

DemoDiffusion: One-Shot Human Imitation Using Pre-Trained Diffusion Policy

Sungjae Park, Homanga Bharadhwaj, Shubham Tulsiani
 Carnegie Mellon University

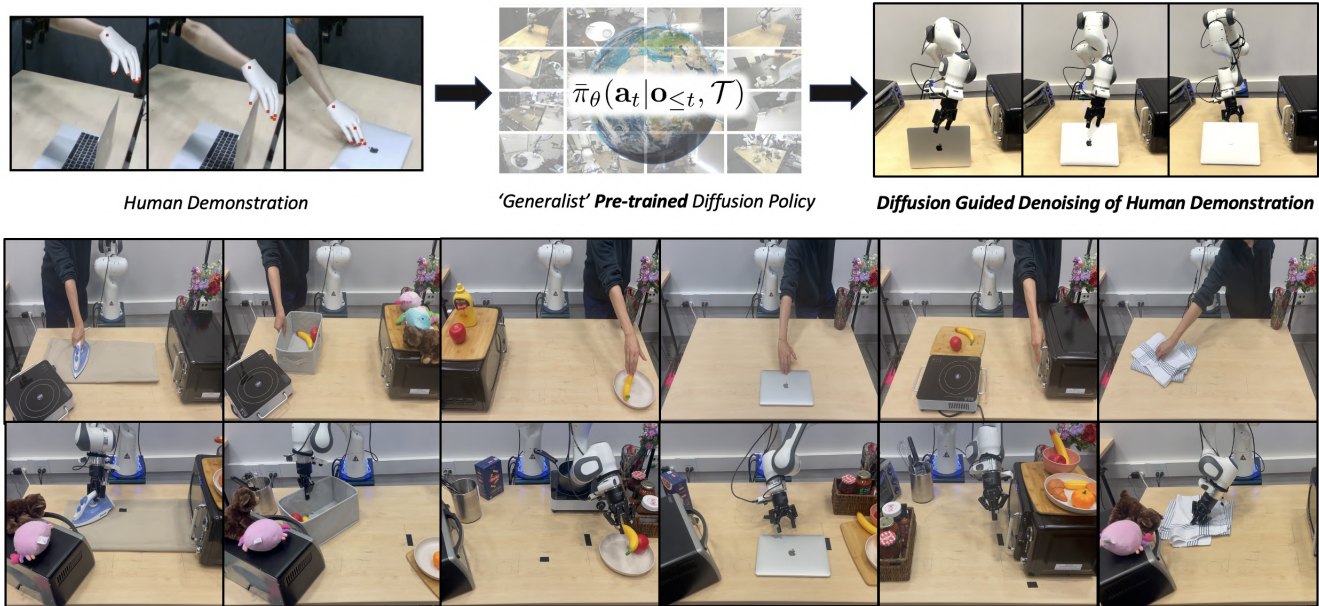


Fig. 1: **Overview of DemoDiffusion.** We leverage a *generalist* pre-trained diffusion policy to refine a single human demonstration into executable robot actions, enabling one-shot deployment without task-specific fine-tuning. Our real-world results span diverse manipulation tasks involving everyday objects.

Abstract—We propose *DemoDiffusion*, a simple method for enabling robots to perform manipulation tasks by imitating a single human demonstration, without requiring task-specific training or paired human-robot data. Our approach is based on two insights. First, the hand motion in a human demonstration provides a useful prior for the robot’s end-effector trajectory, which we can convert into a rough open-loop robot motion trajectory via kinematic retargeting. Second, while this retargeted motion captures the overall structure of the task, it may not align well with plausible robot actions in-context. To address this, we leverage a pre-trained *generalist* diffusion policy to modify the trajectory, ensuring it both follows the human motion and remains within the distribution of plausible robot actions. Unlike approaches based on online reinforcement learning or paired human-robot data, our method enables robust adaptation to new tasks and scenes with minimal effort. In real-world experiments across 8 diverse manipulation tasks, *DemoDiffusion* achieves 83.8% average success rate, compared to 13.8% for the pre-trained policy and 52.5% for kinematic retargeting, succeeding even on tasks where the pre-trained *generalist* policy fails entirely. Project page: <https://demodiffusion.github.io/>

I. INTRODUCTION

How do we build robot manipulation systems that can be readily deployed in unstructured human environments? One possible answer is to learn ‘generalist’ policies that are capable of accomplishing any generic task (specified via some description *e.g.* language or image goal) in any environment. Indeed, there is a general optimism about this paradigm,

reflected in several ongoing efforts to collect large-scale demonstration datasets to train such policies. While these efforts have produced unified policies capable of diverse tasks [1]–[5], these policies still struggle to perform meaningfully when deployed zero-shot to novel environments or asked to perform unseen tasks. Deployment thus often requires additional fine-tuning using task-and-scene-specific robot demonstrations [1], but this is a non-trivial overhead as collecting robot demonstrations in the real world can be time-consuming and beyond the expertise of an average user. In this work, we propose an alternate deployment mechanism – leveraging a pre-trained generalist policy to perform a task via imitating a single *human* demonstration.

We are of course not the first to consider this goal of allowing robots to imitate any given human demonstration. One common approach [6]–[9] is to instantiate this as a kinematic retargeting task and compute open-loop robot actions that maximize a manually defined similarity between the achieved robot end-effector configuration and observed human hand poses (*e.g.* matching locations of fingertips). However, the human-robot embodiment mismatch makes it difficult for the retargeted actions to achieve precisely the same effects as the human ones, and the open-loop execution further makes the approach brittle to noise and scene variations. Another line of work for human imitation attempts to learn robot policies through online reinforcement learning [10]–[14], where the human demonstration helps define reward functions. While

this can overcome the embodiment gap, such test-time online RL requires hours of online interaction and resets, making it difficult to adopt for generic real-world manipulation tasks, especially in safety-critical scenarios. We seek an approach that, like kinematic retargeting, can be deployed one-shot (*i.e.* without test-time training), but still benefits from learned priors for precise closed-loop interaction.

Our approach builds on the insight that pre-trained diffusion policies can act as priors for robot action. Inspired by prior work in leveraging pre-trained diffusion models for image editing [15], we present *DemoDiffusion* – a formulation to utilize a diffusion policy trained on robot interaction data for synthesizing coherent robot actions from a human demonstration. Specifically, we first perform kinematic retargeting by extracting human hand poses from the demonstration and obtain an open-loop robot action trajectory. While this trajectory is typically suboptimal due to embodiment differences and lack of closed-loop feedback, it serves as an effective initialization that can be improved via a diffusion policy. We do so by injecting Gaussian noise and applying the pre-trained diffusion policy to iteratively denoise the trajectory conditioned on robot observations, yielding a refined, executable sequence of robot actions. *DemoDiffusion* thus enables the robot to use the pre-trained policy as a prior and adapt the human-derived trajectory to its own embodiment and environment in a closed-loop manner.

In summary, we propose *DemoDiffusion*: a framework for robotic manipulation that allows a robot to perform generic tasks in natural environments by following a human demonstration with guidance from a pre-trained diffusion policy. Importantly, our framework doesn’t require any robot demonstration of the target task within the target environment, online interaction, or fine-tuning. Our experiments across simulation and real-world environments show that *DemoDiffusion* surpasses both the base policy and retargeted trajectory, achieving 83.8% success across 8 real-world tasks – even those where the pre-trained *generalist* policy fails entirely.

II. RELATED WORK

Generalist Manipulation Policies. There has been a growing trend in developing ‘*generalist*’ robotic policies that can perform multiple tasks based on a specified goal in the form of an image or a language instruction. Collecting large-scale robot interaction datasets [16], [17] combined with behavior cloning and vision-language pre-training [1]–[5], [18], [19] is the predominant recipe for training such multi-task policies. However, since collecting large-scale robot interaction datasets in natural settings like homes and offices is challenging due to operational constraints, current best models still struggle to perform manipulation tasks zero-shot and thus are not yet deployable in-the-wild. Our work enables using such *generalist* diffusion policies for following a human demonstration, enabling it to perform tasks that the pre-trained policy struggles to zero-shot.

Robotic Manipulation with Non-Robot Datasets. Instead of requiring large-scale robot interaction datasets,

a growing body of works have begun utilizing human videos and large-scale web videos for robotics. These approaches have been enabled by recent advances in computer vision for representation learning [20]–[22], predicting tracks [23], [24], and reconstructing hand-object interactions from monocular videos [25]–[28], [28]. Earlier works learned self-supervised visual representations from such non-robotic datasets [29]–[31], which can serve as the visual backbone of robotic policies [32]–[37]. Recent works in this paradigm predict manipulation-relevant cues from web videos in the form of motion trajectories [38]–[40] and object affordances [41]–[44] and combine these predictive models with a limited amount of robot interaction data for training policies via conditional behavior cloning. Our work is orthogonal to these approaches in that instead of using human videos for training, we require a single human demonstration during deployment for guiding a pre-trained generalist policy to perform a new task.

One-Shot Imitation from Robot Demonstration. One-shot imitation aims to enable robots to perform a new task with guidance from a single demonstration. Prior works have explored this in the context of a robot demonstration provided at test time, via training demonstration-conditioned policies [45]. Recent works have developed learning-based visual servoing systems [46] and algorithms for identifying object invariances [47] to replicate the robot end-effector actions from the demonstration onto novel configurations of objects. However, requiring a robot demonstration for imitation is restrictive for ubiquitous deployment as end-users might find it challenging to tele-operate the robot to collect demonstrations in natural human environments.

One-Shot Imitation from Human Demonstration. To imitate a human demonstration for robot manipulation, a straightforward approach is kinematic retargeting of the human hand pose to the robot end-effector pose per time-step, as demonstrated by some recent works [6], [7], [48]. Although simple to implement, the human-robot embodiment mismatch typically introduces errors in the retargeting and the open-loop execution is brittle to object pose variations during deployment. To remedy this, another line of works seek to learn this retargeting implicitly by training a closed-loop policy from paired human-robot datasets [49], [50] enabling more flexible deployments. Since such paired datasets are difficult to collect, other works [10]–[12], [51] perform reinforcement learning with reward functions derived from video comparison. However, such methods require online interaction and resets, limiting their practicality. We propose an alternate paradigm: given the human demonstration, we refine retargeted human behavior using a pre-trained robot policy, avoiding brittle replays and enabling closed-loop control. This formulation avoids both expensive paired human-robot data collection and cumbersome test-time fine-tuning via interaction.

III. METHOD

We target the problem of one-shot visual imitation. Given a human demonstration \mathbf{D} depicting a manipulation task with

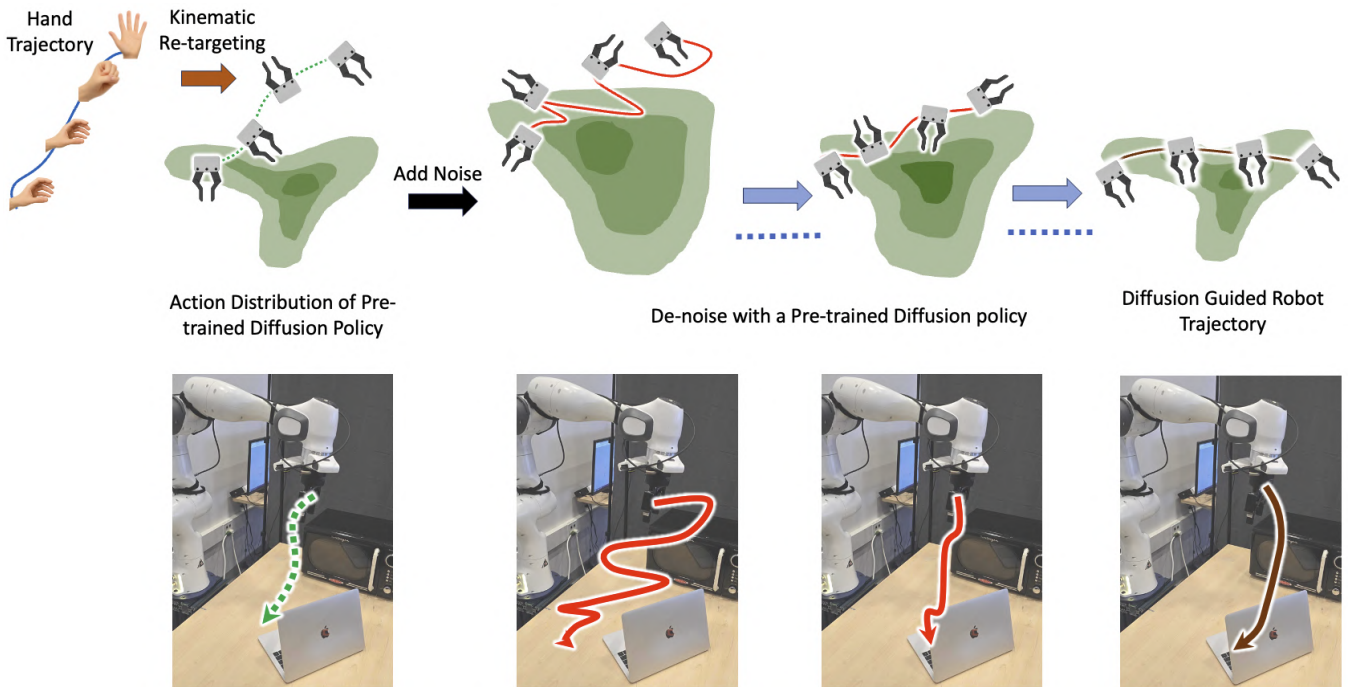


Fig. 2: **Retargeted human hand trajectory to closed-loop robot action sequence, for the task \mathcal{T} : “shut down the laptop”**. The dotted line shows the trajectory of robot end-effector poses after kinematic retargeting. The olive contour plot depicts the distribution of trajectories from a pre-trained diffusion policy. Given a kinematic retargeting, we first perturb it with Gaussian noise and progressively remove the noise by simulating the reverse SDE with the diffusion policy. This process gradually projects a potentially unfeasible but approximately correct retargeting to the manifold of plausible robot actions that can perform real-world manipulation, in this case closing the laptop without missing the edge.

description \mathcal{T} , we want a robot manipulator to perform the same task. Unlike prior work that relies on test-time RL training [10] or paired human-robot datasets [49], we wish to enable such imitation ‘one-shot’ assuming access to a closed-loop diffusion policy $\tilde{\pi}_\theta(\mathbf{a}_t|\mathbf{o}_{\leq t}, \mathcal{T})$ pre-trained on some broad robot interaction dataset \mathcal{D}_{robot} . Our assumption on the human demonstration \mathbf{D} is also very minimal: it can either be an RGBD video or a multi-view video of a human doing a task, such that 3D hand poses per-timestep can be reliably extracted from it.

A. Overview

Our approach is based on two insights. The first is that the trajectory of the hand pose in the human demonstration \mathbf{D} provides useful information of the approximate trajectory the robot end-effector should follow, and we can perform *kinematic retargeting* of the hand trajectory $\{\mathbf{h}_t\}_{t=0}^T$ to an open-loop robot end-effector trajectory $\{\hat{\mathbf{a}}_t\}_{t=0}^T$. The second insight is that the kinematically retargeted robot trajectory has the correct form of motion, but these actions may not be very precise in the distribution of plausible robot actions given the current observation. A diffusion policy models this likelihood, and we can use a pre-trained ‘generalist’ diffusion policy $\tilde{\pi}_\theta(\mathbf{a}_t|\mathbf{o}_{\leq t}, \mathcal{T})$ to refine the retargeted robot actions in a closed loop manner, thus inferring actions that, while still similar to the human demonstration, are more likely under the policy and better aligned with the robot embodiment (see

Fig. 2). We describe both these steps in detail below.

B. Kinematic retargeting of Human Hand Trajectories to Robot End-Effector Poses

Given a human demonstration \mathbf{D} of a manipulation task \mathcal{T} , our first step is to extract the 3D hand pose trajectory $\{\mathbf{h}_t\}_{t=0}^T$. Each hand pose $\mathbf{h}_t \in \mathbb{R}^{3 \times J}$ corresponds to the 3D locations of J keypoints (e.g., wrist and fingertips), which we estimate using a pre-trained monocular hand pose estimator applied to each video frame [52]. These keypoints encode the motion of the human performing the task and form the basis for retargeting to a robotic end-effector.

To translate this human motion into a robot-executable trajectory, we define a simple geometric mapping function $f_{\text{retarget}}: \mathbb{R}^{3 \times J} \rightarrow \mathbb{R}^6$ that converts the human hand pose \mathbf{h}_t into a robot configuration $\hat{\mathbf{a}}_t = f_{\text{retarget}}(\mathbf{h}_t)$. The mapping aligns the wrist pose of the human to the robot end-effector pose. For a two-finger gripper, we use the distance between the thumb and the remaining fingers of the hand mesh to infer a binary robot grasp, and for a dexterous robotic hand we match the robot hand’s fingertip positions to those of a human hand using inverse kinematics. We set the initial configuration of the robot as the kinematically retargeted initial configuration from the human demonstration.

This kinematic retargeting procedure yields a full trajectory $\{\hat{\mathbf{a}}_t\}_{t=0}^T$ in the robot’s configuration space, which can be executed as an open-loop policy. However, due to differences




Fig. 3: **Dexterous Grasping Results in Simulation.** (Left) A human demonstration and *DemoDiffusion* rollout for dexterous grasping. (Right) Success rates (mean \pm std over 3 seeds) as a function of diffusion step s^* . Here, $s^*/S = 0$ corresponds to kinematic retargeting and $s^*/S = 1$ corresponds to the robot policy.

in morphology and embodiment between humans and robots, the absence of environment feedback, and the inaccuracy of the hand estimation module, this trajectory often leads to suboptimal or unstable behavior during open-loop execution. We thus treat this as the basis for searching for a plausible robot action trajectory via denoising with a (pre-trained) closed-loop diffusion policy.

C. Closed-Loop Denoising of Robot Actions with a Pre-Trained Diffusion Policy

A diffusion policy $\bar{\pi}_\theta(\mathbf{a}_t | \mathbf{o}_{\leq t}, \mathcal{T})$ [53], [54] pre-trained on diverse offline robot interaction data reliably models the distribution of plausible robot actions \mathbf{a}_t given previous observations $\mathbf{o}_{\leq t}$. To predict actions \mathbf{a}_t , diffusion policies start with Gaussian noise $\tilde{\mathbf{a}}_t^{(S)} \sim \mathbf{N}(0, I)$ and iteratively denoise it at different diffusion steps s :

$$\tilde{\mathbf{a}}_t^{(s-1)} = \bar{\pi}_\theta(\tilde{\mathbf{a}}_t^{(s)}, \mathbf{o}_{\leq t}), \quad s = S, S-1, \dots, 0 \quad (1)$$

We use this to modify the kinematically retargeted trajectory $\{\hat{\mathbf{a}}_t\}_{t=0}^T$ for obtaining robot actions that follow the high-level motion in the human trajectory and still lie within the distribution of plausible actions under $\bar{\pi}_\theta$. To do this, we modify this typical reverse diffusion process so that instead of starting with pure noise $\tilde{\mathbf{a}}_t^{(S)}$ at step S , we start from an intermediate diffusion step s^* such that $0 < s^* < S$. We define $\{\tilde{\mathbf{a}}_t^{(s^*)}\}_{t=0}^T$ to be a noisy version of the kinematically retargeted trajectory as follows:

$$\tilde{\mathbf{a}}_t^{(s^*)} = \sqrt{\alpha_{s^*}} \hat{\mathbf{a}}_t + \sqrt{1 - \alpha_{s^*}} \varepsilon_t, \quad \varepsilon_t \sim \mathbf{N}(0, I) \quad (2)$$

Here α corresponds to the diffusion schedule of the pre-trained policy. This procedure, inspired by SDEdit [15] which adopted a similar approach for image editing, relies on the assumption that $\{\hat{\mathbf{a}}_t\}_{t=0}^T$ is an approximate version of the ideal action trajectory that should be executed by the robot, and it potentially lies outside the distribution of feasible actions under the pre-trained diffusion policy $\bar{\pi}_\theta$. Intuitively, adding moderate noise to the retargeted trajectory and then

Group	Robot Policy	Kinematic	<i>DemoDiffusion</i>
<i>Small</i>	25.4	2.6	31.8
<i>Medium</i>	26.6	0.7	31.6
<i>Large</i>	27.4	1.5	29.6
Average	26.5	1.6	31.0

TABLE I: **Simulated Dexterous Grasping Results.** The numbers show the average of success rates over 3 seeds. We use $s^*/S = 0.2$, where the robot policy uses $S = 1000$. *DemoDiffusion* consistently outperforms both baselines, across all groups.

denoising preserves the high-level structure of the human demonstration while allowing the policy to correct low-level infeasibilities from the embodiment gap. The diffusion policy $\bar{\pi}_\theta$ then performs iterative denoising steps, conditioned on the robot’s observations $\mathbf{o}_{\leq t}$, to refine this noisy trajectory into feasible robot actions, based on equation (2).

After s^* denoising steps, the final output $\mathbf{a}_t = \tilde{\mathbf{a}}_t^{(0)}$ is deployed on the robot. Importantly, this process is carried out in a closed-loop manner: the policy uses real-time observations from the cameras in the scene to iteratively improve its predictions, thereby compensating for embodiment mismatch and external perturbations (e.g., object slippage or occlusion). The key hyperparameter in this process for *DemoDiffusion* is the diffusion step s^* that trades off between the faithfulness to the demonstration and the likelihood under the robot policy – in the limit $s^* = S$ we recover a rollout from the base policy $\bar{\pi}_\theta$ and in the limit $s^* = 0$ we recover the kinematically retargeted trajectory $\{\hat{\mathbf{a}}_t\}_{t=0}^T$. While derived in the context of diffusion, note that this iterative refinement perspective also applies to flow-matching [55], which we show in real-world experiments.

IV. EXPERIMENTS

We evaluate our approach on dexterous grasping in simulation and across diverse real-world tasks comprising prehensile and non-prehensile table-top manipulation. Through

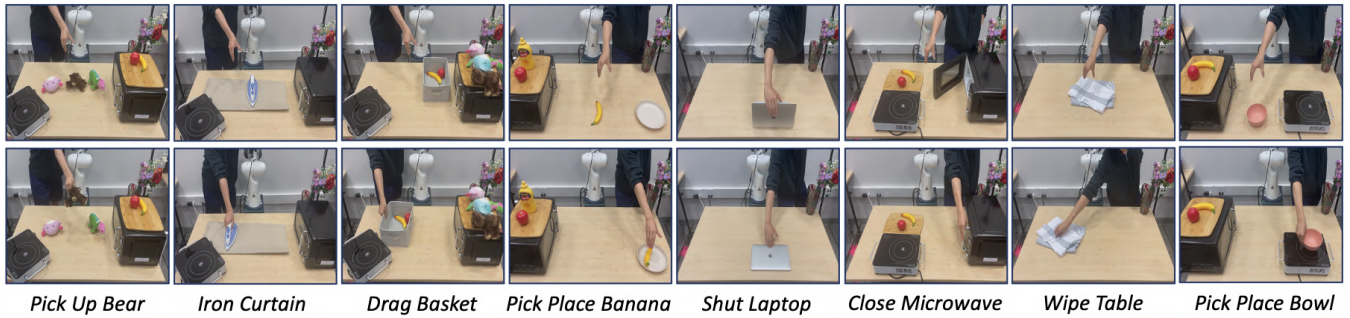


Fig. 4: **Real-World Manipulation Tasks.** Human demonstrations for the 8 evaluation tasks, shown as two frames per task. Tasks span prehensile and non-prehensile manipulation, including grasping, pushing, closing, wiping, and placing.

experiments, we aim to understand the following research questions:

- Can *DemoDiffusion* outperform pure kinematic retargeting from the human demonstration?
- Do the human demonstrations allow *DemoDiffusion* to perform new tasks where the pre-trained diffusion policy fails?
- How to effectively tradeoff between faithfulness to the human demonstration and performing the task reliably, with varying noise level s^*/S ?
- How robust is *DemoDiffusion* to noise in human demonstration, or different kinematic retargeting?

A. Dexterous Grasping in Simulation

To verify our intuition, we first consider a simulation environment where the target task is restricted to picking up a generic object with a 16-DOF four-fingered Allegro hand. Specifically, we train a dexterous grasping policy across a small set of generic objects in simulation, and test our method on human hand grasping trajectories on a different set of objects. This serves as the pre-trained diffusion model for this experiment.

For training the robot policy, we collect a total of 985 grasping trajectories of Allegro hand over 58 training objects [56] (randomly sampled 26 and 32 objects from ShapeNet [57] and PartNet [58], respectively), resulting in 22.2% success rate. We use a variant of 3D Diffusion Policy [59], which takes 3D point tracks of an object instead of point clouds. At test time, we provide human grasping trajectories on a subset of Objaverse dataset [60], which are unseen during training. The test set contains 1220 objects, with 1 human grasping trajectory per object, provided by GraspXL [56]. Note that there is no error from hand estimation or the initial configuration of the object in this setup. We use ground-truth human hand 3D keypoints, and the object is located at the same location as the human demonstration.

Results. As in GraspXL [56], we group the objects based on their size, resulting in small, medium, and large objects. We compare *DemoDiffusion* with two baselines. The first is open-loop kinematic retargeting obtained from the human demonstration. The second is directly deploying the robot policy we trained.




Fig. 5: **Workspace with 5 Cameras.** We use the four external cameras for triangulation to obtain the global pose of the hand mesh from a human demonstration. The pre-trained policy uses the two cameras marked in purple.

Table I shows the main results. Our method outperforms baselines, with the largest gains on small objects. We hypothesize this is because larger objects are easier to grasp, supported by base robot policy performance improving for larger objects. This experiment shows that even in a highly controlled setup where the robot environment is initialized identically to that of human demonstration, kinematic retargeting cannot solve the task, while slightly refining with a generalist grasping policy significantly improves the performance. Using kinematic retargeting to initiate the denoising process also helps the denoising process itself, outperforming the robot policy. Additionally, *DemoDiffusion* achieves higher inference speed, by the factor of S/s^* .

Ablation studies. We analyze the influence of s^* on *DemoDiffusion* in Fig. 3. Overall, we observe a consistent trend of performance gain from $s^*/S = 1$ to $s^*/S = 0.2$, showing that *DemoDiffusion* is robust to the choice of hyperparameter.

B. Real-World Manipulation

For the real-world experiments, we use a Franka Emika Panda arm equipped with a two-finger gripper from Robotiq. We use the publicly available pre-trained flow-matching policy $\bar{\pi}_\theta(\mathbf{a}_t | \mathbf{o}_{\leq t}, \mathcal{T})$ called Pi-0 released by Physical In-

Task	Pi-0	Kinematic	<i>DemoDiffusion</i>
Shut Down Laptop	20	10	60
Close Microwave	10	80	90
Drag Basket	10	80	80
Wipe Table	50	0	100
Iron Curtain	20	70	90
Pick up Bear	0	40	60
Pick Place Bowl	0	70	100
Pick Place Banana	0	70	90
Average	13.8	52.5	83.8

TABLE II: **Quantitative Results for Real-World Manipulation.** Success rates (%) are over 10 trials per task.

telligence [1], trained on the DROID dataset [61]. This policy takes as input a language task specification \mathcal{T} , and observations from two cameras in the scene, and outputs per-timestep joint velocity. This policy was trained on a large offline dataset of tele-operated robot demonstrations, and is a *generalist* policy that can perform a broad set of tasks. We do not fine-tune or adapt this policy in any way and directly use it for *DemoDiffusion*. We use an off-the-shelf monocular hand reconstruction model, Hamer [52] to obtain per-timestep hand mesh reconstructions from each of the external cameras given a human demonstration in the scene. Our experiments follow the protocol of a human first demonstrating an object manipulation in the scene, followed by robot execution after re-setting the scene.

Workspace. Fig. 5 describes our real-robot manipulation setup. The scene has 5 cameras including 4 Realsense external cameras and 1 Zed-Mini wrist camera mounted on the robot gripper. We calibrate the four external cameras and use them for triangulating the 3D hand pose so that we can obtain the position and orientation of a human hand in the world coordinates, with origin at the base of the robot. The pre-trained policy Pi-0 requires just two cameras (marked in purple) and the policy does not use any calibration information. While Pi-0 uses Zed 2 camera for external view, we use Realsense camera due to hardware availability, which introduces additional difficulty.

Tasks. We perform evaluations for 8 different manipulation tasks, the language description and the human demonstration for which are described in Fig. 4. Each of these tasks requires a different manipulation strategy involving either prehensile or non-prehensile manipulation. We collect just a single human demonstration per task but for robot execution, we make organic variations in the location of objects in the scene. Additionally, we add other objects in the scene as visual distractors, to see whether *DemoDiffusion* can generalize under such settings. All variations are consistent across the baselines and our method.

Baselines. For the real-world experiments we compare *DemoDiffusion* with two baselines. The first is open-loop kinematic retargeting obtained from the human demonstration. The second is directly deploying the Pi-0 model with the language instruction for the respective task. Since Pi-0 operates in the joint velocity space, for *DemoDiffusion* we



Fig. 6: **Qualitative comparisons for real-robot manipulation.** Rollout progressions (start, intermediate, final frame) for two tasks. Kinematic retargeting (top) produces plausible motion but loses contact before completing the task. Pi-0 (middle) performs general reaching but fails to manipulate the correct object. *DemoDiffusion* (bottom) reaches the object and maintains contact through task completion. See project page for full videos.

	Shut Laptop	Pick Place Banana
<i>DemoDiffusion</i>	60	90
+5cm noise	40	70
with Thumb+Index	30	80
Thumb+Index	0	30

TABLE III: **Robustness of *DemoDiffusion*.** We ablate *DemoDiffusion* under noisy 3D keypoints (2nd row) and with different kinematic retargeting (3rd row). For noisy 3D keypoints, we randomly shift the tracked hand 3D keypoints by 5cm. For different kinematic retargeting, we use the distance between the thumb tip and the index fingertip instead of all fingers.

first convert the kinematically retargeted trajectory to the joint velocity space with IK and then add noise to the trajectory, followed by denoising with Pi-0. We deploy *DemoDiffusion* with a heuristically selected noise level s^*/S depending on the task – preferring the kinematic trajectory with $s^*/S = 0.2$ when Pi-0 has a zero success rate and $s^*/S = 0.4$ otherwise. Note that in the limit of 0 noise added, we recover the first baseline, kinematic retargeting, and in the limit of maximum noise 1 we obtain the second baseline Pi-0.

Qualitative Visualizations. In Fig. 6, we show rollouts for two tasks across all the methods. We observe that Pi-0 can only perform general reaching behaviors: it either fails to follow the manipulation trajectory needed for the task or confuses the target object with a distractor. On the other hand, we observe that (open-loop) kinematic retargeting from the human demonstration has a plausible motion of the robot end-effector, but it completely misses grasping the iron, and loses contact with the laptop before fully closing it. Compared to these, *DemoDiffusion* reaches the respective object and also follows the approximate motion of the retargeted trajectory while remaining in contact with the object to solve the tasks.

Quantitative Comparisons. In Table II, we compare *DemoDiffusion* and the baselines across all the real-world manipulation tasks. *DemoDiffusion* matches or outperforms both baselines across all 8 tasks, with the largest gains on tasks requiring precise contact – e.g., 60% vs. 10–20% on shutting the laptop and 100% vs. 0–50% on wiping the table. Notably, even for tasks where both baselines fail, *DemoDiffusion* succeeds by operating at an intermediate noise level s^* that combines the trajectory structure from the human demonstration with the policy’s closed-loop refinement.

Noisy Human Demonstrations. To evaluate the robustness of *DemoDiffusion* to imperfect inputs, we injected noise into the tracked human trajectories and evaluated *DemoDiffusion*. Specifically, we place the object at the same location as the human demonstration, and randomly shift the tracked hand 3D keypoints by 5cm. The results in Table III indicate that *DemoDiffusion* maintains high-quality performance despite the presence of tracking noise, thereby reducing the reliance on highly accurate 3D annotations in human demonstration.

Different Kinematic Retargeting. Lastly, we further assessed the sensitivity of our method to the choice of retargeting by evaluating an alternative approach (‘Thumb+Index’). Specifically, we use the distance between thumb and index finger [48], instead of all fingers, to infer the grasp. Table III shows the results. While this alternative retargeting itself shows degraded performance compared to the previous retargeting, *DemoDiffusion* consistently improves task performance despite suboptimal retargeting.

V. DISCUSSION

We presented *DemoDiffusion*, a simple approach to leverage *generalist* diffusion models for imitating a human demonstration. By retargeting hand motion from the human demonstration into an approximate robot end-effector trajectory and modifying it through de-noising with a pre-trained diffusion policy, *DemoDiffusion* enables the robot to produce actions that are both aligned with the demonstrated intent and contextually plausible. Our method consistently outperforms both the base diffusion policy and the raw retargeted trajectory across diverse tasks in simulation and the real world, without requiring additional training, paired demonstrations, or reward functions. We envision that this can serve as a starting point for future efforts at human imitation across generic scenarios and perhaps also yield an improved exploration strategy for methods pursuing policy adaptation with online RL.

LIMITATIONS

Our approach has several limitations. First, it assumes that the robot should act similarly to the human to successfully complete the task, which may not hold in scenarios requiring different strategies due to embodiment or environmental constraints. Second, while effective for one-shot imitation (not requiring training / test-time fine-tuning), the method does not produce a reusable task policy that can generalize

across arbitrary variations of the task. Third, the quality of the retargeted trajectory is crucial: accurate 3D human motion capture is challenging and errors in retargeting can impact downstream performance, though our method shows a degree of robustness to such noise. Additionally, the method implicitly assumes that the timing and speed of human and robot actions are aligned; extending the approach to allow for temporal alignment at test time is a promising direction of future work.

ACKNOWLEDGEMENTS

We appreciate the helpful discussions with Yanbo Xu, Qitao Zhao, and Lucas Wu. We would also like to thank Kenny Shaw, Tony Tao, Jiahui Yang, Andrew Wang, Jason Liu, Hengkai Pan, Mohan Kumar Srirama, Peiyuan Yang for helping set up the hardware. This work was supported by gift awards from CISCO, Google, and NSF Award IIS-2442282.

REFERENCES

- [1] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. $\pi 0$: A vision-language-action flow model for general robot control. In *RSS*, 2025.
- [2] Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *CoRL*, 2023.
- [3] Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey Levine, and Chelsea Finn. Be-z: Zero-shot task generalization with robotic imitation learning. In *CoRL*, 2024.
- [4] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. In *CoRL*, 2024.
- [5] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. In *RSS*, 2023.
- [6] Jinhan Li, Yifeng Zhu, Yuqi Xie, Zhenyu Jiang, Mingyo Seo, Georgios Pavlakos, and Yuke Zhu. Okami: Teaching humanoid robots manipulation skills through single video imitation. In *CoRL*, 2024.
- [7] Georgios Papagiannis, Norman Di Palo, Pietro Vitiello, and Edward Johns. R+ x: Retrieval and execution from everyday human videos. In *ICRA*, 2025.
- [8] Chen Wang, Haochen Shi, Weizhuo Wang, Ruohan Zhang, Li Fei-Fei, and C Karen Liu. Dexcap: Scalable and portable mocap data collection system for dexterous manipulation. In *RSS*, 2024.
- [9] Sirui Chen, Chen Wang, Kaden Nguyen, Li Fei-Fei, and C Karen Liu. Arcap: Collecting high-quality human demonstrations for robot learning with augmented reality feedback. In *ICRA*, 2025.
- [10] Shikhar Bahl, Abhinav Gupta, and Deepak Pathak. Human-to-robot imitation in the wild. In *RSS*, 2022.
- [11] Irmak Guzey, Yinlong Dai, Georgy Savva, Raunaq Bhirangi, and Lerrel Pinto. Bridging the human to robot dexterity gap through object-oriented rewards. In *ICRA*, 2025.
- [12] Yuanpei Chen, Chen Wang, Yaodong Yang, and C Karen Liu. Object-centric dexterous manipulation from human motion data. In *CoRL*, 2024.
- [13] Xueyi Liu, Jianibieke Adalibieke, Qianwei Han, Yuzhe Qin, and Li Yi. Dextrack: Towards generalizable neural tracking control for dexterous manipulation from human references. In *ICLR*, 2025.
- [14] Xueyi Liu, Kangbo Lyu, Jieqiong Zhang, Tao Du, and Li Yi. Quasisim: Parameterized quasi-physical simulators for dexterous manipulations transfer. In *ECCV*, 2024.
- [15] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *ICLR*, 2022.
- [16] Hao-Shu Fang, Hongjie Fang, Zhenyu Tang, Jirong Liu, Junbo Wang, Haoyi Zhu, and Cewu Lu. Rh20t: A robotic dataset for learning diverse skills in one-shot. In *ICRA*, 2024.

- [17] Homer Rich Walke, Kevin Black, Tony Z Zhao, Quan Vuong, Chongyi Zheng, Philippe Hansen-Estruch, Andre Wang He, Vivek Myers, Moo Jin Kim, Max Du, et al. Bridgedata v2: A dataset for robot learning at scale. In *CoRL*, 2023.
- [18] Homanga Bharadhwaj, Jay Vakil, Mohit Sharma, Abhinav Gupta, Shubham Tulsiani, and Vikash Kumar. Roboagent: Generalization and efficiency in robot manipulation via semantic augmentations and action chunking. In *ICRA*, 2024.
- [19] Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, et al. $\pi_{0.5}$: a vision-language-action model with open-world generalization. In *CoRL*, 2025.
- [20] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022.
- [21] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.
- [22] Lingting Zhu, Guying Lin, Jinnan Chen, Xinjie Zhang, Zhenchao Jin, Zhao Wang, and Lequan Yu. Large images are gaussians: High-quality large image representation with levels of 2d gaussian splatting. In *AAAI*, 2025.
- [23] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Cotracker: It is better to track together. In *ECCV*, 2024.
- [24] Carl Doersch, Ankush Gupta, Larisa Markeeva, Adrià Recasens, Lucas Smaira, Yusuf Aytar, João Carreira, Andrew Zisserman, and Yi Yang. Tap-vid: A benchmark for tracking any point in a video. In *NeurIPS*, 2022.
- [25] Yu Rong, Takaaki Shiratori, and Hanbyul Joo. Frankmocap: A monocular 3d whole-body pose estimation system via regression and integration. In *CVPR*, 2021.
- [26] Georgios Pavlakos, Dandan Shan, Ilija Radosavovic, Angjoo Kanazawa, David Fouhey, and Jitendra Malik. Reconstructing hands in 3d with transformers. In *CVPR*, 2024.
- [27] Yufei Ye, Abhinav Gupta, and Shubham Tulsiani. What's in your hands? 3d reconstruction of generic objects in hands. In *CVPR*, 2022.
- [28] Yufei Ye, Poorvi Hebbbar, Abhinav Gupta, and Shubham Tulsiani. Diffusion-guided reconstruction of everyday hand-object interaction clips. In *ICCV*, 2023.
- [29] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The epic-kitchens dataset. In *ECCV*, 2018.
- [30] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *CVPR*, 2022.
- [31] Kristen Grauman, Andrew Westbury, Lorenzo Torresani, Kris Kitani, Jitendra Malik, Triantafyllos Afouras, Kumar Ashutosh, Vijay Baiyya, Siddhant Bansal, Bikram Boote, et al. Ego-exo4d: Understanding skilled human activity from first-and third-person perspectives. In *CVPR*, 2024.
- [32] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3m: A universal visual representation for robot manipulation. In *CoRL*, 2022.
- [33] Simone Parisi, Aravind Rajeswaran, Senthil Purushwalkam, and Abhinav Gupta. The unsurprising effectiveness of pre-trained vision models for control. In *ICML*, 2022.
- [34] Siddharth Karamcheti, Suraj Nair, Annie S Chen, Thomas Kollar, Chelsea Finn, Dorsa Sadigh, and Percy Liang. Language-driven representation learning for robotics. In *RSS*, 2023.
- [35] Tete Xiao, Ilija Radosavovic, Trevor Darrell, and Jitendra Malik. Masked visual pre-training for motor control. *arXiv preprint arXiv:2203.06173*, 2022.
- [36] Yecheng Jason Ma, Shagun Sodhani, Dinesh Jayaraman, Osbert Bastani, Vikash Kumar, and Amy Zhang. Vip: Towards universal visual reward and representation via value-implicit pre-training. In *ICLR*, 2023.
- [37] Yecheng Jason Ma, Vikash Kumar, Amy Zhang, Osbert Bastani, and Dinesh Jayaraman. Liv: Language-image representations and rewards for robotic control. In *ICML*, 2023.
- [38] Homanga Bharadhwaj, Roozbeh Mottaghi, Abhinav Gupta, and Shubham Tulsiani. Track2act: Predicting point tracks from internet videos enables generalizable robot manipulation. In *ECCV*, 2024.
- [39] Kenneth Shaw, Shikhar Bahl, and Deepak Pathak. Videodex: Learning dexterity from internet videos. In *CoRL*, 2023.
- [40] Tony Tao, Mohan Kumar Srirama, Jason Jingzhou Liu, Kenneth Shaw, and Deepak Pathak. Dexwild: Dexterous human interactions for in-the-wild robot policies. *arXiv preprint arXiv:2505.07813*, 2025.
- [41] Mohan Kumar Srirama, Sudeep Dasari, Shikhar Bahl, and Abhinav Gupta. Hrp: Human affordances for robotic pre-training. In *RSS*, 2024.
- [42] Shikhar Bahl, Russell Mendonca, Lili Chen, Unnat Jain, and Deepak Pathak. Affordances from human videos as a versatile representation for robotics. In *CVPR*, 2023.
- [43] Russell Mendonca, Shikhar Bahl, and Deepak Pathak. Structured world models from human videos. In *RSS*, 2023.
- [44] Junyao Shi, Zhuolun Zhao, Tianyou Wang, Ian Pedroza, Amy Luo, Jie Wang, Jason Ma, and Dinesh Jayaraman. Zeromimic: Distilling robotic manipulation skills from web videos. In *ICRA*, 2025.
- [45] Sudeep Dasari and Abhinav Gupta. Transformers for one-shot visual imitation. In *CoRL*, 2020.
- [46] Eugene Valassakis, Georgios Papagiannis, Norman Di Palo, and Edward Johns. Demonstrate once, imitate immediately (dome): Learning visual servoing for one-shot imitation learning. In *IROS*, 2022.
- [47] Xinyu Zhang and Abdeslam Boularias. One-shot imitation learning with invariance matching for robotic manipulation. In *RSS*, 2024.
- [48] Marion Lepert, Jiaying Fang, and Jeannette Bohg. Phantom: Training robots without robots using only human videos. In *CoRL*, 2025.
- [49] Vidhi Jain, Maria Attarian, Nikhil J Joshi, Ayzaan Wahid, Danny Driess, Quan Vuong, Pannag R Sanketi, Pierre Sermanet, Stefan Welker, Christine Chan, et al. Vid2robot: End-to-end video-conditioned policy learning with cross-attention transformers. In *RSS*, 2024.
- [50] Sungjae Park, Seunggho Lee, Mingi Choi, Jiye Lee, Jeonghwan Kim, Jisoo Kim, and Hanbyul Joo. Learning to transfer human hand skills for robot manipulations. *arXiv preprint arXiv:2501.04169*, 2025.
- [51] Tyler Ga Wei Lum, Olivia Y Lee, C Karen Liu, and Jeannette Bohg. Crossing the human-robot embodiment gap with sim-to-real rl using one human demonstration. *arXiv preprint arXiv:2504.12609*, 2025.
- [52] Georgios Pavlakos, Dandan Shan, Ilija Radosavovic, Angjoo Kanazawa, David Fouhey, and Jitendra Malik. Reconstructing hands in 3d with transformers. In *CVPR*, 2024.
- [53] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *IJRR*, 2023.
- [54] Moritz Reuss, Maximilian Li, Xiaogang Jia, and Rudolf Lioutikov. Goal conditioned imitation learning using score-based diffusion policies. In *RSS*, 2023.
- [55] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *ICLR*, 2023.
- [56] Hui Zhang, Sammy Christen, Zicong Fan, Otmar Hilliges, and Jie Song. Grasppl: Generating grasping motions for diverse objects at scale. In *ECCV*, 2024.
- [57] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- [58] Kaichun Mo, Shilin Zhu, Angel X Chang, Li Yi, Subarna Tripathi, Leonidas J Guibas, and Hao Su. Partnet: A large-scale benchmark for fine-grained and hierarchical part-level 3d object understanding. In *CVPR*, 2019.
- [59] Yanjie Ze, Gu Zhang, Kangning Zhang, Chenyuan Hu, Muhan Wang, and Huazhe Xu. 3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations. In *RSS*, 2024.
- [60] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *CVPR*, 2023.
- [61] Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis, et al. Droid: A large-scale in-the-wild robot manipulation dataset. In *RSS*, 2024.
- [62] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. In *Transactions on Graphics*, 2017.