

Control-DINO: Feature Space Conditioning for Controllable Image-to-Video Diffusion

Edoardo A. Dominici^{1*}, Thomas Deixelberger^{2,3*}, Konstantinos Vardis¹, and Markus Steinberger^{2,3}

¹ Huawei, Switzerland

² Huawei, Austria

³ Graz University of Technology, Austria

<https://dedoardo.github.io/projects/Control-DINO>




Fig. 1: We introduce a DINO based ControlNet for I2V diffusion models and apply it to video-to-video and video-from-3D tasks. It allows using dense DINOv3 features for structural and semantic guidance, in the image we show by row: (1) Style transfer on first frame propagated to later timesteps; (2) video-from-3D obtained by using only the voxel as conditioning with no first frame and with first frame conditioning; (3) VKITTI weather transfer where the first frame is obtained with FLUX.1 KREA [21].

Abstract. Video diffusion models have recently been applied with success to problems in content generation, novel view synthesis, and, more broadly, world simulation. Many applications in generation and transfer rely on conditioning these models, typically through perceptual, geometric, or simple semantic signals, fundamentally using them as generative renderers. At the same time, high-dimensional features obtained from large-scale self-supervised learning on images or point clouds are increasingly used as a general-purpose interface for vision models. The connection between the two has been explored for subject specific editing, aligning and training video diffusion models, but not in the role of a

* The authors contributed equally to this work.

dense conditioning signal for pretrained video diffusion models. Features obtained through self-supervised learning like DINOv3, contain a lot of entangled information about style, lighting and semantics of the scene. This makes them great at reconstruction tasks but limits their generative capabilities. In this paper, we show how we can use the features for tasks such as video domain transfer and video-from-3D generation. We introduce a lightweight control architecture and training strategy that decouples appearance from other features that we wish to preserve, enabling robust control for appearance changes such as stylization and relighting. Furthermore, we show that low spatial resolution can be compensated by higher feature dimensionality, improving controllability in generative rendering from explicit spatial representations.

1 Introduction

Generative video models have advanced rapidly and now serve as strong foundations for editing, stylization, simulation, and world modeling [27, 30, 40, 52]. In many practical scenarios, we are not generating from scratch: we already have partial information about the scene, such as a reference video, sparse views, or 3D geometry, and we want generation to preserve this structure while still allowing controlled variation.

Most current control pipelines rely on low-level geometric or perceptual inputs, including edges, depth maps, colored points, pose skeletons and semantic masks [26, 27]. These controls are often effective, but each captures only a narrow aspect of scene structure, and different tasks tend to require diverse control types. This fragmentation makes it difficult to build a unified conditioning interface that transfers cleanly across tasks like video restyling, video-from-3D, and long-horizon view synthesis.

At the same time, universal self-supervised features such as DINO [4] are increasingly used as a common interface across vision models. Trained on large scale 2D and 3D data, these high-dimensional features encode semantics, geometry, and appearance while having some invariance to lighting thanks to the pre-training augmentations. They achieve strong performance in downstream tasks and have also shown promise as latent spaces for prediction [1, 52] or completion [15]. This makes them an attractive candidate for conditioning video diffusion models with richer and more reusable control signals than traditional low-level inputs. However, feature conditioning is not straightforward. The same expressivity that makes foundation features useful also makes them difficult to tame: appearance, geometry and other properties are often entangled, making the process challenging to generalize, especially when training data is limited (Fig. 2). Recent work has shown promising subject-specific video transfer with DINO features [13], but robust general-purpose conditioning remains underexplored.

We study DINO-based conditioning for video diffusion in a unified setting that spans both 2D and 3D feature sources. To extract structural guidance, the key idea is to use features of the original video as conditioning but train the

diffusion model to denoise appearance-augmented versions of the same scene. Because the same features correspond to multiple appearances, the model learns to rely on spatial and semantic information rather than reproducing appearance cues. We further observe that low-spatial-resolution but high-dimensional features can provide robust control and similarly enable us to bridge 3D mesh renderings or other spatial representations to real videos.

Our contributions are as follows:

- We introduce a general-purpose DINO-conditioned ControlNet for video diffusion that retains structural and semantic guidance while decoupling appearance from the conditioning signal.
- We apply the ControlNet to video transfer and video-from-3D tasks showing its versatility and robustness.
- We demonstrate new applications such as rendering low-resolution 3D voxel structures and connecting foundational 3D models to video diffusion.

2 Related Work

Controllable diffusion via spatial conditioning. Modern generative video models [2, 12] perform denoising in the latent space of a variational autoencoder, either with UNets or more recently diffusion transformers [33, 41, 44]. Within this paradigm, controllable generation is commonly achieved by injecting auxiliary signals into the denoising network interleaved between the transformer blocks. ControlNet [49] introduced a widely adopted design in which a trainable copy of the encoder injects residuals into the frozen generator via zero-initialized convolutions. Lighter-weight variants such as T2I-Adapter [28] and IP-Adapter [45] offer more parameter-efficient or reference-based alternatives. Several works extend spatial conditioning to video by adapting image-level control architectures to temporal backbones [5, 11, 16, 22, 23], while alternative designs such as FullDiT [17] tokenize all conditions for joint self-attention at the cost of full retraining. To date, these methods have predominantly been applied with geometric signals (edges, depth, pose), yet foundation model features remain underexplored as conditioning inputs. Recent methods that generate videos from 3D representations [18] also rely on edge conditioning for structure transfer. For single image domain transfer, a discriminator can also be adopted to finetune the model and decoder to generate samples closer to the target domain [32]. For a broader overview of controllable video generation, see [26].

Self-supervised features for generation. Self-supervised vision models such as DINOv2 [31] learn patch-level representations that capture object semantics and support tasks from segmentation to depth estimation, all from a single frozen backbone. DINOv3 [35] scales this family to denser features that preserve fine-grained local structure across model sizes. Multiple studies observe that self-supervised and diffusion features capture complementary information: the former provide high-level semantic structure, while the latter encode dense spatial



Fig. 2: Naively training a ControlNet on top of an Image-to-Video model leaves DINO features prone to overfitting. We isolate the conditioning effect by dropping the first-frame conditioning, relying solely on the text embedding. With the prompt set to "blue", our method significantly reduces the ControlNet’s bias toward the training data.

coherence [25, 37, 39, 48]. This motivates conditioning diffusion models on foundation features to supply semantic guidance beyond what the generation process captures on its own. REPA [47] further shows that frozen DINOv2 features align with internal DiT representations, though as a training regularizer rather than an explicit conditioning signal. Directly conditioning on such features was first explored by Bordes et al. [3] using global DINO embeddings via FiLM [34] modulation. Graikos et al. [10] scale this to local, domain-specific SSL embeddings with cross-attention for large-image generation in medical and remote-sensing domains. AnyDoor [6] injects DINOv2 patch tokens into a pretrained backbone for object identity, yet like the preceding works, targets images and not scene-level spatial control. Self-supervised representations have also been extended to 3D, with methods such as Concerto [51] and Sonata [43] learning dense, multi-view consistent descriptors from point clouds that remain close to DINO-style embeddings. We leverage this property by rendering such 3D features into 2D conditioning maps for our video diffusion adapter.

Foundation features for video. Recent work has begun to explore foundation features as conditioning signals for video generation and editing. CAGE [9] introduces a video model trained jointly with sparse DINOv2 patches to enable conditioning for unsupervised scene composition and object animation on synthetic scenes. DIVE [13] leverages DINOv2 features as implicit motion correspondences for subject-driven video editing through per-video MLP optimization, targeting appearance transfer rather than scene-level spatial control. Concurrent to our work, Driving with DINO (DwD) [7] injects DINOv3 features through a ControlNet-style branch into a frozen video backbone for sim-to-real driving video translation, introducing PCA-based spectral pruning tailored to its synthetic-to-real setting. However, none of these works targets dense foundation feature conditioning for different types of video transfer and rendering operations. We explore this direction with a ControlNet-style adapter that operates on spatially dense features without per-video optimization, and evaluate it across video transfer and 3D-guided video generation.




Fig. 3: (Left) During training, DINOv3 features from the original video condition the trainable Control-DINO branch, while the frozen backbone denoises appearance-augmented latents of the same scene. (Right) At inference, conditioning features (2D or 3D-rendered) guide generation through Control-DINO. A transferred first frame sets the target appearance.

3 Method

The goal of our model is to condition video diffusion on dense features while making their control practical (Fig. 3). We achieve this by injecting foundation features into a frozen video transformer backbone through a ControlNet-style residual adapter (Sec. 3.1). We train on paired augmented data where the conditioning features come from the original video but the denoising target is a structurally similar version of the same scene, enabling the adapter to rely on geometry and semantics alone (Sec. 3.2). A property of feature-level conditioning is that it operates at low spatial resolution but high channel dimensionality, which means the same interface naturally extends to coarse 3D representations where pixel-level detail is unavailable. Since conditioned video models are increasingly used for video-from-3D generation [18], we describe how features can be obtained from 3D data and used as conditioning inputs to our adapter (Sec. 3.3).

3.1 Diffusion and Conditioning

Let $x_{1:T}$ denote an RGB video and $z_{1:T}$ its latent representation under the video VAE encoder. We follow the standard diffusion formulation:

$$z_t = \alpha_t z_0 + \sigma_t \epsilon, \quad \epsilon \sim \mathcal{N}(0, I), \quad (1)$$

where t denotes the diffusion timestep and subscript ranges $1:T$ index video frames.

The denoiser v_θ , a video transformer, is trained to predict the velocity $v_t = \alpha_t \epsilon - \sigma_t z_0$ by minimizing $\mathcal{L} = \mathbb{E}_{z_0, \epsilon, t} [\|v_t - v_\theta(z_t, t, c)\|_2^2]$, where c denotes conditioning. We consider per-frame conditioning signals $c_{1:T}$ (edges, depth, semantic maps, or feature maps) pixel aligned with the output frames. Conditioning signals that are encoded with the video VAE (e.g., depth or edges) match the latent-space temporal resolution directly, while signals that bypass the VAE are

rescaled spatially and temporally. With the spatial resolution adapted, we use a temporal adapter with the same causal architecture as the video encoder to map the conditioning temporal length.

We condition the video diffusion model by injecting residual connections from the output of the transformer block into its first 16 layers. This stream concatenates the noisy latent with the conditioning features along the feature dimension and processes them with a copy of the transformer block($\mathcal{Z}$) at each layer $\mathcal{Z}_\ell^{light}$, yielding the conditioning hidden states h_ℓ^c . The projections $\mathcal{Z}_\ell^{light}$ are zero initialized. We optionally modulate each injection by a mask $M \in [0, 1]^{60 \times 90}$, broadcast over the feature dimension, where $M = 1$ recovers the unmasked case; M can encode, for example, per-pixel visibility obtained by rendering geometry.

$$h_{\ell+1} = \mathcal{Z}_\ell(h_\ell) + M \odot \mathcal{Z}_\ell^{light}(h_\ell \cdot c), \quad \ell = 1, \dots, L. \quad (2)$$

Here $h_{\ell+1}$ is the current hidden state being denoised.

3.2 Feature Conditioning

A side effect of DINO features encoding rich appearance information is that structural signals can be entangled with color, texture, and style. We therefore aim for the ControlNet to faithfully follow the spatial structure depicted in the conditioning features while remaining agnostic to the visual style of the source, enabling domain transfer applications.

We extract dense patch-level features from a frozen DINOv3 vision encoder D applied to each frame of the source video. For $x_{1:T} \in \mathbb{R}^{T \times 3 \times H \times W}$ we obtain per-frame feature maps

$$F_i = D(x_i) \in \mathbb{R}^{384 \times 60 \times 90}, \quad i = 1, \dots, T, \quad (3)$$

for patch size 16. To align conditioning with the video diffusion latent grid, we bicubic-upsample each input frame before DINO encoding so that D operates at the transformer latent spatial scale, which coincides with the latent resolution of the video model’s VAE. These feature maps serve as the conditioning signal $c = A(F_{1:T})$ for the ControlNet adapter described in Sec. 3.1, where A is the temporal adapter.

Training with augmentation. To force style-invariant conditioning, we construct training pairs $(F_{1:T}, \tilde{z}_{1:T})$, where $F_{1:T}$ are DINO features from the original video and $\tilde{z}_{1:T}$ are VAE latents of an appearance-augmented version of the same scene. We train with $\mathcal{L} = \mathbb{E}_{\epsilon, t} [\|v_t - v_\theta(z_t, t, F_{1:T})\|_2^2]$, where z_t is obtained from $\tilde{z}_{1:T}$ using the forward diffusion process in Sec. 3.1. During training we employ a hierarchy of augmentation strategies at varying levels of severity to mimic appearance changes expected at test time: (i) photometric augmentations (hue, grayscale, saturation, brightness, contrast, gamma, etc.) applied uniformly across frames; (ii) five GAN-based neural style transfer methods (cartoon, anime, etc.) [8, 38]; and (iii) small blur to simulate quality degradations. During training, we sample from a mixture uniformly distributed among real data and these

three augmentation groups, all paired with features extracted from the original unstylized video. At inference, we apply a more generic transfer function only to the first frame while conditioning on features from the original source domain (real or rendered) [42, 50]. As long as the first-frame transfer preserves structure and semantics, the model can follow the target appearance consistently over time. Because the same features correspond to multiple appearances, the model learns to rely on spatial and semantic information rather than appearance cues in the features. We compare this strategy against PCA-based feature projection in Sec. 4.6.

3.3 3D Feature Conditioning

Although DINO features are extracted on coarse patch grids, each token is high-dimensional and semantically rich, such that an $h \times w$ grid with $C=384$ channels (here 60×90) preserves object identity, part-level layout, and scene semantics even when fine pixel detail is absent. This makes the features complementary to geometry controls: geometry anchors *where* content appears, while DINO channels encode *what* it is and how regions relate semantically. We therefore align both signals at the same low spatial resolution and condition on them jointly. In our method the geometric prior is a voxel grid, which is stable across views and time: voxel geometry enforces camera-consistent structure while DINO features stabilize subject and background semantics. Because both inputs are low-resolution and high-level, conditioning remains inexpensive while still guiding temporally coherent generation. High-resolution geometry is not required, and low-resolution voxel cues from LiDAR reconstructions or occupancy prediction models suffice.

DINO-like features from uncolored geometry. Importantly, useful feature priors do not require RGB textures at inference time. Recent 2D-3D self-supervised methods, such as Concerto [51], can produce dense geometric descriptors from uncolored geometry like point clouds that remain close to DINO-style semantic embeddings after projection. We therefore treat these geometry-derived descriptors as a drop-in conditioning signal: they can be rendered to the target view, projected to the ControlNet token grid, and fused with standard controls exactly as in Fig. 5. This enables controllable generation in settings where only geometry is available, while retaining semantic consistency typically associated with appearance-based foundation features.

4 Experiments

4.1 Datasets

Our work combines video diffusion conditioning from both 2D and 3D representations, with a focus on generative rendering. To emphasize generation quality and cross-dataset generalization, we train on DL3DV [24] and evaluate primarily



Fig. 4: We show here results for DINO stylization and dramatic light changes (GT Frame is in the bottom left of the image). Our method is able to strongly retain the original geometry and semantics while allowing for appearance and color to change. For brevity, *Canny* refers to Wan 2.2 Fun Canny and *Depth* to Wan 2.2 Fun Depth.




Fig. 5: 3D structure visualization. Top: RGB voxels. Bottom: mesh renderings.

on Tanks and Temples (T&T) [19], which provides longer and smoother camera trajectories. From these trajectories we sample two evaluation settings, stride 8 and stride 16. For Table 3, we also report results on DL3DV140 [24]. In total, we train on about 4.5M unique frames and evaluate on 10k unique frames. The larger camera motion in Tanks and Temples stresses both generative quality and conditioning robustness. For 3D rendering experiments, we fine-tune the same base model on ScanNet++ [46] and evaluate on parts of its official validation split. On ScanNet++, we sample training clips at stride 1, yielding nearly 8k clips, and evaluate video-from-3D pipelines that construct DINO and DINO-like conditioning features from explicit 3D geometry (Table 4).

4.2 Model Details

We implement our approach on CogVideoX-5B-I2V [44] with residual connections [49] into the first $L=16$ transformer blocks, using dense DINOv3-ViT-S/16 patch features as input. These features are bicubically upsampled by a factor of 2 and then encoded. To align the latent space dimension to the one of the hidden

Table 1: Quantitative comparison on T&T Appearance Augmented. All models are trained on DL3DV and evaluated on T&T (for two frame sampling strides) with a combination of 15 lighting or stylistic variations per scene different from the training data. Overall our method produces the most 3D consistent videos, especially for longer camera trajectories, while maintaining a competitive CLIP similarity with the first frame in the sequence. Our same model but trained only on real data performs much worse in style following.

Tanks and Temples - Out of domain - Transfer							
Method	Consistency		Video Quality (VBench)			CLIP Similarity	COLMAP
	Subject $\uparrow$	Background $\uparrow$	Motion Smoothness $\uparrow$	Aesthetic $\uparrow$	Imaging $\uparrow$	CLIP Sim. $\uparrow$	Registration $\uparrow$
Stride = 8							
Wan 2.2 Fun Depth [41]	0.8853	0.8916	0.9421	0.5388	0.6880	0.4626	0.8430
Wan 2.2 Fun Canny [41]	0.8778	0.8848	0.9599	0.5325	0.7568	0.5020	0.9770
AnyV2V [20]	0.8124	0.8818	0.9258	0.5390	0.6590	0.4875	0.5470
Control-DINO Real Only	0.8767	0.8465	0.9680	0.5533	0.7372	0.4408	0.8840
Control-DINO	0.8913	0.9060	0.9640	0.5610	0.7146	0.4717	0.9530
Stride = 16							
Wan 2.2 Fun Depth [41]	0.8222	0.8527	0.9410	0.5044	0.6497	0.4804	0.4956
Wan 2.2 Fun Canny [41]	0.8226	0.8512	0.9357	0.5104	0.7235	0.4630	0.5918
AnyV2V [20]	0.7621	0.8579	0.9162	0.5072	0.6299	0.4864	0.1875
Control-DINO Real Only	0.8544	0.8906	0.9396	0.5485	0.6908	0.4406	0.9911
Control-DINO	0.8461	0.8862	0.9361	0.5508	0.7109	0.4707	0.9932

Table 2: We compare against the only publicly available segmentation ControlNet. Our method produces competitive results, especially under larger camera motions, even against models larger in scale, training data, and architecture.

Tanks and Temples - Out of domain - Transfer							
Method	Consistency		Video Quality (VBench)			CLIP Similarity	Coverage
	Subject $\uparrow$	Background $\uparrow$	Motion Smoothness $\uparrow$	Aesthetic $\uparrow$	Imaging $\uparrow$	CLIP Sim. $\uparrow$	RegRate $\uparrow$
Cosmos-Segmentation [29] (S=8)	0.9128	0.9191	0.9678	0.4987	0.6666	0.7306	0.9374
Control-DINO (S=8)	0.8913	0.9060	0.9640	0.5610	0.7146	0.4717	0.9530
Cosmos-Segmentation [29] (S=16)	0.8315	0.8958	0.9475	0.4547	0.6190	0.7147	0.7755
Control-DINO (S=16)	0.8461	0.8862	0.9361	0.5508	0.7109	0.4707	0.9932

Table 3: We illustrate the tradeoff between structural control signals and generation determinism, by analyzing how different seeds affect generation. Features jointly encode geometry and appearance, constraining the output more than other signals.

DL3DV140 - In domain - Real							
Method	Reconstruction			Video Quality (VBench)			
	FID $\downarrow$	Latent Var.	CLIP Var.	mIOU $\uparrow$	Flickering $\uparrow$	Aesthetic $\uparrow$	Imaging $\uparrow$
Wan 2.2 Depth [41]	47.98	0.5115	0.0151	0.0362	0.9192	0.4343	0.5469
Wan 2.2 Canny [41]	37.65	0.4787	0.0150	0.0322	0.9154	0.4496	0.6169
Control-DINO	20.23	0.1944	0.0051	0.2992	0.9208	0.5133	0.6283

state. The temporal adapter A is composed of two causal Conv3D layers that temporally compress the sequence from 49 to 13 frames, matching the VAE latent temporal resolution (and architecture) and the result is concatenated channel-wise with the VAE latents before being fed into the conditioning branch. We set

Table 4: Video-from-3D comparison on ScanNet++. We evaluate geometry-driven 3D-derived conditioning methods. We compare quantitative results for image-space baselines (Wan and 2D Control-DINO) and report results for 3D-derived methods.

ScanNet++ - In domain - Video-from-3D										
Method	Reconstruction				Video Quality (VBench)					
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FID $\downarrow$	Subject $\uparrow$	Background $\uparrow$	Flickering $\uparrow$	Motion Smooth $\uparrow$	Aesthetic $\uparrow$	Imaging $\uparrow$
Image-space baselines										
Wan 2.2 Depth [41]	18.23	0.7237	0.2782	92.1072	0.9027	0.9234	0.9680	0.9881	0.4111	0.5714
Wan 2.2 Canny [41]	19.82	0.7968	0.1987	72.1929	0.9001	0.9251	0.9590	0.9894	0.4100	0.6625
Control-DINO from 2D (reference)	27.93	0.9084	0.0963	38.5853	0.9183	0.9256	0.9662	0.9931	0.4187	0.5821
3D-derived methods (ours)										
Control-DINO from Mesh	21.41	0.8010	0.1807	67.1730	0.9145	0.9242	0.9647	0.9919	0.4198	0.5852
Control-DINO from Voxel	22.80	0.8153	0.1572	65.3810	0.9216	0.9293	0.9658	0.9928	0.4188	0.5979
Control-DINO Concerto	21.19	0.7805	0.1752	67.9680	0.9195	0.9263	0.9638	0.9926	0.4200	0.6120
Control-DINO Concerto (inpainted)	19.66	0.7542	0.2060	78.8411	0.9122	0.9267	0.9635	0.9917	0.4184	0.6326

the hidden dimension of the conditioning branch to 256 to further reduce the risk of overfitting. Zero-initialized output projections inject additive residuals at each layer of the base CogVideoX transformer, ensuring stable training from initialization. At inference time, we scale the conditioning residuals by 0.8.

We train the temporal adapter and lightweight transformer blocks from scratch with AdamW ($\beta_1=0.9$, $\beta_2=0.95$) at a peak learning rate of 2×10^{-4} with cosine annealing and 500 warmup steps, a batch size of 8, and gradient clipping at 1.0. We apply conditioning dropout during training, setting the conditioning input to zero with probability 0.1. In the supplemental we show results for different video models.

4.3 Metrics

For the main benchmark on out-of-domain evaluation (Table 1) we assess video quality using VBench [14], reporting consistency, motion and imaging scores but also measure 3D consistency through COLMAP-based reconstruction errors and registration success rate. We also quantify style transfer fidelity via CLIP style similarity with the first frame image. On the validation split of the training set (Table 3) we additionally evaluate model generalizability by looking at the distribution of the output latent space and features. For video-from-3D, we additionally report standard reconstruction metrics (Table 4).

4.4 Baselines

Our main baselines use alternative conditioning signals on image-to-video diffusion models with similar architecture and scale. We use the official Wan2.2-Fun-5B-Control (I2V) checkpoints with Canny-edge and depth conditioning [41]. We also compare against the publicly available semantic ControlNet from Cosmos-Transfer-2.5 [29], while noting that Cosmos base model and ControlNet are substantially larger models. For domain transfer, our setup is closest to AnyV2V [20], which uses DDIM-inverted latents [36] for appearance transfer during diffusion. For first-frame appearance transfer baselines, we use IC-Light for relighting [50],

Table 5: We measure feature disentanglement of different projection strategies, measuring cosine similarity as agreement between real and stylized scene features after projection, and retained feature variance. The metrics are computed from samples of the training dataset.

Projection basis disentanglement		
Method	Cosine Similarity $\uparrow$	Explained Variance (%) $\uparrow$
No projection (384-d)	0.9194	100.0
Sim2Real reference (DwD [7])	0.9598	–
Standard PCA	0.9110	72.5
Style-Invariant PCA	0.9481	30.4
Bottom Eigenvectors	0.9929	2.9
Random Orthogonal	0.9226	16.7

Table 6: We ablate our architectural design choices on DL3DV140. We observe that retaining the raw high-dimensional features outperform PCA-projected or tail-dropped variants, and that removing the spatial mixing step substantially degrades quality.

Architectural ablation on DL3DV140.					
Method	FID $\downarrow$	Subj. Cons. $\uparrow$	BG Cons. $\uparrow$	Aesthetic $\uparrow$	Imaging $\uparrow$
Control-DINO	71.67	0.959	0.955	0.504	0.542
w/o Spatial Mixing	177.80	0.919	0.933	0.474	0.495
w/ Style-Invariant PCA	89.14	0.940	0.940	0.479	0.524
w/ Tail Drop (k=64)	108.87	0.918	0.920	0.459	0.494
w/ Tail Drop (k=8)	156.67	0.845	0.908	0.468	0.490

InstantStyle for style transfer [42] picking 10 styles from the default ones, and five additional cartoon neural styles that were excluded from training.

4.5 Experiments and Results

Representation and transfer We evaluate representation-level conditioning for long-horizon transfer on the out-of-domain Tanks and Temples benchmark. Quantitative results in Table 1 and Table 2 show that DINO-based conditioning provides strong and stable guidance under large camera motion, with improved temporal and structural consistency across both settings. As also illustrated in Fig. 9, our method has lower generative flexibility than weaker controls, but this is compensated by substantially higher output quality and consistency. Camera control is stronger because the model can offload part of the generative burden to the base diffusion prior, allowing the diffusion pathway to focus on finer structural details (Fig. 4).

Our appearance-decoupled training makes this representation practical for transfer: the same conditioning interface supports RGB-to-stylized transfer and geometry-driven transfer (e.g., mesh-to-real), as discussed in Sec. 3.3. Qualitative examples are shown in Fig. 7 and Fig. 9, with additional results in Fig. 8. We observe particularly strong relighting behavior, consistent with the built-in invariances of DINO features and the pixel-aligned geometric constraints of the task.

Video-from-3D We evaluate three 3D-to-video conditioning pipelines that all use a ControlNet branch driven by DINO-like feature maps. For the **mesh** setting, we render untextured meshes from target camera trajectories, upscale the rendered views, and extract DINOv3-S features as conditioning maps with shape $384 \times 60 \times 90$. Because mesh renderings contain holes and missing regions, we additionally encode a binary visibility mask with the CogVideoX VAE and provide it as an auxiliary conditioning signal, so the diffusion model can explicitly inpaint invalid regions instead of overfitting to artifacts.




Fig. 6: Qualitative comparison of Control-DINO conditioning methods at matched time steps for ScanNet++. We show generated outputs at $t=35$ and $t=49$ for *2D*, *Voxel*, *Concerto*, and *Mesh* conditioning. In the bottom row $t=0$, we show the conditioning DINO feature signal for each method, while the GT frame at $t=0$ is the given starting frame.

For the **voxel** setting, we investigate whether low-resolution geometry with color cues is sufficient. Starting from an input video sequence, we reconstruct geometry from a dense point cloud and associate DINO-encoded image features to 3D points, then voxelize this feature-augmented point cloud. The resulting voxel grid is rendered back to the target views to produce DINO-like conditioning maps; as in the mesh case, we propagate hole masks so the model can ignore unobserved areas during denoising, visualized in Fig. 5.

Regarding video-from-3D, for **Concerto-based** conditioning we target the realistic case where only sparse geometry is available and RGB coverage is limited to one input view. We extract Concerto point-cloud descriptors, voxelize them, and render Concerto feature maps as the control signal. This setup allows generation from sparse, color-free 3D observations while preserving semantic consistency. Quantitative comparisons are reported in Table 4, and qualitative conditioning examples are shown in Fig. 6.



Fig. 7: We evaluate different methods on a reconstruction task on the DL3DV dataset, featuring high-frequency geometry such as foliage and books. We have found that our method tends to replicate the input more accurately, despite color and lighting shifts caused by the augmentation.




Fig. 8: Additional 2D transfer results on T&T scenes with different validation styles.

For image-based conditioning, our 2D Control-DINO reference outperforms state-of-the-art depth- and edge-conditioned baselines (notably Wan2.2) by a large margin in reconstruction quality; our 3D-derived variants remain competitive with these baselines while operating in a different application setting. Control-DINO performs best when full color information is available, consistent with the 2D setting where appearance cues are strong. Variants based on untextured geometry show lower reconstruction quality, as expected: missing appearance reduces feature fidelity and weakens structural disambiguation, and sparse 3D observations leave holes that carry no conditioning signal.

Concerto-based conditioning is reported in two regimes (Table 4): with color (*Concerto*) and without color (*Concerto inpainted*, the uncolored point-cloud setting). Despite the stronger input sparsity, the color-free Concerto setup still provides usable appearance priors and enables plausible scene generation, high-




Fig. 9: We show intermediate frames for a relighting task compared to Light-a-Video [53] that applies the same transfer function (IC-Light) to all (key)frames as guidance, while we just apply it to the first frame. We obtain similar results qualitatively and in particular observe high similarity when evaluating on the Light-a-Video dataset, this is because both methods share the same frozen backbone.

lighting its practical value for generation-focused use-cases rather than strict reconstruction.

4.6 Ablations

We first ablate the effect of training data composition (real-only vs. mixed real+augmented). As shown in Table 1, mixing data can slightly reduce generation quality, while improving transfer fidelity.

We ablate the feature conditioning against two PCA decomposition strategies (Table 6): (1) Given paired real and stylized DINOv3 features f_{real} and f_{style} , we compute difference vectors $d = f_{\text{real}} - f_{\text{style}}$, perform eigendecomposition on their covariance $S = (1/M)D^T D$ to obtain the top- K style directions V_k , project them out via $P = I - V_k V_k^T$, then run standard PCA on the cleaned features $P f_{\text{real}}$ to extract the top- D structure directions as our final projection basis; (2) A tail dropping strategy described in [13] where we first compute a 64-component PCA basis during training, then randomly retain only the first k components with $k \in \{8, 16, 32, 64\}$. At inference, we ablate against $k = 8$ and $k = 64$. As can be seen in Table 5, PCA decomposition, in part because of the already small DINO model we are using, is quite disruptive of the information signal. Even when computing

a style invariant decomposition with paired data the amount of information explained remains below 40%, which indicates substantial information loss.

We consider the necessity of spatially mixing the signal with the hidden state in a transformer block by replacing it with a simple linear projection directly into the residual (similar to [13]), we have found that the quality of the images remains quite high but the model has difficulty following the conditioning under large camera motion.

5 Discussion

We introduced a DINO-based conditioning mechanism for video diffusion. Dense self-supervised features provide robust structural control while allowing flexible appearance variation and feature-space resampling at generation time. The same interface spans 2D transfer (stylization, relighting) and 3D-guided generation from meshes, voxels, and point clouds. This is not without limitations, as appearance and semantics are not always perfectly disentangled: the model occasionally leaks source appearance, particularly for styles far from the training distribution. Lower guidance mitigates this, but we report all results at a fixed guidance scale for fairness; Table 3 reports the control to diversity behavior across scales. More broadly, using high-dimensional features as conditioning entails a trade-off with generative flexibility: stronger control restricts the diversity of plausible outputs. For 3D-derived conditioning, generation quality additionally depends on the fidelity of the underlying reconstruction and on how well the projected features approximate the distribution of 2D DINO features the model was trained on.

We believe that learning spatial weighting of the conditioning signal, as well as extending to other foundation feature spaces, could further improve the balance between control and generation quality. We release source code and weights for the community to build upon.

References

1. Baldassarre, F., Szafraniec, M., Terver, B., Khalidov, V., Massa, F., LeCun, Y., Labatut, P., Seitzer, M., Bojanowski, P.: Back to the features: Dino as a foundation for video world models. arXiv preprint arXiv:2507.19468 (2025) 2
2. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., Jampani, V., Rombach, R.: Stable video diffusion: Scaling latent video diffusion models to large datasets (2023) 3
3. Bordes, F., Balestriero, R., Vincent, P.: High fidelity visualization of what your self-supervised representation knows about. Transactions on Machine Learning Research (2022) 4
4. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9650–9660 (October 2021) 2

5. Chen, W., Wu, J., Xie, P., Wu, H., Li, J., Xia, X., Xiao, X., Lin, L.: Control-a-video: Controllable text-to-video generation with diffusion models (2023) [3](#)
6. Chen, X., Huang, L., Liu, Y., Shen, Y., Zhao, D., Zhao, H.: Anydoor: Zero-shot object-level image customization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6593–6602 (June 2024) [4](#)
7. Chen, X., Zhang, C., Fu, C., Yang, Z., Zhou, K., Zhang, Y., He, J., Zhang, Y., Sun, M., Wang, Z., Dong, Z., Long, X., Meng, L.: Driving with dino: Vision foundation features as a unified bridge for sim-to-real generation in autonomous driving. arXiv preprint arXiv:2602.06159 (2026) [4](#), [11](#)
8. Chen, Y., Lai, Y.K., Liu, Y.J.: Cartoongan: Generative adversarial networks for photo cartoonization. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [6](#)
9. Davtyan, A., Sameni, S., Ommer, B., Favaro, P.: Cage: Unsupervised visual composition and animation for controllable video generation. Proceedings of the AAAI Conference on Artificial Intelligence 39(15), 16163–16171 (Apr 2025) [4](#)
10. Graikos, A., Yellapragada, S., Le, M.Q., Kapse, S., Prasanna, P., Saltz, J., Samaras, D.: Learned representation-guided diffusion models for large-image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8532–8542 (2024) [4](#)
11. Guo, Y., Yang, C., Rao, A., Agrawala, M., Lin, D., Dai, B.: Sparsectrl: Adding sparse controls to text-to-video diffusion models. In: Leonardis, A., Ricci, E., Roth, S., Russakovsky, O., Sattler, T., Varol, G. (eds.) Computer Vision – ECCV 2024. pp. 330–348. Springer Nature Switzerland, Cham (2025) [3](#)
12. Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., Fleet, D.J.: Video diffusion models. In: Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., Oh, A. (eds.) Advances in Neural Information Processing Systems. vol. 35, pp. 8633–8646. Curran Associates, Inc. (2022) [3](#)
13. Huang, Y., Xiong, W., Zhang, H., Chen, C., Liu, J., Yan, M., Chen, S.: Dive: Taming dino for subject-driven video editing. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 16004–16014 (October 2025) [2](#), [4](#), [14](#), [15](#)
14. Huang, Z., He, Y., Yu, J., Zhang, F., Si, C., Jiang, Y., Zhang, Y., Wu, T., Jin, Q., Chanpaisit, N., Wang, Y., Chen, X., Wang, L., Lin, D., Qiao, Y., Liu, Z.: VBench: Comprehensive Benchmark Suite for Video Generative Models. In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 21807–21818. IEEE Computer Society, Los Alamitos, CA, USA (Jun 2024) [10](#)
15. Jevtić, A., Reich, C., Wimbauer, F., Hahn, O., Rupprecht, C., Roth, S., Cremers, D.: Feed-forward SceneDINO for unsupervised semantic scene completion. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2025) [2](#)
16. Jiang, Z., Han, Z., Mao, C., Zhang, J., Pan, Y., Liu, Y.: Vace: All-in-one video creation and editing. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 17191–17202 (October 2025) [3](#)
17. Ju, X., Ye, W., Liu, Q., Wang, Q., Wang, X., Wan, P., Zhang, D., Gai, K., Xu, Q.: Fulldit: Video generative foundation models with multimodal control via full attention. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 15737–15747 (October 2025) [3](#)
18. Kim, G., Han, J., Cho, S.: Videofrom3d: 3d scene video generation via complementary image and video diffusion models. In: SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers ’25). pp. 1–11. ACM, Hong Kong, Hong Kong (2025) [3](#), [5](#)

19. Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: Benchmarking large-scale scene reconstruction. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2017) 8
20. Ku, M., Wei, C., Ren, W., Yang, H., Chen, W.: Anyv2v: A tuning-free framework for any video-to-video editing tasks. Transactions on Machine Learning Research (2024), reproducibility Certification 9, 10
21. Lee, S., Ebbecke, T., Millon, E., Beddow, W., Zhuo, L., García-Ferrero, I., Esparraguera, L., Petrescu, M., Saß, G., Menezes, G., Perez, V.: Flux.1 krea [dev]. <https://github.com/krea-ai/flux-krea> (2025), accessed: 2026-06-23 1
22. Liang, F., Wu, B., Wang, J., Yu, L., Li, K., Zhao, Y., Misra, I., Huang, J.B., Zhang, P., Vajda, P., Marculescu, D.: Flowvid: Taming imperfect optical flows for consistent video-to-video synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8207–8216 (June 2024) 3
23. Lin, H., Cho, J., Zala, A., Bansal, M.: Ctrl-adapter: An efficient and versatile framework for adapting diverse controls to any diffusion model. In: The Thirteenth International Conference on Learning Representations (2025) 3
24. Ling, L., Sheng, Y., Tu, Z., Zhao, W., Xin, C., Wan, K., Yu, L., Guo, Q., Yu, Z., Lu, Y., Li, X., Sun, X., Ashok, R., Mukherjee, A., Kang, H., Kong, X., Hua, G., Zhang, T., Benes, B., Bera, A.: DL3dv-10k: A large-scale scene dataset for deep learning-based 3d vision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 22160–22169 (June 2024) 7, 8
25. Luo, G., Dunlap, L., Park, D.H., Holynski, A., Darrell, T.: Diffusion hyperfeatures: Searching through time and space for semantic correspondence. In: Advances in Neural Information Processing Systems (2023) 4
26. Ma, Y., Feng, K., Hu, Z., Wang, X., Wang, Y., Zheng, M., Wang, B., Wang, Q., He, X., Wang, H., Zhu, C., Liu, H., He, Y., Wang, Z., Li, Z., Li, X., Han, S., Guo, Y., Liu, W., Xu, D., Zhang, L., Chen, Q.: Controllable video generation: A survey (2026) 2, 3
27. Melnik, A., Ljubljanac, M., Lu, C., Yan, Q., Ren, W., Ritter, H.: Video diffusion models: A survey. Transactions on Machine Learning Research (2024), survey Certification 2
28. Mou, C., Wang, X., Xie, L., Wu, Y., Zhang, J., Qi, Z., Shan, Y., Qie, X.: T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models (2023) 3
29. NVIDIA, :, Alhaija, H.A., Alvarez, J., Bala, M., Cai, T., Cao, T., Cha, L., Chen, J., Chen, M., Ferroni, F., Fidler, S., Fox, D., Ge, Y., Gu, J., Hassani, A., Isaev, M., Jannaty, P., Lan, S., Lasser, T., Ling, H., Liu, M.Y., Liu, X., Lu, Y., Luo, A., Ma, Q., Mao, H., Ramos, F., Ren, X., Shen, T., Sun, X., Tang, S., Wang, T.C., Wu, J., Xu, J., Xu, S., Xie, K., Ye, Y., Yang, X., Zeng, X., Zeng, Y.: Cosmos-transfer1: Conditional world generation with adaptive multimodal control (2025) 9, 10
30. NVIDIA: Cosmos world foundation model platform for physical ai (2025) 2
31. Oquab, M., Darcet, T., Moutakanni, T., Vo, H.V., Szafraniec, M., Khalidov, V., Fernandez, P., HAZIZA, D., Massa, F., El-Nouby, A., Assran, M., Ballas, N., Galuba, W., Howes, R., Huang, P.Y., Li, S.W., Misra, I., Rabbat, M., Sharma, V., Synnaeve, G., Xu, H., Jegou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.: DINOv2: Learning robust visual features without supervision. Transactions on Machine Learning Research (2024), featured Certification 3
32. Parmar, G., Park, T., Narasimhan, S., Zhu, J.Y.: One-step image translation with text-to-image models. arXiv preprint arXiv:2403.12036 (2024) 3

33. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 4195–4205 (October 2023) [3](#)
34. Perez, E., Strub, F., de Vries, H., Dumoulin, V., Courville, A.C.: Film: Visual reasoning with a general conditioning layer. In: AAAI (2018) [4](#)
35. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., Massa, F., Haziza, D., Wehrstedt, L., Wang, J., Darcet, T., Moutakanni, T., Sentana, L., Roberts, C., Vedaldi, A., Tolan, J., Brandt, J., Couprie, C., Mairal, J., Jégou, H., Labatut, P., Bojanowski, P.: Dinov3 (2025) [3](#)
36. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (2021) [10](#)
37. Stracke, N., Baumann, S.A., Bauer, K., Fundel, F., Ommers, B.: Cleandift: Diffusion features without noise. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2025) [4](#)
38. Tachibana, Y., Tomita, S., Tanno, T., Kaji, H., Nagai, T.: Animeganv2: Rethinking the high-quality anime image generation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (2021) [6](#)
39. Tang, L., Jia, M., Wang, Q., Phoo, C.P., Hariharan, B.: Emergent correspondence from image diffusion. In: Thirty-seventh Conference on Neural Information Processing Systems (2023) [4](#)
40. Team, R., Gao, Z., Wang, Q., Zeng, Y., Zhu, J., Cheng, K.L., Li, Y., Wang, H., Xu, Y., Ma, S., Chen, Y., Liu, J., Cheng, Y., Yao, Y., Zhu, J., Meng, Y., Zheng, K., Bai, Q., Chen, J., Shen, Z., Yu, Y., Zhu, X., Shen, Y., Ouyang, H.: Advancing open-source world models (2026) [2](#)
41. Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., Zeng, J., Wang, J., Zhang, J., Zhou, J., Wang, J., Chen, J., Zhu, K., Zhao, K., Yan, K., Huang, L., Feng, M., Zhang, N., Li, P., Wu, P., Chu, R., Feng, R., Zhang, S., Sun, S., Fang, T., Wang, T., Gui, T., Weng, T., Shen, T., Lin, W., Wang, W., Wang, W., Zhou, W., Wang, W., Shen, W., Yu, W., Shi, X., Huang, X., Xu, X., Kou, Y., Lv, Y., Li, Y., Liu, Y., Wang, Y., Zhang, Y., Huang, Y., Li, Y., Wu, Y., Liu, Y., Pan, Y., Zheng, Y., Hong, Y., Shi, Y., Feng, Y., Jiang, Z., Han, Z., Wu, Z.F., Liu, Z.: Wan: Open and advanced large-scale video generative models (2025) [3](#), [9](#), [10](#)
42. Wang, H., Spinelli, M., Wang, Q., Bai, X., Qin, Z., Chen, A.: Instantstyle: Free lunch towards style-preserving in text-to-image generation (2024) [7](#), [11](#)
43. Wu, X., DeTone, D., Frost, D., Shen, T., Xie, C., Yang, N., Engel, J., Newcombe, R., Zhao, H., Straub, J.: Sonata: Self-supervised learning of reliable point representations. In: Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR). pp. 22193–22204 (June 2025) [4](#)
44. Yang, Z., Teng, J., Zheng, W., Ding, M., Huang, S., Xu, J., Yang, Y., Hong, W., Zhang, X., Feng, G., Yin, D., Yuxuan Zhang, Wang, W., Cheng, Y., Xu, B., Gu, X., Dong, Y., Tang, J.: Cogvideox: Text-to-video diffusion models with an expert transformer. In: The Thirteenth International Conference on Learning Representations (2025) [3](#), [8](#)
45. Ye, H., Zhang, J., Liu, S., Han, X., Yang, W.: Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models (2023) [3](#)
46. Yeshwanth, C., Liu, Y.C., Nießner, M., Dai, A.: Scannet++: A high-fidelity dataset of 3d indoor scenes. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12–22 (2023) [8](#)

47. Yu, S., Kwak, S., Jang, H., Jeong, J., Huang, J., Shin, J., Xie, S.: Representation alignment for generation: Training diffusion transformers is easier than you think. In: The Thirteenth International Conference on Learning Representations (2025) [4](#)
48. Zhang, J., Herrmann, C., Hur, J., Cabrera, L.P., Jampani, V., Sun, D., Yang, M.H.: A tale of two features: Stable diffusion complements DINO for zero-shot semantic correspondence. In: Thirty-seventh Conference on Neural Information Processing Systems (2023) [4](#)
49. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: 2023 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 3813–3824 (2023) [3](#), [8](#)
50. Zhang, L., Rao, A., Agrawala, M.: Scaling in-the-wild training for diffusion-based illumination harmonization and editing by imposing consistent light transport. In: The Thirteenth International Conference on Learning Representations (2025) [7](#), [10](#)
51. Zhang, Y., Wu, X., Lao, Y., Wang, C., Tian, Z., Wang, N., Zhao, H.: Concerto: Joint 2d-3d self-supervised learning emerges spatial representations. In: The Thirty-ninth Annual Conference on Neural Information Processing Systems (2025) [4](#), [7](#)
52. Zhou, G., Pan, H., Lecun, Y., Pinto, L.: DINO-WM: World models on pre-trained visual features enable zero-shot planning. In: Singh, A., Fazel, M., Hsu, D., Lacoste-Julien, S., Berkenkamp, F., Maharaj, T., Wagstaff, K., Zhu, J. (eds.) Proceedings of the 42nd International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 267, pp. 79115–79135. PMLR (13–19 Jul 2025) [2](#)
53. Zhou, Y., Bu, J., Ling, P., Zhang, P., Wu, T., Huang, Q., Li, J., Dong, X., Zang, Y., Cao, Y., Rao, A., Wang, J., Niu, L.: Light-a-video: Training-free video relighting via progressive light fusion. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 13315–13325 (October 2025) [14](#)