

Lux Post Facto: Learning Portrait Performance Relighting with Conditional Video Diffusion and a Hybrid Dataset

Yiqun Mei^{1,2} Mingming He¹ Li Ma¹ Julien Philip¹ Wenqi Xian¹ David M George¹
 Xueming Yu¹ Gabriel Dedic¹ Ahmet Levent Taşel¹ Ning Yu¹ Vishal M. Patel² Paul Debevec¹
¹Netflix Eyeline Studios ²Johns Hopkins University




Figure 1. *Lux Post Facto* offers portrait relighting as a simple post-production process. Users can edit the lighting of portrait images (first row) and videos (second row) with high fidelity using any HDR map. Our method is temporally stable and highly photorealistic.

Abstract

Video portrait relighting remains challenging because the results need to be both photorealistic and temporally stable. This typically requires a strong model design that can capture complex facial reflections as well as intensive training on a high-quality paired video dataset, such as dynamic one-light-at-a-time (OLAT). In this work, we introduce *Lux Post Facto*, a novel portrait video relighting method that produces both photorealistic and temporally consistent lighting effects. From the model side, we design a new conditional video diffusion model built upon state-of-the-art pre-trained video diffusion model, alongside a new lighting injection mechanism to enable precise control. This way we leverage strong spatial and temporal generative capability to generate plausible solutions to the ill-posed

relighting problem. Our technique uses a hybrid dataset consisting of static expression OLAT data and in-the-wild portrait performance videos to jointly learn relighting and temporal modeling. This avoids the need to acquire paired video data in different lighting conditions. Our extensive experiments show that our model produces state-of-the-art results both in terms of photorealism and temporal consistency.

1. Introduction

Lighting is a crucial element in visual storytelling, shaping the scene, adding depth, and creating emotional impact. Because of this, lighting receives considerable attention in the filming process, marked by professional expertise, exper-

sive equipment, and coordinated teamwork – resources beyond the reach of novice users. In this work, we endeavor to design a video relighting method to empower everyday content creators to create creative and compelling lighting in portrait videos as a simple post-production process.

Video relighting requires an accurate simulation of light transport through complex materials across both time and space. Image-based relighting [11] allows a portrait to be relit when the subject has been recorded from a dense array of lighting directions, such as a one-light-at-a-time (OLAT) reflectance field. Using a high-speed camera and a synchronized LED stage, Wenger et al. [71] record OLATs at movie frame rates, allowing for cinematic relighting of a facial performance. However, the technique relies on complex equipment and cannot be generalized to new subjects.

More recent works have attempted to transfer facial reflectance information from one or more subjects to another, for example by tracking an OLAT quotient image [41] onto a new flat-lit performance. This has been generalized to a deep learning context, where OLAT data from a set of training subjects is used to relight a novel subject [25, 30, 40, 61, 79]. However, these models can have issues with temporal consistency, a lack of photorealism, or require difficult-to-obtain OLAT video data, or some combination of the three.

In this work, we propose *Lux Post Facto*, a new video relighting model for relighting arbitrary portrait videos realistically and with temporal consistency. We formulate video relighting as a conditional generation process leveraging a state-of-the-art video diffusion model [3], fine-tuned on a hybrid dataset with new lighting injection and training strategies.

While video diffusion models [3, 74] have the power to generate temporally consistent videos with high-quality lighting from text guidance, they lack the ability to achieve fine-grain lighting control. While we would like to specify the lighting as a high dynamic range (HDR) lighting environment, conventional image conditioning methods [75, 85] like CLIP [47], struggle to represent the HDR pixel values and spatial details in HDR maps. And methods based on pixel-aligned control maps [27, 80] require explicit proxy geometry, prone to estimation errors.

We solve this problem by encoding the original HDR map as a set of “lighting embeddings”, where each embedding represents a single directional light source. Collectively, the embeddings produce a complete lighting environment representation. We pass this representation to the diffusion model through cross-attention to achieve precise lighting control.

Our model is practical to train since it requires only a limited number of *static* OLAT image datasets and a larger set of in-the-wild portrait videos. The static OLATs provide relighting supervision for individual frames, and the in-the-

wild videos train the model to produce temporally stable performances under a wild variety of unknown lighting conditions. Our novel training strategy learns from this combination of datasets to achieve temporally consistent portrait video relighting. We show that our method outperforms current single-image and video relighting models. To summarize, our technical contributions include:

1. A novel video diffusion model for portrait relighting, capable of generating high-fidelity lighting effects on arbitrary portrait videos with state-of-the-art performance.
2. A novel lighting control module that improves the encoding of lighting information and enables precise lighting control.
3. A new hybrid dataset and associated training approach which allows training our video relighting models using static OLAT images and in-the-wild video data in a unified framework.

2. Related Work

Portrait relighting has been a key research area for years. The pioneering light stage work [9] proposes to capture the reflectance field of human faces using OLAT, composing them to novel lighting using the linearity of light transport. Ensuing work [2, 11, 17, 36, 37, 71] use time-multiplexed lighting to capture live-action actors at high frame rates. Moving subjects are then relit using image-based or neural rendering [51, 63]. Alternative multi-view setups (e.g., Wang et al. [68]) have been proposed to create relighting models. However, all these methods require per-subject capture, limiting their generalization.

To bypass the need for light stage data, techniques like quotient images [41, 56], intrinsic decomposition [1, 23, 28, 29, 53, 54], or style transfer [57, 58] have been proposed. Recently, data-driven methods use models trained on multi-illumination datasets to generalize. CNN-based methods have shown their effectiveness on image relighting, achieving novel lighting effects [25, 40, 43, 76], shadow manipulation [13, 16, 23, 32, 42, 61, 67, 70, 77, 83, 86], and interactive lighting editing [34]. GAN-based methods enhance portrait relighting using ratio images [22], or explicit 3D representations [35, 39, 48, 62]. While effective for single images, these methods struggle to maintain temporal consistency when applied to videos.

Given the recent advances in image diffusion models [10, 50], known for their quality and generality, some works apply them to portrait and scene relighting [19, 21, 27, 44, 46, 78], and related tasks like portrait harmonization [49, 81]. The relighting models [19, 27, 44, 78] fine-tune pre-trained latent diffusion models [50] on a dataset with lighting annotations. The harmonization methods [49, 81] propose to adjust foreground lighting to match backgrounds. They both achieve photorealistic results. However, none of these image-based models are temporally stable.

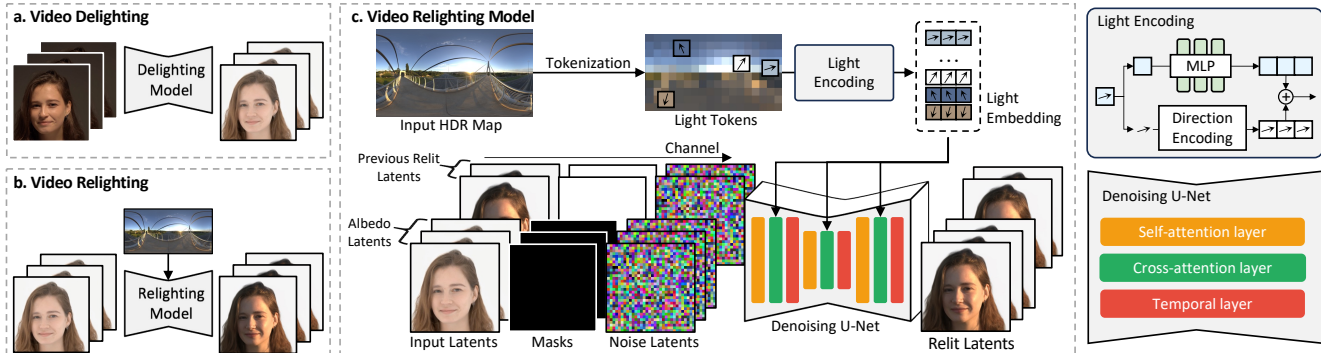


Figure 2. **Overview and model design of *Lux Post Facto*.** To relight an input video, a delighting model predicts an albedo video (a) which is then relit by a relighting model (b). Both models share the same architecture (c) based on stable video diffusion [3] (SVD). We condition the SVD on the input video by concatenating input latents to the Gaussian noise. To support autoregressive prediction for long sequence, we replace the first T frames with previous predictions, indicated with a binary mask concatenated to the input. The output lighting is controlled by an HDR map, converted to a light embedding fed to the U-Net through cross-attention layers. The VAE that encodes and decodes the latents is omitted for clarity.

Video-based Relighting. Video relighting focuses on consistent lighting for dynamic sequences. Recent methods often trade spatial quality for temporal consistency. For example, extensions of image relighting models to video, such as [19, 25], typically apply temporal smoothing to the relit results, leading to blurry shading and averaged details. Some other methods assume simplified lighting conditions, such as diffuse [31], or white and spherical harmonic illumination [6]. NVPR [79], a general face relighting model, integrates a temporal consistency loss to achieve stable portrait relighting under dynamic lighting, but loses some facial details. Some 3D-based methods [30, 73] leverage the advanced 3DGS representations [24] for high-quality facial performance relighting. They require per-subject multi-view capture for both training and inference. These methods rely on scarce dynamic OLAT datasets captured with high-speed cameras. We introduce a hybrid dataset to address the scarcity of high-quality paired video datasets.

Video Diffusion Models. Video diffusion models [3, 4, 18, 60] form the foundation of our video relighting model. These models synthesize realistic and temporally coherent videos from text input. To improve the generation quality, some works propose to use diffusion noise prior [8, 14] or cascading models [20, 66]. To make the model controllable, recent works condition it on control signals [84] such as human poses [7, 65], input images [55], or depth maps [12]. To precisely control the lighting of our video diffusion models, we propose a new light conditioning mechanism.

Preliminary: Our video relighting model is built upon video diffusion models like Stable Video Diffusion (SVD) [3]. These models involve a forward pass that progressively injects Gaussian noise into video sequences and a reverse process with a denoising U-Net learned to predict the noise to reconstruct the original videos. The network

is trained by minimizing the mean squared error between noise predictions and ground truth. To ensure temporal coherence, the models incorporate specialized layers like 3D convolutions and temporal attention, operating in a latent space via a variational autoencoder (VAE) for efficiency.

3. Method

We present *Lux Post Facto*, a video relighting method that can relight in-the-wild portrait videos. Following recent work on portrait relighting [35, 40], our method is composed of two stages: a video delighting pass and a video relighting pass. Both stages are shown in Fig. 2.a and 2.b. First, the delighting model takes a portrait video as input and predicts a shading-free albedo¹ video. Then, the relighting model takes the predicted albedo video as input and an HDR map as lighting condition and generates a relit video under the target lighting condition. The two models are built upon the same video diffusion architecture but differ in input and conditioning.

We train both these models on a hybrid dataset stemming from static OLAT images with precise lighting information and in-the-wild portrait videos with diverse motions. To best use this data, we train our model with the auxiliary task of reference-based appearance copy, improving temporal consistency.

In Sec. 3.1, we first describe both the relighting and delighting model design and our long-sequence inference strategy. Then in Sec. 3.2, we detail our hybrid dataset and how we adapt them to the auxiliary task of appearance copy.

3.1. Model Design

At the core of *Lux Post Facto* is the idea to leverage the strong generative priors from pretrained video diffusion

¹Following [34, 35, 40], we refer to flat-lit images as albedo approximation.

models, adapting them for the relighting task. We thus build our relighting method upon a state-of-the-art open-source video diffusion model, *i.e.*, SVD [3], modified to become a conditional generator. To achieve our relighting goal, we propose a two-stage approach following existing works [35, 40], which first delights the input video to generate its delit albedo video and then relights it. Since both modules share the same backbone, we first introduce the relighting model, then describe how the delighting model differs. Note that both models are trained independently and do not share their weights.

3.1.1. Relighting Model

The model architecture of the relighting module is shown in Fig. 2.c. From a delit albedo video clip - predicted by the delighting module - and potentially previous relit frames, the relighting model predicts a relit video. To adapt the text-guided SVD model to our relighting task, we provide both spatio-temporal conditioning and lighting control.

Spatio-Temporal Conditioning. To support spatio-temporal conditioning on the input frames, we follow previous works [5, 15, 49] and add additional input channels to the first convolution layer of the denoising U-Net. This modification conditions the denoising process on the input video.

For temporal consistency over long sequences during testing, we adopt an iterative mechanism, detailed in Sec. 3.1.3. As such, the model must distinguish the original input frames (*i.e.*, albedo) from previous predictions (*i.e.*, relit frames). To achieve this, we introduce additional binary masks as input to the denoising U-Net. These binary masks M_t indicate the previously generated frames, *i.e.*, they have all 1’s for the frames from the previous window and 0’s otherwise. In summary, the input to the denoising U-Net is a concatenation of input latents, binary masks, and noise latents over time.

Lighting Control. To control the target lighting at test time, we condition the denoising U-Net on an HDR map. Compared to other representations for lighting control, such as background images [49, 81] or text prompts [81] recently used in conditional diffusion models, HDR maps capture a broader range of lighting scenarios and offer more precise control.

To provide the HDR map to the denoising U-Net, we first tokenize it. Each directional “light token” is computed by summing the intensities over a small local area in the HDR map. We then embed these “light tokens” into high dimensional “light embeddings” using a multilayer perceptron (MLP), and concatenate them with positional encodings representing each light’s average direction. The embeddings are then concatenated to form the light embedding \mathcal{L}_e . \mathcal{L}_e is transmitted to the denoising U-Net through cross-attention layers similar to text embeddings. This condition-

ing design is key to achieving high-quality, controllable relighting results and enabling precise control as shown by our experiments.

3.1.2. Delighting Model

The delighting diffusion model uses the same architecture as the relighting model. It differs only in its input and conditioning signals. Instead of taking a delit video and previous relit frames, it takes the original video and previous delit frames as input. The delighting model does not use lighting control, since its target output is always a flat-lit albedo video.

3.1.3. Long Sequence Prediction

Video diffusion models can only predict video frames with a fixed length (*e.g.* L). To support long-video inference, we adopt an iterative scheme following recent works [15, 64]. We generate subsequent frames based on previous predictions. Specifically, for our diffusion models, trained with frame length $L = 30$, we replace the first $T = 4$ frames in the input with previous predictions and update their masks, letting the model predict the subsequent $T - L = 26$ frames. During training, we randomly sample $T \in [0, 4]$ and replace the input frames with ground truth.

3.2. Learning Relighting on a Hybrid Dataset

Training both delighting and relighting models would typically require a paired video dataset consisting of lit videos V_l and their corresponding albedo videos V_a , as well as HDR map conditions E_l . However, collecting such a video dataset is difficult and expensive in practice.

Instead, we use a more accessible hybrid video dataset created from static OLAT images and in-the-wild videos. We first introduce how we obtain training video datasets from both data sources and then present our training scheme, designed to best use the information from each.

3.2.1. Hybrid Video Data Creation

We build our hybrid dataset with two types of videos: videos with diverse illuminations but little motions (denoted as \mathcal{D}_l), synthesized from static OLAT images, and videos with diverse motions but unknown lighting (denoted as \mathcal{D}_m), obtained from in-the-wild videos. A visualization of the raw source data and its corresponding processed training data is shown in Fig. 3.

Lighting-Rich Dataset \mathcal{D}_l . \mathcal{D}_l contains triples of paired video and HDR maps $\{V_l, V_a, E_l\}$. To construct this dataset, we start from static OLAT images paired with matched flat-lit images. We capture a group of subjects with diverse static expressions and poses from multiple views. We use image-based lighting [9, 51] to obtain various relit versions of each expression/pose from their OLATs. We generate the lit images using various HDR maps E_l . For

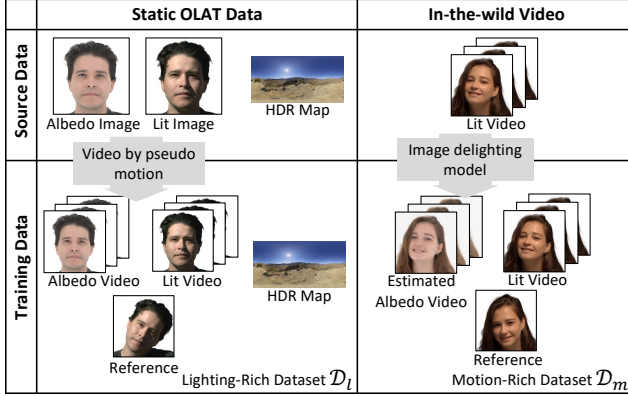


Figure 3. **Hybrid video dataset creation.** We add synthetic camera motion to static OLAT images, and apply the image delighting model to in-the-wild videos, creating our hybrid training data.

their corresponding albedo images, we directly use the flat-lit images [34, 35, 40]. Finally, to lift the data to the video domain we simulate camera motions (e.g., zooming, panning) on the images. This lighting-rich dataset could be used to learn portrait relighting by itself. However, the motions being limited to 2D camera movements, we demonstrate in our experiments that \mathcal{D}_l is not enough to generalize to real in-the-wild motions.

Motion-Rich Dataset \mathcal{D}_m . For \mathcal{D}_m , we use an in-the-wild video dataset [72] containing 15,000 high-quality talking heads with diverse motion patterns. As shown in Fig. 3, to create paired training samples, we first train an *image* delighting model on \mathcal{D}_l and then use this model, we predict pseudo albedos for each video, frame by frame. These albedo videos are thus not temporally consistent. Details about the image delighting model can be found in the supplement. This way, we have not only pseudo-ground truth for the delighting task but also pseudo-input with real ground truth for the relighting task, similar to [27]. Thus, we obtain a dataset of in-the-wild video pairs $\mathcal{D}_m = \{V_l, V_a\}$, but without corresponding HDR map E_l .

3.2.2. Training with Hybrid Data

With a fully-labelled dataset of triples $\{V_l, V_a, E_l\}$, we could supervise both delighting and relighting models to learn the delighting mapping $\mathbf{D} : (V_l) \mapsto V_a$ and relighting mapping $\mathbf{R}_{\text{hdr}} : (V_a | \mathcal{L}_e) \mapsto V_l$, respectively. Where \mathcal{L}_e is the embedding of the HDR map E_l .

However, given our hybrid dataset, we face two issues. First, we lack the HDR map E_l to condition the relighting mapping \mathbf{R}_{hdr} on the motion-rich dataset \mathcal{D}_m . Second, for the delighting mapping \mathbf{D} , \mathcal{D}_m only provides pseudo albedo, which is not temporally consistent enough to be used as supervision. We thus propose to train both models on the auxiliary task of reference-based appearance copy. This lets the networks learn from both datasets, combining accurate lighting control and improved temporal stability.

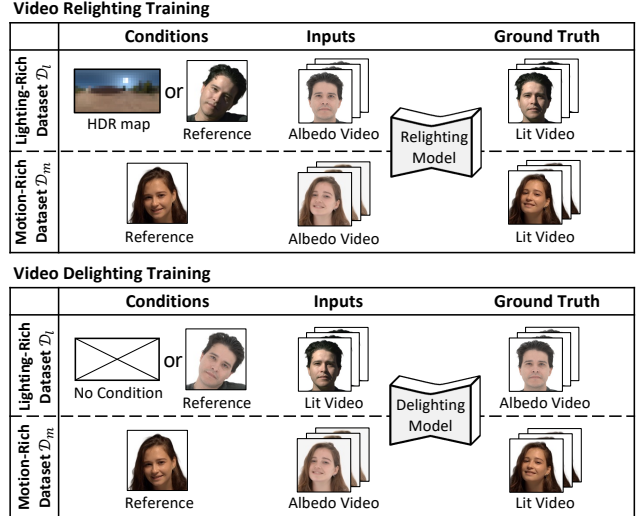


Figure 4. **Training with hybrid data.** To train both the relighting and delighting models, we use the hybrid dataset. We train the models on two tasks simultaneously: HDR-based condition (or no condition) on the OLAT data (i.e. lighting-rich dataset \mathcal{D}_l) and reference-based appearance copy on both datasets (i.e. motion-rich dataset \mathcal{D}_m and lighting-rich dataset \mathcal{D}_l).

Reference-based Appearance Copy. To leverage the motion-rich dataset \mathcal{D}_m , we propose to train our network on an auxiliary task. We choose the task of reference-based appearance copy [57, 59]. It presents similarities with relighting whilst not requiring known HDR maps. Instead of using an HDR map to control lighting, for this task, we use a reference frame F_l^i taken from a lit video V_l . The network is then tasked to “copy” the appearance from the lit reference frame F_l^i to the frames of the corresponding delit clip V_a . F_l^i is randomly selected from V_l and may not overlap with the training subsequences. In practice, to condition the model on the reference frame F_l^i , we encode it as an embedding \mathcal{L}_{ref} using a CNN encoder: $\mathcal{L}_{\text{ref}} = \mathcal{E}(F_l^i)$, thus leading to a reference-based appearance copy mapping $\mathbf{R}_{\text{ref}} : (V_a | \mathcal{L}_{\text{ref}}) \mapsto V_l$. To support both HDR- and reference-based lighting control, we adapt the light encoder for different inputs using masks. More details are included in the supplement. Next, we introduce how we train both video relighting and delighting using this auxiliary task.

Video Relighting. Our relighting model can now leverage both datasets by training HDR-based relighting and reference-based appearance copy simultaneously, as shown in Fig. 4 (Video Relighting). Formally, $\mathbf{R} = \mathbf{R}_{\text{hdr}} \cup \mathbf{R}_{\text{ref}} : (V_a | \mathcal{L}_e \vee \mathcal{L}_{\text{ref}}) \mapsto V_l$.

For videos from the motion-rich dataset \mathcal{D}_m , we only use the reference-based conditioning, while we randomly condition on either or both for videos from the lighting-rich dataset \mathcal{D}_l . Training with the motion-rich dataset \mathcal{D}_m through reference-based appearance copying effectively im-



Figure 5. **Comparison against video relighting methods on in-the-wild portrait videos.** For each sequence we show three input frames, with the target HDR map and reference image rendered with the same HDR map and OLAT data, both shown as inset. Our method produces more faithful lighting effects and is robust to facial expression change (first column) and head motion (last two columns).

proves temporal consistency.

Video Delighting. For the relighting mapping, the motion dataset \mathcal{D}_m provides temporally-inconsistent pseudo-albedo as input, but the relighting model is robust to it since the target is always a temporally consistent real video. To train the delighting model \mathbf{D} , the mapping is inverted. Training with pseudo-albedo from \mathcal{D}_m as *targets* leads to poor temporal consistency, which can affect the following relighting stage. We find that jointly training the delighting model to delight the lighting-rich dataset \mathcal{D}_l , while performing appearance copy on the motion-rich dataset \mathcal{D}_m , similar to the relighting model, yields satisfactory results, as the model can still learn temporal consistency through appearance copy. This process is shown as Video Delighting in Fig. 4. We hypothesize that the main benefit from the appearance copy task is to enforce temporal consistency.

4. Data and Implementation Details

Data Preparation. The hybrid dataset is generated from two sources: static OLAT captures and in-the-wild videos

from VFHQ [72], which contains 15,000 high-quality talking head videos. Our light stage is similar to [40, 61] with 110 programmable LED lights and 36 frontal cameras. We collect data from 67 participants, photographed with diverse head poses, facial expressions and accessories, resulting in a total of 71,280 OLAT sequences. We reserve 7 subjects with varied appearances and genders for evaluation. Following recent works [34, 35, 40], we capture flat-lit images with all lights turned on, approximating diffuse albedo.

To obtain a high-quality dataset with diverse illuminations, we collect 769 HDR maps from PolyHeaven [45], which are randomly paired with OLAT sequences to create the training dataset, using image-based relighting [9, 51]. 600 HDR maps are randomly selected for training and the rest are used for testing. Our final static OLAT dataset contains about 4M rendered images.

Training Details. We initialize both the relighting and delighting models with pre-trained SVD [3]. A training batch consists of 8 video clips at resolution of 512×768 pixels sampled from our hybrid dataset. We train each model, with

a two-stage training schedule. In the first stage, we warm up the model by training on short sequences of 2 frames and optimize all model weights for 200K steps. This speeds-up convergence and the model is quickly tuned to the relighting task. In the second stage, we train with longer sequences of 30 frames over 50K iterations, only optimizing the temporal layers of our models. We optimize both models towards v-prediction objectives [52] using the AdamW [33] optimizer with a learning rate of $1e-5$. More data and training details can be found in the supplement.

5. Experiments

We demonstrate the high-quality video relighting capability of *Lux Post Facto* through extensive evaluations. More video comparison results can be found in the supplement.

Evaluation Metrics. We follow previous work [34, 35] and use PSNR and SSIM [69] to evaluate fidelity and report LPIPS [82] and NIQE [38] scores for perceptual quality. All metrics are computed on the foreground subject, using pre-computed masks from [26].

5.1. Comparison with State-of-the-Art Methods

5.1.1. Video Relighting

We compare *Lux Post Facto* with the state-of-the-art video relighting methods NVPR [79] and SwitchLight [25] on in-the-wild portrait videos. NVPR is a video relighting model trained on a large dynamic OLAT dataset. SwitchLight [25], originally proposes an image relighting model. Recently, the same authors deploy an improved video relighting model commercially, and we compare to this improved version. Since we have no access to ground-truth relit videos, we present a qualitative comparison to both methods in Fig. 5. Our method produces, more realistic highlights and skin tones, finer details and better shadows. In our supplement, we present more video comparisons, highlighting that our method is also more temporally stable.

5.1.2. Image Relighting

Besides video relighting, *Lux Post Facto* can also be used on *single image* portraits, treating images as short static videos. This allows to quantitatively evaluate our model’s performance. To do so, we compare *Lux Post Facto* with the state-of-the-art Total Relighting (TR) [40] method and the SwitchLight image model [25]. We conduct qualitative evaluations on in-the-wild images and both qualitative and quantitative evaluations on our test set, composed of left-out OLATs. The quantitative results are presented in Tab. 1.

While our method is trained with the added temporal consistency constraint, it outperforms both SoTA methods on all metrics by a fair margin. In Fig. 6 & 7, we present qualitative results. In Fig. 6, for which we use subjects from our test set, we observe our method produces results closer




Figure 6. Comparisons to image-based relighting methods on our test set. The first column shows the input and target HDR map. Subsequent columns show results from different methods and ground truth targets. Our results are more faithful and realistic.

Table 1. Quantitative evaluations against image-based relighting methods. Our method outperforms others in all evaluation metrics.

Methods	LPIPS↓	NIQE↓	PSNR↑	SSIM↑
TR [40]	0.1794	6.458	22.44	0.7793
SwitchLight [25]	0.2129	7.166	19.87	0.7481
Ours	0.1158	5.653	24.62	0.8278

to the ground truth. In Fig. 7, we use in-the-wild portraits and show one OLAT captured subject relit by the target lighting as reference. Our method respects skin tones better and produces more plausible highlights and shadows.

5.2. Additional Evaluations

We provide additional evaluations for *Lux Post Facto* to further demonstrate its capabilities.

Consistent Albedo Prediction. Benefiting from our hybrid training strategy, *Lux Post Facto* predicts temporally consistent albedo channel as shown in Fig. 8 (a). The results show that our model removes the existing shading and preserves the facial details.

Lighting Control. *Lux Post Facto* enables precise lighting control thanks to our lighting injection mechanism. We demonstrate its high controllability by relighting portraits using directional light sources. As shown in Fig. 8 (b),




Figure 7. **Comparisons to image-based relighting methods on in-the-wild portrait images.** We present results similarly to Fig. 6, replacing the ground truth with an OLAT computed image as reference. Again, our method has better results with good tone preservation and details.

Table 2. Ablation study on lighting control module.

Methods	LPIPS↓	NIQE↓	PSNR↑	SSIM↑
Clip Image Enc. [47]	0.1747	5.831	20.73	0.7808
Ours	0.1158	5.653	24.62	0.8278

our results showcase complex light transport effects such as specular reflections and cast shadows coherent with the rotating light sources. In Fig. 8 (c), we make the directional light source more diffuse and show that our method renders realistic soft shadows.

5.3. Ablation Studies

We conduct ablation studies on two core designs: lighting control module and hybrid dataset training strategy. We first show that our lighting control mechanism is significantly better than using the commonly used CLIP image encoder [47] for image conditioning (Tab. 2). Then, in Fig. 9 we show that with our hybrid dataset training strategy our model produces sharper and more photorealistic results compared to training only on the lighting-rich data \mathcal{D}_l . Without using it (*i.e.* training solely on OLAT dataset with simulated camera motions), the model cannot handle subject’s movement well and tends to produce blurry results (Fig. 9). More details and qualitative results can be found



Figure 8. **Albedo and directionally lit relighting results.** Our method can predict (a) temporally consistent albedo, and (b) enables precise lighting control. Diffusing the light source effectively leads to softer shadows (c).




Figure 9. **Ablation study on hybrid dataset training.** Without hybrid training on the motion-rich in-the-wild videos, the model produces blurry results.

in the supplement.

6. Conclusion

We presented *Lux Post Facto*, a video relighting method that brings realistic and temporally consistent relighting to standard portrait videos in post-production. By treating video relighting as an HDR map conditioned generation process, and fine-tuning a state-of-the-art video diffusion model, *Lux Post Facto* achieves generalized fine-grained lighting control of portrait videos. *Lux Post Facto* only requires a set of static OLAT images and a larger pool of in-the-wild videos for training. This hybrid dataset effectively provides relighting supervision for individual frames while fostering temporal stability across frames. Discussion of limitations and future work can be found in the supplement.

References

- [1] Jonathan T. Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(8):1670–1687, 2015. 2
- [2] Sai Bi, Stephen Lombardi, Shunsuke Saito, Tomas Simon, Shih-En Wei, Kevyn McPhail, Ravi Ramamoorthi, Yaser Sheikh, and Jason M. Saragih. Deep relightable appearance models for animatable faces. *ACM Trans. Graph.*, 40(4): 89:1–89:15, 2021. 2
- [3] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, Varun Jampani, and Robin Rombach. Stable video diffusion: Scaling latent video diffusion models to large datasets. *CoRR*, abs/2311.15127, 2023. 2, 3, 4, 6
- [4] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 22563–22575. IEEE, 2023. 3
- [5] Tim Brooks, Aleksander Holynski, and Alexei A. Efros. Instructpix2pix: Learning to follow image editing instructions. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 18392–18402. IEEE, 2023. 4
- [6] Ziqi Cai, Kaiwen Jiang, Shu-Yu Chen, Yu-Kun Lai, Hongbo Fu, Boxin Shi, and Lin Gao. Real-time 3d-aware portrait video relighting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 6221–6231. IEEE, 2024. 3
- [7] Di Chang, Yichun Shi, Quankai Gao, Hongyi Xu, Jessica Fu, Guoxian Song, Qing Yan, Yizhe Zhu, Xiao Yang, and Mohammad Soleymani. Magicpose: Realistic human poses and facial expressions retargeting with identity-aware diffusion. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. 3
- [8] Pascal Chang, Jingwei Tang, Markus Gross, and Vinicius C. Azevedo. How I warped your noise: a temporally-correlated noise prior for diffusion models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. 3
- [9] Paul E. 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, SIGGRAPH 2000, New Orleans, LA, USA, July 23-28, 2000*, pages 145–156. ACM, 2000. 2, 4, 6
- [10] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 8780–8794, 2021. 2
- [11] Per Einarsson, Charles-Félix Chabert, Andrew Jones, Wan-Chun Ma, Bruce Lamond, Tim Hawkins, Mark T. Bolas, Sebastian Sylwan, and Paul E. Debevec. Relighting human locomotion with flowed reflectance fields. In *Proceedings of the Eurographics Symposium on Rendering Techniques, Nicosia, Cyprus, 2006*, pages 183–194. Eurographics Association, 2006. 2
- [12] Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. Structure and content-guided video synthesis with diffusion models. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 7312–7322. IEEE, 2023. 3
- [13] David Futschik, Kelvin Ritland, James Vecore, Sean Fanello, Sergio Orts-Escolano, Brian Curless, Daniel Sýkora, and Rohit Pandey. Controllable light diffusion for portraits. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 8412–8421. IEEE, 2023. 2
- [14] Songwei Ge, Seungjun Nah, Guilin Liu, Tyler Poon, Andrew Tao, Bryan Catanzaro, David Jacobs, Jia-Bin Huang, Ming-Yu Liu, and Yogesh Balaji. Preserve your own correlation: A noise prior for video diffusion models. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 22873–22884. IEEE, 2023. 3
- [15] Litong Gong, Yiran Zhu, Weijie Li, Xiaoyang Kang, Biao Wang, Tiezheng Ge, and Bo Zheng. Atomovideo: High fidelity image-to-video generation. *CoRR*, abs/2403.01800, 2024. 4
- [16] David Griffiths, Tobias Ritschel, and Julien Philip. Outcast: Outdoor single-image relighting with cast shadows. *Comput. Graph. Forum*, 41(2):179–193, 2022. 2
- [17] Kaiwen Guo, Peter Lincoln, Philip Davidson, Jay Busch, Xueming Yu, Matt Whalen, Geoff Harvey, Sergio Orts-Escolano, Rohit Pandey, Jason Dourgarian, Danhang Tang, Anastasia Tkach, Adarsh Kowdle, Emily Cooper, Mingsong Dou, Sean Ryan Fanello, Graham Fyffe, Christoph Rhemann, Jonathan Taylor, Paul E. Debevec, and Shahram Izadi. The relightables: volumetric performance capture of humans with realistic relighting. *ACM Trans. Graph.*, 38(6):217:1–217:19, 2019. 2
- [18] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. 3
- [19] Mingming He, Pascal Clausen, Ahmet Levent Taşel, Li Ma, Oliver Pilarski, Wenqi Xian, Laszlo Rikker, Xueming Yu, Ryan Burgert, Ning Yu, and Paul Debevec. Diffrelight: Diffusion-based facial performance relighting. *CoRR*, abs/2410.08188, 2024. 2, 3
- [20] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey A. Gritsenko, Diederik P. Kingma, Ben Poole, Mohammad Norouzi, David J. Fleet, and Tim Sali-

- mans. Imagen video: High definition video generation with diffusion models. *CoRR*, abs/2210.02303, 2022. 3
- [21] Andrew Hou, Zhixin Shu, Xuaner Zhang, He Zhang, Yannick Hold-Geoffroy, Jae Shin Yoon, and Xiaoming Liu. COMPOSE: comprehensive portrait shadow editing. *CoRR*, abs/2408.13922, 2024. 2
- [22] Andrew Z. Hou, Ze Zhang, Michel Sarkis, Ning Bi, Yiyong Tong, and Xiaoming Liu. Towards high fidelity face relighting with realistic shadows. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 14719–14728. Computer Vision Foundation / IEEE, 2021. 2
- [23] Andrew Z. Hou, Michel Sarkis, Ning Bi, Yiyong Tong, and Xiaoming Liu. Face relighting with geometrically consistent shadows. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 4207–4216. IEEE, 2022. 2
- [24] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139:1–139:14, 2023. 3
- [25] Hoon Kim, Minje Jang, Wonjun Yoon, Jisoo Lee, Donghyun Na, and Sanghyun Woo. Switchlight: Co-design of physics-driven architecture and pre-training framework for human portrait relighting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 25096–25106. IEEE, 2024. 2, 3, 7
- [26] Taehun Kim, Kunhee Kim, Joonyeong Lee, Dongmin Cha, Jiho Lee, and Daijin Kim. Revisiting image pyramid structure for high resolution salient object detection. In *Computer Vision - ACCV 2022 - 16th Asian Conference on Computer Vision, Macao, China, December 4-8, 2022, Proceedings, Part VII*, pages 257–273. Springer, 2022. 7
- [27] Peter Kocsis, Julien Philip, Kalyan Sunkavalli, Matthias Nießner, and Yannick Hold-Geoffroy. Lightit: Illumination modeling and control for diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 9359–9369. IEEE, 2024. 2, 5
- [28] Ha A. Le and Ioannis A. Kakadiaris. Illumination-invariant face recognition with deep relit face images. In *IEEE Winter Conference on Applications of Computer Vision, WACV 2019, Waikoloa Village, HI, USA, January 7-11, 2019*, pages 2146–2155. IEEE, 2019. 2
- [29] Chen Li, Kun Zhou, and Stephen Lin. Intrinsic face image decomposition with human face priors. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V*, pages 218–233. Springer, 2014. 2
- [30] Junxuan Li, PChen Cao, Gabriel Schwartz, Rawal Khirrodkar, Christian Richardt, Tomas Simon, Yaser Sheikh, Xueming Yu, Ryan Burgert, and Shunsuke Saito. Ura-vatar: Universal relightable gaussian codec avatars. *CoRR*, abs/2410.24223, 2024. 2, 3
- [31] Min-Hui Lin, Mahesh Reddy, Guillaume Berger, Michel Sarkis, Fatih Porikli, and Ning Bi. Edgerelight360: Text-conditioned 360-degree HDR image generation for real-time on-device video portrait relighting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024 - Workshops, Seattle, WA, USA, June 17-18, 2024*, pages 831–840. IEEE, 2024. 3
- [32] Yaojie Liu, Andrew Z. Hou, Xinyu Huang, Liu Ren, and Xiaoming Liu. Blind removal of facial foreign shadows. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*, page 88. BMVA Press, 2022. 2
- [33] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. 7
- [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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 195–205. IEEE, 2023. 2, 3, 5, 6, 7
- [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. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 4263–4273. IEEE, 2024. 2, 3, 4, 5, 6, 7
- [36] Abhimitra Meka, Christian Häne, Rohit Pandey, Michael Zollhöfer, Sean Ryan Fanello, Graham Fyffe, Adarsh Kowdle, Xueming Yu, Jay Busch, Jason Dourgarian, Peter Denny, Sofien Bouaziz, Peter Lincoln, Matt Whalen, Geoff Harvey, Jonathan Taylor, Shahram Izadi, Andrea Tagliasacchi, Paul E. Debevec, Christian Theobalt, Julien P. C. Valentin, and Christoph Rhemann. Deep reflectance fields: high-quality facial reflectance field inference from color gradient illumination. *ACM Trans. Graph.*, 38(4):77:1–77:12, 2019. 2
- [37] Abhimitra Meka, Rohit Pandey, Christian Häne, Sergio Orts-Escolano, Peter Barnum, Philip Davidson, Daniel Erickson, Yinda Zhang, Jonathan Taylor, Sofien Bouaziz, Chloe LeGendre, Wan-Chun Ma, Ryan S. Overbeck, Thabo Beeler, Paul E. Debevec, Shahram Izadi, Christian Theobalt, Christoph Rhemann, and Sean Ryan Fanello. Deep relightable textures: volumetric performance capture with neural rendering. *ACM Trans. Graph.*, 39(6):259:1–259:21, 2020. 2
- [38] Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal Process. Lett.*, 20(3):209–212, 2013. 7
- [39] Xingang Pan, Xudong Xu, Chen Change Loy, Christian Theobalt, and Bo Dai. A shading-guided generative implicit model for shape-accurate 3d-aware image synthesis. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 20002–20013, 2021. 2
- [40] Rohit Pandey, Sergio Orts-Escolano, Chloe LeGendre, Christian Häne, Sofien Bouaziz, Christoph Rhemann,

- Paul E. Debevec, and Sean Ryan Fanello. Total relighting: learning to relight portraits for background replacement. *ACM Trans. Graph.*, 40(4):43:1–43:21, 2021. 2, 3, 4, 5, 6, 7
- [41] Pieter Peers, Naoki Tamura, Wojciech Matusik, and Paul E. Debevec. Post-production facial performance relighting using reflectance transfer. *ACM Trans. Graph.*, 26(3):52, 2007. 2
- [42] Julien Philip, Michaël Gharbi, Tinghui Zhou, Alexei A. Efros, and George Drettakis. Multi-view relighting using a geometry-aware network. *ACM Trans. Graph.*, 38(4):78:1–78:14, 2019. 2
- [43] Julien Philip, Sébastien Morgenthaler, Michaël Gharbi, and George Drettakis. Free-viewpoint indoor neural relighting from multi-view stereo. *ACM Trans. Graph.*, 40(5):194:1–194:18, 2021. 2
- [44] Yohan Poirier-Ginter, Alban Gauthier, Julien Philip, Jean-François Lalonde, and George Drettakis. A diffusion approach to radiance field relighting using multi-illumination synthesis. *Comput. Graph. Forum*, 43(4):i–x, 2024. 2
- [45] PolyHeaven. Poly heaven. <https://polyhaven.com/>. 6
- [46] Puntawat Ponglertnapakorn, Nontawat Tritrong, and Supasorn Suwajanakorn. Difareli: Diffusion face relighting. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 22589–22600. IEEE, 2023. 2
- [47] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, pages 8748–8763. PMLR, 2021. 2, 8
- [48] Anurag Ranjan, Kwang Moo Yi, Jen-Hao Rick Chang, and Oncel Tuzel. Facelit: Neural 3d relightable faces. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 8619–8628. IEEE, 2023. 2
- [49] Mengwei Ren, Wei Xiong, Jae Shin Yoon, Zhixin Shu, Jianming Zhang, Hyunjoon Jung, Guido Gerig, and He Zhang. Relightful harmonization: Lighting-aware portrait background replacement. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 6452–6462. IEEE, 2024. 2, 4
- [50] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 10674–10685. IEEE, 2022. 2
- [51] Mark Sagar. Reflectance field rendering of human faces for "spider-man 2". In *International Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 2005, Los Angeles, California, USA, July 31 - August 4, 2005, Courses*, page 14. ACM, 2005. 2, 4, 6
- [52] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. 7
- [53] Viraj Shah, Svetlana Lazebnik, and Julien Philip. Join: Joint gans inversion for intrinsic image decomposition. *CoRR*, abs/2305.11321, 2023. 2
- [54] Davoud Shahlaei and Volker Blanz. Realistic inverse lighting from a single 2d image of a face, taken under unknown and complex lighting. In *11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition, FG 2015, Ljubljana, Slovenia, May 4-8, 2015*, pages 1–8. IEEE Computer Society, 2015. 2
- [55] Jiahao Shao, Yuanbo Yang, Hongyu Zhou, Youmin Zhang, Yujun Shen, Matteo Poggi, and Yiyi Liao. Learning temporally consistent video depth from video diffusion priors. *CoRR*, abs/2406.01493, 2024. 3
- [56] Amnon Shashua and Tammy Riklin-Raviv. The quotient image: Class-based re-rendering and recognition with varying illuminations. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(2):129–139, 2001. 2
- [57] Yichang Shih, Sylvain Paris, Connelly Barnes, William T. Freeman, and Frédo Durand. Style transfer for headshot portraits. *ACM Trans. Graph.*, 33(4):148:1–148:14, 2014. 2, 5
- [58] Zhixin Shu, Sunil Hadap, Eli Shechtman, Kalyan Sunkavalli, Sylvain Paris, and Dimitris Samaras. Portrait lighting transfer using a mass transport approach. *ACM Trans. Graph.*, 37(1):2, 2018. 2
- [59] Zhixin Shu, Sunil Hadap, Eli Shechtman, Kalyan Sunkavalli, Sylvain Paris, and Dimitris Samaras. Portrait lighting transfer using a mass transport approach. *ACM Trans. Graph.*, 37(1):2, 2018. 5
- [60] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, Devi Parikh, Sonal Gupta, and Yaniv Taigman. Make-a-video: Text-to-video generation without text-video data. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. 3
- [61] Tiancheng Sun, Jonathan T. Barron, Yun-Ta Tsai, Zexiang Xu, Xueming Yu, Graham Fyffe, Christoph Rhemann, Jay Busch, Paul E. Debevec, and Ravi Ramamoorthi. Single image portrait relighting. *ACM Trans. Graph.*, 38(4):79:1–79:12, 2019. 2, 6
- [62] Feitong Tan, Sean Fanello, Abhimitra Meka, Sergio Orts-Escolano, Danhang Tang, Rohit Pandey, Jonathan Taylor, Ping Tan, and Yinda Zhang. Volux-gan: A generative model for 3d face synthesis with HDRI relighting. In *SIGGRAPH '22: Special Interest Group on Computer Graphics and Interactive Techniques Conference, Vancouver, BC, Canada, August 7 - 11, 2022*, pages 58:1–58:9. ACM, 2022. 2
- [63] Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason M. Saragih, Matthias Nießner, Rohit Pandey, Sean Ryan Fanello, Gordon Wetzstein, Jun-Yan Zhu, Christian Theobalt, Maneesh Agrawala, Eli

- Shechtman, Dan B. Goldman, and Michael Zollhöfer. State of the art on neural rendering. *Comput. Graph. Forum*, 39(2):701–727, 2020. 2
- [64] Linrui Tian, Qi Wang, Bang Zhang, and Liefeng Bo. EMO: emote portrait alive - generating expressive portrait videos with audio2video diffusion model under weak conditions. *CoRR*, abs/2402.17485, 2024. 4
- [65] Xiang Wang, Hangjie Yuan, Shiwei Zhang, Dayou Chen, Jiniu Wang, Yingya Zhang, Yujun Shen, Deli Zhao, and Jingen Zhou. Videocomposer: Compositional video synthesis with motion controllability. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. 3
- [66] Yaohui Wang, Xinyuan Chen, Xin Ma, Shangchen Zhou, Ziqi Huang, Yi Wang, Ceyuan Yang, Yinan He, Jiashuo Yu, Peiqing Yang, Yuwei Guo, Tianxing Wu, Chenyang Si, Yuming Jiang, Cunjian Chen, Chen Change Loy, Bo Dai, Dahua Lin, Yu Qiao, and Ziwei Liu. LAVIE: high-quality video generation with cascaded latent diffusion models. *CoRR*, abs/2309.15103, 2023. 3
- [67] Youjia Wang, Kai He, Taotao Zhou, Kaixin Yao, Nianyi Li, Lan Xu, and Jingyi Yu. Free-view face relighting using a hybrid parametric neural model on a SMALL-OLAT dataset. *Int. J. Comput. Vis.*, 131(4):1002–1021, 2023. 2
- [68] Yifan Wang, Aleksander Holynski, Xiuming Zhang, and Xuaner Zhang. Sunstage: Portrait reconstruction and relighting using the sun as a light stage. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 20792–20802. IEEE, 2023. 2
- [69] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4):600–612, 2004. 7
- [70] Zhibo Wang, Xin Yu, Ming Lu, Quan Wang, Chen Qian, and Feng Xu. Single image portrait relighting via explicit multiple reflectance channel modeling. *ACM Trans. Graph.*, 39(6):220:1–220:13, 2020. 2
- [71] Andreas Wenger, Andrew Gardner, Chris Tchou, Jonas Unger, Tim Hawkins, and Paul E. Debevec. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Trans. Graph.*, 24(3):756–764, 2005. 2
- [72] Liangbin Xie, Xintao Wang, Honglun Zhang, Chao Dong, and Ying Shan. VFHQ: A high-quality dataset and benchmark for video face super-resolution. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2022, New Orleans, LA, USA, June 19-20, 2022*, pages 656–665. IEEE, 2022. 5, 6
- [73] Haotian Yang, Mingwu Zheng, Chongyang Ma, Yu-Kun Lai, Pengfei Wan, and Haibin Huang. VRMM: A volumetric relightable morphable head model. In *ACM SIGGRAPH 2024 Conference Papers, SIGGRAPH 2024, Denver, CO, USA, 27 July 2024- 1 August 2024*, page 46. ACM, 2024. 3
- [74] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, Da Yin, Xiaotao Gu, Yuxuan Zhang, Weihang Wang, Yean Cheng, Ting Liu, Bin Xu, Yuxiao Dong, and Jie Tang. Cogvideox: Text-to-video diffusion models with an expert transformer. *CoRR*, abs/2408.06072, 2024. 2
- [75] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*, 2023. 2
- [76] 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 Trans. Graph.*, 41(6):231:1–231:21, 2022. 2
- [77] Jiangjian Yu, Ling Zhang, Qing Zhang, Qifei Zhang, Daiguo Zhou, Chao Liang, and Chunxia Xiao. Portrait shadow removal using context-aware illumination restoration network. *IEEE Trans. Image Process.*, 34:1–15, 2025. 2
- [78] Chong Zeng, Yue Dong, Pieter Peers, Youkang Kong, Hongzhi Wu, and Xin Tong. Dilightnet: Fine-grained lighting control for diffusion-based image generation. In *ACM SIGGRAPH 2024 Conference Papers, SIGGRAPH 2024, Denver, CO, USA, 27 July 2024- 1 August 2024*, page 73. ACM, 2024. 2
- [79] Longwen Zhang, Qixuan Zhang, Minye Wu, Jingyi Yu, and Lan Xu. Neural video portrait relighting in real-time via consistency modeling. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 782–792. IEEE, 2021. 2, 3, 7
- [80] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 3813–3824. IEEE, 2023. 2
- [81] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Ic-light github page, 2024. 2, 4
- [82] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 586–595. Computer Vision Foundation / IEEE Computer Society, 2018. 7
- [83] Xuaner Cecilia Zhang, Jonathan T. Barron, Yun-Ta Tsai, Rohit Pandey, Xiuming Zhang, Ren Ng, and David E. Jacobs. Portrait shadow manipulation. *ACM Trans. Graph.*, 39(4):78, 2020. 2
- [84] Yabo Zhang, Yuxiang Wei, Dongsheng Jiang, Xiaopeng Zhang, Wangmeng Zuo, and Qi Tian. Controlvideo: Training-free controllable text-to-video generation. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. 3
- [85] Shihao Zhao, Dongdong Chen, Yen-Chun Chen, Jianmin Bao, Shaozhe Hao, Lu Yuan, and Kwan-Yee K Wong. Uni-controlnet: All-in-one control to text-to-image diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024. 2

- [86] Hao Zhou, Sunil Hadap, Kalyan Sunkavalli, and David Jacobs. Deep single-image portrait relighting. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 7193–7201. IEEE, 2019. [2](#)