

Track2Act: Predicting Point Tracks from Internet Videos enables Generalizable Robot Manipulation

Homanga Bharadhwaj¹, Roozbeh Mottaghi^{2,*}, Abhinav Gupta^{1,*}, and Shubham Tulsiani^{1,*}

¹Carnegie Mellon University and ²FAIR at Meta

Abstract. We seek to learn a generalizable goal-conditioned policy that enables *diverse* robot manipulation — interacting with unseen objects in novel scenes without test-time adaptation. While typical approaches rely on a large amount of demonstration data for such generalization, we propose an approach that leverages web videos to predict plausible interaction plans and learns a task-agnostic transformation to obtain robot actions in the real world. Our framework, *Track2Act* predicts tracks of how points in an image should move in future time-steps based on a goal, and can be trained with diverse videos on the web including those of humans and robots manipulating everyday objects. We use these 2D track predictions to infer a sequence of rigid transforms of the object to be manipulated, and obtain robot end-effector poses that can be executed in an open-loop manner. We then refine this open-loop plan by predicting residual actions through a closed loop policy trained with a few embodiment-specific demonstrations. We show that this approach of combining scalably learned track prediction with a residual policy requiring minimal in-domain robot-specific data enables diverse generalizable robot manipulation, and present a wide array of real-world robot manipulation results across unseen tasks, objects, and scenes. <https://homangab.github.io/track2act/>¹

Keywords: Diverse Generalizable Manipulation · Web Videos

1 Introduction

Robots that can be reliably deployed out-of-the-box in new scenarios have the potential for helping humans in everyday tasks. To realize this vision of generalizable robot manipulation, it is crucial to develop *direct* execution capabilities i.e. being able to execute a task out-of-the-box without requiring any test-time training through demonstrations or self-practice before solving a specified task. This is an important desiderata for the system to be repeatedly usable without any downtime, and safe to work alongside humans without performing any exploratory actions. We pursue the goal of developing such directly executable

¹ * equal contribution. Correspondence to Homanga B. hbharadh@cs.cmu.edu



Fig. 1: Glimpse of some of the diverse robot manipulation capabilities across physical office and kitchen scenes enabled by our framework. We learn to predict point tracks from web videos for learning interaction plans that can be used for inferring robot actions in unseen scenarios. This enables a *common* goal-conditioned policy to perform everyday tasks like closing microwaves, pulling out drawers, flipping open toasters, pouring from jars etc. Columns show first and last images of rollouts from our policy.

robot manipulation systems that can perform a broad set of everyday tasks. In addition to being deployable directly, to be widely accessible, we aim to make the robot manipulators generalizable to diverse offices, and kitchens in the real world.

Developing such directly executable manipulation capabilities has been attempted by prior works, through behavior cloning on robot interaction datasets [5, 8, 28, 73]. While this approach is in-principle scalable with data, collecting diverse real-world robot interaction data is challenging due to operational constraints. Indeed, recent works that have attempted to scale robot datasets, including cross-robot and cross-domain datasets [5, 15, 46] still suffer from task diversity issues and are mostly restricted to lab-like structured scenarios. Instead of learning a single-policy that can be directly deployed, some recent works aimed at in-the-wild deployment have adopted the method of test-time training [2, 40]. They require either a video of a human performing the task [2] followed by on-line exploration, or a demo through a robot end-effector held by a human [40]. These approaches are not very convenient for diverse deployments because they require a human to solve the task first, and several hours after that for the robot to learn how to solve that exact task in the exact scene. Thus, such approaches are not directly deployable for new tasks in new scenes.

Our insight to develop an in-the-wild manipulation strategy that is also directly deployable is to factorize a manipulation policy into an *interaction-plan* that can leverage diverse large-scale video sources on the web of humans and robots manipulating everyday objects and a residual policy that requires a small amount of embodiment-specific robot interaction data. Such a factorized structure is inspired by prior works (e.g. [4]), however, unlike hand-object masks in [4], we instantiate this *interaction-plan* in an embodiment agnostic-manner by predicting how points in an image of the initial scene move in future frames.

This choice of an *interaction-plan* is more expressive compared to hand-object masks adopted by Bharadhwaj et al. [4], as it directly captures point correspondences across time, while at the same time being easier to compute than full RGB frames [14]. Given an initial image of the scene, a goal image defining the task to be performed, and a random set of points in the initial image, we define the interaction plan to be a 2D trajectory of the locations of the points in future frames, such that the goal is achieved. Importantly, we can train this model purely from the abundantly available human and robot videos on the web without any data specific to the deployment robot embodiment, by using off-the-shelf point-tracking approaches [29] for generating the ground-truth point trajectories. For deployment in a robot’s environment, we can convert the 2D interaction-plan to a sequence of 3D end-effector poses, by having a depth image of the initial scene as an additional input and solving an optimization problem to obtain rigid transforms of the object being manipulated. Finally, with a small amount of embodiment-specific robot interaction data for different tasks (~ 400 trajectories overall), we can learn a goal-conditioned residual policy that corrects for errors in the predicted plan at each time-step and allows for closed-loop deployment.

In summary, we develop *Track2Act* with the following contributions:

- We develop a framework for predicting embodiment-agnostic *interaction-plans* in the form of point tracks from diverse web videos.
- We show how the interaction-plan prediction model can be used for obtaining 3D rigid transforms in a robot’s environment for direct manipulation without using any robot data or online exploration.
- Given a few (~ 400) embodiment-specific task demonstrations, we show how to learn a goal-conditioned residual policy that can correct for errors in the predicted plan at each time-step. The interaction-plan prediction model combined with the residual policy correction can then be used for closed-loop deployment for new tasks in new scenes.

Our real-world robot manipulation results with a Spot robot (highlighted in Fig. 1) show broad generalization across diverse tasks involving unseen objects in unseen scenes, and demonstrate the potential for leveraging easily available passive videos on the web for learning embodiment-agnostic interaction plans. This is significant as it enables robot manipulation with a *common* goal-conditioned policy, that generalizes to unseen tasks without requiring collection of large scale in-domain manipulation datasets.

2 Related Works

Understanding interactions from videos. Several computer vision methods have investigated deciphering interactions between hands and objects across various daily activities through curation of large-scale video datasets [10, 11, 21, 22, 35, 55, 59], hand pose estimation [1, 6, 18, 24, 27, 34, 37, 53, 58, 72], object pose estimation [25, 26, 31, 52, 64], interaction hotspot/grasp prediction [7, 20, 38, 43, 44].




Fig. 2: Illustration of the pipeline for learning track prediction from web video datasets, inferring rigid transforms of objects based on the predicted tracks in a robot’s environment, and fine-tuning with a residual policy learned with limited robot data. This approach allows us to learn a single goal-conditioned policy for diverse (unseen) tasks.

Some approaches have investigated *generating* videos given a description of the task, and often conditioned on a scene [14, 23, 68, 69]. Other works have attempted understanding generic videos by identifying visual correspondences between frames [19, 36]. Recent approaches have made significant advances on this problem by developing general-purpose video tracking systems that can track points specified in a frame, across other frames in the video [13, 29]. Our track prediction model is inspired by these developments, and is based on leveraging the video tracking approaches to generate ground-truth tracks from web videos, and training a model to *predict* future points tracks given an initial image and a goal.

Learning Visual Representations for Manipulation. Visual imitation is a promising technique for generalizable robot manipulation [16, 42, 70]. Recent works that have scaled this approach for learning large-scale models for manipulation require extremely high number of expert robot trajectories, often demanding years for collection [5, 8, 73], and still suffer from limited generalization to unseen scenarios for novel objects. Going beyond image observations, prior works have also investigated structured representations like point-clouds [9, 47, 54] and keypoints [51] for manipulation, but are restricted to tasks in structured table-top scenarios. Some of these that predict action in the form of flow-based representations [19, 54] require 3D datasets of robot interactions (often from simulation) which constrain them from generalizable real-world deployments. More recently, Vecerik et al. [60] use point tracking for visual servoing, and the setup requires structured multi-stage definitions of the task and is limited to only minor test-time variations compared to training data. Concurrent work [63] that improves upon [60] by predicting future tracks of points in the current image can learn a policy by combining in-domain human videos with in-domain robot videos. However, the framework is not directly amenable for leveraging web videos because the policy relies on per-step image observations for track prediction. Compared to this, and developed independently, we learn to predict trajectories of arbitrary points from web videos given just an initial image and a goal, we show how we can use these predicted tracks to infer rigid transforms of objects for open-loop execution, and further improve the open-loop plan by predicting residuals over the actions, for closed-loop deployment. This enables

much diverse robot manipulation behaviors with a single model, that generalizes to unseen novel objects and scenes in-the-wild.

Leveraging Non-Robot Datasets for Manipulation. One common way of using data beyond robot interactions for efficient learning is to pre-train the visual representations which serve as backbones for the policy models [39, 41, 45, 48, 65] with passive human videos [22, 30] and image data [12]. However, these methods still crucially rely on a lot of in-domain robot data or deployment-time training, and are restricted to learning task-specific policies. Some works that do not require deployment-time training, go beyond visual representations and use curated data of human videos to leverage human hand motion information [50, 56] for learning task-specific policies (instead of a single model across generic tasks). Others that train a single policy across tasks require large in-domain perfectly aligned human-robot data [57, 62, 66] and are not capable of leveraging web data. Towards learning structure more directly related to manipulation from web videos, some works try to predict visual affordances in the form of where to interact in an image, and local information of how to interact [3, 20, 38, 43]. While these could serve as good initializations for a robotic policy, they are not sufficient on their own for accomplishing tasks, and so are typically used in conjunction with online learning, requiring several hours of deployment-time training and robot data [2, 3]. Others learn to predict masks of hand and objects in the scene [4] for conditional behavior cloning and are unable to leverage information of accurate object state changes that is usually ambiguous with a mask. Our work differs from these in terms of predicting an approximate motion of how objects in the scene move in the future through point tracks for the entire trajectory and is directly executable in terms of not requiring any deployment-time training.

3 Approach

We aim to develop a generalizable robot manipulation system that can scalably leverage diverse video data for generalizable real-world manipulation. Our key insight (Fig. 2) is to have a factorized policy for 1) learning embodiment-agnostic *interactions plans* of how points in an image of a scene should move in subsequent time-steps to realize a specified goal, followed by 2) inferring robot actions based on the interaction plan through a residual policy. We show how this approach allows us to generalize to diverse scenarios involving unseen tasks and objects, since the prediction model by virtue of being trained on web data generalizes well to new scenes, and the residual policy has a much simpler task of correcting the robot actions derived from the interaction plan.

3.1 Overview and Setup

Given a scene specified by an RGB image \mathbf{I}_0 and a goal image \mathcal{G} denoting what task should be performed, we want to have a robot manipulator execute actions $\mathbf{a}_{1:H}$ in the scene to achieve the desired goal. To achieve this in unseen scenarios, we leverage web video data by learning a model $\tau = \mathcal{V}_\theta(\mathbf{I}_0, \mathcal{G}, P_0)$




Fig. 3: Architecture of the Diffusion Transformer \mathcal{V}_θ for denoising track predictions given initial image \mathbf{I}_0 , goal \mathcal{G} , and an initial set of p points P_0 .

Algorithm 1 Predicting Rigid Transforms from Point Tracks

- 1: **procedure** RIGID TRANSFORMS($\tau, \mathbf{I}_0, \mathcal{G}, P_0^{3D}, \mathbf{K}, H$)
 - 2: $\{\{(x_t^i, y_t^i)\}_{i=1}^p\}_{t=1}^H = \tau_{obj} = \text{filter}(\tau)$ ▷ Filter moving point tracks
 - 3: Unknown rigid transforms $[\mathbf{T}_t]_{t=1}^H$ ▷ \mathbf{T}_t has dimension 3×4
 - 4: Run RANSAC on τ_{obj} to filter outliers ▷ optional
 - 5: **for** $t \leftarrow 1$ to H **do**
 - 6: $\mathbf{T}_t = \arg \min_{\mathbf{T}_t} \sum_i^N (||x_t^i - u_t^i|| + ||y_t^i - v_t^i||)$
 - 7: where $(u_t^i, v_t^i, 1) \simeq \mathbf{K} \mathbf{T}_t P_t$ ▷ projections in homogeneous coordinates
 - 8: **return** $\{\mathbf{T}_t = (\mathbf{R}_t, \mathbf{t}_t)\}_{t=1}^T$
-

to predict future locations (tracks) of p random points P_0 in the initial image. Given a depth image for the initial frame, we leverage a subset of the predicted tracks τ_{obj} (corresponding to moving points) to infer rigid-transforms of the object being manipulated $[\mathbf{T}_t]_{t=1}^H$ and show that these allow obtaining an open-loop plan in the form of robot end-effector poses $[\bar{\mathbf{a}}_t]_{t=1}^H$. Finally, we consider training a closed-loop residual policy $\pi_{\text{res}}(\mathbf{I}_t, \mathcal{G}, \tau, [\bar{\mathbf{a}}_t]_{t=1}^H)$ that corrects the open-loop action sequence $[\bar{\mathbf{a}}_t]_{t=1}^H$ by predicting residual actions at each timestep $\Delta \mathbf{a}_t$, such that the executed action sequence is $[\mathbf{a}_t = \bar{\mathbf{a}}_t + \Delta \mathbf{a}_t]_{t=1}^H$. In the subsequent sections, we explain the architecture and algorithm design for each of the three stages in our approach.

3.2 Point Track Prediction from Web Videos

We instantiate track prediction as a denoising process through a DiT based diffusion model [49]. Let \mathbf{I}_0 denote the first frame of a video, and \mathcal{G} denote the goal, which we consider to be the last frame of the video. For longer videos, we obtain multiple video clips of 4-5 seconds each for training. Let there be p points in the initial frame to be tracked, such that P_0 denotes the set of those points and let H be the prediction horizon. $[P_t]_{t=1}^H$ denotes the future locations of those points in the subsequent time-steps that we want to predict. In the forward diffusion process, all the points P_t are corrupted by incrementally adding noise

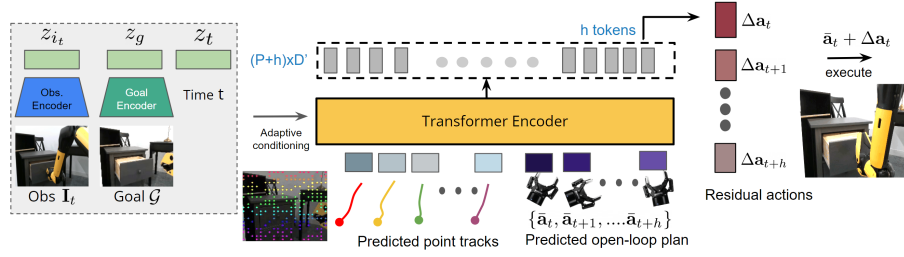


Fig. 4: Architecture of the residual policy that predicts corrections at each time-step over the predicted open-loop plan, and enables closed-loop deployment.

ϵ_k (k denotes the diffusion time-step), to obtain \tilde{P}_t , and converging to a unit Gaussian distribution $N(\mathbf{0}, \mathbf{I})$. New samples can be generated by reversing the forward diffusion process, by going from Gaussian noise back to the space of point locations. To solve the reverse diffusion process, we need to train a noise predictor $\mathcal{V}_\theta(\mathbf{I}_0, \mathcal{G}, P_0, k)$. We design a DiT Transformer based architecture [49] for \mathcal{V}_θ illustrated visually in Fig. 3. Different from the original DiT model, we condition on embeddings of initial (z_0) and goal (z_g) images in addition to that of the diffusion step (z_k). The input to the Transformer in each batch is a sequence of p tokens corresponding the number of points specified for tracking. The initial P_0 points are not noisy, as is the convention in training conditional diffusion models on time-series data.

We train the prediction model with web videos by considering variable number of initial points p that need to be tracked. For flexible modeling, the locations of the p points are also randomized, such that at test-time any set of points in the initial image can be specified. We do not make any assumptions on objects to be tracked or camera motions in the videos, and do not curate the training videos in any way apart from ensuring they are of 4-5 second duration. If the goal image is such that multiple objects have moved from the initial scene, or the camera has moved, the track prediction model will predict different groups of motions for different objects and also predict motions of background points to account for camera motion. However, for robot experiments, we consider only a single object to be manipulated at a time, which is indeed the case with several diverse real-world tasks.

3.3 Inferring Coarse Manipulator Trajectory from Interaction Plan

Given an image of a scene in a robot’s environment \mathbf{I}_0 , a goal \mathcal{G} , and a random set of points P_0 in the initial image, we can use the trained track prediction model to obtain future 2D locations of these points \hat{P}_t . As we consider scenarios with only a single object being manipulated under a fixed camera, only a subset of the points have a large predicted motion. We identify these p points and denote their predicted trajectories as $\tau_{\text{obj}} = [\{(x_t^i, y_t^i)\}_{i=1}^p]_{t=1}^H$. We consider the robot to be equipped with an RGBD camera, so we also have depth for the points P_0 in the first frame. Let us denote these 3D points as $P_0^{3D} = \{(x_0^i, y_0^i, z_0^i)\}_{i=1}^p$.

We seek to infer a (smooth trajectory of) per-time rigid transforms \mathbf{T}_t of the object to be manipulated at time t relative to first frame, given 3D points in the first frame P_0^{3D} , predicted 2D trajectory of points on the object τ_{obj} , and the camera intrinsic matrix \mathbf{K} . As described in Algorithm 1, these can be obtained by ensuring that the projection of the transformed 3D points i.e. $\mathbf{K}\mathbf{T}_tP_0$ matches the predicted 2D tracks $\{(x_t^i, y_t^i)\}_{i=1}^p$ as closely as possible at each time-step t . Let $\mathbf{K}\mathbf{T}_tP_0 \simeq \{(u_t^i, v_t^i, 1)\}_{i=1}^p$. So the 2D projection of the i^{th} point at time t is (u_t^i, v_t^i) . Alternatively, we have the same coordinate for the point from the predicted track i.e. (x_t^i, y_t^i) . In order to determine the rigid transforms \mathbf{T}_t , we can solve the optimization problem in line 6 of Algorithm 1 (e.g. with PnP solvers). This optimization is not under-constrained because the same 3D rigid transform \mathbf{T}_i must explain the 2D motions of several points P_0 in the initial scene. Note that the obtained rigid transforms are *embodiment-agnostic* and describe how the object should move in the scene.

Now, to actually manipulate the object in the scene, we need to bring the robot end-effector² near the object, and optionally execute a grasp to hold on to the object, followed by transforming the end-effector based on the predicted rigid transforms $[\mathbf{T}_t]_{t=1}^H$. For the first step, we use a heuristic such that given initial end-effector pose \mathbf{e}_0 we define the first transform \mathbf{T}_0 to be such that the end-effector moves to the center of the 3D points $\{(x_t^i, y_t^i)\}_{i=0}^p$ with the same orientation as \mathbf{e}_0 . After moving the end-effector to this pose \mathbf{e}_1 we execute a grasp to hold the object. We obtain subsequent end-effector poses (open-loop action trajectory) by applying the rigid transforms $\bar{\mathbf{a}}_t = \mathbf{T}_t\mathbf{e}_1$.

3.4 Closed-loop Manipulation with Residual Policy Correction

The open-loop execution of the predicted 3D end-effector transforms described in the last section $[\bar{\mathbf{a}}_t]_{t=1}^H$ might fail due to small errors in the prediction. In addition, since the approach does not use any embodiment-specific data, it does not have accurate information for reasoning about contact with objects and might suffer from failures like being unable to grasp the object, in spite of executing the rest of the predicted trajectory correctly. To remedy this, we propose learning a residual policy $\pi_{\text{res}}(\mathbf{I}_t, \mathcal{G}, \tau, [\bar{\mathbf{a}}_t]_{t=1}^H)$ shown in Fig. 4 to correct the predicted end-effector poses in each time-step. So the end-effector pose at time t is

$$\hat{\mathbf{a}}_t = \bar{\mathbf{a}}_t + \Delta\mathbf{a}_t \quad ; \quad \text{where } \Delta\mathbf{a}_t = \pi_{\text{res}}(\mathbf{I}_t, \mathcal{G}, \tau, [\bar{\mathbf{a}}_t]_{t=1}^H) \quad (1)$$

Instead of predicting just a single residual action $\Delta\mathbf{a}_t$ we predict residuals h steps in the future $\Delta\mathbf{a}_{t:t+h}$ and during deployment execute just the first action. This multi-step prediction has been shown to mitigate compounding errors in behavior-cloning based training [5, 71]. We can learn the residual policy with a small amount of robot demonstrations (~ 400 trajectories overall) of representative tasks through behavior cloning. The data for each trajectory consists

² by end-effector we mean the part of the robot that interacts with an object

of observation-action pairs of the form $[(\mathbf{I}_t, \mathbf{a}_t)]_{t=1}^H$. Here, \mathbf{I}_t denotes images observed from the robot’s camera and \mathbf{a}_t denotes actions in the form of end-effector poses.

Crucially, since the aim of this policy is to learn only small corrections to the predicted waypoints $[\bar{\mathbf{a}}_t]_{t=1}^H$, we do not need to learn this policy with data from the exact scenarios that the system will be deployed in and the prediction model is expected to generalize to unseen scenarios by virtue of diverse training. The rationale is that having some embodiment-specific demonstration data in a few scenarios will help correct for the open-loop predictions from web-only data. For evaluation, we consider different levels of generalization with unseen object instances and completely unseen objects in unseen scenes.

4 Experiment Setup

We focus our experiments on in-the-wild manipulation scenarios where a mobile manipulator needs to manipulate objects in different living rooms, offices, and kitchens based on specified goals. For all the robot experiments, we use a Boston Dynamics Spot robot equipped with a manipulator (hand) and a front facing Intel RealSense camera [32]. We manipulate the arm through end-effector control based on the outputs of our policy.

4.1 Evaluation Details

Track Prediction. For quantitative evaluation of the track prediction model, we adopt a modification of the metric developed by prior works [13, 29], δ_t^x . For evaluation videos, we consider the output of Co-Tracker [29] to be the ground-truth and compare the difference with respect to the predictions, based on the δ_t^x metric. We define δ_t^x to be the fraction of points that are within a threshold pixel distance of x of their ground truth in a time-step t . We report the *area under the curve* Δ with δ_t^x by varying x from 1 to $N = 10$ and taking the average across the prediction horizon H i.e. $\Delta = (\sum_{t=1}^H \sum_{x=1}^N \delta_t^x) / H$. Hence, Δ can vary from 0 to 1 with higher being better.

Track Prediction. As is the convention in goal-conditioned robot learning, we perform evaluations by quantifying success rate, where a successful trajectory is defined to be one where the final pose of the object in the scene to be manipulated is identical to the pose of the object in the goal image. We categorize results based on different levels of generalization, the definitions of which are inspired by prior works [4, 5, 8, 73]:

- Mild Generalization (**MG**): unseen configurations of seen object instances in seen scenes; organic scene variations like lighting and background changes
- Standard Generalization (**G**): unseen object instances in seen/unseen scenes
- Combinatorial Generalization (**CG**): unseen activity-object type combinations in seen/unseen scenes
- Type Generalization (**TG**): completely unseen object types, or completely unseen activities, in unseen scenes

4.2 Baselines and Comparisons

For quantitative evaluations, we first compare our track prediction approach with other related baselines and then perform comparisons with baselines for robot manipulation experiments.

Track Prediction. We perform comparisons with two baselines that have the same inputs as our track prediction model, i.e. an initial image, a goal image and points specified on the initial image, and the same output type i.e. point tracks in between the initial and goal images. We compare with a flow-based baseline that directly predicts flow between the initial and goal images, and then performs a per-timestep interpolation of the flow vectors [67]. The second baseline performs video-infilling given initial and goal images [17], and then uses Co-Tracker [29] to obtain tracks on the generated video.

Robot Experiments. We perform several comparisons with baselines and ablation studies for goal-conditioned robot manipulation. For baselines, we use the same embodiment-specific demonstrations as *Ours*, the goal-conditioned policy that predicts residuals over open-loop actions at each time-step.

- *Goal-Conditioned BC* is a baseline for multi-task policy learning, similar to prior works [5, 8, 73].
- *Affordance-Conditioned BC* is the approach from [3] that conditions the policy on predicted affordances in the initial image.
- *Video-Conditioned BC* based on [14, 17, 33] first predicts RGB video and then does tracking on top of it.
- *Hand-Object Mask Conditioned BC* from [4] conditions the policy on a predicted interaction plan consisting of masks of hands and objects.

Ours (Open Loop) is the approach for track prediction followed by open-loop execution as described in Algorithm 1. This does not use any embodiment-specific data for training. To understand the benefit of predicting residuals over actions as opposed to predicting complete actions, we compare with an ablated variant *Ours (actions; not residuals)* that predicts actions $\hat{\mathbf{a}}_t$ directly without predicting residuals $\Delta\mathbf{a}_t$ and not relying on an open-loop plan as input.

4.3 Training Data

For training the track prediction model, we leverage diverse passive videos available on the web that are not collected by us. Specifically, we use human video clips from EpicKitchens [10] (clipping videos to ensure they are of 4-5 seconds duration), and large-scale robot videos released in RT1 data [8] and Bridge-Data [61]. To obtain ground-truth tracks for training the prediction model, we run Co-Tracker [29] on the resulting 400,000 video clips. Note that the robot datasets (RT1 and Bridge) are on completely different robots and scenarios than the robot we use for experiments (Spot). For training the residual policy, the embodiment-specific data we collect consists of ~ 400 trajectories obtained by tele-operating the Spot, for solving 10 tasks of manipulating everyday objects like doors, drawers, bottles, jugs. Note that this embodiment-specific data we collect is 3-4 orders of magnitude less than that what related works [5, 8, 73] require for policy learning.

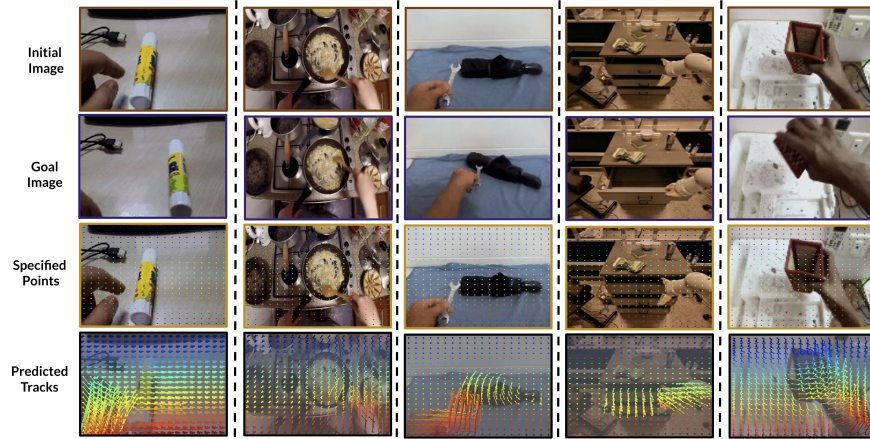


Fig. 5: We show qualitative results of the track predictions on unseen initial and goal images across diverse datasets. Given specified points on the initial image we predict future tracks of these points, corresponding to the goal image. We can see that the predictions are plausible and correspond to manipulating the object(s) in the scene.

5 Results

We present qualitative results of the predicted tracks, and robot evaluations, followed by quantitative comparisons with the metrics defined in section 4.1. Please refer to the supplementary zip for detailed qualitative results and robot evaluation videos.

5.1 Point Track Prediction Results

Table 1: Evaluation of track prediction performance on held-out videos from different datasets on the web. EpicKitchens [10] and SmthSmthv2 [21] are datasets of human videos, and BridgeData [61] and RT1 data [8] are datasets of robot videos. Note that we train a *single* model that we evaluate on these different datasets. The metric Δ is defined in section 4.1. Higher is better and the range is from 0 to 1.

	EpicKitchens [10]	SmthSmthv2 [21]	BridgeData [61]	RT1 Data [8]
Flow [67]	0.21	0.27	0.42	0.38
Video [17]	0.30	0.17	-	-
Ours	0.67	0.70	0.77	0.75

We first look at some qualitative results of the track prediction model in different unseen scenes. In Fig. 5 we show visualization of track predictions on unseen initial and goal images across diverse datasets. We choose points on a grid in the initial frame, as shown in the third row. The prediction model is conditioned on the initial image, the goal image, and the set of points in the initial image whose future tracks are to be predicted. We can see that the predictions




Fig. 6: We show visualizations of point track predictions for different tasks, followed by closed-loop execution with the residual policy. We can see that the predictions are plausible and the robot execution successfully realizes the predictions to complete the respective tasks specified by the goal images. The bottom row shows the generalization level for each task, defined in section 4.1.

(shown in the fourth row) are plausible and correspond to manipulating the objects in the scene as described by the respective goal images. We can also see that when multiple entities (e.g. human and object or robot and object) or the camera moves between the initial and goal images, there are different sets of point tracks predicting the respective motions.

In Table 1 we perform evaluations for track prediction by comparing with the flow-based [67] and video-based [17] baselines. We can see that both the baselines have much lower accuracy compared to our approach of predicting point tracks. This is because flow is too coarse to capture large non-linear state changes in between the initial and goal images. Whereas, predicting an RGB video followed by tracking suffers due to issues of implausible generation because video generation is a much more complex task than predicting the tracks of a set of points where the details about appearance, texture etc. are abstracted out. For reference, not predicting any movement for any point at all time-steps scores 0.03, 0.05, 0.36, 0.28. This suggests the benefit of directly predicting future point tracks as done by our approach if the aim is to capture motion of objects in the scene between an initial and a goal image.

5.2 Robot Manipulation Results

We visualize results of point track predictions using the trained prediction model in a robot’s environment, followed by residual policy executions based on the predictions. In Fig. 6 we show the track predictions overlaid on the initial image,

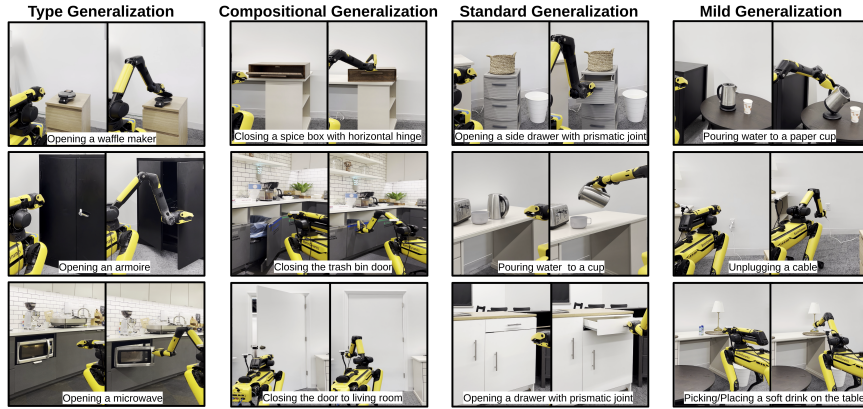


Fig. 7: Qualitative results showing robot executions (from a third person view) with the residual policy for different tasks with respect to the generalization levels defined in section 4.1. We show the first and last images of a rollout. The robot executions are best viewed as videos in the supplementary zip.

corresponding to the goal image shown in the top row. The bottom row shows robot execution with the first frame (1), two intermediate frames (2,3) and the last frame (4) of a rollout. We can see that the predicted tracks correspond to points on the object moving in a way that satisfies the goal, and the policy is able to manipulate the object to the desired goal configuration. Since the camera doesn’t move between the initial and goal images, we can see that background (non object) pixels remain stationary in the predictions, which is useful for accurate prediction of rigid transforms of the object.

In Table 2 we show comparisons for robot manipulation experiments, respectively for each level of generalization. We evaluate each approach for 20 rollouts in each level, across a total of 25 tasks in 5 different physical kitchen, office, and living room locations. We first note that our residual policy outperforms our approach for directly executing an open-loop plan based on predicted rigid transforms. This suggests that the residual policy is able to correct for inaccuracies in the open-loop plan by virtue of leveraging some embodiment-specific data that helps in performing accurate grasps on objects and recovering from potential failures during a trajectory.

We observe that for mild generalization (MG), the goal-conditioned BC baseline has slightly lower success rate compared to our residual policy, and significantly lower (or zero) success rates for standard (G), compositional (CG), and type (TG) generalization. This suggests the benefit of leveraging web video data for learning interaction plans that helps our approach generalize effectively. Finally, compared to baselines that also leverage web data like affordance-conditioned BC, video-conditioned BC, and hand-object mask-conditioned BC, we observe significant gains from our approach in the higher levels of generalization (CG and TG). This suggests that predicting static affordances without reasoning about motion trajectories, hallucinating RGB videos that suffer from incorrect generations and produce implausible artifacts in the scene, or predict-

Table 2: Evaluation of goal-conditioned robot manipulation experiments, per the protocol described in section 4.1. The numbers denote success rate averaged over 20 rollouts for different tasks within each generalization axis (Higher is better). Detailed list of tasks are in the Supplementary pdf. Refer to Fig. 7 for visualizations of some task rollouts corresponding to each of the four generalization axes.

	MG	G	CG	TG
Behavior Cloning (BC)	60%	20%	0%	0%
Affordance-Conditioned	65%	30%	10%	5%
Video-Conditioned	60%	25%	0%	0%
Hand-Object Mask-Conditioned	70%	40%	25%	20%
Ours (Open-Loop)	35%	25%	30%	25%
Ours (Ablation; actions not residuals)	70%	45%	30%	30%
Ours	70%	60%	55%	40%

ing 2D masks of hands and objects without reasoning about correspondences are insufficient cues for effectively leveraging web videos. Compared to these, our interaction plan learned through track prediction provides sufficient cues for solving unseen manipulation tasks by virtue of allowing inference of 3D rigid transforms, and the residual policy helps correct for errors in the predictions.

5.3 Analysis of Failures

Here we discuss the failures displayed by our framework. For the open-loop plan based on residual transforms, the main failure modes we observe are inability to grasp the object at the right location, and inability to recover from intermediate failures. The residual policy corrects for these behaviors by virtue of leveraging some embodiment-specific data, and thus has higher success rates. We note that the success rate for higher levels of generalization (CG and TG) is still not very high since these are very challenging settings and the residual policy sometimes fails by incorrectly grasping the object, getting stuck during the execution by trying to execute a non-feasible motion, or by executing a trajectory that does not conform with the goal image specified.

6 Discussion and Conclusion

In this paper, we developed a framework for generalizable robot manipulation by leveraging large-scale web video data to learn embodiment agnostic plans of how objects should be manipulated in a scene to satisfy a goal. We combined this with a small amount of embodiment-specific data to learn residual corrections over the predicted plans through a closed-loop policy. Our real world manipulation results across a range of diverse tasks with varying levels of generalization demonstrate the potential of scalably leveraging web data to predict plans for object manipulation. While our framework allows for strong generalization to unseen tasks in-the-wild, the tasks are still of short-horizon and involve manipulating a single object in the scene. It would be an interesting direction of future work to extend our framework for tackling long-horizon tasks that involve successive manipulations of multiple objects in the scene.

Acknowledgement

We thank Yufei Ye, Himangi Mittal, Devendra Chaplot, Abitha Thankaraj, Tarasha Khurana, Akash Sharma, Sally Chen, Jay Vakil, Chen Bao, Unnat Jain, Swaminathan Gurumurthy for helpful discussions and feedback. We thank Carl Doersch and Nikita Karaev for insightful discussions about point tracking. This research was partially supported by a Google gift award.

References

1. Baek, S., Kim, K.I., Kim, T.K.: Pushing the envelope for rgb-based dense 3d hand pose estimation via neural rendering. In: CVPR (2019)
2. Bahl, S., Gupta, A., Pathak, D.: Human-to-robot imitation in the wild. RSS (2022)
3. Bahl, S., Mendonca, R., Chen, L., Jain, U., Pathak, D.: Affordances from human videos as a versatile representation for robotics. In: CVPR (2023)
4. Bharadhwaj, H., Gupta, A., Kumar, V., Tulsiani, S.: Towards generalizable zero-shot manipulation via translating human interaction plans. In: 2024 IEEE International Conference on Robotics and Automation (ICRA) (2024)
5. Bharadhwaj, H., Vakil, J., Sharma, M., Gupta, A., Tulsiani, S., Kumar, V.: Roboagent: Generalization and efficiency in robot manipulation via semantic augmentations and action chunking. In: 2024 IEEE International Conference on Robotics and Automation (ICRA) (2024)
6. Boukhayma, A., Bem, R.d., Torr, P.H.: 3d hand shape and pose from images in the wild. In: CVPR (2019)
7. Brahmabhatt, S., Handa, A., Hays, J., Fox, D.: Contactgrasp: Functional multi-finger grasp synthesis from contact. arXiv (2019)
8. Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Dabis, J., Finn, C., Gopalakrishnan, K., Hausman, K., Herzog, A., Hsu, J., et al.: Rt-1: Robotics transformer for real-world control at scale. arXiv preprint arXiv:2212.06817 (2022)
9. Byravan, A., Fox, D.: Se3-nets: Learning rigid body motion using deep neural networks. In: 2017 IEEE International Conference on Robotics and Automation (ICRA). pp. 173–180. IEEE (2017)
10. Damen, D., Doughty, H., Farinella, G.M., Fidler, S., Furnari, A., Kazakos, E., Moltisanti, D., Munro, J., Perrett, T., Price, W., et al.: Scaling egocentric vision: The epic-kitchens dataset. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 720–736 (2018)
11. Das, P., Xu, C., Doell, R.F., Corso, J.J.: A thousand frames in just a few words: Lingual description of videos through latent topics and sparse object stitching. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2634–2641 (2013)
12. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
13. Doersch, C., Gupta, A., Markeeva, L., Recasens, A., Smaira, L., Aytar, Y., Carreira, J., Zisserman, A., Yang, Y.: Tap-vid: A benchmark for tracking any point in a video. Advances in Neural Information Processing Systems **35**, 13610–13626 (2022)
14. Du, Y., Yang, S., Dai, B., Dai, H., Nachum, O., Tenenbaum, J., Schuurmans, D., Abbeel, P.: Learning universal policies via text-guided video generation. Advances in Neural Information Processing Systems **36** (2024)

15. Fang, H.S., Fang, H., Tang, Z., Liu, J., Wang, J., Zhu, H., Lu, C.: Rh20t: A robotic dataset for learning diverse skills in one-shot. *arXiv preprint arXiv:2307.00595* (2023)
16. Finn, C., Yu, T., Zhang, T., Abbeel, P., Levine, S.: One-shot visual imitation learning via meta-learning. In: *Conference on robot learning*. pp. 357–368. PMLR (2017)
17. Fu, T.J., Yu, L., Zhang, N., Fu, C.Y., Su, J.C., Wang, W.Y., Bell, S.: Tell me what happened: Unifying text-guided video completion via multimodal masked video generation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10681–10692 (2023)
18. Ge, L., Ren, Z., Li, Y., Xue, Z., Wang, Y., Cai, J., Yuan, J.: 3d hand shape and pose estimation from a single rgb image. In: *CVPR* (2019)
19. Goyal, A., Mousavian, A., Paxton, C., Chao, Y.W., Okorn, B., Deng, J., Fox, D.: Ifor: Iterative flow minimization for robotic object rearrangement. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14787–14797 (2022)
20. Goyal, M., Modi, S., Goyal, R., Gupta, S.: Human hands as probes for interactive object understanding. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3293–3303 (2022)
21. Goyal, R., Ebrahimi Kahou, S., Michalski, V., Materzynska, J., Westphal, S., Kim, H., Haenel, V., Fruend, I., Yianilos, P., Mueller-Freitag, M., et al.: The "something something" video database for learning and evaluating visual common sense. In: *Proceedings of the IEEE international conference on computer vision*. pp. 5842–5850 (2017)
22. Grauman, K., Westbury, A., Byrne, E., Chavis, Z., Furnari, A., Girdhar, R., Hamburger, J., Jiang, H., Liu, M., Liu, X., et al.: Ego4d: Around the world in 3,000 hours of egocentric video. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 18995–19012 (2022)
23. Gupta, A., Yu, L., Sohn, K., Gu, X., Hahn, M., Fei-Fei, L., Essa, I., Jiang, L., Lezama, J.: Photorealistic video generation with diffusion models. *arXiv preprint arXiv:2312.06662* (2023)
24. Hasson, Y., Varol, G., Tzionas, D., Kalevatykh, I., Black, M.J., Laptev, I., Schmid, C.: Learning joint reconstruction of hands and manipulated objects. In: *CVPR* (2019)
25. He, Y., Sun, W., Huang, H., Liu, J., Fan, H., Sun, J.: Pvn3d: A deep point-wise 3d keypoints voting network for 6dof pose estimation. *CVPR* (2020)
26. Hu, Y., Hugonot, J., Fua, P., Salzmann, M.: Segmentation-driven 6d object pose estimation. *CVPR* (2019)
27. Iqbal, U., Molchanov, P., Breuel Juergen Gall, T., Kautz, J.: Hand pose estimation via latent 2.5 d heatmap regression. In: *ECCV* (2018)
28. Jang, E., Irpan, A., Khansari, M., Kappler, D., Ebert, F., Lynch, C., Levine, S., Finn, C.: Bc-z: Zero-shot task generalization with robotic imitation learning. In: *Conference on Robot Learning*. pp. 991–1002. PMLR (2022)
29. Karaev, N., Rocco, I., Graham, B., Neverova, N., Vedaldi, A., Ruppel, C.: Cotracker: It is better to track together. *arXiv preprint arXiv:2307.07635* (2023)
30. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950* (2017)
31. Kehl, W., Manhardt, F., Tombari, F., Ilic, S., Navab, N.: Ssd-6d: Making rgb-based 3d detection and 6d pose estimation great again. *ICCV* (2017)

32. Keselman, L., Iselin Woodfill, J., Grunnet-Jepsen, A., Bhowmik, A.: Intel realsense stereoscopic depth cameras. In: *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. pp. 1–10 (2017)
33. Ko, P.C., Mao, J., Du, Y., Sun, S.H., Tenenbaum, J.B.: Learning to act from actionless videos through dense correspondences. *arXiv preprint arXiv:2310.08576* (2023)
34. Kulon, D., Guler, R.A., Kokkinos, I., Bronstein, M.M., Zafeiriou, S.: Weakly-supervised mesh-convolutional hand reconstruction in the wild. In: *CVPR* (2020)
35. Li, Y., Liu, M., Rehg, J.M.: In the eye of beholder: Joint learning of gaze and actions in first person video. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 619–635 (2018)
36. Liu, C., Yuen, J., Torralba, A., Sivic, J., Freeman, W.T.: Sift flow: Dense correspondence across different scenes. In: *Computer Vision–ECCV 2008: 10th European Conference on Computer Vision, Marseille, France, October 12–18, 2008, Proceedings, Part III 10*. pp. 28–42. Springer (2008)
37. Liu, S., Jiang, H., Xu, J., Liu, S., Wang, X.: Semi-supervised 3d hand-object poses estimation with interactions in time. In: *CVPR* (2021)
38. Liu, S., Tripathi, S., Majumdar, S., Wang, X.: Joint hand motion and interaction hotspots prediction from egocentric videos. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3282–3292 (2022)
39. Ma, Y.J., Sodhani, S., Jayaraman, D., Bastani, O., Kumar, V., Zhang, A.: Vip: Towards universal visual reward and representation via value-implicit pre-training. *arXiv preprint arXiv:2210.00030* (2022)
40. Mahi Shafiullah, N.M., Rai, A., Etukuru, H., Liu, Y., Misra, I., Chintala, S., Pinto, L.: On bringing robots home. *arXiv e-prints* pp. arXiv–2311 (2023)
41. Majumdar, A., Yadav, K., Arnaud, S., Ma, Y.J., Chen, C., Silwal, S., Jain, A., Berges, V.P., Abbeel, P., Malik, J., et al.: Where are we in the search for an artificial visual cortex for embodied intelligence? *arXiv preprint arXiv:2303.18240* (2023)
42. Mandlekar, A., Zhu, Y., Garg, A., Booher, J., Spero, M., Tung, A., Gao, J., Emmons, J., Gupta, A., Orbay, E., et al.: Roboturk: A crowdsourcing platform for robotic skill learning through imitation. In: *Conference on Robot Learning*. pp. 879–893. PMLR (2018)
43. Mo, K., Guibas, L.J., Mukadam, M., Gupta, A., Tulsiani, S.: Where2act: From pixels to actions for articulated 3d objects. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 6813–6823 (2021)
44. Nagaraajan, T., Feichtenhofer, C., Grauman, K.: Grounded human-object interaction hotspots from video. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 8688–8697 (2019)
45. Nair, S., Rajeswaran, A., Kumar, V., Finn, C., Gupta, A.: R3m: A universal visual representation for robot manipulation. *arXiv preprint arXiv:2203.12601* (2022)
46. Padalkar, A., Pooley, A., Jain, A., Bewley, A., Herzog, A., Irpan, A., Khazatsky, A., Rai, A., Singh, A., Brohan, A., et al.: Open x-embodiment: Robotic learning datasets and rt-x models. *arXiv preprint arXiv:2310.08864* (2023)
47. Pan, C., Okorn, B., Zhang, H., Eisner, B., Held, D.: Tax-pose: Task-specific cross-pose estimation for robot manipulation. In: *Conference on Robot Learning*. pp. 1783–1792. PMLR (2023)
48. Parisi, S., Rajeswaran, A., Purushwalkam, S., Gupta, A.: The unsurprising effectiveness of pre-trained vision models for control. *arXiv preprint arXiv:2203.03580* (2022)

49. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4195–4205 (2023)
50. Qin, Y., Wu, Y.H., Liu, S., Jiang, H., Yang, R., Fu, Y., Wang, X.: Dexmv: Imitation learning for dexterous manipulation from human videos. arXiv preprint arXiv:2108.05877 (2021)
51. Qin, Z., Fang, K., Zhu, Y., Fei-Fei, L., Savarese, S.: Keto: Learning keypoint representations for tool manipulation. In: 2020 IEEE International Conference on Robotics and Automation (ICRA). pp. 7278–7285. IEEE (2020)
52. Rad, M., Lepetit, V.: Bb8: A scalable, accurate, robust to partial occlusion method for predicting the 3d poses of challenging objects without using depth. ICCV (2017)
53. Rong, Y., Shiratori, T., Joo, H.: Frankmocap: Fast monocular 3d hand and body motion capture by regression and integration. arXiv preprint arXiv:2008.08324 (2020)
54. Seita, D., Wang, Y., Shetty, S.J., Li, E.Y., Erickson, Z., Held, D.: Toolflownet: Robotic manipulation with tools via predicting tool flow from point clouds. In: Conference on Robot Learning. pp. 1038–1049. PMLR (2023)
55. Shan, D., Geng, J., Shu, M., Fouhey, D.: Understanding human hands in contact at internet scale. In: CVPR (2020)
56. Shaw, K., Bahl, S., Pathak, D.: Videodex: Learning dexterity from internet videos. In: 6th Annual Conference on Robot Learning
57. Smith, L., Dhawan, N., Zhang, M., Abbeel, P., Levine, S.: Avid: Learning multi-stage tasks via pixel-level translation of human videos. arXiv (2019)
58. Spurr, A., Song, J., Park, S., Hilliges, O.: Cross-modal deep variational hand pose estimation. In: CVPR (2018)
59. De la Torre, F., Hodgins, J., Bargteil, A., Martin, X., Macey, J., Collado, A., Beltran, P.: Guide to the carnegie mellon university multimodal activity (cmu-mmact) database (2009)
60. Vecerik, M., Doersch, C., Yang, Y., Davchev, T., Aytar, Y., Zhou, G., Hadsell, R., Agapito, L., Scholz, J.: Robotap: Tracking arbitrary points for few-shot visual imitation. arXiv preprint arXiv:2308.15975 (2023)
61. Walke, H.R., Black, K., Zhao, T.Z., Vuong, Q., Zheng, C., Hansen-Estruch, P., He, A.W., Myers, V., Kim, M.J., Du, M., et al.: Bridgedata v2: A dataset for robot learning at scale. In: Conference on Robot Learning. pp. 1723–1736. PMLR (2023)
62. Wang, C., Fan, L., Sun, J., Zhang, R., Fei-Fei, L., Xu, D., Zhu, Y., Anandkumar, A.: Mimicplay: Long-horizon imitation learning by watching human play. arXiv preprint arXiv:2302.12422 (2023)
63. Wen, C., Lin, X., So, J., Chen, K., Dou, Q., Gao, Y., Abbeel, P.: Any-point trajectory modeling for policy learning. arXiv preprint arXiv:2401.00025 (2023)
64. Xiang, Y., Schmidt, T., Narayanan, V., Fox, D.: Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. arXiv (2018)
65. Xiao, T., Radosavovic, I., Darrell, T., Malik, J.: Masked visual pre-training for motor control. arXiv preprint arXiv:2203.06173 (2022)
66. Xiong, H., Li, Q., Chen, Y.C., Bharadhwaj, H., Sinha, S., Garg, A.: Learning by watching: Physical imitation of manipulation skills from human videos. arXiv (2021)
67. Xu, H., Zhang, J., Cai, J., Rezatofghi, H., Yu, F., Tao, D., Geiger, A.: Unifying flow, stereo and depth estimation. IEEE Transactions on Pattern Analysis and Machine Intelligence (2023)
68. Yan, W., Hafner, D., James, S., Abbeel, P.: Temporally consistent transformers for video generation. arXiv preprint arXiv:2210.02396 (2022)

- 69. Yang, M., Du, Y., Ghasemipour, K., Tompson, J., Schuurmans, D., Abbeel, P.: Learning interactive real-world simulators. arXiv preprint arXiv:2310.06114 (2023)
- 70. Young, S., Gandhi, D., Tulsiani, S., Gupta, A., Abbeel, P., Pinto, L.: Visual imitation made easy. In: Conference on Robot Learning (CoRL) (2020)
- 71. Zhao, T.Z., Kumar, V., Levine, S., Finn, C.: Learning fine-grained bimanual manipulation with low-cost hardware. arXiv preprint arXiv:2304.13705 (2023)
- 72. Zimmermann, C., Brox, T.: Learning to estimate 3d hand pose from single rgb images. In: CVPR (2017)
- 73. Zitkovich, B., Yu, T., Xu, S., Xu, P., Xiao, T., Xia, F., Wu, J., Wohlhart, P., Welker, S., Wahid, A., et al.: Rt-2: Vision-language-action models transfer web knowledge to robotic control. In: Conference on Robot Learning. pp. 2165–2183. PMLR (2023)