

A Unified Approach for Text- and Image-guided 4D Scene Generation

Yufeng Zheng^{1,2,3}, Xueting Li¹, Koki Nagano¹, Sifei Liu¹, Otmar Hilliges², Shalini De Mello¹
¹NVIDIA, ²ETH Zurich, ³Max Planck Institute for Intelligent Systems

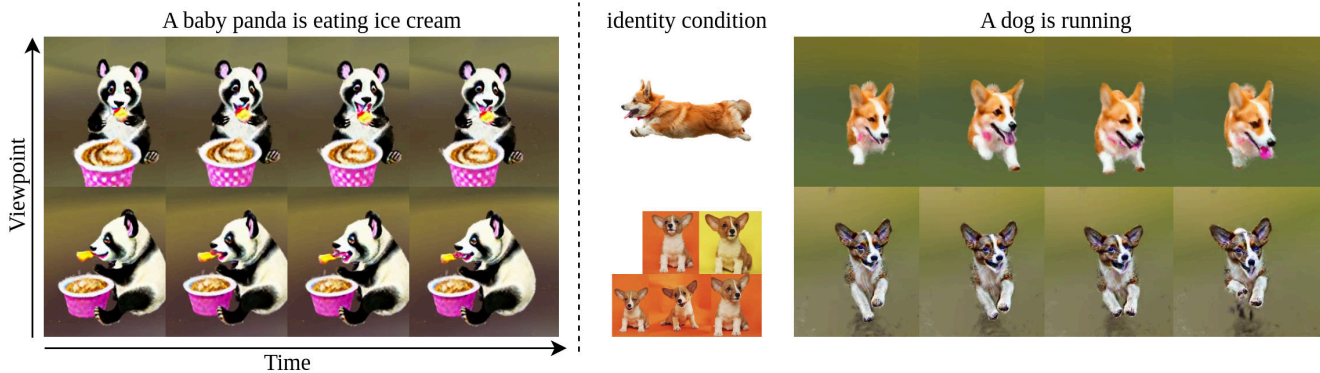


Figure 1. **Text-to-4D**. Our method provides a unified approach for generating 4D dynamic content from a text prompt with diffusion guidance, supporting both unconstrained generation and controllable generation, where appearance is defined by one or multiple images.

Abstract


Large-scale diffusion generative models are greatly simplifying image, video and 3D asset creation from user-provided text prompts and images. However, the challenging problem of text-to-4D dynamic 3D scene generation with diffusion guidance remains largely unexplored. We propose *Dream-in-4D*, which features a novel two-stage approach for text-to-4D synthesis, leveraging (1) 3D and 2D diffusion guidance to effectively learn a high-quality static 3D asset in the first stage; (2) a deformable neural radiance field that explicitly disentangles the learned static asset from its deformation, preserving quality during motion learning; and (3) a multi-resolution feature grid for the deformation field with a displacement total variation loss to effectively learn motion with video diffusion guidance in the second stage. Through a user preference study, we demonstrate that our approach significantly advances image and motion quality, 3D consistency and text fidelity for text-to-4D generation compared to baseline approaches. Thanks to its motion-disentangled representation, *Dream-in-4D* can also be easily adapted for controllable generation where appearance is defined by one or multiple images, without the need to modify the motion learning stage. Thus, our method offers, for the first time, a unified approach for text-to-4D, image-to-4D and personalized 4D generation tasks.

1. Introduction

The advent of large-scale text-conditioned diffusion-based generative models for images has ushered in a new era of imaginative, high-quality image synthesis [1, 28, 30]. Their simple and intuitive conditioning in the form of text prompts are game-changers in democratizing visual content creation for non-expert users. Subsequently, these developments have also led to impressive progress in (1) text- or image-conditioned static 3D content creation [15, 20, 24, 26], achieved by leveraging guidance from generic and 3D-aware [16, 17, 31] image diffusion models, and (2) video content creation via video diffusion models [2–4, 32].

However, for various real-world applications such as gaming, AR/VR, and advertising, synthesizing static 3D assets alone does not suffice. It is desirable to also animate 3D assets using intuitive user-provided text prompts to further save animators’ time and level-of-expertise. To go beyond static 3D content creation, we delve into the largely unexplored problem of text-conditioned 4D scene generation, a.k.a., *text-to-4D* synthesis with diffusion guidance. This is a challenging problem encompassing both text-to-3D and text-to-video synthesis. It requires learning not only a 3D-consistent representation of a static scene capable of free-view rendering, but also its plausible and semantically-

Work was partially done during an internship at NVIDIA. Project page: <https://research.nvidia.com/labs/nxp/dream-in-4d/>



(a) Multi-view images of a generated 3D asset without 3D diffusion. It suffers from the Janus problem. (b) The presence of the Janus problem, further hinders learning of the guidance. It suffers from the Janus cape’s correct temporal motion (top row). Ours

Figure 2. We show the importance of a high-quality static asset for 4D content generation. The prompt for this scene is ‘A superhero dog with a red cape is flying through the sky’.

correct dynamic 3D motion over time.

The pioneering work MAV3D [33] is the first attempt to address this problem. It proposes a two-stage approach: the first to learn a static 3D asset, and the second to optimize its full dynamic representation with guidance from a video diffusion model [32]. It further models the dynamic representation via a neural hexplane [5]. While impressive in demonstrating feasibility, this early work leaves much room for improvement in terms of robustness, quality and realism. Additionally, it does not solve the problems of image-to-4D or personalized-4D content creation, wherein, in addition to a text prompt, an image or a set of images is provided as input to control the appearance of the 4D outcome.

To address these challenges, we propose a novel method for text-to-4D dynamic scene synthesis, named *Dream-in-4D*. It employs a two-stage approach, to first learn a static scene representation and then its motion. Our first primary insight is that achieving high-quality, 3D-consistent static reconstruction in the first stage is crucial for successfully learning motion in the second stage. For example, in Fig. 2 we show a multi-view *inconsistent* 3D dog with two heads due to the Janus problem learned in the first stage. It introduces significant ambiguity for the second dynamic stage and substantially undermines the quality of the learned motion (in the dog’s cape). However, relying solely on guidance from image or video diffusion models for static text-to-3D synthesis, as proposed in [33], easily encounters the Janus problem (see the first row of Fig. 4). Therefore, we leverage 3D-aware [16, 31] and standard image diffusion models [1, 28] in stage-one to achieve high-quality view-consistent text-to-3D synthesis, along with video diffusion guidance [3] in stage-two to learn realistic motion. This forms our first key contribution, which is to leverage guidance from a carefully-designed combination of image, 3D, and video diffusion models to effectively solve the task of text-to-4D synthesis.

Our approach is further motivated by the observation that fine-tuning the static model with video diffusion models in stage-two leads to lower visual quality and prompt fidelity,

primarily because these models are trained with lower quality videos compared to image diffusion models. To address this problem, our insight is to decompose the synthesis process into two *distinct* training stages, the first of which is designed to learn a high-quality static 3D asset and the second dedicated to effectively animating it with the provided text prompt, while keeping the pre-trained static 3D asset unchanged. This, in turn, requires a 4D neural representation that fully disentangles the canonical static representation and its motion. However, this cannot be achieved with the hexplane [5] representation proposed in [33], which entangles the static representation and its motion. To this end, we propose to use a variant of a deformable neural radiance field (D-NeRF) [25] for the task of 4D content generation. A D-NeRF consists of a canonical 3D NeRF [21] and a 4D deformation MLP that maps time-dependent deformed space to the common canonical static space. With the proposed disentangled representation for 4D scene synthesis, we can freeze the pre-trained high-quality static 3D model from stage-one and only optimize the deformation field using video diffusion guidance in stage-two. To successfully learn detailed and realistic motion, we further encode the 4D deformation field with multi-resolution feature grids and regularize motion using a novel total variation loss on the rendered displacement maps. We find that the former enhances detailed motion while the latter reduces spatial and temporal jitter.

Through a user preference study on diverse text prompts, we show that our algorithm achieves significant improvements in visual quality, 3D consistency, prompt matching and motion quality compared to alternate baselines. Furthermore, the ability to disentangle the canonical and motion representation allows for easy adaptation to image-conditioned 4D generation, without requiring modifications to the motion learning stage. Thus, we demonstrate image-to-4D generation given a single-view image, and personalized 4D generation using 4-6 casually captured images of a subject (See Fig. 1) with our unified *Dream-in-4D* method.

In summary, our key contributions include:

1. We propose to combine image, 3D-aware and video diffusion priors for the text-to-4D task, significantly improving the visual quality, 3D consistency and text-fidelity of the learned static assets in the first stage.
2. By explicitly disentangling the static representation from its deformation, our method preserves the high-quality static asset during motion learning.
3. We propose to use a multi-resolution feature grid and a total variation loss on the deformation field to effectively learn motion with video diffusion guidance.
4. We demonstrate that our method offers, for the first time, a unified approach for text-to-4D, image-to-4D and personalized 4D generation tasks.

2. Related Work

Dynamic neural radiance fields. Modeling dynamic 3D content with NeRFs has been extensively studied in the novel-view synthesis literature. To extend NeRFs to dynamic scene modeling, previous works either learn a high-dimensional radiance field conditioned on temporal embeddings [14, 18], or a separate deformation mapping to model motion [23, 25]. To speed up training and inference, plane- and voxel-based feature grids are combined with MLPs to formulate efficient hybrid NeRF representations [6, 11, 22], which are extended to dynamic scene modeling by learning additional planes for the temporal dimension [5, 12]. In this work, we leverage a deformable NeRF representation where the canonical geometry and 4D deformation field are both encoded by multi-resolution feature grids [22]. As a result, the geometry and motion are fully disentangled, which not only eases motion learning but also allows easy adaptation to various applications such as image(s)-to-4D video generation.

Diffusion models. Recently, diffusion models have revolutionized the computer vision community by showing remarkable advancements in image, video or novel view image synthesis. Seminal works such as Stable Diffusion [28] and DeepFloyd [1] take a text prompt as input and produce high-quality images that align with the prompt. Leveraging large-scale image datasets, these diffusion models learn various prior knowledge ranging from object appearance to complex scene layout. One line of subsequent works [2–4] fine-tunes text-to-image diffusion models on video datasets, successfully extending them to generate realistic videos matching both the object and motion described by the input prompt. Another line of methods [16, 17] learns 3D-aware diffusion models with images rendered from synthetic objects [8, 9]. By conditioning on camera parameters, these methods produce novel view images of an object that are consistent with each other and align with the observed view.

Text/image(s) to 3D with diffusion priors. Beyond direct sampling from the diffusion models, several works employ image diffusion priors as an optimization signal for 3D generation. The pioneering work, DreamFusion [24], optimizes a 3D model by presenting its renderings to a text-to-image diffusion model and acquiring gradient supervision through Score Distillation Sampling (SDS). Subsequent works enhance the synthesis quality and speed by incorporating the mesh representation [15, 35], advancing score distillation [36–38], exploring representations in the latent space [20], or disentangling geometry and texture [7]. Yet, these methods suffer from the Janus problem due to the lack of 3D prior in the text-to-image diffusion models. To overcome this limitation, several works [13, 16, 17, 31] leverage 3D-aware diffusion models as supervisions, generating 3D objects that are consistent across multiple views. In addition to text, a few approaches

further take one or multiple images as an input and reconstruct a 3D object matching both the prompt and the image(s). The former [19, 26, 34] simultaneously utilize text-to-image and novel view diffusion models [16], while the latter [27, 31] overfit a diffusion model on a few images depicting the same subject to achieve personalized diffusion guidance.

Text-to-4D with diffusion priors. In this paper, we go beyond 2D/3D generation and aim to synthesize a 4D video given a text prompt. This is hitherto a highly challenging and under-explored domain. The most relevant work to ours is MAV3D [33], which optimizes a 4D scene by leveraging a pre-trained video diffusion model. However, due to the entanglement of geometry and motion, as well as a lack of a 3D-aware prior, this method suffers from the Janus problem and produces low quality texture. To resolve these limitations, we leverage a carefully designed combination of image, video and 3D-aware diffusion models and fully disentangle the geometry and motion. Our method synthesizes multi-view consistent 4D videos with realistic appearance and motion. Furthermore, the disentanglement of canonical and motion representations readily enables novel applications such as image-to-4D and personalized 4D generation.

3. Method

Given a text prompt and optionally one or a few images to specify the object’s appearance, we aim to generate a 4D video that matches both the object and the motion described in the prompt. To this end, we propose a two-stage training pipeline. In the first static stage (Sec. 3.1), we synthesize a high-quality static 3D scene using both 2D and 3D diffusion priors. In the second dynamic stage (Sec. 3.2), we learn the 3D motion of the scene using a video diffusion model, while keeping the static scene representation intact.

3.1. Static Stage

The goal of the static stage is to generate a high-quality 3D scene that aligns with the text prompt. To efficiently learn a canonical model of a 3D scene, we opt for the NeRF [21] representation with multi-resolution hash-encoded features [22], which is extensively used in previous text-to-3D methods [15, 20, 24, 37, 38]. This static 3D model is accompanied by a deformation field to represent the dynamic motion of the 3D scene in the subsequent motion learning stage. Two elements of the static stage are crucial for 3D asset quality and play an important role in facilitating motion learning in the subsequent dynamic stage: the generated 3D object(s) (1) should be view-consistent (*i.e.*, free of the Janus problem) and (2) should follow the spatial composition described in the text prompt. Intuitively, the former reduces contradictory gradients from different views in deformation optimization while the latter eases motion

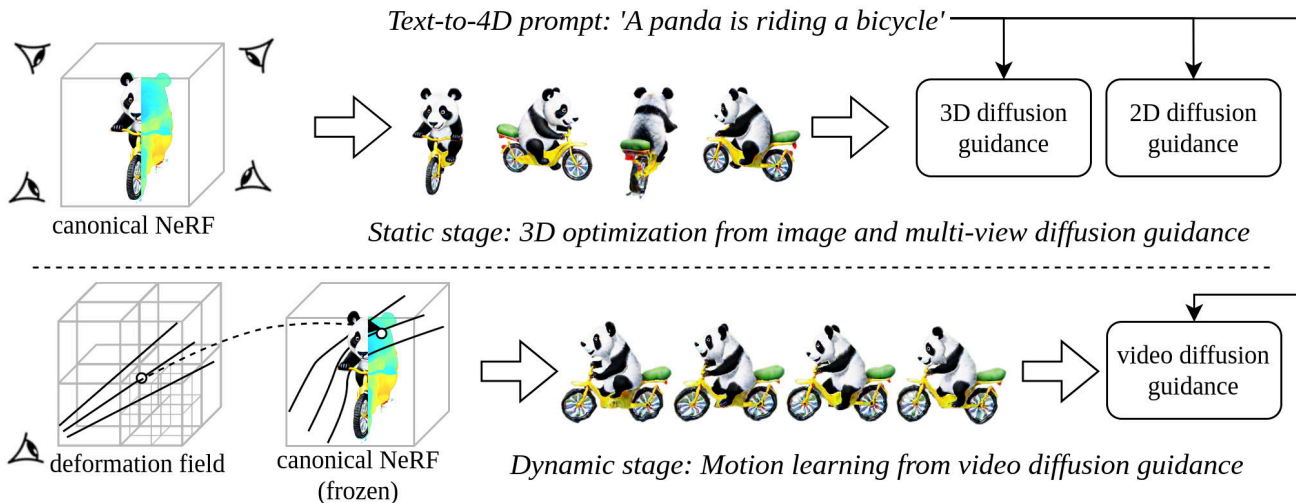


Figure 3. **Method overview.** Adopting a two-stage approach, *Dream-in-4D* first utilizes 3D and 2D diffusion guidance to learn a static 3D asset based on the provided text prompt (top). Then, it optimizes a deformation field using video diffusion guidance to model the motion described in the text prompt (bottom). Featuring a motion-disentangled D-NeRF representation, our method freezes the pre-trained static canonical asset while optimizing for the motion, achieving high quality view-consistent 4D dynamic content with realistic motion.

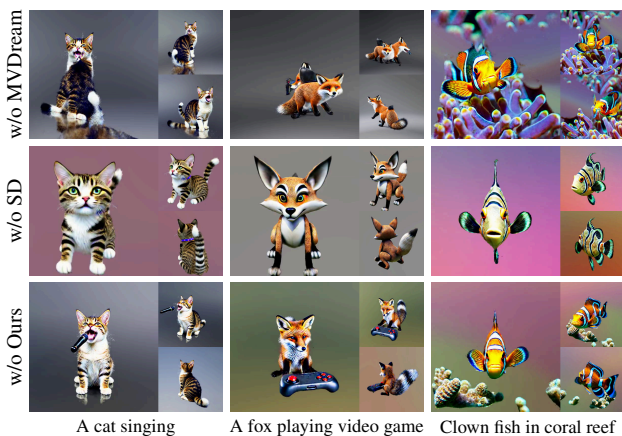


Figure 4. **Static stage.** Without StableDiffusion guidance, the learned static model fails to learn the correct composition. Without MVDream guidance, the learned assets suffer from the Janus problem and contain multiple faces. Using guidance from both StableDiffusion and MVDream, results in the best text prompt fidelity and 3D consistency.

learning by presenting a reasonable spatial layout of multiple objects. For instance, the panda sitting on top of a bike in Fig. 3 sets a good starting point for the dynamic stage to learn the “riding” motion. To achieve these goals, we propose to utilize both 3D and generic 2D diffusion models for the static stage. This is, in spirit, similar to prior work [26] for image-to-3D synthesis. In the following, we introduce the 3D and 2D diffusion guidance used for stage-one.

3D diffusion models [16, 17, 31] take camera parameters with a text prompt or an image as inputs and synthesize

novel view images of the target object. Fine-tuned from image diffusion models with rendered images of synthetic 3D data [10], 3D diffusion models provide valuable prior knowledge of the 3D world and enforce different views of a 3D object to be consistent. For text to 3D generation, we adopt the MVDream [31] model to provide a 3D prior. Specifically, we render the synthesized 3D object from four different viewpoints (*i.e.*, front, back, and two side views) and obtain guidance from a pre-trained MVDream model through the SDS loss [24]. We use a reconstruction formulation of the SDS loss, similar to [31]. The 3D guidance loss is denoted as $\mathcal{L}_{3D}(I)$.

However, due to the limited scale and synthetic nature of the 3D training datasets, static NeRF models optimized using MVDream alone tend to have synthetic-looking texture [31], and occasionally fail to produce realistic scene layouts (see the second row of Fig. 4 where objects in the prompt are missing from the scene). Meanwhile, we observe that image diffusion models trained with large-scale 2D images encourage both realistic appearance and reasonable scene layouts, but by themselves, easily suffer from the Janus problem (see the first row of Fig. 4). Thus, we propose to combine 2D diffusion guidance with 3D guidance in stage-one. Specifically, we use StableDiffusion-v2.1 with the SDS objective (denoted as $\mathcal{L}_{2D}(I)$) to provide 2D guidance. The overall objective for stage-one is:

$$\mathcal{L} = \lambda_{2D}\mathcal{L}_{2D}(I) + \lambda_{3D}\mathcal{L}_{3D}(I),$$

where I denotes the set of rendered images from the sampled camera viewpoints, and $\lambda_{2D/3D}$ are the weights for the 2D and 3D guidances (see Supplement for the loss weights).




Figure 5. **Hexplane v.s. deformable NeRF.** With a hexplane representation, even though the static stage successfully learns a high-quality 3D asset (column 1), its motion learning stage with video diffusion guidance still leads to degradation in texture and re-appearance of the Janus problem (column 2).

As shown in Fig. 4, by combining both the 2D and 3D diffusion models for guidance, our static stage generates 3D-consistent object(s) with realistic texture and plausible scene layouts.

3.2. Dynamic Stage

In the dynamic stage, our goal is to learn a deformation field that animates the 3D scene generated in the static stage using guidance from a video diffusion model. As aforementioned, our key observation is that although video diffusion models provide a valuable motion prior, they are not 3D-aware and tend to produce unappealing visual results (see Fig. 5, column 2). Therefore, we propose to fully disentangle the static model and the motion by freezing the NeRF network learned in the static stage and only learn the deformation field to match the motion described in the text prompt in the dynamic stage. Such a design brings two advantages: (1) it preserves the view consistency and high-quality texture learned in the static stage and (2) it readily enables applications such as image-to-4D and personalized 4D generation (see Sec. 3.3).

Motion-disentangled 4D representation. Our dynamic 4D representation consists of a canonical 3D radiance field (as described in Sec. 3.1) and a deformation field. The deformation field is a 4D to 3D time-dependent mapping $D(\mathbf{x}_d, t) \rightarrow \mathbf{x}_c$, where \mathbf{x}_d is a 3D point’s location in deformed space at time t , and \mathbf{x}_c is its corresponding canonical location. Our insight is that the deformable field should be smooth both spatially and temporally due to the limited elasticity and velocity of the object. As a result, the deformation field does not require as high-resolution a feature grid as its static canonical 3D counterpart. Therefore, we utilize a 4D multi-resolution hash-encoded feature grid with a maximum resolution of 232 for the deformation field, in contrast to the maximum resolution of 4096 for the canonical static NeRF representation. Additionally, we found the usage of multi-resolution features to be crucial for learning correct local motion (see Fig. 6).

Motion optimization with video diffusion models. The deformation field is optimized via score distillation sampling using a video diffusion model. Specifically, we

sample a static camera parameter, and render a 24-frame video \mathbf{V} from our 4D representation. The time stamps are sampled evenly and the length of the video is randomly chosen between 0.8 and 1 (assuming that the full length is 1). We leverage a variant of the SDS loss [38] for video diffusion guidance, where we predict the original video with 1-step denoising, and use a combination of a latent feature loss and a decoded RGB space loss. The video diffusion guidance loss can be expressed as $\mathcal{L}_{video}(\mathbf{V}) = \mathcal{L}_{latent}(\mathbf{V}) + \lambda_{dec}\mathcal{L}_{dec}(\mathbf{V})$, where $\lambda_{dec} = 0.1$. We choose to use the Zeroscope [3] video diffusion model as our motion prior, but our method is robust to other models such as Modelscope [2] (see our project webpage for results). We found that matching the resolution of the learned videos to that of the video diffusion models to be important for successfully distilling motion priors. For Zeroscope, we render videos at a resolution of 144×72 and upsample them to 576×320 when training with video diffusion guidance.

Total variation motion regularization. To reduce temporal and spatial jitter in motion, we propose to use a novel total variation loss for the learned deformation (see Fig. 6). Specifically, in addition to the RGB video \mathbf{V} , we also render a video of 3D displacements \mathbf{D} . The total variation loss on the rendered displacement video \mathbf{D} can be expressed as:

$$\mathcal{L}_{TV}(D) = \sum_{x,y,t} (||\mathbf{D}_{x-1,y,t} - \mathbf{D}_{x,y,t}||_2^2 + ||\mathbf{D}_{x,y-1,t} - \mathbf{D}_{x,y,t}||_2^2 + ||\mathbf{D}_{x,y,t-1} - \mathbf{D}_{x,y,t}||_2^2).$$

The overall objective function for the second stage is then:

$$\mathcal{L} = \mathcal{L}_{video}(\mathbf{V}) + \lambda_{TV}\mathcal{L}_{TV}(\mathbf{D}),$$

where $\lambda_{TV} = 1000$.

3.3. 4D Generation Given One or Multiple Images

While text-to-4D generation is useful for many scenarios, there is a common desire to create content that features a specific object. However, language alone may be insufficient to describe the unique appearance of a given object. Thanks to the full disentanglement of its static and dynamic parts, our method can be easily extended to image-guided 4D generation, without modifying the motion learning stage. In the following, we show that this can be done by simply replacing the diffusion models used in the static stage of our method.

Image-to-4D generation. Given a single image, we reconstruct the corresponding 3D asset by replacing MVDream with an image-conditioned 3D diffusion model. Specifically, we use zero123-xl [16] as our 3D diffusion model and DeepfloydIF [1] as our 2D diffusion model. Additionally, we supervise the reference view with the




Figure 6. **Deformation learning.** The deformation MLP equipped with positional encoding instead of multi-resolution feature grids cannot capture local motions (mouth and fin in row 1). Without the proposed total variation loss on the displacement, the learned deformation contains substantial noise (row 2). Our approach, with both, results in the best quality (row 3).

given image and its estimated foreground mask, similarly to [16, 26]. Fig. 8 shows examples of synthesized 4D videos from a single-view image.

Personalized 4D generation. Given a few casually captured images of an object, Dreambooth [29] finetunes image diffusion models to generate personalized images of this object given a text prompt. By replacing our generic image diffusion model with a finetuned, personalized version, we can create personalized 3D assets given a text prompt and a few casual images. Specifically, we use personalized StableDiffusion together with MVDream for this task. We show synthesized 4D videos in Fig. 9.

4. Results

4.1. Text-to-4D Generation

In Fig. 7, we show qualitative results of our method on text-to-4D generation. Video results are displayed on the project webpage for better assessment of the motion.

User study. We carry out a user preference study to evaluate sample quality along the dimensions of (1) alignment to the input text prompt, (2) motion quality, and (3) 3D consistency and visual quality. For each vote, we present the participant with results from our method as well as from the baseline(s), and ask the participant to pick the best method given one of the three evaluation metrics described above.

Comparison with MAV3D. Since there is no publicly available implementation of MAV3D [33], we compare our method against the 28 visual results displayed on their website. Through a user study with 13 users (results are shown

Metric	visual & 3D	text alignment	motion
Ours	82.4%	65.4%	61.8%

Table 1. **Comparison with MAV3D.** The numbers indicate the percentage of users who prefer our results over MAV3D’s.

in Tab 1), our method outperforms MAV3D in all metrics.

Additionally, we compare with several ablative baselines following a similar user study protocol with 18 users. We detail the ablation study in the following paragraphs.

Metric	w/o 2D guidance	Ours
text alignment	11.67%	88.33%

(a) **Text alignment ablation.** Without 2D diffusion guidance, the learned 3D asset might fail to generate all the required components of a scene or fail to produce a plausible layout (see Fig. 4).

Metric	w/o Multi-res Grid	w/o TV loss	Ours
motion quality	42.22%	2.78%	55.00%

(b) **Motion quality ablation.** Without the multi-resolution feature grid, detailed local motion cannot be learned. Without the proposed TV loss, the generated motion contains substantial noise (see Fig. 6).

Metric	w/o 3D guidance	Hexplane	Ours
visual&3D	6.12%	0.00%	93.88%

(c) **Visual quality and 3D consistency ablation.** Without 3D diffusion guidance in the first stage, the learned 3D assets suffer from the Janus problem (see Fig. 4). With a hexplane 4D representation, the learned high-quality 3D asset cannot be preserved in the dynamic stage, leading to lower visual quality and re-appearance of the Janus problem (see Fig. 5).

Table 2. **Comparison with ablative baselines.**

3D and 2D diffusion guidance in the static stage. Our method combines 3D and 2D diffusion guidance to learn the static model. Fig. 4 shows qualitative comparisons, ablating the guidance used in the static stage. Without 2D diffusion guidance, the method often fails to produce the correct layout of the scene, and sometimes produces synthetic-looking texture. This is also reflected by the user study in Tab. 2a. Without 3D diffusion guidance, the learned assets suffer from severe Janus problems and do not have plausible shapes (see row 1 of Fig. 4 and ‘w/o 3D guidance’ in Tab. 2c). By combining both 3D and 2D guidance, our method reconstructs 3D-consistent static scenes with plausible compositions and realistic textures.

Deformation field and motion regularization. To learn better motion, we propose to use a multi-resolution hash-encoded 4D feature grid for the deformation MLP. We ablate this choice against a baseline MLP with positional encoding [21]. In Fig. 6, our method learns more local motion around the mouth and fin areas of the clown fish. We also ablate the total variation (TV) loss on the dis-

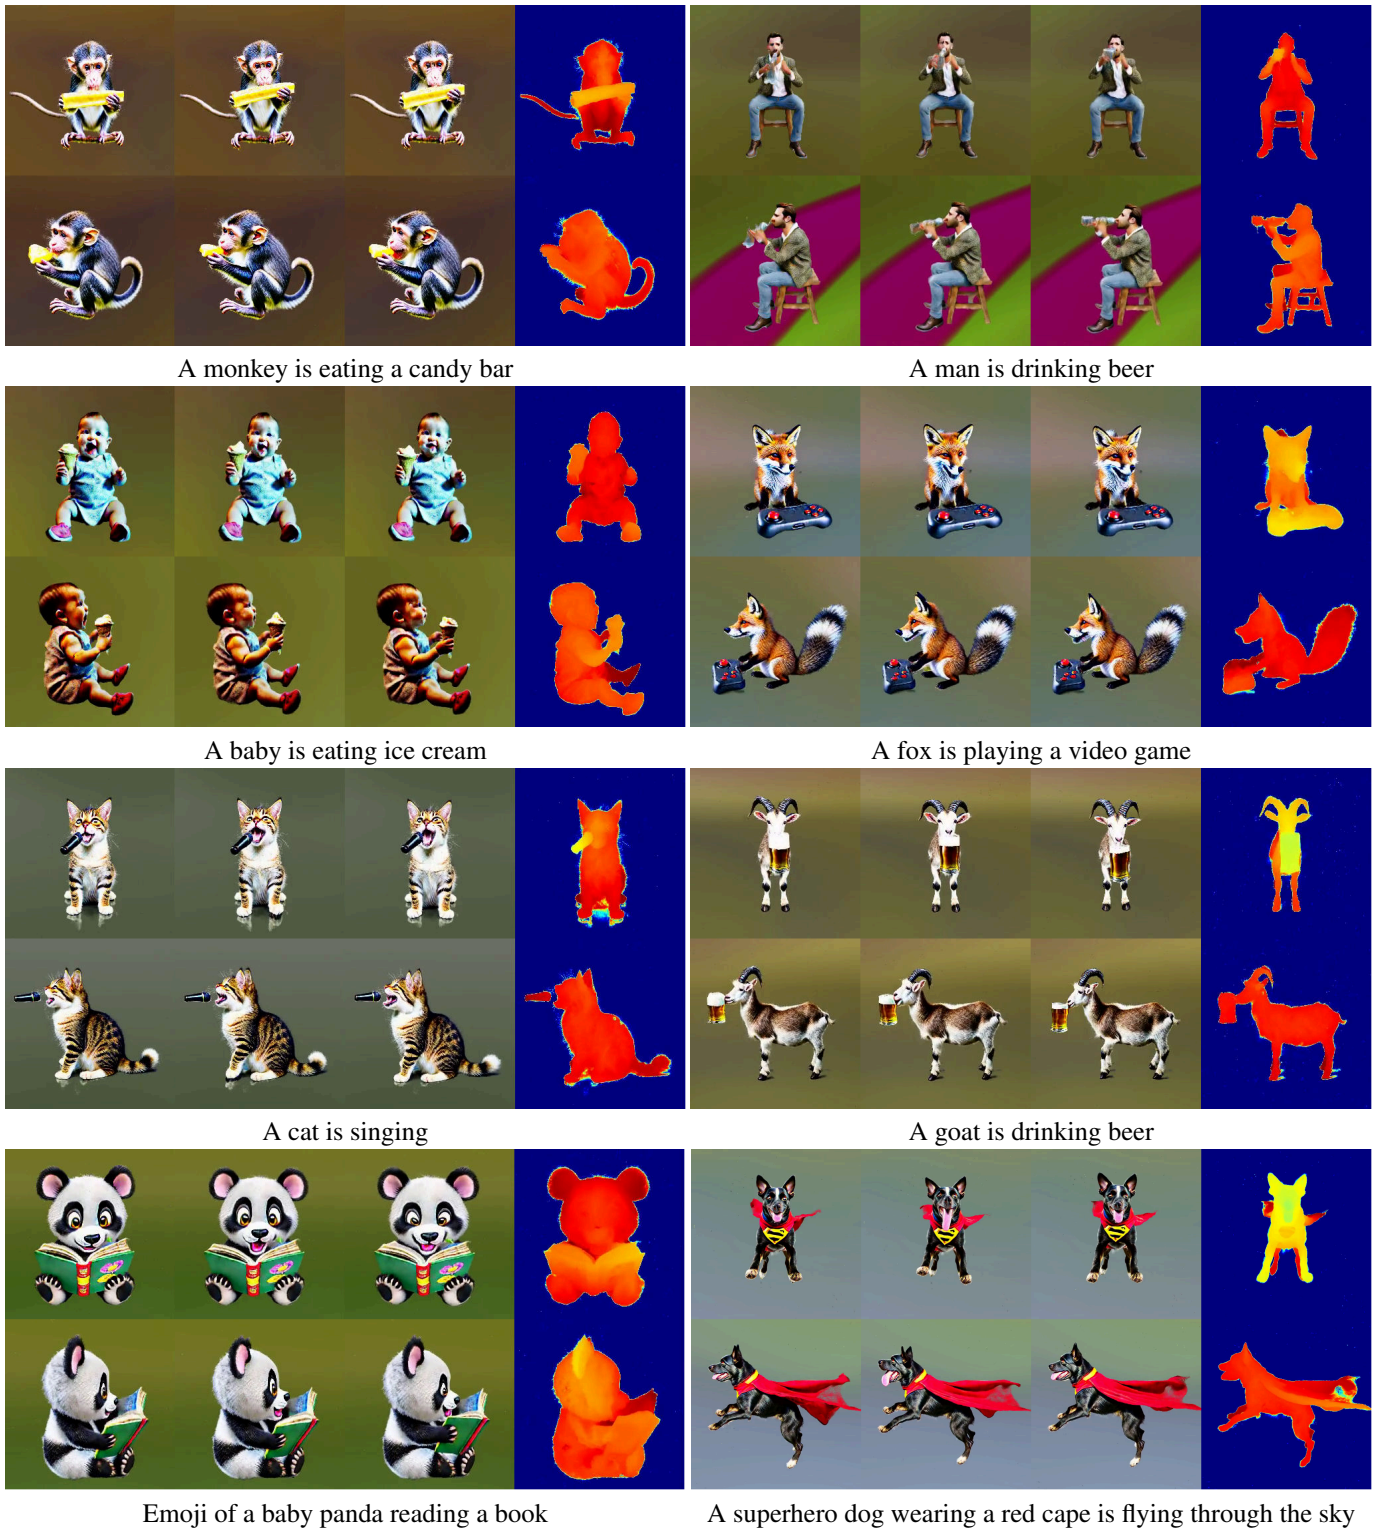


Figure 7. **Text-to-4D generation.** We show qualitative results of text-to-4D generation, demonstrating high visual quality, multi-view consistency, plausible composition and realistic motion. Video results are available in the Supplement.

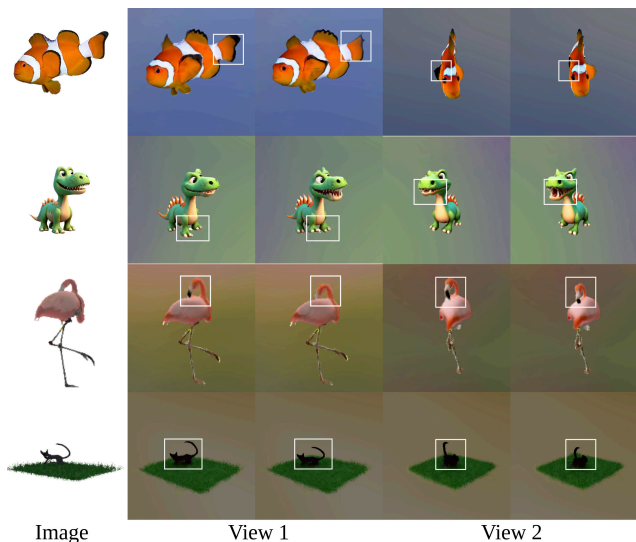


Figure 8. **Image-to-4D generation.** Given an input image, our method reconstructs and animates 3D assets. Prompts used for motion are ‘A clown fish swimming’, ‘A cartoon dragon running’, ‘A flamingo scratching its neck’, and ‘A cat walking on grass’. Video results can be found in our Supplement.

placement map and show that the learned motion presents substantial noise when not using the TV loss (Fig. 6). These observations about the learned motion quality are also reflected by the user study in Tab. 2b.

D-NeRF v.s. hexplane representation We also ablate our deformable NeRF representation against the hexplane 4D representation [5] used previously in [33]. Due to the entanglement of geometry and motion in hexplanes, it is not trivial to keep the static geometry parts frozen during the motion learning stage, which leads to lower visual quality and reappearance of the Janus problem (See Fig. 5 and Tab. 2c, ‘Hexplane’ versus Ours). In comparison, our dynamic representation fully disentangles the canonical model and the deformation field, successfully preserving the static model while learning its motion.

4.2. Controllable 4D Generation

We show qualitative results of text-to-4D generation where the object appearance is defined by one or multiple user-defined images. In Fig. 8, given a single image, our method can preserve the identity and appearance details of the input image and successfully learn the animation specified in the text prompt. It is preferred by users over MAV3D 98.3% of the time for its image texture preservation and 100.0% of the time for its alignment to the text. In Fig. 9, we show personalized 4D generation. Given a few casually captured images of a subject, our method can generate 4D content of the subject under various motion conditions, *e.g.*, eating food or ice cream. More qualitative results can be found on our project page.

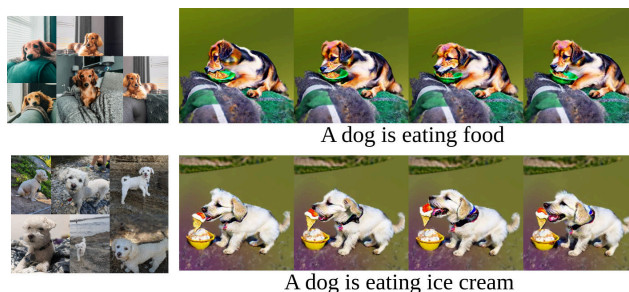


Figure 9. **Personalized 4D.** Our method can generate dynamic 3D scenes of a subject given a text prompt and 4-6 causally captured images of the subject. Videos are available in our Supplement.




Figure 10. **Failure case.** Despite combining 3D and 2D diffusion guidance, our method fails to reconstruct ‘A robot is playing the violin’. The static stage fails to learn a view-consistent violin, and the robot’s hand position is incorrect. In the second stage, our method cannot correct such errors or learn plausible arm motion.

5. Discussion

Conclusion We propose Dream-in-4D, a unified approach to 4D scene synthesis from a text prompt and optionally one or more input images. By leveraging 3D and 2D diffusion priors, our method first learns a high-quality static asset, offering a good starting point for deformation optimization. Then, our motion-disentangled 4D representation allows us to learn motion with video diffusion guidance while maintaining the quality of the static asset. We introduce multi-resolution feature grids and a TV loss for the deformation field, resulting in more realistic motion. Dream-in-4D achieves better visual quality, 3D consistency, motion and spatial layout on the text-to-4D task compared to the baselines, while also enabling image-to-4D and personalized 4D generation.

Limitation The combination of 3D and 2D diffusion priors sometimes fails to learn correct static 3D representations for some difficult prompts (*e.g.*, a robot playing a violin in Fig. 10). In the dynamic stage, our method cannot recover from the wrong static representation and fails to learn correct motion given the wrong position of the hands. We believe this problem could potentially be solved with further advances in 3D and 2D diffusion models.

Acknowledgement We sincerely thank all participants of our user study for their help. Yufeng Zheng is partially supported by the Max Planck ETH Center for Learning Systems.

References

- [1] Deepfloyd if. <https://github.com/deep-floyd/IF>. 1, 2, 3, 5
- [2] Modelscope. <https://huggingface.co/damo-vilab/text-to-video-ms-1.7b>. 1, 3, 5
- [3] Zeroscope. https://huggingface.co/cerspense/zeroscope_v2_576w. 2, 5
- [4] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *CVPR*, 2023. 1, 3
- [5] Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. *CVPR*, 2023. 2, 3, 8
- [6] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16123–16133, 2022. 3
- [7] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. In *ICCV*, 2023. 3
- [8] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. *arXiv preprint arXiv:2212.08051*, 2022. 3
- [9] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, Eli VanderBilt, Aniruddha Kembhavi, Carl Vondrick, Georgia Gkioxari, Kiana Ehsani, Ludwig Schmidt, and Ali Farhadi. Objaverse-xl: A universe of 10m+ 3d objects. *arXiv preprint arXiv:2307.05663*, 2023. 3
- [10] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13142–13153, 2023. 4
- [11] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinlong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022. 3
- [12] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12479–12488, 2023. 3
- [13] 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. 3
- [14] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, et al. Neural 3d video synthesis from multi-view video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5521–5531, 2022. 3
- [15] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 300–309, 2023. 1, 3
- [16] 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. 1, 2, 3, 4, 5, 6
- [17] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Learning to generate multiview-consistent images from a single-view image. *arXiv preprint arXiv:2309.03453*, 2023. 1, 3, 4
- [18] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021. 3
- [19] Luke Melas-Kyriazi, Christian Ruppert, Iro Laina, and Andrea Vedaldi. Realfusion: 360 reconstruction of any object from a single image. In *CVPR*, 2023. 3
- [20] Gal Metzer, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12663–12673, 2023. 1, 3
- [21] 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. *Commun. ACM*, 65(1):99–106, 2021. 2, 3, 6
- [22] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. 3
- [23] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021. 3
- [24] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 1, 3, 4
- [25] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 2, 3
- [26] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Sko-rokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123:

- One image to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint arXiv:2306.17843*, 2023. 1, 3, 4, 6
- [27] Amit Raj, Srinivas Kaza, Ben Poole, Michael Niemeyer, Ben Mildenhall, Nataniel Ruiz, Shiran Zada, Kfir Aberman, Michael Rubenstein, Jonathan Barron, Yuanzhen Li, and Varun Jampani. Dreambooth3d: Subject-driven text-to-3d generation. *ICCV*, 2023. 3
- [28] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. 1, 2, 3
- [29] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510, 2023. 6
- [30] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. 1
- [31] Yichun Shi, Peng Wang, Jianguo Ye, Long Mai, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. *arXiv:2308.16512*, 2023. 1, 2, 3, 4
- [32] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022. 1, 2
- [33] Uriel Singer, Shelly Sheynin, Adam Polyak, Oron Ashual, Iurii Makarov, Filippos Kokkinos, Naman Goyal, Andrea Vedaldi, Devi Parikh, Justin Johnson, and Yaniv Taigman. Text-to-4d dynamic scene generation. In *Proceedings of the 40th International Conference on Machine Learning*, 2023. 2, 3, 6, 8
- [34] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-it-3d: High-fidelity 3d creation from a single image with diffusion prior, 2023. 3
- [35] Christina Tsalicoglou, Fabian Manhardt, Alessio Tonioni, Michael Niemeyer, and Federico Tombari. Textmesh: Generation of realistic 3d meshes from text prompts. *arXiv preprint arXiv:2304.12439*, 2023. 3
- [36] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A. Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. *arXiv preprint arXiv:2212.00774*, 2022. 3
- [37] 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. 3
- [38] Joseph Zhu and Peiye Zhuang. Hifa: High-fidelity text-to-3d with advanced diffusion guidance. *arXiv preprint arXiv:2305.18766*, 2023. 3, 5