

# MagicMirror: Fast and High-Quality Avatar Generation with a Constrained Search Space

Armand Comas-Massagué<sup>1,2</sup>, Di Qiu<sup>1</sup>, Menglei Chai<sup>1</sup>, Marcel Bühler<sup>1,3</sup>,  
Amit Raj<sup>1</sup>, Ruiqi Gao<sup>4</sup>, Qiangeng Xu<sup>1</sup>, Mark Matthews<sup>1</sup>, Paulo Gotardo<sup>1</sup>,  
Sergio Orts-Escolano<sup>1</sup>, and Thabo Beeler<sup>1</sup>

<sup>1</sup> Google AR

<sup>2</sup> Northeastern University

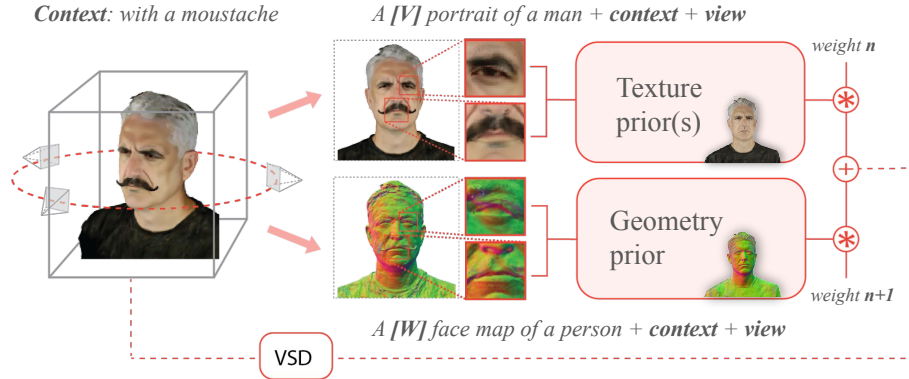
<sup>3</sup> ETH Zürich

<sup>4</sup> Google DeepMind



**Fig. 1:** We propose MagicMirror, a method for fast text-guided 3D avatar head generation, with the option of subject personalization. (left) We illustrate how given subject pictures, MagicMirror can generate a 3D avatar with the subject’s stylized appearance by following text descriptions. Avatars exhibit high-quality in both geometry and texture, with significant alterations while preserving the identity of the subject. (right) It can also generate well-known characters by only employing a text prompt.

**Abstract.** We introduce a novel framework for 3D human avatar generation and personalization, leveraging text prompts to enhance user engagement and customization. Central to our approach are key innovations aimed at overcoming the challenges in photo-realistic avatar synthesis. Firstly, we utilize a conditional Neural Radiance Fields (NeRF) model, trained on a large-scale unannotated multi-view dataset, to create a versatile initial solution space that accelerates and diversifies avatar generation. Secondly, we develop a geometric prior, leveraging the capabilities of Text-to-Image Diffusion Models, to ensure superior view invariance and enable direct optimization of avatar geometry. These foundational ideas are complemented by our optimization pipeline built on Variational Score Distillation (VSD), which mitigates texture loss and over-saturation issues. As supported by our extensive experiments, these strategies collectively enable the creation of custom avatars with unparalleled visual quality and better adherence to input text prompts. You can find more results and videos in our website: [syntec-research.github.io/MagicMirror](https://syntec-research.github.io/MagicMirror)

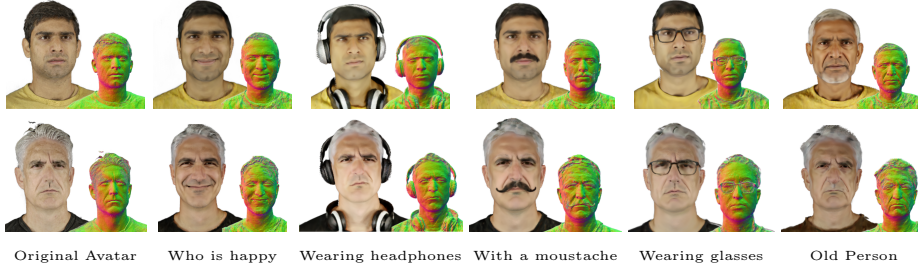


**Fig. 2:** Our two pipelines for 3D head avatar generation and customization follow the same structure: a pre-trained conditional NeRF model serves as 3D prior for fast avatar generation. Our pipelines additionally leverage two pre-trained text-to-image diffusion models as texture and geometry priors, allowing for distillation-based customization of both these components based on input text prompts with state-of-the-art quality.

## 1 Introduction

Customizable 3D human avatars are central for many experiences such as gaming, v-tubing, augmented and virtual reality (AR/VR), or telepresence applications. Intuitive editing and personalization techniques for such avatars are highly desirable, as customized avatars provide a greater sense of engagement, ownership and aid adoption of the aforementioned technologies. Traditional CGI editing techniques, however, are still difficult, non-intuitive, and laborious for the average user. Recently, text-prompting has emerged as a natural and intuitive interface to control the creation and customization of highly complex generative outputs, due to the impressive progress of Language-Image modeling [53] and Text-to-Image Diffusion Models [24]. Two main approaches have emerged for generative modeling of 3D assets: direct 3D modeling, and neural rendering techniques leveraging 2D images.

Direct 3D models largely conform to the text-to-image paradigm which train a generative model from a large dataset of labeled 3D assets [29]. Sourcing such data at scale, however, is difficult and expensive [17]. 3D assets are nowhere nearly as abundant as the 2D images easily available on the Internet. Furthermore, 3D assets that are available typically lack the rich semantic information that often accompanies Internet images. Consequently, the results from this category typically lack diversity and quality compared to their 2D large-scale counterparts. The second category of methods leverage the implicit 3D knowledge within 2D generative models, lifting 2D outputs onto 3D via differential rendering and novel objective functions. Several designs of these objective functions have been proposed, including simple reconstruction losses based on transformed 2D images [23], high-level text-image misalignment scores [67], and model distillation [52]. These methods work best if the outputs of the 2D model are multi-view consistent, which is usually not the case, leading to non-convergence and infamous "Janus



**Fig. 3:** Our novel framework, MagicMirror, can successfully change facial expressions, features, and add accessories or specific styles to the person.

face" artifacts [52]. Although appealing, this second category of methods has several key challenges.

First, despite their ability to generate large amounts of text-guided 2D images and supervisory signals, they are not guaranteed to be multi-view consistent. Consequently, 3D optimization suffers from conflicting supervision. This issue can be mitigated by reducing the amount of overlap between views, through reduced view count, and evenly distributing views. However, a possible unfortunate effect is in making the problem ill-posed, thus leading to poor results. An approach to improve reconstructions might then be to improve the multi-view consistency of existing 2D image generators. This could be achieved by training on multi-view data and sharing information across views [46, 60]. While hopeful, this approach still bears the burden of sourcing multi-view data, often as difficult as sourcing 3D assets described above. Other approaches based on model distillation require the use of a high classifier-free guidance weight [25], causing textureless and over-saturated results [52] and, more importantly, reducing diversity [71].

Another approach is to constrain the space of 3D objects of interest and their representation, similar to popular parametric blendshapes models that enable shape reconstruction given only partial information such as monocular landmarks [16]. For text-guided avatar generation and editing, existing works typically employ an object-specific parametric 3D morphable model (3DMM) [5] as underlying geometry proxy. However, avatar customization remains challenging because it requires creating novel, semantically meaningful geometric structures that need to introduce new out-of-model elements. So far, because of the crucial dependency on multi-view consistent supervision, it is still challenging to obtain high-quality avatar customizations that closely follow their associated text prompts.

This paper presents MagicMirror, our novel framework for text-guided 3D head avatar generation and editing whose visual quality improves upon the current state of the art. Our key idea is to derive constraints and priors that make the test-time optimization problem easier, and less dependent on photometric consistency. This idea is implemented via the following important framework components:

1. A constrained **initial solution space** is first learned as a conditional NeRF model trained on an unannotated multi-view dataset of human heads; this

flexible model can express a wide range of head appearances and geometries and subsequently facilitates fast avatar generation and editing.

2. Leveraging a pre-trained text-to-image diffusion model and its ability to learn new concepts, we build a geometric prior by teaching this model to generate normal maps. This additional **geometry prior** encourages better view invariance, direct geometry optimization, and also largely mitigates the photometric inconsistency problem from conventional multi-view supervision.
3. When optimizing our conditional NeRF, Score Distillation Sampling (SDS) [52] can lead to artifacts such as lack of texture and over saturation. We overcome these issues by adopting Variational Score Distillation (VSD) [71], allowing us to optimize both appearance and geometry with higher quality.

As demonstrated next, our framework generates custom avatars following specific text instructions with a high level of faithfulness and visual detail. As in Dream-Booth [54, 57], we leverage a re-contextualization technique that allows users to personalize their avatars with ease and high fidelity to their own identity, while making it fun to create and explore.

## 2 Related Work

### 2.1 3D Representations for Photorealistic Avatars

The significance of 3D human modeling has spurred thorough exploration into proper avatar representations. Early methods [2, 9, 13, 21, 28, 32, 39, 63, 64, 73] adopt explicit geometry and appearance, particularly parametric human prior models [5, 47]. However, these approaches struggle with limited representation capabilities.

Lately, the rapid progress in volumetric neural rendering like NeRF [50] and 3DGS [38] has promoted implicit avatar modeling, owing to its rendering quality and comprehensive representation. Nevertheless, training such a model typically demands substantial multi-view data for a single subject. To enable monocular inputs and facilitate animation, various human priors have been explored.

One approach involves hybrid representations that leverages morphable models, such as NerFACE [19], RigNeRF [1], IMAvatar [77], and MonoAvatar [3]. While efficient monocular avatar rendering and animation can be achieved, quality is often compromised due to the limitations of explicit models. Another strategy relies on generative human priors, capable of reconstructing high-quality implicit avatars from sparse inputs. For example, PVA [55], CodecAvatar [8], Live3DPortrait [65], and Preface [7]. In this work, we follow the latter approach and demonstrate that such a prior not only assists with monocular avatar modeling but also text-driven avatar synthesis.

### 2.2 Text-Guided Avatar Generation and Editing

Generative models have enabled identity sampling within the 2D [36] and 3D [11] latent space. Nonetheless, there is a general preference for better controllability. Among various control modalities, such as scribbles [15], semantic attributes [59], and image references [79], text prompts in natural language are more widely accepted for a broad range of tasks.



The emergence of the language-vision model CLIP [53] has made text-guided avatar editing feasible. In 2D, pioneering work such as StyleGAN-NADA [20] transfers pre-trained StyleGAN2 [37] models to the target style domain described by a textual prompt. This capability extends to 3D as well, where CLIP supervision is integrated with explicit [26] or implicit [68] human models. However, these models often encounter limitations in expressing full 3D complexity, primarily due to the restricted capacity of CLIP in comprehending intricate prompts.

With the recent advancements in 3D-aware diffusion models, diffusion-based text-guided avatar synthesis have garnered increased attention. DreamFace [76] and HeadSculpt [22] introduce coarse-to-fine pipelines to enhance identity-awareness and achieve fine-grained text-driven head avatar creation. Human-Norm [29] presents an explicit human generation pipeline, employing normal- and depth-adapted diffusion for geometry generation and a normal-aligned diffusion for texture generation. In a similar two-stage pipeline, SEEAAvatar [74] and HeadArtist [22] evolve geometry generation from a template human prior and represent appearance through neural texture fields. Meanwhile, AvatarBooth [75], AvatarCraft [33], DreamAvatar [10], DreamHuman [41], and DreamWaltz [31] propose text-driven avatar creation utilizing implicit surface representation, parameterized with morphable models for easy animation. In terms of editing, AvatarStudio [49] achieves personalized NeRF-based avatar stylization through view-and-time-aware SDS on dynamic multi-view inputs.

### 2.3 3D-Aware Diffusion Models

The success of text-to-image diffusion models [56] naturally encourages researchers to explore 3D-aware diffusion. Building from 2D diffusion, many studies have concentrated on synthesizing consistent novel 2D views of 3D objects, such as 3DiM [72], SparseFusion [78], and GeNVS [12]. Zero-1-to-3 [45] proposes a pipeline that fine-tunes a pre-trained diffusion model with a large-scale synthetic 3D dataset. SyncDreamer [46] further improves the cross-view consistency.

Direct 3D generation have also been explored across various representations, including point clouds [48, 51], feature grids [35], tri-planes [61, 70], and radiance fields [34]. However, due to the complexity of representations, heavy architecture, and shortage of large-scale 3D data, 3D diffusion often suffers from poor generalization and low quality results.

Compared to 3D diffusion models, lifting 2D diffusion for 3D generation is more appealing, spearheaded by the pioneering works of DreamFusion [52] and SJC [69]. At the heart of these approaches lies score distillation sampling (SDS), which employs 2D diffusion models as score functions on sampled renderings, providing supervision for optimizing the underlying 3D representations. Subsequent works like DreamTime [30], MVDream [60], and ProlificDreamer [71] refine the architectural design with better sampling strategy, loss design, and multi-view prior. Meanwhile, Magic3D [42], TextMesh [66], Make-It-3D [62], and Fantasia3D [14] extend the approach to other representations such as textured meshes and point clouds. Notably, variational score distillation (VSD) is proposed in ProlificDreamer to address oversaturation and texture-less issues of SDS. Our method also adopts VSD to enhance the quality of the generated results.

### 3 Method

We present two similar pipelines for (P1) text-driven generation and (P2) personalized 3D head avatars editing. Both pipelines have the same structure, illustrated in Fig. 2. We start from a NeRF avatar initialization, and provide as input a text prompt. We render random views from our avatar, chosen randomly from a set of orbit renders. The avatar is parameterized by a conditional NeRF model and initialized with any latent identity code (Sec. 3.1). We then employ a distillation approach with a *geometry prior* to optimize the initial avatar’s NeRF appearance and density, following the methodology described in Section 3.3.

In both pipelines, the geometry is captured by a diffusion prior, which is fine-tuned to capture facial geometry features from a single avatar (Sec. 3.2). More specifically, besides the conditional NeRF, in Pipeline P1 our method mainly leverages a pre-trained text-to-image diffusion models that captures the distribution of real RGB images. The diffusion model and the geometric prior allows us to customize both the appearance and geometry of our initial NeRF avatar as guided by an input text prompt in the form: “*A portrait of a [source description]*”. There is no personalization element in P1. Thus, the prompt does not require a subject identifier. Avatar customization is done using a distillation-based objective function derived from Variational Score Distillation (VSD) [71] (Sec. 3.3).

In Pipeline P2, we personalize on a particular subject by first conditioning on user-provided 2D images in multiple views, or by rendering images from a reconstructed digital asset of the target subject. This subject is associated with a unique identity token  $[V]$  and  $[source\ description]$ , using DreamBooth to fine-tune our text-to-image diffusion model with the text prompt “*A [V] portrait of [source description]*”. The user can then supply a new  $[target\ description]$  prompt to guide avatar stylization to their preference, which is achieved by optimizing the conditional NeRF using the objective defined in equation 4, with the text embedding corresponding to “*A [V] portrait of [target description]*”.

Finally, a user can combine multiple priors in parallel to achieve different objectives. Updates from subject-aware texture priors can be blended with updates from text-to-image generic diffusion priors, prompted with multiple context prompts.

#### 3.1 Constraining the solution space

To lift the partial 2D information from text-to-image diffusion models, a 3D prior model is needed to constrain the optimization domain. We thus parameterize the subspace of 3D head avatars using a conditional NeRF, which simplifies the optimization while remaining flexible enough to accommodate variations outside the training data.

To learn our solution subspace, we leverage Preface [7], a conditional model that extends Mip-NeRF360 [4] with a conditioning (identity) latent code concatenated to the inputs of each MLP layer. Here, we briefly describe how to train this conditional NeRF on our multiview dataset of 1450 human faces with a neutral expression, captured by 13 synchronized cameras under uniform in-studio

lighting. More details can be found in the Supplementary material. Each face is assigned a learnable latent code [6] that is optimized together with the model weights, under the supervision of pixelwise reconstruction loss only. Each training batch randomly samples pixel rays from all subjects and cameras, promoting generalization over the space of human faces rather than over-fitting to a few subjects. The importance of the diversity of training faces is highlighted in our ablation study in Sec. 3.5.

### 3.2 View-invariant geometric prior



**Fig. 4:** Text-to-Image Diffusion Model have the remarkable ability to re-contextualize new concepts. We show the generated normal maps under new text prompts. Note that they are *not* rendered from a NeRF.

Given the above conditional NeRF that models diverse 3D head avatars, we now turn to incorporating high-frequency geometric details that are (1) authentic to specific target user and (2) consistent with a given text prompt. To this end, we show that an additional pretrained text-to-image diffusion model  $\mathcal{D}_{\text{color}}$  can be used to capture new appearance and geometry concepts without any architectural change. Instead of retraining the diffusion model to encourage multiview consistency, we propose a novel and effective solution that teaches the model to also generate normal maps of human heads, effectively deriving a second model  $\mathcal{D}_{\text{normal}}$ .

Our solution leverages the few-shot learning technique proposed by DreamBooth [57]: given 60 world-space surface normal renderings from different camera views of an avatar, we pair them with text descriptions such as "A  $[W]$  face map of a man/woman  $[source\ description]$ ", where " $[W]$ " is the unique identifier for the new concept of *surface normal*.

Next, we fine-tune the text-to-image diffusion model using these text-annotated surface normal renderings. As a result, the fine-tuned diffusion model  $\mathcal{D}_{\text{normal}}$  can now also predict reasonable surface normal maps for new heads, even when re-contextualized to the rest of the text prompt description (Fig. 4). This new capability provides additional geometric critique for avatar generation and editing, driving the optimization to go beyond simple color edits to effectively improve the geometry.

When defining the input data for fine-tuning the diffusion model, it is crucial that surface normals be defined in a fixed world coordinate system that is aligned with our solution space, hence independent of camera viewpoint. The source normal maps can then be obtained from any subject, as evidenced in Sec. 3.5.

### 3.3 Test-time optimization objective

During our test-time optimization, the fine-tuned diffusion model is re-contextualized with the prompt "A  $[V]$  portrait of a  $[target\ description]$ ". Such geometric prior

is then incorporated via a Variational Score Distillation (VSD) [71] optimization objective, as described in the following.

To derive our solution, we first review Score Distillation Sampling (SDS) [52], which minimizes the following loss function:

$$\mathcal{L}_{\text{SDS}}(\text{sg}(\mathcal{D}), I, \epsilon, T, t) = \omega(t) \|\text{sg}(\mathcal{D}(I, \epsilon, T, t)) - I\|^2, \quad (1)$$

where  $\mathcal{D}$  represents the text-to-image diffusion model that outputs the denoised image by processing the NeRF rendering  $I$ , Gaussian noise  $\epsilon$ , a fixed target text embedding  $T$ , and a time parameter  $t$  that follows certain annealing schedule  $t \rightarrow 0$ .  $\text{sg}(\cdot)$  denotes the stop gradient operator, and  $\omega(t)$  is a time-dependant weighting factor. In what follows, some or all of the loss arguments may be omitted when the context is clear.

It is commonly observed that SDS, with its default high Classifier-Free Guidance (CFG) weight [25], often leads to textureless and over-saturated outputs that significantly impact photorealism and diversity, while a lower CFG weight tends to underperform with SDS. Variational Score Distillation (VSD) [71] introduces a proxy of  $\mathcal{D}$ , denoted as  $\mathcal{D}'$ , which is optimized under the following loss function:

$$\mathcal{L}_{\text{proxy}}(\mathcal{D}', \text{sg}(I)) = \omega(t) \|\mathcal{D}'(I, \epsilon, T, t) - \text{sg}(I)\|^2. \quad (2)$$

Typically,  $\mathcal{D}'$  is selected to be the Low Rank Adaptation (LoRA) [27] of  $\mathcal{D}$ , having identical outputs to  $\mathcal{D}$  at the beginning of the optimization. Simultaneously, VSD also optimizes the NeRF parameters by minimizing  $\mathcal{L}_{\text{SDS}}(I) - \mathcal{L}_{\text{proxy}}(\text{sg}(\mathcal{D}'), I)$ . The full VSD objective is formulated as:

$$\mathcal{L}_{\text{VSD}}(\mathcal{D}', I) = \mathcal{L}_{\text{SDS}}(I) - \mathcal{L}_{\text{proxy}}(\text{sg}(\mathcal{D}'), I) + \mathcal{L}_{\text{proxy}}(\mathcal{D}', \text{sg}(I)). \quad (3)$$

Leveraging the formulation above, our overall test-time optimization objective is:

$$\mathcal{L}_{\text{ours}} = \mathcal{L}_{\text{VSD}}(\mathcal{D}'_{\text{color}}, I_{\text{color}}) + \lambda \mathcal{L}_{\text{VSD}}(\mathcal{D}'_{\text{normal}}, I_{\text{normal}}), \quad (4)$$

where two fixed text-to-image diffusion models,  $\mathcal{D}_{\text{color}}$  and  $\mathcal{D}_{\text{normal}}$ , process color images  $I_{\text{color}}$  and normal maps  $I_{\text{normal}}$ , respectively. The normal map is computed through analytic gradients of the NeRF density. During VSD, we also optimize their LoRAs,  $\mathcal{D}'_{\text{color}}$  and  $\mathcal{D}'_{\text{normal}}$ . Our observations suggest that VSD allows for a smaller CFG weight, generally enhancing convergence and output quality.

*Mixing and weighting concepts:* Through this distillation approach, we can additionally compose and modulate different concepts expressed by the text embedding. We propose to perform compositionality using a perspective akin to findings in the context of Energy-Based Models [18, 43] and Diffusion Models [44]. This notion allows us to generate a variety of results by mixing and weighting two or more concepts, including removal of certain concepts or semantic interpolation, thereby enriching the user experience. More generally, we can optimize the conditional NeRF from initialization with a combination of objectives:

$$\mathcal{L}_{\text{composed}} = \sum_{T \in \text{Positive}} \alpha_T \cdot \mathcal{L}_{\text{ours}}(T) - \sum_{T \in \text{Negative}} \beta_T \cdot \mathcal{L}_{\text{ours}}(T) \quad (5)$$

with  $\{\alpha_T, \beta_T\}$  being the positive modulation constants, balancing the importance of different concepts. A probabilistic interpretation is provided in the supplementary material. The associated updates to the NeRF parameters  $\theta$  are thus

$$\nabla_{\theta} \mathcal{L}_{\text{composed}}(T, \theta) = \sum_{T \in \text{Positive}} \alpha_T \cdot \nabla_{\theta} \mathcal{L}_{\text{ours}}(T, \theta) - \sum_{T \in \text{Negative}} \beta_T \cdot \nabla_{\theta} \mathcal{L}_{\text{ours}}(T, \theta) \quad (6)$$

which is an expression reminiscent of the concept-compositional sampling by means of Langevin Dynamics in Energy-Based Models [18].

Besides the composability interpretation, we can generate smooth interpolations by simply switching from one concept to another. This is, once we already obtained the result from one concept, we can directly apply the optimization with the objective associated to an alternate concept. We observe that the optimization trajectory tends to remain within distribution, assuming the two concepts do not introduce significant changes, such as the opening of the mouth or the addition of extra geometry from accessories. We illustrate these interesting findings in Sec. 4 and hypothesize that incorporating additional data into training our 3D prior model could lead to more meaningful optimization trajectories. Experiments on both methodologies can be found in Sec. 4.3.

### 3.4 Implementation details

For a single prompt, MagicMirror needs a maximum of 1k iterations for both geometry and texture generation, which are performed simultaneously. We utilize 4 TPUs of 96 GB of memory, with a batch sample of  $128 \times 128$  resolution per device. Each device may leverage a different set of weights for its diffusion prior, all of them implemented with Imagen 2.2.3. Our use of Imagen is primarily based on its availability and compatibility with our pipeline. There is no evidence to suggest that Imagen outperforms public models, especially considering recent updates to and their ability to support higher resolutions than the  $128 \times 128$  limit of Imagen 2.2.3. The entire generation process takes about 15 minutes. Additional details can be found in the supplementary.

### 3.5 Ablation studies

**Role of (personalized) geometric prior.** One noteworthy property of the geometry supervision is its robustness. Fig. 5a (top row) illustrates the crucial role played by the geometry prior. Without it, the normal map appears noisy and distorted, with out-of-face structures like headphones poorly constructed. Figure 5a (bottom row) shows that despite our geometry prior being trained on a single avatar, its identity has a negligible impact on the final results. In this comparison, we analyze the effects of a geometry prior trained on the original normals of the subject (shown on the right) against one trained on normals from a random female (shown on the left). The results are visually indistinguishable, with both facial features and overall geometry accurately generated.

**Test-time NeRF initializations.** We test different latent codes to initialize our conditional NeRF for test-time avatar optimization. As can be seen in 5b,

for personalized avatars, different latent codes does not yield significant variance in the final results. The non-personalized avatars can have more variance since there are no more constraints on the appearance than to follow the text prompt.

**Diversity of training data for conditional NeRF.** The performance of MagicMirror significantly benefits from the diversity of the search space. As aforementioned, employing a conditional NeRF trained on multi-view datasets comprising multiple subjects has proven to be valuable. As illustrated in Figure 5c, the final results are notably influenced by the number of subjects used for training the conditional NeRF. Training with a single subject tends to yield very rigid geometry, much more challenging to modify compared to the texture. Utilizing 350 subjects allow for modification in geometry but often results in rough normals and a lack of fine details. Conversely, a more diverse training set of 1450 subjects leads to substantially smoother and more precise geometry. We choose the complete set for all of other experiments.

**Impact of the avatar’s identity latent code** In Fig. 5d, we demonstrate the results of optimizing the conditional latent code in isolation. In this experiment, we perform VSD while freezing the remaining parameters of the network. According to the results, while substantial modifications to the overall geometry and texture are achievable through NeRF conditioning, the solution space is inherently limited to human-like faces, as dictated by the training set. The Figure illustrates how latent inversion fails to accurately capture the green color and distinct facial features of a fantasy character "the Grinch".

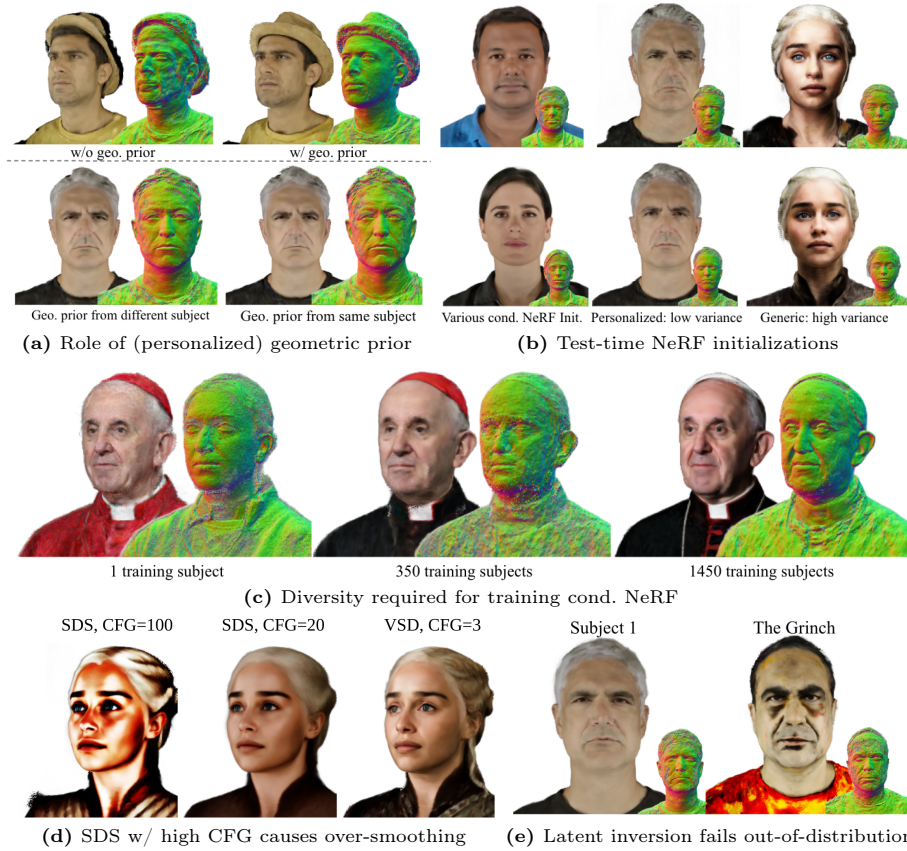
**VSD vs. SDS.** In this final assessment, we evaluate the advantages of VSD over SDS regarding avatar generation in our system. Fig. 5e presents results for both a real captured individual and a fictional character, employing CFG weights of 20 and 100. These comparisons lead us to observe that SDS tends to produce avatars with an overly smooth and saturated appearance that is deficient in fine detail, as extensively documented in existing literature. In contrast, our implementation of VSD yields avatars that exhibit significantly improved realism and finer details, demonstrating the superiority of VSD in our method.

## 4 Experiments

### 4.1 Metrics and Evaluation

It is well known that evaluation is challenging when it comes to 3D generation. Hence, we resort to human preference to assess the quality of our model. We create a human study, where we show 3 rendered views of 18 generated recognizable subjects for all (anonymous) methods to 36 people and ask them to rank them from 1 (low) to 5 (high) in two dimensions. We ask about (i) *visual quality* as a measure of realism of both shape and appearance of the generated avatars as a generic human being; and (ii) *similarity to the real person* as a measure of alignment to the real target person’s identity. Finally, we utilize the same set of views to provide results for PickScore [40], which quantifies human preference.





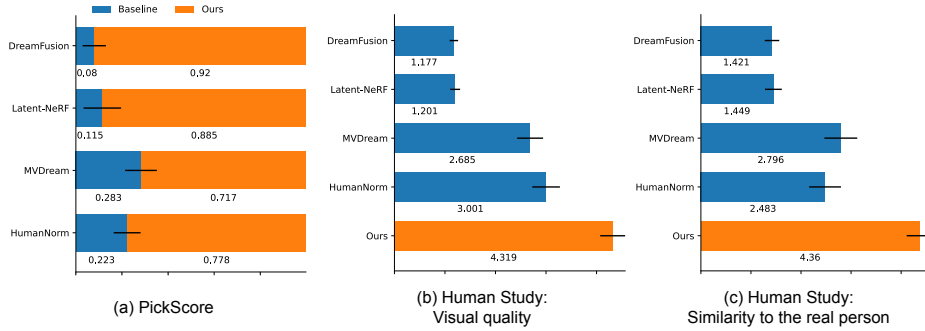
**Fig. 5:** Ablation studies. We show that a geometric prior (5a) improves the results (top-left) even when the geometry prior comes from a different subject (center-left). (top-right) Our method yields very similar results in the personalized setting, even for very different NeRF initializations (5b). A sufficiently diverse prior (5c) is required for convincing results (middle). (bottom-left) we demonstrate the effectiveness of VSD instead of SDS. (bottom-right) we show how inverting the latents works to a certain extent but it fails for out-of distribution cases.

## 4.2 Quantitative Results

We observe in Fig. 6(bc) that MagicMirror outperforms baselines by a large margin ( $> 1.5$  compared to the best baseline), achieving very high marks on both questions. In addition, Fig. 6(a) illustrates a similar trend as seen in the human evaluation. Our method is chosen over baselines the majority of times.

## 4.3 Qualitative Results

**Mixture of concepts** We illustrate the mixture of concepts technique through mixture of objectives described in Sec. 3.3. In Fig. 7a we show the optimized NeRF under various modulation weights for the "happy" and "sad" concepts, as well as an example of removing "green" from "joker". All NeRF are initialized



**Fig. 6:** Quantitative evaluations of our method. (a) We compute PickScore [40] against each baseline. The Bars indicate the percentage of times that our avatar is preferred w.r.t. the baselines. (bc) We report average human study scores regarding visual quality (b) and similarity to the real person (c) of all baselines. Scores range 1 to 5.



(a) The final optimized results under various modulation weights for two different concepts. Removing a certain visual elements is also possible with this method, for example removing the green color from the Joker re-contextualization for the given identity.



(b) The optimization trajectory from one concept to another. This works well when there are no drastic change in geometry.

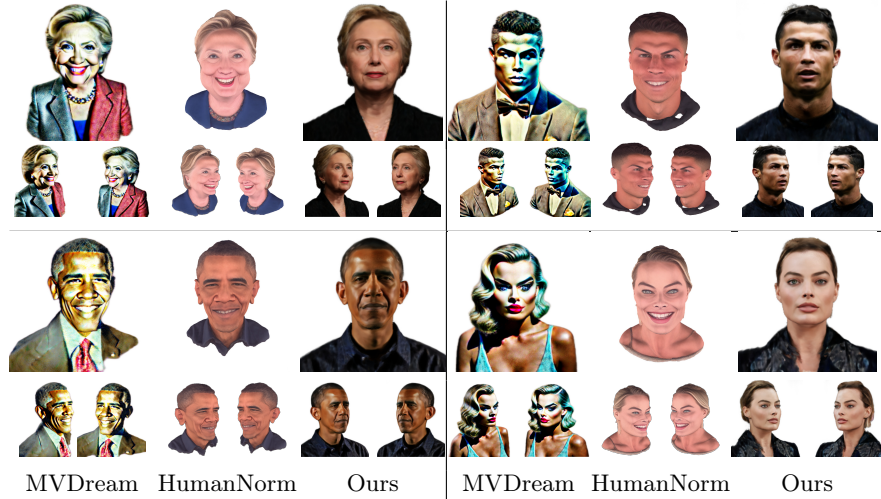
**Fig. 7:** Applying our editing framework with mixture of concepts.

from the common initialization. We can see the mixture results in natural and plausible appearance while retaining the same quality of a single concept. In Fig. 7b we show that it's also possible to move from the concept "young" to "old" by optimizing the target objective based on the optimized source objective, although the trajectory may not always make sense if the intermediate states are too out-of-distribution. Recall that our conditional NeRF is only trained on neutral expression.

**Identity preserving editing with text prompt** AvatarStudio [49] is a recently proposed text guided avatar editing method. It aims to modify the appearance and geometry of the 3D avatar using the SDS technique. The 3D avatar is represented as a conditional NeRF while the conditioning on the time in a expression performance. Thus unlike our approach, there is no modelling of different identities. We compare our results using the same identity and text



**Fig. 8:** We compare our method with AvatarStudio [49]. Please zoom in to see more details.



**Fig. 9:** We compare with MVDream [60] and HumanNorm [29]. MVDream suffers from color over-saturation and HumanNorm struggles with teeth and eyes and yields a cartoon-ish result.

prompt. It’s worth mentioning that we do not have to reconstruct the user at test time, since we only require photos to be used by DreamBooth. Thus unlike AvatarStudio, we don’t need the camera pose estimation that may be hard to obtain with user’s casually captured photos. We observe that we achieve significant improvement in both visual detail and realism, largely benefit from our constrained solution space.

**Non-personalized generation with text prompt** We now switch to celebrity avatar generation using prompt only, and compare with two recent Text-to-3D methods, MVDream [60] and HumanNorm [29]. MVDream integrates multi-view attention mechanism in Stable Diffusion and fine-tunes the model on large 3D

datasets Objaverse. It aims to generate 3D assets of various categories, including humans. HumanNorm focuses on human avatar generation, which fine-tunes Stable Diffusion on depth and normal renderings of 3D human models to guide the geometry generation. Both of them employ DMTet [58] to extract mesh geometry to have better decoupling of geometry and texture, and start the test-time optimization from scratch. The results are visually compared in Fig. 9. We see that MagicMirror achieves much higher realism and fidelity.

## 5 Limitations, Impact, and Conclusion

Although we do not require large scale 3D human data, collecting them for hundreds or thousands of subjects can still be a relatively expensive and time consuming effort. From another point of view, the data we used to constrain the solution space also limit us in the sense that certain extremely out-of-distribution modification is hard to achieve. Our approach can be also limited by the computational resources, since we need multiple Text-to-Image Diffusion Models, at least each for color and normal, and more if we want to perform mixture of concepts. Future research can be invested in more modular design and more direct approach to achieve fast and efficient generation and editing. For its wider adoption, as with all other technologies, we must ensure its development and application to satisfy the security and privacy of the user, and minimize any negative social impact. In particular, we believe the alignment of the pretrained large Text-to-Image Diffusion Models with human values is becoming ever more important given their growing capability and popularity.

We have presented MagicMirror, a next generation text guided 3D avatar generation and editing framework. Through constraining the solution space, looking for a good geometric prior and choosing a good test-time optimization objective, we have achieved a new level of visual quality, diversity and faithfulness. The effectiveness of each component has been demonstrated in our thorough ablation and comparison study. We believe we have made an important step towards an avatar system that people will find it easy and fun to use.

## 6 Acknowledgements

We would like to thank Prof. Octavia Camps and ONR grant N00014-21-1-2431 from NCI for their support.

## References

1. Athar, S., Xu, Z., Sunkavalli, K., Shechtman, E., Shu, Z.: Rignerf: Fully controllable neural 3d portraits. In: CVPR 2022. pp. 20332–20341 (2022) [4](#)
2. Bai, Z., Cui, Z., Liu, X., Tan, P.: Riggable 3d face reconstruction via in-network optimization. In: CVPR 2021. pp. 6216–6225 (2021) [4](#)
3. Bai, Z., Tan, F., Huang, Z., Sarkar, K., Tang, D., Qiu, D., Meka, A., Du, R., Dou, M., Orts-Escolano, S., Pandey, R., Tan, P., Beeler, T., Fanello, S., Zhang, Y.: Learning personalized high quality volumetric head avatars from monocular RGB videos. In: CVPR 2023. pp. 16890–16900 (2023) [4](#)
4. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5470–5479 (2022) [6](#)
5. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: SIGGRAPH 1999. pp. 187–194 (1999) [3](#), [4](#)
6. Bojanowski, P., Joulin, A., Lopez-Paz, D., Szlam, A.: Optimizing the latent space of generative networks. In: Proceedings of the 35th International Conference on Machine Learning. pp. 2640–3498 (2018) [7](#)
7. Bühler, M.C., Sarkar, K., Shah, T., Li, G., Wang, D., Helminger, L., Orts-Escolano, S., Lagun, D., Hilliges, O., Beeler, T., Meka, A.: Preface: A data-driven volumetric prior for few-shot ultra high-resolution face synthesis. In: ICCV 2023. pp. 3379–3390 (2023) [4](#), [6](#)
8. Cao, C., Simon, T., Kim, J.K., Schwartz, G., Zollhöfer, M., Saito, S., Lombardi, S., Wei, S., Belko, D., Yu, S., Sheikh, Y., Saragih, J.M.: Authentic volumetric avatars from a phone scan. *ACM Trans. Graph.* **41**(4), 163:1–163:19 (2022) [4](#)
9. Cao, C., Wu, H., Weng, Y., Shao, T., Zhou, K.: Real-time facial animation with image-based dynamic avatars. *ACM Trans. Graph.* **35**(4), 126:1–126:12 (2016) [4](#)
10. Cao, Y., Cao, Y., Han, K., Shan, Y., Wong, K.K.: Dreamavatar: Text-and-shape guided 3d human avatar generation via diffusion models. *CoRR* **abs/2304.00916** (2023) [5](#)
11. Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., Mello, S.D., Gallo, O., Guibas, L.J., Tremblay, J., Khamis, S., Karras, T., Wetzstein, G.: Efficient geometry-aware 3d generative adversarial networks. In: CVPR 2022. pp. 16102–16112 (2022) [4](#)
12. Chan, E.R., Nagano, K., Chan, M.A., Bergman, A.W., Park, J.J., Levy, A., Aittala, M., Mello, S.D., Karras, T., Wetzstein, G.: Generative novel view synthesis with 3d-aware diffusion models. In: ICCV 2023. pp. 4194–4206 (2023) [5](#)
13. Chaudhuri, B., Vedapant, N., Shapiro, L.G., Wang, B.: Personalized face modeling for improved face reconstruction and motion retargeting. In: ECCV 2020. vol. 12350, pp. 142–160 (2020) [4](#)
14. Chen, R., Chen, Y., Jiao, N., Jia, K.: Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. In: ICCV 2023. pp. 22189–22199 (2023) [5](#)
15. Chen, S., Liu, F., Lai, Y., Rosin, P.L., Li, C., Fu, H., Gao, L.: Deepfaceediting: deep face generation and editing with disentangled geometry and appearance control. *ACM Trans. Graph.* **40**(4), 90:1–90:15 (2021) [4](#)
16. Daněček, R., Black, M.J., Bolkart, T.: EMOCA: Emotion Driven Monocular Face Capture and Animation. In: CVPR (2022) [3](#)
17. Deitke, M., Schwenk, D., Salvador, J., Weihs, L., Michel, O., VanderBilt, E., Schmidt, L., Ehsani, K., Kembhavi, A., Farhadi, A.: Objaverse: A Universe of Annotated 3D Objects. In: CVPR (2023) [2](#)



18. Du, Y., Li, S., Mordatch, I.: Compositional visual generation with energy based models. In: Neural Information Processing Systems (2020), <https://api.semanticscholar.org/CorpusID:214223619> 8, 9
19. Gafni, G., Thies, J., Zollhöfer, M., Nießner, M.: Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In: CVPR 2021. pp. 8649–8658 (2021) 4
20. Gal, R., Patashnik, O., Maron, H., Bermano, A.H., Chechik, G., Cohen-Or, D.: Stylegan-nada: Clip-guided domain adaptation of image generators. *ACM Trans. Graph.* **41**(4), 141:1–141:13 (2022) 5
21. Garrido, P., Zollhöfer, M., Casas, D., Valgaerts, L., Varanasi, K., Pérez, P., Theobalt, C.: Reconstruction of personalized 3d face rigs from monocular video. *ACM Trans. Graph.* **35**(3), 28:1–28:15 (2016) 4
22. Han, X., Cao, Y., Han, K., Zhu, X., Deng, J., Song, Y., Xiang, T., Wong, K.K.: Headsculpt: Crafting 3d head avatars with text. In: NeurIPS 2023 (2023) 5
23. Haque, A., Tancik, M., Efros, A.A., Holynski, A., Kanazawa, A.: Instruct-NeRF2NeRF: Editing 3D Scenes with Instructions. In: CVPR (2023) 2
24. Ho, J., Jain, A., Abbeel, P.: Denoising Diffusion Probabilistic Models (2020) 2
25. Ho, J., Salimans, T.: Classifier-free diffusion guidance. *CoRR* **abs/2207.12598** (2022) 3, 8
26. Hong, F., Zhang, M., Pan, L., Cai, Z., Yang, L., Liu, Z.: Avatarclip: zero-shot text-driven generation and animation of 3d avatars. *ACM Trans. Graph.* **41**(4), 161:1–161:19 (2022) 5
27. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. In: ICLR 2022 (2022) 8
28. Hu, L., Saito, S., Wei, L., Nagano, K., Seo, J., Fursund, J., Sadeghi, I., Sun, C., Chen, Y., Li, H.: Avatar digitization from a single image for real-time rendering. *ACM Trans. Graph.* **36**(6), 195:1–195:14 (2017) 4
29. Huang, X., Shao, R., Zhang, Q., Zhang, H., Feng, Y., Liu, Y., Wang, Q.: Humannorm: Learning normal diffusion model for high-quality and realistic 3d human generation. *CoRR* **abs/2310.01406** (2023) 2, 5, 13
30. Huang, Y., Wang, J., Shi, Y., Qi, X., Zha, Z., Zhang, L.: Dreamtime: An improved optimization strategy for text-to-3d content creation. *CoRR* **abs/2306.12422** (2023) 5
31. Huang, Y., Wang, J., Zeng, A., Cao, H., Qi, X., Shi, Y., Zha, Z., Zhang, L.: Dreamwaltz: Make a scene with complex 3d animatable avatars. In: NeurIPS 2023 (2023) 5
32. Ichim, A.E., Bouaziz, S., Pauly, M.: Dynamic 3d avatar creation from hand-held video input. *ACM Trans. Graph.* **34**(4), 45:1–45:14 (2015) 4
33. Jiang, R., Wang, C., Zhang, J., Chai, M., He, M., Chen, D., Liao, J.: Avatarcraft: Transforming text into neural human avatars with parameterized shape and pose control. In: ICCV 2023. pp. 14325–14336 (2023) 5
34. Jun, H., Nichol, A.: Shap-e: Generating conditional 3d implicit functions. *CoRR* **abs/2305.02463** (2023) 5
35. Karnewar, A., Vedaldi, A., Novotný, D., Mitra, N.J.: HOLODIFFUSION: training a 3d diffusion model using 2d images. In: CVPR 2023. pp. 18423–18433 (2023) 5
36. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **43**(12), 4217–4228 (2021) 4
37. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: CVPR 2020. pp. 8107–8116 (2020) 5



38. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* **42**(4), 139:1–139:14 (2023) [4](#)
39. Kim, H., Garrido, P., Tewari, A., Xu, W., Thies, J., Nießner, M., Pérez, P., Richardt, C., Zollhöfer, M., Theobalt, C.: Deep video portraits. *ACM Trans. Graph.* **37**(4), 163 (2018) [4](#)
40. Kirstain, Y., Polyak, A., Singer, U., Matiana, S., Penna, J., Levy, O.: Pick-a-pic: An open dataset of user preferences for text-to-image generation. *ArXiv abs/2305.01569* (2023), <https://api.semanticscholar.org/CorpusID:258437096> [10](#), [12](#)
41. Kolotouros, N., Alldieck, T., Zanfir, A., Bazavan, E.G., Fieraru, M., Sminchisescu, C.: Dreamhuman: Animatable 3d avatars from text. In: *NeurIPS 2023* (2023) [5](#)
42. Lin, C., Gao, J., Tang, L., Takikawa, T., Zeng, X., Huang, X., Kreis, K., Fidler, S., Liu, M., Lin, T.: Magic3d: High-resolution text-to-3d content creation. In: *CVPR 2023*. pp. 300–309 (2023) [5](#)
43. Liu, N., Li, S., Du, Y., Tenenbaum, J.B., Torralba, A.: Learning to compose visual relations. *ArXiv abs/2111.09297* (2021), <https://api.semanticscholar.org/CorpusID:244270027> [8](#)
44. Liu, N., Li, S., Du, Y., Torralba, A., Tenenbaum, J.B.: Compositional visual generation with composable diffusion models. *ArXiv abs/2206.01714* (2022), <https://api.semanticscholar.org/CorpusID:249375227> [8](#)
45. Liu, R., Wu, R., Hoorick, B.V., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. In: *ICCV 2023*. pp. 9264–9275 (2023) [5](#)
46. Liu, Y., Lin, C., Zeng, Z., Long, X., Liu, L., Komura, T., Wang, W.: Syncdreamer: Generating multiview-consistent images from a single-view image. *CoRR abs/2309.03453* (2023) [3](#), [5](#)
47. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: a skinned multi-person linear model. *ACM Trans. Graph.* **34**(6), 248:1–248:16 (2015) [4](#)
48. Luo, S., Hu, W.: Diffusion probabilistic models for 3d point cloud generation. In: *CVPR 2021*. pp. 2837–2845 (2021) [5](#)
49. Mendiratta, M., Pan, X., Elgharib, M., Teotia, K., R., M.B., Tewari, A., Golyanik, V., Kortylewski, A., Theobalt, C.: Avatarstudio: Text-driven editing of 3d dynamic human head avatars. *ACM Trans. Graph.* **42**(6), 226:1–226:18 (2023) [5](#), [12](#), [13](#)
50. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: *ECCV 2020*. vol. 12346, pp. 405–421 (2020) [4](#)
51. Nichol, A., Jun, H., Dhariwal, P., Mishkin, P., Chen, M.: Point-e: A system for generating 3d point clouds from complex prompts. *CoRR abs/2212.08751* (2022) [5](#)
52. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. In: *ICLR 2023* (2023) [2](#), [3](#), [4](#), [5](#), [8](#)
53. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: *ICML 2021*. vol. 139, pp. 8748–8763 (2021) [2](#), [5](#)
54. Raj, A., Kaza, S., Poole, B., Niemeyer, M., Ruiz, N., Mildenhall, B., Zada, S., Aberman, K., Rubinstein, M., Barron, J., et al.: Dreambooth3d: Subject-driven text-to-3d generation (2023) [4](#)

55. Raj, A., Zollhofer, M., Simon, T., Saragih, J., Saito, S., Hays, J., Lombardi, S.: Pixel-aligned volumetric avatars. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11733–11742 (2021) [4](#)
56. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *CVPR 2022*. pp. 10674–10685 (2022) [5](#)
57. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In: *CVPR 2023*. pp. 22500–22510 (2023) [4](#), [7](#)
58. Shen, T., Gao, J., Yin, K., Liu, M., Fidler, S.: Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. In: *NeurIPS 2021*. pp. 6087–6101 (2021) [14](#)
59. Shen, Y., Yang, C., Tang, X., Zhou, B.: Interfacegan: Interpreting the disentangled face representation learned by gans. *IEEE Trans. Pattern Anal. Mach. Intell.* **44**(4), 2004–2018 (2022) [4](#)
60. Shi, Y., Wang, P., Ye, J., Long, M., Li, K., Yang, X.: Mvdream: Multi-view diffusion for 3d generation. *CoRR* **abs/2308.16512** (2023) [3](#), [5](#), [13](#)
61. Shue, J.R., Chan, E.R., Po, R., Ankner, Z., Wu, J., Wetzstein, G.: 3d neural field generation using triplane diffusion. In: *CVPR 2023*. pp. 20875–20886 (2023) [5](#)
62. Tang, J., Wang, T., Zhang, B., Zhang, T., Yi, R., Ma, L., Chen, D.: Make-it-3d: High-fidelity 3d creation from A single image with diffusion prior. In: *ICCV 2023*. pp. 22762–22772 (2023) [5](#)
63. Tewari, A., Bernard, F., Garrido, P., Bharaj, G., Elgharib, M., Seidel, H., Pérez, P., Zollhöfer, M., Theobalt, C.: FML: face model learning from videos. In: *CVPR 2019*. pp. 10812–10822 (2019) [4](#)
64. Thies, J., Zollhöfer, M., Stamminger, M., Theobalt, C., Nießner, M.: Face2face: Real-time face capture and reenactment of RGB videos. In: *CVPR 2016*. pp. 2387–2395 (2016) [4](#)
65. Trevithick, A., Chan, M.A., Stengel, M., Chan, E.R., Liu, C., Yu, Z., Khamis, S., Chandraker, M., Ramamoorthi, R., Nagano, K.: Real-time radiance fields for single-image portrait view synthesis. *ACM Trans. Graph.* **42**(4), 135:1–135:15 (2023) [4](#)
66. Tsalicoglou, C., Manhardt, F., Tonioni, A., Niemeyer, M., Tombari, F.: Textmesh: Generation of realistic 3d meshes from text prompts. *CoRR* **abs/2304.12439** (2023) [5](#)
67. Wang, C., Jiang, R., Chai, M., He, M., Chen, D., Liao, J.: NeRF-Art: Text-Driven Neural Radiance Fields Stylization. *IEEE TVCG* (2023) [2](#)
68. Wang, C., Jiang, R., Chai, M., He, M., Chen, D., Liao, J.: Nerf-art: Text-driven neural radiance fields stylization. *IEEE Trans. Vis. Comput. Graph.* (01), 1–15 (2023) [5](#)
69. Wang, H., Du, X., Li, J., Yeh, R.A., Shakhnarovich, G.: Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In: *CVPR 2023*. pp. 12619–12629 (2023) [5](#)
70. Wang, T., Zhang, B., Zhang, T., Gu, S., Bao, J., Baltrusaitis, T., Shen, J., Chen, D., Wen, F., Chen, Q., Guo, B.: RODIN: A generative model for sculpting 3d digital avatars using diffusion. In: *CVPR 2023*. pp. 4563–4573 (2023) [5](#)
71. Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J.: Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. In: *NeurIPS 2023* (2023) [3](#), [4](#), [5](#), [6](#), [8](#)
72. Watson, D., Chan, W., Martin-Brualla, R., Ho, J., Tagliasacchi, A., Norouzi, M.: Novel view synthesis with diffusion models. In: *ICLR 2023* (2023) [5](#)

73. Weise, T., Bouaziz, S., Li, H., Pauly, M.: Realtime performance-based facial animation. *ACM Trans. Graph.* **30**(4), 77 (2011) [4](#)
74. Xu, Y., Yang, Z., Yang, Y.: Seeavatar: Photorealistic text-to-3d avatar generation with constrained geometry and appearance. *CoRR* **abs/2312.08889** (2023) [5](#)
75. Zeng, Y., Lu, Y., Ji, X., Yao, Y., Zhu, H., Cao, X.: Avatarbooth: High-quality and customizable 3d human avatar generation. *CoRR* **abs/2306.09864** (2023) [5](#)
76. Zhang, L., Qiu, Q., Lin, H., Zhang, Q., Shi, C., Yang, W., Shi, Y., Yang, S., Xu, L., Yu, J.: Dreamface: Progressive generation of animatable 3d faces under text guidance. *ACM Trans. Graph.* **42**(4), 138:1–138:16 (2023) [5](#)
77. Zheng, Y., Abrevaya, V.F., Bühler, M.C., Chen, X., Black, M.J., Hilliges, O.: I M avatar: Implicit morphable head avatars from videos. In: *CVPR 2022*. pp. 13535–13545 (2022) [4](#)
78. Zhou, Z., Tulsiani, S.: Sparsefusion: Distilling view-conditioned diffusion for 3d reconstruction. In: *CVPR 2023*. pp. 12588–12597 (2023) [5](#)
79. Zhu, P., Abdal, R., Femiani, J., Wonka, P.: Mind the gap: Domain gap control for single shot domain adaptation for generative adversarial networks. In: *ICLR 2022* (2022) [4](#)