

# Dream2Flow: Bridging Video Generation and Open-World Manipulation with 3D Object Flow

Karthik Dharmarajan, Wenlong Huang, Jiajun Wu, Li Fei-Fei\*, Ruohan Zhang\*  
Stanford University

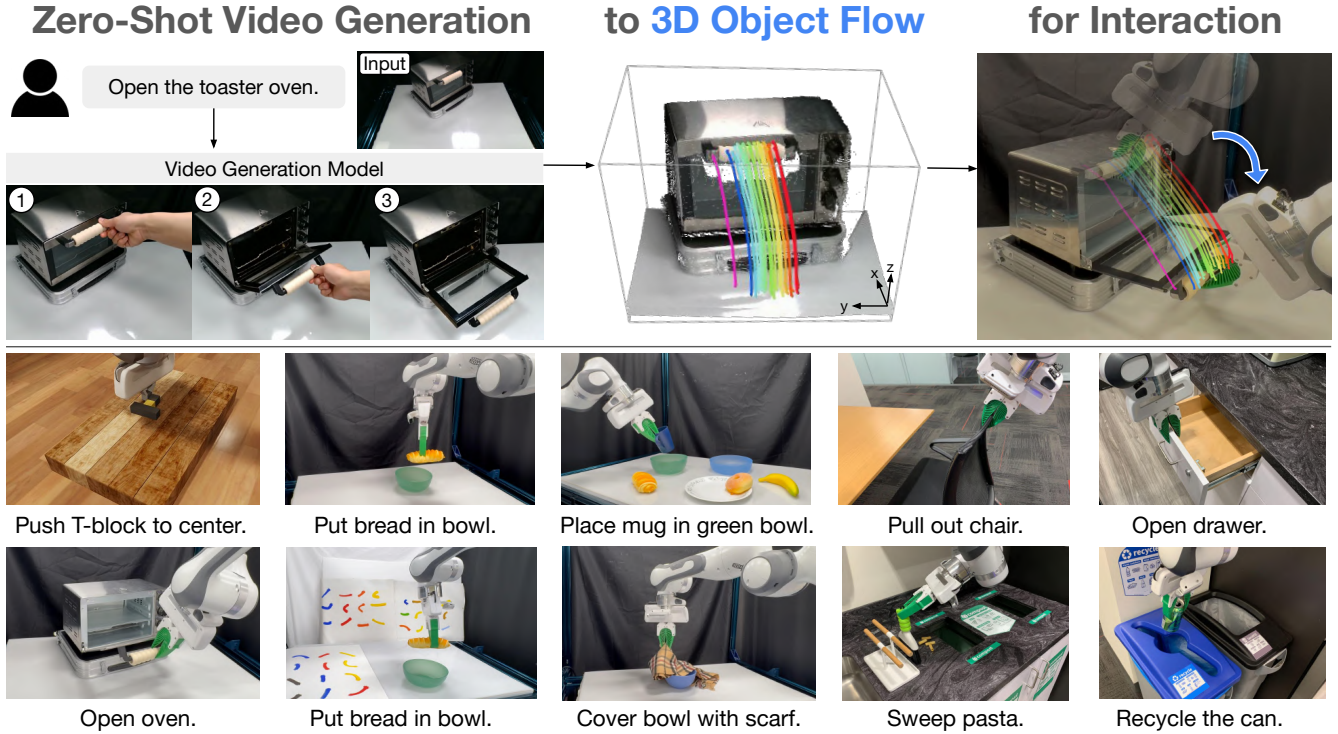


Fig. 1: **Dream2Flow** leverages off-the-shelf video generation models to produce videos of the task being performed in the same scene of the robot. Dream2Flow then extracts a 3D object flow from the motion in the video, allowing for downstream planning and execution with a robot across a wide variety of tasks.

**Abstract**—Generative video modeling has emerged as a compelling tool to zero-shot reason about plausible physical interactions for open-world manipulation. Yet, it remains a challenge to translate such human-led motions into the low-level actions demanded by robotic systems. We observe that given an initial image and task instruction, these models excel at synthesizing sensible object motions. Thus, we introduce Dream2Flow, a framework that bridges video generation and robotic control through 3D object flow as an intermediate representation. Our method reconstructs 3D object motions from generated videos and formulates manipulation as object trajectory tracking. By separating the state changes from the actuators that realize those changes, Dream2Flow overcomes the embodiment gap and enables zero-shot guidance from pre-trained video models to manipulate objects of diverse categories—including rigid, articulated, deformable, and granular. Through trajectory optimization or reinforcement learning, Dream2Flow converts reconstructed 3D object flow into executable low-level commands without task-specific demonstrations. Simulation and real-world experiments highlight 3D object flow as a general and scalable interface for adapting video generation models to open-world robotic manipulation. Videos, visualizations, and appendix are available at <https://dream2flow.github.io/>.

## I. INTRODUCTION

Robotic manipulation in the open world could greatly benefit from visual world models that predict how an environment would evolve given an agent’s interactions. Recent advances in generative video modeling have produced systems capable of zero-shot synthesizing minute-long, high-fidelity clips of physical interactions in pixel space, conditioned on an unseen initial image and an open-ended task instruction [1]. Such video models implicitly capture intuitive physics and rich priors of object properties and interactions, making them compelling for open-world manipulation settings where a robot is tasked to complete novel tasks in unseen environments through partial observations.

Despite their promise, it remains unclear what role such models should serve in a robot manipulation system. Most frontier video generators produce the best interaction clips with a human embodiment. This reflects where supervision is most abundant as human interactions are significantly more

\*Equal Advising. Correspondence: Wenlong Huang

broadly documented than those of robots. But this becomes a challenge for using them in robotic manipulation due to the embodiment gap and hence the different action spaces.

We propose extracting actionable signals from their visual predictions of human interactions, which will then be enacted by a robot. Our proposed method, **Dream2Flow**, employs **3D object flow** as an intermediate interface that bridges high-level video simulation with low-level robot actions. Dream2Flow works because, despite occasional visual artifacts, state-of-the-art video generation models often predict physically plausible object motions that align with the task intent in open-world manipulation tasks. Then, given generated videos, instead of trying to directly mimic the human motions for completing a given task, we focus on reconstructing and reproducing the object flows in 3D.

The problem is thus reduced to object trajectory tracking: the robot’s job is to manipulate the object to closely follow the generated flow that the video model imagined. This approach cleanly separates what needs to happen (i.e., state changes in an environment) from how a particular embodiment achieves it with respect to its kinematic and dynamic constraints (i.e., actions). Importantly, it seamlessly interfaces with both motion planners and sensorimotor policies—the extracted object motion in 3D serves as the tracking goal for trajectory optimization or a reinforcement learning policy, which then yields a sequence of low-level robot joint commands.

Leveraging off-the-shelf models and tools, we demonstrate an autonomous pipeline that 1) generates a text-conditioned video of plausible interaction [1], 2) obtains 3D object flow by performing depth estimation and point tracking using vision foundation models [2, 3], and 3) synthesizes robot actions that realize this flow using trajectory optimization and reinforcement learning. Notably, this design enjoys a number of desirable benefits in manipulation tasks. By first leveraging video models—pre-trained on large corpora of human activities—to interpret and ground open-ended language commands in visual predictions, our system inherits a scalable mechanism for task specification. These predictions are then distilled into reconstructed 3D object flow, a representation that naturally captures diverse object interactions spanning rigid, articulated, deformable, and granular objects. Together, this synergy enables an end-to-end pipeline that performs open-world manipulation directly from visual perception and language, without task-specific data or training.

In summary, our key contributions are:

- We propose 3D object flow as an interface for adapting off-the-shelf video generation models for open-world manipulation by formulating it as an object trajectory tracking problem.
- We demonstrate its effectiveness by implementing the approach in both simulated and real domains, which performs diverse tasks given only RGB-D observations and language instructions in a zero-shot manner.
- We examine the properties of 3D object flow by comparing it with alternative intermediate representations and by studying its key design choices as well as

generalization properties.

## II. RELATED WORKS

### A. Task Specification in Manipulation

Specifying tasks for the wide range of manipulation problems spans symbolic, learning-based, outcome-driven, and object-centric interfaces. Classical approaches encode goals and constraints with symbolic formalisms such as PDDL and temporal logics, or optimize cost-augmented formulations [4–6]. Learning-based systems specify tasks often through language and perception, mapping instructions to actions via language-conditioned visuomotor policies and vision–language–action models [7–10]. Outcome-based specification sets goals by example observations, e.g., image goals with goal-conditioned policies [11–16]. Object-centric alternatives rely on descriptors or keypoints to capture task-relevant structure [17]. Recently, foundation models enable higher-level interfaces that compile intent into actionable specifications via code [18], 3D value maps akin to potential fields [19], keypoint relations [20], or affordance maps [21].

### B. 2D/3D Flow in Robotics

Dense motion fields—optical flow, point tracks, and 3D scene/object flow—provide an embodiment-agnostic, mid-level interface for manipulation [22, 23]. In scene-centric formulations, policies parameterize or condition on motion in 2D or 3D to decide actions, with point/keypoint interfaces unifying perception and action and with action-flow improving precision [24–30]. In object-centric formulations, desired object motion is specified independently of embodiment and then converted into actions via policy inference, planning, and optimization [31–36]. When actions are absent, retargeting with predicted tracks and dense correspondences offers a practical bridge across embodiments [37–39]. Advances in perception and tracking such as RGB-D motion-based segmentation, 3D scene flow in point clouds, zero-shot monocular and ToF-based scene flow, rigid-motion learning, refractive flow for transparent objects, deformable and articulated reconstruction, and structured scene representations [40–50] make the interface reliable. Flow-derived supervision further supports learning, and video generation supplies plausible visual rollouts for planning and imitation [51–54]. Our approach follows the object-centric path by reconstructing 3D object flow from language-conditioned generations and tracking it under embodiment constraints, complementing other flow-conditioned policy representations [24, 26, 31–33].

### C. Video Models for Robotics

Recent work increasingly integrates video models across robotic tasks in various ways [55, 56]. They can serve as auxiliary training objectives [57–62], as reward models [63–65], as policies [66, 67], or as a simulator for the environments [68, 69]. Notably, predictive modeling in robotics can leverage video frame prediction as a form of world model. By simulating future visual observations, these models can enable visual planning and manipulation by anticipating how the environment will evolve. For example, a video generative

model can simulate long-horizon task outcomes, effectively acting as a visual planner [70]. Video models in this way can also serve as dynamics models [67, 71–75]. Another notable direction is that video generation can directly provide new training data for robot learning. Several recent works devise frameworks to imagine new trajectories in the form of videos and use them to train or finetune policies, particularly for imitation learning [76, 77].

### III. METHOD

Herein, we introduce the problem formulation of Dream2Flow in Sec. III-A. Leveraging 3D object flow as an interface, we subsequently discuss how to extract 3D object flow from video generations in Sec. III-B and how to plan actions with 3D object flow for manipulation in Sec. III-C.

#### A. Problem Formulation

Given a task instruction  $\ell$ , an initial RGB-D observation ( $I_0 \in \mathbb{R}^{H \times W \times 3}$ ,  $D_0 \in \mathbb{R}^{H \times W}$ ), and a known camera projection  $\Pi$  (intrinsic and extrinsic to the robot frame), our goal is to output an action sequence  $u_{0:H-1} \in \mathcal{U}^H$  that accomplishes the task by following an object motion inferred from a generated video. We make no assumption about a specific action parameterization:  $\mathcal{U}$  may represent motion primitives, end-effector poses, or low-level controls.

**Extracting 3D Object Flow.** From  $(I_0, \ell)$ , an image-to-video model produces frames  $\{V_t\}_{t=1}^T$ , and a video-depth estimator provides a per-frame depth sequence  $\{Z_t\}_{t=1}^T$ . Using a binary mask  $M$  of the task-relevant object, and the projection  $\Pi$ , we lift masked image points with  $Z_{1:T}$  to obtain an object-centric 3D trajectory  $P_{1:T} \in \mathbb{R}^{T \times n \times 3}$  in the robot frame; we refer to  $P_{1:T}$  as the 3D object flow.

**Action Inference with 3D Object Flow.** We represent state as the task-relevant object and the robot:  $x_t = (x_t^{\text{obj}}, r_t)$ , where  $x_t^{\text{obj}} \in \mathbb{R}^{n \times 3}$  are object points and  $r_t$  denotes the robot state. Let  $f$  be a dynamics model and  $\hat{x}_{t+1} = f(\hat{x}_t, u_t)$  with  $\hat{x}_0 = x_0$ . At each planning step  $t$ , we use a time-aligned target  $\tilde{P}_t \in \mathbb{R}^{n \times 3}$  derived from the video object flow (e.g., via uniform time-warping or nearest-shape matching). We formulate action inference as an optimization problem:

$$\begin{aligned} \min_{\{u_t \in \mathcal{U}\}} & \sum_{t=0}^{H-1} \lambda_{\text{task}}(\hat{x}_t^{\text{obj}}, \tilde{P}_t) + \lambda_{\text{control}}(\hat{x}_t, u_t) \\ \text{s.t.} & \hat{x}_{t+1} = f(\hat{x}_t, u_t), \quad \hat{x}_0 = x_0, \\ & \lambda_{\text{task}}(\hat{x}_t^{\text{obj}}, \tilde{P}_t) = \sum_{i=1}^n \|\hat{x}_t^{\text{obj}}[i] - \tilde{P}_t[i]\|_2^2, \end{aligned}$$

Section III-C instantiates  $\mathcal{U}$  and  $f$  for different domains.

#### B. Extracting 3D Object Flows from Videos

**Video Generation:** Given the task language instruction  $\ell$  and an RGB image of the workspace without the robot visible  $I_0 \in \mathbb{R}^{H \times W \times 3}$ , Dream2Flow uses an off-the-shelf image-to-video generation model to produce an RGB video  $\{V_t\}_{t=1}^T$  with  $V_t \in \mathbb{R}^{H \times W \times 3}$ , showing the task being performed. We do not include the robot in the initial frame or mention the robot in the text prompt, as we empirically find that current

image-to-video generation models not specifically finetuned on robotics data tend to produce physically implausible fine-grained interactions, and consequently have worse object trajectories (for details see Appendix A).

**Video Depth Estimation:** Given the generated video, Dream2Flow leverages SpatialTrackerV2 [78, 79] to estimate per-frame depth  $\{\tilde{Z}_t\}_{t=1}^T$ ,  $\tilde{Z}_t \in \mathbb{R}^{H \times W}$ . Due to the scale-shift ambiguity of monocular video, we compute global  $(s^*, b^*)$  by aligning the first frame to the initial depth  $D_0$  from the robot, and obtain calibrated depths  $Z_t = s^* \tilde{Z}_t + b^*$ .

**3D Object Flow Extraction:** 3D object flow aims to produce 3D trajectories  $P_{1:T} \in \mathbb{R}^{T \times n \times 3}$  with visibilities  $V \in \{0, 1\}^{T \times n}$  for the task-relevant object. We first localize the relevant object using Grounding DINO [80] to produce a bounding box from  $(I_0, \ell)$ , and then use the box to prompt SAM 2 [2] for a binary mask. From the masked region at  $t=1$ , we sample  $n$  pixels and track them across the video with CoTracker3 [3] to obtain 2D trajectories  $c_i^t$  and visibilities  $v_i^t$ . Visible points are lifted to 3D using the calibrated depths and camera intrinsic/extrinsic, producing  $P_{1:T}$  as in Sec. III-A.

#### C. Action Inference with 3D Object Flow

**Simulated Push-T Domain.** For tasks involving non-prehensile manipulation, such as the Push-T task, Dream2Flow uses a push skill primitive parameterized with the start push position on a flat table  $(c_x, c_y)$ , a unit push direction  $(\Delta c_x, \Delta c_y)$ , and the distance of the push  $d$ . In this setting, we learn a forward dynamics model that takes as input a feature-augmented particles  $\tilde{x}_t \in \mathbb{R}^{N \times 14}$  of the whole scene and produces  $\Delta \hat{x}_{t+1}$ , the delta of positions of each point for the next timestep. The features associated with each point consist of the position, RGB color, and normal vector as determined from camera observations along with the push parameters (as further elaborated in Appendix B).

To optimize the 3D object flow following cost with the learned dynamics model, we use random-shooting, where  $r$  push skill parameters are randomly sampled such that all pushes will make contact with the object of interest at different points and directions. Then, we select the push skill parameter out of  $r$  ones that has the least cost according to the predicted point positions from the dynamics model. For determining which timestep from the video should be used in the cost function, we find the timestep  $t^*$  where the points of the relevant object in the trajectory are closest to the current observed points with further details in Appendix C.

**Real-World Domain.** We use absolute end-effector poses as the action space and a rigid-grasp dynamics model for the real robot domain. With this combination, we first proceed to grasp the desired part from the relevant object, and then use the dynamics model and point-flow following objective to move the end-effector such that the grasped part moves in a fashion similar to the video. We use AnyGrasp [81] to propose candidate grasps on the object of interest, but these grasps may not lie exactly on the desired part. We select the grasp closest to the thumb as detected from the video by HaMer [82], as we observe that in generated videos, the




Fig. 2: **An overview of Dream2Flow.** Given a task instruction and an initial RGB-D observation, an image-to-video model synthesizes video frames conditioned on the instruction. We additionally obtain object masks, video depth, and point tracking from vision foundation models, which are used to reconstruct 3D object flow. Finally, a robot policy generates executable actions that track the 3D object flow using trajectory optimization or reinforcement learning.

hand tends to interact with the relevant part of an object, such as the handle.

The rigid-grasp dynamics model assumes that the grasped part is rigid; consequently, the prediction of point positions to the next timestep can be described by a sequence of rigid transformations for the grasped subset, while non-grasped points remain unchanged. To produce a trajectory of end-effector poses, Dream2Flow directly optimizes the objective using PyRoki [83], incorporating pose smoothness and reachability costs as  $\lambda_{\text{control}}$  in addition to the 3D object-flow following cost  $\lambda_{\text{task}}$  as in Appendix F.

**Simulated Door Opening Domain.** For the Door Opening task, we use reinforcement learning to learn a sensimotor policy which moves the object according to the 3D object flow with SAC [84]. This approach can be viewed as using the simulator as a dynamics model for compiling the optimization process prescribed in Eq. III-A into a parametric policy in an offline fashion, using 3D object flow as the reward function. Depending on the embodiment used, the action space consists of delta end-effector poses and delta joint angles for the gripper or dexterous hand. The reward function used consists of one term which encourages the end-effector to move toward the current mean object particle position along with another term for encouraging matching the 3D object flow,  $\frac{t^*}{t_{\text{end}}}$ , where  $t^*$  is the timestep with the closest particle positions to the 3D object flow and  $t_{\text{end}}$  is the last timestep of the object motion (further details in Appendix H). By training policies with such an object-centric reward, we observe different strategies emerge to accomplish the same object motion across embodiments, including quadruped manipulators, humanoids with dexterous hands, and fixed-base arms with parallel grippers.

#### IV. EXPERIMENTS

We seek to answer the following research questions through our experiments: *Q1*: What are properties of 3D object flow when used as an interface to bridge videos and robot control? *Q2*: How does Dream2Flow perform compared to alternative interfaces? *Q3*: How effective is 3D object flow as a reward for learning sensimotor policies? *Q4*: How does the choice of video model affect Dream2Flow in simulation and in real-world tasks? *Q5*: How does the choice

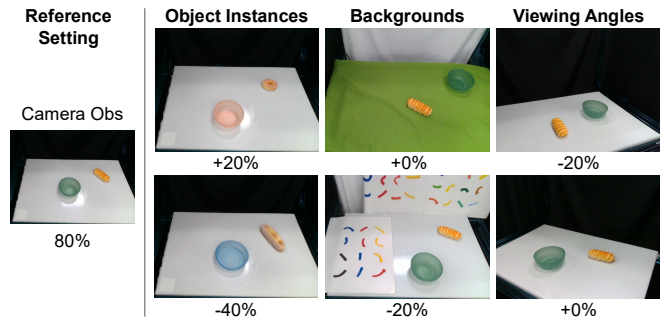


Fig. 3: **Robustness evaluations.** Relative performance across instance, background, and task variations, showing Dream2Flow remains robust under various different settings.

of dynamics model affect the performance of Dream2Flow?

For evaluation, we consider several tasks involving different types of objects and manipulation strategies, such as pushing a T-block to a goal, putting a piece of bread in a bowl, opening a toaster oven, covering a bowl with a scarf, and opening a door from Robosuite [85], as shown in Figure 1.

##### A. Properties of 3D Object Flow as a Video-Control Interface

We run Dream2Flow on the simulated Push-T task with Wan2.1 [86] as the video generation model. For this task only, we allow prompting with a goal image of the T-block in the center, as this task requires accurate movement. We consider 10 different initial states with the same final state. We then proceed to run 10 trials for each initial configuration on different seeds to properly evaluate performance due to the stochastic nature of random shooting, yielding 100 total trials. The result appears in Table IV. We note that in this task, 6 generated videos included substantial morphing of the T-block, which consequently ruined the tracking and downstream execution.

We evaluate Dream2Flow’s ability to perform manipulation with different types of objects with Veo 3 [87] as the video generation model in the real world for 10 trials per task, and report the results in Table I.

To assess Dream2Flow’s generalization to different object instances, backgrounds, and viewing angles, we conduct an additional five trials each for six different scenarios, as shown

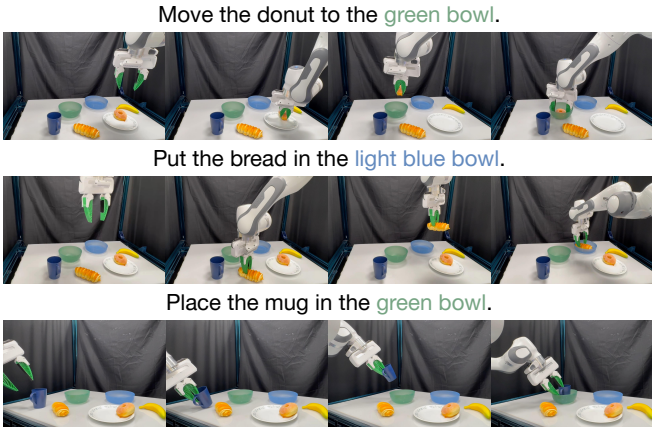


Fig. 4: **Multiple tasks in the same scene.** With different language goals, Dream2Flow adapts object-flow targets to produce distinct behaviors in the same environment.

Task	AVDC	RIGVID	Dream2Flow
Bread in Bowl	7/10	6/10	<b>8/10</b>
Open Oven	0/10	6/10	<b>8/10</b>
Cover Bowl	2/10	1/10	<b>3/10</b>

TABLE I: **Comparisons of intermediate representations on real robot.** Dream2Flow outperforms AVDC and RIGVID across three tasks by following 3D object flow rather than rigid transforms alone.

in Fig. 3. Compared to the reference setting, with exception of putting a large piece of bread, there is not a significant drop off in performance, suggesting that Dream2Flow inherits some generalization through the use of video generation models. We additionally show that Dream2Flow can perform different tasks from the same scene in Fig. 4 due to the ability of video generation models to follow different task instructions given the same input image. We have additional case studies demonstrating that 3D object flow can be used for downstream manipulation in in-the-wild settings for tasks such as pulling a chair, opening a drawer, sweeping pasta, and recycling a can, as shown in Fig. 1 and Appendix G.

For all 60 trials of Dream2Flow executed in the real world, we provide a breakdown of failures in Fig. 6. In the 12 video generation failures, for half the time, the generated video either morphs an object in an implausible way or hallucinates new objects, causing tracking to unreasonably fail or making the robot move an object to an incorrect 3D location. The four flow extraction failures occurred because of severe rotations or objects temporarily going out of view of the camera, leading to tracks with no visibility in the end. The four robot execution failures occurred in the Bowl Covering task, where the robot either did not grasp at the correct point or did not move enough.

### B. How does Dream2Flow perform compared to alternative interfaces?

For the three tasks in the real world, we consider the following alternative interfaces related to extracting object trajectories from videos:

1) *AVDC*: AVDC [38] leverages generated videos by computing dense optical flow between frames to track points on a rigid object. Then, using the initial depth and the point correspondences, it solves for a sequence of rigid transforms of the relevant object, allowing for trajectory playback after a grasp has occurred. In our implementation, we utilize the video depth corresponding to the tracks from the optical flow, and optimize for rigid transforms relative to the initial frame, as we empirically found that to be less noisy than the original optimization procedure.

2) *RIGVID*: RIGVID [54] uses 6D object pose tracking to generate a rigid object trajectory from a generated video. Since it is ill-defined to have such a pose for deformable objects, we do not use a 6D pose tracker, but instead adapt their approach by solving for a rigid pose transformation between the initial 3D points and visible 3D points from our 3D Object Flow computation in the same manner as AVDC.

We present the results in Table I. While AVDC does a reasonable job at tracking the bread, the dense optical flow does not keep up with the motion of the oven, resulting in insufficient motion. Under certain circumstances, there are only a few visible points for RIGVID and AVDC, making transform estimation noisy. Attempting to follow the noisy transform trajectories leads to execution failures or optimization instability. Dream2Flow is less affected by this issue, because there is typically not a heavy cost when most of the points are occluded, allowing the planned end-effector poses to smoothly move between areas of high point visibility. The Cover Bowl task remains a challenge for AVDC and RIGVID, as in addition to video generation and tracking failures, the transform estimates are incorrect since the points are now on a deformable object.

### C. How effective is 3D object flow as a reward for learning sensimotor policies?

The 3D object flow extracted by Dream2Flow can be used in an RL reward for training policies across different embodiments. We evaluate SAC [84] policies trained using handcrafted object state reward from Robosuite and 3D object flow rewards on a Franka Panda, a Spot with floating base, and a GR1 with the right arm only across 100 random door positions, and report the results in Table II with corresponding episode visualizations in Figure 5. The policies trained with the object state and 3D object flow reward have comparable performance across all the embodiments. The learned strategies are different between the embodiments, as the Spot is able to move its base for better reachability and kinematic range, and the GR1 uses the area between the fingers and palm for pulling the door as opposed to individual fingers for better stability.

### D. How does the choice of video model affect Dream2Flow in simulation and in real-world tasks?

To evaluate how well different video generation models capture physically plausible object trajectories, we run Dream2Flow on the simulated Push-T task and the real Open Oven task using three models: Wan2.1 [86], Kling 2.1, and

Reward Type	Franka	Spot	GR1
Object State	99/100	99/100	<b>96/100</b>
3D Object Flow	<b>100/100</b>	<b>100/100</b>	94/100

TABLE II: **Comparison of policies trained using different rewards.** The policies trained using the 3D object flow reward perform comparably to those trained with the object state reward across different embodiments.



Fig. 5: **Rollouts from policies trained using 3D object flow as a reward.** Different embodiments such as the (a) Panda, (b) Spot, or (c) GR1 use different strategies to open the door. The Spot is able move its base for better reachability while the GR1 uses the area between its fingers and palm to pull for better stability.

Video Generation Model	Push-T	Open Oven
Wan2.1 [86]	<b>52/100</b>	2/10
Kling 2.1	31/100	4/10
Veo 3 [87]	-	<b>8/10</b>

TABLE III: **Effect of video generator.** VEO 3 excels on real-world domains such as “Open Oven”, while Wan 2.1 performs better on simulated domains such as “Push-T”.

VeO 3 [87]. The results are shown in Table III. Note that there are no results for VeO 3 on Push-T because at the time of evaluation, VeO 3 did not support prompting with a goal image.

For the Push-T task, Kling 2.1 had more videos that had substantial morphing, throwing off the tracking, resulting in more downstream failures. For the Open Oven task, Wan2.1 tends to produce more videos with substantial camera motion, which violates the still camera assumption. Additionally, both Kling 2.1 and Wan2.1 produce videos where the direction of articulation is incorrect, such as revolving around the wrong axis, leading to far more failures than VeO 3.

#### E. How does the choice of dynamics model affect the performance of Dream2Flow?

For the Push-T task, we compare the effect of different dynamics models used for planning. In addition to the particle-based dynamics model, we consider another learned dynamics model which takes in the block pose and push skill

Dynamics Model Type	Success Rate
Pose	12/100
Heuristic	17/100
Particle	<b>52/100</b>

TABLE IV: **Dynamics model ablation.** Particle dynamics substantially outperform pose and heuristic models, highlighting the importance of per-point predictions.

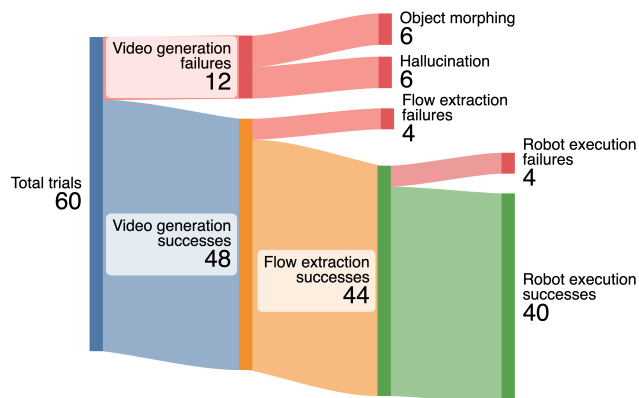


Fig. 6: **Failure breakdown on real-robot experiments.** Common causes include video artifacts (object morphing, object hallucination), tracking errors, and grasp selection mismatches.

parameters and predicts the delta pose of the T-block, trained on the same data as the particle based dynamics model, as well as a heuristic dynamics model, which translates the points of the T-block in the direction and amount of the push without any rotation. We present the results in Table IV, and find that the particle representation for this dynamics model is crucial to ensure success, as despite having the same 3D object flow guidance, the pose and heuristic based dynamics models could not sufficiently account for the rotation needed.

## V. CONCLUSION

We presented Dream2Flow, a simple, general interface that turns text-conditioned video predictions into executable robotic actions by reconstructing and tracking 3D object flow. Built entirely from off-the-shelf video generation and perception tools, Dream2Flow solves open-world manipulation tasks in simulation and on real robots across rigid, articulated, deformable, and granular objects using only RGB-D observations and language, without task-specific demonstrations. Experiments show consistent gains over trajectory baselines derived from dense optical flow or rigid pose transforms, robustness to variations in instances, backgrounds, and viewpoints, and the importance of both the upstream video model and downstream dynamics choice (with particle dynamics proving most reliable). A failure analysis highlights current bottlenecks, such as video artifacts (morphing, hallucinations), occlusion-induced tracking dropouts, and grasp selection mismatches as concrete directions for improvement. Overall, our results indicate that 3D object flow is a scalable bridge from open-ended video generation to robot control in unstructured environments.

## ACKNOWLEDGMENT

This work is in part supported by the Stanford Institute for Human-Centered AI (HAI), the Schmidt Futures Senior Fellows grant, ONR MURI N00014-21-1-2801, ONR MURI N00014-22-1-2740, ONR MURI N00014-24-1-2748, and NSF RI #2338203.

## REFERENCES

- [1] T. Brooks, B. Peebles, C. Holmes, W. DePue, Y. Guo, L. Jing, D. Schnurr, J. Taylor, T. Luhman, E. Luhman *et al.*, “Video generation models as world simulators,” *OpenAI Blog*, vol. 1, no. 8, p. 1, 2024.
- [2] N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. Ryal, 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. [Online]. Available: <https://arxiv.org/abs/2408.00714>
- [3] N. Kararov, I. Makarov, J. Wang, N. Neverova, A. Vedaldi, and C. Rupprecht, “CoTracker3: Simpler and better point tracking by pseudo-labelling real videos,” *arXiv*, 2024.
- [4] C. R. Garrett, R. Chitnis, R. Holladay, B. Kim, T. Silver, L. P. Kaelbling, and T. Lozano-Pérez, “Integrated task and motion planning,” *Annual review of control, robotics, and autonomous systems*, vol. 4, no. 1, pp. 265–293, 2021.
- [5] Z. Zhao, S. Cheng, Y. Ding, Z. Zhou, S. Zhang, D. Xu, and Y. Zhao, “A survey of optimization-based task and motion planning: From classical to learning approaches,” *IEEE/ASME Transactions on Mechatronics*, 2024.
- [6] M. Toussaint, “Logic-geometric programming: An optimization-based approach to combined task and motion planning,” in *IJCAI*, 2015, pp. 1930–1936.
- [7] M. Shridhar, L. Manuelli, and D. Fox, “Cliport: What and where pathways for robotic manipulation,” in *Conference on robot learning*. PMLR, 2022, pp. 894–906.
- [8] —, “Perceiver-actor: A multi-task transformer for robotic manipulation,” in *Conference on Robot Learning*. PMLR, 2023, pp. 785–799.
- [9] B. Zitkovich, T. Yu, S. Xu, P. Xu, T. Xiao, F. Xia, J. Wu, P. Wohlhart, S. Welker, A. Wahid *et al.*, “Rt-2: Vision-language-action models transfer web knowledge to robotic control,” in *Conference on Robot Learning*. PMLR, 2023, pp. 2165–2183.
- [10] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. P. Foster, P. R. Sanketi, Q. Vuong *et al.*, “Openvla: An open-source vision-language-action model,” in *Conference on Robot Learning*. PMLR, 2025, pp. 2679–2713.
- [11] C. Lynch, M. Khansari, T. Xiao, V. Kumar, J. Tompson, S. Levine, and P. Sermanet, “Learning latent plans from play,” in *Conference on robot learning*. Pmlr, 2020, pp. 1113–1132.
- [12] F. Ebert, C. Finn, S. Dasari, A. Xie, A. Lee, and S. Levine, “Visual foresight: Model-based deep reinforcement learning for vision-based robotic control,” *arXiv preprint arXiv:1812.00568*, 2018.
- [13] K. Black, M. Nakamoto, P. Atreya, H. Walke, C. Finn, A. Kumar, and S. Levine, “Zero-shot robotic manipulation with pretrained image-editing diffusion models,” *arXiv preprint arXiv:2310.10639*, 2023.
- [14] A. Xie, A. Singh, S. Levine, and C. Finn, “Few-shot goal inference for visuomotor learning and planning,” in *Conference on Robot Learning*. PMLR, 2018, pp. 40–52.
- [15] P. Sharma, D. Pathak, and A. Gupta, “Third-person visual imitation learning via decoupled hierarchical controller,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [16] R. Mendonca, O. Rybkin, K. Daniilidis, D. Hafner, and D. Pathak, “Discovering and achieving goals via world models,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 24 379–24 391, 2021.
- [17] A. Simeonov, Y. Du, A. Tagliasacchi, J. B. Tenenbaum, A. Rodriguez, P. Agrawal, and V. Sitzmann, “Neural descriptor fields: Se (3)-equivariant object representations for manipulation,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 6394–6400.
- [18] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. Florence, and A. Zeng, “Code as policies: Language model programs for embodied control,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 9493–9500.
- [19] W. Huang, C. Wang, R. Zhang, Y. Li, J. Wu, and L. Fei-Fei, “Voxposer: Composable 3d value maps for robotic manipulation with language models,” *Proceedings of Machine Learning Research*, vol. 229, 2023.
- [20] W. Huang, C. Wang, Y. Li, R. Zhang, and L. Fei-Fei, “Rekep: Spatio-temporal reasoning of relational keypoint constraints for robotic manipulation,” in *Conference on Robot Learning*. PMLR, 2025, pp. 4573–4602.
- [21] Y. Tang, W. Huang, Y. Wang, C. Li, R. Yuan, R. Zhang, J. Wu, and L. Fei-Fei, “Uad: Unsupervised affordance distillation for generalization in robotic manipulation,” *arXiv preprint arXiv:2506.09284*, 2025.
- [22] M. Xu, Z. Xu, Y. Xu, C. Chi, G. Wetzstein, M. Veloso, and S. Song, “Flow as the cross-domain manipulation interface,” in *Conference on Robot Learning*. PMLR, 2025, pp. 2475–2499.
- [23] C. Yuan, C. Wen, T. Zhang, and Y. Gao, “General flow as foundation affordance for scalable robot learning,” in *Conference on Robot Learning*. PMLR, 2025, pp. 1541–1566.
- [24] T. Weng, S. Bajracharya, Y. Wang, K. Agrawal, and D. Held, “Fabricflownet: Bimanual cloth manipulation with a flow-based policy,” *arXiv preprint arXiv:2111.05623*, 2021.
- [25] A. Goyal, A. Mousavian, C. Paxton, Y.-W. Chao, B. Okorn, J. Deng, and D. Fox, “Ifor: Iterative flow minimization for robotic object rearrangement,” *arXiv preprint arXiv:2202.00732*, 2022.
- [26] D. Seita, Y. Wang, S. J. Shetty, E. Y. Li, Z. Erickson, and D. Held, “Toolflownet: Robotic manipulation with tools via predicting tool flow from point clouds,” *arXiv preprint arXiv:2211.09006*, 2022.
- [27] T. Chen, Y. Mu, Z. Liang, Z. Chen, S. Peng, Q. Chen, M. Xu, R. Hu, H. Zhang, X. Li, and P. Luo, “G3flow: Generative 3d semantic flow for pose-aware and generalizable object manipulation,” *arXiv preprint arXiv:2411.18369*, 2024.
- [28] S. Wang, J. You, Y. Hu, J. Li, and Y. Gao, “Skil: Semantic keypoint imitation learning for generalizable data-efficient manipulation,” in *Robotics: Science and Systems (RSS)*, 2025.
- [29] S. Haldar and L. Pinto, “Point policy: Unifying observations and actions with key points for robot manipulation,” *arXiv preprint arXiv:2502.20391*, 2025.
- [30] S. Guo, X. Liang, J. Lin, Y. Zhuang, L. Lin, and X. Liang, “Actionsink: Toward precise robot manipulation with dynamic integration of action flow,” *arXiv preprint arXiv:2508.03218*, 2025.
- [31] B. Eisner, H. Zhang, and D. Held, “Flowbot3d: Learning 3d articulation flow to manipulate articulated objects,” in *Robotics: Science and Systems (RSS)*, 2022.
- [32] C. Gao, H. Zhang, Z. Xu, Z. Cai, and L. Shao, “Flip: Flow-centric generative planning as general-purpose manipulation world model,” *arXiv preprint arXiv:2412.08261*, 2024.
- [33] J. Guo, X. Ma, Y. Wang, M. Yang, H. Liu, and Q. Li, “Flowdreamer: A rgb-d world model with flow-based motion representations for robot manipulation,” *arXiv preprint arXiv:2505.10075*, 2025.
- [34] H. Zhi, P. Chen, S. Zhou, Y. Dong, Q. Wu, L. Han, and M. Tan, “3dflowaction: Learning cross-embodiment manipulation from 3d flow world model,” *arXiv preprint arXiv:2506.06199*, 2025.
- [35] Y. He and Q. Nie, “Manitrend: Bridging future generation and action prediction with 3d flow for robotic manipulation,” *arXiv preprint arXiv:2502.10028*, 2025.
- [36] Z.-H. Yin, S. Yang, and P. Abbeel, “Object-centric 3d motion field for robot learning from human videos,” *arXiv preprint arXiv:2506.04227*, 2025.
- [37] 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*. Springer, 2024, pp. 306–324.
- [38] P.-C. Ko, J. Mao, Y. Du, S.-H. Sun, and J. B. Tenenbaum, “Learning to act from actionless videos through dense correspondences,” *arXiv preprint arXiv:2310.08576*, 2023.
- [39] Y. Chen, P. Li, Y. Huang, J. Yang, K. Chen, and L. Wang, “Ec-flow: Enabling versatile robotic manipulation from action-unlabeled videos via embodiment-centric flow,” *arXiv preprint arXiv:2507.06224*, 2025.
- [40] L. Shao, P. Shah, V. Dwaracherla, and J. Bohg, “Motion-based object segmentation based on dense rgb-d scene flow,” *arXiv preprint arXiv:1804.05195*, 2018.
- [41] X. Liu, C. R. Qi, and L. J. Guibas, “Flownet3d: Learning scene flow in 3d point clouds,” *arXiv preprint arXiv:1806.01411*, 2018.
- [42] “Zero-shot monocular scene flow estimation in the wild,” *CVPR 2025*, CVF Open Access, 2025.
- [43] “Zeromsf: Zero-shot monocular scene flow in the wild,” *NVIDIA Research Project Page*, 2025.

- [44] J. Sander, G. Caroleo, A. Albin, and P. Maiolino, "Estimating scene flow in robot surroundings with distributed miniaturized time-of-flight sensors," *arXiv preprint arXiv:2504.02439*, 2025.
- [45] A. Byravan and D. Fox, "Se3-nets: Learning rigid body motion using deep neural networks," *arXiv preprint arXiv:1606.02378*, 2016.
- [46] T. Tang, J. Liu, J. Zhang, H. Fu, W. Xu, and C. Lu, "Rfrans: Leveraging refractive flow of transparent objects for surface normal estimation and manipulation," *arXiv preprint arXiv:2311.12398*, 2023.
- [47] B. P. Duisterhof, Z. Mandi, Y. Yao, J.-W. Liu, J. Seidenschwarz, M. Z. Shou, D. Ramanan, S. Song, S. Birchfield, B. Wen, and J. Ichnowski, "Deformgs: Scene flow in highly deformable scenes for deformable object manipulation," *arXiv preprint arXiv:2312.00583*, 2024.
- [48] J. Kerr, C. M. Kim, M. Wu, B. Yi, Q. Wang, K. Goldberg, and A. Kanazawa, "Robot see robot do: Imitating articulated object manipulation with monocular 4d reconstruction," *arXiv preprint arXiv:2409.18121*, 2024.
- [49] H. Jiang, B. Huang, R. Wu, Z. Li, S. Garg, H. Nayyeri, S. Wang, and Y. Li, "Roboexp: Action-conditioned scene graph via interactive exploration for robotic manipulation," *arXiv preprint arXiv:2402.15487*, 2024.
- [50] O. Shorinwa, J. Tucker, A. Smith, A. Swann, T. Chen, R. Firoozi, M. Kennedy III, and M. Schwager, "Splat-mover: Multi-stage, open-vocabulary robotic manipulation via editable gaussian splatting," *arXiv preprint arXiv:2405.04378*, 2024.
- [51] I. Guzey, Y. Dai, G. Savva, R. Bhirangi, and L. Pinto, "Bridging the human to robot dexterity gap through object-oriented rewards," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 3344–3351.
- [52] K. Yu, S. Zhang, H. Soora, F. Huang, H. Huang, P. Tokekar, and R. Gao, "Genflowr: Shaping rewards with generative object-centric flow in visual reinforcement learning," *arXiv preprint arXiv:2508.11049*, 2025.
- [53] H. Bharadhwaj, D. Dwibedi, A. Gupta, S. Tulsiani, C. Doersch, T. Xiao, D. Shah, F. Xia *et al.*, "Gen2act: Human video generation in novel scenarios enables generalizable robot manipulation," *arXiv preprint arXiv:2409.16283*, 2024.
- [54] S. Patel, S. Mohan, H. Mai, U. Jain, S. Lazebnik, and Y. Li, "Robotic manipulation by imitating generated videos without physical demonstrations," *arXiv preprint arXiv:2507.00990*, 2025.
- [55] S. Yang, J. Walker, J. Parker-Holder, Y. Du, J. Bruce, A. Barreto, P. Abbeel, and D. Schuurmans, "Video as the new language for real-world decision making," *arXiv preprint arXiv:2402.17139*, 2024.
- [56] R. McCarthy, D. C. Tan, D. Schmidt, F. Acero, N. Herr, Y. Du, T. G. Thuruthel, and Z. Li, "Towards generalist robot learning from internet video: A survey," *Journal of Artificial Intelligence Research*, vol. 83, 2025.
- [57] H. Wu, Y. Jing, C. Cheang, G. Chen, J. Xu, X. Li, M. Liu, H. Li, and T. Kong, "Unleashing large-scale video generative pre-training for visual robot manipulation," *arXiv preprint arXiv:2312.13139*, 2023.
- [58] Y. Seo, K. Lee, S. L. James, and P. Abbeel, "Reinforcement learning with action-free pre-training from videos," in *International Conference on Machine Learning*. PMLR, 2022, pp. 19 561–19 579.
- [59] J. Wu, H. Ma, C. Deng, and M. Long, "Pre-training contextualized world models with in-the-wild videos for reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 36, pp. 39 719–39 743, 2023.
- [60] J. Yang, B. Liu, J. Fu, B. Pan, G. Wu, and L. Wang, "Spatiotemporal predictive pre-training for robotic motor control," *arXiv preprint arXiv:2403.05304*, 2024.
- [61] Y. Hu, Y. Guo, P. Wang, X. Chen, Y.-J. Wang, J. Zhang, K. Sreenath, C. Lu, and J. Chen, "Video prediction policy: A generalist robot policy with predictive visual representations," *arXiv preprint arXiv:2412.14803*, 2024.
- [62] S. Li, Y. Gao, D. Sadigh, and S. Song, "Unified video action model," *arXiv preprint arXiv:2503.00200*, 2025.
- [63] T. Huang, G. Jiang, Y. Ze, and H. Xu, "Diffusion reward: Learning rewards via conditional video diffusion," in *European Conference on Computer Vision*. Springer, 2024, pp. 478–495.
- [64] A. Escontrela, A. Adeniji, W. Yan, A. Jain, X. B. Peng, K. Goldberg, Y. Lee, D. Hafner, and P. Abbeel, "Video prediction models as rewards for reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 36, pp. 68 760–68 783, 2023.
- [65] A. S. Chen, S. Nair, and C. Finn, "Learning generalizable robotic reward functions from in-the-wild human videos," *arXiv preprint arXiv:2103.16817*, 2021.
- [66] A. Ajay, S. Han, Y. Du, S. Li, A. Gupta, T. Jaakkola, J. Tenenbaum, L. Kaelbling, A. Srivastava, and P. Agrawal, "Compositional foundation models for hierarchical planning," *Advances in Neural Information Processing Systems*, vol. 36, pp. 22 304–22 325, 2023.
- [67] Y. Du, S. Yang, B. Dai, H. Dai, O. Nachum, J. Tenenbaum, D. Schuurmans, and P. Abbeel, "Learning universal policies via text-guided video generation," *Advances in neural information processing systems*, vol. 36, pp. 9156–9172, 2023.
- [68] D. Valevski, Y. Leviathan, M. Arar, and S. Fruchter, "Diffusion models are real-time game engines," *arXiv preprint arXiv:2408.14837*, 2024.
- [69] J. Bruce, M. D. Dennis, A. Edwards, J. Parker-Holder, Y. Shi, E. Hughes, M. Lai, A. Mavalankar, R. Steigerwald, C. Apps *et al.*, "Genie: Generative interactive environments," in *Forty-first International Conference on Machine Learning*, 2024.
- [70] Y. Luo and Y. Du, "Grounding video models to actions through goal conditioned exploration," *arXiv preprint arXiv:2411.07223*, 2024.
- [71] M. Yang, Y. Du, K. Ghasemipour, J. Tompson, D. Schuurmans, and P. Abbeel, "Learning interactive real-world simulators," *arXiv preprint arXiv:2310.06114*, vol. 1, no. 2, p. 6, 2023.
- [72] S. Zhou, Y. Du, J. Chen, Y. Li, D.-Y. Yeung, and C. Gan, "Robodreamer: Learning compositional world models for robot imagination," *arXiv preprint arXiv:2404.12377*, 2024.
- [73] O. Rybkin, K. Pertsch, K. G. Derpanis, K. Daniilidis, and A. Jaegle, "Learning what you can do before doing anything," *arXiv preprint arXiv:1806.09655*, 2018.
- [74] R. Mendonca, S. Bahl, and D. Pathak, "Structured world models from human videos," *arXiv preprint arXiv:2308.10901*, 2023.
- [75] H. Che, X. He, Q. Liu, C. Jin, and H. Chen, "Gamegen-x: Interactive open-world game video generation," *arXiv preprint arXiv:2411.00769*, 2024.
- [76] J. Jang, S. Ye, Z. Lin, J. Xiang, J. Bjorck, Y. Fang, F. Hu, S. Huang, K. Kundalia, Y.-C. Lin *et al.*, "Dreamgen: Unlocking generalization in robot learning through neural trajectories," *arXiv e-prints*, pp. arXiv–2505, 2025.
- [77] J. Liang, R. Liu, E. Ozguroglu, S. Sudhakar, A. Dave, P. Tokmakov, S. Song, and C. Vondrick, "Dreamitate: Real-world visuomotor policy learning via video generation," *arXiv preprint arXiv:2406.16862*, 2024.
- [78] Y. Xiao, J. Wang, N. Xue, N. Karaev, I. Makarov, B. Kang, X. Zhu, H. Bao, Y. Shen, and X. Zhou, "Spatialtrackerv2: 3d point tracking made easy," in *ICCV*, 2025.
- [79] L. Yang, B. Kang, Z. Huang, Z. Zhao, X. Xu, J. Feng, and H. Zhao, "Depth anything v2," *Advances in Neural Information Processing Systems*, vol. 37, pp. 21 875–21 911, 2024.
- [80] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, C. Li, J. Yang, H. Su, J. Zhu *et al.*, "Grounding dino: Marrying dino with grounded pre-training for open-set object detection," *arXiv preprint arXiv:2303.05499*, 2023.
- [81] H.-S. Fang, C. Wang, H. Fang, M. Gou, J. Liu, H. Yan, W. Liu, Y. Xie, and C. Lu, "Anygrasp: Robust and efficient grasp perception in spatial and temporal domains," *IEEE Transactions on Robotics (T-RO)*, 2023.
- [82] G. Pavlakos, D. Shan, I. Radosavovic, A. Kanazawa, D. Fouhey, and J. Malik, "Reconstructing hands in 3D with transformers," in *CVPR*, 2024.
- [83] C. M. Kim\*, B. Yi\*, H. Choi, Y. Ma, K. Goldberg, and A. Kanazawa, "Pyroki: A modular toolkit for robot kinematic optimization," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025.
- [84] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 80. PMLR, 10–15 Jul 2018, pp. 1861–1870.
- [85] Y. Zhu, J. Wong, A. Mandlekar, R. Martín-Martín, A. Joshi, K. Lin, S. Nasiriany, and Y. Zhu, "robosuite: A modular simulation framework and benchmark for robot learning," in *arXiv preprint arXiv:2009.12293*, 2020.
- [86] T. Wan, A. Wang, B. Ai, B. Wen, C. Mao, and et al., "Wan: Open and advanced large-scale video generative models," *arXiv preprint arXiv:2503.20314*, 2025.
- [87] Google DeepMind, "Veo-3 Technical Report," 2025. [Online]. Available: <https://storage.googleapis.com/deepmind-media/veo/Veo-3-Tech-Report.pdf>