

# Precise Action-to-Video Generation Through Visual Action Prompts

Yuang Wang<sup>2\*</sup> Chao Wen<sup>3\*</sup> Haoyu Guo<sup>2\*</sup> Sida Peng<sup>2</sup> Minghan Qin<sup>4</sup>  
 Hujun Bao<sup>2</sup> Xiaowei Zhou<sup>2</sup> Ruizhen Hu<sup>1,5†</sup>

<sup>1</sup>Xiangjiang Lab <sup>2</sup>Zhejiang University <sup>3</sup>Fudan University  
<sup>4</sup>Tsinghua University <sup>5</sup>Shenzhen University

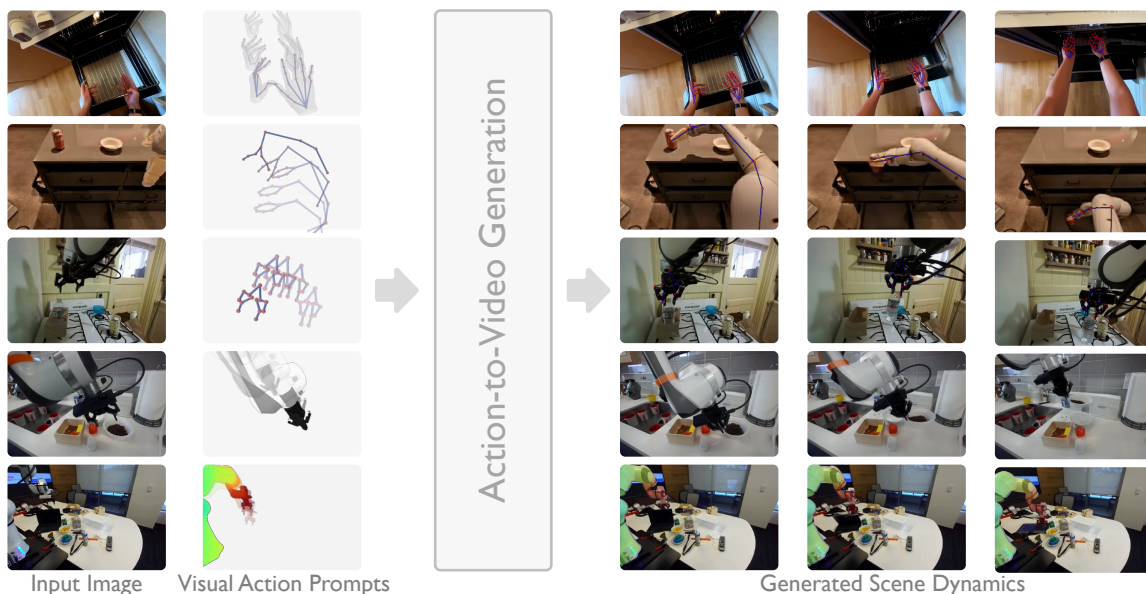


Figure 1. We propose visual action prompts as unified representations for complex, high-degree-of-freedom actions (*e.g.*, simulating scene dynamics driven by human hands or robotic grippers). Visual action prompts are “renderings” of subjects’ action-induced 3D structures, among which we use skeleton as the primary representation for its acquisition efficiency. This paradigm enables training action-driven video generation models across heterogeneous datasets while facilitating cross-domain knowledge transfer.

## Abstract

We present visual action prompts, a unified action representation for action-to-video generation of complex high-DoF interactions while maintaining transferable visual dynamics across domains. Action-driven video generation faces a precision-generality tradeoff: existing methods using text, primitive actions, or coarse masks offer generality but lack precision, while agent-centric action signals provide precision at the cost of cross-domain transferability. To balance action precision and dynamic transferability, we propose to “render” actions into precise visual prompts as domain-agnostic representations that preserve both geometric precision and cross-domain adaptability for complex

actions; specifically, we choose visual skeletons for their generality and accessibility. We propose robust pipelines to construct skeletons from two interaction-rich data sources — human-object interactions (HOI) and dexterous robotic manipulation — enabling cross-domain training of action-driven generative models. By integrating visual skeletons into pretrained video generation models via lightweight fine-tuning, we enable precise action control of complex interaction while preserving the learning of cross-domain dynamics. Experiments on EgoVid [64], RT-1 [11] and DROID [35] demonstrate the effectiveness of our proposed approach.

## 1. Introduction

With improvements in quality and controllability of visual generative models [7, 46, 50, 72], action-driven generative models are now widely applied in gaming [12, 14, 21, 61,

\*Equal contribution. †Corresponding author: Ruizhen Hu.

74], decision-making [19, 69], robot learning and simulation [2, 4, 67, 81]. These frameworks utilize action sequences as conditional inputs to video generation models, producing video frames that depict the outcomes of those actions. This paper focuses on action-driven generative models under complex, high-DoF action control, such as simulation of scene dynamics governed by human hands and robotic grippers operations.

The primary challenge lies in the absence of a unified yet precise action representation to effectively model high-DoF and heterogeneous actions across diverse fields and applications, impeding the training of a unified model facilitating knowledge transfer across domains. Different action representations have been proposed in diverse domains, such as text [69], high-level primitive skills [12, 14, 23, 61], and low-level states of specific agent configurations [2, 67, 81]. However, action representations face a precision-generality trade-off. Text, though universally applicable, can only present the high-level intention of an action. Pre-defined skills like character movement and gaming actions (e.g., shooting [4, 14]) are similar to text, both of which present high-level action primitives, while general, fail to represent complex, low-level character motions and intricate interactions with the environment that are critical in applications like motion sensing gaming. The low-level state of specific agent configuration (e.g., a robot arm’s end-effector 6D pose and gripper openness) is an almost lossless action representation. It is adopted in domains requiring the utmost precision like robotic simulation and planning [2, 67, 81]. However, it tightly couples the action signal to specific embodiments, lacking generality.

To achieve a balanced precision-generality tradeoff, we propose using *precise visual action prompts* as unified control signals for interactive generative models driven by complex actions, while retaining generality across domains. Visual action prompts are produced by “rendering” the 3D structure of actions-induced agent states into image space, which can be in different forms, such as coarse masks, colored renderings or depth maps, and 2D skeletons. They can effectively represent actions of high-DoF subjects — such as human hands, robot grippers, and dexterous hands — with high precision.

Some forms of visual action prompts come with notable drawbacks. Mesh-rendering-based approaches require reconstruction of complete meshes, which are challenging to scale up with in-the-wild data. Coarse subject masks [3, 57] are easier to recover but suffer from occlusions; their limited precision is also problematic in fine-grained tasks like robot simulation. To balance the ease of recovery and action precision, we adopt *skeletons* as the unified control signal, which have long been a universal tool in animation [8, 31, 38, 45, 48]. They can be robustly recovered from in-the-wild data [1, 18, 29, 53, 56, 59], facilitating large-scale training across domains.

To demonstrate the effectiveness of visual action prompts, we propose scalable strategies to recover complete skeletons of human hands and robot grippers on datasets including EgoVid [64], RT-1 [11] and DROID [35]. Then, we fine-tune a base video generation model [72] to adapt it to an action-controllable model supporting intricate interaction. We further show that visual action prompts serve as a more precise and easier to learn control signal of action for video models, compared to text and agent-centric states. Moreover, we show that visual action prompts enable training a unified model on multi-domain data including HOI and robot manipulation, which facilitates cross-domain knowledge transfer of interaction-driven dynamics.

In summary, we make the following contributions:

- We propose using precise visual action prompts, specifically skeletons, as the unified action representation for action-driven generative models in scenarios involving complex, high-DoF actions.
- We introduce scalable strategies to recover skeleton-based visual action prompts on interaction-rich datasets including HOI and robot manipulation.
- We demonstrate visual action prompts’ advantages including ease of learning, precision and generality, which enable joint training on heterogeneous data and facilitate knowledge transfer.

## 2. Related Work

**Action-to-video generation.** Recent works have been pursuing action control of video models to enable interaction-rich applications like gaming and agent training [4, 12, 14, 17, 61, 74], robotic simulation/learning [2, 67, 81], and general decision making [70]. Typical game actions involve predefined primitives where state transitions follow explicit rules or rules combined with physics simulation. Video-driven game research explores various action representations: Genie [12] learns a discrete set of latent actions from unlabeled videos to drive generation; GameNGen [61] and DIAMOND [4] directly map primitive actions (e.g., directional moves, shooting) to video frames. Recent efforts further scale these approaches [14, 16, 21, 74], yet remain constrained by primitive action representations that limit complex environmental interactions. Video generation for world model and RL agents training [4, 12, 17, 67] are also still limited to learning predefined primitive actions, incapable of simulating high-DoF embodiments interacting with complex environments. In simulation, UniSim [69] represents high-level action intentions with texts, which is general but lacks precision. Text as action is more plausible for general decision making like generating how-to guides [55] and high-level policies for robotics [9, 10, 19, 37]. In the other end, IRASim [81] and Cosmos [2] employ agent-centric actions like 7-DoF end-effector states, which is almost lossless for positional control but not general. Recently, CosHand [57] and Inter-

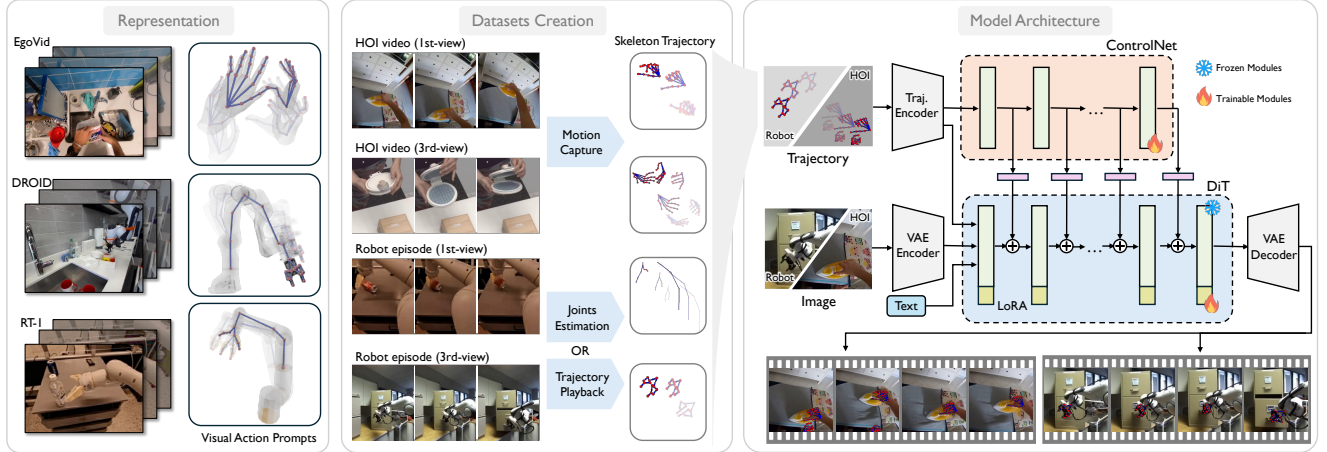


Figure 2. **Action-to-video generation with visual action prompts.** We project action-induced 3D structural dynamics of diverse agents into 2D visual action prompts, primarily 2D skeletons, establishing a unified control signal for action-conditioned video generation. We design data creation pipelines for HOI and robot videos to robustly recover their 2D skeletons. The constructed visual action prompts are injected into a pretrained video generation model for fine-tuning and generating plausible interaction-driven visual dynamics.

Dyn [3] explore generating future state images or videos using coarse hand masks as indications of actions. However, mask-based representations are fragile to occlusions, and segmented masks are often imprecise, which hinders applications requiring highly accurate actions.

**Motion control and character animation.** Controllability of video generation models is critical for downstream applications. Extensive research focuses on achieving control over the synthesized content in different levels, especially for the target agent. Region-based methods [54, 62] provide high-level guidance for local motion. Sparse/dense trajectory-based methods [24, 66, 73, 79] govern local object movements or camera motion. Character animation techniques further enable precise control through skeletons [31, 43, 78], mesh renderings [80, 82], or reference videos [32, 63], all requiring explicit dynamic specifications for the control targets. In contrast, we aim to generate interaction-induced dynamics of the whole environment by providing the model with action signals only, without relying on its inherent dynamics.

**Dynamic-rich datasets.** Effectively orchestrating diverse datasets with rich interaction and interaction-driven dynamics lays the groundwork for action-driven video generation. Video datasets about human-object-interaction (HOI) are natural sources for learning interaction-driven dynamics. Early works like SSV2 and Kinetics [25, 34] collect datasets of humans performing basic actions with everyday objects. Despite their substantial scale, their relatively low video quality falls short of the standards required by modern video generation models. Recent advancements have introduced large-scale egocentric human activity datasets [15, 26, 27, 41] to advance behavioral understanding and learning. A notable example, Ego4D, offers 3,670 hours of daily activity videos across diverse scenarios. However, its lengthy sequences are suboptimal for generation tasks. EgoVid-5M [64] ad-

dresses this by curating trimmed, filtered, and captioned clips from Ego4D, making them more plausible for generative models. Beyond video data, specialized datasets for HOI motion research provide 3D hand-object annotations [6, 20, 22, 40, 75]. While valuable for high-precision 3D-controlled fine-tuning, their limited diversity restricts utility in foundational interaction-driven dynamics learning. Complementing HOI resources, embodied AI and robotics research has yielded high-quality interaction data as well. Open X-Embodiment [47] aggregates multi-task datasets of complex robotic interactions across embodiments. We select RT-1 [11] and DROID [35] from this collection for their scale and relatively precise camera calibration.

### 3. Method

Given an image observation as the initial frame and a sequence of actions from human hands or robot grippers, our goal is to generate videos that accurately depict the interaction outcomes under precise action control. To achieve this, we introduce a general and precise visual action prompt for video generation models. Fig. 2 illustrates our framework, which includes visual action representation, dataset construction, and visual dynamic modeling.

To maximize data scale and interaction relevance, we focus on two primary agents: human hands and robotic grippers. Their actions — despite kinematic differences — are uniformly encoded as skeletons, as our visual prompts (Sec. 3.1). To train our model, we process and annotate two types of datasets (Sec. 3.2): (1) human hand skeletons extracted via motion capture from HOI videos; and (2) robotic gripper skeletons are synthesized through joint-state rendering from robotic episodes. Leveraging these large-scale (skeleton, video) pairs, we fine-tune video generation models



to enable visual action prompt control (Sec. 3.3).

### 3.1. Visual action prompts

Our goal is to develop a generalizable video generation model capable of synthesizing plausible scene dynamics and interaction outcomes  $\mathbf{s}_{1:t} \in \mathcal{S}$ , given the initial observation  $\mathbf{s}_0$  and driven by user-specified complex action sequences  $\mathbf{a}_{0:t-1} \in \mathcal{A}$ . The problem can be formulated as learning the conditional probability distribution  $\mathcal{P}$  governing the state representation:

$$\mathbf{s}_{1:t} \sim P(\mathbf{s}_{1:t} | \mathbf{s}_0, \mathbf{a}_{0:t-1}). \quad (1)$$

While the formulation is brief, there are certain challenges in practice: (1) accommodating diverse configurations of intricate action spaces while ensuring compatibility across tasks, and (2) retaining the transferability of visual dynamics under precise action control, thus enabling the model to learn from scalable datasets composed of diverse domains.

To resolve those challenges, our core insight is to map the action sequence  $\mathbf{a}_{0:t-1}$  to *visual action prompts*  $\mathbf{v}_{1:t}$  as follows:

$$\mathbf{v}_{1:t} = \mathcal{R}(\mathbf{a}_{0:t-1}) \in \mathbb{R}^{T \times H \times W \times C}, \quad (2)$$

where  $T$  represents the action trajectory length,  $H$  and  $W$  represent the image height and width,  $C$  is the number of channels, determined by the specific visual representation, and  $\mathcal{R}$  indicates rendering operation according to known camera parameters. We consider mesh-based renderings (e.g., colored images, depth maps) and 2D skeletons as *precise* visual action prompts. Given the challenges associated with recovering fine-grained meshes from in-the-wild data, we opt to employ 2D skeletons as our primary representation.

### 3.2. Scalable dataset creation

To enable the learning of transferable visual dynamics under precise action conditioning, we construct large-scale (skeleton, video) pairs from two distinct sources: skeletons are estimated from in-the-wild HOI videos via our proposed pipeline (Fig. 3a), while for robot manipulation episodes, they are rendered directly from state logs, followed by an optional correction step to ensure tight alignment with the visual observations (Fig. 3b).

**In-the-wild HOI videos.** We leverage in-the-wild HOI-centric videos for their rich hand-object interactions and scene dynamics, which is ideal for learning visual dynamic models. However, severe occlusions in these videos make direct 2D pose estimation unreliable [36, 42, 71]. To robustly extract 3D hand mesh trajectories, we introduce a four-stage pipeline that addresses common failures like missed detections and temporal jitter (see Algorithm 1 and Fig. 3a): (1) Initialization: We first detect all potential hands per-frame using 3d hand mesh recovery method Wilor [51]. (2) Temporal Stabilization: We then form consistent tracklets and correct handedness errors using SAMURAI [68]. (3) Refinement:

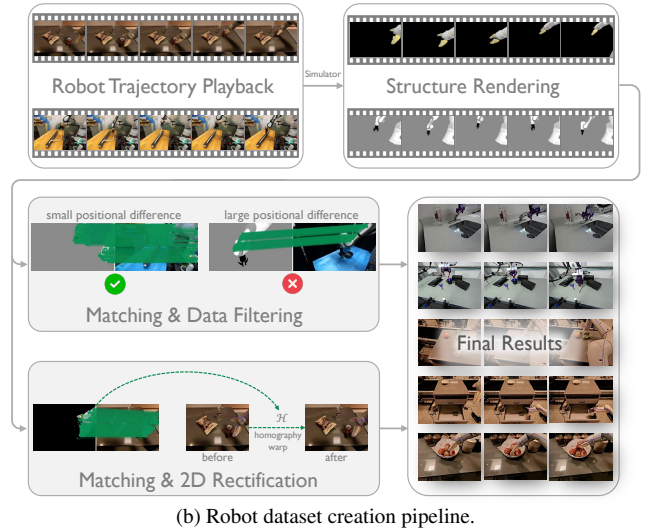
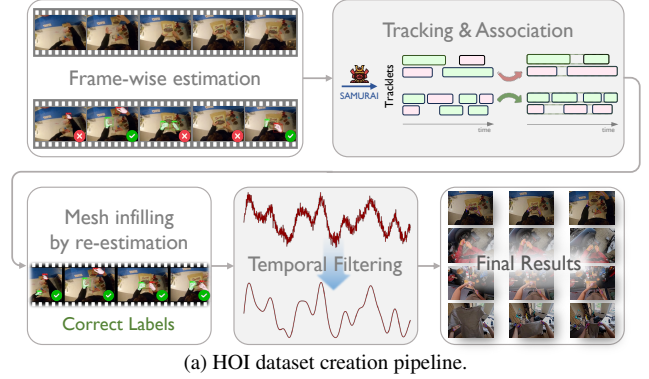


Figure 3. Pipelines of recovering 2D skeletons for visual action prompts from human-object interaction (HOI) and robotic datasets.

Missing meshes within these tracklets are re-estimated. (4) Smoothing: Finally, we apply a OneEuro filter [13] to the MANO trajectories to eliminate jitter.

**Robot manipulation episodes.** We also utilize robotic manipulation datasets (DROID [35], RT-1 [11]), which offer focused interactions and scene dynamics while simplifying 3D skeleton extraction via robot state logs. However, camera calibration errors and temporal drift are common issues. To ensure precise 2D skeleton alignment with video observations, we implement a vision-based correction pipeline (Fig. 3b): (1) Episode Filtering: We use MatchAnything [28] to match rendered robot meshes against real observations, discarding episodes with significant matching coordinates discrepancy. (2) Homography Rectification: For the remaining episodes with camera drift, we apply per-frame homography warping to adjust the initial skeleton renderings in 2D, also guided by image matching, ensuring precise alignment with real-world observations.



---

**Algorithm 1** Tracking and Association

---

```
1:  $\mathcal{B} \leftarrow \emptyset$  ▷ Initialize the bounding box set
2:  $\mathcal{T} \leftarrow \emptyset$  ▷ Initialize the tracklet set
3:  $\mathcal{A} \leftarrow \{\}$  ▷ Initialize the association dictionary
4:
5: for all frame  $F$  in video do
6:    $\mathcal{B} \leftarrow \mathcal{B} \cup \text{Detect}(F)$  ▷ Based on Wilor
7: end for
8:  $\mathcal{B}_{\text{untracked}} \leftarrow \text{Sort}(\mathcal{B})$  ▷ By descending confidence
9: while  $\mathcal{B}_{\text{untracked}} \neq \emptyset$  do
10:   $B \leftarrow \mathcal{B}_{\text{untracked}}.\text{pop}(0)$ 
11:  if  $\max_{T \in \mathcal{T}} \text{IoU}(B, T) < \theta_{\text{IoU}}$  then
12:     $T_{\text{new}} \leftarrow \text{Track}(B)$  ▷ Based on SAMURAI
13:     $\mathcal{T}.\text{append}(T_{\text{new}})$ 
14:  end if
15: end while
16: for all  $(B, T) \in \mathcal{B} \times \mathcal{T}$  do
17:  if  $\text{IoU}(B, T) \geq \theta_{\text{IoU}}$  then
18:     $\mathcal{A}[T].\text{append}(B)$ 
19:  end if
20: end for
21:  $\mathcal{T} \leftarrow \text{HandnessFilter}(\mathcal{T}, \mathcal{A})$ 
22:  $\mathcal{T} \leftarrow \text{Merge}(\mathcal{T}, \mathcal{A})$ 
23:  $\mathcal{T} \leftarrow \text{NumberOfHandsFilter}(\mathcal{T}, \mathcal{A})$ 
```

---

### 3.3. Visual dynamics model with precise control

We build our model based on CogVideoX [72], which is a text-to-video generation model pretrained on large-scale (text, video) pairs and further fine-tuned with (text, initial frame, video) triplets to a (text, image)-to-video model. Its architecture includes: a pretrained text encoder, a video VAE, and a DiT [49] model with FullAttention for spatio-temporal video tokens and text token processing. Due to the high data demands of training visual dynamics models from scratch, we leverage CogVideoX as a pretrained base model.

To integrate visual action prompts, as shown in Fig. 2, we first encode the control signals. Specifically, for controls in the form of skeleton or mesh, we render them as sequences of RGB images  $\mathbf{v}_{1:t} \in \mathbb{R}^{T \times H \times W \times C}$ , where  $C = 3$ . These sequences are then fed into an 3D convolutional trajectory encoder to latent states  $\mathbf{s}_{1:t/4} \in \mathbb{R}^{\frac{T}{4} \times \frac{H}{8} \times \frac{W}{8} \times 16}$ . For depth control, we directly feed  $\mathbf{v}_{1:t}$  with  $C = 1$  into the encoder with the same architecture.

Direct supervised fine-tuning of the video generation model pretrained on massive datasets may lead to overfitting or the loss of generalized knowledge. Therefore we leverage ControlNet [76] to inject the visual action prompts. Specifically, we create trainable copies of the first 14 blocks of the pretrained DiT with zero-initialized linear layers, and inject visual action prompts  $\mathbf{s}_{1:t/4}$  in these blocks. Moreover, following Wonderland [39], we adopt a dual-branch condi-

tioning mechanism by injecting  $\mathbf{s}_{1:t/4}$  also in the main DiT, through merging video and action tokens, and fine-tune the DiT backbone with LoRA [30].

During training, we amplify loss values around hand / gripper regions to prioritize learning the interaction and its induced dynamics. To mitigate dominance of self-motion over interaction dynamics in robot videos with lengthy tasks, we sample more clips around timestamps where gripper state changes.

## 4. Experiment

Our experiments aim to validate two core claims: (1) visual action prompts outperform alternative control signals, *e.g.*, text or agent-centric raw actions / states (Sec. 4.1) in driving interaction-aware scene dynamics; (2) the generality of visual action prompts across agent configurations and the effect of joint training on diverse datasets (Sec. 4.2). Moreover, ablation studies (Sec. 4.3) demonstrate the effectiveness of our model design and present results of different visual action prompts.

**Implementation details.** We curate three datasets with distinct characteristics. **EgoVid** [64]: A subset of 200k training clips (from 1M clips) containing around 120 frames, 30 fps videos of diverse daily activities with hand skeletons. Clips with significant viewpoint changes are filtered via optical flow [58] and point tracking [33]. We manually select 32 clips including direct/indirect dynamics for evaluation. **DROID** [35]: A subset of 47k training clips of random third-person perspectives (from 76k episodes collected across 13 labs), with one lab’s data reserved for evaluation on novel scenes. Tasks related to cleaning are retained for analyzing of novel skills. A total of 234 clips are used for quantitative assessment. **RT-1** [11]: A subset of 57k training clips from 6 basic skills, two skills (“close drawer” and “place object upright”) are held out for evaluation. Notably, unlike previous works such as IRASim [81], which focuses only on simulation of in-domain skills and scenes, we emphasize evaluating the interaction-driven dynamics of novel skills. We caption all video clips with scene-centric captions via Qwen2.5-VL [5], only including scene arrangements and appearances while excluding action/dynamic descriptions. For training text-as-action models, we regenerate captions with explicit action annotations. During training, we resize all video clips to the resolution of  $720 \times 480$  and subsample video frames to 25 with variable fps in a plausible range.

**Metrics.** We utilize multiple metrics to evaluate the generated videos. To evaluate the *visual similarity* between generated and ground truth videos, we report PSNR, SSIM [65], and LPIPS [77]. To evaluate *visual quality and temporal consistency*, we report the Fréchet Video Distance (FVD) [60]. Finally, to explicitly evaluate *dynamic correctness* of action and its impact on scene dynamics, we report the Spatio-temporal IoU proposed in Physics-IQ [44].




Figure 4. **Comparison of different action control signals.** Text as action (TI2V [72]) leads to ambiguity. Raw 7-DoF states/actions (IRASim [81]) leads to inaccurate control. Visual action prompts (Ours) facilitate dynamic learning under precise control.

#### 4.1. Agent-specific control

We demonstrate the superiority of visual action prompts over text-based and agent-centric action representations in control fidelity, scene dynamics plausibility, and learning efficiency. We conduct comparisons on two robot manipulation datasets RT-1 [11] and DROID [35]. For text-based control, we implement high-level action guidance by finetuning the (text,image)-to-video model CogVideoX [72]. For agent-centric control, IRASim [81] directly employs 7-DoF robot-specific actions (end-effector poses and gripper states) as action trajectories to drive video generation. For a fair comparison, we reimplement IRASim’s adaLN-based action injection upon CogVideoX and finetune the base model with 7-DoF state trajectories. We present quantitative comparisons between this reimplementation and the pretrained IRASim model in supplemental material to demonstrate its effectiveness.

As shown in Fig. 4, visual action prompts achieve better action fidelity through direct and precise action rendering, while text descriptions suffer from action ambiguity and less plausible generation. Although IRASim achieves effective control on RT-1 with fixed viewpoints, it fails to provide precise control on DROID with random third-person perspectives. Quantitative comparisons in Tab. 1 further confirm the

Dataset	Action Repr.	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FVD $\downarrow$	ST-IoU $\uparrow$
RT-1 [11]	Text (CogVideoX [72])	18.87	0.761	0.241	642.3	0.267
	Raw State (IRA-Sim [81])	23.96	0.854	0.127	302.2	0.507
	Skeleton (Ours)	<b>25.98</b>	0.859	<b>0.110</b>	288.6	<b>0.604</b>
	Skeleton (Ours unified)	24.90	0.847	0.121	<b>258.1</b>	0.576
DROID [35]	Text (CogVideoX)	18.10	0.790	0.200	248.3	0.239
	Raw State (IRA-Sim)	20.13	0.825	0.146	151.2	0.365
	Skeleton (Ours)	21.26	0.834	0.132	141.8	0.450
	Skeleton (Ours unified)	<b>21.58</b>	<b>0.836</b>	<b>0.126</b>	<b>124.4</b>	<b>0.478</b>
EgoVid [64]	Text (CogVideoX)	13.44	0.440	0.503	1638.6	–
	Skeleton (Ours)	14.71	0.482	0.430	1243.6	–
	Skeleton (Ours unified)	<b>14.93</b>	<b>0.486</b>	<b>0.421</b>	<b>1142.3</b>	–

Table 1. **Quantitative comparison on different datasets.** Visual action prompts (Ours) consistently outperform text-specified actions and raw agent-centric states. Joint training on all three datasets with a unified model leads to improved or comparable results across datasets.

Action Repr.	Known Lab & Skill	Novel Lab	Novel Skill
Mask IoU $\uparrow$ / Boundary IoU $\uparrow$ / $\mathcal{J}\&\mathcal{F}$ $\uparrow$			
<b>Single dataset</b>			
Text (CogVideoX)	34.1 / 31.9 / 33.0	22.8 / 30.9 / 26.9	20.5 / 25.2 / 22.9
Raw State (IRASim)	49.1 / 48.9 / 49.0	25.8 / 40.9 / 33.4	34.9 / 40.6 / 37.8
Skeleton (Ours)	<b>53.5 / 53.6 / 53.6</b>	<b>43.8 / 61.2 / 52.5</b>	<b>47.4 / 55.0 / 51.2</b>
<b>Joint training</b>			
Skeleton (Ours)	<b>58.9 / 60.4 / 60.0</b>	<b>46.5 / 63.3 / 54.9</b>	<b>49.9 / 57.9 / 53.9</b>

Table 2. **Quantitative comparison on different subset of the DROID [35] dataset.** The manipulated object is annotated with point prompts in the first frame and tracked with SAM 2 [52]. We report metrics between the masks extracted from the generated and ground truth videos.

comprehensive superiority of visual action prompts across all metrics.

Due to the small proportion of interactive foreground in robot manipulation data, photometric metrics fail to effectively highlight differences between methods in generating interaction-driven dynamics, we conduct additional dynamic-centric evaluations on DROID. We manually annotate point prompts for interacted objects in the first frame, ensuring SAM 2 [52] successfully tracks them in the ground truth videos. Using the same prompts, we apply SAM-2 on generated videos and compute the  $\mathcal{J}\&\mathcal{F}$  metric between generated tracking masks and ground truth masks to assess the model’s quality in generating dynamics for interacted objects. As shown in Tab. 2, visual action prompts achieve significant improvements across different DROID subsets.

#### 4.2. Agent-agnostic control

Thanks to visual action prompts’ balance over action precision and generality, we can integrate data from diverse domains of different agents to train a unified action-driven generative model. In this experiment, we demonstrate the effectiveness of visual action prompts for agent-agnostic control, where a single model drives the self motion and generates dynamics for agents with distinct configurations.

Specifically, we train a unified model on robotic datasets

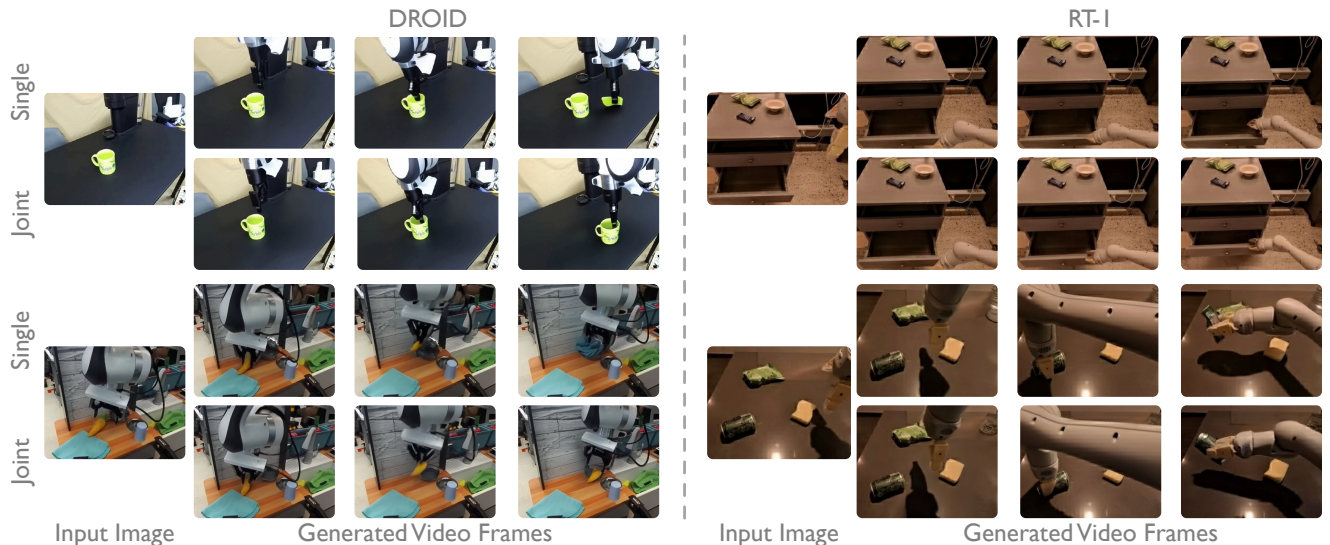


Figure 5. **Effect of joint training with visual action prompts.** Compared to single-dataset training, joint training leads to better object consistency on DROID [35] and enables held-out skill execution (e.g., “close the drawer”) on RT-1 [11].

(RT-1 [11], DROID [35]) and human egocentric video data (EgoVid [64], based on Ego4D [26]). Conventional precise action representations struggle with multi-domain joint training: DROID and RT-1 feature incompatible robot configurations, and it is hard to map agent-centric signals like end-effector poses to DROID’s random third-person camera viewpoints. EgoVid presents even greater challenges with dynamic egocentric views and complex human hand actions, which cannot be handled by previous action representations.

As shown in Tab. 1, joint training with visual action prompts achieves comparable or superior performance in generation quality and dynamic accuracy compared to single-domain training. Dynamic-centric evaluations on DROID in Tab. 2 further confirm its superiority. Fig. 5 highlights its specific advantages: joint training improves object consistency in DROID manipulations and enables novel skill generalization (e.g., closing drawers) on RT-1’s unseen skill subset, which single-domain models fail to achieve. We attribute these benefits to visual action prompts simplifying the learning objective — models bypass learning mapping from abstract action representations to agent motion and focus directly on learning interaction-driven dynamics induced by actions. We present more results of the unified model in Fig. 6. We further illustrate in Fig. 7 that our model can generate diverse interaction outcomes aligned with different action trajectories under identical initial images, indicating its potential for downstream applications like simulation, planning, and robotic policy evaluation.

### 4.3. Ablation studies

**Different forms of visual action prompts.** In Fig. 8 and Tab. 3, we compare different visual action prompt forms:

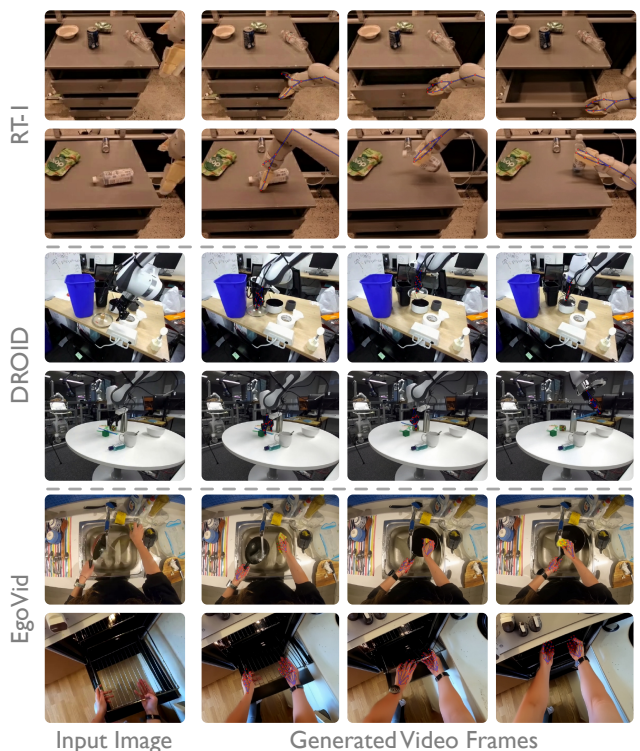


Figure 6. **Action-to-video generation of a unified model.** Visual action prompts enable joint training on diverse datasets and facilitate interaction-driven dynamic generation. The overlaid skeletons are only for visualization, demonstrating accurate action control.

mesh rendering, depth maps, and our primary skeleton-based approach. As shown in Tab. 3, visual prompts with more de-






Figure 7. **Diverse generation results under different action trajectories.** Our model can simulate diverse actions and their visual outcomes from the same initial frame.

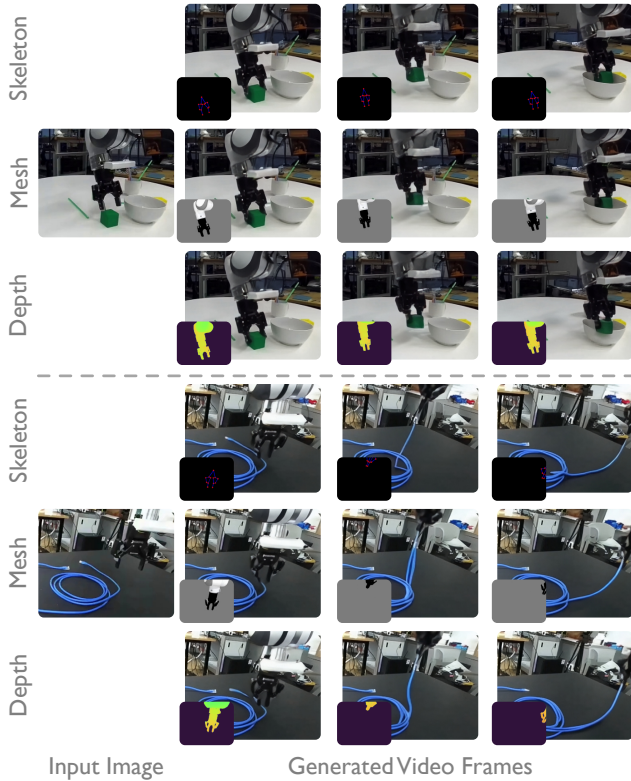


Figure 8. **Comparison of different forms of visual action prompts.** All three forms of visual action prompts can precisely represent delicate actions and drive plausible interactions. We prefer the skeleton-based prompt for its acquisition efficiency.

tails (mesh/depth) improve generation quality and dynamic accuracy compared to skeletons. Fig. 8 demonstrates that

Control Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FVD $\downarrow$	ST-IoU $\uparrow$
Skeleton	21.26	0.834	0.132	141.8	0.450
Mesh	<b>23.51</b>	<b>0.859</b>	<b>0.106</b>	120.4	<b>0.586</b>
Depth	23.41	0.858	<b>0.106</b>	<b>119.7</b>	0.581

Table 3. **Ablation study on different visual action prompts.** Representations with more details (mesh, depth) perform marginally better. Skeleton is preferred for its acquisition efficiency.

Model Variant	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FVD $\downarrow$	ST-IoU $\uparrow$
w/o ControlNet	20.19	0.819	0.151	165.2	0.396
w/o Main Branch	21.09	0.830	0.138	146.9	0.442
<b>Ours (full)</b>	<b>21.26</b>	<b>0.834</b>	<b>0.132</b>	<b>141.8</b>	<b>0.450</b>

Table 4. **Ablation study on model architecture.**

all three forms effectively drive plausible interaction-aware dynamics. Given skeleton’s lower acquisition cost for in-the-wild data and compatibility with additional sparse 3D information, we advocate its use as a unified action representation for large-scale training. For precision-critical applications, skeleton-driven models can be fine-tuned with higher-fidelity action prompts to fully leverage mixed-quality data.

**Model architecture.** We evaluate different modules’ contributions by training models on the DROID [35] dataset. As shown in Tab. 4, ControlNet [76] plays a more critical role in both generation quality and dynamic accuracy. Injecting visual action prompts to the main DiT and utilize LoRA-based fine-tuning of the DiT backbone [30, 72] is also effective, which yields marginal gains.

## 5. Conclusion

In this paper, we propose visual action prompts as universal action representations for action-to-video generation which effectively represent complex high-DoF actions and retain the cross-domain dynamic transfer capability of video generation models at the same time. We design robust pipelines for building visual action prompts from heterogeneous data sources for training, and utilize lightweight fine-tuning to inject visual action prompts into a pretrained video generation model. Our method demonstrates improvements in both interaction fidelity and domain adaptability, with experimental results validating the model’s effectiveness.

**Limitations and future works.** Our method still faces two main limitations. First, current visual action prompts represent actions in 2D, offering limited 3D cues. Integrating additional sparse 3D information could introduce better 3D awareness. Moreover, the base model is pre-trained on text-to-video tasks where motion is explicitly specified through texts, which hasn’t been effectively utilized. Future works could adapt the attention between video-text tokens to video-action tokens, injecting action control more effectively.

**Acknowledgement.** This work was partially supported by the Major Program of Xiangjiang Laboratory (No. 24XJJ-CYJ01004), NSFC (NO. 62322207, NO. U24B20154), and Information Technology Center and State Key Lab of CAD&CG, Zhejiang University.

## References

- [1] Easymocap - make human motion capture easier. Github, 2021. 2
- [2] Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, Yin Cui, Yifan Ding, et al. Cosmos world foundation model platform for physical ai. *arXiv preprint arXiv:2501.03575*, 2025. 2
- [3] Rick Akkerman, Haiwen Feng, Michael J Black, Dimitrios Tzionas, and Victoria Fernández Abrevaya. Interdyn: Controllable interactive dynamics with video diffusion models. *arXiv preprint arXiv:2412.11785*, 2024. 2, 3
- [4] Eloi Alonso, Adam Jelley, Vincent Micheli, Anssi Kanervisto, Amos J Storkey, Tim Pearce, and François Fleuret. Diffusion for world modeling: Visual details matter in atari. *NeurIPS*, 2025. 2
- [5] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 5
- [6] Prithviraj Banerjee, Sindi Shkodrani, Pierre Moulon, Shreyas Hampali, Fan Zhang, Jade Fountain, Edward Miller, Selen Basol, Richard Newcombe, Robert Wang, Jakob Julian Engel, and Tomas Hodan. Introducing hot3d: An egocentric dataset for 3d hand and object tracking, 2024. 3
- [7] Omer Bar-Tal, Hila Chefer, Omer Tov, Charles Herrmann, Roni Paiss, Shiran Zada, Ariel Ephrat, Junhwa Hur, Guanghui Liu, Amit Raj, et al. Lumiere: A space-time diffusion model for video generation. In *SIGGRAPH Asia 2024 Conference Papers*, 2024. 1
- [8] Ilya Baran and Jovan Popović. Automatic rigging and animation of 3d characters. *ACM TOG*, 2007. 2
- [9] Homanga Bharadhwaj, Debidatta Dwibedi, Abhinav Gupta, Shubham Tulsiani, Carl Doersch, Ted Xiao, Dhruv Shah, Fei Xia, Dorsa Sadigh, and Sean Kirmani. Gen2act: Human video generation in novel scenarios enables generalizable robot manipulation. *arXiv preprint arXiv:2409.16283*, 2024. 2
- [10] Kevin Black, Mitsuhiko Nakamoto, Pranav Atreya, Homer Walke, Chelsea Finn, Aviral Kumar, and Sergey Levine. Zero-shot robotic manipulation with pretrained image-editing diffusion models. *arXiv preprint arXiv:2310.10639*, 2023. 2
- [11] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022. 1, 2, 3, 4, 5, 6, 7
- [12] Jake Bruce, Michael D Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, et al. Genie: Generative interactive environments. In *ICML*, 2024. 1, 2
- [13] Géry Casiez, Nicolas Roussel, and Daniel Vogel. 1€ filter: a simple speed-based low-pass filter for noisy input in interactive systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2012. 4
- [14] Haoxuan Che, Xuanhua He, Quande Liu, Cheng Jin, and Hao Chen. Gamegen-x: Interactive open-world game video generation. *arXiv preprint arXiv:2411.00769*, 2024. 1, 2
- [15] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The epic-kitchens dataset. In *ECCV*, 2018. 3
- [16] Etched Decart, Spruce Campbell, Quinn McIntyre, Xinlei Chen, and Julian Quevedo. Oasis: A universe in a transformer, 2024. 2
- [17] DeepMind. Genie 2: A large-scale foundation world model, 2024. 2
- [18] Junting Dong, Qing Shuai, Yuanqing Zhang, Xian Liu, Xiaowei Zhou, and Hujun Bao. Motion capture from internet videos. In *ECCV*, 2020. 2
- [19] Yilun Du, Mengjiao Yang, Bo Dai, Hanjun Dai, Ofir Nachum, Joshua B. Tenenbaum, Dale Schuurmans, and Pieter Abbeel. Learning universal policies via text-guided video generation, 2023. 2
- [20] Zicong Fan, Omid Taheri, Dimitrios Tzionas, Muhammed Kocabas, Manuel Kaufmann, Michael J Black, and Otmar Hilliges. Arctic: A dataset for dexterous bimanual hand-object manipulation. In *CVPR*, 2023. 3
- [21] Ruili Feng, Han Zhang, Zhantao Yang, Jie Xiao, Zhilei Shu, Zhiheng Liu, Andy Zheng, Yukun Huang, Yu Liu, and Hongyang Zhang. The matrix: Infinite-horizon world generation with real-time moving control. *arXiv preprint arXiv:2412.03568*, 2024. 1, 2
- [22] Rao Fu, Dingxi Zhang, Alex Jiang, Wanjia Fu, Austin Funk, Daniel Ritchie, and Srinath Sridhar. Gigahands: A massive annotated dataset of bimanual hand activities, 2024. 3
- [23] Shenyuan Gao, Jiazhi Yang, Li Chen, Kashyap Chitta, Yihang Qiu, Andreas Geiger, Jun Zhang, and Hongyang Li. Vista: A generalizable driving world model with high fidelity and versatile controllability. *NeurIPS*, 2025. 2
- [24] Daniel Geng, Charles Herrmann, Junhwa Hur, Forrester Cole, Serena Zhang, Tobias Pfaff, Tatiana Lopez-Guevara, Carl Doersch, Yusuf Aytar, Michael Rubinstein, et al. Motion prompting: Controlling video generation with motion trajectories. *arXiv preprint arXiv:2412.02700*, 2024. 3
- [25] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzyńska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Freund, Peter Yianilos, Moritz Mueller-Freitag, Florian Hoppe, Christian Thureau, Ingo Bax, and Roland Memisevic. The "something something" video database for learning and evaluating visual common sense, 2017. 3
- [26] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *CVPR*, 2022. 3, 7
- [27] Kristen Grauman, Andrew Westbury, Lorenzo Torresani, Kris Kitani, Jitendra Malik, Triantafyllos Afouras, Kumar

- Ashutosh, Vijay Baiyya, Siddhant Bansal, Bikram Boote, et al. Ego-exo4d: Understanding skilled human activity from first- and third-person perspectives. In *CVPR*, 2024. 3
- [28] Xingyi He, Hao Yu, Sida Peng, Dongli Tan, Zehong Shen, Hujun Bao, and Xiaowei Zhou. Matchanything: Universal cross-modality image matching with large-scale pre-training. *arXiv preprint arXiv:2501.07556*, 2025. 4
- [29] Radu Horaud and Fadi Dornaika. Hand-eye calibration. *The international journal of robotics research*, 1995. 2
- [30] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *ICLR*, 2022. 5, 8
- [31] Li Hu. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In *CVPR*, 2024. 2, 3
- [32] Li Hu, Guangyuan Wang, Zhen Shen, Xin Gao, Dechao Meng, Lian Zhuo, Peng Zhang, Bang Zhang, and Liefeng Bo. Animate anyone 2: High-fidelity character image animation with environment affordance. *arXiv preprint arXiv:2502.06145*, 2025. 3
- [33] Nikita Karaev, Iurii Makarov, Jianyuan Wang, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Co-tracker3: Simpler and better point tracking by pseudo-labelling real videos. *arXiv preprint arXiv:2410.11831*, 2024. 5
- [34] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset, 2017. 3
- [35] Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis, et al. Droid: A large-scale in-the-wild robot manipulation dataset. *arXiv preprint arXiv:2403.12945*, 2024. 1, 2, 3, 4, 5, 6, 7, 8
- [36] Rawal Khirodkar, Timur Bagautdinov, Julieta Martinez, Su Zhaoen, Austin James, Peter Selednik, Stuart Anderson, and Shunsuke Saito. Sapiens: Foundation for human vision models. In *ECCV*, 2024. 4
- [37] Po-Chen Ko, Jiayuan Mao, Yilun Du, Shao-Hua Sun, and Joshua B Tenenbaum. Learning to act from actionless videos through dense correspondences. *arXiv preprint arXiv:2310.08576*, 2023. 2
- [38] John P Lewis, Matt Cordner, and Nickson Fong. Pose space deformation: a unified approach to shape interpolation and skeleton-driven deformation. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*. 2023. 2
- [39] Hanwen Liang, Junli Cao, Vidit Goel, Guocheng Qian, Sergei Korolev, Demetri Terzopoulos, Konstantinos N Plataniotis, Sergey Tulyakov, and Jian Ren. Wonderland: Navigating 3d scenes from a single image. *arXiv preprint arXiv:2412.12091*, 2024. 5
- [40] Yun Liu, Haolin Yang, Xu Si, Ling Liu, Zipeng Li, Yuxiang Zhang, Yebin Liu, and Li Yi. Taco: Benchmarking generalizable bimanual tool-action-object understanding. In *CVPR*, 2024. 3
- [41] Bria Long, Violet Xiang, Stefan Stojanov, Robert Z Sparks, Zi Yin, Grace E Keene, Alvin WM Tan, Steven Y Feng, Chengxu Zhuang, Virginia A Marchman, et al. The babyview dataset: High-resolution egocentric videos of infants' and young children's everyday experiences. *arXiv preprint arXiv:2406.10447*, 2024. 3
- [42] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, et al. Mediapipe: A framework for building perception pipelines. *arXiv preprint arXiv:1906.08172*, 2019. 4
- [43] Yifang Men, Yuan Yao, Miaomiao Cui, and Liefeng Bo. Mimo: Controllable character video synthesis with spatial decomposed modeling. *arXiv preprint arXiv:2409.16160*, 2024. 3
- [44] Saman Motamed, Laura Culp, Kevin Swersky, Priyank Jaini, and Robert Geirhos. Do generative video models understand physical principles? *arXiv preprint arXiv:2501.09038*, 2025. 5
- [45] Lucas Mourot, Ludovic Hoyet, François Le Clerc, François Schnitzler, and Pierre Hellier. A survey on deep learning for skeleton-based human animation. In *Computer Graphics Forum*, 2022. 2
- [46] OpenAI. Openai sora, 2023. 1
- [47] Abby O'Neill, Abdul Rehman, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, et al. Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0. In *ICRA*, 2024. 3
- [48] Rick Parent. *Computer animation: algorithms and techniques*. 2012. 2
- [49] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *ICCV*, 2023. 5
- [50] Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, David Yan, Dhruv Choudhary, Dingkang Wang, Geet Sethi, Guan Pang, Haoyu Ma, Ishan Misra, Ji Hou, Jialiang Wang, Kiran Jagadeesh, Kunpeng Li, Luxin Zhang, Mannat Singh, Mary Williamson, Matt Le, Matthew Yu, Mitesh Kumar Singh, Peizhao Zhang, Peter Vajda, Quentin Duval, Rohit Girdhar, Roshan Sumbaly, Sai Saketh Rambhatla, Sam Tsai, Samaneh Azadi, Samyak Datta, Sanyuan Chen, Sean Bell, Sharadh Ramaswamy, Shelly Sheynin, Siddharth Bhattacharya, Simran Motwani, Tao Xu, Tianhe Li, Tingbo Hou, Wei-Ning Hsu, Xi Yin, Xiaoliang Dai, Yaniv Taigman, Yaqiao Luo, Yen-Cheng Liu, Yi-Chiao Wu, Yue Zhao, Yuval Kirstain, Zecheng He, Zijian He, Albert Pumarola, Ali Thabet, Artsiom Sanakoyeu, Arun Mallya, Baishan Guo, Boris Araya, Breena Kerr, Carleigh Wood, Ce Liu, Cen Peng, Dmitry Vengertsev, Edgar Schonfeld, Elliot Blanchard, Felix Juefei-Xu, Fraylie Nord, Jeff Liang, John Hoffman, Jonas Kohler, Kaolin Fire, Karthik Sivakumar, Lawrence Chen, Licheng Yu, Luya Gao, Markos Georgopoulos, Rashel Moritz, Sara K. Sampson, Shikai Li, Simone Parmeggiani, Steve Fine, Tara Fowler, Vladan Petrovic, and Yuming Du. Movie gen: A cast of media foundation models, 2025. 1



- [51] Rolandos Alexandros Potamias, Jinglei Zhang, Jiankang Deng, and Stefanos Zafeiriou. Wilor: End-to-end 3d hand localization and reconstruction in-the-wild. *arXiv preprint arXiv:2409.12259*, 2024. 4
- [52] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 6
- [53] Zehong Shen, Huaijin Pi, Yan Xia, Zhi Cen, Sida Peng, Zechen Hu, Hujun Bao, Ruizhen Hu, and Xiaowei Zhou. World-grounded human motion recovery via gravity-view coordinates. In *SIGGRAPH Asia 2024 Conference Papers*, 2024. 2
- [54] Xiaoyu Shi, Zhaoyang Huang, Fu-Yun Wang, Weikang Bian, Dasong Li, Yi Zhang, Manyuan Zhang, Ka Chun Cheung, Simon See, Hongwei Qin, Jifeng Dai, and Hongsheng Li. Motion-i2v: Consistent and controllable image-to-video generation with explicit motion modeling, 2024. 3
- [55] Tomáš Souček, Dima Damen, Michael Wray, Ivan Laptev, and Josef Sivic. Genhowto: Learning to generate actions and state transformations from instructional videos. In *CVPR*, 2024. 2
- [56] Klaus H Strobl and Gerd Hirzinger. Optimal hand-eye calibration. In *2006 IEEE/RSJ international conference on intelligent robots and systems*, 2006. 2
- [57] Sruthi Sudhakar, Ruoshi Liu, Basile Van Hoorick, Carl Vondrick, and Richard Zemel. Controlling the world by sleight of hand. In *ECCV*, 2024. 2
- [58] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *European conference on computer vision*, pages 402–419. Springer, 2020. 5
- [59] Roger Y Tsai, Reimar K Lenz, et al. A new technique for fully autonomous and efficient 3 d robotics hand/eye calibration. *IEEE Transactions on robotics and automation*, 1989. 2
- [60] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphaël Marinier, Marcin Michalski, and Sylvain Gelly. Fvd: A new metric for video generation. *ICLR workshop*, 2019. 5
- [61] Dani Valevski, Yaniv Leviathan, Moab Arar, and Shlomi Fruchter. Diffusion models are real-time game engines. *arXiv preprint arXiv:2408.14837*, 2024. 1, 2
- [62] 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. 3
- [63] Luozhou Wang, Ziyang Mai, Guibao Shen, Yixun Liang, Xin Tao, Pengfei Wan, Di Zhang, Yijun Li, and Yingcong Chen. Motion inversion for video customization. *arXiv preprint arXiv:2403.20193*, 2024. 3
- [64] Xiaofeng Wang, Kang Zhao, Feng Liu, Jiayu Wang, Guosheng Zhao, Xiaoyi Bao, Zheng Zhu, Yingya Zhang, and Xingang Wang. Egovid-5m: A large-scale video-action dataset for egocentric video generation. *arXiv preprint arXiv:2411.08380*, 2024. 1, 2, 3, 5, 6, 7
- [65] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 2004. 5
- [66] Zhouxia Wang, Ziyang Yuan, Xintao Wang, Yaowei Li, Tianshui Chen, Menghan Xia, Ping Luo, and Ying Shan. Motionctrl: A unified and flexible motion controller for video generation. In *ACM SIGGRAPH 2024 Conference Papers*, 2024. 3
- [67] Jialong Wu, Shaofeng Yin, Ningya Feng, Xu He, Dong Li, Jianye Hao, and Mingsheng Long. ivideogpt: Interactive videogpts are scalable world models. *NeurIPS*, 2025. 2
- [68] Cheng-Yen Yang, Hsiang-Wei Huang, Wenhao Chai, Zhongyu Jiang, and Jenq-Neng Hwang. Samurai: Adapting segment anything model for zero-shot visual tracking with motion-aware memory, 2024. 4
- [69] Mengjiao Yang, Yilun Du, Kamyar Ghasemipour, Jonathan Tompson, Dale Schuurmans, and Pieter Abbeel. Learning interactive real-world simulators. *arXiv preprint arXiv:2310.06114*, 2023. 2
- [70] Sherry Yang, Jacob Walker, Jack Parker-Holder, Yilun Du, Jake Bruce, Andre Barreto, Pieter Abbeel, and Dale Schuurmans. Video as the new language for real-world decision making, 2024. 2
- [71] Zhendong Yang, Ailing Zeng, Chun Yuan, and Yu Li. Effective whole-body pose estimation with two-stages distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 4
- [72] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu 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. 1, 2, 5, 6, 8
- [73] 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. 3
- [74] Jiwen Yu, Yiran Qin, Xintao Wang, Pengfei Wan, Di Zhang, and Xihui Liu. Gamefactory: Creating new games with generative interactive videos. *arXiv preprint arXiv:2501.08325*, 2025. 2
- [75] Xinyu Zhan, Lixin Yang, Yifei Zhao, Kangrui Mao, Hanlin Xu, Zenan Lin, Kailin Li, and Cewu Lu. Oakink2: A dataset of bimanual hands-object manipulation in complex task completion. In *CVPR*, 2024. 3
- [76] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. 5, 8
- [77] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 5
- [78] Yuang Zhang, Jiaxi Gu, Li-Wen Wang, Han Wang, Junqi Cheng, Yuefeng Zhu, and Fangyuan Zou. Mimicmotion: High-quality human motion video generation with confidence-aware pose guidance. *arXiv preprint arXiv:2406.19680*, 2024. 3
- [79] Zhenghao Zhang, Junchao Liao, Menghao Li, Zuozhuo Dai, Bingxue Qiu, Siyu Zhu, Long Qin, and Weizhi Wang. Torat: Trajectory-oriented diffusion transformer for video generation. *arXiv preprint arXiv:2407.21705*, 2024. 3
- [80] Jingkai Zhou, Benzhi Wang, Weihua Chen, Jingqi Bai, Dongyang Li, Aixi Zhang, Hao Xu, Mingyang Yang, and

- Fan Wang. Realisdance: Equip controllable character animation with realistic hands, 2024. [3](#)
- [81] Fangqi Zhu, Hongtao Wu, Song Guo, Yuxiao Liu, Chilam Cheang, and Tao Kong. Irasim: Learning interactive real-robot action simulators. *arXiv preprint arXiv:2406.14540*, 2024. [2](#), [5](#), [6](#)
- [82] Shenhao Zhu, Junming Leo Chen, Zuozhuo Dai, Qingkun Su, Yinghui Xu, Xun Cao, Yao Yao, Hao Zhu, and Siyu Zhu. Champ: Controllable and consistent human image animation with 3d parametric guidance, 2024. [3](#)