
ID-to-3D: Expressive ID-guided 3D Heads via Score Distillation Sampling

Francesca Babiloni, Alexandros Lattas, Jiankang Deng, Stefanos Zafeiriou
Imperial College London, UK
<https://idto3d.github.io>

Abstract

We propose ID-to-3D, a method to generate identity- and text-guided 3D human heads with disentangled expressions, starting from even a single casually captured in-the-wild image of a subject. The foundation of our approach is anchored in compositionality, alongside the use of task-specific 2D diffusion models as priors for optimization. First, we extend a foundational model with a lightweight expression-aware and ID-aware architecture, and create 2D priors for geometry and texture generation, via fine-tuning only 0.2% of its available training parameters. Then, we jointly leverage a neural parametric representation for the expressions of each subject and a multi-stage generation of highly detailed geometry and albedo texture. This combination of strong face identity embeddings and our neural representation enables accurate reconstruction of not only facial features but also accessories and hair and can be meshed to provide render-ready assets for gaming and telepresence. Our results achieve an unprecedented level of identity-consistent and high-quality texture and geometry generation, generalizing to a “world” of unseen 3D identities, without relying on large 3D captured datasets of human assets.



Figure 1: ID-to-3D leverages identity conditioning and score distillation sampling on large diffusion models, achieving high-quality 3D human asset generation from “in-the-wild” images, without training on large scanned datasets. From left to right: a) renderings, b) input images, c) normals.

1 Introduction

The remarkable ability of humans to discern facial characteristics and emotional cues in others makes the development of high-quality 3D head avatars a challenging yet foundational task for a diverse array of emerging applications, including digital telepresence, game character generation, and the creation of virtual and augmented reality experiences. However, the acquisition of such 3D human assets remains a daunting task, that requires either manual work typically performed by graphic artists, or expensive and laborious scanning. High-quality 3D facial scanning, achieved initially

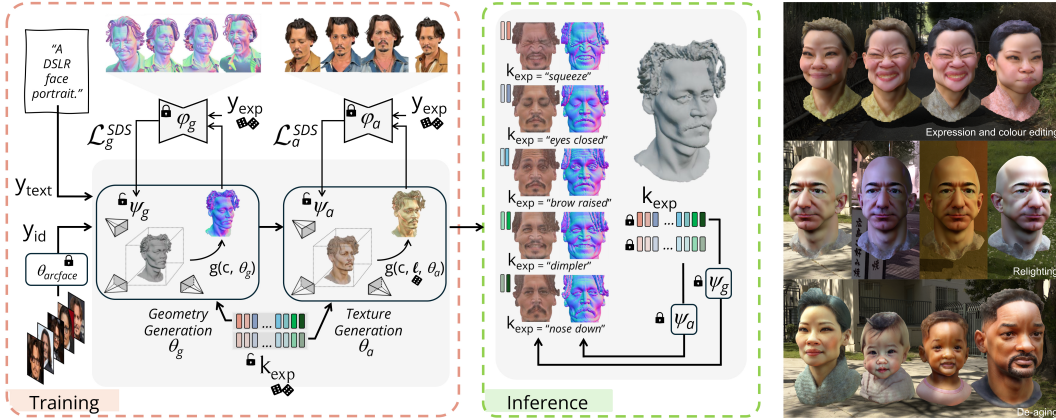


Figure 2: **(Left) Overall pipeline.** ID-to-3D generates expressive 3D head avatars via ArcFace y_{id} and textual y_{text} conditioning. It uses as prior geometry-oriented ψ_g and albedo oriented ϕ_a pretrained models. **Training** The training phase uses SDS to optimize 3D geometry ψ_g , texture ψ_a , and a set of expressions latent codes k_{exp} . It also leverages random lighting l and random expression conditioning y_{exp} . **Inference** At deployment time, ID-to-3D extracts high-quality identity-aware expressive 3D meshes. **(Right) ID-consistent expressive 3D heads** generated by our method. ID-to-3D creates 3D assets that support relighting, ID-consistent editing, and physical simulation.

by hardware with controlled polarized illumination [15, 51, 23], has been simplified tremendously using simpler devices, color-space methods, and inverse rendering [25, 39, 3]. Nevertheless, the requirements for expert knowledge, hardware, and computational time do not allow wide adoption or mass usage.

Consequently, statistical modeling and the rise of deep learning investigated techniques to reconstruct 3D face assets from casually captured images of a subject in arbitrary poses, lighting conditions, and occlusions (referred to as “in-the-wild”). 3D Morphable Models (3DMMs) and Generative Adversarial Networks (GANs) can be used to model facial geometry [4, 19, 56, 24], while GANs and diffusion models have achieved state-of-the-art modeling and reconstruction of facial appearance [40, 22, 92, 60].

However, a common denominator in all the above is the requirement for vast datasets of scanned facial shapes and appearances, in order to avoid limited generalization, ethnicity under-representation, and oversmoothed geometries, requiring up to 10,000 scans for 3DMMs training [6] or facial appearance modeling [22]. Moreover, the requirement for aligned data in 3D, and also in UV texture maps, limits the utilized area to the facial region, and introduces registration errors and additional costs. Our proposed method bypasses these significant issues by utilizing large 2D generative models, pre-trained on vast and easy-to-acquire 2D images, and using only a small set of rasterized 3D data for finetuning, while achieving state-of-the-art 3D head generation.

More recently, the use of Score Distillation Sampling (SDS) [63] and large-scale diffusion models explored the automatic generation of 3D content thanks to text [85, 10, 45, 44, 31, 92, 37, 26] or image prompts [66, 91, 64, 79, 80]. Despite promising results, leveraging 2D priors to generate realistic 3D head avatars remains challenging: **(1)** Large text-to-image foundational models are usually trained to generate realistic RGB images, where geometry, texture, and lighting cannot be individually segregated. This narrow “rendering knowledge” in the 2D guidance compromises the 3D fidelity, texture quality, and consistency of the generated assets, leading to unrealistic geometries and distorted textures with recognizable artifacts such as Janus problems, incorrect proportions, oversaturated albedo, and mismatching texture and geometry details. **(2)** It is difficult to precisely control facial attributes solely using typical prompting methods [13, 59, 88, 84]. Textual prompts lack the granularity to single out the specificity of a subject’s identity and facial expression that might be lengthy and complex to convey in natural language. Moreover, methods that leverage image prompts lack the capability to capture features that represent the facial identities of a subject independently of pose, expression, or contextual scene information (i.e., ID embeddings). The lack of reliable control over identity and facial expression prevents the personalization of 3D head avatars, resulting in a limited range of possible output expressions or custom attributes without incurring identity drifting.

In this work, we present ID-to-3D, a new method to generate expressive and identity-consistent 3D human heads using a text prompt and a small set of 1-5 casually captured, in-the-wild, images of a subject. ID-to-3D creates a variety of separated yet ID-consistent expressions in a single optimization, leveraging as priors compositionality and task-specific 2D diffusion models. To ensure consistency, the identity of each subject is encoded via facial embedding features and, to encourage expressivity, emotions are disentangled via a novel ID-specific neural parametric representation. The generation of each 3D asset leverages score distillation sampling and a two-stage pipeline to create shape details, texture features, and materials. During optimization, the guidance is given by a foundational model extended into an a) identity-aware, b) task-aware and c) expression-aware variant, working as 2D prior for either geometry or texture generation. Each stage guidance is trained only once for all the potential identities, with a lightweight fine-tuning strategy changing only 0.2% of all the available training parameters.

Overall, we present the following contributions: **(1) ArcFace-Conditioned 3D Head Asset Generation:** we introduce the first method for arcface-conditioned generation of 3D head assets using SDS. **(2) Novel ID-Conditioned Expressive Model:** our model creates an identity-conditioned expressive representation for each subject, enabling the generation of up to 13 unique and ID-consistent expressions captured by latent codes and associated with a set of 3D assets with separate geometric, albedo, and material information. **(3) Novel Text-to-2D-Normals and Text-to-2D-Albedo Models:** we present a novel approach to creating ID-conditioned and expression-conditioned text-to-image models capable of generating realistically plausible normals and albedo images from a small set of 3D assets. Extensive experiments confirm that our method outperforms text-based and image-based SDS baselines by producing relight-able 3D head assets with unprecedented geometric details, superior texture quality, and able to exhibit a wide variety of expressions. Fig. 1 displays 3D heads generated with ID-to-3D.

2 Related Work

3D Human Generation and Reconstruction. Human modelling and reconstruction in the 3D domain is typically based on 3D Morphable Models (3DMMs) [4, 6, 43, 78], which model variations in human shape and appearance using PCA. The advent of deep learning enabled very expressive reconstructions [14, 68], but also extended this line of research beyond linear spaces. Neural Parametric Head models [24] explored the use of signed distance function and deformation field to generate expressive geometries. Generative Adversarial Networks have also been proposed to generate [42, 21] and reconstruct faces from “in-the-wild” images [22, 50, 40]. Diffusion models showed impressive results in modeling skin textures [92, 60] especially when paired with large-scale, high-quality datasets of real scans.

Text-to-3D Human Generation. The creation of 3D human assets via text conditioning has seen significant progress, building on the foundations laid by advances in text-to-2D generation [65, 69, 72, 67]. Initial attempts [86, 29, 54, 73, 55, 32, 11] utilized the CLIP language model to optimize implicit or explicit 3D representations. The seminal work of DreamFusion [63] introduced the Score Distillation Sampling (SDS) loss, which uses a pre-trained 2D diffusion as prior for 3D generation. This work led to a revolution in text-to-3D generation [53, 57, 27, 85, 45, 48, 52, 47, 12] and text-to-3D human generation [9, 76, 45, 53, 26, 46, 49]. DreamAvatar [8], TADA [44] and Headevolver [82] build upon the use of a template [61, 43, 5] to create 3D human avatars with controllable shapes and poses. Fantasia3D [10] separated geometry and texture training, employing DMTET [77] and PBR texture [58] for the 3D representation. HumanNorm [31] introduced the idea of training a normal diffusion model to guide high-quality geometry generation. Despite promising results, these models face challenges in generating highly detailed and expressive 3D heads, due to the inherent limitations of natural language conditioning.

Personalized Generation with Diffusion Models. Controllable generation is crucial to develop widely applicable generative models. Work in this domain includes the customization of GANs [74, 33, 35, 36, 34] and research dedicated to steer the generation of large-scale foundation models with additional control signals [20, 70, 30, 93, 88, 71]. The use of ID embeddings as an alternative to text prompts showed promising results in 2D generation [62, 87, 84, 59], but remains underexplored in the creation of 3D human avatars with SDS. Related research supplements traditional text prompts in SDS pipelines with image-based prompts. DreamBooth3D [66] and Avatarbooth [91] proposed to train an ad hoc model to use as a guide in the generation of 3D objects or avatars. Magic123 [64] incorporate 3D and 2D priors in SDS generation. DreamCraft3D [79] proposed a hierarchical 3D

content pipeline to generate textured 3D meshes from a single unposed image. Despite encouraging results, the creation of high-quality head assets with these methods remains challenging, due to the difficulty of extracting appropriate facial features using only naive image prompting.

3 ID-to-3D

We propose a novel method for identity-driven human head generation, which utilizes a pre-trained 2D model to distill expressive head avatars with high geometric details and high-fidelity textures, avoiding the need for large-scale training on 3D datasets. As illustrated in Figure 2, starting from a subject’s identity embeddings, our method trains a set of latent expression representations using a two-stage SDS pipeline. After convergence, the learned 3D representation can be used to create ID-driven expressive heads, that are ready to be used in common rendering engines.

3.1 3D Head Optimization Objective

A Score Distillation Sampling generation pipeline optimizes a 3D representation θ using a pre-trained 2D diffusion model ϕ as guidance. The pipeline optimization objective is to align the distribution of 3D asset renderings with the target distribution $p(\mathbf{x}_0|y_{\text{text}})$, created by the 2D diffusion model conditioned on an input text y_{text} . Given the distribution of renderings under various camera conditions $q^\theta(\mathbf{x}_0|y_{\text{text}}) = \int q^\theta(\mathbf{x}_0|y_{\text{text}}, \mathbf{c})p(\mathbf{c})d\mathbf{c}$, the optimization objective reads:

$$\min_{\theta} D_{KL}(q^\theta(\mathbf{x}_0|y_{\text{text}}) \parallel p(\mathbf{x}_0|y_{\text{text}})). \quad (1)$$

The target distribution $p(\mathbf{x}_0|y_{\text{text}})$ is typically estimated by a foundational text-to-image diffusion model that approximates the distribution of natural RGB images [69] (i.e., trained over large and uncurated datasets such as LAION-5B [75]). Despite its indisputable success in creating a variety of assets, using the above objective and guidance to generate detailed 3D heads remains a complicated task. First, this target distribution might drift significantly from the distribution of natural heads rendered in realistic light and camera conditions. Second, a general one-shot guidance model does not have explicit ways to differentiate texture and geometric characteristics, which makes the creation of light-independent texture and high-quality geometry extremely challenging. Third, the naive use of text prompts limits the control over the generated head assets, since textual prompts cannot easily or exhaustively capture facial and expression features. In our pipeline’s generation, we decompose the Obj. 1 into two smaller and more controllable objectives:

$$\begin{aligned} \min_{\theta_g, \theta_a} & \underbrace{D_{KL}(q^{\theta_g}(\mathbf{z}_0^n|\mathbf{c}, y_{\text{text}}, y_{\text{exp}}, y_{\text{id}}) \parallel p(\mathbf{z}_0^n|\mathbf{c}, y_{\text{text}}, y_{\text{exp}}, y_{\text{id}}))}_{\text{geometry generation objective}} \\ & + \underbrace{D_{KL}(q^{\theta_a}(\mathbf{z}_0^a|\mathbf{c}, \mathbf{l}, y_{\text{text}}, y_{\text{exp}}, y_{\text{id}}) \parallel p(\mathbf{z}_0^a|\mathbf{c}, y_{\text{text}}, y_{\text{exp}}, y_{\text{id}}))}_{\text{texture generation objective}}. \end{aligned} \quad (2)$$

In the above equation, θ_g represents the parameterization of the 3D geometry, \mathbf{z}_0^n denotes normal maps, θ_a denotes the parameterization of the 3D textures and \mathbf{z}_0^a albedo textures. Conditioning is introduced in the form of textual prompt y_{text} , identity condition y_{id} , and expression condition y_{exp} . The letter \mathbf{l} denotes the lighting condition of the rendered image. The target distributions for geometry and texture generation refer to the ideal distribution of head-specific normal and texture maps, which are in practice estimated via geometry-oriented and albedo-oriented models guided by face-specific conditioning.

3.2 2D Guidance

To initiate the 3D head reconstruction, we start from the development of 2D priors capable of accurately separating texture and geometric details while, at the same time, consistently representing the facial characteristic of a subject under various expressions, conveying different emotional states. The difficulty of this task lies in its nuanced nature, exacerbated by the lack of large-scale 3D human scan datasets, which makes the capture of detailed face features and the generalization to new identities particularly difficult when training from scratch or even when naively fine-tuning from a large-scale model. To solve this challenge, we propose to explicitly model geometry and appearance domains, identity conditioning y_{id} and expression conditioning y_{exp} , achieving a modular separation

of otherwise entangled-together information. To overcome the need for a large-scale dataset, we leverage a small dataset of human heads with different expressions (NHMP) [24], a pre-trained stable diffusion model (SD) [69], and a selective fine-tuning strategy that affects only a minimal number of parameters, needed to accommodate these new conditionings. In practice, we use rasterized normals as a 2D proxy for geometric information and rasterized albedo as a representation of appearance information. We treat the shift from the natural image distribution towards normal maps and albedo textures as “style-transfer” tasks, aiming to leave the content of the SD features unchanged while modifying their self-similarity information. We use Low-Rank adaptation matrices (LoRA) [30] to adjust Query \mathbf{Q} , Key \mathbf{K} and Value \mathbf{V} features of the self-attention to work in the adjusted normal and albedo domains. The normal-adapted self-attention equation becomes the following:

$$\mathbf{Z}_{\text{SA}}^n = \text{Att}(\mathbf{Q}^n, \mathbf{K}^n, \mathbf{V}^n), \mathbf{Q}^n = \mathbf{X}\mathbf{W}_Q + \mathbf{X}\mathbf{W}_Q^n, \mathbf{K}^n = \mathbf{X}\mathbf{W}_K + \mathbf{X}\mathbf{W}_K^n, \mathbf{V}^n = \mathbf{X}\mathbf{W}_V + \mathbf{X}\mathbf{W}_V^n \quad (3)$$

while the albedo-adapted self-attention can be read by changing the superscript n to a . As identity representation we select the identity embeddings \mathbf{y}_{id} , from a state-of-the-art face recognition network [17, 16, 94], a compact vector of facial features extracted from “in-the-wild” images of a subject. As expression conditioning, we use \mathbf{y}_{exp} CLIP embeddings [65] for the textual descriptor of the 23 FACS coded expressions proposed in FaceWareHouse [7]. To ensure control over the generated head and face representation during deployment, we treat the integration of identity and expression information in the baseline architecture as “multimodal conditioning”, by including their contribution in the SD cross-attention layers via IP-Adapter [88] strategy:

$$\mathbf{Z}_{\text{CA}}^n = \text{Att}(\mathbf{Q}, \mathbf{K}^{\text{text}}, \mathbf{V}^{\text{text}}) + \lambda_{\text{id}} \cdot \text{Att}(\mathbf{Q}, \mathbf{K}^{\text{id}}, \mathbf{V}^{\text{id}}) + \lambda_{\text{exp}} \cdot \text{Att}(\mathbf{Q}, \mathbf{K}^{\text{exp}}, \mathbf{V}^{\text{exp}}) \quad (4)$$

In the above equation, λ_{id} and λ_{exp} control the contribution of identity and expression conditioning. The $\mathbf{Q}=\mathbf{x}\mathbf{w}_Q$ term represents the Query extracted from SD features, $\mathbf{K}^{\text{exp}}=\mathbf{y}_{\text{exp}}\mathbf{w}_K^{\text{exp}}$ and $\mathbf{V}^{\text{exp}}=\mathbf{y}_{\text{exp}}\mathbf{w}_V^{\text{exp}}$ the Key and Values extracted from the expression embedding, and $\mathbf{K}^{\text{id}}=\mathbf{y}_{\text{id}}\mathbf{w}_K^{\text{id}}$ and $\mathbf{V}^{\text{id}}=\mathbf{y}_{\text{id}}\mathbf{w}_V^{\text{id}}$ the Key and Values extracted from the identity embedding. We train only our additional parameters, leaving the rest of the model frozen, targeting 0.2% of the overall trainable parameters. We separately optimize a 2D prior for geometry ϕ_g and textures ϕ_a , using image-conditioning pairs created from the renders of the NHMP dataset under various camera poses \mathbf{c} . The training objective for the geometric 2D prior follows the same training objective as a traditional SD model [69]:

$$L_{\text{simple}} = \mathbb{E}_{\mathbf{z}_0^n, \epsilon, \mathbf{c}, t, \mathbf{y}_{\text{text}}, \mathbf{y}_{\text{id}}, \mathbf{y}_{\text{exp}}} \|\epsilon - \epsilon_{\phi_g}(\mathbf{z}_t^n, t, \mathbf{c}, \mathbf{y}_{\text{text}}, \mathbf{y}_{\text{id}}, \mathbf{y}_{\text{exp}})\|^2. \quad (5)$$

while the analogous training objective for the 2D texture prior can be derived by changing both superscript n and subscript g to a . Note that in Equation 5, \mathbf{y}_{exp} indicates CLIP embeddings for the textual descriptor of the renders (i.e., camera view and subject’s attributes) and does not convey face-identity. \mathbf{z}_t^n indicates \mathbf{z}_0^n noised at timestep t . Examples of generated prompt-to-images can be seen in Figure 2 as well as additional materials.

3.3 Geometry Generation

We represent an identity-specific geometry as a neural parametric head model composed of a deep marching tetrahedra (DMTET [77]) representation, additionally coupled with a set of identity-dependent facial expression latent codes. Our lightweight geometric representation produces highly detailed geometry and a broad range of expressions without notable identity drift, as shown by our experiments. Compared to explicit template-based approaches [92, 44, 43], it has the flexibility to dynamically modulate the local resolution of the mesh to capture high-frequency geometrical details. In other words, it adjusts the mesh resolution of specific regions of the face to adapt to the given subject and expression.

The DMTET geometry representation uses a deformable tetrahedral grid Γ and a network Ψ_g to generate a 3D asset [77]. We extend DMTET to learn multiple expressions at the same time. We model each facial expression with a learnable latent code $\mathbf{k}_{\text{exp}}^n \in \mathbb{R}^{d_{\text{exp}}}$ and design the network $\Psi_g(\Gamma, \mathbf{k}_{\text{exp}}^n)$ as a Transformer [81], parameterized by ψ_g learnable parameters, that processes both the deformable grid and the expression information. During the training phase, we randomly select one of the potential expressions, estimate the signed distance function (SDF) for the underlying head, and use a differentiable marching tetrahedra layer to convert the implicit representation into the explicit surface mesh for that ID and expression, compelling the Ψ_g model to learn a diverse set of expressions that are consistent with the identity at hand. We supervise optimization through an

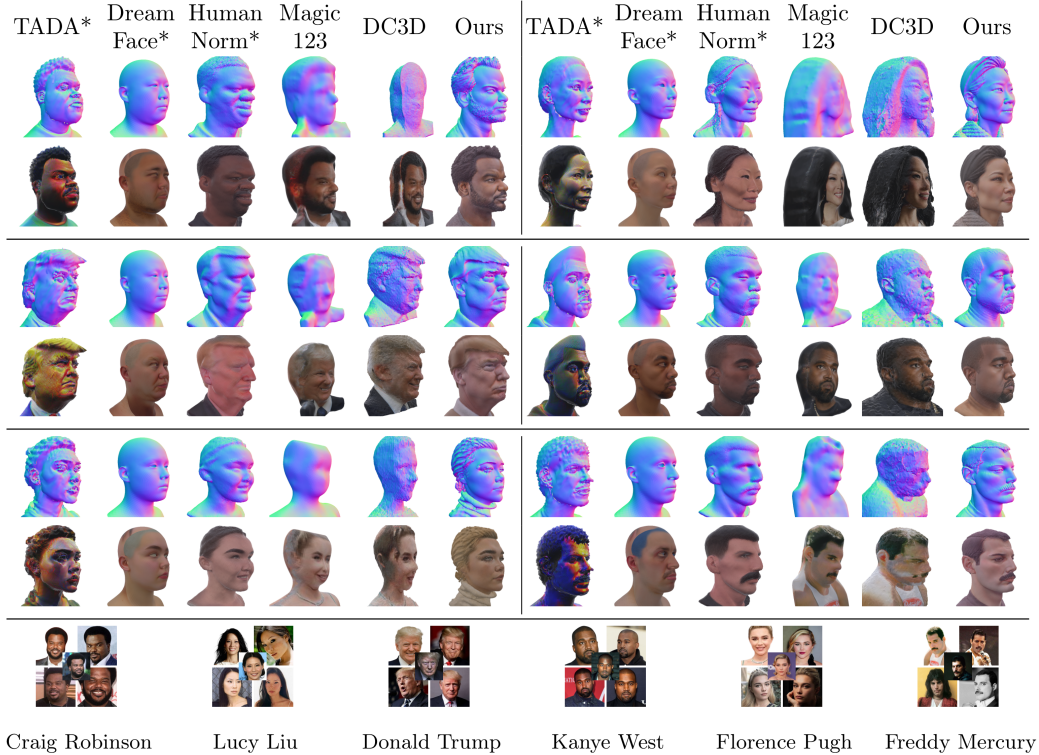


Figure 3: **Qualitative results for text-to-3D (*) and image-to-3D methods.** Methods are evaluated under the same text prompt and rendering conditions. DreamCraft3D is reported as DC3D. Geometry is displayed via normal maps in camera coordinates. Using only a small set of 5 images as conditioning, ID-to-3D achieves high geometric quality and realistic textures.

SDS loss [63], computed using the rasterized geometry model ϕ_g as 2D prior. We optimize the 3D representation θ_g as follows:

$$\nabla \mathcal{L}_{SDS}(\theta_g) = \mathbb{E}_{\mathbf{c}, t, \epsilon} \left[\omega(t) (\epsilon_{\phi_g}(\mathbf{z}_t^n, \mathbf{y}_{id}, \mathbf{y}_{exp}, \mathbf{y}_{text}, t) - \epsilon) \frac{\partial g(\theta_g, \mathbf{c})}{\partial \theta_g} \right]; \theta_g = [\mathbf{k}_{exp}^n, \psi_g]. \quad (6)$$

In the above equation, $g(\theta_g, \mathbf{c})$ represents the normals of the rendered image, created using the differentiable render g and the camera pose \mathbf{c} , and \mathbf{k}_{exp}^n is the randomly sampled expression code. Lastly, we combine the SDS loss with a Laplacian regularizer, to encourage smooth surfaces.

3.4 Texture Generation

Given a trained geometry model θ_g , we model the texture appearance θ_a . To ensure ID-aligned expression generation, we follow an analogous parameterization for the expressions in the texture domain. We represent an identity-specific appearance as a neural parametric head model composed of a pseudo-albedo prediction network Ψ_a , coupled with a set of ID-dependent facial expression latent codes. We instantiate each expression as a learnable latent code $\mathbf{k}_{exp}^a \in R^{d_{exp}}$ and model $\Psi_a(\mathbf{k}_{exp}^a)$ as a Transformer trained to predict the spatially-varying reflectance in a UV-map representation [89], namely the albedo, roughness and specularity, for each ID and expression-specific texture. We use an off-the-shelf physically-based renderer [38]. Reflectance disentanglement is an ill-posed problem, and in the absence of prior data, we use camera and illumination randomization as a regularization constraint [18, 41]. On each iteration, we sample random environment illumination maps, augmented with random Y-axis rotations. During training, we randomize the sampling of the latent expression and deploy the albedo model ϕ_a SDS loss to optimize θ_a :

$$\nabla \mathcal{L}_{SDS}(\theta_a) = \mathbb{E}_{\mathbf{c}, t, \epsilon} \left[\omega(t) (\epsilon_{\phi_a}(\mathbf{z}_t^a, \mathbf{y}_{id}, \mathbf{y}_{exp}, \mathbf{y}_{text}, t) - \epsilon) \frac{\partial g(\theta_a, \mathbf{c}, \mathbf{l})}{\partial \theta_a} \right]; \theta_a = [\mathbf{k}_{exp}^a, \psi_a]. \quad (7)$$




Figure 4: **(Left) Identity Similarity Distribution** between “in-the-wild” images and renderings of 3D heads. **(Right) Comparative Preference Survey** on texture quality (outside) and geometry quality (in). We report % of preferences.

	Texture ↓	Geometry ↓
Fantasia3D* [10]	252	159
DreamFace* [92]	188	145
TADA* [44]	215	118
HumanNorm* [31]	164	98
Magic123 [64]	162	180
DreamCraft3D [79]	205	165
Ours	153	85

Table 1: **Fréchet Inception Distance** measures the geometric and texture quality for the generated 3D heads. “*” indicates text-to-3D SDS pipelines.

where $g(\theta_a, \mathbf{c}, \mathbf{l})$ represents the pseudo-albedo maps created using a differentiable render g and a sample camera pose and lighting condition.

4 Experiments

We assess the efficacy of ID-to-3D as a specialized method for ID-driven expressive human face generation in different scenarios and report comparative analysis against state-of-the-art text-to-3D and image-to-3D generation pipelines. Further analysis, implementation details and discussion can be found in additional material.

4.1 Identity Generation

We benchmark ID-to-3D against several state-of-the-art methods in the domain of SDS-based 3D face asset generation. Specifically, we consider Fantasia3D [10], three recent and similar methods specialized in text-to-human avatar generation (i.e., Human-Norm [31], TADA [44], DreamFace [92]) and two methods that leverage both text and images to create 3D assets (i.e., Magic123 [64], DreamCraft3D [79]). In order to compare together text-to-3D and image-to-3D methods, we select a test benchmark of 40 celebrity names, suggested by ChatGPT, covering actors, sports, and media personalities. We automatically download 25 images for each identity from BingImages [1]: 5 images to use as input and 20 to use as references for our comparisons. For all methods, we use the same textual prompt “a DSLR face portrait of...” and the same input images.

Qualitative Comparisons. Results of existing methods are reported in Figure 3 under the same lighting and rendering conditions. The 3D assets created by ID-to-3D show realistic texture, sharp fine-grained details, and ID fidelity, capturing facial characteristics of the input identity without relying heavily on often ambiguous text prompts or naively lifting 2D images in the 3D domain.

Quantitative Comparisons: Identity Similarity Distribution. We perform a quantitative evaluation on all the evaluated models using similarity ID. We assess the ID fidelity of the 3D assets using the CosFace similarity metric [83, 40]. We center and align the 3D objects and collect renders in a wide range of camera positions (i.e. elevation $[-15^\circ, +15^\circ]$ and rotation $[-65^\circ, +65^\circ]$). We measure the cosine similarity between the ID features of each render and a set of 20 in-the-wild images of the same identity, used as reference. We report the distribution of the identity similarity for each method in Figure 4 (Left). Note that the variance of the distribution correlates with the ID consistency of the 3D object across viewpoints. Despite being able to generate realistic skin textures, DreamFace cannot create hair or eyes by design, resulting in the lowest average similarity score. DreamCraft3D and Magic123 leverage the input image to achieve a photorealistic front-facing camera render, but struggle to create 3D consistent heads, resulting in distributions associated with the highest variance. As clearly visible from Figure 4, ID-to-3D is capable of creating realistic and consistent 3D heads, reporting the lowest variance and the highest similarity score.

Quantitative Comparisons: FID and User-Study. The evaluation of the generated 3D geometries and textures is performed using the Fréchet Inception Distance (FID) metric [28]. For texture quality, the FID is calculated between the renderings and the images from Stable Diffusion V1.5 [69]. For geometry quality, FID is determined by comparing normal maps with those extracted from the NHPM test set [24]. To further substantiate our analysis, we conduct two user preference surveys, comparing our method with the 4 strongest baselines for the generation of texture and geometry. We compare

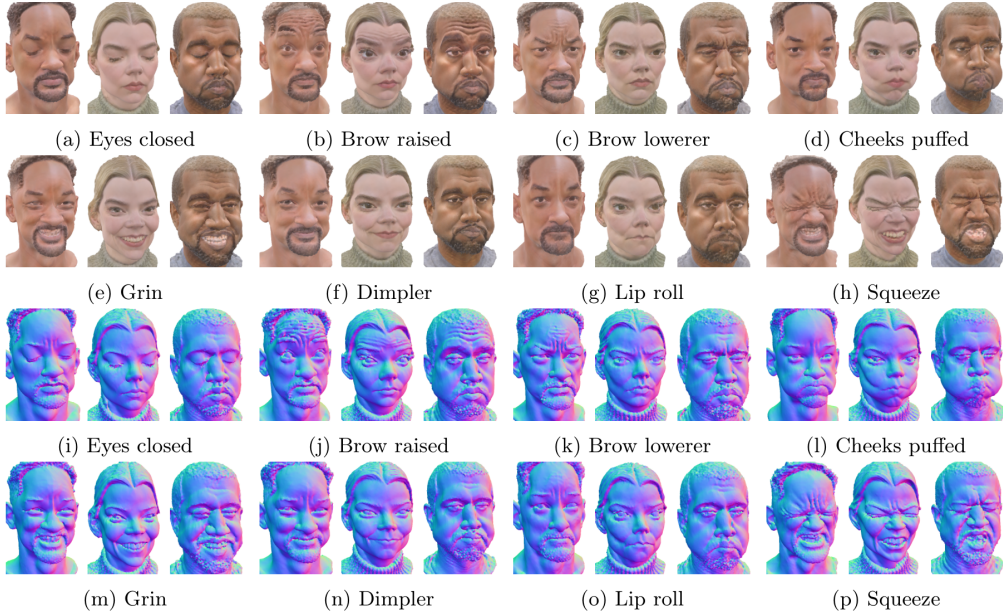


Figure 5: **ID-to-3D expression diversity.** Renderings and normal maps in camera coordinates are taken for 3 identities: *Will Smith*, *Anya Taylor Joy*, and *Kanye West*. Our method achieves fine-grained geometry carving and high-quality texture generation, realistically reproducing various skin tones.

the best 2 performing text-to-3D and image-to-3D methods against ID-to-3D in the user survey, as this enabled us to gather more responses. As shown in Table 1 and Figure 4 (Right), FID metrics and user evaluations report aligned results. Our model achieves the lowest FID scores and the highest user preference in both geometry and texture generation, showcasing together the overall stronger performance of ID-to-3D as a human-specific geometry and texture generator.

4.2 Expressive ID-conditioned Generation

ID-to-3D is specifically designed to create complex, uncommon, and subtle expressions with a level of details not previously achievable using existing SDS methods. In Figure 5 we showcase the unique ability of our method to generate a wide range of expressions that remain identifiable and yet identity-consistent.

Quantitative Analysis: Identity Similarity Distribution Across Expressions. We test the ID consistency across views and expressions by reporting the cosine similarity between the ID features of a reference render (i.e., neutral expression and front-facing camera) and the ID features captured across the remaining expressions and view points. Figure 6 (Right) shows the distribution of ID similarity computed for a set of 10 subjects. ID-to-3D consistently produces 3D assets with high ID similarity and low variance.

Quantitative Analysis: Expression Variety Visualization. A core characteristic of ID-to-3D is its robust ability to produce a wide range of unique and vivid expressions. In this section, we provide a visualization of this expression diversity. We select a subset of 13 expressions and 10 subjects, extracting for each 3D head a set of renders in a range of 9 camera poses. Then we extract the ID features for each render and project them into 2D points using t-SNE. As visible in Figure 6 (Middle), plotting these points clearly shows the heterogeneity of the generated expressions and identities. Firstly, renders associated with the same identity are grouped in distinct clusters. Furthermore, within each cluster, we observe how there is a local variation of ID features as they adapt their response to different expressions.

4.3 ID-conditioned Text-based Customization

The unique characteristics of our model allow for the customization of 3D objects without identity or expression drift. Figure 7 and 2 (Right) showcase how a variety of alterations can be imposed on the

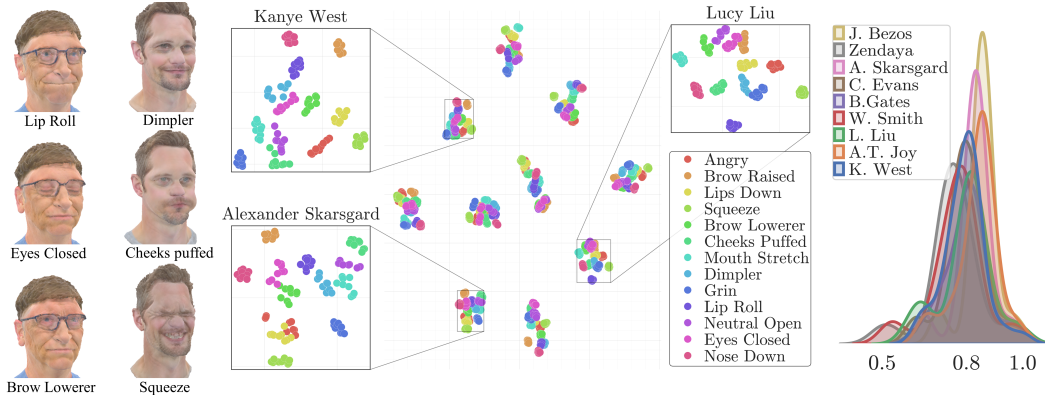


Figure 6: **Expressions analysis.** ID-to-3D creates a variety of expressions with robust ID consistency. **(Left) Visualizations** of different expressions for 2 identities (i.e. *Bill Gates*, *Alexander Skarsgard*). **(Middle) Expression diversity.** t-SNE plot visualizing the ID embeddings computed considering different camera poses, expressions, and subjects. Different identities and expressions are clustered separately. **(Right) Identity similarity distributions** between neutral-pose and remaining expressions. Rendered with [2].



Figure 7: **Identity-Consistent Editing.** **(Left) De-aged 3D heads** generated using different identity conditioning and the textual prompt: "...as a cute baby". Normal maps are displayed in world coordinates next to photorealistic renderings. **(Right) Geometry and texture editing** with text prompts. ID-to-3D edits appearance and geometric features in an ID-consistent manner.

object’s geometry, given the appropriate text conditioning while maintaining the subject’s general physiognomy. In particular, ID-to-3D displays aligned geometry and appearance even after editing, preserving the ability to convey different expressions (e.g. from “eyes closed” to “brow lowerer”, from “neutral” to “squeeze”) even after facial hair textures have been altered via textual prompt.

Figure 8 provides evidence on the text-guided editing capability of our model for richer and more complex textual prompts. In particular, we showcase text prompts associated with various hairstyles, head accessories, and face shape changes driven by different ethnic backgrounds. Note that our method can interpret and exploit text-based inputs not addressed in previous works (e.g. id-driven changes in “aging”, “gender”, “heritage”). Even when using the exact same text prompt (e.g. “A woman with African-American heritage, ... 70’s hairstyle.”) our model generates unique identity-consistent assets that simultaneously align with the text and retain the characteristic facial features of the input ID. Our approach is able to address practical scenarios and opens new avenues for expressive text-guided editing of 3D assets.

5 Conclusion and Ethical Considerations

Limitations. Despite setting a new state-of-the-art, we acknowledge ID-to-3D limitations: **(1)** The generalization capacity is constrained by the employed face embedding network [17], the pretrained

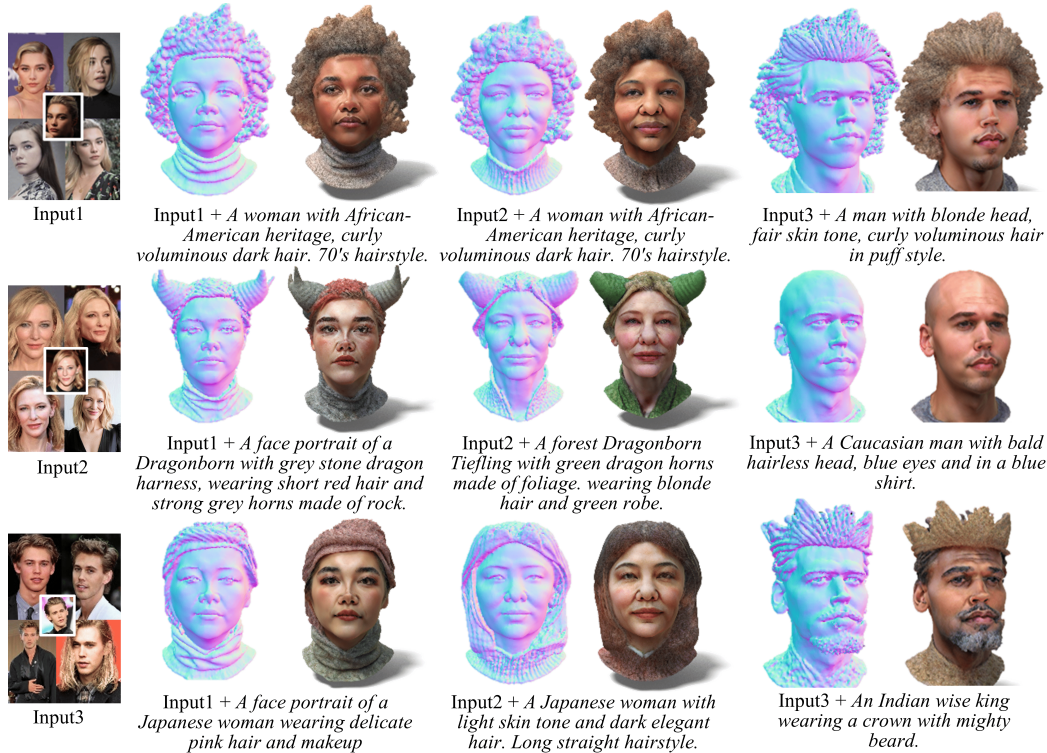


Figure 8: **Identity-Consistent Editing with Rich Textual Prompts.** ID-to-3D generates ID-consistent assets that reflect both subtle and significant changes in geometry and appearance described in a detailed textual input. Normal maps are displayed in world coordinates next to photorealistic renderings and input prompts.

large-scale diffusion models [69, 88], and the finetuning of the 2D guidance models on a dataset of relatively small size, which might result in inaccurate representation biases and biases due to dataset imbalances; (2) the texture-guidance model results might come short of photorealism. The generated textures mimic the albedos from the [24] dataset, which contains expression-rich data but only low-resolution UV diffuse albedo maps; (3) the lack of specific optimization for physically bounded textures and geometries might occasionally produce unnatural exaggerated facial characteristics, hence we call them pseudo-albedo, despite producing render-ready assets; (4) the computational resources needed for PBR rendering hinder potential applications in video-driven problems.

Societal Impact. Technological advancements in automatic 3D human generation have various beneficial applications, but also raise important ethical considerations about representation and potential misuses. We advocate for responsible research and take the following steps to mitigate unauthentic reconstructions: (1) we restrict training data to the facial region only. (2) we advocate for the replacement of text-based prompts with larger ID embeddings that are sometimes lacking [90] but entail significantly less bias than methods trained only on celebrity datasets and text [65].

Conclusion. In this work, we present ID-to-3D, a novel method for expressive 3D head asset generation from one or more face images. Our method deploys a novel human parametric expression model in tandem with specialized geometry and albedo guidance, not only to create intricately detailed head avatars with realistic textures but also to achieve strikingly ID-consistent results across a wide range of expressions, setting a new benchmark in comparison to existing SDS techniques. Without having to rely on 3D captured datasets that are expensive to collect and typically biased, and without being constrained on a specific geometry template, our method can be employed by a broad range of subjects, with different features such as skin tone and hairstyle.

Acknowledgements. S. Zafeiriou and part of the research was funded by the EPSRC Fellowship DEFORM (EP/S010203/1), EPSRC Project GNOMON (EP/X011364/1) and Turing AI Fellowship (EP/Z534699/1).

References

- [1] <https://www.bing.com/images/feed>.
- [2] Marmoset toolbag. <https://marmoset.co/toolbag/>. Accessed: 2022-05-20.
- [3] Dejan Azinović, Olivier Maury, Christophe Hery, Matthias Nießner, and Justus Thies. High-res facial appearance capture from polarized smartphone images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16836–16846, 2023.
- [4] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *SIGGRAPH '99*, 1999.
- [5] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part V 14*, pages 561–578. Springer, 2016.
- [6] James Booth, Anastasios Roussos, Stefanos Zafeiriou, Allan Ponniah, and David Dunaway. A 3d morphable model learnt from 10,000 faces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5543–5552, 2016.
- [7] Chen Cao, Yanlin Weng, Shun Zhou, Yiyong Tong, and Kun Zhou. Facewarehouse: A 3d facial expression database for visual computing. *IEEE Transactions on Visualization and Computer Graphics*, 20(3):413–425, 2013.
- [8] Yukang Cao, Yan-Pei Cao, Kai Han, Ying Shan, and Kwan-Yee K Wong. Dreamavatar: Text-and-shape guided 3d human avatar generation via diffusion models. *arXiv preprint arXiv:2304.00916*, 2023.
- [9] Yukang Cao, Yan-Pei Cao, Kai Han, Ying Shan, and Kwan-Yee K Wong. Guide3d: Create 3d avatars from text and image guidance. *arXiv preprint arXiv:2308.09705*, 2023.
- [10] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22246–22256, 2023.
- [11] Yongwei Chen, Rui Chen, Jiabao Lei, Yabin Zhang, and Kui Jia. Tango: Text-driven photorealistic and robust 3d stylization via lighting decomposition. *Advances in Neural Information Processing Systems*, 35:30923–30936, 2022.
- [12] Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G Schwing, and Liang-Yan Gui. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4456–4465, 2023.
- [13] Siying Cui, Jiankang Deng, Jia Guo, Xiang An, Yongle Zhao, Xinyu Wei, and Ziyong Feng. Idadapter: Learning mixed features for tuning-free personalization of text-to-image models. *arXiv preprint arXiv:2403.13535*, 2024.
- [14] Radek Daněček, Michael J Black, and Timo Bolkart. Emoca: Emotion driven monocular face capture and animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20311–20322, 2022.
- [15] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 145–156, 2000.
- [16] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *CVPR*, 2020.
- [17] 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.
- [18] Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau. Single-image svbrdf capture with a rendering-aware deep network. *ACM Transactions on Graphics (ToG)*, 37(4):1–15, 2018.
- [19] Bernhard Egger, William AP Smith, Ayush Tewari, Stefanie Wuhler, Michael Zollhoefer, Thabo Beeler, Florian Bernard, Timo Bolkart, Adam Kortylewski, Sami Romdhani, et al. 3d morphable face models—past, present, and future. *ACM Transactions on Graphics (ToG)*, 39(5):1–38, 2020.
- [20] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022.
- [21] Baris Gecer, Alexandros Lattas, Stylianos Ploumpis, Jiankang Deng, Athanasios Papaioannou, Stylianos Moschoglou, and Stefanos Zafeiriou. Synthesizing coupled 3d face modalities by trunk-branch generative adversarial networks. In *European conference on computer vision*, pages 415–433. Springer, 2020.
- [22] Baris Gecer, Stylianos Ploumpis, Irene Kotsia, and Stefanos Zafeiriou. Ganfit: Generative adversarial network fitting for high fidelity 3d face reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1155–1164, 2019.
- [23] Abhijeet Ghosh, Graham Fyffe, Borom Tunwattanapong, Jay Busch, Xueming Yu, and Paul Debevec. Multiview face capture using polarized spherical gradient illumination. In *Proceedings of the 2011 SIGGRAPH Asia Conference*, pages 1–10, 2011.
- [24] Simon Giebenhain, Tobias Kirschstein, Markos Georgopoulos, Martin Rünz, Lourdes Agapito, and Matthias Nießner. Learning neural parametric head models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21003–21012, 2023.

- [25] Paulo Gotardo, Jérémy Riviere, Derek Bradley, Abhijeet Ghosh, and Thabo Beeler. Practical dynamic facial appearance modeling and acquisition. *ACM Transactions on Graphics (ToG)*, 37(6):1–13, 2018.
- [26] Xiao Han, Yukang Cao, Kai Han, Xiatian Zhu, Jiankang Deng, Yi-Zhe Song, Tao Xiang, and Kwan-Yee K Wong. Headsculpt: Crafting 3d head avatars with text. *Advances in Neural Information Processing Systems*, 36, 2024.
- [27] Ayaan Haque, Matthew Tancik, Alexei A Efros, Aleksander Holynski, and Angjoo Kanazawa. Instruct-nerf2nerf: Editing 3d scenes with instructions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19740–19750, 2023.
- [28] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [29] Fangzhou Hong, Mingyuan Zhang, Liang Pan, Zhongang Cai, Lei Yang, and Ziwei Liu. Avatarclip: Zero-shot text-driven generation and animation of 3d avatars. *arXiv preprint arXiv:2205.08535*, 2022.
- [30] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- [31] Xin Huang, Ruizhi Shao, Qi Zhang, Hongwen Zhang, Ying Feng, Yebin Liu, and Qing Wang. Humannorm: Learning normal diffusion model for high-quality and realistic 3d human generation. *arXiv preprint arXiv:2310.01406*, 2023.
- [32] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 867–876, 2022.
- [33] Minguk Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. Scaling up gans for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10124–10134, 2023.
- [34] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in neural information processing systems*, 34:852–863, 2021.
- [35] 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.
- [36] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020.
- [37] Nikos Kolotouros, Thimo Alldieck, Andrei Zanfir, Eduard Bazavan, Mihai Fieraru, and Cristian Sminchisescu. Dreamhuman: Animatable 3d avatars from text. *Advances in Neural Information Processing Systems*, 36, 2024.
- [38] Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. Modular primitives for high-performance differentiable rendering. *ACM Transactions on Graphics (ToG)*, 39(6):1–14, 2020.
- [39] Alexandros Lattas, Yiming Lin, Jayanth Kannan, Ekin Ozturk, Luca Filipi, Giuseppe Claudio Guarnera, Gaurav Chawla, and Abhijeet Ghosh. Practical and scalable desktop-based high-quality facial capture. In *European Conference on Computer Vision*, pages 522–537. Springer, 2022.
- [40] Alexandros Lattas, Stylianos Moschoglou, Stylianos Ploumpis, Baris Gecer, Jiankang Deng, and Stefanos Zafeiriou. Fitme: Deep photorealistic 3d morphable model avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8629–8640, 2023.
- [41] Alexandros Lattas, Stylianos Moschoglou, Stylianos Ploumpis, Baris Gecer, Abhijeet Ghosh, and Stefanos Zafeiriou. Avatarme++: Facial shape and brdf inference with photorealistic rendering-aware gans. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9269–9284, 2021.
- [42] Ruilong Li, Karl Bladin, Yajie Zhao, Chinmay Chinara, Owen Ingraham, Pengda Xiang, Xinglei Ren, Pratusha Prasad, Bipin Kishore, Jun Xing, et al. Learning formation of physically-based face attributes. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3410–3419, 2020.
- [43] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4d scans. *ACM Trans. Graph.*, 36(6):194–1, 2017.
- [44] Tingting Liao, Hongwei Yi, Yulian Xiu, Jiayang Tang, Yangyi Huang, Justus Thies, and Michael J Black. Tada! text to animatable digital avatars. *arXiv preprint arXiv:2308.10899*, 2023.
- [45] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 300–309, 2023.
- [46] Hongyu Liu, Xuan Wang, Ziyu Wan, Yujun Shen, Yibing Song, Jing Liao, and Qifeng Chen. Headartist: Text-conditioned 3d head generation with self score distillation. *arXiv preprint arXiv:2312.07539*, 2023.
- [47] Minghua Liu, Chao Xu, Haiyan Jin, Linghao Chen, Mukund Varma T, Zexiang Xu, and Hao Su. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. *Advances in Neural Information Processing Systems*, 36, 2024.

- [48] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9298–9309, 2023.
- [49] Xian Liu, Xiaohang Zhan, Jiayang Tang, Ying Shan, Gang Zeng, Dahua Lin, Xihui Liu, and Ziwei Liu. Humangaussian: Text-driven 3d human generation with gaussian splatting. *arXiv preprint arXiv:2311.17061*, 2023.
- [50] Huiwen Luo, Koki Nagano, Han-Wei Kung, Qingguo Xu, Zejian Wang, Lingyu Wei, Liwen Hu, and Hao Li. Normalized avatar synthesis using stylegan and perceptual refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11662–11672, 2021.
- [51] Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, Paul E Debevec, et al. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. *Rendering Techniques*, 9(10):2, 2007.
- [52] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Realfusion: 360deg reconstruction of any object from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8446–8455, 2023.
- [53] Gal Metzger, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12663–12673, 2023.
- [54] Oscar Michel, Roi Bar-On, Richard Liu, Sagie Benaim, and Rana Hanocka. Text2mesh: Text-driven neural stylization for meshes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13492–13502, 2022.
- [55] Nasir Mohammad Khalid, Tianhao Xie, Eugene Belilovsky, and Tiberiu Popa. Clip-mesh: Generating textured meshes from text using pretrained image-text models. In *SIGGRAPH Asia 2022 conference papers*, pages 1–8, 2022.
- [56] Stylianos Moschoglou, Stylianos Ploumpis, Mihalis A Nicolaou, Athanasios Papaioannou, and Stefanos Zafeiriou. 3dfacegan: Adversarial nets for 3d face representation, generation, and translation. *International Journal of Computer Vision*, 128:2534–2551, 2020.
- [57] Norman Müller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Buló, Peter Kotschieder, and Matthias Nießner. Diffrr: Rendering-guided 3d radiance field diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4328–4338, 2023.
- [58] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting triangular 3d models, materials, and lighting from images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8280–8290, 2022.
- [59] Foivos Paraperas Papantoniou, Alexandros Lattas, Stylianos Moschoglou, Jiankang Deng, Bernhard Kainz, and Stefanos Zafeiriou. Arc2face: A foundation model of human faces. *arXiv preprint arXiv:2403.11641*, 2024.
- [60] Foivos Paraperas Papantoniou, Alexandros Lattas, Stylianos Moschoglou, and Stefanos Zafeiriou. Relightify: Relightable 3d faces from a single image via diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8806–8817, 2023.
- [61] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10975–10985, 2019.
- [62] Xu Peng, Junwei Zhu, Boyuan Jiang, Ying Tai, Donghao Luo, Jiangning Zhang, Wei Lin, Taisong Jin, Chengjie Wang, and Rongrong Ji. Portraitbooth: A versatile portrait model for fast identity-preserved personalization. *arXiv preprint arXiv:2312.06354*, 2023.
- [63] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *ICLR*, 2023.
- [64] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint arXiv:2306.17843*, 2023.
- [65] 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.
- [66] Amit Raj, Srinivas Kaza, Ben Poole, Michael Niemeyer, Nataniel Ruiz, Ben Mildenhall, Shiran Zada, Kfir Aberman, Michael Rubinstein, Jonathan Barron, et al. Dreambooth3d: Subject-driven text-to-3d generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2349–2359, 2023.
- [67] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.
- [68] George Retsinas, Panagiotis P Filintisis, Radek Danecek, Victoria F Abrevaya, Anastasios Roussos, Timo Bolkart, and Petros Maragos. 3d facial expressions through analysis-by-neural-synthesis. *arXiv preprint arXiv:2404.04104*, 2024.
- [69] 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.

- [70] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510, 2023.
- [71] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Wei Wei, Tingbo Hou, Yael Pritch, Neal Wadhwa, Michael Rubinstein, and Kfir Aberman. Hyperdreambooth: Hypernetworks for fast personalization of text-to-image models. *arXiv preprint arXiv:2307.06949*, 2023.
- [72] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022.
- [73] Aditya Sanghi, Hang Chu, Joseph G. Lambourne, Ye Wang, Chin-Yi Cheng, Marco Fumero, and Kamal Rahimi Malekshan. Clip-forge: Towards zero-shot text-to-shape generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18603–18613, June 2022.
- [74] Axel Sauer, Tero Karras, Samuli Laine, Andreas Geiger, and Timo Aila. Stylegan-t: Unlocking the power of gans for fast large-scale text-to-image synthesis. In *International conference on machine learning*, pages 30105–30118. PMLR, 2023.
- [75] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022.
- [76] Junyoung Seo, Wooseok Jang, Min-Seop Kwak, Hyeonsu Kim, Jaehoon Ko, Junho Kim, Jin-Hwa Kim, Jiyoung Lee, and Seungryoung Kim. Let 2d diffusion model know 3d-consistency for robust text-to-3d generation. *arXiv preprint arXiv:2303.07937*, 2023.
- [77] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. *Advances in Neural Information Processing Systems*, 34:6087–6101, 2021.
- [78] William AP Smith, Alassane Seck, Hannah Dee, Bernard Tiddeman, Joshua B Tenenbaum, and Bernhard Egger. A morphable face albedo model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5011–5020, 2020.
- [79] Jingxiang Sun, Bo Zhang, Ruizhi Shao, Lizhen Wang, Wen Liu, Zhenda Xie, and Yebin Liu. Dreamcraft3d: Hierarchical 3d generation with bootstrapped diffusion prior. *arXiv preprint arXiv:2310.16818*, 2023.
- [80] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-it-3d: High-fidelity 3d creation from a single image with diffusion prior. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22819–22829, 2023.
- [81] 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.
- [82] Duotun Wang, Hengyu Meng, Zeyu Cai, Zhijing Shao, Qianxi Liu, Lin Wang, Mingming Fan, Ying Shan, Xiaohang Zhan, and Zeyu Wang. Headevolver: Text to head avatars via locally learnable mesh deformation. *arXiv preprint arXiv:2403.09326*, 2024.
- [83] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [84] Qixun Wang, Xu Bai, Haofan Wang, Zekui Qin, and Anthony Chen. Instantid: Zero-shot identity-preserving generation in seconds. *arXiv preprint arXiv:2401.07519*, 2024.
- [85] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. In *Advances in Neural Information Processing Systems*, volume 36, 2023.
- [86] Jiale Xu, Xintao Wang, Weihao Cheng, Yan-Pei Cao, Ying Shan, Xiaohu Qie, and Shenghua Gao. Dream3d: Zero-shot text-to-3d synthesis using 3d shape prior and text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20908–20918, 2023.
- [87] Yuxuan Yan, Chi Zhang, Rui Wang, Yichao Zhou, Gege Zhang, Pei Cheng, Gang Yu, and Bin Fu. Facestudio: Put your face everywhere in seconds. *arXiv preprint arXiv:2312.02663*, 2023.
- [88] 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.
- [89] Jonathan Young. xatlas. <https://github.com/jpcy/xatlas>. Accessed: 2022-05-20.
- [90] Seyma Yucer, Furkan Tektas, Noura Al Moubayed, and Toby P Breckon. Measuring hidden bias within face recognition via racial phenotypes. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 995–1004, 2022.
- [91] Yifei Zeng, Yuanxun Lu, Xinya Ji, Yao Yao, Hao Zhu, and Xun Cao. Avatarbooth: High-quality and customizable 3d human avatar generation. *arXiv preprint arXiv:2306.09864*, 2023.
- [92] Longwen Zhang, Qiwei Qiu, Hongyang Lin, Qixuan Zhang, Cheng Shi, Wei Yang, Ye Shi, Sibi Yang, Lan Xu, and Jingyi Yu. Dreamface: Progressive generation of animatable 3d faces under text guidance. *arXiv preprint arXiv:2304.03117*, 2023.
- [93] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847,

- 2023.
- [94] Zheng Zhu, Guan Huang, Jiankang Deng, Yun Ye, Junjie Huang, Xinze Chen, Jiagang Zhu, Tian Yang, Dalong Du, Jiwen Lu, and Jie Zhou. Webface260m: A benchmark for million-scale deep face recognition. *TPAMI*, 2022.

1. **Claims**

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: [Yes]

Justification: The claims are listed in both the abstract and introduction, reflect the paper’s method and are justified by the experimental section.

2. **Limitations**

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The limitations are listed in the limitations section.

3. **Theory Assumptions and Proofs**

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: [NA]

4. **Experimental Result Reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We base our method on public code, report our technical details in the supplemental materials, and will be releasing our code.

5. **Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We use already publicly available data, and will be open-sourcing our code.

6. **Experimental Setting/Details**

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We include all the above details in the supplemental materials.

7. **Experiment Statistical Significance**

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Error bars are not reported because it would be too computationally expensive.

8. **Experiments Compute Resources**

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We have included the above details in the supplemental materials.

9. **Code Of Ethics**

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification:

10. **Broader Impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: Yes

Justification: We discuss both in the Conclusion and Ethical Considerations section.

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [Yes]

Justification: We discuss both in the Conclusion and Ethical Considerations section.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification:

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: the paper does not release new assets.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: