

DiFaReli: Diffusion Face Relighting

Puntawat Ponglertnapakorn

Nontawat Tritrong

Supasorn Suwajanakorn

VISTEC, Thailand




Figure 1: **Relighting and shadow manipulation.** Our method can realistically relight an input image using different target lights shown in small shading references. We can add or remove cast shadows, as well as control their intensity.

Abstract

We present a novel approach to single-view face relighting in the wild. Handling non-diffuse effects, such as global illumination or cast shadows, has long been a challenge in face relighting. Prior work often assumes Lambertian surfaces, simplified lighting models or involves estimating 3D shape, albedo, or a shadow map. This estimation, however, is error-prone and requires many training examples with lighting ground truth to generalize well. Our work bypasses the need for accurate estimation of intrinsic components and can be trained solely on 2D images without any light stage data, multi-view images, or lighting ground truth. Our key idea is to leverage a conditional diffusion implicit model (DDIM) for decoding a disentangled light encoding along with other encodings related to 3D shape and facial identity inferred from off-the-shelf estimators. We also propose a novel conditioning technique that eases the modeling of the complex interaction between light and geometry by using a rendered shading reference to spatially modulate the DDIM. We achieve state-of-the-art performance on standard benchmark Multi-PIE and can photorealistically relight in-the-wild images. Please visit our page: <https://diffusion-face-relighting.github.io>.

1. Introduction

The ability to relight face images under any lighting condition has a wide range of applications, such as in Augmented Reality, where consistent lighting for all individuals in the scene is essential to achieve realism. Another use is in portrait photography, where one may aim to soften cast shadows to create a more pleasing, diffuse appearance. Yet, relighting single-view face images remains unsolved.

Relighting a face image requires modeling the physical interactions between the geometry, material, and lighting, which are not inherently present in a 2D image and difficult to estimate accurately. Earlier work [4, 53, 26, 68, 56] thus often assumes Lambertian surfaces and a simplified lighting model, which struggle to model complex light interactions like global illumination, subsurface scattering, or cast shadows. Using multi-view, multi-illumination data from a light stage or a simulation, [38, 72] proposed relighting pipelines that predict surface normals, albedo, and a set of diffuse and specular maps with neural networks given a target HDR map. Some recent methods aim to specifically model cast shadows by predicting a shadow map with a neural network [23, 35] or rendering a shadow map through physical ray tracing with estimated geometry [22].

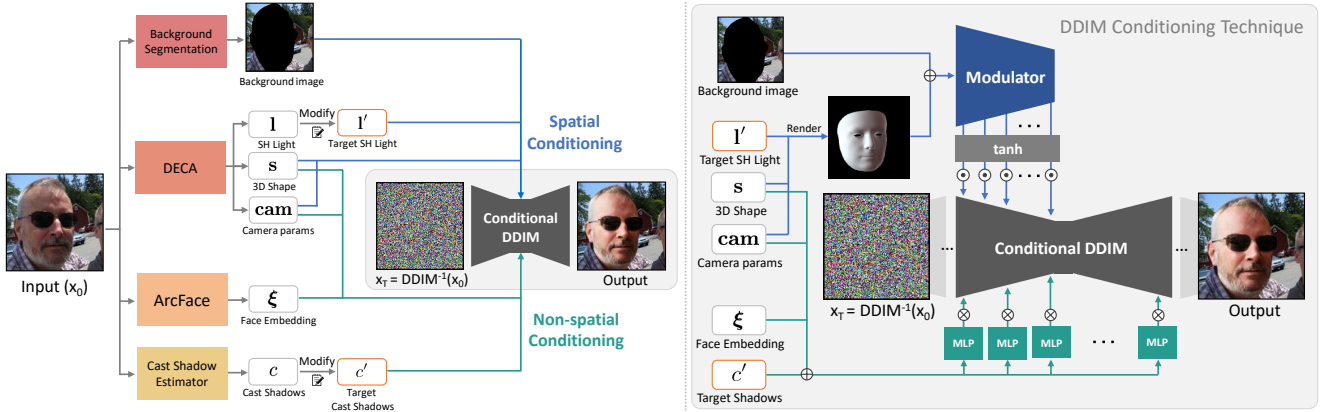


Figure 2: **Pipeline overview.** We use off-the-shelf estimators to encode the input image into encodings of light, shape, camera, face embedding, shadow scalar, and background image, which are then fed to DDIM via “spatial” and “non-spatial” conditioning techniques. For spatial conditioning, the modified SH, 3D shape, and camera encodings are rendered to a shading reference, which is then concatenated with the background image. This concatenated image is fed into *Modulator* to produce spatial modulation weights for DDIM’s first half. Non-spatial conditioning feeds a stack of 3D shape, camera, face embedding, and a modified shadow scalar to a set of MLPs for modulating the DDIM with our modified version of adaptive group normalization (AdaGN).

These approaches share a common scheme in which they first intrinsically decompose the face image into its surface normals, albedo, and lighting parameters, then use them along with a shadow or visibility map to render a relit output. However, one major issue of this scheme stems from its over-reliance on the accuracy of the estimated components, which are difficult to estimate correctly in real-world scenarios. For instance, when an input image contains cast shadows that need to be removed, these approaches often leave behind shadow residuals in the predicted albedo map, which in turn produces artifacts in the final output (Figure 3). Estimating the geometry for other areas like hair and ears is also extremely challenging, and they are often omitted from relighting pipelines, resulting in unrealistic final composites (Figure 3, 4 and 5).

This paper introduces an alternative approach that does not rely on accurate intrinsic decomposition of the face and can be trained exclusively on 2D images, without any 3D face scan, multi-view images, or lighting ground truth, once given a few off-the-shelf estimators. The general idea of our method is simple: we first encode the input image into a feature vector that disentangles the light information from other information about the input image. Then, we modify the light encoding in the feature vector and decode it. The challenge, however, is how to disentangle the light encoding well enough so that the decoding will only affect the shading without altering the person’s shape and identity. Our key idea is to leverage a conditional diffusion implicit model [59] with a novel conditioning technique for this task and learn the complex light interactions implicitly via the generative model trained on a real-world 2D face dataset.

Our method relies on mechanisms recently introduced

in Denoising Diffusion Implicit Models (DDIM) [59] and Diffusion Autoencoders (DiffAE) [41]. By exploiting the deterministic reversal process of DDIM proposed by Song et al. [59], DiffAE shows how one can encode an image into a meaningful semantic code and disentangle it from other information, which includes stochastic variations. By modifying the semantic code and decoding it, DiffAE can manipulate semantic attributes in a real image. Relighting can be thought of as a manipulation of the “light” attribute in the input image. But unlike DiffAE, which discovers semantic attributes automatically and encodes them in a *latent* code, our method requires an explicit and interpretable light encoding that facilitates lighting manipulation by the user.

To solve this problem without access to the lighting ground truth, we use an off-the-shelf estimator, DECA [13], to encode the lighting information as spherical harmonic (SH) coefficients and rely on a conditional DDIM to decode and learn to disentangle the light information in the process. Unlike prior work, our use of SH lighting is not for direct rendering of the output shading, as this would be restricted by the limited capacity of SH lighting to express complex illumination. Rather, it is used to condition a generative process that learns the complex shading prior to reproduce real-world 2D face images. To help preserve the input’s identity during relighting, we also condition the DDIM on other attributes, such as the face shape and deep feature embeddings from a face recognition model, ArcFace [8].

Another key component is our novel technique for conditioning the DDIM. Instead of treating the SH lighting as a global, non-spatial condition vector as in DiffAE or other diffusion models, we render a shading reference using the known SH equation and feed it to another network called

Modulator, which computes layer-wise spatial modulation weights for the DDIM. This conditioning technique helps retain spatial information in the shading reference and provides an easy-to-learn conditioning signal as the pixel intensities in the shading reference correlate more directly with the output RGB pixels.

With our novel framework, the visibility of cast shadows can also be modeled with a simple modification: add one conditioning scalar that indicates the “degree” of cast shadows to the DDIM. At test time, we can strengthen or attenuate cast shadows by modifying this scalar. Since our diffusion-based framework does not directly use this flag or the shape and SH parameters in a physical image formation model, imprecise estimation of these parameters can be tolerated and does not significantly compromise our quality.

Our method produces highly plausible and photorealistic results and can convincingly strengthen or attenuate cast shadows. Moreover, we can reproduce the original facial details with high fidelity, which is difficult for competing methods that predict an albedo map with neural networks. We conduct qualitative and quantitative evaluations and achieve state-of-the-art performance on a standard benchmark, Multi-Pie [17]. To summarize, our contributions are:

- A state-of-the-art face relighting framework based on a conditional DDIM that produces photorealistic shading without requiring accurate intrinsic decomposition or 3D and lighting ground truth.
- A novel conditioning technique that converts a shading reference rendered from the estimated light and shape parameters into layer-wise spatial modulation weights.

2. Related work

A common approach to face relighting [4, 65, 26, 53, 68, 56] is to decompose an input image into multiple intrinsic components (e.g., lighting, albedo, surface normals) and recompose the image back with modified light-related components. The decomposition can be done by regularized optimization [4], fitting a morphable model [5, 68], or predicted from a neural network [53, 26, 65, 35, 69, 38, 56]. Most earlier methods [4, 53, 26, 68, 56] assume Lambertian surfaces, a simplified lighting model, such as second-order spherical harmonics, and a physical image formation model based on these simplified assumptions. Thus, they cannot handle non-diffuse effects, such as specular highlights or cast shadows, commonly occur in real-world scenarios.

Rather than decomposing an image into physical components, some techniques [77, 61] rely on an encoder-decoder network with a bottleneck layer that holds a latent lighting representation. Zhou et al. [77] force such a latent code to be predictive of the SH lighting and train another regressor that can map the SH lighting of a reference image back to

a latent code for relighting. Sun et al. [61] rely on a similar idea but use a low-resolution illumination map, obtained from a light stage, e.g., [70], instead of the SH lighting. In principle, these learning techniques can learn to handle hard shadows and specularities, given sufficient examples. However, in practice, these approaches still struggle to model those effects due to their small light stage data [61] or limited variations in their synthetic dataset [77]. In contrast, our framework can be trained on 2D face images, which are cheaply available and cover far more diverse scenarios.

Handling non-diffuse components. Relighting non-diffuse components has long been a challenge. Nestmeyer et al. [35] propose a two-stage framework to predict non-diffuse components as a residual correction of a diffuse rendering from their first stage. Cast shadows are predicted separately as a visibility map, which is multiplied to the output. Wang et al. [69] propose a technique based on intrinsic decomposition that predicts shadow and specular maps by learning from their own large-scale relighting dataset. Pandey et al. [38] introduces a pipeline that predicts a set of specular maps with varying degrees of Phong exponents using estimated surface normals and an input HDR environment map. These maps along with diffuse and albedo maps are used to predict a relit image with a UNet. Yeh et al. [72] uses a pipeline similar to [38] but with synthetic light stage data generated from 3D face scans, and an albedo refinement step to reduce the domain gap between synthetic and real data. Hou et al. (2021) [23] compute a shadow map based on a morphable model fitted to the input and standard ray tracing, then use it to help predict the ratio of pixel luminance changes for relighting. Hou et al. (2022) [22] predict a shadow mask via ray tracing based on their estimated depth map and render a relit image with estimated albedo and shading maps from neural networks.

While these methods [35, 23, 22] produce promising non-diffuse effects, their physical image formation model makes it difficult to operate in the wild when the estimated geometry is inaccurate. As a result, some [23, 22] can only relight the face region but not the ears or hair and still struggle to handle in-the-wild cast shadows (Figure 3). Neural rendering approaches [69, 38] can tolerate some estimation error, but high-frequency details are often lost, even when predicted by a UNet [69] and still require light stage data. The synthetic light stage data of [72] shows great potential but currently relies on 3D face scans to generate, which are difficult to obtain compared to 2D images.

Style transfer-based methods. Another class of relighting approaches is based on style transfer. Although some of these methods [28, 31, 54] do not directly solve face relighting, they can be adapted for this task by transferring the lighting and shading styles from one image to another. However, the style representation used in these methods captures broad information beyond the lighting condition

and cannot produce accurate relit results. Shu et al. [55] solves relighting using color histogram matching that is redesigned to be spatially varying and dependent on the face geometry and can be solved as a mass transport problem. However, this technique does not model self-occlusion required for handling cast shadows and can easily suffer from occlusion by hair or accessories.

GAN-based methods. A few techniques use GANs [16] to solve relighting [63, 32]. Tewari et al. [63] rely on StyleRig [64], a technique to enable semantic control of StyleGAN [25] by mapping a set of morphable model parameters along with an initial StyleGAN latent code to a new one representing the target parameters. Specifically, they extend StyleRig, which only works on synthetic images, to real images by optimizing a latent code that reproduces the input image and use StyleRig to manipulate the lighting condition. Similarly, Mallikarjun et al. [32] maps a target illumination and a StyleGAN latent code predicted from pSp network [45] to a new code that represents a relit image. However, these techniques tend to change the identity and facial details of the input person due to the imperfect GAN inversion. New GAN inversion techniques [12, 46] are promising, but no relighting results with these techniques have been demonstrated. Our solution overcomes this issue by leveraging DDIM’s near-perfect inversion and produces high-fidelity results that preserve the original details.

3. Approach

Given an input face image, we seek to relight this image under a target lighting condition, described by spherical harmonic coefficients and an additional scalar representing the “degree” of visible cast shadows. To explain our method, we first cover relevant background on DDIM [59] and a key finding from DiffAE [41] that shows how a conditional DDIM can perform attribute manipulation on real images by acting as both a decoder and a “stochastic” encoder.

3.1. Background: Conditional DDIM & DiffAE

Our method relies on a conditional Denoising Diffusion Implicit Model (DDIM) [59], which is a variant of diffusion models [58, 19, 60]. (For a full review and notation convention, please refer to [59].) Unlike standard diffusion models, DDIM uses a non-Markovian inference process that relies on the conditional distribution $q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0)$ that is conditioned on \mathbf{x}_0 (the original image) in addition to \mathbf{x}_t . One important implication is that the generative process can be made deterministic, allowing us to deterministically map $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ to \mathbf{x}_0 and vice versa. Here the mapping from \mathbf{x}_0 to \mathbf{x}_T can be viewed as the encoding of an input image \mathbf{x}_0 to a latent variable \mathbf{x}_T .

Diffusion Autoencoders (DiffAE) [41] show that such image encoding yields \mathbf{x}_T that contains little semantic information about the input image \mathbf{x}_0 and propose to condi-

tion the DDIM also on a learnable latent variable \mathbf{z} predicted from a separate image encoder. By jointly training the image encoder and the DDIM, the encoded \mathbf{z} now captures meaningful semantics, while the encoded \mathbf{x}_T , inferred by reversing the deterministic generative process of the DDIM, captures the rest of the information not encoded in \mathbf{z} , such as stochastic variations. The resulting latent code $(\mathbf{z}, \mathbf{x}_T)$ can also be decoded back to the input image near-perfectly using the same conditional DDIM. By modifying the semantic latent variable \mathbf{z} and decoding the new $(\mathbf{z}', \mathbf{x}_T)$, DiffAE can manipulate semantic attributes of a real input image—a capability that inspires our work.

3.2. Method overview

The general idea of our method is to encode the input image into a feature vector that disentangles the light information from other information about the input image. Then, the relit image is produced by modifying the light encoding in the feature vector and decoding the resulting vector with a conditional DDIM (see Figure 2). This process is similar to how DiffAE performs attribute manipulation; however, our task requires well-disentangled and interpretable light encoding that facilitates lighting manipulation by the user.

To solve this, we use off-the-shelf estimators to encode an input image into light, shape, and camera encodings, as well as a face embedding, a shadow scalar, and a background image (Section 3.3). Then, these encodings are used to condition our DDIM decoder (Section 3.4) with a novel conditioning technique (Section 3.5). For training, we use a standard diffusion objective to reconstruct training images (Section 3.6). To relight, we reverse the generative process of the DDIM conditioned on the input’s encodings to obtain \mathbf{x}_T , modify the light encoding, and decode \mathbf{x}_T using the modified encodings (Section 3.7).

3.3. Encoding

The goal of this step is to encode the input face image $I \in \mathbb{R}^{H \times W \times 3}$ into a feature vector:

$$\mathbf{f} = (\mathbf{l}, \mathbf{s}, \mathbf{cam}, \xi, c, \mathbf{bg}), \quad (1)$$

where $\mathbf{l} \in \mathbb{R}^{9 \times 3}$ represents 2nd-order spherical harmonic lighting coefficients, $\mathbf{s} \in \mathbb{R}^{|\mathbf{s}|}$ represents parameterized face shape, $\mathbf{cam} \in \mathbb{R}^{1+2}$ represents orthographic camera parameters, $\xi \in \mathbb{R}^{512}$ is a deep feature embedding based on ArcFace [8], c is a scalar that indicates the degree of visible cast shadows, and $\mathbf{bg} \in \mathbb{R}^{H \times W \times 3}$ contains the background pixels with the face, hair, neck masked out. These variables will be inferred using off-the-shelf or pretrained estimators.

Light, shape, & camera encodings ($\mathbf{l}, \mathbf{s}, \mathbf{cam}$). We use an off-the-shelf single-view 3D face reconstruction method, DECA [13]. Given a face image, DECA predicts the 3D face shape, camera pose, albedo map, and spherical harmonic lighting (SH) coefficients.

For our light encoding \mathbf{l} , we directly use the SH coefficients from DECA, consisting of 9 coefficients for each channel of the RGB. DECA’s 3D face shape is parameterized based on FLAME model [27] as blendshapes with three linear bases for identity shape, pose, and expression. Their respective coefficients are denoted by β, θ, ψ . Our face shape encoding \mathbf{s} is the combined $(\beta, \theta, \psi) \in \mathbb{R}^{|\beta|+|\theta|+|\psi|}$. DECA assumes orthographic projection and models the camera pose with isotropic scaling and 2D translation. We combine the scaling and translation parameters into $\mathbf{cam} \in \mathbb{R}^{1+2}$. Note that we do not use the predicted albedo map because its estimation by DECA can be unreliable and we found it empirically unnecessary.

Identity encoding (ξ). To compute our deep feature embedding that helps preserve the input’s identity, we use ArcFace[8], a pre-trained face recognition model based on ResNet [18]. This model has been shown to produce discriminative and identity-preserving feature embeddings.

Cast shadow encoding (c). This scalar describes the degree of visible cast shadows, typically caused by a dominant point or directional light source, such as the sun.

We trained a model to estimate c from a face image ourselves and fixed this pretrained estimator. To do this, we manually labeled around 1,000 face images with binary flags indicating whether cast shadows are visible. Following a technique proposed in DiffAE [41], we first use DiffAE’s pretrained encoder to map each face image to a semantically meaningful latent code \mathbf{z} and train a logistic regression classifier on \mathbf{z} to predict the flag. c is then computed as the logit value of the logistic regression. As shown in [41], this technique helps reduce the number of training examples required to achieve good accuracy, but we note that c can be estimated in other ways, such as with a CNN.

Background encoding (\mathbf{bg}). To help fix the background during relighting, we condition the DDIM with an image of the input’s background. The background region is detected using a face segmentation algorithm [73]. The ears, hair, and neck are not part of the background and can be relit by our algorithm (see Figure 7.)

3.4. DDIM decoder & Modulator network

Our main network is a conditional DDIM that decodes our feature vector (with modified lighting information) to a relit version of the input image. In practice, the feature vector is used to *condition* the DDIM that maps $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ to the original input \mathbf{x}_0 during training or maps $\mathbf{x}_T = \text{DDIM}^{-1}(\mathbf{x}_0)$ from reversing the generative process to the relit output during relighting (Section 3.7). This conditioning involves another network called *Modulator* network, which converts the light, shape, and camera encodings into spatial modulation weights for the DDIM decoder.

The architecture of the DDIM decoder is based on Dhariwal et al. [10], which is a modified UNet built from a stack

of residual blocks interleaved with self-attention layers. We provide full details in Appendix C. Our only differences are that 1) the output of each residual block in the first half of the UNet will be modulated by the signal from the Modulator network and 2) we use our own version of adaptive group normalization. Our Modulator network has the same architecture as the first half of our DDIM’s UNet, but they do not share weights.

3.5. Conditioning DDIM decoder

Conditioning a diffusion model on a condition vector can be done in various ways, such as through adaptive group normalization [71, 41, 10] or attention-based mechanisms [36, 47], among others. In our problem, the lighting information is encoded explicitly as SH coefficients and their interaction with 3D shape, specifically the surface normals, can be precisely modeled with the SH lighting equation. Our idea is to ease the modeling of the known interaction by rendering a shading reference of the target relit face. The primary goal of this reference is to convey the information about the target lighting and shading in a spatially-aligned manner, not the geometry or the exact shading intensities. The following sections detail this “spatial” conditioning technique as well as a standard non-spatial conditioning technique used for other encodings.

Spatial conditioning. This technique is used for the light, shape, camera and background encodings ($\mathbf{l}, \mathbf{s}, \mathbf{cam}, \mathbf{bg}$). Given the face shape \mathbf{s} , we first convert it to a triangle mesh using the three linear bases of FLAME model [27] and remove the ears, eyeballs, neck, and scalp from the mesh to retain only the face region (See Figure 2). We remove those parts because they are often inaccurate and hard to estimate correctly (e.g., occluded ears behind hair). We assume a constant gray albedo (0.7, 0.7, 0.7) and render this mesh in the camera pose described by \mathbf{cam} with surface colors computed with \mathbf{l} using the standard SH lighting equation. The details are in Appendix F, and we discuss this albedo choice and the inherent albedo-light ambiguity in Section 5.

Then, this shading reference R , which shows a shaded face in the shape and pose of the input person under the target lighting, is concatenated with the background image \mathbf{bg} and fed to the Modulator network. Let us denote the output of each residual block i in the Modulator network by $\mathbf{m}_i \in \mathbb{R}^{H_i \times W_i \times D_i}$, and the output of the corresponding residual block in the identical DDIM’s first half by $\mathbf{o}_i \in \mathbb{R}^{H_i \times W_i \times D_i}$. In the DDIM, we take each residual block’s output \mathbf{o}_i and replace it with \mathbf{o}'_i , which will be used as input to the subsequent layer in the network:

$$\mathbf{o}'_i = \mathbf{o}_i \odot \tanh(\mathbf{m}_i), \quad (2)$$

where \odot is the element-wise multiplication. This conditioning technique allows the shaded image R and the background to retain their spatial structure and facilitate local

conditioning of the generation as they are spatially aligned with the input (e.g., their facial parts and background are in the same positions).

Non-spatial conditioning. This technique is used for $(\mathbf{s}, \mathbf{cam}, \boldsymbol{\xi}, c)$. The direct use of \mathbf{s}, \mathbf{cam} again in this technique is empirically found to be helpful, in addition to their indirect use through the shading reference. We use a similar conditioning technique as used in [10, 41] based on adaptive group normalization (AdaGN) [71] for these encodings and also for the time embedding in the standard diffusion model training $\gamma(t)$, where γ is a sinusoidal encoding function [10]. Given an input feature map $\mathbf{h}_j \in \mathbb{R}^{H_j \times W_j \times D_j}$, we compute

$$\text{AdaGN}_j(\mathbf{h}_j, \mathbf{s}, \mathbf{cam}, \boldsymbol{\xi}, c, t) = \mathbf{k}_j(\mathbf{t}_j^s \text{GN}(\mathbf{h}_j) + \mathbf{t}_j^b), \quad (3)$$

where $\mathbf{k}_j = \text{MLP}_j^3(\text{Concat}(\mathbf{s}, \mathbf{cam}, \boldsymbol{\xi}, c)) \in \mathbb{R}^{D_j}$ is the output of a 3-layer MLP with the SiLU activation [11], and $(\mathbf{t}_j^s, \mathbf{t}_j^b) \in \mathbb{R}^{2 \times D_j} = \text{MLP}_j^1(\gamma(t))$ is the output from a single-layer MLP also with the SiLU activation. GN is the standard group normalization. We apply our AdaGN in place of all the AdaGNs in the original architecture of [10], which occur throughout the UNet. (Details in Appendix C.)

3.6. Training

We jointly train the DDIM decoder, parameterized as a noise prediction network ϵ_θ , and the Modulator network $M_\phi(\mathbf{l}, \mathbf{s}, \mathbf{cam}, \mathbf{bg})$ using standard diffusion training [19, 59, 41]. Here we consider the MLPs in Figure 2 as part of the DDIM. We adopt the simplified, re-weighted version of the variational lower bound with ϵ parameterization:

$$L_{\text{simple}} = \mathbb{E}_{t, \mathbf{x}_0, \epsilon} \|\epsilon_\theta(\mathbf{x}_t, t, M_\phi, \mathbf{s}, \mathbf{cam}, \boldsymbol{\xi}, c) - \epsilon\|_2^2,$$

where ϵ_θ is trained to predict the added noise $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ in $\mathbf{x}_t = \sqrt{\alpha_t} \mathbf{x}_0 + \sqrt{1 - \alpha_t} \epsilon$, given a training image \mathbf{x}_0 . We define α_t as $\prod_{s=1}^t (1 - \beta_s)$, where β_t is the noise level at timestep t in the Gaussian diffusion process $q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I})$. We use a linear noise schedule and a total step $T = 1000$. Note that we do not reverse $\mathbf{x}_T = \text{DDIM}^{-1}(\mathbf{x}_0)$ during training.

3.7. Relighting

To relight an input image, we first encode the input image into our feature vector \mathbf{f} (Equation 1), then reverse the deterministic generative process of our DDIM conditioned on \mathbf{f} , starting from the input image \mathbf{x}_0 to $\mathbf{x}_{T=1000}$.

$$\mathbf{x}_{t+1} = \sqrt{\alpha_{t+1}} \mathbf{g}_\theta(\mathbf{x}_t, t, \mathbf{f}) + \sqrt{1 - \alpha_{t+1}} \epsilon_\theta(\mathbf{x}_t, t, \mathbf{f}), \quad (4)$$

where \mathbf{g}_θ represents the predicted \mathbf{x}_0 , which is reparameterized from ϵ_θ and is computed by:

$$\mathbf{g}_\theta(\mathbf{x}_t, t, \mathbf{f}) = \frac{1}{\sqrt{\alpha_t}} (\mathbf{x}_t - \sqrt{1 - \alpha_t} \epsilon_\theta(\mathbf{x}_t, t, \mathbf{f})). \quad (5)$$

After obtaining \mathbf{x}_T , we modify the SH light encoding \mathbf{l} and the cast shadow flag c to the target \mathbf{l}' and c' , which can be set manually or inferred from a reference lighting image using DECA and our cast shadow estimator. Then, we decode the modified $\mathbf{f}' = (\mathbf{l}', \mathbf{s}, \mathbf{cam}, \boldsymbol{\xi}, c', \mathbf{bg})$ using the reverse of Equation 4, starting from \mathbf{x}_T to produce the final output.

The reverse process to obtain \mathbf{x}_T is key to reproducing high-frequency details from the input image. As demonstrated in DiffAE [41], DDIM will encode any information not captured in the conditioning feature vector \mathbf{f} in the noise map \mathbf{x}_T . This information includes high-frequency details, such as the hair pattern or skin texture.

Improved DDIM sampling with mean-matching. We observe that when the input image contains a background with extreme intensities (e.g., too dark or too bright), DDIM can produce results with a slight change in the overall brightness. We alleviate this issue by computing the mean pixel difference between each \mathbf{x}_t during DDIM’s generative reversal ($\mathbf{x}_0 \rightarrow \mathbf{x}_T$) and \mathbf{x}_t from self-decoding of the reversed noise \mathbf{x}_T . This sequence of mean differences is then applied to the decoding for relighting (Appendix B).

4. Experiments

In this section, we present quantitative and qualitative results of our proposed method. We provide a comparison of our relighting performance (Section 4.1) to the state of the art on Multi-PIE dataset and ablation studies (Section 4.3) on the non-spatial and light conditioning. Implementation details, dataset details, and runtime are in Appendix B.

Evaluation metrics. We use DSSIM [35], LPIPS [75], and MSE. DSSIM measures the structure dissimilarity, and LPIPS measures the perceptual quality. All metrics are computed between each relit image and its ground-truth image only on the face region following [22, 23] using the same face parsing algorithm [73].

4.1. Relighting performance

We evaluate our relighting performance on Multi-PIE dataset [17] against recent state-of-the-art methods [23, 22, 35, 38]. Note that Pandey et al. [38] solve a different problem setup (also [69, 72]) and require an HDR environment map as input, which has to be first estimated from a target image, making a comparison with [38] not entirely apples-to-apples. The results of [38] in our experiment were generated by the authors themselves, including the HDR maps. Other test sets and code of [38] were not released.

Our experiment has two setups where the target lighting is taken either from **i). the same person**. This setup uses the same test set as [23], which contains 826 testing samples from 329 subjects. Or **ii). a different person**. This setup contains 200 random triplets of input, target, and ground-truth images, where the target image is of a different person.

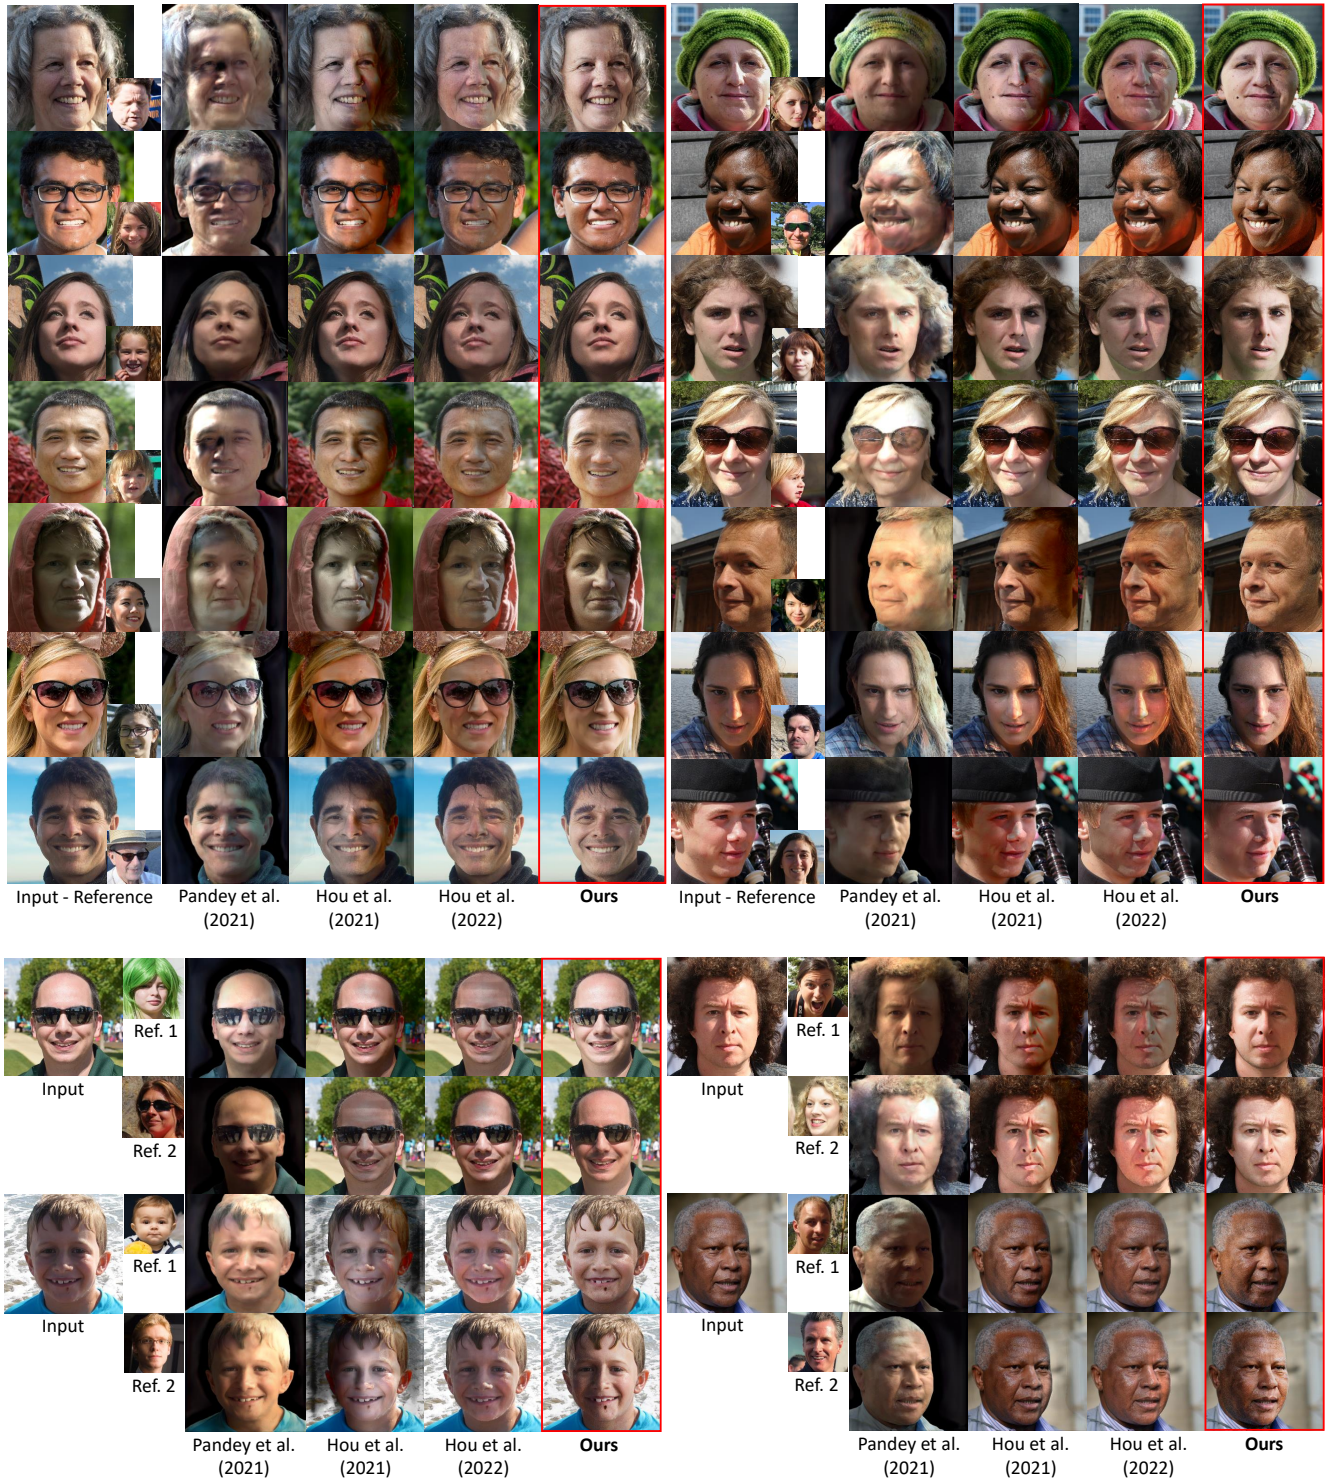


Figure 3: **Relit images on FFHQ [25]**. The FFHQ dataset contains a variety of face images captured in real-world environments. Our method produces more realistic relit images than previous methods and can add or remove cast shadows and highlights to match the reference lighting.

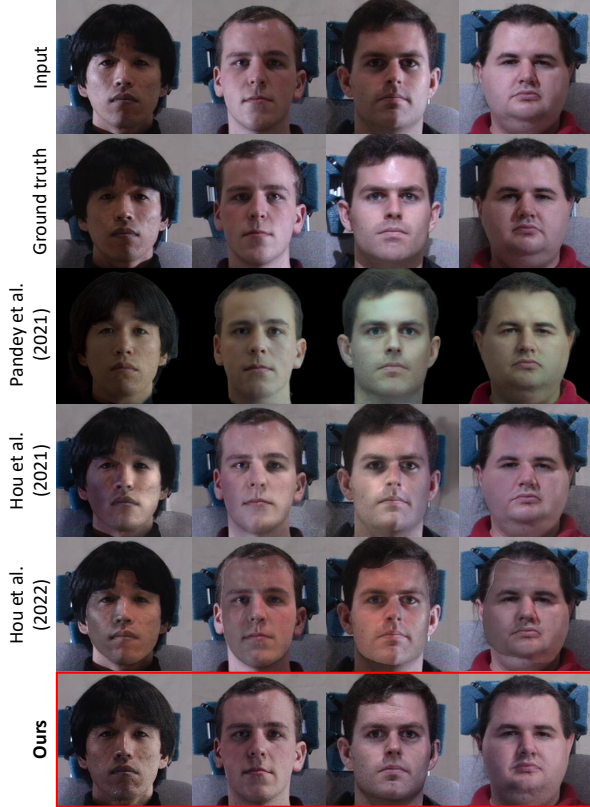


Figure 4: Relighting results on Multi-PIE when the target lighting comes from the same person.

The results are shown in Table 1, Figure 4, and Figure 5. For both setups, our method achieves the best performance across all metrics with minimal artifacts and can convincingly relight the neck and ears or remove cast shadows, e.g., from the nose of the lady in Figure 5. We include a comparison with [53, 77, 61] and more qualitative results of [23, 22, 38] in Appendix E.

Table 1: State-of-the-art comparison on Multi-PIE.

Method	DDSIM↓	MSE↓	LPIPS↓
i). Same subject as target lighting			
Nestmayer et al. [35]	0.2226	0.0588	0.3795
Pandey et al. [38]	0.0875	0.0165	0.2010
Hou et al. (CVPR’21) [23]	0.1186	0.0303	0.2013
Hou et al. (CVPR’22) [22]	0.0990	0.0150	0.1622
Ours	0.0711	0.0122	0.1370
ii). Different subject as target lighting			
Pandey et al. [38]	0.1000	0.0252	0.2053
Hou et al. (CVPR’21) [23]	0.1056	0.0247	0.1989
Hou et al. (CVPR’22) [22]	0.1150	0.0238	0.2215
Ours	0.0969	0.0215	0.1669

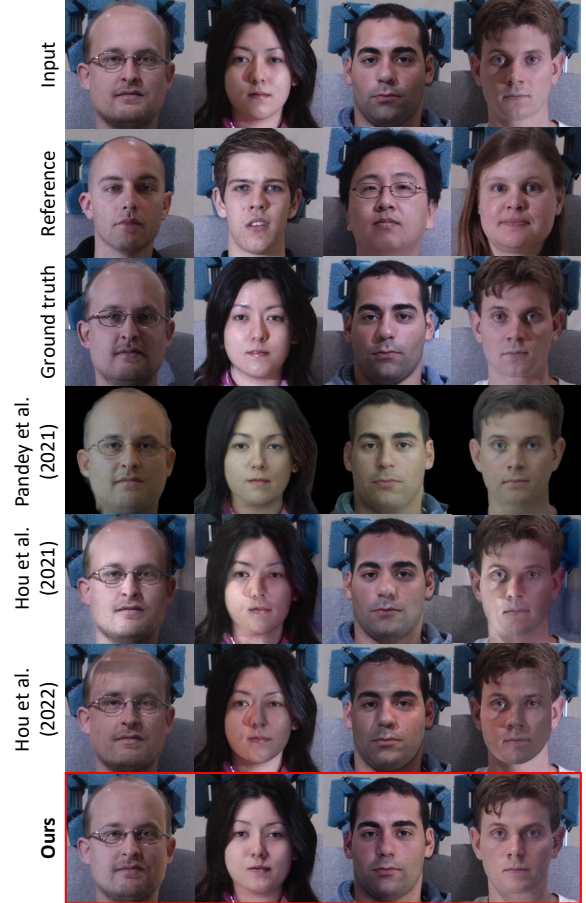


Figure 5: Relighting results on Multi-PIE when the target lighting comes from a different person.

4.2. Qualitative evaluations

Relighting on FFHQ dataset. In Figure 3, we provide a qualitative comparison with three recent SOTAs [38, 23, 22] on subjects with different head poses, genders, races, accessories. Our approach produces highly realistic results and can synthesize new highlights and eliminate hard shadows, while the competing methods often leave behind shadow or shading residuals due to the inaccurate albedo prediction of these in-the-wild images.

Shadow flag condition. We show a novel ability to change the strength of cast shadows on FFHQ [25] in Figure 6. We generate these results by varying the cast shadow’s logit value (c). Our method can realistically remove shadows (e.g., those cast by eyeglasses or face geometry) or intensify their effects. Figure 1 demonstrates how the direction and appearance of cast shadows can change according to a new target lighting condition (the bottom guy’s chin).

4.3. Ablation studies

Light conditioning. We compare our full pipeline against two alternatives for conditioning the DDIM on the



Figure 6: **Varying degrees of cast shadow.** We show the ability to change the degree of cast shadows by adjusting the scalar c and decode the modified feature vector.



Figure 7: Without **bg** conditioning, the background and hat of the input person are not well preserved.

light encoding: a) We do not use our Modulator network and instead feed the reference shading directly to the DDIM by concatenating it with each x_t in every timestep. b) We do not use the shading reference and instead concatenate the light encoding l with (s, cam, ξ, c) in the non-spatial conditioning technique.

We report the results in Table 2 and show a qualitative comparison in Appendix E. Using the light encoding as part of a non-spatial vector (b) performs worst among all three, whereas feeding the shading reference directly to the DDIM without our Modulator (a) improves the results but still lacks behind our proposed pipeline.

Non-spatial conditioning. In this section, we study the benefits of non-spatial, face-related conditions extracted from ArcFace (ξ) and DECA (s, cam) by evaluating the re-light performance on: c) Our method with no s, cam, ξ . d) Our method with no s, cam . e) Our method with no ξ .

We report the results in Table 2 and a qualitative comparison in Appendix E. Removing all of s, cam, ξ performs



Figure 8: **Failure cases.** (Left) Shadows cast by external objects are not relit correctly. (Right) The sunglasses, which resemble cast shadows, are mistakenly removed.

the worst, whereas removing s, cam but retaining ξ obtains a better MSE score. On the other hand, our full pipeline outperforms these alternatives on both DDSIM and LPIPS metrics, which agree with human perception.

Table 2: Ablation study on conditioning methods.

Method	DDSIM↓	MSE↓	LPIPS↓
Light conditioning			
a) No <i>Modulator</i>	0.0749	0.0081	0.0868
b) Used as non-spatial	0.0885	0.0098	0.0947
Ours	0.0670	0.0077	0.0789
Non-spatial condition vector			
c) No s, cam, ξ	0.0713	0.0082	0.0909
d) No s, cam	0.0674	0.0063	0.0846
e) No ξ	0.0686	0.0074	0.0847
Ours	0.0670	0.0077	0.0789

5. Limitations & Discussion

While our results look photorealistic and plausible, some aspects such as cast shadows may not be physically accurate. There is room for improvement in terms of temporal consistency when the result is rendered as a video with moving lights (best observed on our website). Other limitations include: shadows cast by external objects can be incorrectly relit, and sunglasses, which resemble cast shadows, can sometimes be removed (Figure 8). Relighting to match a reference lighting image can be inaccurate as we rely on a light estimator, which is susceptible to the ambiguity where it is unclear if, e.g., a dark appearance is caused by the skin tone or dim lighting. Incorporating physical rendering to help guide the appearance of cast shadows or allowing other lighting models are interesting future work. Our diffusion-based model requires multiple network passes and is currently slower than other GAN-based methods.

In conclusion, we have presented a diffusion-based face relighting method that eliminates the need for accurate intrinsic decomposition and can be trained on 2D images without any 3D or lighting ground truth. Our key component is a conditional diffusion implicit model and a novel conditioning technique that maps a disentangled light representation to a relit image. This enables our method to achieve new state-of-the-art performance and produce highly photorealistic results for real-world scenarios.

References

- [1] Tomer Amit, Eliya Nachmani, Tal Shaharabany, and Lior Wolf. Segdiff: Image segmentation with diffusion probabilistic models. *arXiv preprint arXiv:2112.00390*, 2021. [14](#)
- [2] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*, 2022. [14](#)
- [3] Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khrukov, and Artem Babenko. Label-efficient semantic segmentation with diffusion models. *arXiv preprint arXiv:2112.03126*, 2021. [14](#)
- [4] Jonathan T Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. *IEEE transactions on pattern analysis and machine intelligence*, 37(8):1670–1687, 2014. [1, 3](#)
- [5] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, 1999. [3, 15](#)
- [6] Chen-Hao Chao, Wei-Fang Sun, Bo-Wun Cheng, Yi-Chen Lo, Chia-Che Chang, Yu-Lun Liu, Yu-Lin Chang, Chia-Ping Chen, and Chun-Yi Lee. Denoising likelihood score matching for conditional score-based data generation. *arXiv preprint arXiv:2203.14206*, 2022. [14](#)
- [7] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *arXiv preprint arXiv:2209.04747*, 2022. [14](#)
- [8] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2019. [2, 4, 5, 15](#)
- [9] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. [15](#)
- [10] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. [5, 6, 13, 14, 15](#)
- [11] Stefan Elfving, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107:3–11, 2018. [6](#)
- [12] Qianli Feng, Viraj Shah, Raghudeep Gadde, Pietro Perona, and Aleix Martinez. Near perfect gan inversion. *arXiv preprint arXiv:2202.11833*, 2022. [4](#)
- [13] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3D face model from in-the-wild images. volume 40, 2021. [2, 4](#)
- [14] 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. [15](#)
- [15] Kyle Genova, Forrester Cole, Aaron Maschinot, Aaron Sarna, Daniel Vlasic, and William T Freeman. Unsupervised training for 3d morphable model regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8377–8386, 2018. [15](#)
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. [4](#)
- [17] Ralph Gross, Iain Matthews, Jeffrey Cohn, Takeo Kanade, and Simon Baker. Multi-pie. *Image and vision computing*, 28(5):807–813, 2010. [3, 6, 13](#)
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. [5](#)
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. [4, 6, 14](#)
- [20] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *J. Mach. Learn. Res.*, 23:47–1, 2022. [14](#)
- [21] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. [14](#)
- [22] Andrew Hou, Michel Sarkis, Ning Bi, Yiyong Tong, and Xiaoming Liu. Face relighting with geometrically consistent shadows. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4217–4226, 2022. [1, 3, 6, 8, 13, 15](#)
- [23] Andrew Hou, Ze Zhang, Michel Sarkis, Ning Bi, Yiyong Tong, and Xiaoming Liu. Towards high fidelity face relighting with realistic shadows. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14719–14728, 2021. [1, 3, 6, 8, 15](#)
- [24] Luo Jiang, Juyong Zhang, Bailin Deng, Hao Li, and Ligang Liu. 3d face reconstruction with geometry details from a single image. *IEEE Transactions on Image Processing*, 27(10):4756–4770, 2018. [15](#)
- [25] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. [4, 7, 8, 13, 17, 18](#)
- [26] Ha A Le and Ioannis A Kakadiaris. Illumination-invariant face recognition with deep relit face images. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2146–2155. IEEE, 2019. [1, 3](#)
- [27] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 36(6):194:1–194:17, 2017. [5, 15](#)
- [28] Yijun Li, Ming-Yu Liu, Xueting Li, Ming-Hsuan Yang, and Jan Kautz. A closed-form solution to photorealistic image stylization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 453–468, 2018. [3](#)

- [29] Yue Li, Liqian Ma, Haoqiang Fan, and Kenny Mitchell. Feature-preserving detailed 3d face reconstruction from a single image. In *Proceedings of the 15th ACM SIGGRAPH European Conference on Visual Media Production*, pages 1–9, 2018. 15
- [30] Jingtuo Liu, Yafeng Deng, Tao Bai, Zhengping Wei, and Chang Huang. Targeting ultimate accuracy: Face recognition via deep embedding. *arXiv preprint arXiv:1506.07310*, 2015. 15
- [31] Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala. Deep photo style transfer. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4990–4998, 2017. 3
- [32] BR Mallikarjun, Ayush Tewari, Abdallah Dib, Tim Weyrich, Bernd Bickel, Hans Peter Seidel, Hanspeter Pfister, Wojciech Matusik, Louis Chevallier, Mohamed A Elgharib, et al. Photoapp: Photorealistic appearance editing of head portraits. *ACM Transactions on Graphics*, 40(4), 2021. 4
- [33] Iacopo Masi, Anh Tun Trn, Tal Hassner, Gozde Sahin, and Gérard Medioni. Face-specific data augmentation for unconstrained face recognition. *International Journal of Computer Vision*, 127(6):642–667, 2019. 15
- [34] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2021. 14
- [35] Thomas Nestmeyer, Jean-François Lalonde, Iain Matthews, and Andreas Lehrmann. Learning physics-guided face relighting under directional light. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5124–5133, 2020. 1, 3, 6, 8, 15
- [36] 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. 5, 14
- [37] Kushagra Pandey, Avideep Mukherjee, Piyush Rai, and Abhishek Kumar. Vaes meet diffusion models: Efficient and high-fidelity generation. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. 14
- [38] 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. 1, 3, 6, 8, 14, 15
- [39] Omkar M Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. 2015. 15
- [40] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 14
- [41] Konpat Preechakul, Nattanat Chatthee, Suttisak Wizadwongsa, and Supasorn Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10619–10629, 2022. 2, 4, 5, 6, 14, 15
- [42] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 14
- [43] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020. 14
- [44] 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. 14
- [45] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2287–2296, 2021. 4
- [46] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Transactions on Graphics (TOG)*, 42(1):1–13, 2022. 4
- [47] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. 5
- [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*, pages 10684–10695, 2022. 14
- [49] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. *arXiv preprint arXiv:2208.12242*, 2022. 14
- [50] 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. 14
- [51] 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. 14
- [52] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. 15
- [53] Soumyadip Sengupta, Angjoo Kanazawa, Carlos D Castillo, and David W Jacobs. Sfsnet: Learning shape, reflectance and illuminance of faces in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6296–6305, 2018. 1, 3, 8, 14, 15

- [54] YiChang Shih, Sylvain Paris, Connelly Barnes, William T Freeman, and Frédo Durand. Style transfer for headshot portraits. 2014. **3**
- [55] Zhixin Shu, Sunil Hadap, Eli Shechtman, Kalyan Sunkavalli, Sylvain Paris, and Dimitris Samaras. Portrait lighting transfer using a mass transport approach. *ACM Transactions on Graphics (TOG)*, 36(4):1, 2017. **4, 14**
- [56] Zhixin Shu, Ersin Yumer, Sunil Hadap, Kalyan Sunkavalli, Eli Shechtman, and Dimitris Samaras. Neural face editing with intrinsic image disentangling. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5541–5550, 2017. **1, 3**
- [57] Abhishek Sinha, Jiaming Song, Chenlin Meng, and Stefano Ermon. D2c: Diffusion-decoding models for few-shot conditional generation. *Advances in Neural Information Processing Systems*, 34:12533–12548, 2021. **14**
- [58] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR. **4, 14**
- [59] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. **2, 4, 6**
- [60] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. **4, 14**
- [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, 2019. **3, 8, 14, 15**
- [62] Yaniv Taigman, Ming Yang, Marc Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1701–1708, 2014. **15**
- [63] Ayush Tewari, Mohamed Elgharib, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhöfer, and Christian Theobalt. Pie: Portrait image embedding for semantic control. *ACM Transactions on Graphics (TOG)*, 39(6):1–14, 2020. **4**
- [64] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zollhofer, and Christian Theobalt. Stylerig: Rigging stylegan for 3d control over portrait images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6142–6151, 2020. **4**
- [65] Ayush Tewari, Tae-Hyun Oh, Tim Weyrich, Bernd Bickel, Hans-Peter Seidel, Hanspeter Pfister, Wojciech Matusik, Mohamed Elgharib, Christian Theobalt, et al. Monocular reconstruction of neural face reflectance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4791–4800, 2021. **3**
- [66] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. **14**
- [67] Mei Wang and Weihong Deng. Deep face recognition: A survey. *Neurocomputing*, 429:215–244, 2021. **15**
- [68] Yang Wang, Lei Zhang, Zicheng Liu, Gang Hua, Zhen Wen, Zhengyou Zhang, and Dimitris Samaras. Face relighting from a single image under arbitrary unknown lighting conditions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(11):1968–1984, 2008. **1, 3**
- [69] Zhibo Wang, Xin Yu, Ming Lu, Quan Wang, Chen Qian, and Feng Xu. Single image portrait relighting via explicit multiple reflectance channel modeling. *ACM Transactions on Graphics (TOG)*, 39(6):1–13, 2020. **3, 6**
- [70] Andreas Wenger, Andrew Gardner, Chris Tchou, Jonas Unger, Tim Hawkins, and Paul Debevec. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Transactions on Graphics (TOG)*, 24(3):756–764, 2005. **3**
- [71] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. **5, 6**
- [72] 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 Transactions on Graphics (TOG)*, 41(6):1–21, 2022. **1, 3, 6, 14**
- [73] Changqian Yu, Jingbo Wang, Chao Peng, Changxin Gao, Gang Yu, and Nong Sang. Bisenet: Bilateral segmentation network for real-time semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 325–341, 2018. **5, 6**
- [74] Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*, 2023. **14**
- [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*, pages 586–595, 2018. **6**
- [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. **14**
- [77] 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. **3, 8, 15**