

# DITTO: Demonstration Imitation by Trajectory Transformation

Nick Heppert, Max Argus, Tim Welschehold, Thomas Brox, Abhinav Valada

**Abstract**—Teaching robots new skills quickly and conveniently is crucial for the broader adoption of robotic systems. In this work, we address the problem of one-shot imitation from a single human demonstration, given by an RGB-D video recording. We propose a two-stage process. In the first stage we extract the demonstration trajectory offline. This entails segmenting manipulated objects and determining their relative motion in relation to secondary objects such as containers. In the online trajectory generation stage, we first re-detect all objects, then warp the demonstration trajectory to the current scene and execute it on the robot. To complete these steps, our method leverages several ancillary models, including those for segmentation, relative object pose estimation, and grasp prediction. We systematically evaluate different combinations of correspondence and re-detection methods to validate our design decision across a diverse range of tasks. Specifically, we collect and quantitatively test on demonstrations of ten different tasks including pick-and-place tasks as well as articulated object manipulation. Finally, we perform extensive evaluations on a real robot system to demonstrate the effectiveness and utility of our approach in real-world scenarios. We make the code publicly available at <http://ditto.cs.uni-freiburg.de>.

## I. INTRODUCTION

Humans are remarkably good at learning new motion skills from just a few, or even a single demonstration, given by other humans. Similarly, a common paradigm to teach a robot is imitation learning [1]. Here, a human actively demonstrates a skill to a robot either directly on the robot by teleoperating it [2], by kinesthetic teaching [3] or alternatively, by performing the task themselves [4]. Teleoperating a robot to collect demonstrations is possible with various different input devices [2], [5]. Despite this variety, collecting robot demonstrations remains difficult as these devices typically do not share the morphology of the robot and thus, leads to a significant training phase for the human operator or requires an expert teacher. To circumvent these problems, kinesthetic teaching is an appealing alternative. While this reduces the training time for the operator, it is not always beneficial as the operator has to be present in the scene, which can introduce various challenges such as occlusion or restricting the robot’s workspace due to safety constraints.

In contrast to these on-robot approaches, in this paper, we propose a novel way to teach robots new tasks by letting them passively observe a human performing a task only once. We move away from end-effector action representations and

Department of Computer Science, University of Freiburg, Germany.  
This work was funded by the Carl Zeiss Foundation with the ReScaLe project and the German Research Foundation (DFG): 417962828, 401269959.  
The authors would like to thank Martin Büchner for helping creating the figures and Vilja Lott for data sanity checking.

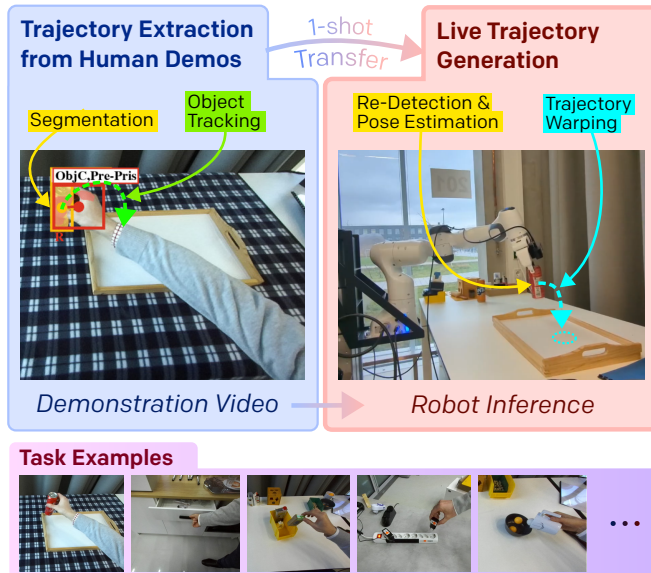


Fig. 1: Human demonstration of manipulation actions are transferred to new scenes so that a robot can replicate the manipulation action. For this, we use a two-stage process, first extracting object trajectories by segmenting and tracking objects. Then, we transfer the trajectory to a new scene by re-detection and trajectory transformation according to the re-detected positions. The proposed method is then evaluated on several different tasks.

towards an object-pose-centric perspective [6] in which we represent the trajectory as a sequence of poses of the object. This allows us to collect demonstration data independent of the embodiment, and we only later perform the transfer from human to robot.

Learning from human demonstrations is preferable over learning from robot demonstrations in many settings as human demonstrations can be collected most naturally and in principle be performed by non-expert users. Thus, datasets of humans performing tasks are much easier to collect and more diverse than datasets of robot demonstrations, making them appealing for general learning-based approaches. However, learning from human demonstrations introduces additional challenges. The embodiment gap between humans and robots have to be handled with respect to e.g. grasping or kinematic constraints. Furthermore, often there are differences in the visual observation space. Human demonstrations are typically provided from the 3D-person perspective whereas robot demonstrations and skill execution are usually performed from the same perspective.

We present Demonstration Imitation by Trajectory Transformation (DITTO), illustrated in Fig. 1. DITTO consists of two stages; first, a trajectory extraction stage in which we leverage recent object-hand-segmenters [7] and

correspondence detectors [8] to extract relevant objects and calculate their movement in 3D throughout the demonstration. In the second stage, trajectory generation, we present the robot with a novel scene with the same objects, re-detect the objects, and estimate their relative poses. Based on these poses, we warp and interpolate the trajectory. To eventually execute the task, we use an off-the-shelf grasping method [9] and motion planning algorithms to execute the final trajectory. We extensively evaluate each phase of our pipeline in an offline procedure and test it on a real robot to examine failure cases. We show exemplary produced trajectories in Fig. 4.

In summary, we make the following main contributions:

- A novel, modular method for 1-shot transfer from RGB-D human manipulation demonstration videos to robots.
- Experiments validating the method and its ablations, conducted both in an offline manner and on a real robot.
- Open source data and code is publicly available at <http://ditto.cs.uni-freiburg.de>.

## II. RELATED WORK

Imitation learning is a common paradigm to teach a robot a new task [1], [10]. The numerous approaches to this problem can be characterized based on different factors, e.g. by the number of samples used for imitation learning or whether robot or human demonstrations are collected. The following section highlights recent advancements in learning from a few robot demonstrations as well as learning from human demonstrations.

### A. Imitation Learning from Robot Demonstrations

A typical approach to collect demonstrations is through teleoperating a robot [11] for example directly through an external controller [2], [5] or a tracked human hand [12]. Nonetheless, as the human and robot morphology is vastly different, teleoperating a robot can be tedious. Thus, recently, researchers started to investigate how to reduce the amount of needed demonstrations [2].

*One- and Few-Shot Robot Imitation:* Imitation learning from a single or few robot demonstration is a challenging endeavour as it is difficult to separate the intrinsic geometric invariances that define a task from coincidental ones. Few-shot methods often make use of sparse representations to learn invariances more efficiently, examples of this include explicit object proposals [5] or keypoint trajectories [13]. Another strategy is to make use of heterogeneous demonstrations from different tasks [14], [15]. One-shot methods often compensate for the availability of other demonstrations by requiring additional inputs such as foreground object segmentation mask [6], [16], [17] or demonstrations with singulated objects [18].

### B. Imitation Learning from Human Demonstrations

While imitation learning from human demonstration videos is compelling, human demonstrations suffer from a substantial embodiment shift between humans and robots,

even in the case of humanoid robots. Representation learning from human demonstrations can occur on different levels, starting with visual feature learning, as is done by R3M [19]. In more explicitly structured systems, there are the options of learning visual affordances [20], value functions [21] or category-level representations [22], [23]. One work made use of eye-in-hand type human demonstration data [24]. Similar to DITTO, WHIRL [25] extracts a human prior from the given demonstration video by using an off-the-shelf hand pose detector. A number of other works also use large numbers of human demonstrations to learn generative models, which generate intermediate representations such as segmentation maps [26] of hands, flow [27], or trajectories [28] on which a policy is based.

*One- and few-shot Human Imitation:* One and few-shot imitation from human demonstrations is particularly challenging as it compounds the problem of identifying invariances with an embodiment gap. Examples of few-shot imitation from human demonstrations include the work by Kyriazi *et al.* [29] and EquivAct [30]. Similar to these and our approach, Zimmermann *et al.* [4] make use of human and fiducial marker-based pose estimation to learn manipulation trajectories from few demonstrations. One-shot imitation from human demonstrations was done by following a meta-learning based approach [31], as well as translating tasks to a shared latent space and generating actions by either using the latent representations as inputs for a reinforcement learning policy [32] or behavioral cloning [33]. Finally, an applied system that explicitly models contacts and computes relative poses was presented by Guo *et al.* [34]. In contrast, DITTO does not require an explicit optimization or learning step and directly transfers to the robot.

### C. Segmentation & Human Action Understanding

Semantic image segmentation is a well established task in computer vision [35]. Recently, generic segmentation models that segment all objects such as SAM [36] or segment with few annotated labels such as SPINO [37] have become widely used. Other works, such as Hands23 [38] focus on human interactions with objects by estimating bounding boxes and segmentation masks for hands, manipulated objects, as well as secondary objects (containers). Other manipulation datasets such as Ego4D [39] have led to the learning of object-centric video representations [40].

### D. Correspondences and Detection

The core of DITTO strongly relies on a robust relative pose estimation. Classically, given two unordered point clouds, if correspondences between points are unknown, the iterative closest point (ICP) method can be used to estimate the relative pose between the point clouds. Alternatively, given point correspondences it is possible to directly estimate the relative pose using singular value decomposition [41]. Leveraging available RGB-D images, these point correspondences can first be obtained via pixel correspondences, for which we highlight a multitude of methods in the following.

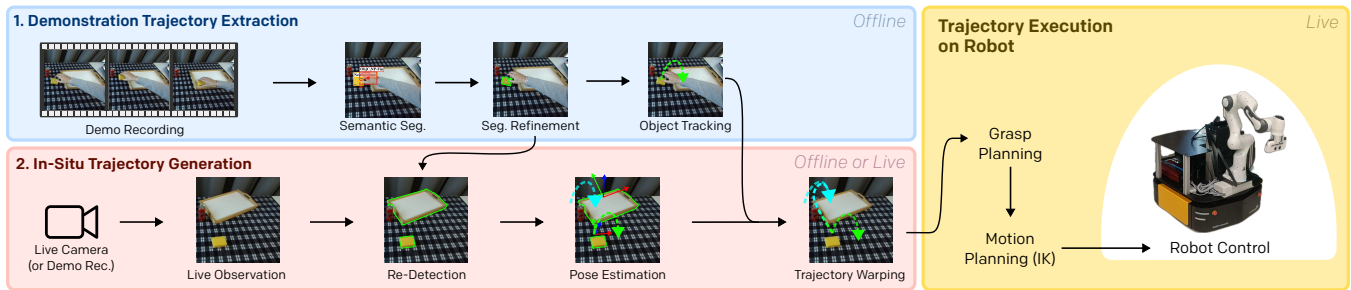


Fig. 2: Our method first computes masks and trajectories from demonstration videos, then maps these onto new live observations by accounting for the change in object poses. These warped trajectories can either be evaluated separately or executed on a robot by using grasp planning and IK trajectory solvers.

There is a wide range of semi-dense correspondence methods available that are typically used for ego-pose estimation problems, e.g. in SLAM. These methods can either be based on analytical features such as SIFT [42] and ORB [43] or learned features such as SuperGlue [44] and LoFTR [8].

Dense methods that are used for estimating optical flow, such as RAFT [45] or UniMatch [46], can also be used to establish correspondences. Unlike semi-dense methods, flow-prediction methods have the advantage of guaranteeing the existence of correspondences. However, as flow is trained on a data distribution characterized by small rotations and translations, it lacks the capability to establish global correspondences or correspondences between strongly rotated objects. The classical pose estimation problem is not further discussed here, as it requires known object CAD models.

Most pose estimation systems rely on upstream detection methods, which aim to identify the rough location of a relevant object. A simple but yet competitive detection method is CNOS [47], which re-detects template images by comparing DINOv2 [48] descriptor features that have been pooled according to SAM masks, yielding detections in the form of segmentation masks. This problem is also addressed in other works [49].

### E. Grasp Generation

Generating stable grasps based on image or point cloud observations is a well studied problem [50]. Grasp generation methods depend on the gripper geometry. While some methods generate grasps for humanoid hands [51], [52], most relevant to our work, are Contact-GraspNet [9] and Anygrasp [53]. Both methods generate grasps for two-finger grippers over whole scenes. Other methods extend these generic setups by making the grasping closed-loop [54] or by combining it with shape reconstruction [55].

## III. DITTO METHOD

Given a demonstration sequence of RGB-D observations  $O = \{o_d^1, \dots, o_d^T\}$ , with length  $T$  and a live observation,  $o_l$ , we aim to infer a robot trajectory  $J_l$  that upon execution will complete the task shown in the demonstration sequence. For this, we take an object-centric approach in which manipulation actions consist of robot end-effector poses that are represented in the manipulated objects frame or a secondary objects frame (e.g. containers). As outlined

in Sec. III-B, our method consists of three stages: a demonstration trajectory extraction stage that can be run as offline pre-processing, an in-situ trajectory generation, and trajectory execution, when running online on a robot. We describe these in detail in this section.

### A. Demonstration Trajectory Extraction

In the first stage, given a sequence of RGB-D observations,  $O$ , we extract the demonstration trajectory  $J_d$  based on the relative object transformations.

**Object-Hand Segmentation:** For all demonstration observations,  $o_d^t \forall t$ , we use Hands23 [38] to compute the segmentation masks for the manipulated object and if present, of a secondary object (e.g. a container). While Hands23 provides state-of-the-art results on their benchmark, we still occasionally observe suboptimal segmentation masks. This may be due to Hands23 design, which processes a single image rather than a full sequence. However, improved versions of this method are very likely to emerge in the near future and we thus circumvent this problem by manually discarding frames for which the segmentation masks are poor, leaving more robust object-hand detection to future research.

**Secondary Object Segmentation:** To obtain segmentation masks for the secondary objects, we follow a similar procedure. We first make use of the secondary object segmentations provided by Hands23. If these are not present, we fall back onto a heuristic (specified in App. A) to identify the secondary object mask. The results are again manually verified and bad masks are discarded.

**Object Pose Extraction:** Subsequently, given object segmentation masks throughout the sequence, we compute the trajectory of the manipulated object. This is done by performing relative pose estimation between pairs of subsequent time steps  $(t, t+1)$ . As the translation and rotation between subsequent observations is small and due to hand occlusions the amount of visible object points is low, we prefer methods with a high recall over precision. Thus, from the relative pose estimation methods outlined in Sec. II-D, we chose to estimate correspondences with the flow estimation method RAFT [45].

After establishing correspondences, we filter them based on the object mask. Given the depth images, we then lift the correspondences to 3D and compute the least-squares rigid transformation  ${}^t\mathbf{T}_{t+1} \in SE(3)$  using singular value

decomposition [41]. To make this process robust to outliers, we estimate inliers through a RANSAC [56] procedure. We perform relative pose estimation for all pairs of subsequent images and aggregate them in our demonstration trajectory  $J_d = \{{}^1\mathbf{T}_2, \dots, {}^{T-1}\mathbf{T}_T\}$ .

**Hand Position Extraction:** Lastly, we also estimate the hand position in relation to the manipulated object. At the time step where the hand is just about to grasp the object to be manipulated, we lift the center of the 2D hand mask to 3D, resulting in  ${}^c\mathbf{p}_h$ . For lifting, we use the given depth image. The hand position relative to the manipulated object  ${}^c\mathbf{T}_o$  is then given by

$${}^o\mathbf{p}_h = ({}^c\mathbf{T}_o)^{-1} {}^c\mathbf{p}_h \quad (1)$$

where  ${}^c\mathbf{T}_o$  is a canonical frame of the manipulated object.

### B. In-Situ Trajectory Generation

In the second stage, given a live RGB-D observation  $o_l$  with the same objects visible, we will first warp the previously extracted trajectory  $J_d$  and then execute it. Similar to before, we first estimate the relative pose of the manipulated object between the first demonstration observation  $o_d^1$  and the live observation  $o_l$ . If applicable, we do the same for the secondary object. Given the results, we warp the demonstration trajectory  $J_d$  and retrieve a resulting live trajectory  $J_l$  which is passed to the robot for execution.

**Re-Detection and Relative-Pose Estimation:** While pose estimation systems work on full images, it is common practice to first run detection systems to extract a region containing the relevant objects [47]. This is particularly useful in our case as we use local flow-based correspondence methods. Thus, to improve robustness, we propose to use a modified version of CNOS [47].\*

The method originally assumed known CAD models, which are rendered to provide template views for re-detection. We replace the template views with actual views, cropped from demonstration images. Re-detection allows us to create tight crops around the object of interest for both the demonstration observation  $o_d$  and the live observation  $o_l$ . Similar to Sec. III-A, we perform relative pose estimation on the cropped observations. One drawback of using such a two-step approach is the fact that if the mask re-detection fails the correspondence estimation will also fail as there is no way to retrieve information from the cropped image parts.

Alternatively, referring to Sec. II-D, we also propose to replace the inherently local flow estimation with a semi-dense, global method, LoFTR [8] which does not require an additional detection step. This decision is motivated by the fact that we are faced with the vice-versa case of the previously discussed trajectory extraction. During trajectory generation, we are faced with potentially strong rotations but very little occlusions, thus we can sacrifice a lower recall for higher precision of LoFTR compared to flow-based methods. When using LoFTR we calculate a re-detection mask by fitting a bounding box around all detected correspondences.

\*For simplicity, this modified version is also referred to as CNOS in the paper as the changes are minor.

Nonetheless, we perform the relative pose estimation step for the manipulated object  ${}^d\mathbf{T}_l^o \in SE(3)$  and if applicable for the secondary object  ${}^d\mathbf{T}_l^s \in SE(3)$ .

**Trajectory Warping:** In the next step, the demonstration trajectory  $J_d$  is warped to the live scene, yielding the object trajectory in the live scene. In the simpler case, with no secondary object present, we use the relative pose change of the object  ${}^d\mathbf{T}_l^o$  and apply it to the demonstration trajectory  $J_d$  as

$${}^{l,o}J_d = \{{}^t\mathbf{T}_{t+1} {}^d\mathbf{T}_l^o \mid \forall {}^t\mathbf{T}_{t+1} \in J_d\}. \quad (2)$$

In the extended case, if a secondary object is present, we perform the same application as in Eq. (2) but with  ${}^d\mathbf{T}_l^s$

$${}^{l,s}J_d = \{{}^t\mathbf{T}_{t+1} {}^d\mathbf{T}_l^s \mid \forall {}^t\mathbf{T}_{t+1} \in J_d\}. \quad (3)$$

We have two potential live trajectories, one based on the object's location  ${}^{l,o}J_d$  and one based on the secondary's object location  ${}^{l,s}J_d$ . To obtain a single final trajectory we smoothly interpolate them<sup>†</sup> using slerp [57]

$${}^lJ_d = \left\{ \alpha(t) {}^{l,o}J_d \oplus (1 - \alpha(t)) {}^{l,s}J_d \right\} \quad (4)$$

with Gaussian weights

$$\alpha(t) = G(t \mid 0, \sigma(T - 1)) \in \mathbb{R} \quad (5)$$

as detailed in App. B.

### C. Trajectory Execution on Robot

Until now we focused on object motion. The next sections describe how to generate robot motion for a specific robot morphology under the assumption of a stable grasp. This addresses the problem of differing embodiments between humans and robots.

**Grasp Generation and Selection:** We use Contact-GraspNet [9] as an off-the-shelf grasping method to detect possible grasps  $G$  in the live scene. The grasps are computed using the initial live observation and then filtered using the object mask from re-detection to give the subset of grasps only on the object to be manipulated. We further filter grasps via inverse kinematic computation (see below) based on reachability and the potentially resulting full robot trajectory.

Additionally, we use the estimated object pose and the relative hand position to conduct a type of affordance transformation by choosing the grasp with the smallest distance to the hand detection  ${}^c\mathbf{T}_g \in G$ . For this, we leverage the previously estimated relative transformation  ${}^d\mathbf{T}_l^o$  to transform the hand position  ${}^o\mathbf{p}_h$  back to the live camera frame  ${}^o\mathbf{p}_c^l$ .

**Motion Planning and Robot Control:** Based on the grasp, we then compute the end-effector joint trajectory that yields our desired warped object trajectory. We then calculate a full robot joint trajectory for our end-effector pose sequence which includes pre-grasp pose, grasp pose, and all poses of the generated trajectory using KDL kinematics<sup>‡</sup>. For the execution, we plan and execute the grasp and the generated

<sup>†</sup> $\alpha\mathbf{A} \oplus (1 - \alpha)\mathbf{A}$  is a generalized addition in the  $SE(3)$ -space.

<sup>‡</sup>Default ROS MoveIt Solver



Fig. 3: Robot setup showing the Franka manipulator, with an end-of-arm depth camera, mounted onto a mobile base.

trajectory separately as we stop in between to close the gripper and confirm the grasp. Note that since the relative pose changes in the generated trajectory are quite small ( $\sim 0.05$  [m]/  $\sim 0.15$  [rad]), the motion planning algorithm is heavily restricted in its search space, potentially inducing failures.

#### IV. EXPERIMENTS

We evaluate our approach in three different configurations: through live real robot executions (see Sec. IV-D) using the robot shown in Fig. 3, and on offline data by predicting correspondences (see Sec. IV-B) and the object trajectories (see Sec. IV-C) of demonstration episodes. While the offline evaluations have several advantages, the most important are speed of evaluation and the comparability of results, it also remains an inherently approximate evaluation, see Sec. IV-B and Sec. IV-C respectively.

##### A. Experimental Setup

We perform experiments using a Franka Panda robot arm mounted on a mobile robot base as shown in Fig. 3. We consider a mixture of table-top manipulation tasks along with manipulation of articulated kitchen furniture, resulting in a total of 10 tasks. A full list of tasks is given in App. C, together with example images shown in Fig. 1 and Fig. 4. For each task (except `plug_charger`), we recorded five demonstrations with varying initial poses of both the manipulated objects and secondary objects.

In theory, demonstrations and inference can use RGB-D observations from different cameras, however, for convenience we both record demonstrations and run inference with a StereoLab ZED2i. We record demonstration videos with a static camera position. This allows easy recording of human

Corresp.	Method		Segm. rate [%]	Inlier-num. [N]	Runtime [s]
	Detection				
RAFT [45]	-		88.2	<b>5472</b>	<b>0.56</b>
RAFT [45]	CNOS [47]		77.7	4821	8.18
LoFTR [8]	-		<b>89.4</b>	65	0.80
LoFTR [8]	CNOS [47]		72.8	32	8.29

(a) Tracking evaluation, within single demonstrations.

Corresp.	Method		Segm. rate [%]	Inlier-num. [N]	Runtime [s]
	Detection				
RAFT [45]	-		56.4	2308	<b>0.54</b>
RAFT [45]	CNOS [47]		71.4	<b>2922</b>	7.99
LoFTR [8]	-		<b>79.4</b>	25	0.75
LoFTR [8]	CNOS [47]		63.4	14	8.02

(b) Re-detection evaluation, between different demonstrations.

TABLE I: Tracking and re-detection evaluation based on correspondence methods. We evaluate the performance of correspondence methods by detecting correspondences between a source and a target image given a source mask. We measure the absolute inlier count ( $N$ ) of established correspondences as well as the percentage of inliers (%) that map to the ground truth target mask. We evaluate on two different setups, the first one, being within a demonstration (in Tab. Ia) and the second one, between the initial observation of two demonstrations (in Tab. Ib).

Corresp.	Method		Traj. Pose Errors	
	Detection		Rot. [rad]	Trns. [m]
RAFT [45]	-		0.2243	0.0401
RAFT [45]	CNOS [47]		0.2288	0.0441
LoFTR [8]	-		<b>0.2226</b>	<b>0.0387</b>
LoFTR [8]	CNOS [47]		0.2417	0.0437

TABLE II: Relative Trajectory Pose Estimation Errors. For a set of demonstrations, we assume one of them as our given demonstration. From the set of remaining demonstrations, we will use the first image as a hypothesized live observation. We then calculate the translation and rotation error between the relative change in the demonstration trajectory and the generated trajectory. This assumes that each trajectory is executed with the same movement (direction and speed) and thus, the change between two steps should be similar. Note that this can only be considered a pseudo-error as the assumption can not be strictly enforced, due to human errors. Pose estimation is done using least-squares rigid motion.

demonstrations from a third-person perspective and aligns with the prospective setting of having a robot watch a human demonstration and then being able to imitate it. For the purpose of faster computation, we subsample demonstration videos to a fixed length of  $T = 11$  observations, yielding ten steps, with a linear spacing between frames.

##### B. Demo. Based Tracking & Re-Detection Evaluation

In the first offline experiment, we evaluate the task of finding correspondences in a target image given a source image and a mask from which we want to establish correspondences. This procedure is an integral part of and used twice in DITTO. Once, when tracking the object within a demonstration and once, when the object is re-detected in a live observation. We evaluate the quality of the various correspondence methods outlined Sec. II-D and substantiate our decisions taken in Sec. III. We use the segmentation masks from our pre-processing procedure (refer to Sec. III-A) which includes manual verification of masks. For the evaluation criterion, we count the percentage of correspondence points

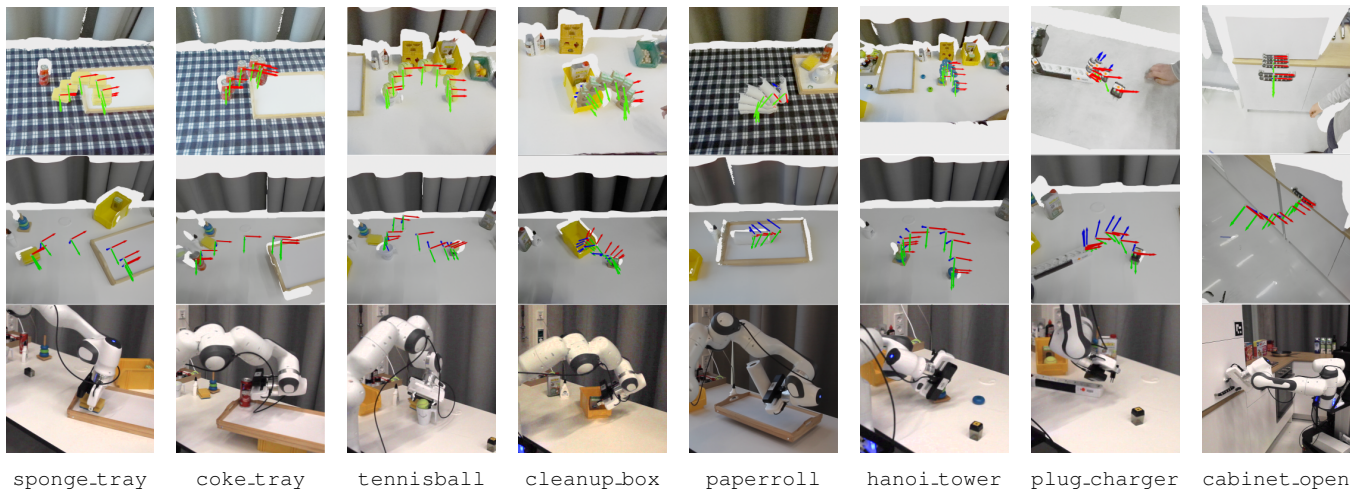


Fig. 4: Examples of trajectory generation, shown for various different tasks. *Top row*: rendered examples of trajectories extracted from human demonstrations, in-painted into the initial demonstration observation. *Middle row*: rendered trajectories that have been generated in-situ for the robot imitation, in-painted into the live robot view. *Bottom row*: images from live robot imitation runs.

that remain within the segmentation masks in the next frame (precision) and their absolute amount (recall). We report the results in Tab. I.

We observe that when establishing correspondences within a single demonstration (see Tab. Ia) there is no significant precision increase of LoFTR (89.4%) over RAFT (88.2%) but RAFT’s recall is  $84\times$  higher. Using CNOS re-detection within a demonstration consistently decreases performance as the re-detection is more likely to fail due to the hand holding and thus, occluding the object heavily from time. In contrast, when estimating correspondences between demonstrations (see Tab. Ib), one observes that using the CNOS mask re-detection helps to overcome the shortcomings of the local flow estimation. On the other hand, the CNOS mask re-detection hurts the global LoFTR method, as it can generate incorrect mask crops and lead to premature failures. Thus, we chose to track the object within a demonstration only using RAFT and between demonstrations LoFTR.

### C. Demo. Based Trajectory Generation Evaluation

In this setup, we compare the relative change of poses between the transformed demonstration trajectory  ${}^lJ_d$  and the pseudo ground truth trajectory  $J_l$ . This comparison shows how well the relative pose estimation and mixing components perform. As before, we compare against various combinations outlined in Sec. II-D. Given two demonstration episodes from our set, we consider one of them  $R_d$  as our input to DITTO and the other one as our pseudo ground truth  $R_l$ . Thus, we perform an offline evaluation under the assumption that the extracted trajectory  $J_l$  in  $R_l$  is valid. As verified in Sec. IV-B this is a reasonable assumption. It is important to note that this experiment cannot yield perfect results unless the human performs the task in the exact same manner, which is nearly impossible. To address these concerns, we took significant care to perform the same movement in all demonstrations.

We compare the difference in relative poses between

trajectories for a given time step. We calculate the translation error through Euclidean distance and the rotation error using angle-axis. Results are shown in Tab. II. As expected, no method achieves an error close to zero, but the overall results of this experiment align closely with those of in the inter-demonstration correspondence evaluation experiment (refer to Tab. Ib)

### D. Real Robot Evaluation

Given the promising results of the previous experiments in Sec. IV-C, we evaluate DITTO on the real-world robot setup. We additionally ablate using the proposed CNOS re-detection step (over LoFTR) to detect grasping regions as well as hand affordance. Given our previously collected ten tasks, we again set these up under similar conditions. We then thoroughly evaluate DITTO on over 150 real-world executions and conclude the following results and drawbacks of our proposed method. Overall, DITTO is able to correctly warp the demonstration trajectory to the live scene in 79% of our evaluation runs. For the remaining runs DITTO fails because no correspondences could be established (e.g. when the objects are heavily rotated). In the case of successful transfers, the majority of subsequent failures are caused by the robot’s kinematic constraints which are limited compared to the human teacher. For a visualization of success and failure cases we refer to Fig. 4.

**In-Depth Task Analysis:** For the easiest pick-and-place tasks, `sponge_tray` and `coke_tray`, we achieve a high success rate even under modifications of the scene (e.g. tray moved up). The more difficult pick-and-place tasks, which require greater precision, `tennisball` and `cleanup_box`, are also executed well when given a similar setup as shown in the demonstration. For the re-orientation task `paperroll`, we frequently observe an execution to almost until the end of the task, just before the robot needs to lower its wrist, at which point it moves into its joint limits. For the precise insertion tasks, `hanoi_ring` and

plug\_charger, as well as the pouring task pour\_cup, despite actual imprecisions in object pose estimation the most common failure case is grasping the object. Our used task objects are quite small and have intricate features. For the articulated object manipulation tasks, cabinet\_open and drawer\_open, we encounter two main problems preventing successful execution, first, due to sensor noise the predicted grasps on the narrow handles are colliding with the environment and second, the inverse kinematics solution often leading the robot into a singularity. This behavior is expected as prior work has previously shown that for articulated object manipulation in the wild, a mobile base is beneficial [58].

**Ablation Results:** We see no significant difference when using the hand affordance-based grasp filtering. Nonetheless, we would expect the filtering to yield an improvement when focusing on tasks where the grasp pose is crucial for task success. Similar to the quantitative evaluation in Sec. IV-B, using CNOS as a pre-detection method sometimes fails catastrophically when similar objects are present in the scene.

## V. CONCLUSION

We present DITTO, a modular method for strong one-shot imitation from human demonstration videos. We evaluated different variations of DITTO in an offline manner, proving its potential use in real-world robotic tasks as well as on a real robot setup, demonstrating its efficacy, and identifying key weaknesses. To facilitate, future research we made the code publicly available. Potential improvements could include setting up a standardized benchmark that allows future researchers to iterate on subcomponents separately in order to improve the overall performance of DITTO. For instance, this could include fine-tuning the segmentation model or locally refining the predicted grasps. The robotic execution can also be improved, for example, by integrating the robot’s mobile base into the motion planning process to tackle tasks that are kinematically more challenging.

## APPENDIX

### A. Secondary Object Segmentation

In cases where Hands23 [38] did not detect a secondary object mask, we propose an alternative approach. At the last time step of a manipulation action, we first use SAM [36] to fully segment the image. Given each mask, we will lift them into 3D, resulting in a set of point clouds. We then chose the point cloud (and consequently the secondary object mask) with smallest minimal distance to the manipulated object point cloud, i.e. if they are in contact the distance is 0.

### B. Combining Trajectories

Given two trajectories, at each timestep  $t$  we interpolate between them using a time-dependent mixing weight  $\alpha(t)$ . The positions are summed and the rotation is interpolated using slerp [57]. The mixing weight

$$\alpha(t) = G(t \mid 0, \sigma(T - 1)) \in \mathbb{R} \quad (6)$$

is a Gaussian distributed coefficient where  $\sigma$  is a hyperparameter controlling the steepness of the mixing curve and

$T - 1$  is the number of trajectory steps. If too steep, i.e. too small  $\sigma \rightarrow 0$ , there will be a sudden jump in the middle of the trajectory, if too flat, i.e. too large  $\sigma \rightarrow \text{inf}$ , two sudden jumps will happen close to the beginning and the end. We chose  $\sigma = 1/2$  as it is a good trade-off between both.

### C. List of Experimental Tasks

A list of our experimental tasks with brief descriptions is detailed in Tab. III.

Name	Object	Secondary Object	Action Type
coke_tray	Coke can	Kitchen tray	Pick and place
sponge_tray	Sponge	Kitchen tray	Pick and place
drawer_open	Drawer handle	-	Articulated obj. man.
cabinet_open	Cabinet handle	-	Articulated obj. man.
hanoi_ring	Hanoi tower ring	Wood peg	Insertion
plug_charger	Phone charger	Socket bar	Insertion
pour_cup	Mug	Gray bowl	Pouring
paperroll	Paper-towel roll	-	Re-orienting
tennisball	Tennis ball	Cup	Pick and place
cleanup_box	Cardboard box	Storage box	Pick and place

TABLE III: Overview of the tasks used in the experiments. Some example images are shown in Fig. 1.

## REFERENCES

- [1] C. Celemin, R. Pérez-Dattari, E. Chisari, G. Franzese, L. de Souza Rosa, R. Prakash, Z. Ajanović, M. Ferraz, A. Valada, J. Kober, *et al.*, “Interactive imitation learning in robotics: A survey,” *Foundations and Trends® in Robotics*, vol. 10, no. 1-2, pp. 1–197, 2022.
- [2] E. Chisari, T. Welschhold, J. Boedecker, W. Burgard, and A. Valada, “Correct me if i am wrong: Interactive learning for robotic manipulation,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3695–3702, 2022.
- [3] J. Zhao, A. Giammarino, E. Lamon, J. M. Gandarias, E. De Momi, and A. Ajoudani, “A hybrid learning and optimization framework to achieve physically interactive tasks with mobile manipulators,” *IEEE Rob. and Auto. Let.*, vol. 7, no. 3, pp. 8036–8043, 2022.
- [4] C. Zimmermann, T. Welschhold, C. Dornhege, W. Burgard, and T. Brox, “3d human pose estimation in rgbd images for robotic task learning,” *Proc. IEEE Int. Conf. on Rob. and Auto.*, pp. 1986–1992, 2018.
- [5] Y. Zhu, A. Joshi, P. Stone, and Y. Zhu, “Viola: Imitation learning for vision-based manipulation with object proposal priors,” *Conf. on Robot Learning*, 2022.
- [6] P. Vitiello, K. Dreczkowski, and E. Johns, “One-shot imitation learning: A pose estimation perspective,” in *Conf. on Robot Learning*, 2023.
- [7] T. Cheng, D. Shan, A. S. Hassen, R. E. L. Higgins, and D. Fouhey, “Towards a richer 2d understanding of hands at scale,” in *Proc. Adv. Neural Inform. Process. Syst.*, 2023.
- [8] J. Sun, Z. Shen, Y. Wang, H. Bao, and X. Zhou, “Loft: Detector-free local feature matching with transformers,” *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 8918–8927, 2021.
- [9] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, “Contact-graspnet: Efficient 6-dof grasp generation in cluttered scenes,” *Proc. IEEE Int. Conf. on Rob. and Auto.*, pp. 13438–13444, 2021.
- [10] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, and J. Peters, “An algorithmic perspective on imitation learning,” *Foundations and Trends® in Robotics*, vol. 7, no. 1-2, pp. 1–179, 2018.
- [11] W. Si, N. Wang, and C. Yang, “A review on manipulation skill acquisition through teleoperation-based learning from demonstration,” *Cognitive Computation and Systems*, vol. 3, no. 1, pp. 1–16, 2021.
- [12] Y. Qin, H. Su, and X. Wang, “From one hand to multiple hands: Imitation learning for dexterous manipulation from single-camera teleoperation,” *IEEE Rob. and Auto. Let.*, vol. 7, no. 4, pp. 10873–10881, 2022.
- [13] M. Vecerík, C. Doersch, Y. Yang, T. Davchev, Y. Aytar, G. Zhou, R. Hadsell, L. de Agapito, and J. Scholz, “Robotap: Tracking arbitrary points for few-shot visual imitation,” *arXiv preprint arXiv:2308.15975*, pp. 8866–8873, 2023.

- [14] J.-F. Yeh, C.-M. Chung, H.-T. Su, Y.-T. Chen, and W. H. Hsu, "Stage conscious attention network (SCAN) : A demonstration-conditioned policy for few-shot imitation," in *AAAI*, 2022.
- [15] L. Wang, J. Zhao, Y. Du, E. H. Adelson, and R. Tedrake, "Poco: Policy composition from and for heterogeneous robot learning," in *Proc. Rob.: Sci. and Syst.*, 2024.
- [16] M. Argus, L. Hermann, J. Long, and T. Brox, "Flowcontrol: Optical flow based visual servoing," *Proc. IEEE Int. Conf. on Intel. Rob. and Syst.*, pp. 7534–7541, 2020.
- [17] Y. Zhu, Z. Jiang, P. Stone, and Y. Zhu, "Learning generalizable manipulation policies with object-centric 3d representations," in *Conf. on Robot Learning*, 2023.
- [18] N. D. Palo and E. Johns, "Dinobot: Robot manipulation via retrieval and alignment with vision foundation models," in *Proc. IEEE Int. Conf. on Rob. and Auto.*, 2024.
- [19] S. Nair, A. Rajeswaran, V. Kumar, C. Finn, and A. Gupta, "R3m: A universal visual representation for robot manipulation," in *Conf. on Robot Learning*, 2022.
- [20] S. Bahl, R. Mendonca, L. Chen, U. Jain, and D. Pathak, "Affordances from human videos as a versatile representation for robotics," *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 01–13, 2023.
- [21] C. Bhateja, D. Guo, D. Ghosh, A. Singh, M. Tomar, Q. H. Vuong, Y. Chebotar, S. Levine, and A. Kumar, "Robotic offline rl from internet videos via value-function pre-training," in *Proc. IEEE Int. Conf. on Rob. and Auto.*, 2024.
- [22] J. Gao, Z. Tao, N. Jaquier, and T. Asfour, "K-vil: Keypoints-based visual imitation learning," *IEEE Trans. on Robotics*, vol. 39, pp. 3888–3908, 2022.
- [23] B. Wen, W. Lian, K. E. Bekris, and S. Schaal, "You only demonstrate once: Category-level manipulation from single visual demonstration," in *Proc. Rob.: Sci. and Syst.*, 2022.
- [24] M. J. Kim, J. Wu, and C. Finn, "Giving robots a hand: Learning generalizable manipulation with eye-in-hand human video demonstrations," *arXiv preprint arXiv:2307.05959*, 2023.
- [25] S. Bahl, A. Gupta, and D. Pathak, "Human-to-robot imitation in the wild," in *Proc. Rob.: Sci. and Syst.*, 2022.
- [26] H. Bharadhwaj, A. Gupta, V. Kumar, and S. Tulsiani, "Towards generalizable zero-shot manipulation via translating human interaction plans," in *Proc. IEEE Int. Conf. on Rob. and Auto.*, 2024.
- [27] P.-C. Ko, J. Mao, Y. Du, S.-H. Sun, and J. B. Tenenbaum, "Learning to Act from Actionless Video through Dense Correspondences," in *Int. Conf. Learn. Represent.*, 2024.
- [28] H. Bharadhwaj, A. Gupta, S. Tulsiani, and V. Kumar, "Zero-shot robot manipulation from passive human videos," *arXiv preprint arXiv:2302.02011*, 2023.
- [29] N. Kyriazis and A. A. Argyros, "Tracking of hands interacting with several objects," in *IEEE Int. Conf. on Computer Vision Workshops*, 2015.
- [30] J. Yang, C. Deng, J. Wu, R. Antonova, L. J. Guibas, and J. Bohg, "Equivact: Sim(3)-equivariant visuomotor policies beyond rigid object manipulation," in *Proc. IEEE Int. Conf. on Rob. and Auto.*, 2024.
- [31] T. Yu, C. Finn, A. Xie, S. Dasari, T. Zhang, P. Abbeel, and S. Levine, "One-shot imitation from observing humans via domain-adaptive meta-learning," in *Proc. Rob.: Sci. and Syst.*, 2018.
- [32] L. Pauly, W. C. Agboh, D. C. Hogg, and R. Fuentes, "O2a: One-shot observational learning with action vectors," *Frontiers in Robotics and AI*, vol. 8, 2018.
- [33] S. Dasari and A. K. Gupta, "Transformers for one-shot visual imitation," in *Proc. Conf. on Rob. Learn.*, 2020.
- [34] D. Guo, "Learning multi-step manipulation tasks from a single human demonstration," *arXiv preprint arXiv:2312.15346*, 2023.
- [35] J. V. Hurtado and A. Valada, "Semantic scene segmentation for robotics," in *Deep learning for robot perception and cognition*, pp. 279–311, Elsevier, 2022.
- [36] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. B. Girshick, "Segment anything," *Proc. Int. Conf. Comput. Vis.*, pp. 3992–4003, 2023.
- [37] M. Käppeler, K. Petek, N. Vödisch, W. Burgard, and A. Valada, "Few-shot panoptic segmentation with foundation models," in *Proc. IEEE Int. Conf. on Rob. and Auto.*, 2024.
- [38] D. Shan, J. Geng, M. Shu, and D. F. Fouhey, "Understanding human hands in contact at internet scale," *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 9866–9875, 2020.
- [39] K. Grauman, A. Westbury, E. Byrne, Z. Chavis, A. Furnari, R. Girdhar, *et al.*, "Ego4d: Around the world in 3,000 hours of egocentric video," *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 18973–18990, 2021.
- [40] C. Zhang, C. Fu, S. Wang, N. Agarwal, K. Lee, C. Choi, and C. Sun, "Object-centric video representation for long-term action anticipation," in *Proc. IEEE Win. Conf. on App. of Comput. Vis.*, pp. 6737–6747, 2024.
- [41] K. S. Arun, T. S. Huang, and S. D. Blostein, "Least-squares fitting of two 3-d point sets," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. PAMI-9, pp. 698–700, 1987.
- [42] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. Journal of Computer Vision*, vol. 60, pp. 91–110, 2004.
- [43] E. Rublee, V. Rabaud, K. Konolige, and G. R. Bradski, "Orb: An efficient alternative to sift or surf," *Proc. Int. Conf. Comput. Vis.*, pp. 2564–2571, 2011.
- [44] P.-E. Sarlin, D. DeTone, T. Malisiewicz, and A. Rabinovich, "Superglue: Learning feature matching with graph neural networks," *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 4937–4946, 2019.
- [45] Z. Teed and J. Deng, "Raft: Recurrent all-pairs field transforms for optical flow," in *Proc. Springer Eur. Conf. Comput. Vis.*, 2020.
- [46] H. Xu, J. Zhang, J. Cai, H. Rezatofghi, F. Yu, D. Tao, and A. Geiger, "Unifying flow, stereo and depth estimation," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 2023.
- [47] V. N. Nguyen, T. Hodan, G. Ponimatkin, T. Groueix, and V. Lepetit, "Cnos: A strong baseline for cad-based novel object segmentation," *IEEE/CVF Int. Conf. on Computer Vision Workshops*, pp. 2126–2132, 2023.
- [48] M. Oquab, T. Darcet, T. Moutakanni, H. Q. Vo, M. Szafraniec, *et al.*, "Dinov2: Learning robust visual features without supervision," *Trans. Mach. Learn. Res.*, vol. 2024, 2024.
- [49] E. Valassakis, G. Papagiannis, N. D. Palo, and E. Johns, "Demonstrate once, imitate immediately (dome): Learning visual servoing for one-shot imitation learning," *Proc. IEEE Int. Conf. on Intel. Rob. and Syst.*, pp. 8614–8621, 2022.
- [50] R. Newbury, M. Gu, L. Chumbley, A. Mousavian, C. Eppner, J. Leitner, J. Bohg, A. Morales, T. Asfour, D. Kragic, *et al.*, "Deep learning approaches to grasp synthesis: A review," *IEEE Transactions on Robotics*, vol. 39, no. 5, pp. 3994–4015, 2023.
- [51] W.-J. Baek, C. Pohl, P. Pelcz, T. Kröger, and T. Asfour, "Improving humanoid grasp success rate based on uncertainty-aware metrics and sensitivity optimization," *IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids)*, pp. 786–793, 2022.
- [52] M. A. Roa, M. J. Argus, D. Leidner, C. W. Borst, and G. Hirzinger, "Power grasp planning for anthropomorphic robot hands," *Proc. IEEE Int. Conf. on Rob. and Auto.*, pp. 563–569, 2012.
- [53] H. Fang, C. Wang, H. Fang, M. Gou, J. Liu, H. Yan, W. Liu, Y. Xie, and C. Lu, "Anygrasp: Robust and efficient grasp perception in spatial and temporal domains," *IEEE Trans. on Robotics*, vol. 39, pp. 3929–3945, 2022.
- [54] P. Piacenza, J. Yuan, J. Huh, and V. Isler, "Vfas-grasp: Closed loop grasping with visual feedback and adaptive sampling," in *Proc. IEEE Int. Conf. on Rob. and Auto.*, pp. 4126–4132, 2024.
- [55] E. Chisari, N. Heppert, T. Welschhold, W. Burgard, and A. Valada, "Centergrasp: Object-aware implicit representation learning for simultaneous shape reconstruction and 6-dof grasp estimation," *IEEE Robotics and Automation Letters*, pp. 5094–5101, 2024.
- [56] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, pp. 381–395, 1981.
- [57] K. Shoemake, "Animating rotation with quaternion curves," in *Proc. SIGGRAPH*, 1985.
- [58] M. Mittal, D. Hoeller, F. Farshidian, M. Hutter, and A. Garg, "Articulated object interaction in unknown scenes with whole-body mobile manipulation," in *Proc. IEEE Int. Conf. on Intel. Rob. and Syst.*, pp. 1647–1654, 2022.