

AC3D: Analyzing and Improving 3D Camera Control in Video Diffusion Transformers

Sherwin Bahmani^{*1,2,3} Ivan Skorokhodov^{*3} Guocheng Qian³ Aliaksandr Siarohin³
Willi Menapace³ Andrea Tagliasacchi^{1,4} David B. Lindell^{1,2} Sergey Tulyakov³

¹University of Toronto ²Vector Institute ³Snap Inc. ⁴SFU

*equal contribution

<https://snap-research.github.io/ac3d>



Figure 1. **Camera-controlled video generation.** Our method enables precise camera controllability in pre-trained video diffusion transformers, allowing joint conditioning of text and camera sequences. We synthesize the same scene with two different camera trajectories as input. The inset images visualize the cameras for the videos in the corresponding columns. The left camera sequence consists of a rotation to the right, while the right camera visualizes a zoom-in, up, and zoom-out trajectory.

Abstract

Numerous works have recently integrated 3D camera control into foundational text-to-video models, but the resulting camera control is often imprecise, and video generation quality suffers. In this work, we analyze camera motion from a first principles perspective, uncovering insights that enable precise 3D camera manipulation without compromising synthesis quality. First, we determine that motion induced by camera movements in videos is low-frequency in nature. This motivates us to adjust train and test pose conditioning schedules, accelerating training convergence while improving visual and motion quality. Then, by probing the representations of an unconditional video diffusion transformer, we observe that they implicitly perform camera pose estimation under the hood, and only a sub-portion of their layers contain the camera information. This suggested us to limit the injection of camera conditioning to a subset of the

architecture to prevent interference with other video features, leading to $4\times$ reduction of training parameters, improved training speed and 10% higher visual quality. Finally, we complement the typical dataset for camera control learning with a curated dataset of 20K diverse dynamic videos with stationary cameras. This helps the model disambiguate the difference between camera and scene motion, and improves the dynamics of generated pose-conditioned videos. We compound these findings to design the Advanced 3D Camera Control (AC3D) architecture, the new state-of-the-art model for generative video modeling with camera control.

1. Introduction

Foundational video diffusion models (VDMs) trained on internet-scale data, acquire abundant knowledge about the physical world [10]. They not only learn appearance and

plausible 2D dynamics, but they also have abundant understanding of 3D structure [7]. However, most of this knowledge is stored *implicitly* within the model, as these models do not expose fine-grained control mechanisms, such as camera motion control. We recently witnessed a surge of works that bring 3D camera control into foundational video models [38, 150, 166], but the control they provide is not very precise, and the synthesis quality is often compromised [6]. We analyze camera motion control in video diffusion models from first principles, and develop several findings that allow us to incorporate *precise* 3D camera conditioning *without* degrading synthesis quality. To perform our analysis we train a 11.5B-parameters VDiT (video latent diffusion transformer) [100] on a dataset of 100M text/video pairs. On this model, we perform *three* key studies. With what we learn, we adapt the camera control solution from VD3D [6] from a pixel-based to latent-based diffusion model, and significantly increase its performance.

1) The spectral properties of camera motion. To study the statistical nature of motion control, we analyze *motion spectral volumes* (MSV) [75] of the videos generated by a large-scale video DiT model. MSVs show the amount of energy in different portions of the frequency spectra (i.e., high energy in the low frequencies indicate smooth motion) and we measure them across 200 generated videos of different types (camera motion, scene motion, scene plus camera motion) and at various stages of the denoising synthesis process. We observe that camera motion mostly affects the lower portion of the spectrum and kicks in very early ($\approx 10\%$) in the denoising trajectory. Then, as diffusion models are inherently coarse-to-fine in nature [25], we restrict our camera conditioning to only being injected on the subset of the denoising steps corresponding to low-frequencies. This results in $\approx 15\%$ higher visual fidelity, $\approx 30\%$ better camera following, and mitigates scene motion degradation.

2) Camera motion knowledge in VDiTs. Then, we consider our text-only VDiT, and determine *whether* such a model possesses knowledge about cameras, and *where* this knowledge is expressed within its architecture. With this objective, we feed the (*unseen* during training) RealEstate10k [198] videos to our VDiT, and perform *linear probing* [27] to determine if camera poses can be recovered from its internal representation. Our analysis revealed that a video DiT implicitly performs camera pose estimation under the hood, and the presence of camera knowledge in a *disentangled* form peaks in its middle layers. This implies that the camera signal emerges in its early blocks to allow the later ones rely on it to build subsequent visual representations. Therefore, we adjust our conditioning scheme to only affect the first 30% of the architecture, leading to a $\approx 4\times$ reduction in training parameters, 15% training and inference acceleration, and 10% improved visual quality.

3) Re-balancing the training distribution. Finally, to supervise camera control architectures, the typical solution is to rely on the camera pose annotations provided by RealEstate10k [198]. However, this dataset contains mostly *static* scenes, which results in significant motion degradation of the fine-tuned video model. To overcome this problem, we curate a subset of 20K diverse videos with dynamic scenes *but* static cameras. As the camera conditioning branch is still activated for these videos, this helps the model disambiguate the camera from scene movement. Our experiments show that this simple adjustment in the data is sufficient to recover the scene dynamism while still enabling an effective pose-conditioned video model.

Contributions. We compound the knowledge gained from these three studies into the design of the Advanced 3D Camera Control (AC3D) method. We perform extensive ablation studies and compare against state-of-the-art models for camera control, including MotionCtrl [150], CameraCtrl [38], and VD3D [6]. We demonstrate 18% higher video fidelity and 25% more precise camera steering in terms of quantitative metrics than a closest competitor, and our generated videos are favored to others in 90% of cases.

2. Related work

Our approach lies at the intersection of text-to-video, text-to-3D, and text-to-4D generation approaches. We refer to recent state-of-the-reports [101, 179] for a more thorough analysis of previous work.

Text-to-video generation. Our approach builds on recent advancements in 2D video diffusion models. One prominent technique in this area enhances text-to-image models by adding temporal layers to support video generation [7, 8, 36, 36, 120, 152]. While these methods use the U-Net architecture, more recent ones [10, 92, 93, 168] have been adapting transformer-based architectures for more scalable, realistic, and highly dynamic scene generation. We are interested in controlling the camera movements during the generation process of recent transformer-based video models based on precise camera extrinsics, i.e., cameras represented as rotation and translation sequences for each frame.

4D generation. Early 4D generation works [2, 164] used 4D GANs to learn category-specific generators with an underlying dynamic 3D representation. More recent approaches [4, 5, 5, 81, 121, 195] have tackled 4D generation by distilling motion priors from pre-trained video diffusion models into an explicit 4D representation, enabling category-agnostic 4D generation. Follow-up works investigate image or video conditioned 4D generation [33, 76, 81, 98, 109, 172, 181, 192, 195] instead of pure text inputs, improving flexibility in the generation process. While most of these works are object-centric, recent approaches [4, 159] shifted towards more complex scenes, including methods [23, 174]

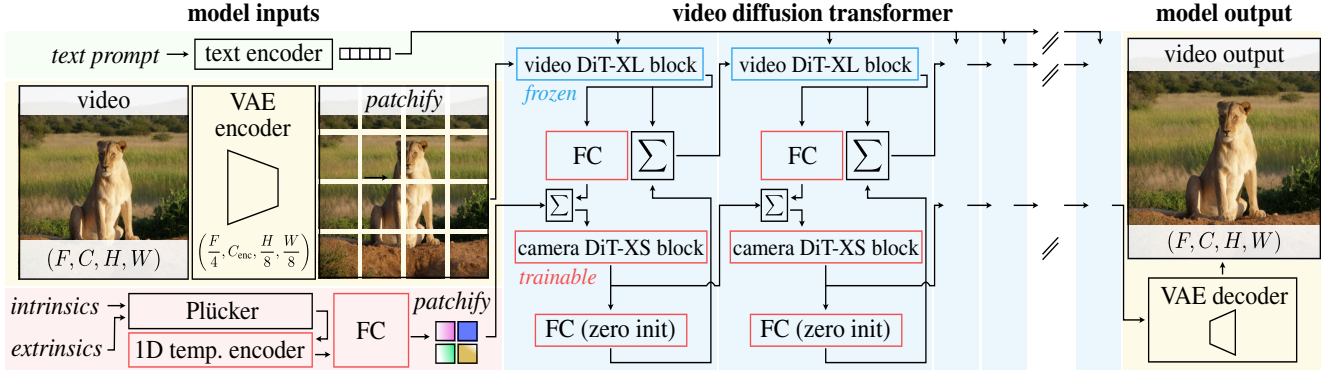


Figure 2. **VDiT-CC** model with ControlNet [71, 187] camera conditioning built on top of **VDiT**. Video synthesis is performed by large 4,096-dimensional **DiT-XL** blocks of the frozen **VDiT** backbone, while **VDiT-CC** only processes and injects the camera information through lightweight 128-dimensional **DiT-XS** blocks (FC stands for fully-connected layers); see Section 3.2 for details.

which model the background. However, all these methods are optimization-based, i.e., each scene is generated independently from scratch. Recently, L4GM [110] proposes a feed-forward 4D generator trained on object-centric synthetic 4D data. While these approaches are explicit and provide space-time control, they are limited in their photo-realism compared to recent 2D video diffusion models. We investigate dynamic 3D scene generation from a different perspective by extending pre-trained video diffusion models with 3D camera control.

Camera control for video models. Recently, there has been significant progress in adding camera control to video diffusion models. As the pioneering work, MotionCtrl [150] learns camera control by conditioning pre-trained video models [7, 17] with extrinsics matrices. Follow-up works [38, 65, 160] further improve the conditioning mechanisms by representing cameras as Plücker coordinates. Another line of work [49, 50, 82, 155] controls camera motion without training additional parameters. However, all of these approaches use U-Net-based architectures as their backbone. More recently, 4DiM [151] trains a space-time diffusion model from scratch for novel view synthesis from a single image input. Closely related to our work, VD3D [6] incorporates camera control into a pre-trained video diffusion transformer. While the motion and camera control improves over U-Net-based approaches, the synthesized motion in the scenes and the visual quality are still degraded compared to the base video model. In contrast to VD3D, we first thoroughly investigate the pre-trained base video model and its knowledge of camera motion. We derive an improved training and architecture design for high-quality and dynamic video generation based on our findings.

Concurrent works. Concurrent approaches [68, 128, 158, 176, 184, 193, 194] further improve camera control in U-Net-based architectures, while another work [22] tackles video diffusion transformer. However, the scene and visual

quality is still limited in that approach. [51] explores pose estimation with a video DiT by pairing it with DUST3R [146] and fine-tuning, while we perform linear probing without any training to assess its existing camera knowledge.

3. Method

We first describe our base video diffusion model (Sec. 3.1), and the baseline camera control method built on top of it (Sec. 3.2). Then, we proceed with the analysis of motion (Sec. 3.3), linear probing (Sec. 3.4) and dataset biases (Sec. 3.5), and additional insights on how to build an effective camera control (Sec. 3.6)

3.1. Base model (VDiT)

Following Sora [10], most modern foundational video generators use the diffusion framework [45, 123] to train a large-scale transformer [140] in the latent space of a variational autoencoder [64, 111, 112] to generate videos from text descriptions. We adopt the same design and, for a base video model, pre-train an 11.5B-parameters Video DiT model [100] with 32-blocks of 4,096 hidden dimensionality for text-to-video generation in the latent space of CogVideoX [168] 16-channels $4 \times 8 \times 8$ -compression autoencoder with the rectified flow diffusion parametrization [85]. We use the T5 [108] encoder to produce text embeddings, which are passed into VDiT via cross-attention. We train our base model on a large-scale dataset of images and videos with text annotations, with resolutions ranging from $17 \times 144 \times 256$ to $121 \times 576 \times 1024$. Such design is fairly standard and followed by many existing works with little deviation [32, 102, 168, 196], and we describe our particular architectural and training setups in detail in Appendix C.

3.2. VDiT with Camera Control (VDiT-CC)

To construct a baseline architecture for camera control, we implement ControlNet [18, 187] conditioning on top VDiT.




Figure 3. Comparing motion spectral volumes for scenes with different motion types. Videos with camera motion (purple) exhibit stronger overall motion than the videos with scene motion (orange), especially for the low-frequency range, suggesting that the motion induced by camera transitions is heavily biased towards low-frequency components.

Similarly to [6, 38, 150], we use the RealEstate10k [198] dataset, consisting of 65k (text, video, camera trajectory) triplets $(x_n, y_n, c_n)_{n=1}^N$ and train a new set of model parameters to input the camera information into the model. Camera trajectories $c \in \mathbb{R}^{f \times 25}$ are provided in the form of camera extrinsics $C_f \in \mathbb{R}^{4 \times 4}$ and intrinsics $K_f \in \mathbb{R}^{3 \times 3}$ for each f -th frame x_f .

Camera conditioning. For base camera control, we adapt VD3D [6] since it was designed for transformer-based models and suits our setup the most, while other methods are built on top of UNet-based [113] backbones. We use Plücker camera representations [6, 16, 38, 59, 122], which are projected to the same dimensionality and resolution as the video tokens via a fully-convolutional encoder to produce camera tokens. These camera tokens are processed by a sequence of lightweight DiT-XS blocks with 128 hidden dimensionality and 4 attention heads each. To mix the camera information with the video tokens of VDiT, we found a summation to be sufficient, which we perform before each main DiT block. We also found it useful to perform cross-attention from video tokens to camera tokens as a form of a feedback connection [71]. We illustrate this architecture VDiT-CC in Figure 2; see implementation details in Appendix C.

Training. Keeping the VDiT backbone frozen, we train the new parameters with a rectified flow objective [85] and standard (location of 0 and scale of 1) logit-normal noise distribution [28]. Similar to prior works [6, 151], we apply a 10% camera dropout to support classifier-free guidance (CFG) [43] later. Notably, we train VDiT-CC only at the 256^2 resolution: since camera motion is a low-frequency type of signal (which is fully unveiled at lower resolutions) and the main VDiT backbone is frozen, we found that such design generalizes to higher resolutions out-of-the-box. During inference, we input text prompts and camera embeddings with classifier-free guidance at each time step.

Model behavior. This baseline model, being built on top of our powerful VDiT, already achieves decent-quality camera control. However, it struggles with degraded visual quality, reduced scene motion, and sometime the camera control gets ignored. To improve the design, we analyze our VDiT backbone to understand how camera motion is modeled and represented by it, and then inspect VDiT-CC’s failure cases and where they arise to rigorously address them.


3.3. How is camera motion modeled by diffusion?

We start by analyzing how camera motion is modeled by a pre-trained video diffusion model (i.e., before camera control is incorporated). We hypothesize that the motion induced by changes in camera pose is a low-frequency type of signal and investigate the *motion spectral volumes* [75] of the generated videos at different steps of the denoising process. To perform this analysis, we generate 200 diverse videos with our VDiT model with 80-steps denoising and *manually* annotate them into four categories: videos with only scene motion, videos with only camera motion, videos with both scene and camera motion, and others; see Appendix D for details. During generation, we save the denoised predictions at *each* denoising step and estimate optical flow to compute the motion spectral volumes.


Analysis. We visualize these lines with 95% confidence intervals in Figure 3. Videos with camera motion exhibit higher amplitudes than scene-motion-only videos for low-frequency components while having similar characteristics for high-frequency ones. This supports the conjecture that the camera motion is a low-frequency type of signal. We also depict an example of a generated video with both scene and camera motion with four denoising steps on Fig. 4a: one can observe that the camera movement has been fully produced by $t=0.9$ (first 10% of the rectified flow denoising process). In contrast, scene motion details like the hand movements of the subjects are not finalized even till $t=0.5$.

Inspired by this finding, we pose the question: *when does exactly a video diffusion model determine the camera?* To answer it, we plot aggregated spectral volumes for different timesteps in Figure 4b, and also their ratio with respect to the last timestep $t=0$ (for which all the motion has been produced) We then inspect when different types of motion kick in during the denoising process. Figure 4b (right) shows that the low-frequency motion components fill up to $\approx 84\%$ at $t=0.9$ (the first 10% of the denoising process), while high-frequency ones are not well complete until $t=0.6$.

An immediate consequence of this observation is that trying to control the camera later in the denoising trajectory is simply *unnecessary* and will not influence the manipulation result. In this way, instead of using the standard logit-normal noise level distribution of SD3 [28] with a location of 0.0 and scale of 1.0 (which we use by default for VDiT), we switch to using truncated normal with a location of 0.8 and



(a) A generated video at different diffusion timesteps. The camera has already been decided by the model even at $t = 0.9$ (first 10% of the denoising process) and does not change after that.



(b) Motion spectral volumes of VDiT’s generated videos for different diffusion timesteps (left) and their ratio w.r.t. the motion spectral volume at $t = 0$ (i.e., a fully denoised video).

Figure 4. How camera motion is modeled by diffusion? As visualized in Figure 4a and Figure 3, the motion induced by camera transitions is a low-frequency type of motion. We observe that a video DiT creates low-frequency motion very early in the denoising trajectory: Figure 4b (left) shows that even at $t=0.96$ (first $\approx 4\%$ of the steps), the low-frequency motion components have already been created, while high frequency ones do not fully unveil even till $t=0.5$. We found that controlling the camera pose later in the denoising trajectory is not only unnecessary but *detrimental* to both scene motion and overall visual quality.

scale of 0.075 on the $[0.6, 1]$ interval to cover the early steps of the denoising rectified flow trajectory. Besides, we apply camera conditioning at inference time on the same $[0.6, 1]$ interval. But surprisingly, we observe that it is unnecessary but also *detrimental* to the scene motion and overall visual quality. Following this insight, we restrict both our train-time noise levels and test-time camera conditioning schedules to cover only the first 40% of the reverse diffusion trajectory. As Sec. 4.3 shows, it improves FID and FVD by 14% on average, and camera following by 30% on MSR-VTT (the dataset which measures generalization to diverse, out-of-fine-tuning-distribution scenes). Besides, it enhances the overall scene motion, which we demonstrate qualitatively.

3.4. What does VDiT know about camera pose?

Foundational video models acquire rich knowledge about the physical world, and we hypothesize that they store information about the camera signal within their representations. To investigate this, we perform linear probing of our base VDiT model on the RealEstate10k [198] dataset (not seen during training) for camera extrinsics.

Specifically, we take 1,000 random 49-frame videos from RealEstate10K, feed them into VDiT under 8 noise levels $(1/8, 2/8, \dots, 1)$, and extract the activations for all 32 DiT blocks. Next, we split them into 900 train and 100 test videos and train a linear ridge regression model to predict the rotation pitch/yaw/roll angles and translation vectors for the entire viewpoint trajectory (49×6 target values in total). This results in 8×32 trained models, and we report the rotation and (normalized) translation errors [38] on a held-out test set of 100 videos in Figure 5. Surprisingly, VDiT can accurately predict the camera signal, achieving minimum test errors of ≈ 0.025 for rotation and for ≈ 0.48 translation

prediction. The knowledge quality increases around layer #9 and peaks in the range of #13-21. We reason that since the camera information in block #13 is stored in such a disentangled manner, then the model is using it to build other representations, and conditioning the camera at this point is risky and unnecessary and would interfere with other visual features. In this way, we propose to input the camera *only* in the first #8 blocks and leave the rest 24 DiT blocks unconditioned. We empirically found in Section 4.3, that this not only reduces the amount of trainable parameters ≈ 4 times and improves training speed $\approx 15\%$, but also enhances the visual quality $\approx 10\%$.

3.5. Mitigating training data limitations

Estimating camera parameters from in-the-wild videos remains challenging, as leading methods like [115, 116, 146, 191] frequently fail when processing videos containing dynamic scene content. This limitation has resulted in camera-annotated datasets being heavily biased toward static scenes, particularly evident in RealEstate10K (RE10K) [198], which is the predominant dataset for training camera-controlled video models [6, 38, 150]. We hypothesize that models fine-tuned on such data interpret camera position information as a signal to suppress scene dynamics. This bias persists even when jointly training on unconstrained 2D video data [151], because the camera conditioning branch is only activated when camera parameters are available, which occurs exclusively for static scenes from RE10K, as static scenes remain the only reliable source for accurate camera annotation.

To address this fundamental limitation, we propose an alternative approach: rather than attempting to annotate dynamic scenes, which proved unsuccessful in our extensive preliminary research, even with state-of-the-art meth-




Figure 5. **Video DiT is secretly a camera pose estimator.** We perform linear probing of camera poses in each of **VDiT** blocks for various noise levels and observe that video DiT performs pose estimation under the hood. Its middle blocks carry the most accurate information about the camera locations and orientations, which indicates that the camera signal emerges in the early layers to help the middle and late blocks render other visual features aligned with the viewpoint.

ods [146], we curate a collection of 20K diverse videos featuring dynamic scenes captured by stationary cameras (see Figure 6). With stationary cameras, the camera position is inherently known, (we can assign fixed arbitrary extrinsics), allowing us to maintain active camera conditioning during training. This approach enables the camera conditioning branch to remain active during training while exposing the model to dynamic content, helping it distinguish between viewpoint conditioning and scene stillness. On top of this secondary dataset, following [151], we remove the scale ambiguity in RE10K by leveraging an off-the-shelf metric depth estimator; see Appendix G. Our experiments in Sec. 4.3 demonstrate that this straightforward yet effective data curation strategy successfully mitigates the distributional limitations of RE10K, restoring much of the lost scene dynamics, while maintaining precise camera control.

3.6. Miscellaneous improvements

In addition to our core analysis, we introduce several auxiliary techniques that enhance model performance.

1D temporal camera encoder. **VDiT-CC** processes videos in the autoencoder’s latent space with $4\times$ temporal compression [168], raising the question of how to incorporate full-resolution camera parameters. While camera poses could be naively downsized (e.g., subsampled or reshaped) to match the latent resolution, this would force small **DiT-XS** blocks to process compressed camera information. Instead, we implement a sequence of causal 1D convolutions that transform a $F\times 6$ sequence of Plücker coordinates for each pixel into a $(F/4)\times 32$ representation; see Appendix C.

Separate text and camera guidance. Text and camera signals require different guidance weights due to their distinct nature, motivating us to separate their classifier-free guidance (CFG) [9, 44]. We formulate the equation as:

$$\begin{aligned} \hat{s}(\mathbf{x}|\mathbf{y}, \mathbf{c}) &= (1 + w_y + w_c)s_\theta(\mathbf{x}|\mathbf{y}, \mathbf{c}) \\ &\quad - w_y s_\theta(\mathbf{x}|\mathbf{c}) - w_c s_\theta(\mathbf{x}|\mathbf{y}), \end{aligned} \quad (1)$$



Figure 6. RealEstate10k [198] videos (upper 2 rows) contain diverse camera trajectories, but are strongly biased towards static scenes. To mitigate this bias and also increase the concepts diversity, we curate 20K videos with stationary cameras, but dynamic content (lower 2 rows). Such dataset is easy to construct, and surprisingly effective. Section 4.3 shows that integrating it into our training allows to improve visual quality on out-of-distribution prompts by 17%.

where $\hat{s}(\cdot)$ denotes the final update direction used during synthesis, s_θ represents the model’s predicted update direction, \mathbf{y} and \mathbf{c} are text and camera conditions, and w_y and w_c are their respective CFG weights.

ControlNet with feedback. Traditional ControlNet [187], used in recent camera control methods [6, 38, 150], only processes conditioning signals without accessing the main branch. Our experiments show that adding cross-attention from video tokens to camera tokens in transformer computations enables better camera representations. This modification behaves as a feedback mechanism [71] provided by the main synthesis branch to the camera processing branch.

Dropping context in the camera branch. We found that applying cross-attention over the context information (text prompts, resolution, etc.) in the camera **DiT-XS** blocks

Method	Human Preference				Overall
	CA	MQ	TA	VQ	
Ours vs. VD3D (FIT)	89.5%	79.0%	87.5%	97.5%	95.0%
Ours vs. VD3D (DiT)	65.0%	87.5%	83.5%	95.0%	92.5%

Table 1. **User study.** We compare our approach to the original VD3D (FIT) and reimplemented VD3D (DiT) on top of our base model. We conduct a user study where participants indicate their preference based on *camera alignment (CA)*, *motion quality (MQ)*, *text alignment (TA)*, *visual quality (VQ)*, and *overall preference (Overall)*.

lead to detrimental visual quality and camera steering due to harmful interference of the context embeddings with camera representations.

4. Experiments

Datasets. Our base **VDiT** model was trained on a large-scale dataset of text-annotated images and videos. **VDiT-CC** is fine-tuned from **VDiT** on RealEstate10K [198], contains $\approx 65\text{M}$ video clips with per-frame camera parameters (extrinsics and intrinsics) since it is the setup used by existing methods [6, 38, 150].

Metrics. To assess the performance, we rely on a wide range of automatic quantitative metrics and user studies. For the metrics, we use FID [42], FVD [137], and CLIP score [41] to evaluate visual quality, and rotation and normalized translation errors [38] of ParticleSfM [191]-reconstructed trajectories to assess camera steerability. We evaluate them both on RE10K and MSR-VTT [161], since the latter allows to assess zero-shot generalization on out-of-distribution data.

In the user study, we engage 10 professional labelers, each of whom evaluates 100 different video pairs. The labelers are asked to choose between two videos based on several preference metrics:

- *Camera Alignment (CA)*: How well the camera trajectory follows the reference video.
- *Motion Quality (MQ)*: Which video has larger and more natural motion.
- *Text Alignment (TA)*: Which video better aligns with the provided reference text prompt.
- *Visual Quality (VQ)*: Which video has a higher overall visual quality.
- *Overall Preference (Overall)*: Which generated video the user would prefer for this task.

4.1. Baselines

We select 3 contemporary camera-control methods: MotionCtrl [150], CameraCtrl [38], and VD3D [6]. MotionCtrl and CameraCtrl use a UNet-based video diffusion backbone [36], while VD3D builds on top of FIT [21, 93] and

as such, is easily extendable to our video DiT [100] setup. That’s why we re-implement it on top of our **VDiT** model to obtain an additional “VD3D+DiT” baseline. Besides, we also employ the unconditional base **VDiT** model as a reference for visual quality degradation.

4.2. Main results

We present the quantitative comparison results against the baselines in Tab. 2. One can observe that just switching from 4B-parameters pixel-space FIT [21] backbone, employed by the original VD3D approach, to our larger 11.5B-parameters latent-space DiT yielded clear gains across the most of the metrics. Next, the results demonstrate that our AC3D establishes a new state-of-the-art performance against all the baselines, and achieves $\approx 18\%$ better visual quality than the closest competitor (VD3D DiT) and $\approx 25\%$ more precise camera steering. As evaluating the quality of camera motion by staring at a small selection of frames in a PDF is incredibly difficult, we instead opt for visualizing all qualitative results in the website provided within our supplementary material. Therein, we can observe that AC3D better follows pose conditioning, and achieves higher visual fidelity. We conduct user studies against VD3D+FIT (the original model) and VD3D+DiT (our improved re-implementation on top of the bigger video transformer). The results are presented in Table 1: our AC3D outperforms them across all the qualitative aspects, achieving 90%+ overall preference score. Finally, we encourage the reader to assess the visual quality by observing the video comparisons in our supplementary materials.

4.3. Ablations

No camera conditioning. The first ablation we conduct is dropping any camera conditioning at all, which makes the model equivalent to vanilla **VDiT**. It is needed to understand the visual quality and text alignment degradation. The results (Tab. 2, row *w/o camera cond*) show that our model loses less only $\approx 7\%$ of the original visual fidelity on MSR-VTT (as measured by FVD), while (as expected) greatly improves on its in-domain RE10K data. In comparison, VD3D-DiT (the closest baseline) loses $\approx 20\%$ of its visual fidelity on MSR-VTT.

Importance of biasing the noise towards higher levels. As Sec. 3.3 tells, we use the truncated normal distribution with location of 0.8 and scale of 0.075 with the [0.6, 1] bounds for training AC3D. We ablated the importance of biasing the noise sampling towards high noise and observe higher motion, visual quality, and camera controllability.

Importance of truncating the noise schedule. We change the training and inference procedure by using no truncation during noise sampling. Instead, we condition the model with camera inputs over the whole noise range and observe decreased visual quality.

Method	RealEstate10K [198]					MSR-VTT [161]				
	TransErr ↓	RotErr ↓	FID ↓	FVD ↓	CLIP ↑	TransErr ↓	RotErr ↓	FID ↓	FVD ↓	CLIP ↑
MotionCtrl [150]	0.477	0.094	2.99	61.70	26.46	0.593	0.137	16.85	283.12	24.11
CameraCtrl [38]	0.465	0.089	2.48	55.64	26.81	0.587	0.132	12.33	201.33	25.05
VD3D (FIT) [6]	0.409	0.043	1.40	42.43	28.07	0.504	0.050	7.80	165.18	26.89
VD3D (DiT)	0.421	0.056	1.21	38.57	28.34	0.486	0.047	6.88	137.62	27.90
AC3D (ours)	0.358	0.035	1.18	36.55	28.76	0.428	0.038	5.34	110.71	28.58
w/o camera cond	+0.233	+0.153	+4.02	+53.83	-1.63	+0.266	+0.157	-0.48	-8.53	+0.35
w/o biasing noise	+0.093	+0.015	+0.02	+1.78	-0.32	+0.138	+0.033	+0.59	+16.92	-0.54
w/o noise truncation	+0.020	-0.003	+0.06	+1.69	-0.20	+0.016	+0.005	+0.76	+6.63	-0.18
w/o camera CFG	+0.014	+0.004	+0.49	+4.57	-0.54	+0.025	+0.003	+0.03	+1.42	-0.27
w/o our dynamic data	-0.005	-0.004	-0.06	+0.22	-0.20	+0.004	-0.001	+0.89	+4.40	-0.55
w/o metric scaled data	+0.013	+0.005	+0.17	+4.65	0.00	+0.023	+0.002	-0.01	0.00	-0.12
w/o dropping camera context	+0.013	+0.001	+0.04	+2.46	-0.65	+0.029	+0.003	+1.25	+7.41	-0.36
w/o limiting camera cond to 8 blocks	-0.001	+0.001	+0.09	+0.56	-0.02	+0.003	0.000	+0.32	+9.23	-0.33
w/ 2D training	+0.129	+0.068	+2.60	+33.85	-1.17	+0.128	+0.093	-0.26	-3.83	+0.21

Table 2. **Quantitative evaluation.** We evaluate all the models using camera pose and visual quality metrics based on unseen camera trajectories. We compute translation and rotation errors based on the estimated camera poses from generations using ParticleSfM [191]. We evaluate both in-distribution with RealEstate10K [198] and out-of-distribution with MSR-VTT [161].

No camera guidance. We assess the importance of classifier-free guidance [43] on the camera conditioning in Tab. 2 (*w/o camera CFG*). It attains the same visual quality on both in-distribution (RE10K) and out-of-distribution (MSR-VTT) data, but degrades camera following, resulting in $\approx 5\%$ higher pose reconstruction errors.

Training without our data with scene motion. To understand how well our curated data with scene motion but stationary cameras mitigate static scene bias, we train AC3D exclusively on RE10K, and report the results in Tab. 2 (*w/o our dynamic data*). It let the model maintain similar visual quality and text alignment on RE10K (in-domain data), but degraded its performance on out-of-distribution concepts of MSR-VTT ($\approx 17\%$ worse FID and $\approx 4\%$ worse FVD). Evaluating the quality of scene motion is only reasonably through qualitative assessment, which is why we refer the reader to our qualitative comparisons in the supplementary which highlight the positive influence of our curated data on scene dynamics.

Importance of metric scaled cameras. We train AC3D using the original RE10K’s camera parameters without our scaling procedure and present the results in Tab. 2 (*w/o metric scaled data*). Since it is a more ambiguous conditioning signal, it made the model lose both visual quality ($\approx 10\%$ FVD on RE10K) and also fail to accurately follow the camera, leading to $\approx 12\%$ higher trajectory reconstruction errors.

Providing context into the camera branch. As discussed in Sec. 3.6, we chose not to input the context information (text embeddings, resolution conditioning, etc.) into the lightweight camera branch to avoid potential interference with the camera representations. As Tab. 2 (*w/o dropping camera context*) shows, providing this information indeed results in $\approx 4\%$ worse camera following and $\approx 15\%$ lower visual quality.

Importance of limiting conditioning to the first 8 VDiT blocks. Following our insights in Sec. 3.4, we condition AC3D only in the first 8 blocks. Trying to condition in all the 32 DiT blocks (*w/o limiting camera cond to 8 blocks*) worsens the visual quality by $\approx 10\%$, while keeping the quality control at the same level. This suggests that the middle and late VDiT layers indeed rely on processed camera information and conditioning them on external camera poses might lead to interference with other visual features.

Joint training with 2D data. To mitigate visual quality and scene motion degradation, we attempted to perform joint training on 2D video data (without camera annotations) which was used in base VDiT training by applying dropout on camera inputs for it. While some prior works showed beneficial performance of such strategy [151] and, as Tab. 2 (*with 2D training*) shows, it indeed helped to maintain slightly higher visual fidelity ($\approx 3\%$ better FVD on MSR-VTT), camera steering severely deteriorated: up to $3\times$ worse results for translation/rotation errors.

5. Conclusions

Our findings demonstrate that principled analysis of camera motion in video diffusion models leads to significant improvements in control precision and efficiency. Through enhanced conditioning schedules, targeted layer-specific camera control, and better-calibrated training data, AC3D achieves state-of-the-art performance in 3D camera-controlled video synthesis while maintaining high visual quality and natural scene dynamics. This work establishes a foundation for more precise and efficient camera control in text-to-video generation. We discuss the limitations of our approach in Appendix B. In future work, we plan to focus on further improving data limitations and develop control mechanisms for camera trajectories far outside of the training distribution.

References

- [1] Jimmy Lei Ba. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 18
- [2] Sherwin Bahmani, Jeong Joon Park, Despoina Paschalidou, Hao Tang, Gordon Wetzstein, Leonidas Guibas, Luc Van Gool, and Radu Timofte. 3d-aware video generation. In *TMLR*, 2023. 2, 20
- [3] Sherwin Bahmani, Jeong Joon Park, Despoina Paschalidou, Xingguang Yan, Gordon Wetzstein, Leonidas Guibas, and Andrea Tagliasacchi. CC3D: Layout-conditioned generation of compositional 3D scenes. *Proc. ICCV*, 2023. 20
- [4] Sherwin Bahmani, Xian Liu, Wang Yifan, Ivan Skorokhodov, Victor Rong, Ziwei Liu, Xihui Liu, Jeong Joon Park, Sergey Tulyakov, Gordon Wetzstein, Andrea Tagliasacchi, and David B. Lindell. Tc4d: Trajectory-conditioned text-to-4d generation. In *Proc. ECCV*, 2024. 2, 20
- [5] Sherwin Bahmani, Ivan Skorokhodov, Victor Rong, Gordon Wetzstein, Leonidas Guibas, Peter Wonka, Sergey Tulyakov, Jeong Joon Park, Andrea Tagliasacchi, and David B. Lindell. 4d-fy: Text-to-4d generation using hybrid score distillation sampling. In *Proc. CVPR*, 2024. 2, 20
- [6] Sherwin Bahmani, Ivan Skorokhodov, Aliaksandr Siarohin, Willi Menapace, Guocheng Qian, Michael Vasilkovsky, Hsin-Ying Lee, Chaoyang Wang, Jiayu Zou, Andrea Tagliasacchi, et al. Vd3d: Taming large video diffusion transformers for 3d camera control. *arXiv preprint arXiv:2407.12781*, 2024. 2, 3, 4, 5, 6, 7, 8, 18
- [7] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendeleevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023. 2, 3
- [8] 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 *Proc. CVPR*, 2023. 2
- [9] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. 6
- [10] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, Clarence Ng, Ricky Wang, and Aditya Ramesh. Video generation models as world simulators. *OpenAI technical reports*, 2024. 1, 2, 3, 17
- [11] Yukang Cao, Liang Pan, Kai Han, Kwan-Yee K Wong, and Ziwei Liu. Avatargo: Zero-shot 4d human-object interaction generation and animation. *arXiv preprint arXiv:2410.07164*, 2024. 20
- [12] Zenghao Chai, Chen Tang, Yongkang Wong, and Mohan Kankanhalli. Star: Skeleton-aware text-based 4d avatar generation with in-network motion retargeting. *arXiv preprint arXiv:2406.04629*, 2024. 20
- [13] 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. *Proc. CVPR*, 2022. 20
- [14] Eric R Chan, Koki Nagano, Matthew A Chan, Alexander W Bergman, Jeong Joon Park, Axel Levy, Miika Aittala, Shalini De Mello, Tero Karras, and Gordon Wetzstein. Generative novel view synthesis with 3d-aware diffusion models. In *Proc. ICCV*, 2023. 20
- [15] Ce Chen, Shaoli Huang, Xuelin Chen, Guangyi Chen, Xiaoguang Han, Kun Zhang, and Mingming Gong. Ct4d: Consistent text-to-4d generation with animatable meshes. *arXiv preprint arXiv:2408.08342*, 2024. 20
- [16] Eric Ming Chen, Sidhanth Holalkere, Ruyu Yan, Kai Zhang, and Abe Davis. Ray conditioning: Trading photo-consistency for photo-realism in multi-view image generation. In *Proc. ICCV*, 2023. 4
- [17] Haoxin Chen, Menghan Xia, Yingqing He, Yong Zhang, Xiaodong Cun, Shaoshu Yang, Jinbo Xing, Yaofang Liu, Qifeng Chen, Xintao Wang, Chao Weng, and Ying Shan. Videocrafter1: Open diffusion models for high-quality video generation. *arXiv preprint arXiv:2310.19512*, 2023. 3
- [18] Junsong Chen, Yue Wu, Simian Luo, Enze Xie, Sayak Paul, Ping Luo, Hang Zhao, and Zhenguo Li. Pixart- $\{\delta\}$: Fast and controllable image generation with latent consistency models. *arXiv preprint arXiv:2401.05252*, 2024. 3
- [19] Kevin Chen, Christopher B Choy, Manolis Savva, Angel X Chang, Thomas Funkhouser, and Silvio Savarese. Text2Shape: Generating shapes from natural language by learning joint embeddings. *Proc. ACCV*, 2018. 20
- [20] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3D: Disentangling geometry and appearance for high-quality text-to-3D content creation. *arXiv preprint arXiv:2303.13873*, 2023. 20
- [21] Ting Chen and Lala Li. Fit: Far-reaching interleaved transformers. *arXiv preprint arXiv:2305.12689*, 2023. 7
- [22] Soon Yau Cheong, Duygu Ceylan, Armin Mustafa, Andrew Gilbert, and Chun-Hao Paul Huang. Boosting camera motion control for video diffusion transformers. *arXiv preprint arXiv:2410.10802*, 2024. 3
- [23] Wen-Hsuan Chu, Lei Ke, and Katerina Fragkiadaki. Dreamscene4d: Dynamic multi-object scene generation from monocular videos. *arXiv preprint arXiv:2405.02280*, 2024. 2, 20
- [24] Terrance DeVries, Miguel Angel Bautista, Nitish Srivastava, Graham W Taylor, and Joshua M Susskind. Unconstrained scene generation with locally conditioned radiance fields. *Proc. ICCV*, 2021. 20
- [25] Sander Dieleman. Perspectives on diffusion, 2023. 2
- [26] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *Proc. ICLR*, 2021. 18
- [27] Mohamed El Banani, Amit Raj, Kevis-Kokitsi Maninis, Abhishek Kar, Yuanzhen Li, Michael Rubinstein, Deqing Sun,

- Leonidas Guibas, Justin Johnson, and Varun Jampani. Probing the 3d awareness of visual foundation models. In *Proc. CVPR*, 2024. 2
- [28] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024. 4, 18
- [29] Gunnar Farnebäck. Two-frame motion estimation based on polynomial expansion. In *Image Analysis: 13th Scandinavian Conference, SCIA 2003 Halmstad, Sweden, June 29–July 2, 2003 Proceedings 13*, pages 363–370. Springer, 2003. 19
- [30] Qijun Feng, Zhen Xing, Zuxuan Wu, and Yu-Gang Jiang. FDGaussian: Fast Gaussian splatting from single image via geometric-aware diffusion model. *arXiv preprint arXiv:2403.10242*, 2024. 20
- [31] Yutao Feng, Yintong Shang, Xiang Feng, Lei Lan, Shandian Zhe, Tianjia Shao, Hongzhi Wu, Kun Zhou, Hao Su, Chenfanfu Jiang, et al. Elastogen: 4d generative elastodynamics. *arXiv preprint arXiv:2405.15056*, 2024. 20
- [32] Peng Gao, Le Zhuo, Ziyi Lin, Dongyang Liu, Ruoyi Du, Xu Luo, Longtian Qiu, Yuhang Zhang, et al. Lumina-t2x: Transforming text into any modality, resolution, and duration via flow-based large diffusion transformers. *arXiv preprint arXiv:2405.05945*, 2024. 3, 17, 18
- [33] Quankai Gao, Qiangeng Xu, Zhe Cao, Ben Mildenhall, Wenchao Ma, Le Chen, Danhang Tang, and Ulrich Neumann. Gaussianflow: Splatting Gaussian dynamics for 4D content creation. *arXiv preprint arXiv:2403.12365*, 2024. 2, 20
- [34] William Gao, Noam Aigerman, Thibault Groueix, Vova Kim, and Rana Hanocka. TextDeformer: Geometry manipulation using text guidance. *Proc. SIGGRAPH*, 2023. 20
- [35] Jiatao Gu, Alex Trevithick, Kai-En Lin, Joshua M Susskind, Christian Theobalt, Lingjie Liu, and Ravi Ramamoorthi. NerfDiff: Single-image view synthesis with NeRF-guided distillation from 3D-aware diffusion. *Proc. ICML*, 2023. 20
- [36] Yuwei Guo, Ceyuan Yang, Anyi Rao, Yaohui Wang, Yu Qiao, Dahua Lin, and Bo Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *Proc. ICLR*, 2024. 2, 7
- [37] Junlin Han, Filippos Kokkinos, and Philip Torr. VFusion3D: Learning scalable 3D generative models from video diffusion models. *arXiv preprint arXiv:2403.12034*, 2024. 20
- [38] Hao He, Yinghao Xu, Yuwei Guo, Gordon Wetzstein, Bo Dai, Hongsheng Li, and Ceyuan Yang. Cameractrl: Enabling camera control for text-to-video generation. *arXiv preprint arXiv:2404.02101*, 2024. 2, 3, 4, 5, 6, 7, 8, 20
- [39] Xianglong He, Junyi Chen, Sida Peng, Di Huang, Yangguang Li, Xiaoshui Huang, Chun Yuan, Wanli Ouyang, and Tong He. GVGEn: Text-to-3D generation with volumetric representation. *arXiv preprint arXiv:2403.12957*, 2024. 20
- [40] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 18
- [41] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. *arXiv preprint arXiv:2104.08718*, 2021. 7
- [42] 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. *Proc. NeurIPS*, 2017. 7
- [43] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 4, 8
- [44] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 6
- [45] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 3
- [46] Lukas Höllein, Ang Cao, Andrew Owens, Justin Johnson, and Matthias Nießner. Text2room: Extracting textured 3d meshes from 2d text-to-image models. In *Proc. ICCV*, 2023. 20
- [47] Lukas Höllein, Aljaž Božič, Norman Müller, David Novotny, Hung-Yu Tseng, Christian Richardt, Michael Zollhöfer, and Matthias Nießner. ViewDiff: 3D-consistent image generation with text-to-image models. *Proc. CVPR*, 2024. 20
- [48] Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. LRM: Large reconstruction model for single image to 3D. *Proc. ICLR*, 2024. 20
- [49] Chen Hou, Guoqiang Wei, Yan Zeng, and Zhibo Chen. Training-free camera control for video generation. *arXiv preprint arXiv:2406.10126*, 2024. 3
- [50] Teng Hu, Jiangning Zhang, Ran Yi, Yating Wang, Hongrui Huang, Jieyu Weng, Yabiao Wang, and Lizhuang Ma. Motionmaster: Training-free camera motion transfer for video generation. *arXiv preprint arXiv:2404.15789*, 2024. 3
- [51] Chun-Hao Paul Huang, Jae Shin Yoon, Hyeonho Jeong, Niloy Mitra, and Duygu Ceylan. Camera pose estimation emerging in video diffusion transformer, 2024. 3
- [52] Tianyu Huang, Yihan Zeng, Hui Li, Wangmeng Zuo, and Rynson WH Lau. Dreamphysics: Learning physical properties of dynamic 3d gaussians with video diffusion priors. *arXiv preprint arXiv:2406.01476*, 2024. 20
- [53] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. *Proc. CVPR*, 2022. 20
- [54] Yash Jain, Anshul Nasery, Vibhav Vineet, and Harkirat Behl. Peekaboo: Interactive video generation via masked-diffusion. In *CVPR*, 2024. 21
- [55] Nikolay Jetchev. ClipMatrix: Text-controlled creation of 3D textured meshes. *arXiv preprint arXiv:2109.12922*, 2021. 20
- [56] Lutao Jiang and Lin Wang. Brightdreamer: Generic 3D Gaussian generative framework for fast text-to-3D synthesis. *arXiv preprint arXiv:2403.11273*, 2024. 20
- [57] Yanqin Jiang, Chaohui Yu, Chenjie Cao, Fan Wang, Weiming Hu, and Jin Gao. Animate3d: Animating any 3d model with multi-view video diffusion. *arXiv preprint arXiv:2407.11398*, 2024. 20
- [58] Yanqin Jiang, Li Zhang, Jin Gao, Weimin Hu, and Yao Yao. Consistent4d: Consistent 360 $\{\deg\}$ dynamic object generation from monocular video. *Proc. ICLR*, 2024. 20

- [59] Yash Kant, Aliaksandr Siarohin, Ziyi Wu, Michael Vasilkovsky, Guocheng Qian, Jian Ren, Riza Alp Guler, Bernard Ghanem, Sergey Tulyakov, and Igor Gilitschenski. Spad: Spatially aware multi-view diffusers. In *Proc. CVPR*, 2024. 4
- [60] Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. Analyzing and improving the training dynamics of diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24174–24184, 2024. 18
- [61] Oren Katzir, Or Patashnik, Daniel Cohen-Or, and Dani Lischinski. Noise-free score distillation. *Proc. ICLR*, 2024. 20
- [62] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM TOG*, 2023. 20
- [63] Seung Wook Kim, Bradley Brown, Kangxue Yin, Karsten Kreis, Katja Schwarz, Daiqing Li, Robin Rombach, Antonio Torralba, and Sanja Fidler. NeuralField-LDM: Scene generation with hierarchical latent diffusion models. *Proc. CVPR*, 2023. 20
- [64] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 3
- [65] Zhengfei Kuang, Shengqu Cai, Hao He, Yinghao Xu, Hongsheng Li, Leonidas Guibas, and Gordon Wetzstein. Collaborative video diffusion: Consistent multi-video generation with camera control. *arXiv preprint arXiv:2405.17414*, 2024. 3
- [66] Kyungmin Lee, Kihyuk Sohn, and Jinwoo Shin. DreamFlow: High-quality text-to-3D generation by approximating probability flow. *Proc. ICLR*, 2024. 20
- [67] Yao-Chih Lee, Yi-Ting Chen, Andrew Wang, Ting-Hsuan Liao, Brandon Y Feng, and Jia-Bin Huang. Vividdream: Generating 3d scene with ambient dynamics. *arXiv preprint arXiv:2405.20334*, 2024. 20
- [68] Guojun Lei, Chi Wang, Hong Li, Rong Zhang, Yikai Wang, and Weiwei Xu. Animateanything: Consistent and controllable animation for video generation. *arXiv preprint arXiv:2411.10836*, 2024. 3
- [69] Bing Li, Cheng Zheng, Wenxuan Zhu, Jinjie Mai, Biao Zhang, Peter Wonka, and Bernard Ghanem. Vivid-zoo: Multi-view video generation with diffusion model. *arXiv preprint arXiv:2406.08659*, 2024. 20
- [70] Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan Sunkavalli, Greg Shakhnarovich, and Sai Bi. Instant3D: Fast text-to-3D with sparse-view generation and large reconstruction model. *Proc. ICLR*, 2024. 20
- [71] Ming Li, Taojiannan Yang, Huafeng Kuang, Jie Wu, Zhaoning Wang, Xuefeng Xiao, and Chen Chen. Controlnet++: Improving conditional controls with efficient consistency feedback. In *European Conference on Computer Vision*, pages 129–147. Springer, 2024. 3, 4, 6, 18
- [72] Renjie Li, Panwang Pan, Bangbang Yang, Dejia Xu, Shijie Zhou, Xuanyang Zhang, Zeming Li, Achuta Kadambi, Zhangyang Wang, and Zhiwen Fan. 4k4dgen: Panoramic 4d generation at 4k resolution. *arXiv preprint arXiv:2406.13527*, 2024. 20
- [73] Zhiqi Li, Yiming Chen, and Peidong Liu. Dreammesh4d: Video-to-4d generation with sparse-controlled gaussian-mesh hybrid representation. *arXiv preprint arXiv:2410.06756*, 2024. 20
- [74] Zhiqi Li, Yiming Chen, Lingzhe Zhao, and Peidong Liu. Controllable text-to-3D generation via surface-aligned Gaussian splatting. *arXiv preprint arXiv:2403.09981*, 2024. 20
- [75] Zhengqi Li, Richard Tucker, Noah Snavely, and Aleksander Holynski. Generative image dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 4, 17, 19
- [76] Hanwen Liang, Yuyang Yin, Dejia Xu, Hanxue Liang, Zhangyang Wang, Konstantinos N Plataniotis, Yao Zhao, and Yunchao Wei. Diffusion4d: Fast spatial-temporal consistent 4d generation via video diffusion models. *arXiv preprint arXiv:2405.16645*, 2024. 2, 20
- [77] Yixun Liang, Xin Yang, Jiantao Lin, Haodong Li, Xiaogang Xu, and Yingcong Chen. Luciddreamer: Towards high-fidelity text-to-3D generation via interval score matching. *arXiv preprint arXiv:2311.11284*, 2023. 20
- [78] 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. *Proc. CVPR*, 2023. 20
- [79] Jiajing Lin, Zhenzhong Wang, Yongjie Hou, Yuzhou Tang, and Min Jiang. Phy124: Fast physics-driven 4d content generation from a single image. *arXiv preprint arXiv:2409.07179*, 2024. 20
- [80] Yukang Lin, Haonan Han, Chaoqun Gong, Zunnan Xu, Yachao Zhang, and Xiu Li. Consistent123: One image to highly consistent 3D asset using case-aware diffusion priors. *arXiv preprint arXiv:2309.17261*, 2023. 20
- [81] Huan Ling, Seung Wook Kim, Antonio Torralba, Sanja Fidler, and Karsten Kreis. Align your gaussians: Text-to-4d with dynamic 3d gaussians and composed diffusion models. In *Proc. CVPR*, 2024. 2, 20
- [82] Pengyang Ling, Jiazi Bu, Pan Zhang, Xiaoyi Dong, Yuhang Zang, Tong Wu, Huaian Chen, Jiaqi Wang, and Yi Jin. Motionclone: Training-free motion cloning for controllable video generation. *arXiv preprint arXiv:2406.05338*, 2024. 3
- [83] Pengkun Liu, Yikai Wang, Fuchun Sun, Jiafang Li, Hang Xiao, Hongxiang Xue, and Xinzhou Wang. Isotropic3D: Image-to-3D generation based on a single clip embedding. *arXiv preprint arXiv:2403.10395*, 2024. 20
- [84] 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. *Proc. ICCV*, 2023. 20
- [85] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. 3, 4
- [86] Xian Liu, Xiaohang Zhan, Jiayang Tang, Ying Shan, Gang Zeng, Dahua Lin, Xihui Liu, and Ziwei Liu. HumanGaussian: Text-driven 3D human generation with Gaussian splatting. *Proc. CVPR*, 2024. 20
- [87] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. SyncDreamer:

- Generating multiview-consistent images from a single-view image. *Proc. ICLR*, 2024. 20
- [88] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3D: Single image to 3D using cross-domain diffusion. *Proc. CVPR*, 2024. 20
- [89] I Loshchilov. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 18
- [90] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 18
- [91] Wan-Duo Kurt Ma, John P Lewis, and W Bastiaan Kleijn. Trailblazer: Trajectory control for diffusion-based video generation. *arXiv preprint arXiv:2401.00896*, 2023. 21
- [92] Xin Ma, Yaohui Wang, Gengyun Jia, Xinyuan Chen, Ziwei Liu, Yuan-Fang Li, Cunjian Chen, and Yu Qiao. Latte: Latent diffusion transformer for video generation. *arXiv preprint arXiv:2401.03048*, 2024. 2
- [93] Willi Menapace, Aliaksandr Siarohin, Ivan Skorokhodov, Ekaterina Deyneka, Tsai-Shien Chen, Anil Kag, Yuwei Fang, Aleksei Stoliar, Elisa Ricci, Jian Ren, et al. Snap video: Scaled spatiotemporal transformers for text-to-video synthesis. *Proc. CVPR*, 2024. 2, 7, 18
- [94] Qiaowei Miao, Yawei Luo, and Yi Yang. Pla4d: Pixel-level alignments for text-to-4d gaussian splatting. *arXiv preprint arXiv:2405.19957*, 2024. 20
- [95] 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 *Proc. ECCV*, 2020. 20
- [96] Koichi Namekata, Sherwin Bahmani, Ziyi Wu, Yash Kant, Igor Gilitschenski, and David B Lindell. Sg-i2v: Self-guided trajectory control in image-to-video generation. *arXiv preprint arXiv:2411.04989*, 2024. 21
- [97] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. StyleSDF: High-resolution 3D-consistent image and geometry generation. *Proc. CVPR*, 2022. 20
- [98] Zijie Pan, Zeyu Yang, Xiatian Zhu, and Li Zhang. Fast dynamic 3d object generation from a single-view video. *arXiv preprint arXiv:2401.08742*, 2024. 2, 20
- [99] Deepak Pathak, Ross Girshick, Piotr Dollár, Trevor Darrell, and Bharath Hariharan. Learning features by watching objects move. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2701–2710, 2017. 19
- [100] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. 2, 3, 7, 17, 18
- [101] Ryan Po, Wang Yifan, Vladislav Golyanik, Kfir Aberman, Jonathan T Barron, Amit H Bermano, Eric Ryan Chan, Tali Dekel, Aleksander Holynski, Angjoo Kanazawa, et al. State of the art on diffusion models for visual computing. *arXiv preprint arXiv:2310.07204*, 2023. 2
- [102] Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, et al. Movie gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720*, 2024. 3, 17, 18
- [103] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. DreamFusion: Text-to-3D using 2D diffusion. In *Proc. ICLR*, 2023. 20
- [104] Guocheng Qian, Junli Cao, Aliaksandr Siarohin, Yash Kant, Chaoyang Wang, Michael Vasilkovsky, Hsin-Ying Lee, Yuwei Fang, Ivan Skorokhodov, Peiye Zhuang, et al. Atom: Amortized text-to-mesh using 2d diffusion. *arXiv preprint arXiv:2402.00867*, 2024. 20
- [105] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image to high-quality 3D object generation using both 2D and 3D diffusion priors. *Proc. ICLR*, 2024. 20
- [106] Haonan Qiu, Zhaoxi Chen, Zhouxia Wang, Yingqing He, Menghan Xia, and Ziwei Liu. Freetraj: Tuning-free trajectory control in video diffusion models. *arXiv preprint arXiv:2406.16863*, 2024. 21
- [107] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. *Proc. ICML*, 2021. 20
- [108] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Proc. JMLR*, 2020. 3, 17
- [109] Jiawei Ren, Liang Pan, Jiayang Tang, Chi Zhang, Ang Cao, Gang Zeng, and Ziwei Liu. DreamGaussian4D: Generative 4D Gaussian splatting. *arXiv preprint arXiv:2312.17142*, 2023. 2, 20
- [110] Jiawei Ren, Kevin Xie, Ashkan Mirzaei, Hanxue Liang, Xiaohui Zeng, Karsten Kreis, Ziwei Liu, Antonio Torralba, Sanja Fidler, Seung Wook Kim, et al. L4gm: Large 4d gaussian reconstruction model. *arXiv preprint arXiv:2406.10324*, 2024. 3, 20
- [111] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 3, 19
- [112] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3, 19
- [113] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. 4

- [114] Aditya Sanghi, Hang Chu, Joseph G Lambourne, Ye Wang, Chin-Yi Cheng, Marco Fumero, and Kamal Rahimi Malekshah. CLIP-Forge: Towards zero-shot text-to-shape generation. *Proc. CVPR*, 2022. 20
- [115] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5, 20
- [116] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016. 5, 20
- [117] Katja Schwarz, Axel Sauer, Michael Niemeyer, Yiyi Liao, and Andreas Geiger. VoxGRAF: Fast 3D-aware image synthesis with sparse voxel grids. *Proc. NeurIPS*, 2022. 20
- [118] Yichun Shi, Peng Wang, Jialong Ye, Long Mai, Kejie Li, and Xiao Yang. MVDream: Multi-view diffusion for 3D generation. *Proc. ICLR*, 2024. 20
- [119] Jaidev Shriram, Alex Trevithick, Lingjie Liu, and Ravi Ramamoorthi. Realmdreamer: Text-driven 3d scene generation with inpainting and depth diffusion. *arXiv preprint arXiv:2404.07199*, 2024. 20
- [120] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oran Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *Proc. ICLR*, 2023. 2
- [121] Uriel Singer, Shelly Sheynin, Adam Polyak, Oran Ashual, Iurii Makarov, Filippos Kokkinos, Naman Goyal, Andrea Vedaldi, Devi Parikh, Justin Johnson, et al. Text-to-4d dynamic scene generation. In *Proc. ICML*, 2023. 2, 20
- [122] Vincent Sitzmann, Semon Rezchikov, Bill Freeman, Josh Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. *Proc. NeurIPS*, 2021. 4
- [123] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015. 3
- [124] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. 17
- [125] Jingxiang Sun, Bo Zhang, Ruizhi Shao, Lizhen Wang, Wen Liu, Zhenda Xie, and Yebin Liu. DreamCraft3D: Hierarchical 3D generation with bootstrapped diffusion prior. *Proc. ICLR*, 2024. 20
- [126] Keqiang Sun, Dor Litvak, Yunzhi Zhang, Hongsheng Li, Jiajun Wu, and Shangzhe Wu. Ponymation: Learning articulated 3d animal motions from unlabeled online videos. In *European Conference on Computer Vision*, 2024. 20
- [127] Qi Sun, Zhiyang Guo, Ziyu Wan, Jing Nathan Yan, Shengming Yin, Wengang Zhou, Jing Liao, and Houqiang Li. Eg4d: Explicit generation of 4d object without score distillation. *arXiv preprint arXiv:2405.18132*, 2024. 20
- [128] Wenqiang Sun, Shuo Chen, Fangfu Liu, Zilong Chen, Yueqi Duan, Jun Zhang, and Yikai Wang. Dimensionx: Create any 3d and 4d scenes from a single image with controllable video diffusion. *arXiv preprint arXiv:2411.04928*, 2024. 3
- [129] Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. Viewset diffusion:(0-) image-conditioned 3d generative models from 2d data. In *Proc. ICCV*, 2023. 20
- [130] Stanislaw Szymanowicz, Eldar Insafutdinov, Chuanxia Zheng, Dylan Campbell, João F Henriques, Christian Rupprecht, and Andrea Vedaldi. Flash3d: Feed-forward generalisable 3d scene reconstruction from a single image. *arXiv preprint arXiv:2406.04343*, 2024. 20
- [131] Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. Splatter image: Ultra-fast single-view 3D reconstruction. *Proc. CVPR*, 2024. 20
- [132] 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. *arXiv preprint arXiv:2303.14184*, 2023. 20
- [133] Jiayang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. LGM: Large multi-view gaussian model for high-resolution 3d content creation. *Proc. ECCV*, 2024. 20
- [134] Zhenggang Tang, Peiye Zhuang, Chaoyang Wang, Aliaksandr Siarohin, Yash Kant, Alexander Schwing, Sergey Tulyakov, and Hsin-Ying Lee. Pixel-aligned multi-view generation with depth guided decoder. *arXiv preprint arXiv:2408.14016*, 2024. 20
- [135] Ayush Tewari, Tianwei Yin, George Cazenavette, Semon Rezchikov, Josh Tenenbaum, Frédo Durand, Bill Freeman, and Vincent Sitzmann. Diffusion with forward models: Solving stochastic inverse problems without direct supervision. *Proc. NeurIPS*, 2023. 20
- [136] Dmitry Tochilkin, David Pankratz, Zexiang Liu, Zixuan Huang, Adam Letts, Yangguang Li, Ding Liang, Christian Laforte, Varun Jampani, and Yan-Pei Cao. Triposr: Fast 3D object reconstruction from a single image. *arXiv preprint arXiv:2403.02151*, 2024. 20
- [137] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018. 7
- [138] Lukas Uzolas, Elmar Eisemann, and Petr Kellnhofer. Motiondreamer: Zero-shot 3d mesh animation from video diffusion models. *arXiv preprint arXiv:2405.20155*, 2024. 20
- [139] Basile Van Hoorick, Rundi Wu, Ege Ozguroglu, Kyle Sargent, Ruoshi Liu, Pavel Tokmakov, Achal Dave, Changxi Zheng, and Carl Vondrick. Generative camera dolly: Extreme monocular dynamic novel view synthesis. *arXiv preprint arXiv:2405.14868*, 2024. 20
- [140] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Proc. NeurIPS*, 2017. 3
- [141] Vikram Voleti, Chun-Han Yao, Mark Boss, Adam Letts, David Pankratz, Dmitry Tochilkin, Christian Laforte, Robin Rombach, and Varun Jampani. SV3D: Novel multi-view synthesis and 3D generation from a single image using latent video diffusion. *arXiv preprint arXiv:2403.12008*, 2024. 20
- [142] Ziyu Wan, Despoina Paschalidou, Ian Huang, Hongyu Liu, Bokui Shen, Xiaoyu Xiang, Jing Liao, and Leonidas Guibas.

- CAD: Photorealistic 3D generation via adversarial distillation. *Proc. CVPR*, 2024. 20
- [143] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-NeRF: Text-and-image driven manipulation of neural radiance fields. *Proc. CVPR*, 2022. 20
- [144] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score Jacobian chaining: Lifting pretrained 2d diffusion models for 3D generation. *Proc. CVPR*, 2023. 20
- [145] Jiawei Wang, Yuchen Zhang, Jiaxin Zou, Yan Zeng, Guoqiang Wei, Liping Yuan, and Hang Li. Boximator: Generating rich and controllable motions for video synthesis. *arXiv preprint arXiv:2402.01566*, 2024. 21
- [146] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *CVPR*, 2024. 3, 5, 6
- [147] Xiang Wang, Hangjie Yuan, Shiwei Zhang, Dayou Chen, Jiniu Wang, Yingya Zhang, Yujun Shen, Deli Zhao, and Jingren Zhou. Videocomposer: Compositional video synthesis with motion controllability. *arXiv preprint arXiv:2306.02018*, 2023. 21
- [148] Yikai Wang, Xinzhou Wang, Zilong Chen, Zhengyi Wang, Fuchun Sun, and Jun Zhu. Vidu4d: Single generated video to high-fidelity 4d reconstruction with dynamic gaussian surfels. *arXiv preprint arXiv:2405.16822*, 2024. 20
- [149] 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. *Proc. NeurIPS*, 2023. 20
- [150] Zhouxia Wang, Ziyang Yuan, Xintao Wang, Tianshui Chen, Menghan Xia, Ping Luo, and Yin Shan. Motionctrl: A unified and flexible motion controller for video generation. In *arXiv preprint arXiv:2312.03641*, 2023. 2, 3, 4, 5, 6, 7, 8
- [151] Daniel Watson, Saurabh Saxena, Lala Li, Andrea Tagliasacchi, and David J Fleet. Controlling space and time with diffusion models. *arXiv preprint arXiv:2407.07860*, 2024. 3, 4, 5, 6, 8
- [152] Ruiqi Wu, , Liangyu Chen, Tong Yang, Chunle Guo, Chongyi Li, and Xiangyu Zhang. Lamp: Learn a motion pattern for few-shot-based video generation. *arXiv preprint arXiv:2310.10769*, 2023. 2
- [153] Wejia Wu, Zhuang Li, Yuchao Gu, Rui Zhao, Yefei He, David Junhao Zhang, Mike Zheng Shou, Yan Li, Tingting Gao, and Di Zhang. Draganything: Motion control for anything using entity representation. In *ECCV*, 2024. 21
- [154] Zijie Wu, Chaohui Yu, Yanqin Jiang, Chenjie Cao, Fan Wang, and Xiang Bai. Sc4d: Sparse-controlled video-to-4d generation and motion transfer. *arXiv preprint arXiv:2404.03736*, 2024. 20
- [155] Zeqi Xiao, Yifan Zhou, Shuai Yang, and Xingang Pan. Video diffusion models are training-free motion interpreter and controller. *arXiv preprint arXiv:2405.14864*, 2024. 3
- [156] Kevin Xie, Jonathan Lorraine, Tianshi Cao, Jun Gao, James Lucas, Antonio Torralba, Sanja Fidler, and Xiaohui Zeng. LATTE3D: Large-scale amortized text-to-enhanced3D synthesis. *Proc. ECCV*, 2024. 20
- [157] Yiming Xie, Chun-Han Yao, Vikram Voleti, Huaizu Jiang, and Varun Jampani. Sv4d: Dynamic 3d content generation with multi-frame and multi-view consistency. *arXiv preprint arXiv:2407.17470*, 2024. 20
- [158] Dejia Xu, Yifan Jiang, Chen Huang, Liangchen Song, Thorsten Gernoth, Liangliang Cao, Zhangyang Wang, and Hao Tang. Cavia: Camera-controllable multi-view video diffusion with view-integrated attention. *arXiv preprint arXiv:2410.10774*, 2024. 3
- [159] Dejia Xu, Hanwen Liang, Neel P Bhatt, Hezhen Hu, Hanxue Liang, Konstantinos N Plataniotis, and Zhangyang Wang. Comp4d: Llm-guided compositional 4d scene generation. *arXiv preprint arXiv:2403.16993*, 2024. 2, 20
- [160] Dejia Xu, Weili Nie, Chao Liu, Sifei Liu, Jan Kautz, Zhangyang Wang, and Arash Vahdat. Camco: Camera-controllable 3d-consistent image-to-video generation. *arXiv preprint arXiv:2406.02509*, 2024. 3
- [161] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *Proc. CVPR*, 2016. 7, 8
- [162] Yinghao Xu, Zifan Shi, Wang Yifan, Hansheng Chen, Ceyuan Yang, Sida Peng, Yujun Shen, and Gordon Wetzstein. GRM: Large Gaussian reconstruction model for efficient 3D reconstruction and generation. *Proc. ECCV*, 2024. 20
- [163] Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli, Gordon Wetzstein, Zexiang Xu, et al. DMV3D: Denoising multi-view diffusion using 3D large reconstruction model. *Proc. ICLR*, 2024. 20
- [164] Zhongcong Xu, Jianfeng Zhang, Jun Hao Liew, Wenqing Zhang, Song Bai, Jiashi Feng, and Mike Zheng Shou. Pv3d: A 3d generative model for portrait video generation. *Proc. ICLR*, 2023. 2, 20
- [165] Qitong Yang, Mingtao Feng, Zijie Wu, Shijie Sun, Weisheng Dong, Yaonan Wang, and Ajmal Mian. Beyond skeletons: Integrative latent mapping for coherent 4d sequence generation. *arXiv preprint arXiv:2403.13238*, 2024. 20
- [166] Shiyuan Yang, Liang Hou, Haibin Huang, Chongyang Ma, Pengfei Wan, Di Zhang, Xiaodong Chen, and Jing Liao. Direct-a-video: Customized video generation with user-directed camera movement and object motion. *arXiv preprint arXiv:2402.03162*, 2024. 2, 21
- [167] Zeyu Yang, Zijie Pan, Chun Gu, and Li Zhang. Diffusion²: Dynamic 3d content generation via score composition of orthogonal diffusion models. *arXiv preprint 2404.02148*, 2024. 20
- [168] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyuan Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. 2, 3, 6, 17, 18, 19
- [169] Junliang Ye, Fangfu Liu, Qixiu Li, Zhengyi Wang, Yikai Wang, Xinzhou Wang, Yueqi Duan, and Jun Zhu. DreamReward: Text-to-3D generation with human preference. *arXiv preprint arXiv:2403.14613*, 2024. 20
- [170] Shengming Yin, Chenfei Wu, Jian Liang, Jie Shi, Houqiang Li, Gong Ming, and Nan Duan. Dragnuwa: Fine-grained

- control in video generation by integrating text, image, and trajectory. *arXiv preprint arXiv:2308.08089*, 2023. 21
- [171] Wei Yin, Chi Zhang, Hao Chen, Zhipeng Cai, Gang Yu, Kaixuan Wang, Xiaozhi Chen, and Chunhua Shen. Metric3d: Towards zero-shot metric 3d prediction from A single image. In *ICCV*, pages 9009–9019. IEEE, 2023. 20
- [172] Yuyang Yin, Deji Xu, Zhangyang Wang, Yao Zhao, and Yunchao Wei. 4dgen: Grounded 4d content generation with spatial-temporal consistency. *arXiv preprint arXiv:2312.17225*, 2023. 2, 20
- [173] Paul Yoo, Jiaxian Guo, Yutaka Matsuo, and Shixiang Shane Gu. DreamSparse: Escaping from Plato’s cave with 2D diffusion model given sparse views. *arXiv preprint arXiv:2306.03414*, 2023. 20
- [174] Heng Yu, Chaoyang Wang, Peiye Zhuang, Willi Menapace, Aliaksandr Siarohin, Junli Cao, Laszlo A Jeni, Sergey Tulyakov, and Hsin-Ying Lee. 4real: Towards photorealistic 4d scene generation via video diffusion models. *arXiv preprint arXiv:2406.07472*, 2024. 2, 20
- [175] Shoubin Yu, Jacob Zhiyuan Fang, Jian Zheng, Gunnar A Sigurdsson, Vicente Ordonez, Robinson Piramuthu, and Mohit Bansal. Zero-shot controllable image-to-video animation via motion decomposition. In *ACM MM*, 2024. 21
- [176] Wangbo Yu, Jinbo Xing, Li Yuan, Wenbo Hu, Xiaoyu Li, Zhipeng Huang, Xiangjun Gao, Tien-Tsin Wong, Ying Shan, and Yonghong Tian. Viewcrafter: Taming video diffusion models for high-fidelity novel view synthesis. *arXiv preprint arXiv:2409.02048*, 2024. 3
- [177] Xin Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Song-Hai Zhang, and Xiaojuan Qi. Text-to-3D with classifier score distillation. *arXiv preprint arXiv:2310.19415*, 2023. 20
- [178] Yu-Jie Yuan, Leif Kobbelt, Jiwen Liu, Yuan Zhang, Pengfei Wan, Yu-Kun Lai, and Lin Gao. 4dynamic: Text-to-4d generation with hybrid priors. *arXiv preprint arXiv:2407.12684*, 2024. 20
- [179] Raza Yunus, Jan Eric Lenssen, Michael Niemeyer, Yiyi Liao, Christian Rupprecht, Christian Theobalt, Gerard Pons-Moll, Jia-Bin Huang, Vladislav Golyanik, and Eddy Ilg. Recent trends in 3d reconstruction of general non-rigid scenes. In *Computer Graphics Forum*, 2024. 2
- [180] Bohan Zeng, Ling Yang, Siyu Li, Jiaming Liu, Zixiang Zhang, Juanxi Tian, Kaixin Zhu, Yongzhen Guo, Fu-Yun Wang, Minkai Xu, et al. Trans4d: Realistic geometry-aware transition for compositional text-to-4d synthesis. *arXiv preprint arXiv:2410.07155*, 2024. 20
- [181] Yifei Zeng, Yanqin Jiang, Siyu Zhu, Yuanxun Lu, Youtian Lin, Hao Zhu, Weiming Hu, Xun Cao, and Yao Yao. Stag4d: Spatial-temporal anchored generative 4d gaussians. *arXiv preprint arXiv:2403.14939*, 2024. 2, 20
- [182] Biao Zhang and Rico Sennrich. Root mean square layer normalization. *Advances in Neural Information Processing Systems*, 32, 2019. 17
- [183] Bowen Zhang, Tianyu Yang, Yu Li, Lei Zhang, and Xi Zhao. Compress3D: a compressed latent space for 3D generation from a single image. *arXiv preprint arXiv:2403.13524*, 2024. 20
- [184] David Junhao Zhang, Roni Paiss, Shiran Zada, Nikhil Karnad, David E Jacobs, Yael Pritch, Inbar Mosseri, Mike Zheng Shou, Neal Wadhwa, and Nataniel Ruiz. Recapture: Generative video camera controls for user-provided videos using masked video fine-tuning. *arXiv preprint arXiv:2411.05003*, 2024. 3
- [185] Hao Zhang, Di Chang, Fang Li, Mohammad Soleymani, and Narendra Ahuja. Magicpose4d: Crafting articulated models with appearance and motion control. *arXiv preprint arXiv:2405.14017*, 2024. 20
- [186] Haiyu Zhang, Xinyuan Chen, Yaohui Wang, Xihui Liu, Yunhong Wang, and Yu Qiao. 4diffusion: Multi-view video diffusion model for 4d generation. *arXiv preprint arXiv:2405.20674*, 2024. 20
- [187] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. *Proc. ICCV*, 2023. 3, 6, 18
- [188] Tianyuan Zhang, Hong-Xing Yu, Rundi Wu, Brandon Y Feng, Changxi Zheng, Noah Snively, Jiajun Wu, and William T Freeman. Physdreamer: Physics-based interaction with 3d objects via video generation. *arXiv preprint arXiv:2404.13026*, 2024. 20
- [189] Zhichao Zhang, Hui Chen, Jinsheng Deng, Xiaoqing Yin, Xingshen Song, and Ming Xu. Motion4d: A decoupled pipeline for enhanced text-to-4d generation with optimized motion patterns. *SSRN*, 2024. 20
- [190] Zhenghao Zhang, Junchao Liao, Menghao Li, Long Qin, and Weizhi Wang. Tora: Trajectory-oriented diffusion transformer for video generation. *arXiv preprint arXiv:2407.21705*, 2024. 21
- [191] Wang Zhao, Shaohui Liu, Hengkai Guo, Wenping Wang, and Yong-Jin Liu. Particlesfm: Exploiting dense point trajectories for localizing moving cameras in the wild. In *Proc. ECCV*, 2022. 5, 7, 8
- [192] Yuyang Zhao, Zhiwen Yan, Enze Xie, Lanqing Hong, Zhen-guo Li, and Gim Hee Lee. Animate124: Animating one image to 4d dynamic scene. *arXiv preprint arXiv:2311.14603*, 2023. 2, 20
- [193] Yuyang Zhao, Chung-Ching Lin, Kevin Lin, Zhiwen Yan, Linjie Li, Zhengyuan Yang, Jianfeng Wang, Gim Hee Lee, and Lijuan Wang. Genxd: Generating any 3d and 4d scenes. *arXiv preprint arXiv:2411.02319*, 2024. 3
- [194] Guangcong Zheng, Teng Li, Rui Jiang, Yehao Lu, Tao Wu, and Xi Li. Cami2v: Camera-controlled image-to-video diffusion model. *arXiv preprint arXiv:2410.15957*, 2024. 3
- [195] 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 *Proc. CVPR*, 2024. 2, 20
- [196] Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. Open-sora: Democratizing efficient video production for all, 2024. 3, 17
- [197] Haitao Zhou, Chuang Wang, Rui Nie, Jinxiao Lin, Dongdong Yu, Qian Yu, and Changhu Wang. Trackgo: A flexible and efficient method for controllable video generation. *arXiv preprint arXiv:2408.11475*, 2024. 21

- [198] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *ACM TOG*, 2018. [2](#), [4](#), [5](#), [6](#), [7](#), [8](#), [17](#), [18](#), [19](#), [20](#)
- [199] Hanxin Zhu, Tianyu He, Anni Tang, Junliang Guo, Zhibo Chen, and Jiang Bian. Compositional 3d-aware video generation with llm director. *arXiv preprint arXiv:2409.00558*, 2024. [20](#)

AC3D: Analyzing and Improving 3D Camera Control in Video Diffusion Transformers

Supplementary Material

We encourage the reader to inspect our visual results, comparisons with other models and additional visualizations in the accompanying website in <https://snap-research.github.io/ac3d>.

A. Ethics Statement

As with all generative AI technologies, there is the potential for misuse by bad actors. However, we anticipate this technology will advance creative expression, education, and research through:

1. Enhanced creative tools enabling filmmakers and educators to achieve complex camera movements without specialized equipment, democratizing high-quality video production and expanding possibilities for visual storytelling.
2. More realistic synthetic video that better simulates real world camera behaviors, improving applications in training autonomous systems, virtual production, and educational simulations where accurate camera dynamics are crucial.
3. Advancing our technical understanding of how camera motion affects visual perception and generation, contributing to fundamental research in computer vision, graphics, and human visual processing.

B. Limitations

In our work, we substantially advance the quality of 3D camera control of video diffusion models, but our analysis and method are not free from limitations.

OOD trajectories generalization. Both our model and all the baselines struggle to generalize to the camera trajectories, which are far away from the training distribution of RealEstate10K [198]. While, in general, it is an expected behavior, it indicates that the model processes the viewpoint conditioning information in a way that is entangled with the main video representations. Another source of this issue is the pre-training distribution of the base VDiT itself: natural videos typically have simple recording trajectories and rarely exhibit something that looks like 3D scanning. In this way, producing OOD trajectories is not an attempt to control existing knowledge of a video model, but an attempt to induce new knowledge into the model, which should require better and more diverse fine-tuning data.

Motion analysis limitations. As discussed in Appendix D, we estimate the optical flows in the latent space of CogVideoX [168] autoencoder rather than the pixel space,

because it's the space the video DiT operates in and early trajectory steps produce disarranged decoder's outputs. Besides (as also observed by [75]), motion spectral volumes are sensitive to the quality of optical flow estimation. While the key conclusions (of the camera motion being low-frequency and kicking in very early in the diffusion trajectory) would hold since they are evident even with a bare eye from inspecting the denoising process visualization, the exact behavioral details of motion spectral volumes might change when an optical flow estimator is swapped or the analysis is moved from the latent to pixel space.

Linear probing limitations. We inspect the presence of disentangled camera information in the video DiT model of a particular architecture and trained on particular data. For it, we observe that the knowledge starts to arise from the 9-th block, but for a different instance of a video model it can be distributed across the blocks differently. Besides, we only evaluate it on RealEstate10K [198] test-set trajectories, and do not explore classical camera estimation datasets. Such analysis was sufficient to draw actionable conclusions and improve our base method, but there is vast room in making it more rigorous. We expect that video diffusion models are capable of revolutionizing the field of camera registration bringing strong priors about feature correspondences under difficult conditions, like scene motion, changing lighting and occlusions.

C. Implementation details

This section describes the training and architectural details of the base VDiT, VDiT-CC, and our downstream experiments.

C.1. VDiT implementation details

Architecture details. Our video DiT [100] architecture follows a very similar design to the other contemporary video DiT models (e.g., [10, 32, 102, 168, 196]). As the backbone, we used a transformer-based architecture of 32 DiT blocks [100]. Each DiT block consists on a cross-attention layer to read the text embeddings information, produced by the T5 [108] model; a self-attention layer, and a fully-connected network with a $4\times$ dimensionality expansion. Each attention layer has 32 heads and RMSNorm [182] for queries and keys normalization. To encode positional information, we used 3D RoPE [124] attention, where each axis (temporal, vertical, and horizontal) had a fixed dimensionality allocated for it in each attention head (we split the dimensions in the ratio of 2:1:1 for temporal, vertical,

and horizontal axes, respectively). LayerNorm [1] is used to normalize the activations in each DiT block. We used CogVideoX [168] autoencoder which is a causal 3D convolutional autoencoder with $4 \times 8 \times 8$ compression rate and 16 channels for each latent token. The hidden dimensionality of our DiT model is 4,096 and it has 11.5B parameters in total. Similar to DiT [100], we use block modulations to condition the video backbone on the rectified flow timestep information, SiLU [40] activations and 2×2 ViT-like [26] patchification of the input latents to reduce the sequence size.

Training details. We train the model with the AdamW [89] optimizer with the learning rate of 0.0001 and weight decay of 0.01. The model was trained for 750,000 total iterations with the cosine learning rate scheduler [90] in bfloat16. We also incorporate the support of image animation by encoding the first frame with the same CogVideoX encoder, adding random gaussian noise (with the noise level sampled independently from the video noise levels σ_t), projecting with a separate learnable ViT-like [26] patchification layer, repeating sequence-wise to match the video length and summing with the video tokens. During training, we use loss normalization [60]. The model is trained jointly on images and videos of variable resolution (256, 512 and 1024), aspect ratio (16 : 9 and 9 : 16 for videos, and 16 : 9, 9 : 16 and 1 : 1 for images), and video lengths (from 17 frames to 385 frames). The video framerate was set to 24 frames per second and we did not use variable-FPS training as contemporary works [93, 102] since we found it to decrease the performance for a target framerate (at least, without fine-tuning).

Inference details. During inference, we use the standard rectified flow without any stochasticity. We found that 40 steps gives a good trade-off between quality and sampling speed. We follows the same time shifting strategy as Lumina-T2X for higher resolutions and longer video generation [32] with time shifting of $\sqrt{32}$ for the 1024 resolution.

C.2. VDiT-CC implementation details

As being said in Section 3.2, VDiT-CC is a simple ControlNet-like [71, 187] fine-tuning of VDiT for camera control. We use smaller versions of the base VDiT blocks with only a 128 hidden dimensionality and 4 attention heads. Besides, we do not use cross-attention over the context information since we found it to severely decrease the visual quality and camera control precision. For the Plücker encoding computation, we replicate the pipeline of VD3D [6]. Our linear layer that processes them contains 4096 hidden features and SiLU [40] non-linearity.

C.3. AC3D implementation details

We train our complete setup for 6K iterations on the joint dataset of 65K videos from RealEstate10K [198] and 20K




Figure 7. Comparing rectified flow noise schedules: (orange) vanilla standard logit-normal noise schedule proposed by [28] and used for baseline experiments; (purple) biased but non-truncated noise schedule; (pink) biased and truncated noise schedule.

dynamic videos with static cameras (this dataset is described in Section 3.5 and Appendix F). We train with the learning rate of 0.0001 using the AdamW [89] optimizer with weight decay of 0.01 and cosine learning schedule [90].

As described in Section 3.2, since the camera motion is a low-frequency type of signal, we propose to use truncated and biased noise schedule: both at train and inference time. In Figure 7, we visualize three distributions: 1) the standard one used by SD3 [28], our base VDiT, and “w/o biasing noise” experiment (orange); 2) the biased but non-truncated schedule used in the VDiT-CC “w/o noise truncation” ablation experiment (purple), which has some unnecessary and detrimental probability mass in the high-frequency range; and 3) our final schedule (red).

Each ablation experiment from Section 4.3 was trained for 6K iterations on 32 NVidia A100 80GB GPUs.

During training, to plug in the conditioning camera pose for our static dataset, we were randomly sampling extrinsics and intrinsics parameters from a RealEstate10K dataset.

D. Motion analysis details

As discussed in Section 3.3, we perform camera motion analysis of generated videos at different time steps by inspecting their spectral volumes.

For our analysis, we generate 200 random 121-frames videos without time shifting [32] in the 288×512 resolution with VDiT and manually annotate them for the following labels: quality (a score from 1 to 5), scene motion strength (a score from 1 to 5, where a score of 1 corresponds to a completely still scene), camera motion strength (a score from 1 to 5, where a score of 1 corresponds to a completely




Figure 8. Our annotations collected for 200 randomly generated videos from VDiT and used in our camera motion analysis in Section 3.3.

stationary viewpoint), whether the camera is smooth or shaky (a binary flag). We visualize the obtained scores in Fig. 8. Next we discard the videos with the quality score of less than 4, discarding 18 out of 200 videos: “broken” samples (e.g., an artificial animation or a blank black canvas) indicate a complete generation failure which we should exclude from the analysis of the camera motion.

To analyze the spectral volumes differences between scene and camera motion for Fig. 3, we extract three categories of videos from our dataset: 1) scene motion (videos with scene motion, but with no camera motion); 2) camera motion (videos with camera motion, but with no scene motion); and 3) scene and camera motion (videos with scene motion and non-artifactory camera motion). The first category (scene motion) was obtained by selecting the videos with the camera motion strength of 1 (i.e., no camera movement), and the scene motion strength of more or equal than 3; the second category (camera motion) was obtained by selecting the videos with the camera motion strength of more or equal than 3, and the scene motion strength of 1 (i.e., no scene motion); the third category (scene and camera motion) was obtained by selecting the videos with the camera and scene motion strengths of more or equal than 3 and the smooth camera flag being true (to exclude shaky camera movements). To analyze the spectral volumes for Fig. 4, we took the videos which have scene or camera motion strength of more or equal than 3.

To obtain spectral volumes, we need to obtain per-pixel optical flow information. Since our VDiT is following the latent diffusion (LDM) paradigm [111, 112], we opt for performing optical flow estimation in the latent space of the autoencoder. There are two reasons for that: 1) we noticed that it provides more robust flow estimation at earlier denoising timesteps (since the decoder part of CogVideoX [168] autoencoder does not need to operate at out-of-distribution



Figure 9. Frames of the generated videos by the VDiT model (upper row) and the corresponding PCA projections of their latents (lower row).

inputs); and 2) the video model operates in this space. In this way, we used the raw generated latents to estimate the optical flow. Since most of the optical flow algorithms operate in a 3-dimensional RGB space, and our latents are 16-channels, we projected them into 3-dimensional inputs via a PCA, computing it independently for each latent. We found that these representations maintain very strong spatiotemporal resemblance to the original videos, as visualized in Fig. 9.

Following [75], we use PyFlow [99] coarse-to-fine optical flow estimation to obtain more robust results. We attempted to use Farneback [29] optical flow estimation, but observed that it is less accurate and does align less with our visual evaluation of the results. 121-frames 288×512 -resolution videos correspond to 31 latent frames of 36×64 resolution. For each 6-th latent frame, we estimate its flow with respect to each of the 24 subsequent frames. Next, we perform Fast Fourier transform for x and y spatial coordinates independently, compute the amplitudes and average them spatially for each video.

E. Linear probing details

In Section 3.4, we perform linear probing of VDiT for camera pose knowledge. For this, we use 1,000 videos of 49 frames and the 144×256 resolution from the test set of RealEstate10K [198] and extract their internal representations of our VDiT model under various noise levels. We use the noise levels σ_t of $[\frac{1}{8}, \frac{2}{8}, \frac{3}{8}, \dots, \frac{7}{8}, 1]$. The hidden dimensionality of our VDiT model is 4,096 and to reduce the memory requirements and speed up linear probing we project each video representations into 512 dimensionality using PCA. This results in latent representations of $512 \times 13 \times 18 \times 32$ for each block, each video and each timestep. To construct the training features for the linear regression model, we extract the (spatially) middle vector of shape 512×13 and perform spatial average pooling to obtain a context representation of 512×13 . We then unroll and concatenate them to obtain the final training representation of dimensionality 13,312. We split our 1,000 videos into training and test sets as 900 and 100 and then train a ridge linear regression with the regularization weight of 25,000. Our target variable covers the extrinsics parameters of the viewpoints, which are provided by RE10K. We extract rotation angles and translations from the extrinsic matrices and normalize them with respect to the

first frame. Then, we compute the rotation and translation errors on the held-out set of 100 videos using the evaluation pipeline of CameraCtrl [38]. Since we have 8 noise values and 32 VDiT blocks, this resulted into 256 linear regression models in total. It was taking ≈ 5 minutes of CPU time to train each model, and their training was parallelizable across different cores.

F. Dataset construction details

As describe in Section 3.5, we construct a dataset of videos with scene motion but recorded from stationary cameras. One might attempt to build such a dataset in an automated way by estimating the optical flow and checking whether there is non-zero motion in the middle region of the frame, and no motion on the borders. However, this would not work well for a myriad of corner cases, which is why we opted for manual construction. For this, we annotated 110K internal videos for the presence of camera and the presence of scene motion. Each video has a duration of 5–30 seconds and covers various diverse categories of scenes: humans, animals, landscapes, food and other types. Out of these videos, 20K videos turned out to be the necessary ones: with scene motion, but completely static scenes. We proceeded to use them as our training data without further processing. As being discussed in Appendix C, we augmented their camera information artificially by using random extrinsics from RealEstate10K [198].

G. Scaling bias

Classical feature-based camera estimation [115, 116] outputs camera trajectories with arbitrary scale, as it does not possess any priors about the absolute size of objects in the scene. For example, a house might appear identical to a small object like an apple when viewed from the appropriate distance, making it impossible to determine absolute scale from visual features alone. This lack of scale awareness complicates precise user control over camera trajectories: while the model can determine the camera’s direction, it remains unaware of the magnitude of each movement, as demonstrated in Sec. 4. Ideally, all camera trajectories should be aligned to a consistent reference scale. A natural reference would be a metric scale, now achievable due to recent advances in metric depth estimation [171]. We propose a method to rescale camera trajectories across all data using the following steps:

1. Obtain 3D points from the COLMAP output and render the COLMAP depth D_c^f from these points for each frame.
2. Estimate the metric depth D_m^f for each frame using a pre-trained zero-shot metric depth estimator [171].
3. Calculate the re-scaling parameter $\hat{\lambda}$ by solving the optimization problem $\hat{\lambda} = \arg \min_{\lambda} \mathbb{E}_{f \sim F} |\lambda D_c^f - D_m^f|$, where F is the total number of frames.
4. Set the camera translation vector \hat{t}_c to $\hat{\lambda} t_c$.

In Sec. 4.3, we show that training with properly scaled cameras does not lead to visual quality degradation (even improving it slightly), and, as we demonstrate in our supplementary visuals, makes the camera control more predictable and less frustrating for a user by allowing to control the magnitudes of camera transitions.

H. Additional Related Work

Due to space constraints, we summarize related 3D and an extended list of 4D works in the appendix.

3D generation. Early efforts in 3D generation focused on training models for single object categories, extending GANs to 3D by incorporating neural renderers as an inductive bias [3, 13, 24, 97, 117]. As the field progressed, CLIP-based supervision [107] enabled more flexible and diverse 3D asset generation, supporting both text-based generation and editing [19, 34, 53, 55, 114, 143]. Recent advances in diffusion models further enhance generation quality by replacing CLIP with Score Distillation Sampling (SDS) for supervision [20, 39, 61, 66, 74, 77, 78, 86, 103, 125, 144, 149, 169, 177]. To improve the structural coherence of 3D scenes, several approaches generate multiple views of a scene for consistency [30, 47, 63, 80, 83, 84, 118, 134, 141]. Alternatively, iterative inpainting has been explored as a technique for scene generation [46, 119]. Recent works also focus on lifting 2D images to 3D representations, employing methods like NeRF [95], 3D Gaussian Splatting [62], or meshes in combination with diffusion models [14, 35, 87, 88, 105, 129, 132, 135, 142, 173]. Other studies explore fast, feed-forward 3D generation techniques that directly predict 3D models from input images or text [37, 48, 56, 70, 104, 130, 131, 133, 136, 156, 162, 163, 183]. These methods, however, are limited to synthesizing static scenes, in contrast to our approach.

4D generation. There has been significant progress in 4D generation, i.e., dynamic 3D scene generation. These works often rely on input text prompts or images to guide the generation. Since the early advancements in large-scale generative models for this task [121], significant strides have been made in improving both the visual and motion quality of generated scenes [5, 33, 57, 58, 69, 81, 94, 109, 165, 178, 189, 195]. While many of these methods are conditioned on text input, other approaches focus on converting 2D images or videos into dynamic 3D scenes [12, 23, 31, 33, 67, 72, 73, 76, 81, 98, 109, 110, 127, 138, 139, 148, 154, 157, 167, 172, 181, 186, 192, 195]. Recently, several works [52, 79, 188] investigate physics priors in 4D generation pipelines. Other works [15, 126, 185] enhance motion controllability with template-based methods. Another line of work [4, 11, 159, 174, 180, 199] focuses on compositional and interactive 4D generation. Another strand of research extends 3D GANs into the 4D domain by training on 2D video data [2, 164]. However, the quality

of these methods is often constrained by the limited nature of the datasets, which typically focus on single object categories. Moreover, the majority of these approaches tackle object-centric generation. As a result, they typically neglect background elements, and their visual fidelity falls short when compared to the high photorealism achieved by state-of-the-art video generation models, such as those employed in our approach.

Motion-controlled video generation. Orthogonal to camera-controlled video generation methods, another line of work investigates object trajectory control [145, 147, 153, 170, 190, 197]. More recently, several methods [54, 91, 96, 106, 166, 175] focus on object trajectory control without relying on additional external data or additional model fine-tuning.