Dreamitate:

Real-World Visuomotor Policy Learning via Video Generation

¹Columbia University ²Toyota Research Institute ³Stanford University

dreamitate.cs.columbia.edu




Figure 1: **Real-World Visuomotor Policy Learning via Video Generation.** Dreamitate is a visuomotor policy learning framework that fine-tunes a video generative model to synthesize videos (indicated by of humans using tools to complete a task. The tool's trajectory in the generated video is tracked, and the robot executes this trajectory to accomplish the task in the real-world.

Abstract: A key challenge in manipulation is learning a policy that can robustly generalize to diverse visual environments. A promising mechanism for learning robust policies is to leverage video generative models that are pretrained on large-scale datasets of internet videos. In this paper, we propose a visuomotor policy learning framework that fine-tunes a video diffusion model on human demonstrations of a given task. At test time, we generate a video showing an execution of the task conditioned on images of a novel scene, and use this synthesized video to directly control the robot. Our key insight is that using common tools allows us to effortlessly bridge the embodiment gap between the human hand and the robot manipulator. We evaluate our approach on four tasks of increasing complexity and demonstrate that harnessing internet-scale generative models allows the learned policy to achieve a significantly higher degree of generalization than existing behavior cloning approaches.

Keywords: Visuomotor Policy Learning, Imitation Learning, Video Generation

1 Introduction

Learning visuomotor policies that decide actions from perception has been a longstanding challenge in robotics because they need to generalize to diverse physical situations [1]. Recently, behavior cloning from human demonstrations [2, 3] has become the method of choice for many manipulation

^{*} Equal contribution.

tasks, where visuomotor policy learning is formulated as a regression problem to map visual observations to actions supervised by human demonstrations. However, behavior cloning requires ground truth robot actions, which makes scaling to diverse situations challenging.

We aim to learn generalizable visuomotor policies by leveraging video generation. Trained on large-scale datasets of internet videos, video generation provides a promising avenue for generalization because extensive priors from human behavior can be transferred into robot behavior. Many recent methods have also explored how to capitalize on these priors, such as synthesizing videos of human behavior [4] or synthesizing videos of robot behavior [5]. However, while videos of people capture the diversity of behavior, generated human actions are difficult to transfer to robots due to the embodiment gap. While directly synthesizing videos of robots would be more realizable physically, the scale of available data is many orders of magnitude smaller than videos of humans, making the learned policy less robust in an in-the-wild environment.

How else can we leverage video generation for policy learning? In this work, we introduce Dreamitate, ¹ a visuomotor policy that controls robot behavior via conditional video generation. Our key insight is that the end-effector is the most important part of the embodiment for manipulation: its interaction with diverse objects is challenging to model explicitly and can benefit most from commonsense knowledge from video generation. The rest of robot embodiment can be solved through inverse kinematics. Given a visual observation, our policy synthesizes a video of a person performing tasks with a tool. To translate into robot behaviors, we 3D track the tool in the synthesized video, and simply transfer the trajectory into explicit robot actions. Figure 1 shows an example of the actions created from this policy.

Our formulation offers several major advantages compared to traditional visuomotor policies:

- **Generalizability.** Our underlying video generation model is pretrained on a massive amount of videos available on the internet, including a wide variety of human manipulation in all kinds of environments. By formulating policy learning as video model finetuning, the prior knowledge learned from pre-taining can be preserved, allowing our model to learn generalizable manipulation skills with the same amount of demonstration data.
- Scalability. Our finetuning videos are recordings of human demonstrations instead of teleoperation, making our dataset collection more scalable.
- **Interpretability.** During inference time, our video model predicts future execution plans in the form of videos before the actual robot executions are performed. Compared to a black-box end-to-end policy, our formulation offers an intermediate representation of policy that is interpretable to humans, a key feature for human-robot interaction applications.

We evaluated our approach across four real-world tasks, including bimanual manipulation, precise 3D manipulation, and long-horizon tasks, using a small number of expert human demonstrations. We found that the video model consistently outperformed the baseline behavior cloning model in generalizing to unseen scenarios. Additionally, we analyzed the model's performance when scaling down the training dataset and found that the video model maintained strong generalization performance even with fewer demonstrations. We will release the code and data for reproducing our results. Please refer to the supplementary material for the robot videos.

2 Related Work

Behavior Cloning. Behavior cloning (BC) has emerged as a leading approach in manipulation tasks, learning policies from demonstrations through supervised learning. Early BC efforts involved end-to-end models mapping states to actions, but struggled with multimodal behaviors and high-precision tasks [6, 7, 8]. Subsequently, Energy-Based Models (EBMs) were explored, predicting actions by minimizing energy in sequence optimization [9, 10, 11, 12]. Recently, conditional generative models have shown promise in capturing multimodal demonstration behaviors, enhancing

¹Dream then imitate




Figure 2: **Method Overview.** For each task, (Step 1) we capture stereo camera recordings of human demonstrations of tool use from a top and a side view. (Step 2) The video model is fine-tuned to recreate these demonstrations from an initial scene image. (Step 3) In a new, unseen scenario, a stereo image pair is provided to the video model to generate () the manipulation. The tool's trajectory in the generated video is tracked and executed by the robot to complete the task.

task success rates [2, 3, 13]. Different from previous models, we integrate video prediction and 3D tracking to predict action prediction.

Visual Pretraining for Policy Learning. Prior works have extensively explored various ways to pre-train the perception model in a visuomotor policy for learning more robust visual representation. One of the most popular pertaining objectives has been video prediction [14, 15, 16, 17, 18]. By learning to predict future based on the current observation, a model can learn the dynamic and causality of the world crucial for physical interaction. Contrastive learning [15, 19, 20, 21] as well as masked autoencoding [22, 23, 24] are two very popular self-supervised learning objectives for learning visual representation for robotics. Besides, another line of work [25, 26, 27] studies learning a generalizable reward function through visual pertaining for reinforcement learning.

Video Models for Decision-Making. More recently, in light of the fast progress of text-to-video generative models [28, 29, 30, 31, 32], the hope to leverage internet-scale video pertaining for robotics is rekindled [33]. A line of work uses video generative model as a world simulator [34, 5, 28], which can predict future video conditioned on an action. Another line of work uses video-language model for longer-horizon planning [35, 36, 37]. Different from prior works on this topic which typically use video prediction models as a world model or simulator in order to perform planning and decision making, Dreamitate uses video prediction model as part of a visuomotor policy and directly uses the predicted video for action prediction.

3 Approach

3.1 Overview

Given a video frame of the scene v_0 , our goal is to plan and execute robot actions $a_t \in SE(3)$ via video generation. We approach this problem through the framework:

$$a_t = \mathcal{T}(\hat{v}_t)$$
 where $\{\hat{v}_t\} = f_{\theta}(v_0)$ (1)

where $f_{\theta}(\cdot)$ is a generative video model with learnable parameters θ and \hat{v}_t is a synthesized frame. We train f to generate videos of a human doing the task using a tool that we can track in 3D. \mathcal{T} tracks the trajectory of the tool in the generated video \hat{v}_t , which we can directly use to control the endeffector state to perform manipulation tasks. Figure 2 visualizes this framework. This composition of synthesize-then-track facilitates cheap data collection without teleoperation. Moreover, it aligns with pretaining video data, which is composed of mostly videos showing human behavior.

3.2 Video Generation

Starting with a video generator pre-trained on internet videos, we fine-tune it on a small video demonstration dataset in order to combine the priors from internet videos with behaviors from demonstration videos. To collect videos for fine-tuning, we build a tabletop setup with two calibrated cameras positioned 45 degrees apart for better visibility. We capture several pair of stereo videos $(\hat{v}^1, \hat{v}^2) \in \mathcal{V}$ of humans performing the task using a 3D printed tool (see Figure 3 for details).




Figure 3: Aligning Human and Robot Behavior via **Trackable** Tools. We use custom tools with CAD models known for human demonstrations, which facilitate precise 3D tracking, across four real-world manipulation tasks. design unique tools (e.g., a spoon for scooping) to show the generalizability of our approach for varying tasks.

We use these demonstration videos to fine-tune the video generative model. Since we will imitate the trajectory of the synthesized tool in the physical world, we generate stereo videos to make action possible in 3D. Let $(\hat{v}_t^1, \hat{v}_t^2)$ be the predicted stereo video frames at time t and (v_t^1, v_t^2) be the ground truth stereo frames from demonstration videos. We optimize the video prediction objective:

$$\min_{\theta} \mathbb{E}_{v \in \mathcal{V}} \left[\sum_{t=1}^{T} \left| \left| \left(\hat{v_t}^1 - v_t^1 \right) \right| \right|_2 + \left| \left| \left(\hat{v_t}^2 - v_t^2 \right) \right| \right|_2 \right] \quad \text{for} \quad \{ \hat{v}_t \} = f_{\theta}(v_0)$$
 (2)

We initialize θ to the pre-trained weights learned from large-scale internet video datasets (Stable Video Diffusion [29]). We train a separate f_{θ} for each task (e.g. sweeping or scooping).

Following the implementation from [38], the encoder and decoder are frozen such that only the spatial/temporal attention layers are fine-tuned. The per-frame image embedding input to the model is modified based on the viewing angle of the output frame to facilitate stereo video generation. This ensures that the first half of the frames generate the first view and the half generates the second view, using the initial image embedding from the respective view.

At test time, we will be given a stereo image pair of the scene that we wish the robot to manipulate. We apply the fine-tuned video model $f(v_0)$ and obtain a generated stereo video frames \hat{v}_t of the task being performed.

3.3 Track then Act

The synthesized video frames $\{\hat{v}_t\}$ serve as the intermediate representation of our policy to obtain an action trajectory, $\{a_t\}$, that we then execute on the robot. Actions $a_t \in SE(3)$ are represented as the 6D pose of the tool relative to the camera. By using a known CAD model, we can efficiently and accurately track the tool in each frame of the generated videos and therefore obtain precise locations of the tool. Each action a_i in the action trajectory corresponds to one frame from the synthesized video. For 3D consistency, we use the generated stereo pair, along with the calibrated camera parameters to resolve 6D pose of the tool. For tracking the tool from an RGB image, we use Megapose [39] and operate on 768×448 resolution video frames with camera intrinsic derived from the default values of the Intel Realsense camera. Finally, during execution, the customized tool is mounted on the robot arm and actions $a_0, ..., a_T$ are executed by the robot arm. The entire data-collection, training, and fine-tuning pipeline are agent-independent. This is possible because we use tool's as manipulators and track known tools in generated videos, removing the need for learning any agent-specific action spaces.

4 Experiments

In this section, we evaluate the performance of our video model in real-world robot manipulation tasks. The evaluation, summarized in Table 1, encompasses four distinct tasks including object rotation, granular material scooping, table top sweeping and shape pushing (see Figure 3 for a visualization of the used tools). The details of each task are further explained in the corresponding sections below. For an assessment of our model's capabilities, we benchmark its performance against Diffusion Policy [40] as a baseline.




Figure 4: **Rotation Qualitative Results.** For this task, we provide as input the image on the left for each row, with an example of a successful rotation overlaid. Our video generation-based approach succeeds, while Diffusion Policy fails to select stable grasping points on the object, causing it to slip during manipulation. Generated videos are indicated by

4.1 Experimental Setup

To evaluate the model's generalizability, the data was collected by a human demonstrator at different locations for training and evaluation using the same camera setup, though the exact tabletop and lighting conditions varied. To further test the model's robustness in real-world scenarios, we ensured that the sets of objects used in training and evaluation are non-overlapping. To initialize tool tracking with Megapose [39], background subtraction and hand removal based on skin color were employed. In rare cases where this process failed, human corrections were applied to the bounding boxes (both for our method and for the baseline).

We used the same training data to train Diffusion Policy with stereo image input. The videos were preprocessed using Megapose for tool tracking, providing target trajectories for training. We used a pretrained ResNet-18 variant [40] as the baseline, as the pretrained CLIP encoder [41] variant showed lower performance in our tasks. We trained Diffusion Policy for 200 epochs to predict actions for the next 12 time steps, maintaining an open-loop system without subsampling the training data, similar to the video model.

	Training Objects	Demonstrations	Test Objects	Test Trials
Rotation	31	371	10	40
Scooping	17 Bowls, 8 Particles	368	8 Bowls, 4 Particles	40
Sweeping	6 Particles	356	6 Particles	40
Push-Shape	26 Letters	727	8 Shapes	32

Table 1: **Tasks Summary.** We detail the train and test setup for each task above. For each task, we use a distinct set of objects during testing than training.

4.2 Object Rotation Task

Task. We design a rotation task, shown in Figure 4, to test our policy's ability to coordinate end-effectors and choose appropriate grasping points to manipulate real-world objects. In each training demonstration, we randomly place an object on the table, and use a gripper, shown in Figure 3 (top), to grasp the object at appropriate points and rotate it counterclockwise up to 45 degrees.

Setup. We collect training data with 31 objects, with 14 real-world objects (such as boxes or a hammer) and 17 custom colored shapes made out of foam. We use 10 unseen real-world objects for evaluation. We mark a trajectory as successful if the robot makes and maintains contact with the object throughout the rotation, and rotates it at least 25 degrees counterclockwise.

Results. As shown in Table 2, our policy significantly outperforms Diffusion Policy (92.5% vs. 55%). Fig. 4 illustrates cases where our policy succeeds while Diffusion Policy fails. We observe



Figure 5: **Scooping Qualitative Results.** We provide the input image on the left (without the overlays provided for illustration) to our model and Diffusion Policy, and show the output trajectory on the right. Our approach succeeds in most trials, while Diffusion Policy is often distracted by other objects in the scene, placing material in the wrong container and failing to scoop accurately. Generated videos are indicated by

that Diffusion Policy can fail to move the end-effector into contact with the object. In more challenging cases, Diffusion Policy is prone to selecting unstable grasping points causing the object to slip during manipulation. In contrast, our policy consistently makes contact with the object, with limited failures when selecting appropriate grasping points for particularly challenging shapes, such as a transparent bag with toys inside.

	Rotation	Scooping	Sweeping	Push-Shape	
Model	Success Rate	Success Rate	Success Rate	mIoU	Rot. Error
Diffusion Policy [2] Ours	22 / 40 37 / 40	22 / 40 34 / 40	5 / 40 37 / 40	0.550 0.731	48.2° 8.0 °

Table 2: **Quantitative Results.** We compare our method to Diffusion Policy on four tasks quantitatively. We report success rates (successful trials / total trials) for rotation, scooping, and sweeping tasks. For Push-Shape, we report mean intersection-over-union (mIoU) and average rotation error (Rot. Error). Our approach performs well across all tasks, whereas Diffusion Policy shows worse performance overall and degrades in the more challenging sweeping and Push-Shape scenarios.

4.3 Granular Material Scooping Task

Task. This task, shown in Figure 5, requires scooping granular material (e.g., beans) from a full container to an empty one while avoiding distractor objects. This task requires the policy to perform precise manipulation of a scooping tool, identify the full and the empty containers, along with their precise locations, in arbitrary positions, and ignore distractors.

Setup. We collect demonstrations with 17 bowls and 8 colored beans with only 1 distraction object in the scene at a time. At test time, we use an unseen set of 8 bowls, 4 new colored particles, and 5 distractor objects per trial sampled from a fixed set of 15 everyday objects. We mark a trajectory as successful if the robot transfers any particles to the empty bowl.

Results. As shown in Table 2 (middle), our policy significantly outperforms Diffusion Policy (85% vs. 55%). This is a particularly challenging task for our approach, due to the small target (the scooper) for object tracking and video generation. This demonstrates that stereo video generation can accurately determine the object's pose to perform 3D manipulation. By comparison, Diffusion Policy is often distracted by objects in the scene, placing material in the wrong container, and failing more frequently in scooping due to misjudging the bowl's height on the table, as illustrated in Fig. 5.

4.4 Table Top Sweeping Task

Task. The sweeping task, shown in Figure 6, requires policies to use a brush to sweep randomly placed particles to a target location marked by a randomly placed star, while avoiding obstacles. This task is designed to test policy's ability to handle multi-modal distributions in the training data.




Figure 6: **Sweeping Qualitative Results.** We provide the input image on the left (minus the white and green circles overlaid for illustration) to our model and Diffusion Policy, and show the output trajectories on the right. Our approach again achieves high success rate even on this challenging, multi-modal task, whereas Diffusion Policy generates trajectories that often collide with obstacles and fail to sweep to the target. Generated videos are indicated by

Setup. Training data includes 6 colored particles and 25 distraction objects, with 5 to 6 distractions at a time. At test time, we use 6 new particles and 15 unseen distractions, with 2 to 4 beans and 5 distractions randomly placed on the table. The training data contains some variations with multiple ways to achieve the goal, such as choosing which particle to sweep in the scene. Due to the small size of the particles, we remove the final pooling layer from Diffusion Policy ResNet-18 encoder to provide a higher spatial resolution and improve baseline performance. We mark a trajectory as successful if any particle is transferred to be within 50 mm from the target.

Results. In Table 2 (right), we observe that the video model maintains strong performance in this task with a 92.5% success rate. In contrast, Diffusion Policy generates trajectories that often collide with obstacles and fail to sweep to the target (see Fig. 6), achieving only a 12.5% success rate. This experiment demonstrates that the capitalizing on internet-scale pre-trained video generation models allows to better handle multi-modal demonstrations and achieve a much larger degree of generalization in this challenging scenario.

4.5 Push-Shape Task (Long Horizon)

Task. Push-Shape, shown in Figure 7, is a challenging version of the long-horizon Push-T task [2]: we place a foam shape on the table, and task the robot with pushing the object to a specified target goal mask (given as input to the policy) over consecutive steps. This is a challenging task as it requires adjusting the position and orientation of the shape, requiring predicting the shape's movement, which depends on the contact between the object and the end-effector, as well as material properties of the table.

Setup. We train with a set of 26 foam objects each in the shape of a letter from the alphabet, and test on 8 unseen foam shapes (including digits and polygons). As this is a challenging, multi-step task, we score the best of 4 rollouts for each trial, and report the mean intersect-over-union (mIoU) with the target mask as well as the average rotation error from the target.

Results. As shown in Table 2, over the 32 trials, the video model achieved a significantly higher mIoU of 0.731 compared to 0.550 mIoU from Diffusion Policy, and a rotation error of 8.0 degrees compared to 48.2 degrees from Diffusion Policy. We find that Diffusion Policy tends to push the object to the target, but fails to effectively rotate the object to match the mask. By contrast, the video model produces appropriate pushing actions to adjust the object's position and orientation, thereby reducing the rotation error. Qualitative examples of the model rollouts are illustrated in Fig. 7.




Figure 7: **Push-Shape Qualitative Results.** We provide the image on the left, including the gray overlay indicating the target object position and orientation, to our model and Diffusion Policy, and show output trajectories on the right. Our method produces appropriate pushing actions to adjust the object's position. In contrast, Diffusion Policy fails to effectively rotate the object to match the mask. Generated videos (top row, bottom right) are indicated by

4.6 Performance Scaling Curve

We studied the impact of training set size on the generalization ability of our policy and compared it with Diffusion Policy using a rotation task. Both models performed well with the full dataset. When re-trained with two-thirds and one-third of the dataset and tested over 40 trials, the results depicted in Figure 8 indicate that while Diffusion Policy's performance declines significantly with reduced data. In contrast, our model remains stable and maintains high success rates even with only one-third of the data. This highlights our model's superior generalization, which can be attributed to the extensive pre-training on Internet-scale video generation models.

5 Limitations

We highlight a few limitations and directions for future work. By visually tracking the tools used for manipula-




Figure 8: **Number of Demonstrations.** Our approach maintains strong generalization performance with less training data compared to Diffusion Policy.

tion, our implementation is limited to generating visually trackable robot actions. Although advances in video generation and object tracking models will improve accuracy, our approach can fail when the end-effector is heavily occluded. Additionally, reliance on rigid tools limits the applicability of our approach to the task requiring fine-grained control. Finally, video models have higher computational costs, making real-time closed-loop control infeasible, though this can be mitigated with recent advances [42, 43] in accelerating video model inference.

6 Conclusion

In this work, we investigate how video generative models can be used to learn generalizable visuomotor policies. We propose to fine-tune a video diffusion model on human demonstrations to synthesize an execution plan in the form of videos at test-time. These synthesized execution videos are then directly used to control the robot. Our key insight is that using common tools allows us to easily bridge the embodiment gap between the human hand in the demonstrations and the robot manipulator in the real-world. Our experiments validate that capitalizing on Internet-scale video diffusion models allows our approach to achieve a much larger degree of generalization compared to previous behavior cloning methods like Diffusion Policy.

Acknowledgments

This research is based on work partially supported by the Toyota Research Institute and the NSF NRI Award #2132519.

References

- [1] A. Khazatsky, K. Pertsch, S. Nair, A. Balakrishna, S. Dasari, S. Karamcheti, S. Nasiriany, M. K. Srirama, L. Y. Chen, K. Ellis, P. D. Fagan, J. Hejna, M. Itkina, M. Lepert, Y. J. Ma, P. T. Miller, J. Wu, S. Belkhale, S. Dass, H. Ha, A. Jain, A. Lee, Y. Lee, M. Memmel, S. Park, I. Radosavovic, K. Wang, A. Zhan, K. Black, C. Chi, K. B. Hatch, S. Lin, J. Lu, J. Mercat, A. Rehman, P. R. Sanketi, A. Sharma, C. Simpson, Q. Vuong, H. R. Walke, B. Wulfe, T. Xiao, J. H. Yang, A. Yavary, T. Z. Zhao, C. Agia, R. Baijal, M. G. Castro, D. Chen, Q. Chen, T. Chung, J. Drake, E. P. Foster, J. Gao, D. A. Herrera, M. Heo, K. Hsu, J. Hu, D. Jackson, C. Le, Y. Li, K. Lin, R. Lin, Z. Ma, A. Maddukuri, S. Mirchandani, D. Morton, T. Nguyen, A. O'Neill, R. Scalise, D. Seale, V. Son, S. Tian, E. Tran, A. E. Wang, Y. Wu, A. Xie, J. Yang, P. Yin, Y. Zhang, O. Bastani, G. Berseth, J. Bohg, K. Goldberg, A. Gupta, A. Gupta, D. Jayaraman, J. J. Lim, J. Malik, R. Martín-Martín, S. Ramamoorthy, D. Sadigh, S. Song, J. Wu, M. C. Yip, Y. Zhu, T. Kollar, S. Levine, and C. Finn. Droid: A large-scale in-the-wild robot manipulation dataset, 2024.
- [2] 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.
- [3] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.
- [4] H. Bharadhwaj, A. Gupta, V. Kumar, and S. Tulsiani. Towards generalizable zero-shot manipulation via translating human interaction plans. *arXiv* preprint arXiv:2312.00775, 2023.
- [5] Y. Du, M. Yang, B. Dai, H. Dai, O. Nachum, J. B. Tenenbaum, D. Schuurmans, and P. Abbeel. Learning universal policies via text-guided video generation, 2023.
- [6] D. A. Pomerleau. Alvinn: An autonomous land vehicle in a neural network. *Advances in neural information processing systems*, 1, 1988.
- [7] T. Zhang, Z. McCarthy, O. Jow, D. Lee, X. Chen, K. Goldberg, and P. Abbeel. Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. In 2018 IEEE international conference on robotics and automation (ICRA), pages 5628–5635. IEEE, 2018.
- [8] P. Florence, L. Manuelli, and R. Tedrake. Self-supervised correspondence in visuomotor policy learning, 2019.
- [9] Y. LeCun, S. Chopra, R. Hadsell, M. Ranzato, and F. Huang. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006.
- [10] Y. Du and I. Mordatch. Implicit generation and generalization in energy-based models, 2020.
- [11] 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, 2023.
- [12] P. Florence, C. Lynch, A. Zeng, O. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson. Implicit behavioral cloning, 2021.
- [13] S. Lee, Y. Wang, H. Etukuru, H. J. Kim, N. M. M. Shafiullah, and L. Pinto. Behavior generation with latent actions. *arXiv preprint arXiv:2403.03181*, 2024.
- [14] C. Finn, I. Goodfellow, and S. Levine. Unsupervised learning for physical interaction through video prediction, 2016.

- [15] P. Sermanet, C. Lynch, Y. Chebotar, J. Hsu, E. Jang, S. Schaal, and S. Levine. Time-contrastive networks: Self-supervised learning from video, 2018.
- [16] M. Babaeizadeh, C. Finn, D. Erhan, R. H. Campbell, and S. Levine. Stochastic variational video prediction. *arXiv preprint arXiv:1710.11252*, 2017.
- [17] A. X. Lee, R. Zhang, F. Ebert, P. Abbeel, C. Finn, and S. Levine. Stochastic adversarial video prediction, 2018.
- [18] D. Suris, R. Liu, and C. Vondrick. Learning the predictability of the future. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 12607–12617, June 2021.
- [19] A. Srinivas, M. Laskin, and P. Abbeel. Curl: Contrastive unsupervised representations for reinforcement learning, 2020.
- [20] S. Nair, A. Rajeswaran, V. Kumar, C. Finn, and A. Gupta. R3m: A universal visual representation for robot manipulation, 2022.
- [21] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning transferable visual models from natural language supervision, 2021.
- [22] Y. Seo, D. Hafner, H. Liu, F. Liu, S. James, K. Lee, and P. Abbeel. Masked world models for visual control. In K. Liu, D. Kulic, and J. Ichnowski, editors, *Proceedings of The 6th Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research*, pages 1332–1344. PMLR, 14–18 Dec 2023. URL https://proceedings.mlr.press/v205/seo23a.html.
- [23] I. Radosavovic, T. Xiao, S. James, P. Abbeel, J. Malik, and T. Darrell. Real-world robot learning with masked visual pre-training. In K. Liu, D. Kulic, and J. Ichnowski, editors, Proceedings of The 6th Conference on Robot Learning, volume 205 of Proceedings of Machine Learning Research, pages 416–426. PMLR, 14–18 Dec 2023. URL https://proceedings. mlr.press/v205/radosavovic23a.html.
- [24] I. Radosavovic, B. Shi, L. Fu, K. Goldberg, T. Darrell, and J. Malik. Robot learning with sensorimotor pre-training, 2023.
- [25] Y. J. Ma, W. Liang, V. Som, V. Kumar, A. Zhang, O. Bastani, and D. Jayaraman. Liv: Language-image representations and rewards for robotic control, 2023.
- [26] A. S. Chen, S. Nair, and C. Finn. Learning generalizable robotic reward functions from "in-the-wild" human videos, 2021.
- [27] 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, 2023.
- [28] T. Brooks, B. Peebles, C. Holmes, W. DePue, Y. Guo, L. Jing, D. Schnurr, J. Taylor, T. Luhman, E. Luhman, C. Ng, R. Wang, and A. Ramesh. Video generation models as world simulators. 2024. URL https://openai.com/research/video-generation-models-as-world-simulators.
- [29] A. Blattmann, T. Dockhorn, S. Kulal, D. Mendelevitch, M. Kilian, D. Lorenz, Y. Levi, Z. English, V. Voleti, A. Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.
- [30] A. Blattmann, R. Rombach, H. Ling, T. Dockhorn, S. W. Kim, S. Fidler, and K. Kreis. Align your latents: High-resolution video synthesis with latent diffusion models, 2023.

- [31] S. Zhang, J. Wang, Y. Zhang, K. Zhao, H. Yuan, Z. Qin, X. Wang, D. Zhao, and J. Zhou. I2vgen-xl: High-quality image-to-video synthesis via cascaded diffusion models, 2023.
- [32] J. Ho, W. Chan, C. Saharia, J. Whang, R. Gao, A. Gritsenko, D. P. Kingma, B. Poole, M. Norouzi, D. J. Fleet, and T. Salimans. Imagen video: High definition video generation with diffusion models, 2022.
- [33] 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, 2024.
- [34] Z. Yang, Y. Chen, J. Wang, S. Manivasagam, W.-C. Ma, A. J. Yang, and R. Urtasun. Unisim: A neural closed-loop sensor simulator, 2023.
- [35] Y. Du, M. Yang, P. Florence, F. Xia, A. Wahid, B. Ichter, P. Sermanet, T. Yu, P. Abbeel, J. B. Tenenbaum, L. Kaelbling, A. Zeng, and J. Tompson. Video language planning, 2023.
- [36] 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, 2023.
- [37] K. Black, M. Nakamoto, P. Atreya, H. Walke, C. Finn, A. Kumar, and S. Levine. Zeroshot robotic manipulation with pretrained image-editing diffusion models. *arXiv* preprint *arXiv*:2310.10639, 2023.
- [38] B. Van Hoorick, R. Wu, E. Ozguroglu, K. Sargent, R. Liu, P. Tokmakov, A. Dave, C. Zheng, and C. Vondrick. Generative camera dolly: Extreme monocular dynamic novel view synthesis. *arXiv* preprint arXiv:2405.14868, 2024.
- [39] Y. Labbé, L. Manuelli, A. Mousavian, S. Tyree, S. Birchfield, J. Tremblay, J. Carpentier, M. Aubry, D. Fox, and J. Sivic. Megapose: 6d pose estimation of novel objects via render & compare. arXiv preprint arXiv:2212.06870, 2022.
- [40] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song. Diffusion policy: Visuomotor policy learning via action diffusion. *arXiv preprint arXiv:2303.04137*, 2023.
- [41] C. Chi, Z. Xu, C. Pan, E. Cousineau, B. Burchfiel, S. Feng, R. Tedrake, and S. Song. Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots. In *Proceedings of Robotics: Science and Systems (RSS)*, 2024.
- [42] S. Luo, Y. Tan, S. Patil, D. Gu, P. von Platen, A. Passos, L. Huang, J. Li, and H. Zhao. Lcm-lora: A universal stable-diffusion acceleration module. *arXiv preprint arXiv:2311.05556*, 2023.
- [43] X. Li, Y. Liu, L. Lian, H. Yang, Z. Dong, D. Kang, S. Zhang, and K. Keutzer. Q-diffusion: Quantizing diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17535–17545, 2023.
- [44] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

A Video Model Implementation

We adapted the pretrained Stable Video Diffusion model [29], which generates 25-frame videos at a time. In our adaptation, the first 13 frames correspond to the stereo view 1, and the last 12 frames correspond to the stereo view 2, captured from the two cameras. To condition the video model to generate a stereo video, we modified the per-frame image embedding based on the viewing angle of each output frame. Since each frame of the stereo videos should be paired but the video model generates an odd number of frames, the first frame of the video model output is always the same as the input and discarded at test time. The model training hyperparameters are given in Table 3. During inference, we use 30 denoising steps with a constant classifier-free guidance of 1.0. Additional qualitative results for the four tasks are shown in Fig 9, 10, 11 and 12.

H-Param	Res	Lr	Batch Size	Train Steps	Clip Duration	Fps	MotScr
Rotation (Full DS)	768×448	1e-5	4	16384	2.0	6	200
Rotation (2/3 DS)	768×448	1e-5	4	16384	2.0	6	200
Rotation (1/3 DS)	768×448	1e-5	3	15360	2.0	6	200
Scooping	768×448	1e-5	4	16384	3.0	5	200
Sweeping	768×448	1e-5	4	16384	3.0	5	200
Push-Shape	768×448	1e-5	4	17408	2.0	6	200

Table 3: **Hyperparameters for Video Model Training.** Res: image and video resolution, Lr: learning rate, Batch Size: batch size, Training Steps: training steps for the evaluation checkpoint, Clip Duration: single demonstration video length in seconds, Fps: the video sub-sampling frame rate and model fps parameter, MotScr: model motion score parameter.

B Experimental Setup

The stereo camera setup consists of two Intel RealSense D435i cameras spaced approximately 660 mm apart at a 45° angle. The distance between the cameras and the table is about 760 mm. The real-world data collection and the robot experiment setups are shown in Fig. 17 and 18. The training videos are recorded at a resolution of 1280×720 and are then cropped and resized to the appropriate resolution for model input. The table surface used for data collection and experiments is covered with a black cloth, which introduces variations in friction and increases uncertainty in the Push-Shape experiments. For the rotation and scooping experiments, UFACTORY xArm 7 robots are used, while UR5 robots are used for the sweeping and Push-Shape experiments. For calculating the mIoU in the Push-Shape experiment, the view from the stereo camera 1 is used. In each trial with multiple steps, the resulting image with the highest IoU with the target is used to calculate the rotation error. In sweeping and push-shape experiments, the robot end-effector height is limited to avoid robot collision with the table top.

C Data Collection

For all tasks, the first frame of the human demonstration video is an image of the scene. The subsequent frames include the human demonstrator using the tool to perform the manipulation. In the Push-Shape demonstration, an object is pushed to a location in multiple steps. The final position of the object is used as a mask and blended with the entire video for the target position. The objects used in training and testing for different tasks are shown in Fig. 13, 14, 15 and 16.

D Object Tracking

In the videos, the tool is tracked using MegaPose [39]. Utilizing a stereo setup, the center of the tracked object from each camera are projected into 3D space as a straight line. The translation component of the object in 3D space is determined by finding the midpoint between the projected lines from the two cameras. The rotation component of the object is obtained by averaging the

object rotations from the two views. This refined object pose from the stereo setup enhances the accuracy of the object's depth measurement from the cameras. In the scooping task, only the handle of the scooper is tracked to avoid inaccuracies due to occlusion by particles. In the sweeping and Push-Shape tasks, the tool without the handle is tracked, as the handle is occluded by the human hand. To obtain the tool trajectories for training Diffusion Policy, the same stereo tracking is applied to the demonstration videos.

E Diffusion Policy Baseline

We use a CNN-based Diffusion Policy as our baseline, employing two pretrained ResNet-18 [44] image encoders to process the stereo images of the scene. The input images have a resolution of 384×224, similar to the original implementation resolution. We found that higher resolution input images did not improve model performance.



Figure 9: **Additional Rotation Qualitative Results.** The trajectories of the end-effectors are projected onto the input image.




Figure 10: **Additional Scooping Qualitative Results.** The trajectory of the end-effector is projected onto the input image.



Figure 11: Additional Sweeping Qualitative Results. The trajectory of the end-effector is projected onto the input image.




Figure 12: **Additional Push-Shape Qualitative Results.** The trajectory of the end-effector is projected onto the input image.



Figure 13: **Rotation Objects.** The training set includes 14 real-world objects and 17 custom colored shapes made out of foam. The testing set includes 10 challenging real-world objects.



Figure 14: **Scooping Objects.** The training set includes 17 bowls, 8 colored beans and 4 real-world objects. The testing set includes 8 bowls, 4 colored particles and 15 real-world distraction objects. During data collection, both the real-world objects and inverted bowls are used as the single distraction object in the scene.





Figure 15: **Sweeping Objects.** The training set includes 25 distraction objects and 6 colored beans, with 5 to 6 distractions in the scene at a time during data collection. The testing set includes 15 real-world distraction objects and 6 colored particles. The star is used as the target in the experiments.



Figure 16: **Push-Shape Objects.** The training set includes all 26 capital letters of the alphabet, while the testing set consists of 8 shapes, including digits and polygons.






Figure 18: **Real-world Robot Experiment Setup.** The robot experiment setup includes the 3 robots to perform all the experiments.