as shown in Fig. 20a, the color of the subject’s clothing is inaccurately altered during the color harmonization process. Similarly, in Fig. 20b, there is a notable change in hair color (*middle*). Furthermore, in scenarios where the input skin tone is not clearly indicated (*right*), our model occasionally produces ambiguous results in skin tone modification. Additionally, our method does not incorporate intermediate steps like albedo estimation, which can be crucial in handling complex lighting conditions. As a result, in inputs with pronounced cast shadows, our model sometimes fails to eliminate these shadows effectively.




Figure 8. Example comparison results on the real world test set to supplement Fig. 4c in the main paper.



Figure 9. Example comparison results on the real world test set to supplement Fig. 4c in the main paper.



Figure 10. Example comparison results on the real world test set to supplement Fig. 4c in the main paper.



Figure 11. Example comparison results on the light stage test set to supplement Fig. 4a in the main paper. The environment map is shown at the bottom of the ground truth image.




Figure 12. Example comparison results on the natural image test set to supplement Fig. 4b in the main paper.




Figure 13. Example testing results from our ablation on the real image test set. Model 0 to Model 4 correspond to the configurations in Table 3. Our final model (Model 4) presents the best visual quality while maintaining plausible lighting effects.




Figure 14. Example testing results from our ablation on the natural image test set. Model 0 to Model 4 correspond to the configurations in Table 3. Our final model (Model 4) presents the best visual quality while maintaining plausible lighting effects.




Figure 15. Example testing results from our ablation on the light stage test set. Model 0 to Model 4 correspond to the configurations in Table 3. Our final model (Model 4) presents the best visual quality while maintaining plausible lighting effects.

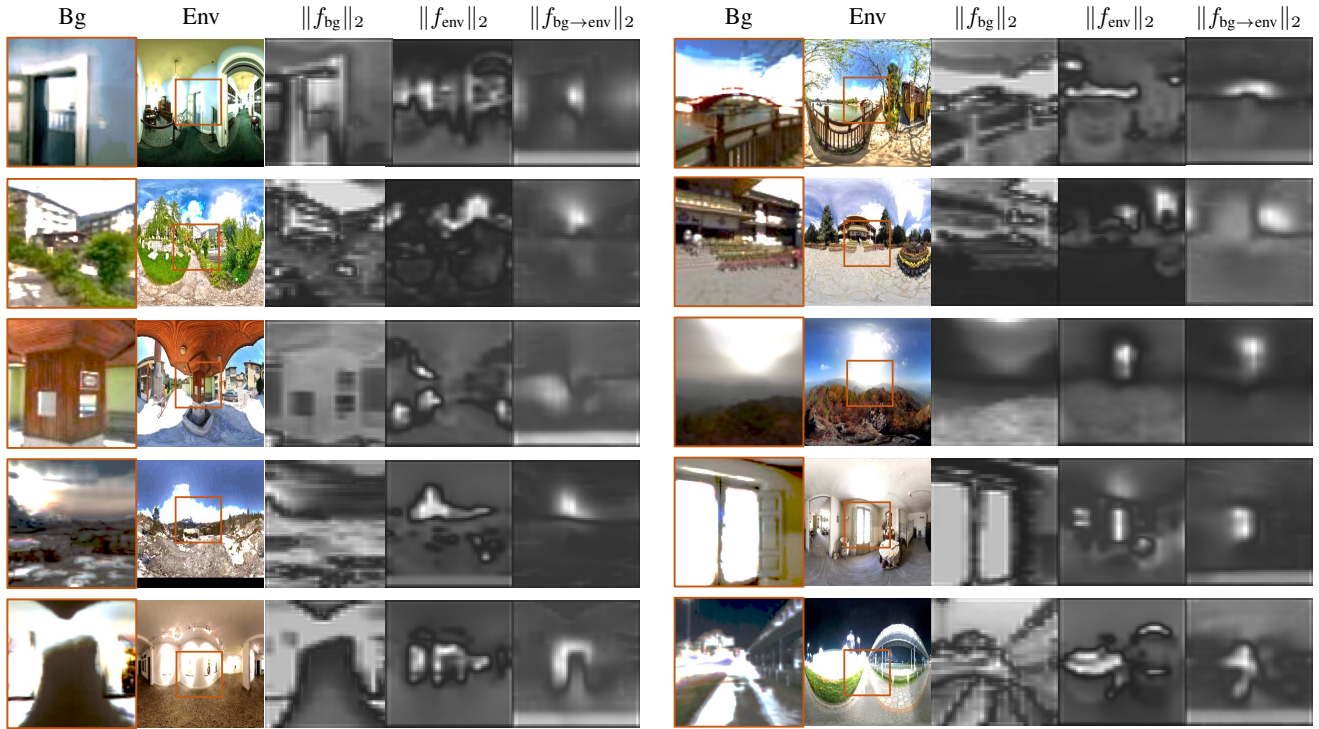


Figure 16. The L_2 norm of learned lighting representations to supplement Fig. 7. The aligned background-derived feature on the right matches the panorama much closer, indicating a better lighting representation.



Figure 17. Visual results on the reference-based harmonization application to supplement Fig. 5c. It allows user images to be blended into scenes from real portraits. This involves removing the subject from the reference image (*upper left*) to create a background (*lower left*) for composition. The harmonized results (*right*) achieve lighting effects closely resembling those in the reference.



Figure 18. Harmonization results when flipping the background, to supplement Fig. 5a.




(a) We create spatially changing lighting conditions by cropping background images (*top*) from a panoramic image. Our model produces visually coherent lighting changes on different portrait images.




(b) We obtain temporally changing lighting conditions by taking multiple screenshots (*top*) from a timelapse video (<https://www.youtube.com/watch?v=CSfri4U9w28>). Our model produces visually reasonable harmonization results.


Figure 19. Harmonization results under the background images where lighting conditions are changing spatially (a) or temporally (b).



(a) In some examples, our model modified the color of the subject clothes due to its harmonization nature.



(b) In some examples, our model may not fully preserve the subject identity, such as the hair color and the skin tone, especially when the input skin tone is ambiguous (right two examples).



(c) In portraits with strong casted shadows, our model may fail to completely remove them.

Figure 20. Failure cases