

Dexterous Manipulation Policies from RGB Human Videos via 3D Hand-Object Trajectory Reconstruction

Hongyi Chen¹, Tony Dong¹, Tiancheng Wu¹, Liquan Wang², Yash Jangir¹, Yaru Niu¹,
Yufei Ye³, Homanga Bharadhwaj¹, Zackory Erickson^{1,†}, Jeffrey Ichnowski^{1,†}

¹ Carnegie Mellon University, ² Georgia Institute of Technology, ³ Stanford University, [†] Equal advising

videomanip.github.io

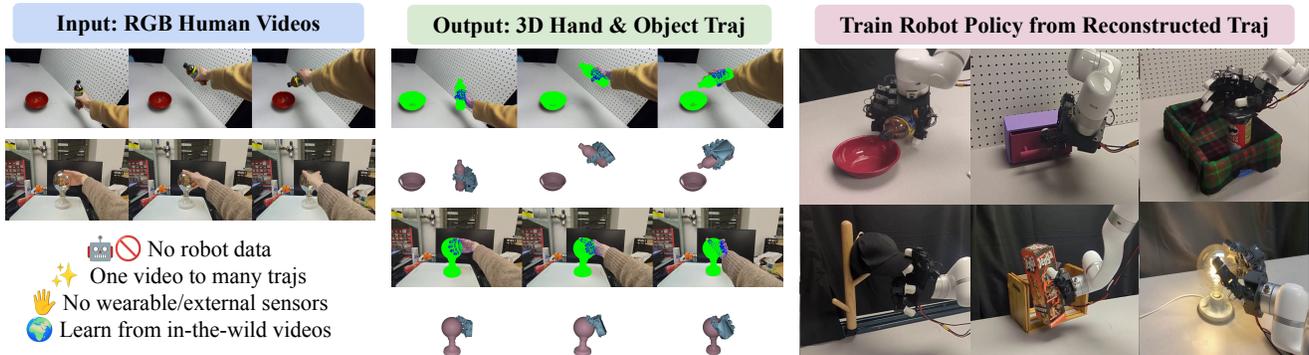


Fig. 1: Given RGB human videos, including both in-scene and in-the-wild recordings (left), the proposed VIDEOMANIP framework reconstructs trajectories of hand and object meshes (middle, projected onto RGB frames) and retargets them into robot actions to train policies. This enables dexterous manipulation learning directly from human videos without robot data or wearable/external sensing devices (right), offering a scalable alternative to conventional multifinger robot data collection pipelines.

Abstract—Multi-finger robotic hand manipulation and grasping are challenging due to the high-dimensional action space and the difficulty of acquiring large-scale training data. Existing approaches largely rely on human teleoperation with wearable devices or specialized sensing equipment to collect robot demonstrations, limiting scalability. In this work, we propose VIDEOMANIP, a device-free framework that learns dexterous multi-finger manipulation directly from RGB human videos. Leveraging recent advances in computer vision, VIDEOMANIP reconstructs explicit robot-object trajectories from monocular videos by estimating human hand poses, object meshes, and retargets the reconstructed human motions to robot hands for manipulation learning. To make the reconstructed robot data suitable for dexterous manipulation training, we introduce hand-object contact optimization with interaction-centric grasp modeling, as well as a demonstration synthesis strategy that generates diverse training trajectories from a single video, enabling generalizable policy learning without additional robot demonstrations. In simulation, the learned grasping model achieves a 70.25% success rate across 20 diverse objects using the Inspire Hand. In the real world, manipulation policies trained from RGB videos achieve an average 62.86% success rate across seven tasks using the LEAP Hand, outperforming retargeting-based methods by 15.87%. Project videos are available at videomanip.github.io.

I. INTRODUCTION

To what extent can dexterous robot manipulation be learned using human video-only supervision? Existing works have trained manipulation policies through human data obtained via hand-held grippers [1], wearables such as smart glasses [2] and headsets [3], [4], or multi-camera studio se-

tups [5]. Although effective, these methods typically require specialized hardware, controlled capture environments, or direct human involvement, limiting scalability. In contrast, RGB videos are ubiquitous and available at internet scale, offering the potential for scalable learning without specialized hardware. However, extracting accurate supervision for robot-object interactions from such data remains challenging, as these videos lack robot actions and precise 3D information. In this work, we propose a framework for learning dexterous manipulation from RGB-only human videos alone. We evaluate this framework under progressively relaxed data conditions, from in-scene videos to in-the-wild videos.

Prior work explored leveraging egocentric human videos for manipulation learning; however, these approaches either have not been successfully deployed on multi-fingered robotic hands [6], [7], rely on robot demonstrations to fine-tune models pre-trained on in-the-wild dataset [8], [9], [10], are trained on videos collected in deployment-specific scenarios [11], or assume access to pre-scanned object models [12]. Several studies reconstruct human hand poses or object meshes from images or videos of humans manipulating objects. However, these methods either have not demonstrated effective use of the reconstructed information for learning dexterous robotic hand manipulation [13], [14], [15], or fail to exploit the reconstructed trajectories for direct imitation learning, instead relying on designed reward functions for manipulation policy learning [16] or requiring additional policy fine-tuning [9]. In parallel, another line of research

focuses on human teleoperation of robotic hands [17], [18], often requiring head-mounted vision devices or specialized sensing equipment to track human hands [4], [3], [19]. These dependencies on specialized hardware limit scalability and restrict data diversity.

To address these challenges, we propose VIDEOMANIP, a framework that learns dexterous grasping and manipulation entirely from RGB human videos, without requiring wearables or external sensor devices, robot demonstrations for post-training, or pre-scanned object meshes (Fig. 1). Leveraging recent advances in computer vision, such as metric depth estimation and image-to-mesh reconstruction, VIDEOMANIP recovers explicit 3D hand-object trajectories from monocular videos by estimating human hand poses, object meshes, and object scales, and subsequently retargets the reconstructed human motions to robot hands. For *in-scene videos*, we assume known camera-robot extrinsic calibration and use it to transform reconstructed trajectories from the camera frame to the robot base frame for policy training. In contrast, for *in-the-wild videos*, which are recorded outside robot setups without extrinsic calibration, we estimate the gravity direction from visual observations and align camera-centric trajectories with a physically meaningful world frame.

To make the reconstructed robot data applicable for dexterous manipulation training, we introduce two components that improve the of grasp-model training and the generalization of manipulation policies. First, we perform (i) differential hand pose optimization by predicting hand-object contact maps to encourage physically plausible interactions, and (ii) interaction-centric grasp modeling to exploit the optimized contact map for valid grasp learning. Second, we leverage DemoGen [20] to synthesize diverse demonstrations from a single reconstructed video trajectory, enabling one-to-many trajectory generation for generalizable manipulation policy training. We evaluate VIDEOMANIP on two task categories: object grasping in simulation, and object manipulation in real-world. Comprehensive evaluations and ablations show that the learned grasp model achieves a 70.25% success rate across 20 objects using the Inspire Hand, while the learned manipulation policies with the LEAP Hand achieve an average 62.86% success rate across seven real-world manipulation tasks, including three tasks learned from in-scene videos and four tasks learned from in-the-wild videos. **Contributions.** (i) VIDEOMANIP: a framework for learning multi-fingered robotic hand grasping and manipulation directly from RGB human videos alone, without requiring robot demonstrations, wearable devices, or external sensing hardware; (ii) An explicit, trainable 3D robot hand-object trajectory reconstruction pipeline from both in-scene and in-the-wild human videos; (iii) Using reconstructed robot data, we enable dexterous robot hands to perform diverse object grasping and generalizable manipulation tasks from a single RGB human video per object or task.

II. RELATED WORK

A. Manipulation Learning from Human Videos

Large-scale RGB videos of humans interacting with the world are abundant, yet learning dexterous manipulation

policies from them remains challenging because human videos do not directly provide robot-executable actions, and hand-object interactions are often ambiguous. Prior works leverages human videos to learn manipulation priors [8], [9], [21], [22], but typically still require robot demonstrations for fine-tuning, limiting generalization from human data alone. Other approaches extract intermediate representations such as object affordances [23], [24], [6], point flows [7], [25], [26], [27], or hand poses [12], [11]. However, these methods are largely limited to coarse manipulation or gripper-based tasks [28], [16], [29], [15].

Another line of works demonstrate dexterous manipulation without robot demonstrations, but rely on specialized sensing hardware. AINA [2] uses smart glasses to capture hand-object point clouds, while related methods depend on wearable headsets [3], [19] or hand motion capture systems [12], [30], introducing additional hardware dependencies. In parallel, world models [31], [32] and video generative models [27], [33], [34] enable zero-shot manipulation by predicting future object trajectories, primarily for gripper-based tasks. LVP [35] extends this approach to dexterous hands by retargeting predicted human hand motions, but it does not explicitly model hand-object contacts and can produce infeasible motions. In contrast, our work emphasizes contact-aware grasp modeling and reconstructs reliable, learnable 3D hand-object trajectories from human videos, while remaining complementary to video generation based methods.

B. Human-Object Interaction Reconstruction from Videos

Reconstructing hands and objects from images and videos is a long-studied research area. Modern approaches leverage powerful models, such as transformers, to infer low-dimensional hand pose and shape representations [36], [37]. With strong 3D supervision, it is feasible to learn shared models across multiple object categories using diverse object representations, including meshes [38], point clouds [39] and primitives [40]. With these advancements, joint reasoning about hand-object interactions from images and videos has become increasingly feasible [41], [13], [42], leveraging methods like diffusion models [43], implicit signed distance fields [44], [45] and differentiable pose estimation to optimize the 3D hand-object contact [46], [47]. Recent works have utilized reconstructed hand-object interactions from images and videos to train dexterous robot grasping and manipulation policies [14], [9], [48], often by explicitly retargeting human hand motions to robot actions [49], [50]. Existing methods, however, either require extensive filtering due to reconstruction inaccuracies [14], incur long reconstruction times [45], or rely on additional robot demonstrations because reconstructed data alone is insufficient [9]. In contrast, VIDEOMANIP learns manipulation policies directly from robot trajectories reconstructed from RGB human videos.

III. METHOD

In this section, we present (i) the overall pipeline for reconstructing hand mesh \mathbf{H} and object mesh \mathbf{O} from RGB human videos in Sec III-A, which yields explicit 3D robot

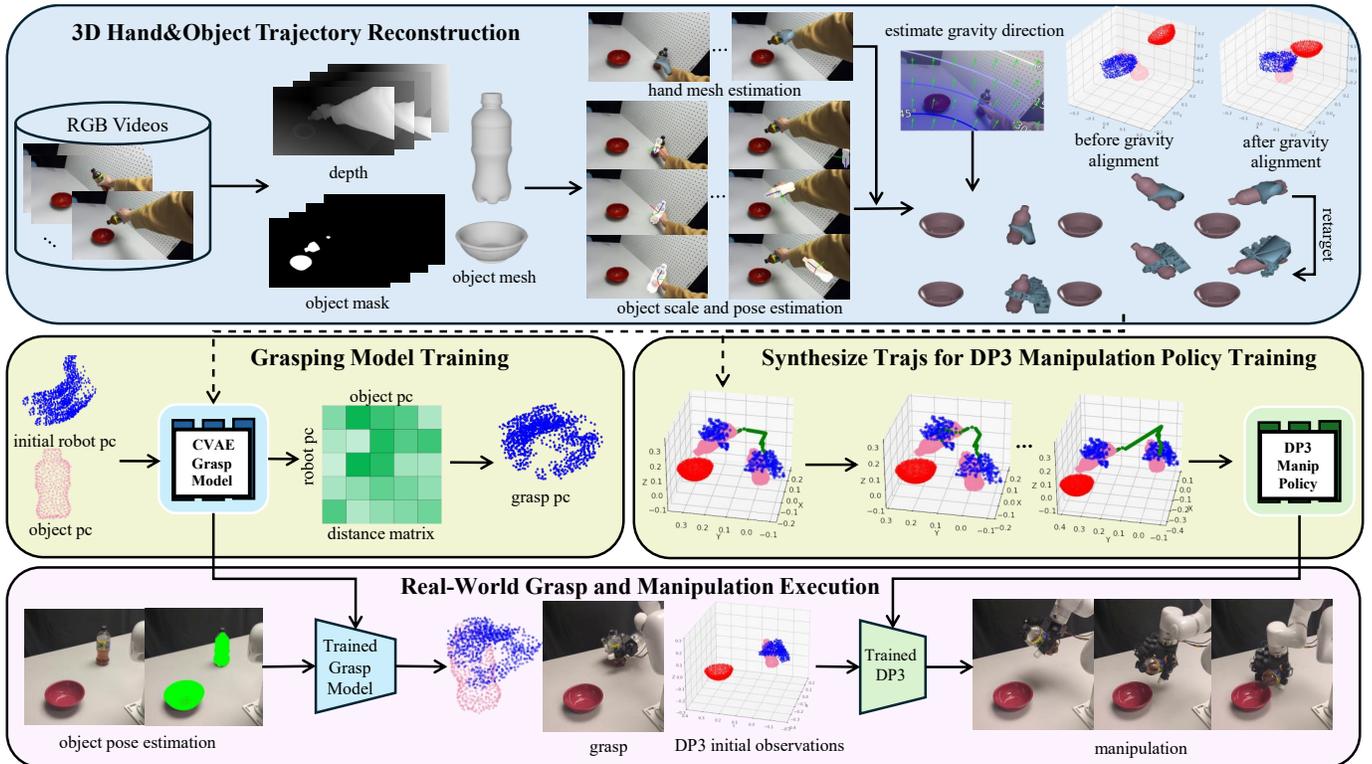


Fig. 2: **Overview of the VIDEOMANIP framework.** We first reconstruct 3D robot–object interaction trajectories from RGB human videos via recent advances in 3D vision (Sec. III-A). To utilize the reconstructed data for dexterous grasping and manipulation learning, we perform grasp contact optimization and interaction-centric grasp modeling, and synthesize trajectories for generalizable manipulation (Sec. III-B). Finally, we deploy the trained models for real-world dexterous grasping and manipulation.

hand trajectories, finger-level manipulations, and the resulting object pose changes; and (ii) the learning of grasping and manipulation policies from these reconstructed hand actions and object point clouds in Sec III-B.

Assumptions. (i) The raw dataset consists exclusively of egocentric human videos captured with a static camera, including both in-scene and in-the-wild recordings. Only *one video* is available for each task or object. (ii) No pre-scanned object meshes, object size information, depth measurements, or camera intrinsics are assumed to be available. (iii) No robot data or wearable and external sensing devices are used.

A. 3D Hand-Object Trajectory Reconstruction from Video

Given an RGB video of a human interacting with objects, $\mathcal{V} = \{I_1, \dots, I_T\}$, with image frames $I_t \in \mathbb{R}^{H \times W \times 3}$, we leverage existing 3D vision techniques to reconstruct the object mesh \mathcal{O} , estimate its 6D pose, and recover the human hand pose, followed by robot hand retargeting. To do so, we first adopt MoGe-2 [51] to estimate metric depth maps and camera intrinsics, which define a joint, metric 3D coordinate frame that enables physically consistent hand-object spatial alignment in all subsequent reconstructions.

Object Mesh Reconstruction and Pose Estimation. Our next step is to reconstruct object mesh \mathcal{O} and its 6D pose from the video. We reconstruct a complete 3D object mesh directly from RGB observations to enable accurate grasp contact modeling. We first identify manipulated objects in

the video and obtain object masks using Segment Anything Model 2 (SAM 2) [52]. We crop the masked regions to produce object-centric images and feed them into MeshyAI for image-to-mesh generation [53]. While the reconstructed object meshes \mathcal{O} capture the objects’ visual appearance and overall shape, they lack accurate metric scale. Pose estimation methods such as FoundationPose [54] assume access to object meshes with accurate real-world dimensions; violating this assumption leads to erroneous pose estimates.

To address the unknown object scale, we adopt a two-stage scale estimation strategy. First, we query the GPT-4.1 language model to obtain a coarse estimate of the object physical dimension and rescale the mesh accordingly. Second, we further refine the object scale by evaluating multiple candidate scalings (e.g., ranging from 0.5 \times to 2 \times) using FoundationPose, leveraging the previously estimated object mask and metric depth information. We select the scaling that minimizes the rendering error, enabling fine-grained scale verification. The rendering error is computed by comparing the rendered object mesh in each frame against the object mask obtained from SAM 2 in the video frames.

Human Hand Mesh Estimation and Robot Hand Retargeting. To extract robot 3D geometry and actions from video, we first recover human hand mesh and then retarget it to the robot hand. We use the hand mesh recovery model HaMeR [37] to reconstruct the hand mesh \mathcal{H} for all video frames. HaMeR represents the hand mesh using a low-dimensional parameterization $\mathbf{h} = (\theta, \beta)$, where θ denotes

the hand pose and β denotes the hand shape. HaMeR uses a weak-perspective camera model, which is inherently depth-ambiguous and sensitive to focal length errors. To align the reconstructed hand mesh \mathcal{H} within a joint coordinate space with the object, we reuse the metric depth maps predicted by MoGe-2, which are consistent with those used for object pose estimation, to ensure that \mathcal{H} and \mathcal{O} share the same depth reference. We then compute the corrected hand depth t'_z by averaging the metric depths at the 2D keypoints predicted by HaMeR. For every frame I_t , The reconstructed human hand poses (θ_t and β_t) are then retargeted to the robot hand configuration q_t , including the wrist pose and finger joint angles, via an optimization that minimizes the error between selected robot link keypoints and their corresponding human hand joints [49]. Given the robot URDF file, the robot mesh \mathcal{R} can be generated for any robot hand configuration q .

In-the-Wild Videos Calibration. Trajectories from in-the-wild videos, unlike those from in-scene videos that can be transformed to the robot base using a calibrated transform ${}^{\text{world}}T_{\text{cam}}$, lack this transformation. Their actions are expressed in the camera frame and are misaligned with the world frame due to unknown camera orientation and pose. As illustrated in Figure 2, before gravity alignment the reconstructed objects (e.g., the bowl and bottle) do not lie on the horizontal plane due to a tilted camera. Policies trained on such data would therefore learn camera-frame specific actions that do not transfer to the physical world frame. To address this issue, we apply GeoCalib [55], a single-image camera calibration method that leverages universal visual cues of 3D geometry to estimate camera orientation. From the first video frame, GeoCalib infers the gravity direction in the camera frame and the corresponding rotation ${}^{\text{grav}}R_{\text{cam}} \in \text{SO}(3)$ that aligns gravity with the negative z -axis, $\mathbf{g}_{\text{cam}} = [0, 0, -1]^\top$. We apply this rotation to all reconstructed meshes \mathcal{R} , \mathcal{O} , and robot configurations q , producing gravity-aligned trajectories and point clouds suitable for policy training. While this does not recover the full camera-to-world transformation ${}^{\text{world}}T_{\text{cam}}$, the estimated ${}^{\text{grav}}R_{\text{cam}}$ is sufficient to align trajectories from egocentric, in-the-wild videos to a shared reference plane with in-scene videos (e.g., the robot-mounted table), yielding trajectories that closely match those reconstructed from in-scene recordings.

B. Dexterous Grasp and Manipulation Learning

Here, we discuss how to exploit the reconstructed trajectories to teach robot dexterous manipulation. To handle visual differences in hand embodiment and background environments between robot and human RGB videos, we adopt point-cloud-based policies by sampling points from meshes’ surfaces. In practice, although the reconstructed trajectories capture the correct motions, the resulting robot–object interactions are not always physically feasible due to reconstruction errors in object geometry and inaccuracies in hand mesh scale, shape, or pose, which can lead to interpenetration or invalid contact (see Figure 5(b) for examples). During execution, such errors may cause the robot to miss the object or fail to grasp it securely. Moreover, although a single human video may include multiple grasp instances, it yields

only one manipulation trajectory, which is insufficient for learning robust manipulation policies.

To improve the physical validity of hand-object interactions and increase trajectory diversity from a single video, we perform differentiable hand–object contact optimization [47] and adopt DemoGen [20] to synthesize spatially randomized manipulation demonstrations from a single reconstructed human trajectory. Following DemoGen’s skill–motion decomposition, we split the reconstructed trajectory into two stages, grasping and manipulation, each learned by a dedicated policy. The grasp stage spans from the time the hand approaches sufficiently close to the object at time t_1 to when a stable grasp is achieved at time t_2 , before any object manipulation begins. The manipulation stage covers $[t_2, T]$, where the hand operates on the grasped object. The grasping policy is responsible for learning correct hand–object interactions, whereas the manipulation policy controls the object’s relative translation and rotation along with fine-grained finger motions.

Contact Optimization and Interaction-Centric Grasp Modeling. To address grasp reconstruction failures, during the grasping stage $[t_1, t_2]$, we refine the grasps using pre-trained ContactOpt [47]. The hand parameters \mathbf{h} determine the geometry of the reconstructed hand mesh \mathcal{H} relative to the object mesh \mathcal{O} . We therefore compute hand pose-dependent contact maps $C_{\mathcal{H}}(\mathbf{h})$ and $C_{\mathcal{O}}(\mathbf{h})$, which assign smooth, distance-based contact values to each vertex based on its proximity to the other mesh. For example, the contact value at an object vertex $v_{\mathcal{O}}^i$ can be expressed as

$$C_{\mathcal{O}}(v_{\mathcal{O}}^i; \mathbf{h}) = \max\left(0, 1 - \frac{\min_j \|v_{\mathcal{H}}^j(\mathbf{h}) - v_{\mathcal{O}}^i\|}{c_{\text{rad}}}\right),$$

where $v_{\mathcal{H}}^j(\mathbf{h})$ are the hand vertices at pose \mathbf{h} and c_{rad} is a radius parameter controlling the falloff. Similarly, $C_{\mathcal{H}}(\mathbf{h})$ is computed for the hand vertices. These maps are differentiable and provide gradients for optimizing the hand pose. ContactOpt then predicts desirable contact regions on the hand mesh $\hat{C}_{\mathcal{H}}$ and object mesh $\hat{C}_{\mathcal{O}}$, and adjusts the hand pose parameters \mathbf{h} to align the current contacts with these targets using the differentiable objective (for the full formulation, please refer to [47]):

$$E(\mathbf{h}) = |C_{\mathcal{O}}(\mathbf{h}) - \hat{C}_{\mathcal{O}}| + |C_{\mathcal{H}}(\mathbf{h}) - \hat{C}_{\mathcal{H}}|.$$

The optimized human hand poses are then retargeted to robot hand configurations $q_t^{\text{opt}}, t \in [t_1, t_2]$, which serve as grasp demonstrations for training the grasping model.

Given the robot mesh \mathcal{R} from q_t^{opt} and the object mesh \mathcal{O} , we then employ the DRO model [56] to capture their interaction for robust grasp modeling. DRO predicts dense point-to-point distances between the robot hand point cloud $\mathbf{P}^{\mathcal{R}} \in \mathbb{R}^{N_{\mathcal{R}} \times 3}$ and the object point cloud $\mathbf{P}^{\mathcal{O}} \in \mathbb{R}^{N_{\mathcal{O}} \times 3}$, effectively capturing the interaction patterns and spatial relationships provided by ContactOpt. Specifically, DRO takes randomly initialized robot hand point clouds $\mathbf{P}_{\text{init}}^{\mathcal{R}}$ and zero-centered object point clouds $\mathbf{P}^{\mathcal{O}}$ sampled from the meshes as input, and predicts the distance matrix $\mathcal{D}(\mathcal{R}, \mathcal{O})^{\text{Pred}} \in \mathbb{R}^{N_{\mathcal{R}} \times N_{\mathcal{O}}}$, as illustrated in Figure 2. The training loss is the difference between the predicted and ground-truth distance matrices with

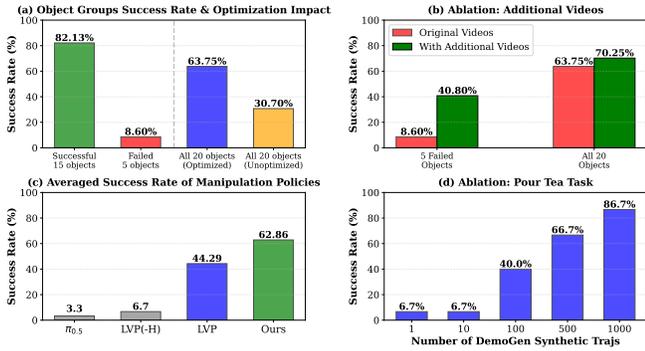


Fig. 3: **Quantitative Results on Grasping and Manipulation.** **Grasping:** (a) Success rates across object groups, with comparison between models trained with and without grasp optimization; (b) Ablation study on incorporating additional videos for previously failed objects. **Manipulation:** (c) Performance comparison between our method and baselines across seven manipulation tasks; (d) Ablation study on the number of DemoGen-synthesized trajectories.

$\mathcal{L}_{L1} \left(D(\mathcal{R}, \mathcal{O})^{\text{Pred}}, D(\mathcal{R}, \mathcal{O})^{\text{GT}} \right)$. Using $D(\mathcal{R}, \mathcal{O})^{\text{Pred}}$ and the object point clouds $\mathbf{P}^{\mathcal{O}}$, we can position the robot point cloud in the target grasp pose $\mathbf{P}^{\mathcal{R}_{\text{grasp}}}$ with a multilateration method [57] and compute the grasp configuration q^{grasp} through optimization, expressed relative to the object.

Manipulation Demonstration Synthesis and Training. Generalizable manipulation policies typically require a large number of demonstrations, whereas our goal is to learn from a single reconstructed human video to avoid repetitive human video collection and reconstruction. To this end, we adopt DemoGen [20] to synthesize additional manipulation demonstrations via objects spatial randomization. DemoGen enforces spatial equivariance by applying consistent SE(3) transformations to both the object point clouds and the associated robot trajectories. Since 3D point clouds and robot proprioceptive states can be transformed directly, this enables the synthesis of diverse trajectories while preserving hand–object contact and fine-grained finger motions. For manipulation policy learning, we adopt the 3D diffusion-based policy (DP3) [58], which takes as input the robot hand point clouds and proprioceptive state at the grasp pose, $\mathbf{P}^{\mathcal{R}_{\text{grasp}}}$ and q^{grasp} , along with the object point clouds $\mathbf{P}^{\mathcal{O}}$ as the initial observation. The policy outputs actions as predicted changes in the robot configuration, Δq , and is executed in a closed-loop manner using updated observations.

IV. EXPERIMENTS

We conduct two sets of experiments with our framework: grasping of objects and manipulating different objects. We present the main results in Fig. 3. Video rollouts for the results are best viewed in our website.

A. Grasping Experiments

Setup. We collect 20 videos of humans grasping objects, with one video per object; the object categories are shown in Fig. 5 (a). Each 30 fps video lasts 3 to 6 seconds. We train our models across all objects and evaluate grasp success in the IsaacGym simulator using an 18-Dof Inspire robot hand, as its size closely matches that of a human hand.

We initialize the simulator with object meshes and poses reconstructed from the recorded videos. Each grasp $q^{\text{grasp}} \in \mathbb{R}^{18}$ undergoes a 300-step disturbance phase in IsaacGym, with forces applied sequentially from six directions: $\pm x$, $\pm y$, and $\pm z$. The applied force magnitude is set to 0.5 times the object’s mass. We consider a grasp successful if the object’s displacement under applied force disturbances stays within 3 cm of its initial position. Through these experiments, we aim to answer the following questions:

- **Q1:** How well does the grasping model generalize across object categories when trained on grasps reconstructed from human RGB videos?
- **Q2:** For objects that are initially challenging to grasp, does incorporating additional videos captured from diverse viewpoints improve performance?

A1: Our grasping model, trained on video-reconstructed grasp data, achieves an average success rate of 82.13% over 15 successfully grasped objects and 63.75% across all objects, including five failure cases, as shown in Fig. 4. The predicted grasps closely mimic the human grasps observed in the input videos; for example, on *Soap Dispenser* and *Spray Bottle*, the middle finger is placed on top of the dispenser and the trigger, respectively. We further ablate the effect of grasp optimization using ContactOpt [47] by retraining the DRO grasping model on all objects with unoptimized grasps, which yields only a 30.7% average success rate across 20 objects, compared to 63.75% when using optimized grasps, as shown in Fig. 3(a). As shown in Fig. 5(b), unoptimized grasps often penetrate object surfaces or fail to establish contact due to reconstruction errors, whereas optimized grasps achieve more accurate hand–object contact, which is reflected in the higher grasp success rate.

A2: However, five objects (highlighted by the red dotted box in Fig. 4) exhibit lower success rates. These failures primarily stem from object pose estimation errors, for example in the *Pan* case where heavy occlusion of the handle by the human hand leads to inaccurate pose estimation, and from a lack of force awareness, in which grasps for the *Glasses Case* and *Hand Bag* are kinematically plausible but fail to achieve force closure, resulting in instability. To mitigate these issues, we examine whether incorporating grasping videos with diverse viewpoints and grasp styles improves performance. The underlying intuition is that reconstruction may fail in some videos due to occlusion or suboptimal grasps, but succeed in others with reduced occlusion or more stable force-closure configurations. Accordingly, for each failed object, we collect two additional human videos and augment the dataset, resulting in 30 videos (20 + 5×2) for retraining. As shown in Tab. I, training with the augmented dataset improves the success rate on the five previously failed objects from 8.6% to 40.8%, and increases the overall success rate from 63.75% to 70.25%. Although the *Hand Bag* does not improve due to its grasp being sensitive to disturbances, the performance gains on the remaining objects indicate that increased video diversity alone can substantially enhance model robustness and generalization.

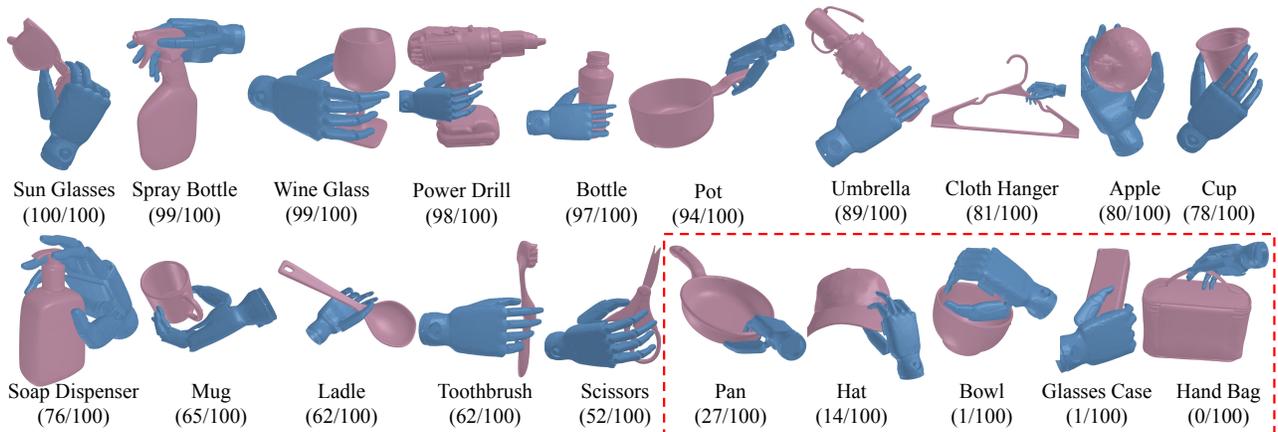
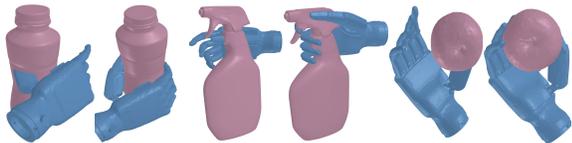


Fig. 4: **Predicted grasps and success rates in IsaacGym.** The DRO grasping model is trained on 20 object categories. Each object is evaluated over 100 trials and sorted by descending success rate; red dotted box denotes failed grasps.



(a) 20 Test Objects



(b) Grasp Optimization Results

Fig. 5: (a) Visualization of the 20 objects used for video collection and grasp model training. (b) Grasp optimization results on *Bottle*, *Spray Bottle*, and *Apple*. Left: unoptimized reconstructed grasps; right: grasps optimized using ContactOpt [47].

TABLE I: **Quantitative ablation of grasp performance with multiple videos.** Success rates on the five initially failed objects in Fig. 4, comparing models trained with the original 20 videos and with additional videos for the failed objects.

	Pan	Hat	Bowl	Glasses Case	Hand bag
Single Video	27	14	1	1	0
Multi Videos	36	52	48	68	0

B. Manipulation Experiments

Setup. We evaluate seven manipulation tasks on a four-finger LEAP Hand mounted on a 7-DoF xArm, using one human demonstration video per task. Three tasks use in-the-scene videos: *Pour Tea* (pick up a bottle and tilt to pour into a bowl), *Close Drawer* (push a drawer fully into a shelf), and *Pick&Place Can* (pick up a can and place it into a box). The remaining four tasks use in-the-wild videos with gravity-based calibration: *Pour Tea* without camera-robot calibration, *Hang Hat* (hang a hat held in the hand onto a rack), *Move Jenga Box* (pick a Jenga box and place it onto a shelf), and *Screw Bulb* (screw a bulb into a socket until it's

TABLE II: **Quantitative comparison of our method and baselines across different video sources and tasks over 10 trials with randomized target object locations.** Our DP3 model is trained on 1000 trajectories synthesized by DemoGen from a single reconstructed source trajectory.

Video Type	Task	$\pi_{0.5}$	LVP(-H)	LVP	Ours
In-scene Videos	Pour Tea	0/10	1/10	7/10	8/10
	Close Drawer	1/10	2/10	6/10	9/10
	Pick&Place Can	-	0/10	4/10	5/10
In-the-wild Videos	Pour Tea	0/10	1/10	7/10	7/10
	Hang Hat	0/10	-	4/10	6/10
	Screw Bulb	1/10	0/10	1/10	4/10
	Move Jenga Box	0/10	0/10	2/10	5/10

on). For *Pick&Place Can* task, the DP3 action Δq_t and state q_t spaces are 22-DoF, comprising a 6-DoF delta wrist pose and 16-DoF delta finger joints to enable object release. For all other tasks, we use a 6-DoF delta wrist action space. We train a single DRO model across all grasp objects, including the *Bottle*, *Can*, *Drawer*, *Hat*, *Jenga Box*, and *Bulb*, and train a separate DP3 manipulation policy for each task. The second object in each task (e.g., *Bowl*, *Box*, *Rack*) is referred to as the target object.

During evaluation, we estimate object poses using FoundationPose [54] with the reconstructed object mesh \mathcal{O} under a calibrated ZED camera. The trained DRO model then predicts a grasp pose from the object point cloud, and the LEAP Hand executes the grasp $q^{\text{grasp}} \in \mathbb{R}^{22}$ in free space. The initial observation for the subsequent manipulation stage comprises the point clouds of the grasped and target object $\mathbf{P}^{\mathcal{O}}$, the robot hand point cloud at the predicted grasp pose $\mathbf{P}_{\text{grasp}}^{\mathcal{R}}$, and the corresponding proprioceptive state q^{grasp} . DP3 is then rolled out in a closed-loop manner based on the current observation. At each timestep, the predicted action Δq_t updates the robot configuration and the associated hand point clouds. During execution, the grasped object point cloud $\mathbf{P}^{\mathcal{O}}$ is often occluded by the large LEAP Hand, making closed-loop DP3 control difficult. We therefore assume that the hand-object relative pose established at grasp remains fixed during rollout, and update the object point cloud using the hand-to-object transformation computed from q^{grasp} . This

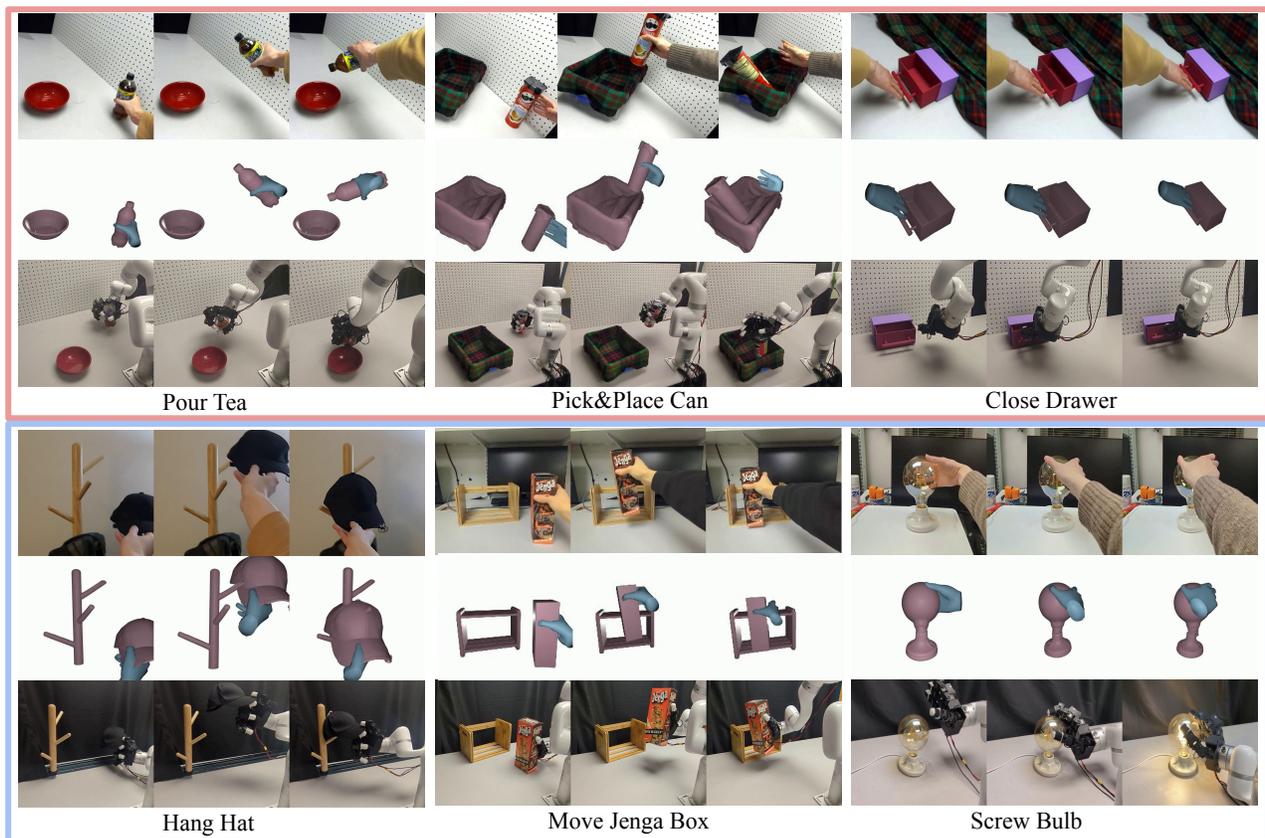


Fig. 6: Visualization of VIDEOMANIP Execution Using In-Scene (top) and In-the-Wild (bottom) Video Data Sources. For each task, given RGB human videos (row 1), we reconstruct 3D trajectories of the human hand and objects (row 2). Trained with these trajectories, we executed on a real-world LEAP Hand (row 3).

approximation suffices for most tasks, while object release in *Pick&Place Can* is handled implicitly during execution rather than explicitly modeled in the point-cloud update.

Since our policies are trained solely from RGB human videos without any robot demonstrations, there are few directly comparable baselines. We therefore consider two representative alternatives: (1) $\pi_{0.5}$, a generalizable vision-language-action model [59] used as a baseline for zero-shot manipulation; and (2) Large Video Planner (LVP), a retargeting-based approach that maps predicted human hand motions from its video generator to robot actions [35]. We fine-tune $\pi_{0.5}$ -DROID using 200 robot demonstrations in our scenes and LEAP Hand on tasks disjoint from the seven tasks. Since $\pi_{0.5}$ -DROID is trained with parallel-jaw grippers, we initialize each rollout with the object already grasped by the LEAP Hand and evaluate only the subsequent manipulation phase. For LVP, we evaluate two variants: LVP, conditioned on an initial observation containing a human hand to predict subsequent hand grasping and manipulation motion, and LVP(-H), conditioned on an initial observation containing only the objects, and is used to evaluate the model’s ability to infer manipulation purely from object observations without human hand cues. Through the experiments, we aim to answer the following questions:

- **Q3:** How well does the manipulation policy perform when trained from a single in-scene and in-the-wild RGB human video?

- **Q4:** How do in-the-wild calibration and the number of synthesized demonstrations affect policy performance?

A3: VIDEOMANIP captures hand-object motion from human videos, providing spatiotemporal supervision for manipulation policy learning. The reconstructed trajectories closely align with the corresponding RGB human videos, as shown in Fig. 6. Leveraging both these reconstructed trajectories and additional synthesized trajectories generated by DemoGen, the trained DP3 policy achieves the highest overall success rate of 62.86% across all tasks (Tab. II). The synthesized trajectories enable our policy to generalize towards different target object locations (videos available on our website). While, DemoGen based trajectory augmentation is less helpful for tasks like *Close Drawer* and *Screw Bulb*, where objects (*Bulb* and *Socket*) are tightly coupled, the variations in object locations are still helpful for generalization of the trained policy.

Besides, our model outperforms the LVP baselines, which fail to consistently generate feasible grasps for the *Jenga Box* and exhibit motion failures in tasks such as *Move Jenga Box* (object not placed on the shelf) and *Screw Bulb* (missing required rotational motion). Moreover, when the initial observation lacks a human right hand, a common scenario in robot manipulation, LVP(-H) frequently generates erroneous future frames even when explicitly instructed to use the right hand. These failures include hallucinating a robot gripper instead of a human hand, failing to execute

the intended task (e.g., remaining static instead of push the drawer), attempting infeasible grasps, confusing left and right hands, or producing undetectable hands. In contrast, our approach reconstructs explicit 3D geometry, grounding supervision in physically meaningful meshes rather than pixel-level video hallucinations. Together with synthesized trajectories, our results demonstrate that reconstruction-based trajectories provide reliable and generalizable supervision for dexterous manipulation policy learning. Finally, $\pi_{0.5}$ achieves limited success because it is trained primarily on parallel-gripper datasets, and the small scale of our LEAP Hand fine-tuning demonstrations is insufficient to meaningfully improve its dexterous manipulation performance.

A4: The *Pour Tea* task is evaluated on both in-scene and in-the-wild video categories using ${}^{\text{world}}T_{\text{cam}}$ and ${}^{\text{gra}}R_{\text{cam}}$, respectively. Table II shows no significant performance difference between the two video sources; however, without applying ${}^{\text{gra}}R_{\text{cam}}$ for the in-the-wild *Pour Tea* task, the success rate drops to 0%, highlighting the importance of gravity-alignment calibration. We further ablate the effect of DemoGen on the in-scene *Pour Tea* task by adding x - y plane perturbations within $[-0.2, 0.2]$ m to the target object *Bowl* position reconstructed from the source video, and synthesizing between 1 and 1000 trajectories. During evaluation, the *Bowl* is placed at 15 different locations within this perturbed region. As the number of synthesized trajectories increases, the success rate improves steadily from 1/15 to 13/15 (Fig. 3(d)), demonstrating that DemoGen enables DP3 to learn actions that generalize across target object locations.

V. CONCLUSION, LIMITATIONS, AND FUTURE WORK

We present a reconstruction-based framework that recovers 3D hand-object trajectories from RGB-only, in-the-wild human videos, without relying on specialized sensing devices. The reconstructed trajectories are shown to be effective for training dexterous grasping and manipulation policies, and provide more physically grounded supervision compared to retargeting-based approaches that rely on generated videos.

Despite these advantages, VIDEOMANIP has several limitations. (1) VIDEOMANIP framework relies on multiple 3D vision models, including image-to-mesh reconstruction and object pose estimation, which may introduce compounding errors across stages. We currently circumvent this issue by recording the human videos from approximately egocentric viewpoints that facilitate reliable hand and object reconstruction, instead of downloading generic videos from the internet. In future work, we plan to incorporate a trajectory verification module to automatically detect and filter erroneous reconstructions, and scale to more in-the-wild videos. (2) The current framework assumes static or approximately static camera setups. We plan to extend our approach to dynamic camera settings by leveraging recent advances in dynamic scene reconstruction [60], [61]. (3) The manipulation policy learning currently relies on synthesized trajectories from DemoGen represented as 3D point clouds. In real-world manipulation, however, object point clouds are often difficult to track due to occlusions caused by the robot hand. We currently assume that the hand-object relative contact es-

tablished during grasping is preserved throughout execution to update the object point clouds along with robot hand point clouds. In future work, we plan to explore alternative trajectory synthesis modalities beyond point clouds, such as image-based representations.

REFERENCES

- [1] K. Liu, Z. Jia, Y. Li, P. Chen, S. Liu, X. Liu, P. Zhang, H. Song, X. Ye, N. Cao *et al.*, “Fastumi-100k: Advancing data-driven robotic manipulation with a large-scale umi-style dataset,” *arXiv preprint arXiv:2510.08022*, 2025.
- [2] I. Guzey, H. Qi, J. Urain, C. Wang, J. Yin, K. Bodduluri, M. Lambeta, L. Pinto, A. Rai, J. Malik *et al.*, “Dexterity from smart lenses: Multi-fingered robot manipulation with in-the-wild human demonstrations,” *arXiv preprint arXiv:2511.16661*, 2025.
- [3] Z. Jiang, Y. Xie, K. Lin, Z. Xu, W. Wan, A. Mandlkar, L. J. Fan, and Y. Zhu, “Dexmimicgen: Automated data generation for bimanual dexterous manipulation via imitation learning,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 16 923–16 930.
- [4] Y. Niu, Y. Zhang, M. Yu, C. Lin, C. Li, Y. Wang, Y. Yang, W. Yu, T. Zhang, Z. Li, J. Francis, B. Chen, J. Tan, and D. Zhao, “Human2locoman: Learning versatile quadrupedal manipulation with human pretraining,” in *Robotics: Science and Systems (RSS)*, 2025.
- [5] R. Fu, D. Zhang, A. Jiang, W. Fu, A. Funk, D. Ritchie, and S. Sridhar, “Gigahands: A massive annotated dataset of bimanual hand activities,” in *CVPR*, 2025, pp. 17 461–17 474.
- [6] C. Bao, J. Xu, X. Wang, A. Gupta, and H. Bharadhwaj, “Handsonvlm: Vision-language models for hand-object interaction prediction,” *arXiv preprint arXiv:2412.13187*, 2024.
- [7] H. Bharadhwaj, R. Mottaghi, A. Gupta, and S. Tulsiani, “Track2act: Predicting point tracks from internet videos enables diverse zero-shot robot manipulation,” *CoRR*, 2024.
- [8] K. Shaw, S. Bahl, A. Sivakumar, A. Kannan, and D. Pathak, “Learning dexterity from human hand motion in internet videos,” *IJRR*, 2022.
- [9] H. G. Singh, A. Loquercio, C. Sferrazza, J. Wu, H. Qi, P. Abbeel, and J. Malik, “Hand-object interaction pretraining from videos,” in *ICRA*, 2025.
- [10] T. Tao, M. K. Srirama, J. J. Liu, K. Shaw, and D. Pathak, “Dexwild: Dexterous human interactions for in-the-wild robot policies,” *Robotics: Science and Systems (RSS)*, 2025.
- [11] T. G. W. Lum, O. Y. Lee, C. K. Liu, and J. Bohg, “Crossing the human-robot embodiment gap with sim-to-real rl using one human demonstration,” *arXiv preprint arXiv:2504.12609*, 2025.
- [12] Y. Qin, Y.-H. Wu, S. Liu, H. Jiang, R. Yang, Y. Fu, and X. Wang, “Dexmv: Imitation learning for dexterous manipulation from human videos,” in *European Conference on Computer Vision*. Springer, 2022, pp. 570–587.
- [13] J. Wu, G. Pavlakos, G. Gkioxari, and J. Malik, “Reconstructing hand-held objects in 3d,” *arXiv preprint arXiv:2404.06507*, 2024.
- [14] H. Chen, Y. Yao, Y. Ye, Z. Xu, H. Bharadhwaj, J. Wang, S. Tulsiani, Z. Erickson, and J. Ichnowski, “Web2grasp: Learning functional grasps from web images of hand-object interactions,” *arXiv preprint arXiv:2505.05517*, 2025.
- [15] J. Yu, L. Fu, H. Huang, K. El-Refai, R. A. Ambrus, R. Cheng, M. Z. Irshad, and K. Goldberg, “Real2render2real: Scaling robot data without dynamics simulation or robot hardware,” *arXiv preprint arXiv:2505.09601*, 2025.
- [16] W. Ye, F. Liu, Z. Ding, Y. Gao, O. Rybkin, and P. Abbeel, “Video2policy: Scaling up manipulation tasks in simulation through internet videos,” *arXiv preprint arXiv:2502.09886*, 2025.
- [17] C. Chi, Z. Xu, C. Pan, E. Cousineau, B. Burchfiel, S. Feng, R. Tedrake, and S. Song, “Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots,” *arXiv preprint arXiv:2402.10329*, 2024.
- [18] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, “Learning fine-grained bimanual manipulation with low-cost hardware,” *RSS*, 2023.
- [19] R.-Z. Qiu, S. Yang, X. Cheng, C. Chawla, J. Li, T. He, G. Yan, D. J. Yoon, R. Hoque, L. Paulsen *et al.*, “Humanoid policy” human policy,” *arXiv preprint arXiv:2503.13441*, 2025.
- [20] Z. Xue, S. Deng, Z. Chen, Y. Wang, Z. Yuan, and H. Xu, “Demogen: Synthetic demonstration generation for data-efficient visuomotor policy learning,” *arXiv preprint arXiv:2502.16932*, 2025.

- [21] S. Nair, A. Rajeswaran, V. Kumar, C. Finn, and A. Gupta, "R3m: A universal visual representation for robot manipulation," in *6th Annual Conference on Robot Learning*.
- [22] Y. J. Ma, S. Sodhani, D. Jayaraman, O. Bastani, V. Kumar, and A. Zhang, "Vip: Towards universal visual reward and representation via value-implicit pre-training," in *The Eleventh ICLR*.
- [23] J. Shi, Z. Zhao, T. Wang, I. Pedroza, A. Luo, J. Wang, J. Ma, and D. Jayaraman, "Zeromimic: Distilling robotic manipulation skills from web videos," *arXiv preprint arXiv:2503.23877*, 2025.
- [24] A. Agarwal, S. Uppal, K. Shaw, and D. Pathak, "Dexterous functional grasping," in *7th Annual Conference on Robot Learning*.
- [25] A. Goyal, A. Mousavian, C. Paxton, Y.-W. Chao, B. Okorn, J. Deng, and D. Fox, "Ifor: Iterative flow minimization for robotic object rearrangement," in *CVPR*, 2022, pp. 14 787–14 797.
- [26] C. Wang, L. Fan, J. Sun, R. Zhang, L. Fei-Fei, D. Xu, Y. Zhu, and A. Anandkumar, "Mimicplay: Long-horizon imitation learning by watching human play," in *7th Annual Conference on Robot Learning*.
- [27] H. Li, L. Sun, Y. Hu, D. Ta, J. Barry, G. Konidaris, and J. Fu, "Novaflo: Zero-shot manipulation via actionable flow from generated videos," *arXiv preprint arXiv:2510.08568*, 2025.
- [28] V. Jain, M. Attarian, N. J. Joshi, A. Wahid, D. Driess, Q. Vuong, P. R. Sanketi, P. Sermanet, S. Welker, C. Chan *et al.*, "Vid2robot: End-to-end video-conditioned policy learning with cross-attention transformers," *arXiv preprint arXiv:2403.12943*, 2024.
- [29] C. Tang, A. Xiao, Y. Deng, T. Hu, W. Dong, H. Zhang, D. Hsu, and H. Zhang, "Mimicfunc: Imitating tool manipulation from a single human video via functional correspondence," *arXiv preprint arXiv:2508.13534*, 2025.
- [30] C. Wang, H. Shi, W. Wang, R. Zhang, L. Fei-Fei, and C. K. Liu, "Dexcap: Scalable and portable mocap data collection system for dexterous manipulation," *arXiv preprint arXiv:2403.07788*, 2024.
- [31] R. G. Goswami, A. Bar, D. Fan, T.-Y. Yang, G. Zhou, P. Krishna-murthy, M. Rabbat, F. Khorrani, and Y. LeCun, "World models can leverage human videos for dexterous manipulation," *arXiv preprint arXiv:2512.13644*, 2025.
- [32] W. Huang, Y.-W. Chao, A. Mousavian, M.-Y. Liu, D. Fox, K. Mo, and L. Fei-Fei, "Pointworld: Scaling 3d world models for in-the-wild robotic manipulation," *arXiv preprint arXiv:2601.03782*, 2026.
- [33] J. Liang, P. Tokmakov, R. Liu, S. Sudhakar, P. Shah, R. Ambrus, and C. Vondrick, "Video generators are robot policies," *arXiv preprint arXiv:2508.00795*, 2025.
- [34] S. Patel, S. Mohan, H. Mai, U. Jain, S. Lazebnik, and Y. Li, "Robotic manipulation by imitating generated videos without physical demonstrations," *arXiv preprint arXiv:2507.00990*, 2025.
- [35] B. Chen, T. Zhang, H. Geng, K. Song, C. Zhang, P. Li, W. T. Freeman, J. Malik, P. Abbeel, R. Tedrake *et al.*, "Large video planner enables generalizable robot control," *arXiv preprint arXiv:2512.15840*, 2025.
- [36] Y. Rong, T. Shiratori, and H. Joo, "Frankmocap: Fast monocular 3d hand and body motion capture by regression and integration," *arXiv preprint arXiv:2008.08324*, 2020.
- [37] G. Pavlakos, D. Shan, I. Radosavovic, A. Kanazawa, D. Fouhey, and J. Malik, "Reconstructing hands in 3d with transformers," in *CVPR*, 2024, pp. 9826–9836.
- [38] G. Gkioxari, J. Malik, and J. Johnson, "Mesh r-cnn," in *ICCV*, 2019.
- [39] C.-H. Lin, C. Kong, and S. Lucey, "Learning efficient point cloud generation for dense 3d object reconstruction," in *AAAI*, 2018.
- [40] B. Deng, K. Genova, S. Yazdani, S. Bouaziz, G. Hinton, and A. Tagliasacchi, "Cvxnet: Learnable convex decomposition," in *CVPR*, 2020.
- [41] Y. Hasson, G. Varol, D. Tzionas, I. Kalevatykh, M. J. Black, I. Laptev, and C. Schmid, "Learning joint reconstruction of hands and manipulated objects," in *CVPR*, 2019, pp. 11 807–11 816.
- [42] Y. Liu, X. Long, Z. Yang, Y. Liu, M. Habermann, C. Theobalt, Y. Ma, and W. Wang, "Easyhoi: Unleashing the power of large models for reconstructing hand-object interactions in the wild," *arXiv preprint arXiv:2411.14280*, 2024.
- [43] Y. Ye, P. Hebbbar, A. Gupta, and S. Tulsiani, "Diffusion-guided reconstruction of everyday hand-object interaction clips," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 19 717–19 728.
- [44] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "DeepSDF: Learning continuous signed distance functions for shape representation," in *CVPR*, 2019.
- [45] Z. Fan, M. Parelli, M. E. Kadoglou, X. Chen, M. Kocabas, M. J. Black, and O. Hilliges, "Hold: Category-agnostic 3d reconstruction of interacting hands and objects from video," in *CVPR*, 2024, pp. 494–504.
- [46] Z. Cao, I. Radosavovic, A. Kanazawa, and J. Malik, "Reconstructing hand-object interactions in the wild," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 12 417–12 426.
- [47] P. Grady, C. Tang, C. D. Twigg, M. Vo, S. Brahmhatt, and C. C. Kemp, "Contactopt: Optimizing contact to improve grasps," in *CVPR*, 2021, pp. 1471–1481.
- [48] K. Li, P. Li, T. Liu, Y. Li, and S. Huang, "Maniptrans: Efficient dexterous bimanual manipulation transfer via residual learning," in *CVPR*, 2025, pp. 6991–7003.
- [49] Y. Qin, W. Yang, B. Huang, K. Van Wyk, H. Su, X. Wang, Y.-W. Chao, and D. Fox, "Anyteleop: A general vision-based dexterous robot arm-hand teleoperation system," in *Robotics: Science and Systems*, 2023.
- [50] C. Pan, C. Wang, H. Qi, Z. Liu, H. Bharadhwaj, A. Sharma, T. Wu, G. Shi, J. Malik, and F. Hogan, "Spider: Scalable physics-informed dexterous retargeting," *arXiv preprint arXiv:2511.09484*, 2025.
- [51] R. Wang, S. Xu, Y. Dong, Y. Deng, J. Xiang, Z. Lv, G. Sun, X. Tong, and J. Yang, "Moge-2: Accurate monocular geometry with metric scale and sharp details," *arXiv preprint arXiv:2507.02546*, 2025.
- [52] N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. Ryal, T. Ma, H. Khedr, R. Rädle, C. Rolland, L. Gustafson *et al.*, "Sam 2: Segment anything in images and videos," in *The Thirteenth ICLR*.
- [53] Meshy AI, "Meshy ai: The #1 ai 3d model generator for creators," 2025, accessed: 2025-04-17. [Online]. Available: <https://www.meshy.ai/>
- [54] B. Wen, W. Yang, J. Kautz, and S. Birchfield, "Foundationpose: Unified 6d pose estimation and tracking of novel objects," in *CVPR*, 2024, pp. 17 868–17 879.
- [55] A. Veicht, P.-E. Sarlin, P. Lindenberger, and M. Pollefeys, "Geocalib: Learning single-image calibration with geometric optimization," in *European Conference on Computer Vision*. Springer, 2024, pp. 1–20.
- [56] Z. Wei, Z. Xu, J. Guo, Y. Hou, C. Gao, Z. Cai, J. Luo, and L. Shao, "D (r, o) grasp: A unified representation of robot and object interaction for cross-embodiment dexterous grasping," in *2025 ICRA*. IEEE, 2025, pp. 4982–4988.
- [57] A. Norrdine, "An algebraic solution to the multilateration problem," in *Proceedings of the 15th international conference on indoor positioning and indoor navigation, Sydney, Australia*, vol. 1315, 2012.
- [58] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu, "3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations," in *RSS*, 2024.
- [59] P. Intelligence, K. Black, N. Brown, J. Darpinian, K. Dhabalia, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai *et al.*, "pi.0.5: a vision-language-action model with open-world generalization," *arXiv preprint arXiv:2504.16054*, 2025.
- [60] S. Wang, Z. Jiang, X. Yang, and X. Wang, "C4d: 4d made from 3d through dual correspondences," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025, pp. 7570–7580.
- [61] Z. Yu, S. Zafeiriou, and T. Birdal, "Dyn-hamr: Recovering 4d interacting hand motion from a dynamic camera," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 27 716–27 726.