

# MolmoMotion

## Forecasting Point Trajectories in 3D with Language Instruction

Jianing Zhang<sup>♥1,2\*</sup> Chenhao Zheng<sup>♥1,2\*</sup> Yajun Yang<sup>♥2</sup>

Max Argus<sup>1</sup> Rustin Soraki<sup>1,2</sup> Winson Han<sup>1</sup> Taira Anderson<sup>1</sup> Chun-Liang Li<sup>2</sup> Shuo Liu<sup>1,2</sup>  
Jiafei Duan<sup>1,2</sup>

Zhongzheng Ren<sup>♥1,2,3</sup> Jieyu Zhang<sup>♥1,2</sup> Ranjay Krishna<sup>♥1,2</sup>

<sup>1</sup>Allen Institute for AI <sup>2</sup>University of Washington <sup>3</sup>UNC-Chapel Hill

\* Equal contribution. ♥ Core contributors.

🤖 **Models:** MolmoMotion

📊 **Data:** MolmoMotion-1M PointMotionBench

🔄 **Code:** <https://github.com/allenai/molmo-motion>

🌐 **Websites:** Technical-Website Blog


### Abstract



Motion forecasting is central to visual intelligence: agents must anticipate how objects will move in order to plan actions, reason about physical interactions, and synthesize realistic futures. We argue that 3D points in world coordinates provide a general representation that is class-agnostic, view-stable, compact, and directly useful for downstream tasks. We formalize the task of *goal-conditioned 3D point motion forecasting*: given a short visual history, a set of 3D query points on an object of interest, and a language description of the intended goal, the model predicts the future 3D trajectory of each point. We introduce a full stack to study this task at scale: (1) **MolmoMotion-1M** is a large corpus of action-described, object-grounded 3D point trajectory dataset annotated from 1.16M unconstrained videos; (2) **PointMotionBench** is a human-verified benchmark spanning 111 object categories and 61 motion types; and (3) **MolmoMotion** is a general motion forecasting model that supports both autoregressive coordinate prediction and flow-matching-based trajectory generation. MolmoMotion is able to accurately predicts diverse motion patterns with different language instructions, and significantly outperforms all existing motion prediction baselines on PointMotionBench. Finally, we show that the learned 3D motion prior transfers well to downstream applications: it improves training efficiency and generalization for robot manipulation, and its predicted trajectories provide effective motion guidance for generative models to synthesize videos with more realistic object motion.

## 1 Introduction

Psychologists like James J. Gibson have long argued that motion is core to perception, hypothesizing that motion informs an observer how they and other objects move through space, explains object occlusion and permanence, and identifies affordances [27]. In the 70s, Ullman formalized motion perception as a computational problem [66], with Lucas and Kanade providing an algorithm for estimating motion to enable tracking as optical flow [49]. Although such methods have improved [63, 20, 37], they remain primarily focused on estimating motion that has *already* occurred. Many real-world applications instead require *forecasting* how motion will unfold. In robotics, an agent must anticipate how its actions will move objects through the scene [25, 7]. In video generation, realistic synthesis requires precise forecasting of the future motion of



**Figure 1 Overview.** We introduce the task of goal-conditioned 3D point motion prediction. Given initial 3D query points on an object, history RGB observations, and a language description of the future action, our model predicts the future 3D positions of all queried points in a metric world coordinate frame. We show that pretraining this motion prediction task produces a transferable motion representation for downstream applications, including robotics planning and video generation.

objects [1, 11].

Building a general motion forecasting model imposes several requirements on the representation used for motion. First, the representation should be *class-agnostic*: it should not depend on templates tailored to humans, hands, rigid objects, or any other fixed categories. Second, it should be *view-stable*: the same underlying motion should be represented consistently across different cameras, from static surveillance footage to egocentric robot videos and moving outdoor platforms. Third, it should expose *physical structure* in a form that downstream systems can use directly. Existing approaches of motion prediction only partially satisfy these requirements. While *pixels* provide a rich rendering of possible futures [67, 1, 11], they are expensive to generate and often difficult to utilize directly in downstream applications. Many applications instead require explicit geometric and physical quantities, such as object pose [78, 68, 13] or particle-level dynamics [44, 60]. *Parametric 3D models* (for humans [14], hands [8], or rigid objects [61]) are useful for such applications but remain restricted to specific categories and embodiments. *2D point trajectories* are more category-agnostic, but image-plane coordinates entangle object motion with camera ego-motion and viewpoint change, making them difficult to disentangle from videos and transfer across domains [7, 64].

We argue that *object-attached 3D points in world coordinates* provide a suitable representation for general motion forecasting. A sparse set of surface points moving through 3D space can describe the motion of rigid, articulated, or deformable objects without assuming a category-specific template. Because 3D points share a world frame, the same physical motion can remain stable across cameras, viewpoints, and capture settings. Additionally, many task-relevant quantities can be expressed as the motion of 3D points, from the pose change of a robot gripper [13] to particle trajectories in physical simulation [44, 60]. Finally, by focusing only on object-attached points, the prediction target remains compact while still capturing the motion relevant for physical interaction.

We formalize this idea as the task of *goal-conditioned 3D point motion prediction* (Fig. 1). Given a short history of visual observations, a set of initial 3D query points on an object of interest, and a language description of the intended goal, the model predicts the future 3D trajectory of each query point over time. The language instruction disambiguates among plausible futures and reduces the search space of possible future states.

We present the full stack needed to study this task at scale: **a scalable data curation pipeline, a novel prediction model, and a human-verified benchmark**. The first challenge is 3D motion supervision data. Existing video datasets with 3D capture are small and domain-limited [5, 41, 47], while internet videos provide scale and diversity but lack 3D annotations. We therefore develop an automatic annotation pipeline for robustly extracting object-grounded 3D point trajectories from unconstrained videos. Applying this pipeline to roughly 1.16M video clips yields MolmoMotion-1M, the largest corpus of action-described, object-grounded 3D point trajectories. We additionally introduce PointMotionBench, a benchmark for 3D motion forecasting spanning

111 object categories and 61 motion types. The benchmark uses ground-truth 3D capture where available and human-verified 3D tracks otherwise, providing a reliable testbed for evaluating motion prediction.

We introduce MolmoMotion, a general motion prediction model pretrained on the MolmoMotion-1M dataset. We train two complementary classes of motion prediction models. The first predicts future trajectories autoregressively as coordinate sequences; the second predicts future motion with flow matching as a continuous trajectory distribution. Experiments show MolmoMotion accurately forecast diverse motion types and generalize across a wide range of scenes and instructions, and significantly outperforms all existing motion prediction methods in PointMotionBench.

Finally, we validate 3D point forecasting as a useful task: MolmoMotion transfers effectively to downstream applications. Because object trajectories in 3D world space are largely embodiment-agnostic [7, 19], the motion prior learned from internet-scale human video can transfer naturally to robotics planning. We show that this prior improves sample efficiency, closed-loop success, and generalization to unseen objects and scenes in the MolmoSpaces pick-and-place task [40], and adapts efficiently to real-world robot videos from DROID [38]. MolmoMotion can also guide video generation. When its predicted trajectories are used to condition a lightweight 3D-track-conditioned image-to-video model [29], the generated videos exhibit more realistic motion than those produced directly by much larger image-to-video models [67], while also improving multiple quantitative video-quality metrics.

## 2 MolmoMotion

We formulate general motion prediction as the task of predicting future 3D trajectories of points attached to an object of interest, conditioned on an action description. In this section, we formulate the problem (Sec. 2.1) and introduce our model architecture (Sec. 2.2).

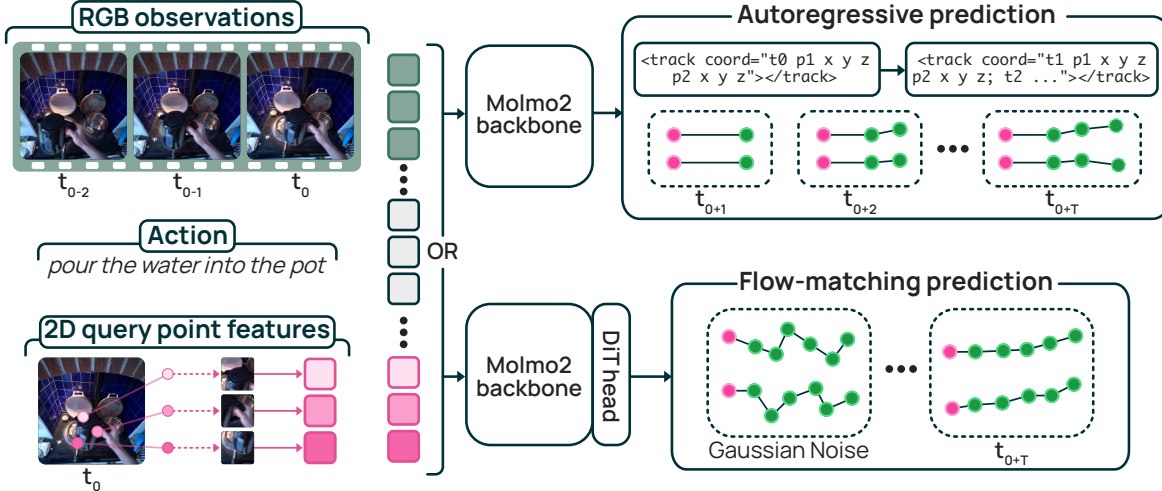
### 2.1 Problem formulation

We define the model input as follows. At a reference time  $t_0$ , the model is given  $N$  user-specified 2D query points on an object of interest in the image,  $\{\mathbf{q}_{t_0}^n \in \mathbb{R}^2\}_{n=1}^N$ . It is also given the corresponding initial 3D positions  $\{\mathbf{p}_{t_0}^n \in \mathbb{R}^3\}_{n=1}^N$ , expressed in the camera coordinate frame at  $t_0$ . In practice, these 3D positions can be obtained by lifting the 2D query points using estimated or measured depth together with known camera intrinsics. The model also receives a short history of RGB observations,  $I_{t_s:t_0} = \{I_{t_s}, \dots, I_{t_0}\}$ , and a language description  $a$  of the intended action. The goal is to predict the future 3D positions of all query points, over a horizon of  $T$  steps:  $\{\{\hat{\mathbf{p}}_t^n \in \mathbb{R}^3\}_{t=t_0+1}^{t_0+T}\}_{n=1}^N$ . All future coordinates are expressed in a world coordinate frame anchored at the camera at time  $t_0$ . This choice makes the prediction independent of future camera motion.

### 2.2 MolmoMotion architecture

We implement two classes of trajectory predictors: one with an autoregressive objective and the other with a flow-matching objective. They share the same input encoding of the RGB observation, action description, and 2D query-point visual features. They differ in how the initial 3D query coordinates are encoded and how the future 3D trajectories are decoded (Fig. 2). We study both objectives because they offer complementary modeling biases: autoregressive decoding conditions each prediction on the previously generated trajectory, which encourages smooth temporal evolution, while flow-matching models a distribution over future trajectories and is better suited for capturing motion uncertainty [15].

**Input encoding.** Predicting object motion from a language instruction first requires grounding the relevant object in the image and understanding the instruction. We therefore use Molmo2 [17] as the vision-language backbone for input processing, leveraging its strong object-grounding capability. Given the RGB observation history  $I_{t_s:t_0}$ , the Molmo2 vision encoder produces image tokens  $\mathcal{T}_{\text{img}}$ . The action description  $a$  is tokenized into language tokens  $\mathcal{T}_{\text{text}}$ . To condition on the query points, we additionally insert one visual point token for each 2D query coordinate. Let  $F_{t_0}$  be the anchor-frame feature map produced by the vision encoder. For each query point  $\mathbf{q}_{t_0}^n$ , we bilinearly sample the anchor-frame feature map  $F_{t_0}$  at its 2D location to obtain a point feature  $\mathbf{e}_{\text{pt}}^n$ . These features form the point-token sequence  $\mathcal{T}_{\text{pt}} = \{\mathbf{e}_{\text{pt}}^1, \dots, \mathbf{e}_{\text{pt}}^N\}$ . We concatenate the image, text, and point tokens as  $\mathcal{C} = [\mathcal{T}_{\text{img}}, \mathcal{T}_{\text{text}}, \mathcal{T}_{\text{pt}}]$  and process them with the Molmo2 language model component.



**Figure 2 MolmoMotion architecture.** The shared input to Molmo2 [17] backbone consists of image tokens of RGB observations, text tokens of action description, and 2D query point feature tokens sampled from Molmo2 vision encoder. The autoregressive variant encodes the initial 3D query coordinates and decodes future trajectories as quantized coordinate text, while the flow-matching variant represents them directly in continuous 3D coordinate space.

For both prediction variants, we represent 3D coordinates relative to the first query point at  $t_0$ . Let  $\mathbf{p}_{\text{anc}} = \mathbf{p}_{t_0}^1$  denote this anchor point. For each point  $n$  and time  $t$ , our models represent 3d point coordinates as  $\delta_t^n = \mathbf{p}_t^n - \mathbf{p}_{\text{anc}}$ . All 3D point coordinates are in metric scale, in units of *meter*.

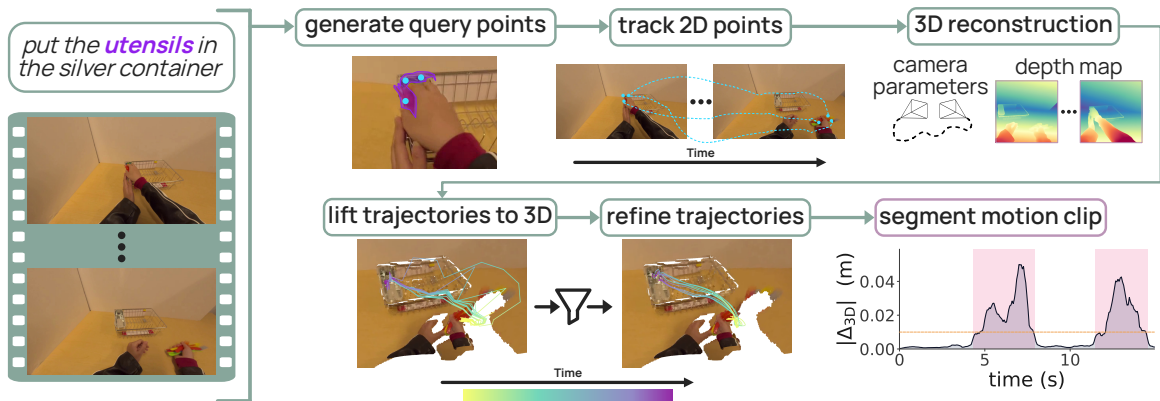
**Autoregressive objective.** The autoregressive variant follows Molmo2’s coordinate representation, encoding both the initial 3D query coordinates and the future trajectory as structured text. Each anchor-relative coordinate is discretized into *millimeter* bins ( $\bar{\delta}_t^n = \text{round}(1000 \delta_t^n)$ ) and serialized as timestamped point-coordinate tuples. Fig. 2 shows an example of the coordinate encoding format. The input prompt contains the visual-language conditioning  $\mathcal{C}$  together with the serialized initial query coordinates  $\{\bar{\delta}_{t_0}^n\}_{n=1}^N$ . The output is the serialized future trajectory  $y_{1:L}$ , generated in temporal order. An example of text pieces encoded and decoded can be found at Fig. 2. We train the autoregressive model with the standard next-token objective. At inference time, the model decodes the trajectory string autoregressively and the coordinates are parsed back into  $\hat{\mathbf{P}}_{t_0+1:t_0+T}$ . Since coordinates are emitted in temporal order, each future timestamp is conditioned on all earlier generated coordinates, giving the model a direct mechanism for modeling temporal dependence and smooth rollouts.

**Flow-matching objective.** The flow-matching variant predicts future trajectories in continuous 3D coordinate space. Following a MolmoBot-style design [18], we use a DiT decoder [55] conditioned on Molmo2 features from all layers, with lightweight MLPs for coordinate encoding and decoding. We concatenate the clean initial 3D query coordinates with a noised version of the future coordinates and project them into point tokens. We use RoPE along both the point and time axes to distinguish point identity and timestamp, analogous to the Multimodal RoPE used in Qwen2.5-VL [62, 4].

The decoder is trained with the standard flow-matching objective [46]. Specifically, we sample Gaussian noise with the same shape as the future trajectory, linearly interpolate between the noise and the ground-truth trajectory, and train the decoder to predict the corresponding velocity field from the noised trajectory to the clean trajectory. At inference time, the model starts from Gaussian noise and integrates the learned velocity field with 10 Euler steps to obtain the predicted future trajectory.

### 3 MolmoMotion-1M and PointMotionBench

Existing 3D point-track datasets are small, domain-limited, and often lack object grounding or language annotations. We therefore build an automatic annotation pipeline that extracts object-grounded 3D point trajectories from unconstrained videos. Applying it to 1.16M public videos yields **MolmoMotion-1M**, the



**Figure 3 Overview of data annotation pipeline.** Given a video of an action event and its description, we first ground the moving object and sample query points on it. We then track dense 2D points on the object, lift these tracks into a shared metric 3D frame, and use object-level spatial and temporal consistency priors to filter unreliable trajectories. Finally, we clip the video around intervals where the grounded object undergoes meaningful motion.

largest action-described, object-grounded 3D point trajectory dataset. We also introduce **PointMotionBench**, a held-out benchmark with verified 3D point trajectories for object-centric 3D motion forecasting.

### 3.1 MolmoMotion-1M Annotation

We start with public video datasets that provide action descriptions or task captions [31, 56, 59, 36]. Our goal is to localize the described object, select query points, and track them in 3D world coordinates. The pipeline consists of five stages, shown in Fig. 3: semantic object grounding, temporal point correspondence, metric 3D lifting, trajectory-level filtering, and video-level clipping.

**Semantic object grounding.** Given an action description  $a$ , we use an LLM [70] to extract the manipulated or moving entity as a short object phrase. Instead of prompting SAM3 [12] to directly produce an object mask from this phrase, we first use MolmoPoint [16] to localize the entity as a 2D point in the reference frame  $t_0$ . We find MolmoPoint more reliable for vague phrases like “an object on top of the table,” because the prompt can specify the object by its action, e.g., the object being moved, rather than by appearance alone. We then use the point prompt with SAM3 to obtain the object mask  $M_{t_0}$ , and sample  $N$  query points  $\{\mathbf{q}_{t_0}^n\}_{n=1}^N$  inside  $M_{t_0}$  using K-means cluster centers.

**2D point tracking and metric 3D lifting.** We next establish point correspondence across frames by tracking the query points through the video. We run AllTracker [30] to obtain temporally persistent 2D tracks  $\{\mathbf{q}_t^n\}_{t=1}^L$  and visibility masks  $\{m_t^n\}_{t=1}^L$ . To lift these tracks into 3D, we run ViPE [33] on the monocular video to estimate per-frame metric depth and camera geometry. We empirically find that this paradigm produces more accurate 3D trajectories than current end-to-end 3D point trackers such as SpatialTrackerV2 [69]. ViPE also provides metric scale output, allowing the resulting trajectories to be expressed in physical 3D units rather than a arbitrary relative scale. Using the estimated depth, intrinsics, and camera poses, we back-project each visible 2D track point into a world frame anchored at the first-frame camera, producing metric 3D tracks  $\{\tilde{\mathbf{p}}_t^n \in \mathbb{R}^3\}_{t=1}^L$ .

**Trajectory-level filtering and smoothing.** We empirically find that some lifted trajectories can be corrupted by 2D tracking drift, depth noise, and camera estimation errors. We therefore both filter outlier tracks and smooth remaining tracks using the object-level prior that sampled points should move coherently as parts of the same physical entity. We remove tracks with consistently low trust using a MAD-based outlier criterion [42]. For retained tracks, we follow the smoothing algorithm in Stereo4D [36] to smooth their depth value along each camera ray. More details can be found in Appendix B.4. This removes high-frequency depth jitter in annotated 3D tracks.

**Video-level clipping.** Event-level video clips often contain long static intervals before or after the described motion, which provide little supervision for learning dynamics. We therefore re-clip videos around intervals



Paradigm	Model	Inputs		HOT3D [5]			WorldTrack [24]			DAVIS [57]		
		Frames	Text	ADE↓	FDE↓	PWT↑	ADE↓	FDE↓	PWT↑	ADE↓	FDE↓	PWT↑
Non-parametric	Static	1	✗	0.180	0.316	0.293	0.167	0.317	0.390	2.281	4.360	0.085
	Extrapolate	3	✗	0.159	0.309	0.351	0.184	0.432	0.436	2.683	5.741	0.104
Pixel-space	Wan2.2-5B [67]	1	✓	0.200	0.308	0.253	0.852	1.046	0.090	3.074	5.192	0.051
	Cosmos Predict [1]	5	✓	0.225	0.294	0.199	0.831	0.988	0.072	4.191	6.368	0.033
3D model	ObjectForesight [61]	3	✗	<u>0.129</u>	<u>0.192</u>	0.353	–	–	–	–	–	–
	EgoScaler [75]	1	✓	0.170	<b>0.179</b>	0.218	–	–	–	–	–	–
	Robot4DGen [48]	3	✗	0.212	0.271	0.112	0.548	0.704	0.121	2.120	3.382	0.081
2D track	Track2Act [7]	1	✗	0.294	0.413	0.202	1.230	1.567	0.053	4.853	8.110	0.018
<b>3D track (Ours)</b>	<b>MolmoMotion-FM</b>	1	✓	0.183	0.311	0.286	0.165	0.305	0.401	1.380	2.205	<u>0.165</u>
	<b>MolmoMotion-FM</b>	3	✓	0.135	0.255	<u>0.382</u>	0.158	0.295	<u>0.438</u>	1.480	2.520	0.130
	<b>MolmoMotion-AR</b>	1	✓	0.157	0.290	0.303	<u>0.148</u>	<u>0.269</u>	0.424	<b>1.146</b>	<b>1.843</b>	<b>0.199</b>
	<b>MolmoMotion-AR</b>	3	✓	<b>0.109</b>	0.217	<b>0.444</b>	<b>0.143</b>	<b>0.261</b>	<b>0.445</b>	<u>1.227</u>	<u>2.108</u>	0.153

**Table 1 3D point trajectory prediction on PointMotionBench.** We report 3D ADE, FDE, and average PWT in meters. Input columns indicate the number of past observations consumed and whether the model accepts a natural-language action description. *MolmoMotion-AR* denotes our autoregressive variant, while *MolmoMotion-FM* denotes the flow-matching variant.

further include DAVIS [57] to cover outdoor dynamic scenes; for this split, we run our annotation pipeline and manually verify the correctness of each resulting trajectory. In total, the benchmark contains **742** clips spanning **111** object categories and **61** action/motion types. Fig. 4(c) and Fig. 5(c) show the per-cohort distributions of motion verbs and manipulated objects across PointMotionBench.

## 4 Experiments

This section evaluates MolmoMotion on 3D point trajectory prediction (Sec. 4.1) and on two downstream transfer tasks: robotic planning (Sec. 4.2) and trajectory-guided video generation (Sec. 4.3). Model ablations are presented in Appendix E.

### 4.1 3D point motion forecasting

We first evaluate MolmoMotion and existing motion prediction methods on the 3D point trajectory prediction task in PointMotionBench.

**MolmoMotion implementation.** MolmoMotion uses the pretrained 4B Molmo2 [17] as its VLM backbone. Training proceeds in two stages. In the first stage, we randomly sample a start timestamp  $t_0$  from each video clip and sample  $N = 8$  query points from one object. The model is supervised to predict  $T = 8$  future timestamps at 15 fps, giving 64 future 3D point targets per example. We train with history length  $H = 3$  in the first stage, providing a short visual history before  $t_0$ . This stage runs for 40K steps. In the second stage, we continue training for 10K steps while increasing the prediction horizon to  $T = 32$ . In this stage we train two variants of models, with  $H = 3$  and  $H = 1$  respectively, to support both short visual history and single frame history settings.

**Baselines.** Baselines fall into four families. **Non-parametric baselines** include *Static*, which keeps each point fixed at its initial 3D position, and *Extrapolate*, which estimates a constant velocity from the history frames and linearly extrapolates into the future. **Pixel-space methods** first generate future RGB video and then recover query-point trajectories using the 3D tracking pipeline from Sec. 3.1; we compare against Wan2.2-I2V-5B [67] and Cosmos-Predict-2.5 [1]. **Parametric 3D prediction methods** predict an intermediate 3D representation; we include methods whose output can be converted to 3D point motion. ObjectForesight [61] and EgoScaler [75] forecast future 6-DoF object poses under a rigid-object assumption, and we use their predicted pose sequence to transform all query points. Robot4DGen predicts future scene flow [48], from which we extract query point



**Figure 6 Qualitative MolmoMotion prediction.** MolmoMotion predicts accurate motion trajectories on diverse motion patterns with different action instructions.

trajectories. **2D point-track methods** predict future image-plane trajectories; we evaluate Track2Act by lifting its predicted 2D tracks to 3D using ground-truth depth [7]. Note that PointWorld [34] and MotionForecast [64] are also closely related works, but their models have not yet been released.


**Evaluation setting and metrics.** We evaluate all methods in PointMotionBench. Each method predicts future motion for up to 2 seconds at 15 fps; if the ground truth clip has a shorter valid future horizon, we evaluate only on the available frames. Since baselines vary in the number of input frames  $H$  and whether they accept text condition, we follow each baseline’s native setting where applicable. Following prior work [64], we use best-of-5 evaluation for each sample. We follow the point-track forecasting metrics used in [64], adapted from 2D pixel distance to 3D metric distance. In our setting, one unit corresponds to one meter in 3D world coordinates. ADE measures the mean displacement error across all visible query points and predicted timesteps, while FDE measures the final-timestep displacement error. PWT is the average fraction of predicted points within  $\{0.01, 0.02, 0.05, 0.10, 0.20\}$  meters of the ground truth.

**Results.** Tab. 1 shows the quantitative results. Note that ObjectForesight [61] and EgoScaler [75] are evaluated only on the HOT3D subset because they require object mesh inputs, which are available only for HOT3D. MolmoMotion outperforms prior methods by a large margin in almost all subsets of PointMotionBench, with the autoregressive variant achieving the strongest overall performance. The autoregressive model performs better than the flow-matching model under deterministic trajectory metrics, likely because conditioning on the previously generated coordinate sequence encourages temporally smooth predictions. Using three RGB observation frames generally improves over the single-frame setting, as additional history provides velocity cues for future motion. A notable finding is that simple non-parametric baselines are competitive with, or stronger than, several learned baselines, including pixel-space video prediction methods, suggesting that visually plausible RGB futures do not necessarily recover accurate metric point motion. We also show qualitative examples in Fig. 6, where MolmoMotion accurately predicts 3D motion trajectories across diverse motion patterns and language instructions.

## 4.2 MolmoMotion transfers effectively to robotics planning

The representation learned by MolmoMotion is well-suited for robotics planning. The intuition is that object motion in 3D is more transferable than embodiment-specific actions. A human hand and a robot gripper execute differently, but successful manipulation often produces similar object trajectories in 3D space.

We use the pick-and-place task as a controlled transfer setting, focusing on the post-grasp stage where the policy must lift, transport and place the object at the correct target. We train two MolmoBot policies [18] with the same flow-matching action head and 20K released episodes, differing only in backbone initialization:



**Figure 7 MolmoMotion transfers effectively to robotics planning.** (a) Task success rate on the MolmoSpaces Franka Pick-and-place benchmark over robot-finetuning steps (left) and the final step per split breakdown (right) across SS (seen scene, seen object), SU (seen scene, unseen object), US (unseen scene, seen object), and All. (b) Test L2 error of predicted 3D trajectories on held-out DROID clips. In both settings, initializing from MolmoMotion substantially outperforms the baseline.



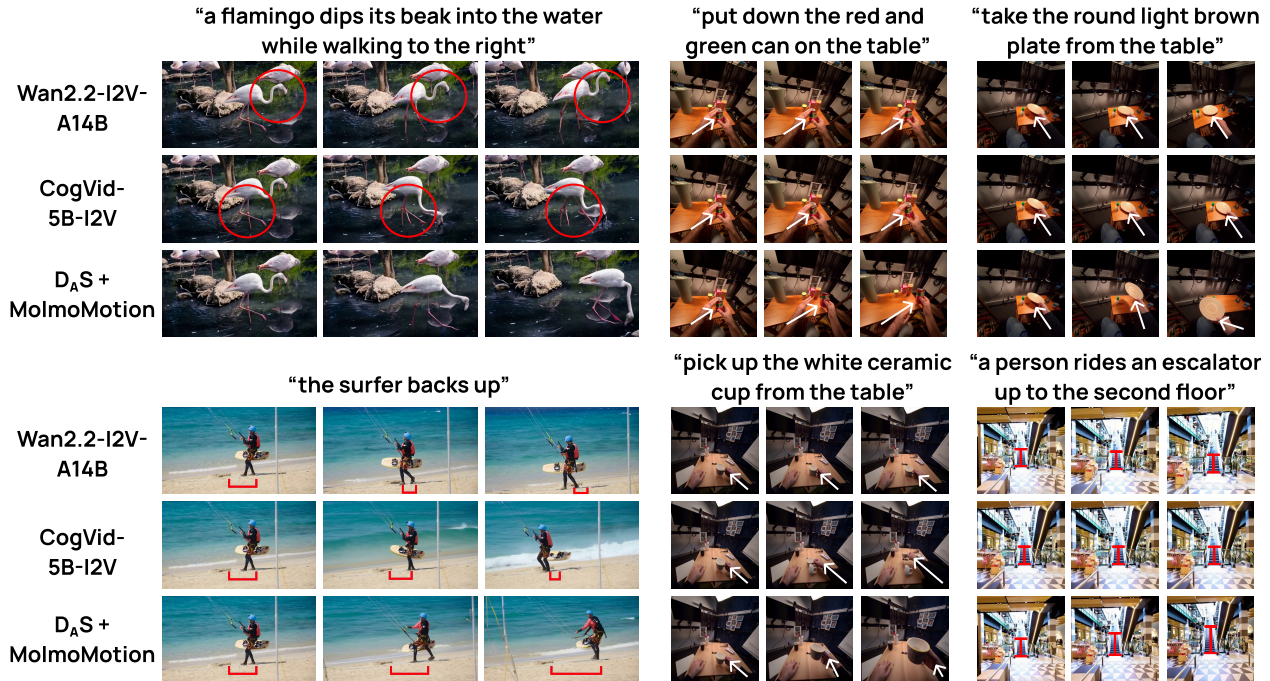
**Figure 8 MolmoMotion trajectory prediction on real-world scenarios.** Predicted future 3D object trajectories on a held-out scenario after finetuning MolmoMotion on DROID videos. MolmoMotion can plan accurate object trajectories for various robotic manipulation tasks.

Molmo2 pretrained weights vs. MolmoMotion-AR. During evaluation, the original MolmoBot policy performs grasping, after which control is handed to our trained policy. On FrankaPickandPlaceDroidMiniBench in MolmoSpaces [40], we report closed-loop success across seen/unseen scene and object splits. As shown in Fig. 7a, MolmoMotion initialization substantially improves training efficiency and final performance: success reaches 51% at 10K steps vs. 19% for Molmo2, and final average success increases from 56.0% to 76.3%. The smaller drop in unseen-object and unseen-scene splits further suggests MolmoMotion improves downstream generalization.

We further evaluate whether the MolmoMotion can be adapted to real-world robot scenarios. We finetune MolmoMotion on DROID single-camera videos [38] using the same 3D object point trajectory prediction task. Fig. 8 shows qualitative examples of object trajectory planned by MolmoMotion across different scenes, objects, and tasks in DROID held out videos. We also compare against training the same architecture with Molmo2 initialization. As shown in Fig. 7b, MolmoMotion starts with substantially lower trajectory error and reaches the best performance much quicker. This suggests that the motion prior learned by MolmoMotion can be efficiently adapted to real-world robot data. We leave closed-loop real-robot evaluation to future work.

### 4.3 MolmoMotion provides controllable motion for video generation

A natural question is whether the 3D point trajectories predicted by MolmoMotion can serve as an explicit motion-control signal for downstream video generation. Given an input image and action description, we first obtain query points on the manipulated object and estimate their initial 3D positions using the pipeline in Sec. 3.1. MolmoMotion then predicts their future 3D trajectories, which we use to guide DaS [29], a 3D point-trajectory-guided image-to-video model built on CogVideoX-5B [73]. We compare against caption-conditioned image-to-video generators without explicit motion guidance, including DaS’s base model CogVideoX-5B and the larger Wan2.2-I2V-A14B [67]. As shown in Tab. 2, DaS+MolmoMotion improves over CogVideoX-5B on all metrics and outperforms Wan2.2-I2V-A14B on four out of five metrics. Qualitatively (Fig. 9), it produces more physically plausible object motion, especially for fine-grained manipulation. It also follows the prompted



**Figure 9 Qualitative Video Generation Comparison.** MolmoMotion-guided videos exhibit more physically plausible object motion and follow the prompted action more faithfully.

Method	Tem-Con	Subj-Cons	M-Smooth	Dyn-Deg	Bg-Cons
CogVideoX-5B [73]	0.964	0.939	<u>0.988</u>	0.861	0.941
Wan-14B [67]	<u>0.965</u>	<u>0.940</u>	0.983	<b>0.908</b>	<u>0.947</u>
DAS [29] + MolmoMotion	<b>0.968</b>	<b>0.950</b>	<b>0.990</b>	<u>0.876</u>	<b>0.948</b>

**Table 2 Video generation quality on PointMotionBench videos.** We report metrics in VBenCh [35] that evaluates motion: temporal consistency, subject consistency, motion smoothness, dynamic degree, and background consistency.

action more faithfully. These results suggest that MolmoMotion predicts 3D trajectories that can serve as an effective control interface for downstream video generation.

## 5 Related work

**Motion prediction models.** Motion prediction models differ mainly in the representation they forecast. Pixel-space methods cast prediction as conditional video generation, from latent-action world models [52, 32] to recent large-scale video generators [11, 1, 67]; these models are general but spend substantial capacity rendering appearance, lighting, and camera motion. Latent prediction methods forecast learned feature states [77, 2], avoiding pixels but producing encoder-tied representations that are difficult to use as physical quantities. Recent point-trajectory forecasting methods predict category-agnostic motion directly [7, 64], but operate in 2D image space, where object motion is entangled with camera motion and viewpoint change. 3D forecasting methods predict physically grounded states such as human motion [50], rigid-object pose [75, 61], or scene-level 3D motion [34], but are often tied to specific object categories or domains. We predict object-attached 3D point trajectories in world coordinates, yielding a category-agnostic and physically grounded representation.

**3D point tracking.** Our data annotation pipeline builds on recent advances in 3D point tracking from monocular video. 2D point trackers estimate temporally persistent pixel locations for queried points [20, 21], with recent models improving long-range tracking [22], occlusion reasoning [37], and dense correspondence over entire videos [30]. Recent monocular reconstruction methods lift image observations into 3D by estimating camera motion and per-frame depth, enabling pixel tracks to be lifted to 3D [33, 45]. There are also direct 3D point

trackers that track points in 3D space [69, 24, 76]. Together, these methods have made it increasingly feasible to recover 3D motion from ordinary RGB videos, though the focus is on tracking observed frames rather than forecasting future motion.

**Human-to-robot transfer.** Prior work transfers non-robot video to robot control through different abstractions. Some methods re-target human hands, endpoints, or skill trajectories as planning scaffolds [3, 13, 23, 78, 6, 43]. Others train generalist vision-language-action policies on cross-embodiment robot data [10, 54, 53, 39, 9], sometimes with human video, but still predict embodiment-specific actions. Closest to ours are methods that learn motion priors or controls from unlabeled video through latent actions or inverse dynamics [51, 28, 26, 74, 72]. In contrast, we use 3D world-frame object point trajectories as the pretraining target itself, yielding a motion prior that is both embodiment-agnostic and directly usable for downstream robot learning.

## 6 Limitations

MolmoMotion requires multiple forward passes to predict dense point tracks on an object, since pretraining uses only 8 query points per example due to Molmo2 context length limit in stage 2 training. This sparse point setting is not sufficient to densely represent object geometry, limiting the model’s understanding of fine-grained structure and complex deformable motion. In addition, more downstream evaluations are needed to fully establish the effectiveness of this motion pretraining task, such as closed-loop real-world robot experiments.

## 7 Conclusions

We introduce MolmoMotion, a language-guided motion predictor that forecasts future trajectories of object-attached points in a 3D world frame. We built MolmoMotion-1M for pretraining and PointMotionBench for 3D motion evaluation. Experiments show that MolmoMotion significantly outperforms prior motion prediction methods and transfers to robot manipulation and video generation.

## Acknowledgements

This work would not be possible without the support of our colleagues at Ai2. We thank David Albright, Stephen Kelman, Kristin Cha, Byron Bischoff, David Everhart, Jon Borchardt, Kyle Wiggers, Will Smith, Peter Clark, Dieter Fox, and Noah Smith for their important work for the MolmoMotion public release. We thank Ropedia for providing access to the Xperience dataset used in this work, and granting permission the release of MolmoMotion under the Apache License 2.0. Chenhao Zheng is partially funded through an Apple grant. We thank Oncel Tuzel, Pavan Kumar, and Rick Chang for the helpful discussion and support on this project.

## References

- [1] N. Agarwal, A. Ali, M. Bala, Y. Balaji, E. Barker, T. Cai, P. Chattopadhyay, Y. Chen, Y. Cui, Y. Ding, et al. Cosmos world foundation model platform for physical ai. *arXiv preprint arXiv:2501.03575*, 2025.
- [2] M. Assran, A. Bardes, D. Fan, Q. Garrido, R. Howes, M. Komeili, M. Muckley, A. Rizvi, C. Roberts, K. Sinha, A. Zholus, S. Arnaud, A. Gejji, A. Martin, F. Robert Hogan, D. Dugas, P. Bojanowski, V. Khalidov, P. Labatut, F. Massa, M. Szafraniec, K. Krishnakumar, Y. Li, X. Ma, S. Chandar, F. Meier, Y. LeCun, M. Rabbat, and N. Ballas. V-jepa 2: Self-supervised video models enable understanding, prediction and planning. Technical report, FAIR at Meta, 2025.
- [3] S. Bahl, R. Mendonca, L. Chen, U. Jain, and D. Pathak. Affordances from human videos as a versatile representation for robotics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [4] S. Bai, K. Chen, X. Liu, J. Wang, W. Ge, S. Song, K. Dang, P. Wang, S. Wang, J. Tang, et al. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- [5] P. Banerjee, S. Shkodrani, P. Moulon, S. Hampali, S. Han, F. Zhang, L. Zhang, J. Fountain, E. Miller, S. Basol, R. Newcombe, R. Wang, J. J. Engel, and T. Hodan. HOT3D: Hand and object tracking in 3D from egocentric multi-view videos. *CVPR*, 2025.
- [6] H. Bharadhwaj, D. Dwibedi, A. Gupta, S. Tulsiani, C. Doersch, T. Xiao, D. Shah, F. Xia, D. Sadigh, and S. Kirmani. Gen2act: Human video generation in novel scenarios enables generalizable robot manipulation. *arXiv preprint arXiv:2409.16283*, 2024.
- [7] H. Bharadhwaj, R. Mottaghi, A. Gupta, and S. Tulsiani. Track2act: Predicting point tracks from internet videos enables generalizable robot manipulation. In *European Conference on Computer Vision (ECCV)*, 2024.
- [8] H. Bi, L. Wu, T. Lin, H. Tan, Z. Su, H. Su, and J. Zhu. H-rdt: Human manipulation enhanced bimanual robotic manipulation, 2025. URL <https://arxiv.org/abs/2507.23523>.
- [9] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, S. Levine, A. Li-Bell, M. Mothukuri, S. Nair, K. Pertsch, L. X. Shi, J. Tanner, Q. Vuong, A. Walling, H. Wang, and U. Zhilinsky.  $\pi_0$ : A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.
- [10] K. Bousmalis, G. Vezzani, D. Rao, C. Devin, A. X. Lee, M. Bauza, T. Davchev, Y. Zhou, A. Gupta, A. Raju, A. Laurens, C. Fantacci, V. Dalibard, M. Zambelli, M. Martins, R. Pevcevičute, M. Blokzijl, M. Denil, N. Batchelor, T. Lampe, E. Parisotto, K. Żołna, S. Reed, S. G. Colmenarejo, J. Scholz, A. Abdolmaleki, O. Groth, J.-B. Regli, O. Sushkov, T. Rothörl, J. E. Chen, Y. Aytar, D. Barker, J. Ortiz, M. Riedmiller, J. T. Springenberg, R. Hadsell, F. Nori, and N. Heess. Robocat: A self-improving generalist agent for robotic manipulation. *arXiv preprint arXiv:2306.11706*, 2023.
- [11] J. Bruce, M. Dennis, A. Edwards, J. Parker-Holder, Y. Shi, E. Hughes, M. Lai, A. Mavalankar, R. Steigerwald, C. Apps, Y. Aytar, S. Bechtel, F. Behbahani, S. Chan, N. Heess, L. Gonzalez, S. Osindero, S. Ozair, S. Reed, J. Zhang, K. Zolna, J. Clune, N. de Freitas, S. Singh, and T. Rocktäschel. Genie: Generative interactive environments. *arXiv preprint arXiv:2402.15391*, 2024.
- [12] N. Carion, L. Gustafson, Y.-T. Hu, S. Debnath, R. Hu, D. Suris, C. Ryali, K. V. Alwala, H. Khedr, A. Huang, J. Lei, T. Ma, B. Guo, A. Kalla, M. Marks, J. Greer, M. Wang, P. Sun, R. Rädle, T. Afouras, E. Mavroudi, K. Xu, T.-H. Wu, Y. Zhou, L. Momeni, R. Hazra, S. Ding, S. Vaze, F. Porcher, F. Li, S. Li, A. Kamath, H. K. Cheng, P. Dollár, N. Ravi, K. Saenko, P. Zhang, and C. Feichtenhofer. Sam 3: Segment anything with concepts. *arXiv preprint arXiv:2511.16719*, 2025.
- [13] H. Chen, B. Sun, A. Zhang, M. Pollefeys, and S. Leutenegger. Vidbot: Learning generalizable 3d actions from in-the-wild 2d human videos for zero-shot robotic manipulation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 27661–27672, 2025.
- [14] L.-H. Chen, J. Zhang, Y. Li, Y. Pang, X. Xia, and T. Liu. Humanmac: Masked motion completion for human motion prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9544–9555, 2023.
- [15] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song. Diffusion policy: Visuomotor policy learning via action diffusion, 2024. URL <https://arxiv.org/abs/2303.04137>.

- [16] C. Clark, Y. Yang, J. S. Park, Z. Ma, J. Zhang, R. Tripathi, M. Salehi, S. Lee, T. Anderson, W. Han, et al. Molmopoint: Better pointing for vlms with grounding tokens. *arXiv preprint arXiv:2603.28069*, 2026.
- [17] C. Clark, J. Zhang, Z. Ma, J. S. Park, M. Salehi, R. Tripathi, S. Lee, Z. Ren, C. D. Kim, Y. Yang, V. Shao, Y. Yang, W. Huang, Z. Gao, T. Anderson, J. Zhang, J. Jain, G. Stoica, W. Han, A. Farhadi, and R. Krishna. Molmo2: Open weights and data for vision-language models with video understanding and grounding. *arXiv preprint arXiv:2601.10611*, 2026.
- [18] A. Deshpande, M. Guru, R. Hendrix, S. Jauhri, A. Eftekhari, R. Tripathi, M. Argus, J. Salvador, H. Fang, M. Wallingford, W. Pumacay, Y. Kim, Q. Pfeifer, Y.-C. Lee, P. Wolters, O. Rayyan, M. Zhang, J. Duan, K. Farley, W. Han, E. VanderBilt, D. Fox, A. Farhadi, G. Chalvatzaki, D. Shah, and R. Krishna. Molmobot: Large-scale simulation enables zero-shot manipulation. *arXiv preprint arXiv:2603.16861*, 2026.
- [19] K. Dharmarajan, W. Huang, J. Wu, L. Fei-Fei, and R. Zhang. Dream2flow: Bridging video generation and open-world manipulation with 3d object flow. *arXiv preprint arXiv:2512.24766*, 2025.
- [20] C. Doersch, A. Gupta, L. Markeeva, A. Recasens, L. Smaira, Y. Aytar, J. Carreira, A. Zisserman, and Y. Yang. Tap-vid: A benchmark for tracking any point in a video. *Advances in Neural Information Processing Systems*, 35: 13610–13626, 2022.
- [21] C. Doersch, Y. Yang, M. Vecerik, D. Gokay, A. Gupta, Y. Aytar, J. Carreira, and A. Zisserman. Tapir: Tracking any point with per-frame initialization and temporal refinement. *arXiv preprint arXiv:2306.08637*, 2023.
- [22] C. Doersch, P. Luc, Y. Yang, D. Gokay, S. Koppula, A. Gupta, J. Heyward, I. Rocco, R. Goroshin, J. Carreira, and A. Zisserman. Bootstap: Bootstrapped training for tracking-any-point. *arXiv preprint arXiv:2402.00847*, 2024.
- [23] H. Fang, J. Duan, D. Clay, S. Wang, S. Liu, W. Huang, X. Fan, W.-C. Tsai, S. Chen, Y. R. Wang, et al. Molmoact2: Action reasoning models for real-world deployment. *arXiv preprint arXiv:2605.02881*, 2026.
- [24] H. Feng, J. Zhang, Q. Wang, Y. Ye, P. Yu, M. J. Black, T. Darrell, and A. Kanazawa. St4rtrack: Simultaneous 4d reconstruction and tracking in the world. *arXiv preprint arXiv:2504.13152*, 2025.
- [25] C. Finn and S. Levine. Deep visual foresight for planning robot motion. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 2786–2793. IEEE, 2017.
- [26] Q. Garrido, T. Nagarajan, B. Terver, N. Ballas, Y. LeCun, and M. Rabbat. Learning latent action world models in the wild. *arXiv preprint arXiv:2601.05230*, 2026.
- [27] J. Gibson. *The Ecological Approach to Visual Perception*. Resources for ecological psychology. Lawrence Erlbaum Associates, 1986. ISBN 9780898599596. URL <https://books.google.com/books?id=DrhCCWmJpWUC>.
- [28] R. G. Goswami, A. Bar, D. Fan, T.-Y. Yang, G. Zhou, P. Krishnamurthy, M. Rabbat, F. Khorrani, and Y. LeCun. World models for learning dexterous hand-object interactions from human videos. *arXiv preprint arXiv:2512.13644*, 2025.
- [29] Z. Gu, R. Yan, J. Lu, P. Li, Z. Dou, C. Si, Z. Dong, Q. Liu, C. Lin, Z. Liu, W. Wang, and Y. Liu. Diffusion as shader: 3d-aware video diffusion for versatile video generation control. *arXiv preprint arXiv:2501.03847*, 2025.
- [30] A. W. Harley, Y. You, X. Sun, Y. Zheng, N. Raghuraman, Y. Gu, S. Liang, W.-H. Chu, A. Dave, P. Tokmakov, S. You, R. Ambrus, K. Fragkiadaki, and L. J. Guibas. Alltracker: Efficient dense point tracking at high resolution. *arXiv preprint arXiv:2506.07310*, 2025.
- [31] R. Hoque, P. Huang, D. J. Yoon, M. Sivapurapu, and J. Zhang. Egodex: Learning dexterous manipulation from large-scale egocentric video. *arXiv preprint arXiv:2505.11709*, 2025.
- [32] A. Hu, L. Russell, H. Yeo, Z. Murez, G. Fedoseev, A. Kendall, J. Shotton, and G. Corrado. Gaia-1: A generative world model for autonomous driving. *arXiv preprint arXiv:2309.17080*, 2023.
- [33] J. Huang, Q. Zhou, H. Rabeti, A. Korovko, H. Ling, X. Ren, T. Shen, J. Gao, D. Slepichev, C.-H. Lin, J. Ren, K. Xie, J. Biswas, L. Leal-Taixé, and S. Fidler. Vipe: Video pose engine for 3d geometric perception. *arXiv preprint arXiv:2508.10934*, 2025.
- [34] W. Huang, Y.-W. Chao, A. Mousavian, M.-Y. Liu, D. Fox, K. Mo, and L. Fei-Fei. Pointworld: Scaling 3d world models for in-the-wild robotic manipulation. *arXiv preprint arXiv:2601.03782*, 2026.

- [35] Z. Huang, Y. He, J. Yu, F. Zhang, C. Si, Y. Jiang, Y. Zhang, T. Wu, Q. Jin, N. Chanpaisit, Y. Wang, X. Chen, L. Wang, D. Lin, Y. Qiao, and Z. Liu. VBench: Comprehensive benchmark suite for video generative models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- [36] L. Jin, R. Tucker, Z. Li, D. Fouhey, N. Snavely, and A. Holynski. Stereo4d: Learning how things move in 3d from internet stereo videos. *arXiv preprint arXiv:2412.09621*, 2024.
- [37] N. Karaev, I. Makarov, J. Wang, N. Neverova, A. Vedaldi, and C. Rupprecht. Cotracker3: Simpler and better point tracking by pseudo-labelling real videos. *arXiv preprint arXiv:2410.11831*, 2024.
- [38] A. Khazatsky, K. Pertsch, S. Nair, A. Balakrishna, S. Dasari, S. Karamcheti, S. Nasiriany, M. Kumar, L. Y. Chen, K. Ellis, et al. Droid: A large-scale in-the-wild robot manipulation dataset. *arXiv preprint arXiv:2403.12945*, 2024.
- [39] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi, Q. Vuong, T. Kollar, B. Burchfiel, R. Tedrake, D. Sadigh, S. Levine, P. Liang, and C. Finn. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.
- [40] Y. Kim, W. Pumacay, O. Rayyan, M. Argus, W. Han, E. VanderBilt, J. Salvador, A. Deshpande, R. Hendrix, S. Jauhri, S. Liu, N. M. M. Shafullah, M. Guru, A. Eftekhari, K. Farley, D. Clay, J. Duan, A. Guru, P. Wolters, A. Herrasti, Y.-C. Lee, G. Chalvatzaki, Y. Cui, A. Farhadi, D. Fox, and R. Krishna. Molmospaces: A large-scale open ecosystem for robot navigation and manipulation. *arXiv preprint arXiv:2602.11337*, 2026.
- [41] S. Koppula, I. Rocco, Y. Yang, J. Heyward, J. Carreira, A. Zisserman, G. Brostow, and C. Doersch. Tapvid-3d: A benchmark for tracking any point in 3d. *Advances in Neural Information Processing Systems*, 37:82149–82165, 2024.
- [42] C. Leys, C. Ley, O. Klein, P. Bernard, and L. Licata. Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology*, 49(4):764–766, 2013.
- [43] G. Li, Y. Lyu, Z. Liu, C. Hou, J. Zhang, and S. Zhang. H2r: A human-to-robot data augmentation for robot pre-training from videos. *arXiv preprint arXiv:2505.11920*, 2025.
- [44] Y. Li, J. Wu, R. Tedrake, J. B. Tenenbaum, and A. Torralba. Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. *arXiv preprint arXiv:1810.01566*, 2018.
- [45] Z. Li, R. Tucker, F. Cole, Q. Wang, L. Jin, V. Ye, A. Kanazawa, A. Holynski, and N. Snavely. Megasam: Accurate, fast and robust structure and motion from casual dynamic videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10486–10496, 2025.
- [46] Y. Lipman, R. T. Chen, H. Ben-Hamu, M. Nickel, and M. Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- [47] Y. Liu, Y. Liu, C. Jiang, K. Lyu, W. Wan, H. Shen, B. Liang, Z. Fu, H. Wang, and L. Yi. Hoi4d: A 4d egocentric dataset for category-level human-object interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21013–21022, June 2022.
- [48] Z. Liu, S. Li, E. Cousineau, S. Feng, B. Burchfiel, and S. Song. Geometry-aware 4d video generation for robot manipulation. *arXiv preprint arXiv:2507.01099*, 2025.
- [49] B. D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI’81*, page 674–679, San Francisco, CA, USA, 1981. Morgan Kaufmann Publishers Inc.
- [50] W. Mao, M. Liu, M. Salzmann, and H. Li. Learning trajectory dependencies for human motion prediction, 2020. URL <https://arxiv.org/abs/1908.05436>.
- [51] R. Mendonca, S. Bahl, and D. Pathak. Structured world models from human videos. *arXiv preprint arXiv:2308.10901*, 2023.
- [52] V. Micheli, E. Alonso, and F. Fleuret. Transformers are sample-efficient world models. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=vhFu1Acb0xb>.
- [53] Octo Model Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, C. Xu, J. Luo, T. Kreiman, Y. Tan, L. Y. Chen, P. Sanketi, Q. Vuong, T. Xiao, D. Sadigh, C. Finn, and S. Levine. Octo: An open-source generalist robot policy. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, 2024.

- [54] A. O’Neill, A. Rehman, A. Maddukuri, A. Gupta, A. Padalkar, A. Lee, A. Pooley, A. Gupta, A. Mandlekar, A. Jain, et al. Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6892–6903. IEEE, 2024.
- [55] W. Peebles and S. Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023.
- [56] T. Perrett, A. Darkhalil, S. Sinha, O. Emara, S. Pollard, K. K. Parida, K. Liu, P. Gatti, S. Bansal, K. Flanagan, et al. Hd-epic: A highly-detailed egocentric video dataset. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 23901–23913, 2025.
- [57] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017.
- [58] N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. Ryali, T. Ma, H. Khedr, R. Rädle, C. Rolland, L. Gustafson, E. Mintun, J. Pan, K. V. Alwala, N. Carion, C.-Y. Wu, R. Girshick, P. Dollár, and C. Feichtenhofer. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024.
- [59] Ropedia AI. Xperience-10m: A large-scale egocentric multimodal dataset with structured 3d/4d annotations. <https://huggingface.co/datasets/ropedia-ai/xperience-10m>, 2026. Hugging Face dataset.
- [60] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. Battaglia. Learning to simulate complex physics with graph networks. In *International conference on machine learning*, pages 8459–8468. PMLR, 2020.
- [61] R. Soraki, H. Bharadhwaj, A. Farhadi, and R. Mottaghi. Objectforesight: Predicting future 3d object trajectories from human videos. *arXiv preprint arXiv:2601.05237*, 2026.
- [62] J. Su, Y. Lu, S. Pan, A. Murtadha, B. Wen, and Y. Liu. Roformer: Enhanced transformer with rotary position embedding, 2023. URL <https://arxiv.org/abs/2104.09864>.
- [63] Z. Teed and J. Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *European conference on computer vision*, pages 402–419. Springer, 2020.
- [64] N. Thakkar, S. Ginosar, J. Walker, J. Malik, J. Carreira, and C. Doersch. Forecasting motion in the wild. *arXiv preprint arXiv:2604.01015*, 2026.
- [65] V. A. Traag, L. Waltman, and N. J. van Eck. From Louvain to Leiden: guaranteeing well-connected communities. *Scientific Reports*, 9(1):5233, 2019.
- [66] S. Ullman. *The Interpretation of Visual Motion*. The MIT Press, 03 1979. ISBN 9780262257121. doi: 10.7551/mitpress/3877.001.0001. URL <https://doi.org/10.7551/mitpress/3877.001.0001>.
- [67] T. Wan, A. Wang, B. Ai, B. Wen, C. Mao, C.-W. Xie, D. Chen, F. Yu, H. Zhao, J. Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.
- [68] B. Wen, J. Tremblay, V. Blukis, S. Tyree, T. Müller, A. Evans, D. Fox, J. Kautz, and S. Birchfield. Bundlesdf: Neural 6-dof tracking and 3d reconstruction of unknown objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 606–617, 2023.
- [69] Y. Xiao, J. Wang, N. Xue, N. Karaev, Y. Makarov, B. Kang, X. Zhu, H. Bao, Y. Shen, and X. Zhou. Spatialtrackerv2: 3d point tracking made easy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025. URL <https://arxiv.org/abs/2507.12462>.
- [70] A. Yang, A. Li, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Gao, C. Huang, C. Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
- [71] L. Yang, Y. Fan, and N. Xu. Video instance segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5188–5197, 2019.
- [72] R. Yang, Q. Yu, Y. Wu, R. Yan, B. Li, A.-C. Cheng, X. Zou, Y. Fang, X. Cheng, R.-Z. Qiu, et al. Egovla: Learning vision-language-action models from egocentric human videos. *arXiv preprint arXiv:2507.12440*, 2025.
- [73] Z. Yang, J. Teng, W. Zheng, M. Ding, S. Huang, J. Xu, Y. Yang, W. Hong, X. Zhang, G. Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024.
- [74] S. Ye, J. Jang, B. Jeon, S. Joo, J. Yang, B. Peng, A. Mandlekar, R. Tan, Y.-W. Chao, B. Y. Lin, L. Liden, K. Lee, J. Gao, L. Zettlemoyer, D. Fox, and M. Seo. Latent action pretraining from videos. In *International Conference on Learning Representations (ICLR)*, 2025.

- [75] T. Yoshida, S. Kurita, T. Nishimura, and S. Mori. Generating 6dof object manipulation trajectories from action description in egocentric vision. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 17370–17382, 2025.
- [76] C. Zhang, G. L. Moing, S. Koppula, I. Rocco, L. Momeni, J. Xie, S. Sun, R. Sukthankar, J. K. Barral, R. Hadsell, Z. Ghahramani, A. Zisserman, J. Zhang, and M. S. M. Sajjadi. Efficiently reconstructing dynamic scenes one d4rt at a time, 2025. URL <https://arxiv.org/abs/2512.08924>.
- [77] G. Zhou, H. Pan, Y. Lecun, and L. Pinto. DINO-WM: World models on pre-trained visual features enable zero-shot planning. In A. Singh, M. Fazel, D. Hsu, S. Lacoste-Julien, F. Berkenkamp, T. Maharaj, K. Wagstaff, and J. Zhu, editors, *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 79115–79135. PMLR, 13–19 Jul 2025. URL <https://proceedings.mlr.press/v267/zhou25t.html>.
- [78] H. Zhou, J. Cao, L. Ma, X. Fang, and G. jun Qi. Traj2action: A co-denoising framework for trajectory-guided human-to-robot skill transfer, 2026. URL <https://arxiv.org/abs/2510.00491>.

# Appendix

<b>A</b>	<b>Qualitative examples</b>	18
<b>B</b>	<b>MolmoMotion-1M Data Generation Details</b>	18
	B.1 Source video corpora	18
	B.2 Recaptioning and object name extraction	18
	B.3 Semantic grounding and 3D lifting	19
	B.4 Trajectory filtering and smoothing	19
<b>C</b>	<b>PointMotionBench</b>	20
	C.1 HOT3D	20
	C.2 WorldTrack	21
	C.3 DAVIS	22
	C.4 Evaluation Protocol	22
<b>D</b>	<b>Model Implementation Details</b>	22
	D.1 Architecture	22
	D.2 Prompt format	23
	D.3 Training hyperparameters	23
	D.4 Flow-matching objective and inference	23
<b>E</b>	<b>Model Ablations</b>	24
<b>F</b>	<b>Robotics Transfer Settings</b>	25
	F.1 MolmoSpaces pick-and-place	25
	F.2 DROID trajectory finetuning	26
<b>G</b>	<b>Video generation experiment details</b>	27
	G.1 Methods compared	27
	G.2 Evaluation protocol	27

## A Qualitative examples

We provide additional qualitative results for two downstream applications of MolmoMotion: real-robot trajectory prediction on DROID, and trajectory-conditioned video generation.

**Real-robot trajectory prediction on DROID.** Fig. 10 shows MolmoMotion’s predicted future 3D trajectories on held-out DROID clips after finetuning on real-robot video. MolmoMotion plans accurate point trajectories across diverse manipulation scenes, objects, and tasks.

**Trajectory-conditioned video generation.** Fig. 11 and Fig. 12 compare videos generated by DaS [29] conditioned on MolmoMotion-predicted 3D trajectories with the unconditioned CogVideoX-5B and Wan-14B baselines on held-out PointMotionBench prompts. MolmoMotion-guided videos exhibit more physically plausible object motion, preserve manipulated-object identity better, and follow the prompted action more faithfully than the unconditioned baselines.

Together these examples illustrate that MolmoMotion’s 3D trajectories transfer to real-world robot data and serve as an effective control signal for downstream video generation, complementing the quantitative results in the main paper.

## B MolmoMotion-1M Data Generation Details

This appendix expands the annotation pipeline summarized in Sec. 3.1. Given a public video and its caption, the pipeline recovers an object phrase that names the moving object, grounds the phrase to a set of query points on the object, and lifts those points into a metric 3D world frame. We describe the source corpora (§B.1), recaptioning and object-name extraction (§B.2), object grounding and 3D lifting (§B.3), and trajectory filtering and smoothing (§B.4).

### B.1 Source video corpora

MolmoMotion-1M aggregates seven source corpora (Tab. 3) that together cover egocentric and third-person human manipulation, simulated and robot manipulation, and in-the-wild scenes. The five manipulation corpora ship with action descriptions or task templates, which we use directly. YT-VIS supplies object masks but no captions, so we re-caption each clip from the video. Stereo4D contributes short third-person clips with metric stereo depth, expanding coverage to outdoor scenes and deformable subjects.

Corpus	Motion clips	Domain	Action-description source
EgoDex	~160K	egocentric, human	VLM re-caption
HD-EPIC	~21K	egocentric, human	narrations
Xperience	~500K	3rd-person, human	metadata
MolmoSpaces	~185K	3rd-person, sim	task templates
DROID	~27K	3rd-person, robot	language instructions
YT-VIS	~2K	3rd-person, wild	VLM caption
Stereo4D	~70K	3rd-person, wild	paper captions

**Table 3** Source corpora used to construct MolmoMotion-1M.

### B.2 Recaptioning and object name extraction

For corpora that do not ship object-level captions (EgoDex, where the original task labels are too coarse to ground a specific entity, and YT-VIS), we generate a one-sentence visual description of each clip with Molmo2-8B [17]. The full 15 FPS re-encoded video is passed to Molmo2-8B. The prompt is:

Watch this video carefully. Describe the manipulation action you observe in exactly this format: [action verb] [specific object with color/material/shape] [preposition and location if present].  
Examples: "pick up red ceramic coffee mug", "place blue plastic

bottle on table". Be specific about the object -- include its color, material, and shape as you see them. Output only the short description, no extra words.

Then we extract the manipulated object as a noun phrase with Qwen3-0.6B [70].

### B.3 Semantic grounding and 3D lifting

Given an object phrase, we localize the entity, segment it, sample query points, track those points across the video, and lift the tracks to a metric 3D world frame.

**Localization with motion-aware prompting.** We first localize the object as a 2D point with MolmoPoint-8B [16], and then convert that point into a segmentation mask with SAM3 [12]. The non-trivial choice here is the prompt given to MolmoPoint. Conditioning the prompt on the agent and the action, we give the model the following prompt: *"point to the {obj} gripped and picked up by the hand"* for human-manipulation corpora, and *"track the {obj}"* for in-the-wild video. This is substantially more reliable than a bare *"point to the {obj}"* prompt. The motion cue disambiguates vague phrases like *"an object on the table"* by biasing MolmoPoint toward the moving entity rather than a static distractor that matches the phrase equally well.

**Point prompt for SAM 3.** We feed the 2D point returned by MolmoPoint to SAM3. We then sample  $N = 100$  query points per mask using  $K$ -means cluster centers on the mask pixel coordinates, so points are spread across the object surface rather than concentrated near the centroid.

**2D tracking and lifting.** We propagate query points through the video with AllTracker [30], which yields temporally persistent 2D tracks and per-frame visibility scores. We run ViPE [33] on the same video to estimate per-frame metric depth, intrinsics, and camera-to-world poses in a single pass. Each visible 2D track location is back-projected with the estimated depth and intrinsics, then transformed by the corresponding camera pose into the world frame anchored at the query-time camera.

### B.4 Trajectory filtering and smoothing

Lifted 3D trajectories are corrupted by noise. We make this prior precise as follows.

**Anchor tracks and inconsistency score.** We select a small set of anchor tracks  $\mathcal{A}$  per object (sixteen in our pipeline) that are visible throughout the clip and have low per-frame velocity, taking these as the most reliable estimate of the object’s body motion. For every other track  $n$  we measure how its distance to each anchor varies over time:

$$e_n(t) = \text{median}_{k \in \mathcal{A}} \left| \|\tilde{\mathbf{p}}_t^n - \tilde{\mathbf{p}}_t^k\|_2 - \bar{d}_{nk} \right|, \quad \bar{d}_{nk} = \text{median}_t \|\tilde{\mathbf{p}}_t^n - \tilde{\mathbf{p}}_t^k\|_2. \quad (1)$$

A rigid or near-rigid object keeps its inter-point distances roughly constant, so  $e_n(t)$  is large precisely at frames where track  $n$  has drifted in 2D, had its depth pulled to the background, or been displaced by a pose error.

**Trust weights and the choice of scale.** We convert the inconsistency score into a per-frame trust weight  $w_{t,n} = \exp(-e_n(t)/s_n)$  with a per-track scale  $s_n$ . We use  $s_n = \text{median}_t(e_n)$ , which behaves well under uniform noise while still suppressing per-frame outliers within a track.

**Spatial auto-split.** A single object phrase sometimes resolves to two physically separate instances, and mixing tracks across instances violates the rigid-object assumption behind  $e_n(t)$ . We therefore cluster points by their temporal-mean 3D position with mean-shift and run anchor selection, trust scoring, and smoothing independently per sub-cluster.

**Track-level outlier drop.** Within each sub-cluster, we drop tracks whose mean trust  $\bar{w}_n$  is a MAD-scaled z-score below the sub-cluster median.

**Depth-ray smoothing.** Surviving tracks are still noisy along the depth axis. Following Stereo4D [36], we re-parametrize the 3D point at frame  $t$  as  $\mathbf{p}_t^n = \mathbf{c}_t + \lambda_t^n \mathbf{r}_t^n$ , where  $\mathbf{c}_t$  is the camera center and  $\mathbf{r}_t^n$  is the unit

ray through the 2D track location, and we optimize the depth scalars  $\{\lambda_t^n\}$  with two competing objectives:

$$\min_{\{\lambda_t^n\}} \underbrace{\sum_{t,n} w_{t,n}^2 \|\mathbf{c}_t + \lambda_t^n \mathbf{r}_t^n - \tilde{\mathbf{p}}_t^n\|_2^2}_{\text{trust-weighted pin to lifted point}} + \beta \underbrace{\sum_n \sum_{\Delta \in \{1,3,5\}} \sum_t \|\mathbf{p}_{t+\Delta}^n - 2\mathbf{p}_t^n + \mathbf{p}_{t-\Delta}^n\|_2^2}_{\text{multi-stride acceleration penalty}}. \quad (2)$$

The pin term is gated by the trust weights, so untrusted frames are free to slide along the ray while trusted frames remain anchored; the acceleration term enforces smooth motion in 3D. Two design choices are worth explaining. First, the acceleration term is summed over multiple strides  $\Delta \in \{1, 3, 5\}$  rather than only  $\Delta = 1$ . A single-stride second difference penalizes only sharp single-frame jitter, but our depth errors include slow ramps over five to ten frames (e.g. when ViPE depth gradually bleeds onto a textureless background); penalizing acceleration at  $\Delta = 3$  and  $\Delta = 5$  in addition to  $\Delta = 1$  catches these slower drifts. Second, we optimize with first-order gradient descent rather than LBFGS. LBFGS converges faster on well-behaved tracks but diverges on near-degenerate tracks, so the more conservative optimizer is the safer default at corpus scale.

## C PointMotionBench

PointMotionBench evaluates object-centric 3D motion forecasting across three diverse datasets; Tab. 4 summarizes the per-source clip counts, resolutions, and frame rates. The following sections detail the data preparation pipelines and evaluation protocol (including metric definitions).

Dataset	Clips	Resolution	FPS
HOT3D	497	1408 × 1408	30
WorldTrack	155	split-dependent <sup>†</sup>	30
DAVIS	90	854 × 480	24

**Table 4** Source datasets used to construct PointMotionBench. <sup>†</sup>adt\_mini 512 × 512; ds\_mini 1280 × 720; po\_mini 960 × 540; pstudio\_mini 640 × 360.

**Exclusions.** We exclude 10 HOT3D clips that contain no ground-truth-visible query points at frame 0 (AllTracker crashes at initialization with an empty point set). For WorldTrack, we exclude the `tum` split entirely (no 3D tracks shipped) and the 27 `po_mini dancingroom1_3rd*` sequences that exhibit cluttered motion patterns. No DAVIS sequences are excluded. The same exclusion lists apply across all three baselines so that comparisons are always over identical clip sets, leaving 497 HOT3D, 155 WorldTrack, and 90 DAVIS clips.

**Human verification.** For HOT3D and WorldTrack, we view each clip alongside the annotated trajectories and confirm that (i) the foreground object is correctly identified and (ii) the one-sentence model-generated description accurately reflects the observed motion and action; clips that fail are re-annotated or excluded. For DAVIS, we directly wrote the descriptions during trajectory review rather than reviewing model-generated captions.

**Contamination check.** PointMotionBench contains no clip-level near-duplicates with the training data of the evaluated models. All per-clip captions are either written directly by human annotators or generated by Molmo2-8B and subsequently verified by annotators, ensuring no captions are sourced from existing dataset metadata that evaluated models may have been trained on.

### C.1 HOT3D

HOT3D [5] Aria provides 1,415 source clips with per-frame 3D object pose annotations as fitted mesh models. We sample 2,000 surface points per object at its first active frame and propagate them forward with per-frame pose transforms, yielding world-space trajectories and their corresponding pixel-space projections.

Each source clip typically contains one or two manipulated objects alongside several static ones. We split every clip into single-moving-object sub-clips by detecting per-object motion windows: for each object, we threshold a per-frame body-speed signal (median surface-point displacement) at 0.005 m/frame ( $\approx 15$  cm/s at

30 fps), bridge single-frame gaps, and drop runs shorter than 0.5 s. This yields 2,534 sub-clips, each covering exactly one continuously moving foreground object. We then uniformly subsample one sub-clip per group of five (random selection within the group), yielding 507 clips; the 10 with no ground-truth-visible query points at frame 0 are excluded, leaving 497.

## C.2 WorldTrack

WorldTrack [24] bundles four indoor motion-capture splits mixing egocentric and third-person viewpoints: `adt_mini`, `ds_mini`, `po_mini`, and `pstudio_mini` (resolutions listed in Tab. 4). Unlike HOT3D, source point tracks are not pre-assigned to objects. We augment the raw data through a four-step pipeline: (1) identify dynamic foreground points, (2) cluster them into per-object groups, (3) segment sequences into motion-coherent sub-clips, and (4) generate a one-sentence caption per clip.

Dataset	Filter type	Window(s)	Threshold
<code>adt_mini</code>	per-frame	3, 5 fr	0.25%
<code>po_mini</code>	global	—	1.0%
<code>pstudio_mini</code>	global+per-fr	3, 15 fr	1.0%
<code>ds_mini</code>	per-frame	10 fr	0.75%

**Table 5** Per-dataset motion-filter settings for dynamic point extraction of WorldTrack data.

**Dynamic point extraction.** We lift camera-space tracks into a shared world coordinate frame using per-frame extrinsics; for `pstudio_mini` (fixed camera), camera-space coordinates serve directly as world-space. A point is classified as dynamic if its world-space displacement over a look-back window exceeds a scene-normalized threshold (a per-dataset fraction of the scene’s 1st-to-99th percentile spatial extent; see Tab. 5). Masks from multiple window lengths are OR-ed together, and each active segment is extended backward to capture motion onset.

**Object clustering.** We cluster the dynamic points into per-object groups in two stages.

**Stage 1 --- Temporal clustering (Leiden [65]).** We build a pairwise affinity graph over the dynamic points. For each pair  $(i, j)$ , let  $\mathcal{T}_{co}$  be the set of co-visible frames and  $\mathbf{v}_i(t)$  the unit-normalized frame-to-frame velocity of point  $i$  at frame  $t$ , with zero-velocity frames contributing 0. The affinity is

$$W(i, j) = \frac{1}{|\mathcal{T}_{co}|} \sum_{t \in \mathcal{T}_{co}} \langle \mathbf{v}_i(t), \mathbf{v}_j(t) \rangle.$$

Pairs with no co-visible support or below a similarity threshold are zeroed out. Additionally, any pair whose 3D world-space distance exceeds a fixed fraction of scene extent at any sampled keyframe is forced apart—a single violating frame is conclusive evidence of distinct objects, since rigid-body points cannot drift far apart in 3D. We run Leiden community detection on the resulting graph and iteratively merge undersized clusters into their most temporally similar neighbor. Per-dataset minimum cluster sizes are 1 (`adt_mini`, `pstudio_mini`), 3 (`ds_mini`), and 5 (`po_mini`), set by visual inspection.

**Stage 2 --- Segmentation merging (SAM2 [58]).** Temporal clustering occasionally splits a single physical object across multiple clusters. For each cluster, we designate a representative point and sample three keyframes evenly across its active window, then query SAM 2. Two clusters are merged only when their representatives are spatially proximate and co-occur within the same SAM 2 mask in at least one keyframe without ever appearing in separate masks. The criterion is intentionally conservative: a single frame of visible separation blocks the merge, since merging distinct objects corrupts ground truth irreversibly.

**Sub-clip extraction and captioning.** We group objects with overlapping active spans, bridge gaps of fewer than three frames, split on longer gaps, and merge sub-clips shorter than 2 s into their nearest neighbor. Each sub-clip is captioned by Molmo2-8B [17], prompted to produce a one-sentence egocentric description (e.g. “*a hand moves a keyboard across the desk*”); captions are subsequently verified by human annotators.

### C.3 DAVIS

DAVIS [57] provides RGB frames and per-object segmentation masks but no 3D ground truth. We generate 3D trajectories by running the MolmoMotion-1M annotation pipeline, seeding query points from the DAVIS masks and lifting them through ViPE depth and camera pose estimation. We verify the resulting tracks and write the per-clip action descriptions directly during trajectory review.

### C.4 Evaluation Protocol

**Conditioning and future split.** All metrics are computed on future frames only. With  $T_{\text{cond}} \in \{1, 3\}$ , the model observes frames  $0, \dots, T_{\text{cond}} - 1$  and predicts frames  $T_{\text{cond}}, \dots, T - 1$ . Scoring is restricted to points visible at frame 0, as these are the only query positions supplied to the model. The evaluation mask is

$$v_{\text{eval}}[t, n] = v[t, n] \wedge v[0, n],$$

applied before every metric. Let  $\mathcal{S} = \{(t, n) : t \geq T_{\text{cond}}, v_{\text{eval}}[t, n] = 1\}$  be the set of scored (frame, point) pairs, and let  $N_{\text{eval}} = \sum_n v[0, n]$  be the per-clip count of evaluated points.

**Temporal alignment.** HOT3D and WorldTrack ground truth tracks are at 30 fps, and DAVIS ground truth tracks are at 24 fps; Wan2.2 and Cosmos predictions are at 24 fps. For HOT3D and Worldtrack, we resample these predictions to the ground truth timebase by linearly interpolating positions at source time  $t_{\text{src}} = 0.8 t_{\text{gt}}$  and nearest-neighbor-resampling visibility; only the  $\min(T_{\text{gt}}, T_{\text{resampled}})$  overlapping frames are scored. Track2Act outputs exactly 8 frames at ground truth indices  $\text{round}(i(T - 1)/7)$  for  $i = 0, \dots, 7$ ; we match ground truth frames at the same indices directly, with no interpolation.

**Metrics.** All metrics are computed in 3D world space. Let  $\hat{p}(t, n)$  and  $q(t, n)$  denote the predicted and ground-truth positions of query point  $n$  at frame  $t$  (in metres; 1 unit = 1 m throughout).

**ADE** (Average Displacement Error):

$$\text{ADE} = \frac{1}{|\mathcal{S}|} \sum_{(t,n) \in \mathcal{S}} \|\hat{p}(t, n) - q(t, n)\|_2$$

**FDE** (Final Displacement Error), evaluated at the last frame  $T - 1$ :

$$\text{FDE} = \frac{1}{|\mathcal{S}_T|} \sum_{n \in \mathcal{S}_T} \|\hat{p}(T - 1, n) - q(T - 1, n)\|_2, \quad \mathcal{S}_T = \{n : v_{\text{eval}}[T - 1, n] = 1\}.$$

**PWT** (Percentage Within Threshold) at threshold  $\delta$ :

$$\text{PWT}(\delta) = \frac{|\{(t, n) \in \mathcal{S} : \|\hat{p}(t, n) - q(t, n)\|_2 < \delta\}|}{|\mathcal{S}|}$$

We report  $\overline{\text{PWT}}$ , the mean of  $\text{PWT}(\delta)$  over  $\delta \in \{0.01, 0.02, 0.05, 0.10, 0.20\}$  m. ADE and FDE are in meters (lower is better);  $\overline{\text{PWT}} \in [0, 1]$  (higher is better).

## D Model Implementation Details

We expand the architecture and training recipe here.

### D.1 Architecture

**Vision-language backbone** We initialize from the public Molmo2-4B-Pretrain checkpoint. The vision encoder is a SigLIP2 ViT operating on  $378 \times 378$  RGB inputs at 14-pixel patch size, producing a  $27 \times 27$  grid of 1152-D patch tokens per frame. The Molmo2 connector pools the patch grid by  $3 \times 3$  and projects the pooled tokens through an MLP into the LLM hidden dimension of 2560. The language model is Qwen3-4B. All backbone parameters are trained end-to-end.

**Decoder heads** The autoregressive variant uses the unmodified Molmo2 LM head plus a small regex parser that maps the answer span back to coordinates at inference. The flow-matching variant uses a DiT trajectory expert with 36 blocks (one per LM layer). Each block applies a self-attention over the trajectory-token tensor of shape  $(N, H+T, 3)$  followed by a cross-attention whose keys and values come from the corresponding LM-layer hidden states and whose queries come from the trajectory tokens. RoPE is applied along both the point-index and the frame-index axes so the same self-attention can distinguish “the same point at different times” from “different points at the same time” without architectural specialization.

## D.2 Prompt format

The autoregressive variant serializes everything as a single multimodal prompt. Image tokens for the anchor frame and the  $H$  history frames are inserted by Molmo2’s video preprocessor at the front; the textual portion follows. With  $N = 8$  and  $T = 32$ , an example prompt is

```
Predict the future 3D point coordinates of 8 points over 32 timestamps,
given action: ‘open the drawer’,
2d history point features: ‘<points anchor 1 <2d_feat_start><|point_feat|><2d_feat_end>
                             2 <2d_feat_start><|point_feat|><2d_feat_end>
                             ... 8 <2d_feat_start><|point_feat|><2d_feat_
end>/>’,
and history 3d point coordinates:
‘<tracks coords=‘0.0 1 0 0 0 2 12 -3 0 ... 8 5 4 -1’>3d object history</tracks>’.
```

The supervised answer span is

```
<tracks coords=‘1.0 1 4 -8 2 2 17 -10 5 ... 8 9 7 -3;
                2.0 1 8 -16 4 ...;
                ...
                32.0 1 ...’>3d object trajectories</tracks>
```

Each per-frame block lists  $(n, q_x, q_y, q_z)$  quadruples where  $n$  is the integer point identifier and  $(q_x, q_y, q_z)$  are the millimetre-quantised anchor-relative deltas  $\delta_t^n$ . Only points visible at frame  $t$  are emitted; occluded points are not imputed. At inference, the model is asked to generate this answer span autoregressively under greedy decoding, and the parser re-assembles  $\hat{\mathbf{P}}_{t_0+1:t_0+T}$  from the recovered quadruples.

## D.3 Training hyperparameters

We train with AdamW using the Molmo2 supervised-fine-tuning defaults ( $\beta_1 = 0.9$ ,  $\beta_2 = 0.95$ , weight decay 0.1). The learning rate is warmed up linearly over the first 1K steps to its peak and then cosine-decayed to  $0.1 \times$  peak. Activations are bf16 with fp32 master weights; the model is distributed with FSDP2 full-shard across 16 H100 GPUs at per-device batch 16 for a global batch size of 256. Gradients are clipped at maximum-norm 1.0. Multi-dataset mixing across the six MolmoMotion-1M sources is square-root-weighted by per-dataset clip count. The autoregressive cross-entropy is computed on the answer span only, with the prompt portion (image, text, point-feature, and history-coordinate tokens) masked out of the loss; the flow-matching MSE is computed on the future positions only, leaving the clean history portion of the trajectory tensor unsupervised.

## D.4 Flow-matching objective and inference

We give the full flow-matching specification briefly described in Sec. 2.2. The flow-matching head predicts the future anchor-relative trajectory tensor  $\delta_{t_0+1:t_0+T} \in \mathbb{R}^{N \times T \times 3}$  in continuous metric coordinates, conditioned on the multimodal context  $\mathcal{C}$  (image, text, point-feature, and history-coordinate tokens) and on the clean initial 3D query coordinates  $\{\delta_{t_0}^n\}_{n=1}^N$ .

**Forward interpolation.** For each training example we draw a flow timestep  $\tau \sim \mathcal{U}(0, 1)$  and a noise tensor  $\epsilon \in \mathbb{R}^{N \times T \times 3}$  whose entries are i.i.d. standard Gaussian. The interpolated trajectory is

$$\delta_\tau = (1 - \tau)\epsilon + \tau\delta_{t_0+1:t_0+T}, \quad (3)$$

which slides linearly from pure Gaussian noise at  $\tau = 0$  to the clean ground-truth future trajectory at  $\tau = 1$ .

Variant	HOT3D			WorldTrack			DAVIS		
	ADE↓	FDE↓	PWT↑	ADE↓	FDE↓	PWT↑	ADE↓	FDE↓	PWT↑
<b>MolmoMotion-AR (<math>H = 3</math>)</b>	<b>0.109</b>	<b>0.217</b>	<b>0.444</b>	<b>0.143</b>	<b>0.261</b>	<b>0.445</b>	<b>1.227</b>	<b>2.108</b>	<b>0.153</b>
without 2D point feature	0.118	0.231	0.421	0.155	0.282	0.418	1.310	2.252	0.143
Absolute coords (no delta)	0.165	0.330	0.276	0.220	0.401	0.288	1.940	3.275	0.082
without language instruction	0.158	0.318	0.291	0.215	0.392	0.305	1.890	3.182	0.092
$N = 16$ query points	0.106	0.212	0.452	0.140	0.255	0.451	1.198	2.061	0.158

**Table 6 Model ablations on PointMotionBench.** All variants use the autoregressive head with the same Molmo2-4B initialization, the same 40K+10K training schedule, and the same evaluation protocol as Tab. 1. The reference row is MolmoMotion-AR at  $H = 3$ . Rows 2-4 remove one design choice; row 5 doubles the per-object query count. Anchor-relative parameterization and language conditioning are the largest contributors; the 2D point feature yields a small but consistent gain; doubling  $N$  improves accuracy slightly at substantial training cost (see text).

**Velocity head.** The DiT decoder  $v_\phi$  predicts the velocity field at  $(\delta_\tau, \tau)$  in the direction of the clean future. It takes three inputs: (i) the noised trajectory tensor with clean history concatenated to the noisy future, RoPE-tagged along both the point-index and the frame-index axes; (ii) the scalar  $\tau$  via a sinusoidal embedding added to every trajectory token; and (iii) the per-layer Molmo2 LM hidden states  $\mathcal{C}$  through the per-block cross-attention described in §D.1.

**Training loss.** The decoder is trained with the standard flow-matching mean-squared error [46],

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{\tau, \epsilon} \left[ \left\| v_\phi(\delta_\tau, \tau, \{\delta_{t_0}^n\}_{n=1}^N, \mathcal{C}) - (\delta_{t_0+1:t_0+T} - \epsilon) \right\|_2^2 \right], \quad (4)$$

where the regression target  $\delta_{t_0+1:t_0+T} - \epsilon$  is the constant velocity that drives  $\delta_\tau$  along the straight-line path of Eq. (3). The loss is masked to the future positions only, leaving the clean history portion of the trajectory tensor unsupervised.

**Inference.** At inference we sample  $\epsilon \sim \mathcal{N}(0, I)$  and integrate the learned velocity field with  $K = 10$  Euler steps, advancing  $\tau$  uniformly from 0 to 1 in increments of  $\Delta\tau = 0.1$ . Each step evaluates  $v_\phi$  once and updates

$$\delta_{\tau+\Delta\tau} = \delta_\tau + \Delta\tau \cdot v_\phi(\delta_\tau, \tau, \{\delta_{t_0}^n\}_{n=1}^N, \mathcal{C}). \quad (5)$$

The final  $\delta_1$  is added to the anchor 3D position  $\mathbf{p}_{\text{anc}}$  to recover the predicted future positions in the world frame,  $\hat{\mathbf{P}}_{t_0+1:t_0+T}$ . Because the DiT block count matches the LM layer count and one  $v_\phi$  evaluation reuses cached LM activations, each Euler step has roughly the cost of one LM forward pass.

**Comparison to autoregressive decoding.** We report flow-matching numbers in Tab. 1 alongside the autoregressive variant. The autoregressive head is stronger on the deterministic point-prediction metrics ADE / FDE / PWT because each predicted coordinate is conditioned in a strict left-to-right sense on every previously emitted coordinate, which encourages temporally smooth rollouts. The flow-matching head, by contrast, samples from the conditional distribution rather than committing to a single mode, which is desirable in settings where the action description leaves multiple plausible futures.

## E Model Ablations

We ablate three load-bearing design choices in the autoregressive variant: the per-query-point 2D feature, the anchor-relative coordinate parameterization, and the language-instruction conditioning. All ablations use the autoregressive variant with everything else held identical to the main model.

**Ablations.** The four ablated variants are constructed as follows. (i) *Without 2D point feature.* The grid-sampled per-query-point feature is omitted to test if 2D point feature is useful. (ii) *Absolute coordinates.* We replace the anchor-relative deltas  $\delta_t^n$  with absolute world-frame positions  $\mathbf{p}_t^n$  in both the prompt’s history

block and the answer span. The 1 mm quantization grid is preserved; only the coordinate origin changes. This isolates the value of the anchor-relative parameterization introduced in Sec. 2.2. (iii) *Without language instruction*. The action caption  $a$  is replaced with a single fixed token (“`motion`”); image, text, point-feature, and history-coordinate tokens are unchanged. This isolates the value of language conditioning. (iv)  $N = 16$  *query points*. We double the number of query points sampled per object from  $N = 8$  to  $N = 16$ , leaving every other hyperparameter unchanged. This tests whether denser per-object coverage improves prediction accuracy and quantifies the cost of going beyond our default.

**Results.** Tab. 6 reports 3D ADE, FDE, and PWT on the three PointMotionBench splits. Removing the anchor-relative parameterization produces the largest single drop ( $\approx 50\%$  on ADE/FDE across all splits), making absolute coordinates the strongest signal of design importance. Removing the language instruction produces a comparable drop, indicating that the action description does substantially more than disambiguate the object: it provides the direction prior the model relies on when the visual context alone leaves the future ambiguous, with the largest hit on DAVIS where intent is hardest to infer from a single anchor frame. The 2D point feature contributes a smaller but consistent gain (5–8% on ADE/FDE,  $\approx 5\%$  on PWT) and is most helpful on DAVIS, where the manipulated object is small relative to the frame. Doubling the per-object query count from  $N = 8$  to  $N = 16$  improves accuracy by roughly 2–3% on ADE/FDE and  $\approx 2\%$  on PWT across all three splits. The gain is small because the eight default points already cover the manipulated object’s surface densely after  $K$ -means selection, and our query points are sampled to be spatially well-distributed. The cost, however, is substantial: with  $T = 32$  in stage 2, the autoregressive answer span is twice as long under  $N = 16$  and exceeds the 4096-token context window of the Qwen3-4B language model that backs Molmo2-4B. We therefore keep  $N = 8$  as the default, which trades a small accuracy improvement for the longer prediction horizon  $T = 32$ . Lifting this constraint — through tokenization schemes that encode multiple coordinates per LM token, or through context-extension recipes for the backbone — is a natural next step for representing dense object coverage at long horizons; we leave it to future work.

**Inference cost.** The two heads share a single Molmo2-4B forward pass over the prompt and diverge only in how they decode the future trajectory. AR emits the answer span as text, so its cost grows linearly in  $N \cdot T$ . FM runs a fixed  $K$  Euler steps through the DiT trajectory expert, so its cost is independent of  $T$ . Tab. 7 reports per-clip latency at the headline ( $N=8, T=32$ ) setting on a single A100. Flow-matching is roughly two orders of magnitude faster than autoregressive decoding, which is the regime that matters in closed-loop robotic control, trajectory-conditioned video generation, and large-scale evaluation.

Variant	Setting	Latency (s / clip)
MolmoMotion-AR	$N=8, T=8$	37.2
MolmoMotion-AR	$N=8, T=32$	148.4
MolmoMotion-FM ( $K=10$ )	$N=8, T=32$	1.1

**Table 7 Inference cost.** Per-clip latency at  $H=3$  on a single A100. AR cost scales linearly in  $N \cdot T$ ; FM cost is independent of  $T$ . Flow-matching is roughly  $150\times$  faster than autoregressive decoding at  $T=32$ , at the modest accuracy cost in Tab. 1.

## F Robotics Transfer Settings

This appendix documents the implementation specifics of the two robotics experiments reported in Sec. 4.2: closed-loop pick-and-place on MolmoSpaces (Fig. 7a) and 3D-trajectory finetuning on DROID (Fig. 7b).

### F.1 MolmoSpaces pick-and-place

**Policy.** The downstream policy follows the MolmoBot [18] prompt-encoder recipe. The Molmo2-4B vision-language backbone is followed by an ActionExpert head: a flow-matching transformer with 36 cross-attention blocks (one per LM layer), action dimension 8 (7 arm joints plus 1 gripper), and action horizon 16. At inference, action chunks are produced by integrating the ActionExpert’s velocity field for 10 Euler steps from Gaussian noise. The backbone is initialised either from MolmoMotion (autoregressive variant after both training stages) or from the public Molmo2-4B-Pretrain checkpoint for the matched control.

Metric	What it measures
Tem-Con $\uparrow$	mean CLIP cosine similarity between adjacent frames
Subj-Cons $\uparrow$	feature-similarity of the moving subject across frames
M-Smooth $\uparrow$	smoothness of the estimated optical-flow residual
Dyn-Deg $\uparrow$	fraction of clips whose flow magnitude exceeds a motion threshold
Bg-Cons $\uparrow$	feature-similarity of the background region across frames

**Table 8** Higher-is-better video-quality metrics used to compare the three image-to-video generators. Subj-Cons, M-Smooth, Dyn-Deg, and Bg-Cons are standard VBench dimensions; Tem-Con is computed separately from CLIP frame embeddings.

**Inputs.** Each step provides three history frames per camera at sim-step deltas  $\{-4, -2, 0\}$  from two cameras (egocentric exo and wrist-mounted), plus the current 8-D robot state normalised to unit-quantile space using statistics fit on the training split. The textual prompt reuses the trajectory-prediction wrapper of §D.2 with an empty answer span; the LM head’s cross-entropy weight is set to zero, so only the ActionExpert flow-matching MSE contributes to the gradient.

**Action representation.** Actions are absolute joint-position targets supervised over a 16-step chunk ( $\approx 1.06$  s of simulated time at 15 fps). At execution time we run 8 steps of the predicted chunk and re-query the policy at sim-step +8, matching the training-time action-time semantics. Per-joint deltas are clamped to  $\pm 0.2$  rad/step at execution as a safety guard against out-of-distribution velocity predictions.

**Training.** We finetune on the 20K pick-and-place episodes released with MolmoBot, restricted to `policy_phase`  $\in \{5, 6, 7, 8\}$  (lift, preplace, place, retreat) so the policy is supervised only on the post-grasp portion of each trajectory. Optimization uses AdamW with the Molmo2 SFT defaults, maximum sequence length 2048, gradient clipping at norm 1.0, bf16 mixed precision, FSDP2 full-shard, per-device batch 2 with 8 gradient-accumulation micro-batches for a global batch size of 128 across 8 H100 GPUs. Training runs for 100K optimization steps; checkpoints are saved every 5K steps and evaluated periodically.

**Hybrid rollout and evaluation.** Closed-loop evaluation uses a hybrid rollout. The released MolmoBot-DROID policy drives every episode through the approach and grasp phases. The first sim step at which the simulator reports `held = true` on the pickup object marks a hand-off trigger; one execute-horizon window (8 sim steps) later, control is handed to the evaluated policy, which performs the lift, transport, and placement. Both policies output absolute joint-position targets. Each rollout runs for at most 600 sim steps ( $\approx 40$  s at 15 fps); an episode succeeds when the simulator’s terminal `success` flag is true, i.e. when the pickup object lies inside the success-position threshold of its target receptacle.

## F.2 DROID trajectory finetuning

**Task and data.** We use the same 3D point-trajectory prediction objective as in Sec. 2.2; no robot-action loss is applied. DROID wrist and shoulder videos are paired with 3D trajectories produced by running the MolmoMotion-1M annotation pipeline of Sec. 3.1 on the DROID corpus, yielding the same triplet of (object mask, action description, dense 3D trajectory) as the pretraining data. We hold out a fixed 10 percent subset of clips for the test set.

**Initialization and finetuning.** The pretraining-init run starts from the MolmoMotion-AR ( $H = 1$ ) checkpoint; the matched control starts from Molmo2-4B-Pretrain and is trained on DROID alone. Both runs use the same hyperparameters: AdamW, 128 global batch on 8 H100s, max-norm 1.0 gradient clipping, bf16 mixed precision. Both runs are trained for 12K finetuning steps; we evaluate every 1K steps and trace the test-loss curves shown in Fig. 7b.

**Evaluation.** We report 3D test L2 on held-out DROID clips, computed identically to the ADE metric of Sec. 4.1 but on the DROID test split. The pretraining-init run starts at substantially lower L2 and reaches the matched-control’s 12K-step error after only  $\approx 2$ K finetuning steps.

## G Video generation experiment details

We expand the video-generation experiment summarized in Sec. 4.3: starting from the first frame and action description of a benchmark clip, we ask three image-to-video generators to synthesize the rest of the clip and compare their outputs on standard video-quality metrics. We describe the methods compared (§G.1) and the evaluation protocol (§G.2).

### G.1 Methods compared

We evaluate three image-to-video generators. DAS [29] is a 3D-track-conditioned image-to-video model built on the CogVideoX-5B backbone [73], and is the only method that consumes MolmoMotion’s predicted 3D trajectories. CogVideoX-5B-I2V is the same backbone without the tracking branch; comparing DAS+MolmoMotion against it isolates the contribution of track conditioning while controlling for the underlying generator. Wan2.2-I2V-A14B [67] is an unconditioned image-to-video baseline at roughly  $2.8\times$  the parameter count of CogVideoX-5B-I2V, and tells us how a much larger generator without explicit motion conditioning compares to a smaller one with it.

### G.2 Evaluation protocol

For each clip we feed the same first-frame image and caption to all three generators; DAS additionally receives the MolmoMotion prediction track. Each generated clip is then transformed to the ground-truth frame count, frame rate, and resolution, so all methods are scored on the same temporal and spatial grid.

We score each generated clip on five higher-is-better video-quality metrics (Tab. 8): four standard VBench dimensions [35]—subject consistency, motion smoothness, dynamic degree, and background consistency—and a CLIP-based temporal-consistency score, reported alongside as a complementary check on frame-to-frame coherence. We aggregate over PointMotionBench (§C).

**“take cloth out of container”**




**“move object to the white box”**




**“move lid on pot”**




**“move tool to bowl”**




**“move purple toy to white dish”**




**“stack red bowl on black one”**




**“put toy in black bowl”**



**“move block to the black bowl”**



**“move object to container”**



**Figure 10 MolmoMotion predictions on held-out DROID clips.**



Figure 11 Video generation comparisons on held-out PointMotionBench prompts (1/2).



Figure 12 Video generation comparisons on held-out PointMotionBench prompts (2/2).