

Joker: Conditional 3D Head Synthesis With Extreme Facial Expressions

Malte Prinzler^{1,2} Egor Zakharov³ Vanessa Sklyarova^{1,2} Berna Kabadayi¹ Justus Thies^{1,4}

¹Max Planck Institute for Intelligent Systems, Tübingen

²Max Planck ETH Center for Intelligent Systems

³ETH Zürich ⁴Technical University of Darmstadt

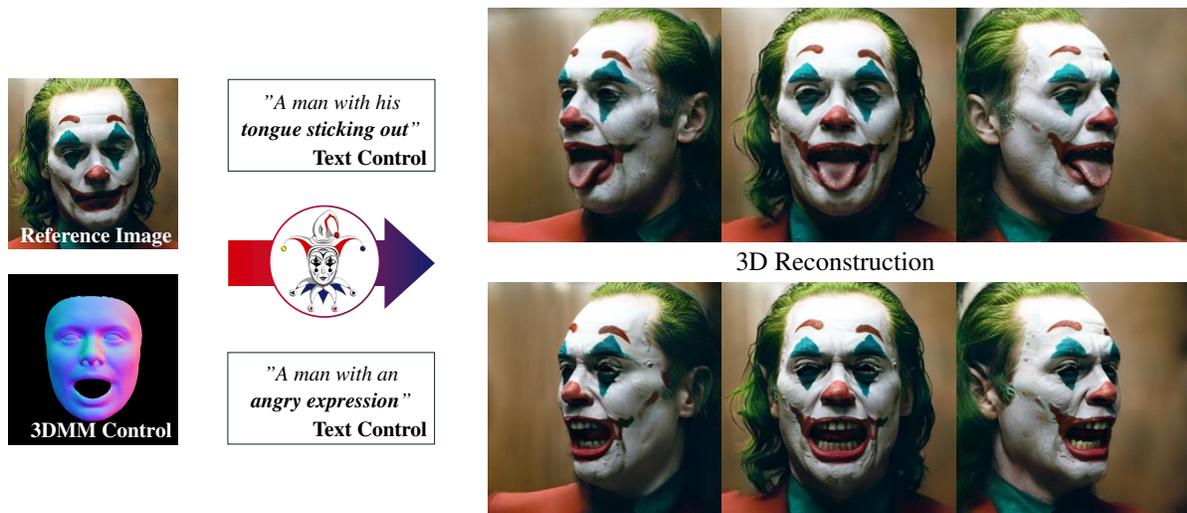


Figure 1. Given a reference image and a target expression defined by a 3DMM and text, Joker generates a 3D model of the subject.

Abstract

We introduce Joker, a new method for the conditional synthesis of 3D human heads with extreme expressions. Given a single reference image of a person, we synthesize a volumetric human head with the reference’s identity and a new expression. We offer control over the expression via a 3D morphable model (3DMM) and textual inputs. This multi-modal conditioning signal is essential since 3DMMs alone fail to define subtle emotional changes and extreme expressions, including those involving the mouth cavity and tongue articulation. Our method is built upon a 2D diffusion-based prior that generalizes well to out-of-domain samples, such as sculptures, heavy makeup, and paintings while achieving high levels of expressiveness. To improve view consistency, we propose a new 3D distillation technique that converts predictions of our 2D prior into a neural radiance field (NeRF). Both the 2D prior and our distillation technique produce state-of-the-art results, which are confirmed by our extensive evaluations. Also, to the best of our knowledge, our method is the first to achieve view-consistent extreme tongue articulation. [Project Page](#)

1. Introduction

Human head avatars have manifold applications in areas such as AR/VR telepresence [33, 42, 53], video games [51, 73], and visual effects [37, 44]. To facilitate downstream applications, techniques for creating 2D head avatars from a single image, often controlled by keypoints, parametric models, and other driving modalities, have been widely studied [14, 15, 25, 49, 57, 60, 63, 64, 66]. However, in many cases, a view-consistent 3D model of a human subject is required to enable rendering in fully virtual environments.

To address this problem, multiple techniques for reconstructing a 3D model of the human subject have been developed [11, 12, 31, 36, 53–55]. Human expressions and appearances, however, have a long-tailed distribution. Mouth cavities and tongues are examples of the areas that are traditionally difficult to capture and, thus, present a challenge for modern avatar systems. Large head rotations are further challenging scenarios, as such examples are missing in most of the existing human-centric datasets [8, 29, 65, 72]. Some approaches [1, 45, 68] addressed this problem by training avatar reconstruction meth-

ods using synthetic datasets of digital human assets rendered via classical graphics pipelines. However, their photorealism and expressiveness still remain subpar. Another group of works [17, 18, 40, 42] extended linear parametric head models [2, 27, 34] with non-linear neural components trained from multi-view datasets [32, 62]. However, collecting such data is expensive, and these approaches can not create avatars with controllable tongues or from a single image. To address these challenges, we propose a novel approach with multi-modal control for extreme expression synthesis. Our method follows an existing line of works on human-centric image synthesis [20, 21, 63] and fine-tunes a pre-trained Stable Diffusion [47] model paired with a ControlNet [69]. We then introduce multi-modal driving inputs to achieve robust and realistic novel-view synthesis of rare and extreme expressions. Our conditioning signal combines the parameters of a 3DMM with a textual prompt. We found that text prompts greatly supplement the control with 3DMM parameters by resolving ambiguities w.r.t. subtle emotional changes and tongue articulations.

View consistency is achieved by optimizing a 3D NeRF [43] from the predictions of our diffusion prior using a novel distillation procedure. Existing 3D distillation approaches [16, 46, 48, 56, 59, 61, 71] exploit 2D diffusion models to predict pseudo-ground-truth target images from noised renders of the current 3D representation. Most of these methods update these target images for every optimization step of the 3D representation – i.e., utilize *dynamic targets*. In a recent concurrent work [16], it was shown that improved performance can be achieved by generating all pseudo ground truth images only once and then optimizing the 3D representation against the generated *static targets*.

We found that neither dynamic nor static target-based approaches yield optimal results for novel-expression 3D distillation. Instead, we propose a new distillation procedure based on *progressively updated targets*. For each time step of a standard DDIM [52] denoising schedule, we use a diffusion-based prior to predict all target images from noised renderings of the 3D representation. We then optimize the 3D representation for several iterations against these target images. Notably, we found it highly beneficial to transition from the dynamically updated target images to static target images at some point during the optimization process. Thus, our proposed progressive distillation method consists of two stages: i) optimization based on dynamically updated targets and ii) optimization based on static targets. We demonstrate that this procedure converges more stably than dynamic-target approaches and is more robust to multi-view inconsistencies than static-target approaches. Ultimately, this yields 3D reconstructions of extreme expressions with high visual fidelity.

To evaluate our method, we collected new benchmark samples that contain extreme expressions in both studio-

capture and in-the-wild environments. Thus, we evaluate our method on three datasets: our proposed extreme expression benchmark, CelebV-Text [65], and NeRsemble [32]. We demonstrate an improved performance compared to existing baselines across all these benchmarks. To train our method, we have also collected new metadata for the above-mentioned datasets, such as textual descriptions of the facial expressions and 3DMM fittings.

In summary, we contribute:

- a 2D diffusion model for single-shot extreme expression synthesis with control through text prompts and 3DMM parameters,
- a 3D distillation approach exploiting progressively updated optimization targets to generate photorealistic 3D reconstructions with extreme expressions,
- a new benchmark and metadata for existing training datasets tailored for extreme expression synthesis.

We plan to make our full codebase, validation dataset, and training dataset metadata publicly available.

2. Related work

We aim for controllable 3D head synthesis, given a reference image and a target expression. This target expression is defined in terms of parameters of a 3DMM [27, 34] with additional textual control to allow the synthesis of expressions and appearances that fall outside the space of the 3DMM. Our approach is based on a learned conditional 2D prior which is used for consistent distillation of the 3D head model.

Conditional 2D Head Synthesis. Modern 2D head synthesis methods rely on human-centric image- and video-based datasets to train generative models capable of directly synthesizing images with novel expressions from a single portrait image. Most of these methods utilize GANs [19] or LDMs [47] to enable high realism of the predictions and generalization capabilities of the trained models. GAN-based reenactment methods [4, 13–15, 49, 50, 57, 66, 67] directly predict an output image given a driving signal and a source image. While achieving high visual quality and expressiveness for frontal-facing images, these methods typically suffer from mode collapse, resulting in low generalization capabilities for extreme expressions and head poses.

Diffusion-based approaches [6, 63] resolve some of these limitations by relying on models that were pre-trained on large-scale data, such as Stable Diffusion [47]. These models can be adapted for human reenactment using a separately trained control network [69]. Such an approach achieves a substantially higher degree of generalization than the GAN-based methods trained from scratch using human-centric image and video data. However, its limitations include a substantial change in the identity of the outputs and a lack of explicit viewpoint control [6, 63]. Moreover,

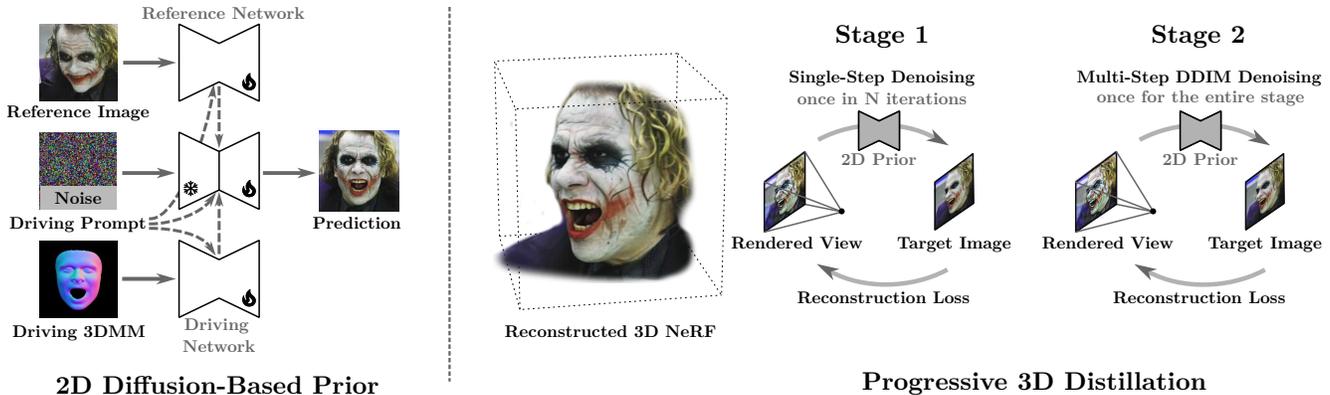


Figure 2. **Method Overview.** We train a 2D diffusion-based prior for novel pose and expression synthesis from a single reference image. It is controlled through text prompts and 3DMM parameters. We leverage this 2D prior to optimize a Neural Radiance Field (NeRF) [43] with a novel two-stage distillation procedure. During Stage 1, the NeRF is optimized against single-step-denoised predictions of the 2D prior that are recalculated every N optimization iterations. In Stage 2, the target images are calculated once in a multi-step denoising process and kept fixed during the NeRF optimization.

these methods use driving modalities such as keypoints [6] or reenacted images produced by a pre-trained GAN-based network [63], which have a limited expressiveness.

Conditional 3D Head Synthesis. 3D-aware head synthesis methods largely address the challenge of view consistency. These methods can also be trained using GAN-based training procedures. Typically, they combine a reconstructed 3D human head model with neural rendering to introduce a high degree of view-consistency [11, 12, 31, 36, 53, 54]. However, these methods have substantial limitations that include a lack of expressiveness and low quality of rendered images. Moreover, they still lack view consistency, especially in high-frequency features, since only the coarse head shape is reconstructed explicitly while super-resolution modules hallucinate the remaining details. These problems are addressed by a growing group of methods that modify pre-trained diffusion models to produce view-consistent renders via viewpoint conditioning [16, 39, 61]. They have been further adapted to human avatar synthesis and can be trained to include explicit pose control [3, 7, 24]. However, since the views are still predicted in image space directly, the 3D consistency of these methods is subpar. While some approaches [20, 63] attempted to resolve this issue with multi-view-aware denoising techniques, the improvements still remain limited.

An alternative approach is the distillation of pre-trained diffusion models into 3D representations using score distillation sampling (SDS) [28, 46, 48, 56, 59], which was used in several previous works on human avatar synthesis [26, 38]. However, these methods either fall short in terms of realism [46] or are unstable w.r.t. the choice of the base diffusion model, i.e. its denoising scheme and hyper-parameters, such as classifier-free guidance scale [56, 59].

Contrary to the existing approaches for novel expression

synthesis, our method utilizes a progressive optimization strategy of the underlying neural radiance field (NeRF) [43]. First, we utilize dynamically updated targets for supervision via single-step denoising, akin to a classical SDS, which results in a blurry yet consistent reconstruction. Once we achieve a coarse reconstruction, we use it to produce fixed optimization targets using multi-step denoising, which helps to complement the missing high-frequency details. Compared to using fixed ground-truth targets throughout the distillation process, as in the concurrent work [16], we achieve a substantially higher degree of view consistency, especially for the mouth cavity and tongue.

3. Method

Our method takes a portrait image as an input and creates a photo-realistic 3D reconstruction with a driving expression specified via parameters of the Basel Face Model [27] and text prompts, see Figure 2. We train a conditional 2D prior that takes an image of a reference person as input and predicts its novel-view renders with novel expressions (Section 3.1). Exploiting this 2D prior and its predictions, we propose a novel 3D distillation pipeline (Section 3.2) to optimize a view-consistent NeRF.

3.1. 2D Prior for Extreme Expression Synthesis

Our prior model is based on a Stable Diffusion [47] backbone. To convert it into a conditional synthesis model, we follow [5, 7, 20, 63] and train a separate *reference network* to input the information from the reference image into the denoising network. Specifically, we share keys and values between the self-attention layers of the reference network and its denoising counterpart [20]. We then train a ControlNet [69], which we refer to as a *driving network*, to condi-

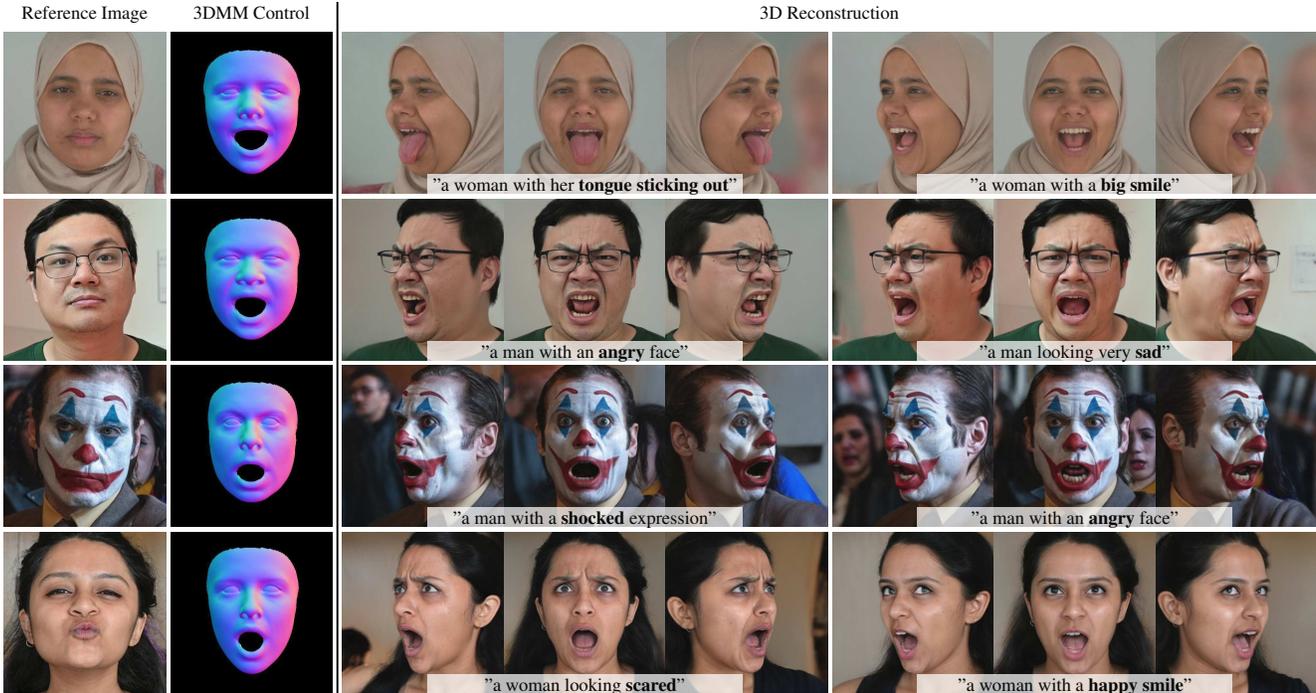


Figure 3. **3DMM- and text-guided 3D reconstruction.** Through text guidance our model resolves ambiguities in the 3DMM control signal, can formulate tongue articulation, and provides fine-grained emotion control. Note that the 3DMM input is kept fixed for both 3D reconstructions of each row and only the text prompt changes.

tion our model on the parameters of a 3DMM. Compared to previous methods [20, 63] and similar to a concurrent work [7], we utilize mesh-based renderers of the normal maps to encode the driving head pose and expression. We have observed that this conditioning style helps to achieve higher view consistency of the results even without view-aware noise or multi-view self-attention techniques [20].

In contrast to existing works [3, 7, 63], we also preserve text-based conditioning from the base denoising model and incorporate it into the reference and driving networks. Text conditioning allows our model to synthesize extreme expressions by supplementing 3DMM-based signals with missing cues, such as tongue movements and emotion-related appearance details. Also, contrary to previous diffusion-based reenactment methods [7, 63], we found it beneficial to fine-tune the decoding part of the denoising network to improve identity preservation. To train this model, we implemented an iterative training pipeline that utilizes a subset of the in-the-wild CelebV-Text dataset [65] followed by a short fine-tuning phase on the NeRSemble dataset [32]. We supplemented these datasets with new metadata that includes textual annotations of the expressions and 3DMM fittings (see suppl. doc.).

3.2. 3D Distillation

Following [48, 56], we use a NeRF [43] formulation to represent our 3D reconstruction. During distillation, the NeRF

is rendered under several target views and encoded into the latent space of the diffusion model. Then, noise is added to these renders, and the 2D diffusion prior denoises the latents and decodes them into the target images \hat{x}_0 that the NeRF is optimized against.

Following previous works [48, 56, 61], we apply the distillation losses directly in the image domain rather than the latent space. The optimization objective follows [61] and consists of a combination of an L1 distance and a perceptual loss \mathcal{L}_p [70] between the rendered images x and targets \hat{x}_0 :

$$\mathcal{L}_{\text{recon}} = \mathbb{E}_{\mathbf{c}} [\|x - \hat{x}_0\|_1 + \mathcal{L}_p(x, \hat{x}_0)], \quad (1)$$

where the expectation is taken over the viewing angles.

Our distillation procedure consists of two stages that differ in the way that the target images are updated. For Stage 1, the target images are repeatedly updated based on the NeRF renderings to improve their multi-view consistency. For Stage 2, the target images are predicted only once and the NeRF is optimized against them until convergence.

Stage 1: Dynamic Target Optimization. During the first stage, we frequently update the target images using the noised NeRF renderings as input to the distillation prior.

We render the NeRF under several target views, apply noise to the renders, and then use the 2D prior to generate the new target images in one denoising step. The NeRF is optimized against the target views for N iterations, after

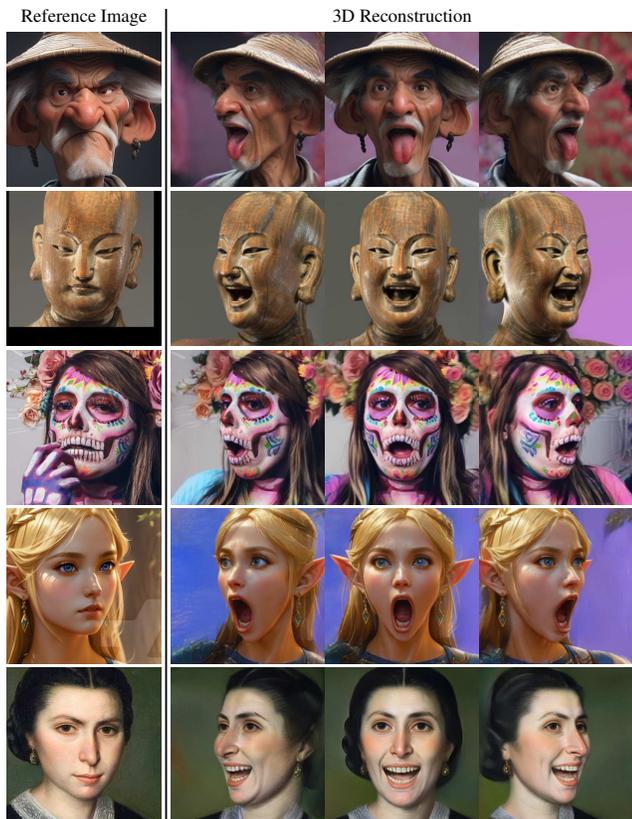


Figure 4. Out-of-distribution 3D reconstruction examples.

which we recalculate the targets and repeat the procedure. We use 3D-consistent NeRF renderings to update the target images to improve their view consistency. The main difference compared to the previous approaches [48, 56, 61] is that we sample the noise levels deterministically following a standard DDIM denoising scheduler with 100 steps, instead of sampling them randomly. Furthermore, the NeRF is optimized against the target images for N iterations before updating them whereas most existing methods update the targets at every NeRF update. We found that this substantially improves the visual quality and consistency of the reconstruction results (see Figure 6).

Stage 2: Fixed Target Optimization. Performing Stage 1 distillation alone tends to drift towards blurry results. The reason for this lies in the NeRF optimization against the target images. During optimization, the NeRF effectively averages over the inconsistencies in the target images and introduces a low-frequency bias. This has a significant impact on the distillation procedure. After the NeRF optimization, its renderings will be more blurry than the target images. The lack of high-frequency details in the NeRF renderings is picked up by the 2D prior and propagates into the updated target images. Optimizing the NeRF against them leads to even more blurry reconstruction results and intro-

duces a positive feedback loop.

To effectively solve this phenomenon, we interrupt Stage 1 at an intermediate noise level (after 60 out of 100 denoising steps) and generate the final target images through standard DDIM sampling with the 40 remaining steps. Afterward, the NeRF is optimized against the final target images until convergence. While the optimization procedure against the fixed target images has a low-frequency bias as well, we avoid the repeated low-frequency feedback loop of Stage 1 and converge to high-quality results.

Note that Stage 2 of our distillation procedure is similar to static-target approaches, with the difference that we start the denoising process not from white noise but from view-consistent renderings of a well-converged NeRF. This largely improves the view consistency of the final target images, while the multi-step denoising procedure generates high-frequency details. As a consequence, the predictions of the 2D prior combine high view consistency and image quality which results in superior NeRF optimization results (see Figure 6).

3.3. Implementation Details

We train the 2D prior starting from pretrained weights of Stable Diffusion v1.5¹. We use the AdamW optimizer [41] with $lr = 10^{-5}$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$, $\lambda = 10^{-2}$. Our model is trained on 8 NVIDIA A100 SXM4-80GB with a per-GPU batch size of 10 for 200,000 iterations on CelebV-Text [65] and 30,000 more iterations on an equal mix of CelebV-Text and NerSemble [32]. We refer to the suppl. doc. for dataset and preprocessing details.

For the 3D distillation, we use the threestudio framework [22] and largely follow the configuration of ImageDream [56]. To generate the target images, we render BFM [27] normal maps on a regular 20×20 grid of the frontal hemisphere with an azimuth $\in [-22.5, 22.5]$ and elevation $\in [-10, 10]$. The images are generated with a classifier-free guidance scale of 19.0. We optimize the NeRF for $N = 130$ iterations between each target image update. During optimization, we randomly sample 64 patches of size 64×64 and gradually increase the resolution of the target images from 64 to 512. The optimization takes approximately 3 hours on a single NVIDIA A100 SXM4-80GB GPU.

4. Experiments

Below we evaluate both our 2D prior and our 3D distillation technique. Please refer to the suppl. doc. for detailed ablation studies and additional experiments. To evaluate the synthesis of the extreme facial expressions, the validation sets of CelebV-Text and NerSemble alone are not suffi-

¹huggingface.co/runwayml/stable-diffusion-v1-5

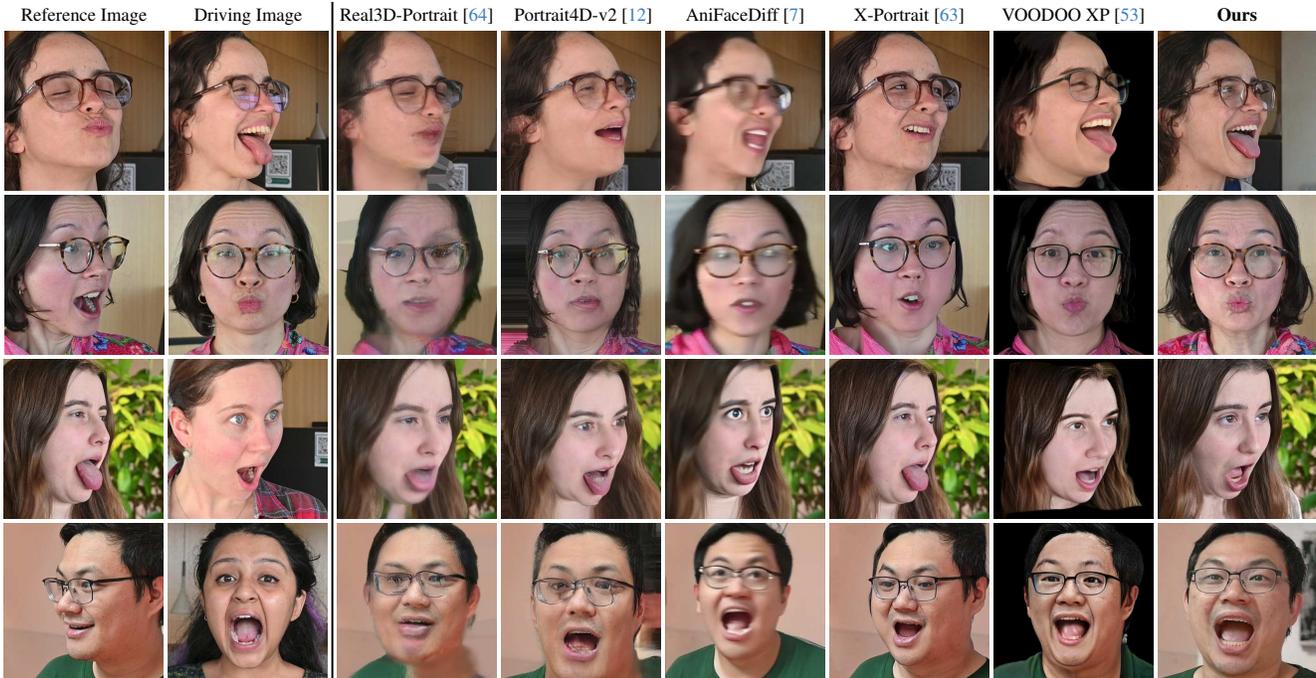


Figure 5. Comparison of our **2D diffusion prior** for self- and cross-reenactment (row 1-2 and 3-4 respectively).

	Self-reenactment								Cross-reenactment			
	PSNR \uparrow	LPIPS \downarrow	SSIM \uparrow	FID \downarrow	CSIM \uparrow	AKD \downarrow	AED \downarrow	APD \downarrow	FID \downarrow	CSIM \uparrow	AED \downarrow	APD \downarrow
GOHA [35]	16.25	0.299	0.572	46.33	0.32	0.015	0.20	0.42	46.82	0.36	0.26	0.51
Real3D-Portrait [64]	16.58	0.322	0.592	31.68	0.46	0.021	0.23	0.47	31.54	0.57	0.30	0.59
Portrait4D-v2 [12]	14.13	0.393	0.494	19.96	0.49	0.081	0.18	0.38	20.71	0.56	0.25	0.52
AniFaceDiff [7]	16.24	0.360	0.577	29.64	0.50	0.034	0.17	0.41	29.34	0.53	0.25	0.52
X-Portrait [63]	14.77	0.357	0.493	10.66	0.61	0.051	0.20	0.68	10.79	0.75*	0.29	0.81
VOODOO 3D [54]	16.11	0.324	0.558	38.63	0.27	0.035	0.19	0.47	39.24	0.30	0.24	0.54
VOODOO XP [53]	13.74	0.397	0.483	24.59	0.49	0.075	0.15	0.45	24.83	0.43	0.21	0.52
Ours	18.63	0.212	0.619	7.57	0.62	0.007	0.11	0.31	8.48	0.57	0.22	0.49

Table 1. Quantitative comparison of our **2D diffusion prior** in self- and cross-reenactment scenarios. *: Note that for extreme pose changes, X-Portrait has a tendency to reproduce the reference image without adopting the driving pose, leading to a high identity similarity score (CSIM) yet poor pose accuracy (APD).

cient: CelebV-Text only contains moderately extreme expressions, and NeRSemble is restricted to uniform lighting and background scenarios. For this reason, we captured the *Joker benchmark* for evaluation of extreme expression synthesis, which will be made publicly available to the research community, see suppl. doc.

4.1. Evaluation of 2D Prior

We compare our 2D diffusion prior against the following baselines: *GOHA* [35], *VOODOO 3D* [54], *VOODOO XP* [53], *Real3D-Portrait* [64], *Portrait4D-v2* [12], *AniFaceDiff* [7], and *X-Portrait* [63]. We refer to the suppl. doc. for a detailed discussion of those and their implementation details. For the baseline results, we used the official code repository of GOHA, VOODOO 3D, Real3D-Portrait,

Portrait4D, and X-Portrait; and obtained the results for AniFaceDiff and VOODOO XP from the authors.

Figure 5 qualitatively compares our 2D diffusion prior with the baselines for the scenario of self- and cross-reenactment. We visualize the best-performing baselines and refer the reader to the suppl. doc. for further results. The comparisons are conducted on the newly captured *Joker benchmark*, which contains in-the-wild scenes with diverse backgrounds, subject ethnicities, and genders. We observe a significant qualitative improvement over all baselines.

These results are confirmed by the quantitative comparison, see Table 1. We evaluate standard metrics such as peak signal-to-noise ratio (PSNR), learned perceptual image patch similarity [70] (LPIPS), structural similarity index measure [58] (SSIM), and Fréchet inception dis-

	LPIPS ↓	SSIM ↑	PSNR ↑	MSE ↓
ProlificDreamer* [59]	0.43	0.52	15.6	0.028
ImageDream* [56]	0.25	0.79	20.2	0.011
Ours, Stage 1 Only	0.27	0.83	20.1	0.012
Ours, Stage 2 Only	0.18	0.81	22.0	0.007
Ours	0.19	0.82	21.5	0.008

Table 2. Quantitative comparison of our **3D distillation** approach. *: Replacing the method’s 2D prior with our model for fair comparison.

tance [23] (FID). We measure identity preservation by comparing the cosine similarity between the embeddings of a face recognition network [9] for the predicted and ground truth images (CSIM). Further, we report the average distance between the extracted keypoints (AKD), expression (AED), and pose (APD) parameters using [10]. For the cross-reenactment scenario where no ground truth data is available, we evaluate the CSIM score between the reference image and the prediction and the AED and APD scores between the driving image and prediction.

Since VOODOO 3D, VOODOO XP, and GOHA do not synthesize the background, we mask out the backgrounds of the other methods using MODNet [30] before evaluation. The quantitative comparison is conducted on 10,000 images for self- and cross-reenactment respectively, evenly sampled from the validation sets of CelebV-Text and NeRSemble, our recordings in a studio environment with uniform lighting and background, and our recordings with in-the-wild backgrounds and lighting. Table 1 shows a consistent improvement over all existing methods. Note that VOODOO XP, VOODOO 3D, Real3D-Portrait, GOHA, and Portrait4D are single-step methods that enable real-time inference, while AniFaceDiff, X-Portrait, and our method are diffusion models that require several denoising steps. Further, X-Portrait was trained on video sequences and with a different crop size than our evaluation samples. However, at the time of writing this paper, no code was available to retrain the model.

Note that the very high CSIM identity similarity score of X-Portrait for the cross-reenactment scenario is misleading. We found that for extreme pose changes, X-Portrait tends to reproduce the reference image without adopting the driving pose (see row 3 of Figure 5). Since for the cross-reenactment scenario, we calculate CSIM between the reference image and the prediction, this artifact results in a significantly overestimated CSIM score. This effect is confirmed by the comparatively high pose reconstruction error (APD) of X-Portrait in Table 1. Our method excels in high-fidelity synthesis and identity preservation while being robust w.r.t. extreme expressions and poses both in the reference and driving image.

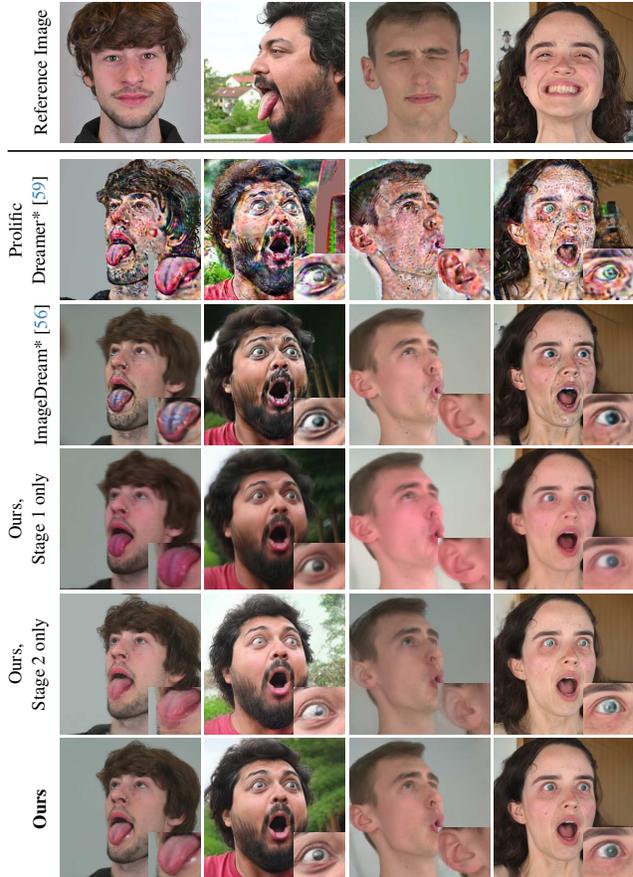


Figure 6. Comparison of **3D reconstructions** from different distillation procedures. *: Replacing the method’s 2D prior with our model for fair comparison.

4.2. Evaluation of 3D Distillation

We compare our novel distillation approach against two state-of-the-art baselines: *ImageDream* [56] and *ProlificDreamer* [59]. *ImageDream* [56] uses a diffusion prior that predicts consistent multi-view images given a reference image and exploits this prior to perform multi-view score distillation sampling. *ProlificDreamer* [59] generalizes score distillation sampling to variational score distillation by treating the NeRF renderings as random variables approximated by a pose-conditioned diffusion model that is fine-tuned on the NeRF renderings during distillation. For *ImageDream*, we use the official code base, for *ProlificDreamer* we use the threestudio implementation [22]. For a fair comparison, we replace the 2D diffusion priors of both baselines with our own prior. We further compare against versions of our method that only use Stage 1 and Stage 2 respectively. Note that the Stage-2-only setting is similar to Cat3D [16]. However, Cat3D only considers a novel view synthesis scenario where reference images of the same scene are given. Instead, in our scenario the refer-

ence and output images differ drastically due to strong pose and expression changes.

Figure 6 presents a qualitative comparison of the distillation procedures on samples from the NerSembler validation set and our self-captured *Joker benchmark* with in-the-wild scenarios. We observe that ProlificDreamer exhibits instabilities during distillation. They are caused by inaccuracies in the fine-tuned diffusion prior that approximates the probability distribution of the NeRF renderings. ImageDream converges more stably, yet artifacts remain since even at the end of the distillation procedure, denoising timesteps are sampled randomly including high noise levels leading to inaccurate estimates of the target images for the optimization of the NeRF. These artifacts are resolved using the Stage 1 optimization of our approach that follows a deterministic denoising schedule. However, using only Stage 1 yields blurry synthesis results. As discussed in Section 3.2, the low-frequency bias of Stage 1 stems from a repeated interplay between NeRF optimization and target image prediction. When performing Stage 2 only, we suppress this effect and can distill a NeRF with high-frequency details. However, note that in Stage 2 all target images are generated at once and the NeRF is optimized against fixed targets. The inconsistencies in these images cause synthesis artifacts like semi-transparencies and misalignment artifacts in the eyes and around the silhouette. Please refer to the suppl. video for a dynamic comparison of the distillation results.

In Table 2, we also perform a quantitative comparison on 30 samples from the NerSembler validation set which provides multi-view ground truth images. Our distillation approach consistently improves over the baselines. We observe that while using both stages of our method qualitatively yielded the best combination of view consistency and high-frequency detail, using only Stage 2 even improves the scores slightly, however, at the cost of view consistency. Please refer to the suppl. video for a dynamic visualization of this effect. The suppl. doc. further provides an evaluation of the impact of the classifier-free guidance scale and the ratio between Stage 1 and Stage 2. We observe that our distillation generates plausible geometry and generalizes well to challenging out-of-distribution samples (see Figure 4).

4.3. Text-Guided Expression Synthesis

Figure 3 demonstrates the effectiveness of using text prompts to control the 3D reconstruction and disambiguate the control through 3DMM parameters: Each row presents two reconstruction results that use the same reference image and 3DMM parameters but different text prompts. We find that text prompts provide an intuitive control mechanism to specify the target emotion and tongue articulation. Please refer to the suppl. doc. for a comparison against a model without text control.



Figure 7. Failure cases of our method.

4.4. Limitations

Figure 7 visualizes failure cases of our method. We observe that implausible colors may be synthesized for challenging out-of-distribution samples in face regions that are not visible in the reference image. In rare cases, we find that even for in-distribution samples the high cfg value (19.0) during distillation causes unnatural colorizations. Please refer to the suppl. doc. for an ablation study on this parameter during distillation. For samples with a uniform background, we observe a tendency of our model to project them to the NerSembler lighting setting: compare rows 2 (NerSembler) and 3 (*Joker benchmark*) of Figure 7. Lastly, particularly for dark curly hair, we observe a low-frequency bias. While using only Stage 2 of our distillation procedure can reduce this effect, this comes at the cost of reduced 3D consistency. Further, note that we only consider static scenes in our work. Extending it to 4D avatar synthesis is a fascinating topic for future research.

5. Conclusion

We introduced *Joker*, a novel method for conditional 3D human head synthesis with extreme expressions from a single reference image. Based on control through 3DMM parameters and text prompts, our method produces high-quality results and generalizes well to out-of-distribution samples. The foundation of this approach is a 2D diffusion-based prior which is learned on in-the-wild imagery of human faces. We leverage this prior to progressively distill a 3D volumetric representation of the target subject with a different facial expression. The textual description allows us to specify the facial expression state beyond the parameters of the 3DMM, including subtle emotional changes, as well as extreme expressions with protruding tongue. We believe that *Joker* is a stepping stone for creating high-resolution 3D content of people with a high degree of identity preservation and emotional expressiveness.

References

- [1] ItSeez3D AvatarSDK, <https://avatarsdk.com>. 1
- [2] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, 1999. 2
- [3] Stella Bounareli, Christos Tzelepis, Vasileios Argyriou, Ioannis Patras, and Georgios Tzimiropoulos. Diffusionact: Controllable diffusion autoencoder for one-shot face reenactment. 2024. 3, 4
- [4] Egor Burkov, Igor Pasechnik, Artur Grigorev, and Victor Lempitsky. Neural head reenactment with latent pose descriptors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13786–13795, 2020. 2
- [5] Di Chang, Yichun Shi, Quankai Gao, Jessica Fu, Hongyi Xu, Guoxian Song, Qing Yan, Xiao Yang, and Mohammad Soleymani. Magicedance: Realistic human dance video generation with motions & facial expressions transfer. *arXiv preprint arXiv:2311.12052*, 2023. 3
- [6] Di Chang, Yichun Shi, Quankai Gao, Jessica Fu, Hongyi Xu, Guoxian Song, Qing Yan, Yizhe Zhu, Xiao Yang, and Mohammad Soleymani. Magicpose: Realistic human poses and facial expressions retargeting with identity-aware diffusion, 2024. 2, 3
- [7] Ken Chen, Sachith Seneviratne, Wei Wang, Dongting Hu, Sanjay Saha, Md. Tarek Hasan, Sanka Rasnayaka, Tamasha Malepathirana, Mingming Gong, and Saman Halgamuge. Anifacediff: High-fidelity face reenactment via facial parametric conditioned diffusion models, 2024. 3, 4, 6
- [8] Joon Son Chung, Arsha Nagrani, and Andrew Senior. Voxceleb2: Deep speaker recognition. *Interspeech 2018*, 2018. 1
- [9] 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. 7
- [10] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *IEEE Computer Vision and Pattern Recognition Workshops*, 2019. 7
- [11] Yu Deng, Duomin Wang, Xiaohang Ren, Xingyu Chen, and Baoyuan Wang. Portrait4d: Learning one-shot 4d head avatar synthesis using synthetic data. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 1, 3
- [12] Yu Deng, Duomin Wang, and Baoyuan Wang. Portrait4d-v2: Pseudo multi-view data creates better 4d head synthesizer. *arXiv preprint arXiv:2403.13570*, 2024. 1, 3, 6
- [13] Michail Christos Doukas, Stefanos Zafeiriou, and Viktoriia Sharmanska. Headgan: One-shot neural head synthesis and editing. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [14] Nikita Drobyshev, Jenya Chelishchev, Taras Khakhulin, Aleksei Ivakhnenko, Victor Lempitsky, and Egor Zakharov. Megaportraits: One-shot megapixel neural head avatars. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 2663–2671, 2022. 1
- [15] Nikita Drobyshev, Antoni Bigata Casademunt, Konstantinos Vougioukas, Zoe Landgraf, Stavros Petridis, and Maja Pantic. Emoportraits: Emotion-enhanced multimodal one-shot head avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1, 2
- [16] Ruiqi Gao*, Aleksander Holynski*, Philipp Henzler, Arthur Brussee, Ricardo Martin-Brualla, Pratul P. Srinivasan, Jonathan T. Barron, and Ben Poole*. Cat3d: Create anything in 3d with multi-view diffusion models. *arXiv*, 2024. 2, 3, 7
- [17] Simon Giebenhain, Tobias Kirschstein, Markos Georgopoulos, Martin Rünz, Lourdes Agapito, and Matthias Nießner. Monophm: Dynamic head reconstruction from monocular videos. *ArXiv*, abs/2312.06740, 2023. 2
- [18] Simon Giebenhain, Tobias Kirschstein, Markos Georgopoulos, Martin Rünz, Lourdes Agapito, and Matthias Nießner. Learning neural parametric head models. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [19] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 2
- [20] Yuming Gu, Hongyi Xu, You Xie, Guoxian Song, Yichun Shi, Di Chang, Jing Yang, and Lingjie Luo. Diffportrait3d: Controllable diffusion for zero-shot portrait view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 3, 4
- [21] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *International Conference on Learning Representations*, 2024. 2
- [22] Yuan-Chen Guo, Ying-Tian Liu, Ruizhi Shao, Christian Laforte, Vikram Voleti, Guan Luo, Chia-Hao Chen, Zi-Xin Zou, Chen Wang, Yan-Pei Cao, and Song-Hai Zhang. threestudio: A unified framework for 3d content generation. <https://github.com/threestudio-project/threestudio>, 2023. 5, 7
- [23] 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. 7
- [24] Hsuan-I Ho, Jie Song, and Otmar Hilliges. Sith: Single-view textured human reconstruction with image-conditioned diffusion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [25] Fa-Ting Hong, Longhao Zhang, Li Shen, and Dan Xu. Depth-aware generative adversarial network for talking head video generation. 2022. 1
- [26] Yangyi Huang, Hongwei Yi, Yuliang Xiu, Tingting Liao, Ji-axiang Tang, Deng Cai, and Justus Thies. TeCH: Text-guided Reconstruction of Lifelike Clothed Humans. In *International Conference on 3D Vision (3DV)*, 2024. 3

- [27] *A 3D Face Model for Pose and Illumination Invariant Face Recognition*, Genova, Italy, 2009. IEEE. 2, 3, 5
- [28] Chenhan Jiang, Yihan Zeng, Tianyang Hu, Songcun Xu, Wei Zhang, Wei Xu, and Dit-Yan Yeung. Jointdreamer: Ensuring geometry consistency and text congruence in text-to-3d generation via joint score distillation. In *ECCV*, 2024. 3
- [29] 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 (CVPR)*, 2019. 1
- [30] Zhanghan Ke, Jiayu Sun, Kaican Li, Qiong Yan, and Rynson W.H. Lau. Modnet: Real-time trimap-free portrait matting via objective decomposition. In *AAAI*, 2022. 7
- [31] Taras Khakhulin, Vanessa Sklyarova, Victor Lempitsky, and Egor Zakharov. Realistic one-shot mesh-based head avatars. In *European Conference of Computer vision (ECCV)*, pages 345–362. Springer, 2022. 1, 3
- [32] Tobias Kirschstein, Shenhan Qian, Simon Giebenhain, Tim Walter, and Matthias Nießner. Nersemble: Multi-view radiance field reconstruction of human heads. *ACM Trans. Graph.*, 42(4), 2023. 2, 4, 5
- [33] Ruilong Li, Kyle Olszewski, Yuliang Xiu, Shunsuke Saito, Zeng Huang, and Hao Li. Volumetric human teleportation. In *ACM SIGGRAPH 2020 Real-Time Live!*, 2020. 1
- [34] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 36(6):194:1–194:17, 2017. 2
- [35] Xueting Li, Shalini De Mello, Sifei Liu, Koki Nagano, Umar Iqbal, and Jan Kautz. Generalizable one-shot neural head avatar. *NeurIPS*, 2023. 6
- [36] Xueting Li, Shalini De Mello, Sifei Liu, Koki Nagano, Umar Iqbal, and Jan Kautz. Generalizable one-shot neural head avatar. *Advances in Neural Information Processing Systems*, 36, 2023. 1, 3
- [37] Yixuan Li, Chao Ma, Yichao Yan, Wenhan Zhu, and Xiaokang Yang. 3d-aware face swapping. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12705–12714, 2023. 1
- [38] Tingting Liao, Hongwei Yi, Yuliang Xiu, Jiayang Tang, Yangyi Huang, Justus Thies, and Michael J Black. Tada! text to animatable digital avatars. *ArXiv*, 2023. 3
- [39] 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, 2023. 3
- [40] Stephen Lombardi, Jason M. Saragih, Tomas Simon, and Yaser Sheikh. Deep appearance models for face rendering. *ACM Transactions on Graphics (TOG)*, 37:1 – 13, 2018. 2
- [41] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization, 2019. 5
- [42] Shugao Ma, Tomas Simon, Jason Saragih, Dawei Wang, Yuecheng Li, Fernando De La Torre, and Yaser Sheikh. Pixel codec avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 64–73, 2021. 1, 2
- [43] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 2, 3, 4
- [44] Jacek Naruniec, Leonhard Helming, Christopher Schroers, and Romann M. Weber. High-resolution neural face swapping for visual effects. *Computer Graphics Forum*, 39, 2020. 1
- [45] Pinscreen, 2024. Pinscreen Avatar Neo, <https://www.avatarnio.com>. 1
- [46] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv*, 2022. 2, 3
- [47] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. 2, 3
- [48] Yichun Shi, Peng Wang, Jianglong Ye, Long Mai, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. *arXiv:2308.16512*, 2023. 2, 3, 4, 5
- [49] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. *Advances in Neural Information Processing Systems*, 32, 2019. 1, 2
- [50] Aliaksandr Siarohin, Oliver Woodford, Jian Ren, Menglei Chai, and Sergey Tulyakov. Motion representations for articulated animation. In *CVPR*, 2021. 2
- [51] Vanessa Sklyarova, Egor Zakharov, Otmar Hilliges, Michael J. Black, and Justus Thies. Haar: Text-conditioned generative model of 3d strand-based human hairstyles. *ArXiv*, abs/2312.11666, 2023. 1
- [52] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. 2
- [53] Phong Tran, Egor Zakharov, Long-Nhat Ho, Liwen Hu, Adilbek Karmanov, Aviral Agarwal, McLean Goldwhite, Ariana Bermudez Venegas, Anh Tuan Tran, and Hao Li. Voodoo xp: Expressive one-shot head reenactment for vr telepresence, 2024. 1, 3, 6
- [54] Phong Tran, Egor Zakharov, Long-Nhat Ho, Anh Tuan Tran, Liwen Hu, and Hao Li. Voodoo 3d: Volumetric portrait disentanglement for one-shot 3d head reenactment. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 3, 6
- [55] Alex Trevithick, Matthew Chan, Michael Stengel, Eric Chan, Chao Liu, Zhiding Yu, Sameh Khamis, Manmohan Chandraker, Ravi Ramamoorthi, and Koki Nagano. Real-time radiance fields for single-image portrait view synthesis. *ACM Transactions on Graphics (TOG)*, 42(4):1–15, 2023. 1
- [56] Peng Wang and Yichun Shi. Imagedream: Image-prompt multi-view diffusion for 3d generation. *arXiv preprint arXiv:2312.02201*, 2023. 2, 3, 4, 5, 7
- [57] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10039–10049, 2021. 1, 2
- [58] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004. 6

- [59] 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. *arXiv preprint arXiv:2305.16213*, 2023. [2](#), [3](#), [7](#)
- [60] Olivia Wiles, A. Sophia Koepke, and Andrew Zisserman. X2face: A network for controlling face generation by using images, audio, and pose codes. *ArXiv*, abs/1807.10550, 2018. [1](#)
- [61] Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P. Srinivasan, Dor Verbin, Jonathan T. Barron, Ben Poole, and Aleksander Holynski. Reconfusion: 3d reconstruction with diffusion priors. *arXiv*, 2023. [2](#), [3](#), [4](#), [5](#)
- [62] Cheng-hsin Wu, Ningyuan Zheng, Scott Ardisson, Rohan Bali, Danielle Belko, Eric Brockmeyer, Lucas Evans, Timothy Godisart, Hyowon Ha, Xuhua Huang, Alexander Hypes, Taylor Koska, Steven Krenn, Stephen Lombardi, Xiaomin Luo, Kevyn McPhail, Laura Millerschoen, Michal Perdoch, Mark Pitts, Alexander Richard, Jason Saragih, Junko Saragih, Takaaki Shiratori, Tomas Simon, Matt Stewart, Autumn Trimble, Xinshuo Weng, David Whitewolf, Chenglei Wu, Shoou-I Yu, and Yaser Sheikh. Multiface: A dataset for neural face rendering. In *arXiv*, 2022. [2](#)
- [63] You Xie, Hongyi Xu, Guoxian Song, Chao Wang, Yichun Shi, and Linjie Luo. X-portrait: Expressive portrait animation with hierarchical motion attention, 2024. [1](#), [2](#), [3](#), [4](#), [6](#)
- [64] Zhenhui Ye, Tianyun Zhong, Yi Ren, Jiaqi Yang, Weichuang Li, Jiangwei Huang, Ziyue Jiang, Jinzheng He, Rongjie Huang, Jinglin Liu, Chen Zhang, Xiang Yin, Zejun Ma, and Zhou Zhao. Real3d-portrait: One-shot realistic 3d talking portrait synthesis. 2024. [1](#), [6](#)
- [65] Jianhui Yu, Hao Zhu, Liming Jiang, Chen Change Loy, Weidong Cai, and Wayne Wu. CelebV-Text: A large-scale facial text-video dataset. In *CVPR*, 2023. [1](#), [2](#), [4](#), [5](#)
- [66] Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-shot adversarial learning of realistic neural talking head models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9459–9468, 2019. [1](#), [2](#)
- [67] Egor Zakharov, Aleksei Ivakhnenko, Aliaksandra Shysheya, and Victor Lempitsky. Fast bi-layer neural synthesis of one-shot realistic head avatars. In *European Conference on Computer Vision*, pages 524–540. Springer, 2020. [2](#)
- [68] Bowen Zhang, Yiji Cheng, Chunyu Wang, Ting Zhang, Jiaolong Yang, Yansong Tang, Feng Zhao, Dong Chen, and Baining Guo. Rodinhd: High-fidelity 3d avatar generation with diffusion models. *arXiv preprint arXiv:2407.06938*, 2024. [1](#)
- [69] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. [2](#), [3](#)
- [70] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [4](#), [6](#)
- [71] Yufeng Zheng, Xueting Li, Koki Nagano, Sifei Liu, Otmar Hilliges, and Shalini De Mello. A unified approach for text- and image-guided 4d scene generation. In *CVPR*, 2024. [2](#)
- [72] Hao Zhu, Wayne Wu, Wentao Zhu, Liming Jiang, Siwei Tang, Li Zhang, Ziwei Liu, and Chen Change Loy. Celebv-hq: A large-scale video facial attributes dataset. In *European conference on computer vision*, pages 650–667. Springer, 2022. [1](#)
- [73] Luyang Zhu, Konstantinos Rematas, Brian Curless, Steven M. Seitz, and Ira Kemelmacher-Shlizerman. Reconstructing nba players. In *European Conference on Computer Vision*, 2020. [1](#)