

Learning Human-to-Robot Handovers from Point Clouds

Sammy Christen^{1,2*}
Otmar Hilliges¹

Wei Yang²
Dieter Fox^{2,3}

Claudia Pérez-D'Arpino²
Yu-Wei Chao²

¹ETH Zurich ²NVIDIA ³University of Washington

{sammy.christen, otmar.hilliges}@inf.ethz.ch {weiy, claudiap, dieterf, ychao}@nvidia.com

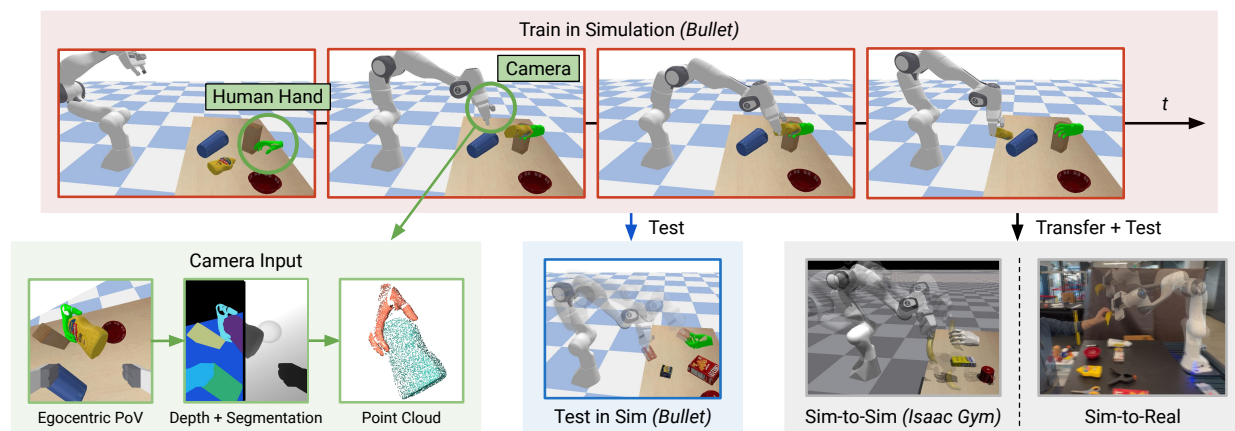


Figure 1. We introduce a framework to learn human-to-robot handover policies from point cloud input. Our policies take input from a wrist mounted camera and directly generate action output for the robot’s end effector. We train our policies in a simulated handover environment, and evaluate on unseen handover motion and poses. We further transfer the model across physics simulators and to a real robotic platform.

Abstract

We propose the first framework to learn control policies for vision-based human-to-robot handovers, a critical task for human-robot interaction. While research in Embodied AI has made significant progress in training robot agents in simulated environments, interacting with humans remains challenging due to the difficulties of simulating humans. Fortunately, recent research has developed realistic simulated environments for human-to-robot handovers. Leveraging this result, we introduce a method that is trained with a human-in-the-loop via a two-stage teacher-student framework that uses motion and grasp planning, reinforcement learning, and self-supervision. We show significant performance gains over baselines on a simulation benchmark, sim-to-sim transfer and sim-to-real transfer. Video and code are available at <https://handover-sim2real.github.io>.

1. Introduction

Handing over objects between humans and robots is an important task for human-robot interaction (HRI) [35]. It

allows robots to assist humans in daily collaborative activities, such as helping to prepare a meal, or to exchange tools and parts with human collaborators in manufacturing settings. To complete these tasks successfully and safely, intricate coordination between human and robot is required. This is challenging, because the robot has to react to human behavior, while only having access to sparse sensory inputs such as a single camera with limited field of view. Therefore, a need for methods that solve interactive tasks such as handovers purely from vision input arises.

Bootstrapping robot training in the real world can be unsafe and time-consuming. Therefore, recent trends in Embodied AI have focused on training agents to act and interact in simulated (sim) environments [11, 12, 19, 43, 45, 46, 51]. With advances in rendering and physics simulation, models have been trained to map raw sensory input to action output, and can even be directly transferred from simulation to the real world [2, 42]. Many successes have been achieved particularly around the suite of tasks of robot navigation, manipulation, or a combination of both. In contrast to these areas, little progress has been made around tasks pertained to HRI. This is largely hindered by the challenges in embedding realistic human agents in these environments, since

*This work was done during an internship at NVIDIA.

modeling and simulating realistic humans is challenging.

Despite the challenges, an increasing number of works have attempted to embed realistic human agents in simulated environments [6, 9, 16, 36–38, 48]. Notably, a recent work has introduced a simulation environment (“HandoverSim”) for human-to-robot handover (H2R) [6]. To ensure a realistic human handover motion, they use a large motion capture dataset [7] to drive the movements of a virtual human in simulation. However, despite the great potential for training robots, the work of [6] only evaluates off-the-shelf models from prior work, and has not explored any policy training with humans in the loop in their environment.

We aim to close this gap by introducing a vision-based learning framework for H2R handovers that is trained with a human-in-the-loop (see Fig. 1). In particular, we propose a novel mixed imitation learning (IL) and reinforcement learning (RL) based approach, trained by interacting with the humans in HandoverSim. Our approach draws inspiration from a recent method for learning policies for grasping static objects from point clouds [50], but proposes several key changes to address the challenges in H2R handovers. In contrast to static object grasping, where the policy only requires object information, we additionally encode human hand information in the policy’s input. Also, compared to static grasping without a human, we explicitly take human collisions into account in the supervision of training. Finally, the key distinction between static object grasping and handovers is the dynamic nature of the hand and object during handover. To excel on the task, the robot needs to react to dynamic human behavior. Prior work typically relies on open-loop motion planners [49] to generate expert demonstrations, which may result in suboptimal supervision for dynamic cases. To this end, we propose a two-stage training framework. In the first stage, we fix the humans to be stationary and train an RL policy that is partially guided by expert demonstrations obtained from a motion and grasp planner. In the second stage, we finetune the RL policy in the original dynamic setting where the human and robot move simultaneously. Instead of relying on a planner, we propose a self-supervision scheme, where the pre-trained RL policy serves as a teacher to the downstream policy.

We evaluate our method in three “worlds” (see Fig. 1). First, we evaluate on the “native” test scenes in HandoverSim [6], which use the same backend physics simulator (Bullet [10]) as training but unseen handover motions from the simulated humans. Next, we perform sim-to-sim evaluation on the test scenes implemented with a different physics simulator (Isaac Gym [29]). Lastly, we investigate sim-to-real transfer by evaluating policies on a real robotic system and demonstrate the benefits of our method.

We contribute: i) the first framework to train human-to-robot handover tasks from vision input with a human-in-the-loop, ii) a novel teacher-student method to train in the

setting of a jointly moving human and robot, iii) an empirical evaluation showing that our approach outperforms baselines on the HandoverSim benchmark, iv) transfer experiments indicating that our method leads to more robust sim-to-sim and sim-to-real transfer compared to baselines.

2. Related Work

Human-to-Robot Handovers Encouraging progress in hand and object pose estimation [22, 26, 27] has been achieved, aided by the introduction of large hand-object interaction datasets [5, 7, 17, 20, 21, 28, 32, 47, 54, 55]. These developments enable applying model-based grasp planning [3, 4, 31], a well-studied approach in which full pose estimation and tracking are needed, to H2R handovers [7, 41]. However, these methods require the 3D shape models of the object and cannot handle unseen objects. Alternatively, some recent works [13, 30, 40, 52, 53] achieve H2R handover by employing learning-based grasp planners to generate grasps for novel objects from raw vision inputs such as images or point clouds [33, 34]. While promising results have been shown, these methods work only on an open-loop sequential setting in which the human hand has to stay still once the robot starts to move [40], or need complex hand-designed cost functions for grasp selection [52] and robot motion planning [30, 53] for reactive handovers, which requires expertise in robot motion and control. Hence, these methods are difficult to reproduce and deploy to new environments. Progress towards dynamic simultaneous motion has been shown by a learning-based method [48], using state inputs, leaving an open challenge for training policies that receive visual input directly. In contrast, we propose to learn control policies together with grasp prediction for handovers in an end-to-end manner from segmented point clouds with a deep neural net. To facilitate easy and fair comparisons among different handover methods, [6] propose a physics-simulated environment with diverse objects and realistic human handover behavior collected by a mocap system [7]. They provide benchmark results of several previous handover systems, including a learning-based grasping policy trained with static objects [50]. However, learning a safe and efficient handover policy is not trivial with a human-in-the-loop, which we address in this work.

Policy Learning for Grasping Object grasping is an essential skill for many robot tasks, including handovers. Prior works usually generate grasp poses given a known 3D object geometry such as object shape or pose [3, 4, 31], which is nontrivial to obtain from real-world sensory input such as images or point clouds. To overcome this, recent works train deep neural networks to predict grasps from sensor data [25] and compute trajectories to reach the predicted grasp pose. Though 3D object geometry is no longer needed, the feasibility is not guaranteed since the grasp

prediction and trajectory planning are computed separately. Some recent works directly learn grasping policies given raw sensor data. [24] propose a self-supervised RL framework based on RGB images to learn a deep Q-function from real-world grasps. To improve data efficiency, [44] use a low-cost handheld device to collect grasping demonstrations with a wrist-mounted camera. They train an RL-based 6-DoF closed-loop grasping policy with these demonstrations. [50] combines imitation learning from expert data with RL to learn a control policy for object grasping from point clouds. Although this method performs well in HandoverSim [6] when the human hand is not moving, it has difficulty coordinating with a dynamic human hand since the policy is learned with static objects. Instead, our policy is directly learned from large-scale dynamic hand-object trajectories obtained from the real world. To facilitate the training for the dynamic case, we propose a two-stage teacher-student framework, that is conceptually inspired by [8], which has been proven critical through experiments.

3. Background

3.1. Reinforcement Learning

MDP We formalize RL as a Markov Decision Process (MDP), that consists of a 5-tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}, \gamma)$, where \mathcal{S} is the state space, \mathcal{A} the action space, \mathcal{R} a scalar reward function, \mathcal{T} a transition function that maps state-action pairs to distributions over states, and γ a discount factor. The goal is to find a policy that maximizes the long-term reward: $\pi^* = \arg \max_{\pi} \mathbb{E} \sum_{t=0}^{t=T} \gamma^t \mathcal{R}(\mathbf{s}_t)$, with $\mathbf{s}_t \sim \mathcal{T}(\mathbf{s}_{t-1}, \mathbf{a}_{t-1})$ and $\mathbf{a}_{t-1} \sim \pi(\mathbf{s}_{t-1})$.

Learning Algorithm In this work, we use TD3 [18], a common algorithm for continuous control. It is an actor-critic method, which consists of a policy $\pi_{\theta}(\mathbf{s})$ (actor) and a Q-function approximator $Q_{\phi}(\mathbf{s}, \mathbf{a})$ (critic) that predicts the expected return from a state-action pair. Both are represented by neural networks with parameters θ and ϕ . TD3 is off-policy, and hence there is a replay buffer in which training transitions are stored. During training, both the actor and critic are updated using samples from the buffer. To update the critic, we minimize the Bellman error:

$$L_{\text{BE}}(\phi) = \mathbb{E}_{\mathcal{M}} \left[\left(Q_{\phi}(\mathbf{s}_t, \mathbf{a}_t) - r(\mathbf{s}_t, \mathbf{a}_t) + \gamma Q_{\phi}(\mathbf{s}_{t+1}, \mathbf{a}_{t+1}) \right)^2 \right] \quad (1)$$

For the actor network, the policy parameters are trained to maximize the Q-values:

$$L_{\text{DDPG}}(\theta) = \mathbb{E}_{\pi} [Q_{\phi}(\mathbf{s}_t, \mathbf{a}_t) | \mathbf{s}_t, \mathbf{a}_t = \pi_{\theta}(\mathbf{s}_t)] \quad (2)$$

For more details, we refer the reader to [18].

3.2. HandoverSim Benchmark

HandoverSim [6] is a benchmark for evaluating H2R handover policies in simulation. The task setting consists of a tabletop with different objects, a Panda 7DoF robotic arm with a gripper and a wrist-mounted RGB-D camera, and a simulated human hand. The task starts with the human grasping an object and moving it to a handover pose. The robot should move to the object and grasp it. The task is successful if the object has been grasped from the human without collision and brought to a designated position without dropping. To accurately model the human, trajectories from the DexYCB dataset [7], which comprises a large amount of human-object interaction sequences, are replayed in simulation. Several baselines [49, 50, 52] are provided for comparison. The setup in HandoverSim has only been used for handover performance evaluation purposes, whereas in this work we utilize it as a learning environment.

4. Method

The overall pipeline is depicted in Fig. 2 and consists of three different modules: perception, vision-based control, and the handover environment. The perception module receives egocentric visual information from the handover environment and processes it into segmented point clouds. The vision-based control module receives the point clouds and predicts the next action for the robot and whether to approach or to grasp the object. This information is passed to the handover environment, which updates the robot state and sends the new visual information to the perception module. Note that the input to our method comes from the wrist-mounted camera, i.e., there is no explicit information, such as object or hand pose, provided to the agent. We will now explain each of the modules of our method in more detail.

4.1. Handover Environment

We split the handover task into two distinct phases (see Fig. 2). First, during the *approaching phase*, the robot moves to a pre-grasp pose that is close to the object by running the learned control policy π . A learned grasp predictor σ continuously computes a grasp probability to determine when the system can proceed to the second phase. Once the pre-grasp pose is reached and the grasp prediction is confident to take over the object from the human, the task will switch to the *grasping phase*, in which the end-effector moves forward to the final grasp pose in open-loop fashion and closes the gripper to grasp the object. Finally, after object grasping, the robot follows a predetermined trajectory to retract to a base position and complete the episode. This task logic is used in both our simulation environment and the real robot deployment. Sequencing based on a pre-grasp pose is widely used in literature for dynamic grasping [1].

We follow the HandoverSim task setup [6], where the hu-

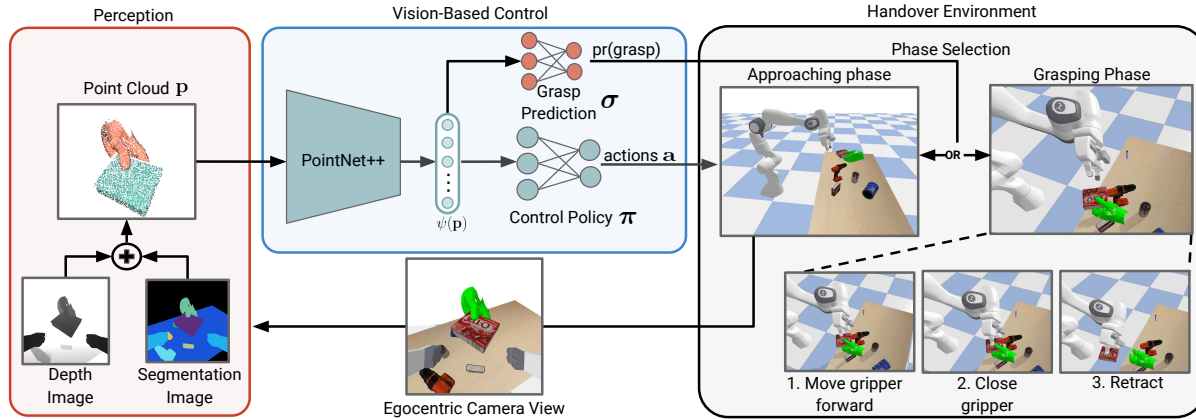


Figure 2. **Method Overview.** The **Perception** module takes egocentric RGB-D and segmentation images from the environment and outputs a hand/object segmented point cloud. Next, the segmented point cloud is passed to the **Vision-based Control** module and processed by PointNet++ [39] to obtain a lower-dimensional representation. This embedding is used as input to both the control policy and the grasp predictor. Each task episode in the **Handover Environment** follows two phases: during the approaching phase, the robot moves towards a pre-grasp pose, driven by the control policy π that outputs end-effector actions \mathbf{a} . A learned grasp predictor monitors the motion and determines when the robot should switch into the grasping phase, which follows the steps: 1. moving the gripper forward from a pre-grasp to a grasping pose 2. closing the gripper 3. retracting the object to a designated location, after which the episode ends.

man hand and objects are simulated by replaying data from the DexYCB dataset [7] (see Sec. 3.2). First, actions \mathbf{a} in the form of the next 6DoF end-effector pose (translation and rotation) are received from the policy $\pi(\mathbf{a}|\mathbf{s})$. We then convert the end-effector pose into a target robot configuration using inverse kinematics. Thereafter, we use PD-controllers to compute torques, which are applied to the robot. Finally, the visual information is rendered from the robot’s wrist-mounted RGB-D camera and sent to the perception module.

4.2. Perception

Our policy network takes a segmented hand and object point cloud as input. In the handover environment, we first render an egocentric RGB-D image from the wrist camera. Then we obtain the object point cloud \mathbf{p}_o and hand point cloud \mathbf{p}_h by overlaying the ground-truth segmentation mask with the RGB-D image. Since the hand and object may not always be visible from the current egocentric view, we keep track of the last available point clouds. The latest available point clouds are then sent to the control module.

4.3. Vision-Based Control

Input Representation Depending on the amount of points contained in the hand point cloud \mathbf{p}_h and object point cloud \mathbf{p}_o , we down- or upsample them into constant size. Next, we concatenate the two point clouds into a single point cloud \mathbf{p} and add two one-hot-encoded vectors to indicate the locations of object and hand points within \mathbf{p} . We then encode the point cloud into a lower dimensional representation $\psi(\mathbf{p})$ by passing it through PointNet++ [39]. Finally, the lower dimensional encoding $\psi(\mathbf{p})$ is passed on to the control policy π and the grasp prediction network σ .

Control Policy The policy network $\pi(\mathbf{a}|\psi(\mathbf{p}))$ is a small, two-layered MLP that takes the PointNet++ embedding as input state ($\mathbf{s} = \psi(\mathbf{p})$) and predicts actions \mathbf{a} that correspond to the change in 6DoF end-effector pose. These are passed on to the handover environment.

Grasp Prediction We introduce a grasp prediction network $\sigma(\psi(\mathbf{p}))$ that predicts when the robot should switch from approaching to executing the grasping motion (cf. Fig. 2). We model grasp prediction as a binary classification task. The input corresponds to the PointNet++ embedding $\psi(\mathbf{p})$, which is fed through a 3-layered MLP. The output is a probability that indicates the likelihood of a successful grasp given the current point cloud feature. If the probability is above a tunable threshold, we execute an open-loop grasping motion. The model is trained offline with pre-grasp poses attained from [15]. We augment the dataset by adding random noise to pre-grasp poses. To determine the labels, we initialize the robot with the pre-grasp poses in the physics simulation and execute the forward grasping motion. The label is one if the grasp is successful, and zero otherwise. We use a binary cross-entropy loss for training.

4.4. Two-Stage Teacher-Student Training

We aim at training a handover policy capable of moving simultaneously with the human. Training this policy directly in the setting of dynamic motion is challenging because expert demonstrations with open-loop planners to guide training can only be obtained when the human is stationary. A key contribution of our work is a two-stage training scheme for handovers that incrementally trains the policy to alleviate this challenge. In the first stage, we pretrain

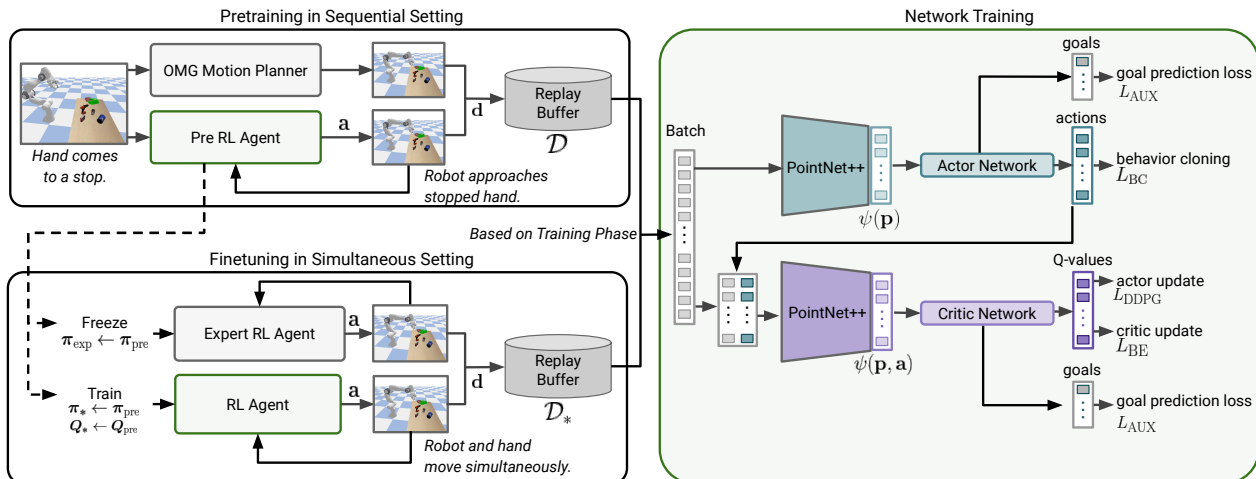


Figure 3. **Training Procedure.** In the **pretraining stage** (top left box), the human hand is stationary. We alternate between collecting expert demonstrations via motion planning and exploration data with the RL policy π_{pre} . Transitions \mathbf{d} are stored in a replay buffer \mathcal{D} . During training (green box, right), a batch of randomly sampled transitions from the replay buffer is passed through PointNet++ and the actor and critic networks. In the **finetuning stage** (bottom left box), the human and robot move concurrently. The expert motion planner is replaced by the expert policy π_{exp} , which shares the weights of the pretrained policy π_{pre} . This policy network will be kept frozen for the rest of training and serves as a regularizer for the RL agent. The RL agent’s actor network π_* and critic network Q_* are also initialized with the weights of pretrained agent’s networks, but the model will be updated during finetuning. In this stage, transitions are stored in a new replay buffer \mathcal{D}_* . Data is sampled solely from this buffer during finetuning.

in a setting where the robot only starts moving once the human has stopped (**sequential**). This pretrained policy is further finetuned in the second stage in which the human and robot move simultaneously (**simultaneous**).

Pretraining in Sequential Setting In the sequential setting, the robot starts moving once the human has come to a stop (see Fig. 3, top left). To grasp the object from the stationary human hand, we leverage motion planning to provide expert demonstrations. During data collection, we alternate between motion planning and RL-based exploration. In both cases, we store the transitions $\mathbf{d}_t = \{\mathbf{p}_t, \mathbf{a}_t, \mathbf{g}_t, \mathbf{r}_t, \mathbf{p}_{t+1}, \mathbf{e}_t\}$ in a replay buffer \mathcal{D} , from which we sample during network training. The term \mathbf{p}_t and \mathbf{p}_{t+1} indicate the point cloud and the next point cloud, \mathbf{a}_t the action, \mathbf{g}_t the pre-grasp goal pose, \mathbf{r}_t the reward, and \mathbf{e}_t an indicator of whether the transition is from the expert.

Inspired by [50], we collect expert trajectories with the OMG planner [49] that leverages ground-truth states. Note that some expert trajectories generated by the planner result in collision with the hand, which is why we introduce an offline pre-filtering scheme. We first parse the ACRONYM dataset [14] for potential grasps. We then run collision checking to filter out grasps where the robot and human hand collide. For the set of remaining collision-free grasps, we plan trajectories to grasp the object and execute them in open-loop fashion. On the other hand, the RL policy π_{pre} explores the environment and receives a sparse reward, i.e., the reward is one if the task is completed successfully, oth-

erwise zero. Hence, collisions with the human will get implicitly penalized by not receiving any positive reward.

Finetuning in Simultaneous Setting In this setting, the human and robot move at the same time. Hence, we cannot rely on motion and grasp planning to guide the policy. On the other hand, simply taking the pre-trained policy π_{pre} from the sequential setting and continue training it without an expert leads to an immediate drop in performance. Hence, we introduce a self-supervision scheme for stability reasons, i.e., we want to keep the finetuning policy close to the pre-trained policy. To this end, we replace the expert planner from the sequential setting by an expert policy π_{exp} , which is initialized with the weights of the pre-trained policy π_{pre} that already provides a reasonable prior policy (see Fig. 3 bottom left). Therefore, we have two policies: i) the expert policy π_{exp} as proxy for the motion and grasping planner. We freeze the network weights of this policy, ii) the finetuning policy π_* and critic Q_* , which are initialized with the weights of the pre-trained policy π_{pre} and critic Q_{pre} , respectively. We proceed to train these two networks using the loss functions which we describe next.

Network Training During training, we sample a batch of random transitions from the replay buffer \mathcal{D} . The policy network is trained using a combination of behavior cloning, RL-based losses and an auxiliary objective. In particular, the policy is updated using the following loss function:

$$L(\theta) = \lambda L_{\text{BC}} + (1 - \lambda) L_{\text{DDPG}} + L_{\text{AUX}}, \quad (3)$$

		success (%)	mean accum time (s)			failure (%)			
			exec	plan	total	contact	drop	timeout	total
Sequential	OMG Planner [49] †	62.50	8.309	1.414	9.722	27.78	8.33	1.39	37.50
	Yang et al. [52] †	64.58	4.864	0.036	4.900	17.36	11.81	6.25	35.42
	GA-DDPG [50]	50.00	7.139	0.142	7.281	4.86	19.44	25.69	50.00
	GA-DDPG [50] finetuned	57.18	6.324	0.086	6.411	6.48	27.08	9.26	42.82
	Ours	75.23	7.743	0.177	7.922	9.26	13.43	2.08	24.77
Simult.	GA-DDPG [50]	36.81	4.664	0.132	4.796	9.03	25.00	29.17	63.19
	GA-DDPG [50] finetuned	54.86	4.832	0.082	4.914	6.71	26.39	12.04	45.14
	Ours	68.75	6.232	0.178	6.411	8.80	17.82	4.63	31.25

Table 1. **HandoverSim Benchmark Evaluation.** Comparison of our method against various baselines from the HandoverSim benchmark [6]. In the sequential setting, we find that our baseline achieves better overall success rates than the baselines. In the simultaneous setting, we outperform the applicable baselines by large margins. The results for our method are averaged across 3 random seeds. †: both methods [49, 52] are evaluated with ground-truth states in [6] and thus are not directly comparable with ours.

where L_{BC} is a behavior cloning loss that keeps the policy close to the expert policy, L_{DDPG} is the standard actor-critic loss described in Eq. 2, and L_{AUX} is an auxiliary objective that predicts the grasping goal pose of the end-effector. The coefficient λ balances the behavior cloning and the RL objective. The critic loss is defined as:

$$L(\phi) = L_{BE} + L_{AUX}, \quad (4)$$

where L_{BE} indicates the Bellman error from Eq. 1 and L_{AUX} is the same auxiliary loss used in Eq. 3. We refer the reader to supplementary material or [50] for more details.

5. Experiments

We first evaluate our approach in simulation using the HandoverSim benchmark (Sec. 5.1). Next, we investigate the performance of sim-to-sim transfer by evaluating the trained models on the test environments powered by a different physics engine (Sec. 5.2). Finally, we apply the trained model to a real-world robotic system and analyze the performance of sim-to-real transfer (Sec. 5.3).

5.1. Simulation Evaluation

Setup HandoverSim [6] contains 1,000 unique H2R handover scenes divided into train, val, and test splits. Each scene contains a unique human handover motion. We evaluate on the “s0” setup which contains 720 training and 144 testing scenes. See the supp. material for evaluations on unseen objects, subjects, and handedness. Following the evaluation of GA-DDPG [50] in [6], we consider two settings: (1) the “sequential” setting where the robot is allowed to move only after the human hand reaches the handover location and remains static there (i.e., “hold” in [6]), and (2) the “simultaneous” setting where the robot is allowed to move from the beginning of the episode (i.e., “w/o hold” in [6]).

Metrics We follow the evaluation protocol in HandoverSim [6]. A handover is considered successful if the robot grasps the object from the human hand and moves it to a designated location. A failure is claimed and the episode

is terminated if any of the following three conditions occur: (1) the robot collides with the hand (*contact*), (2) the robot drops the object (*drop*), or (3) a maximum time limit is reached (*timeout*). Besides efficacy, the benchmark also reports efficiency in time. The time metric is further broken down into (1) the execution time (*exec*), i.e., the time to physically move the robot, and (2) the planning time (*plan*), i.e., the time spent on running the policy. All reported metrics are averaged over the rollouts on the test scenes.

Baselines Our primary baseline is GA-DDPG [50]. Besides comparing with the original model (i.e., trained in [50] for table-top grasping and evaluated in [6]), we additionally compare with a variant finetuned on HandoverSim (“GA-DDPG [50] finetuned”). For completeness, we also include two other baselines from [6]: “OMG Planner [49]” and “Yang et al. [52]”. However, both of them are evaluated with ground-truth state input in [6] and thus are not directly comparable with our method.

Results Tab. 1 reports the evaluation results on the test scenes. In the sequential setting, our method significantly outperforms all the baselines in terms of success rate, even compared to methods that use state-based input. Our method is slightly slower on average than GA-DDPG in terms of total time needed for handovers. In the simultaneous setting, our method clearly outperforms GA-DDPG, which has low success rates. Qualitatively, we observe that GA-DDPG directly tries to grasp the object from the user while it is still moving, while our method follows the hand and finds a feasible grasp once the hand has come to a stop, resulting in a trade-off on the overall execution time. We provide a qualitative example of this behavior in Fig. 4 (a) and in the supplementary video. We also refer to the supp. material for a discussion of limitations and a robustness analysis of our pipeline under noisy observations.

Ablations We evaluate our design choices in an ablation study and report the results in Tab. 2. We analyze the vision backbone by replacing PointNet++ with a ResNet18 [23]

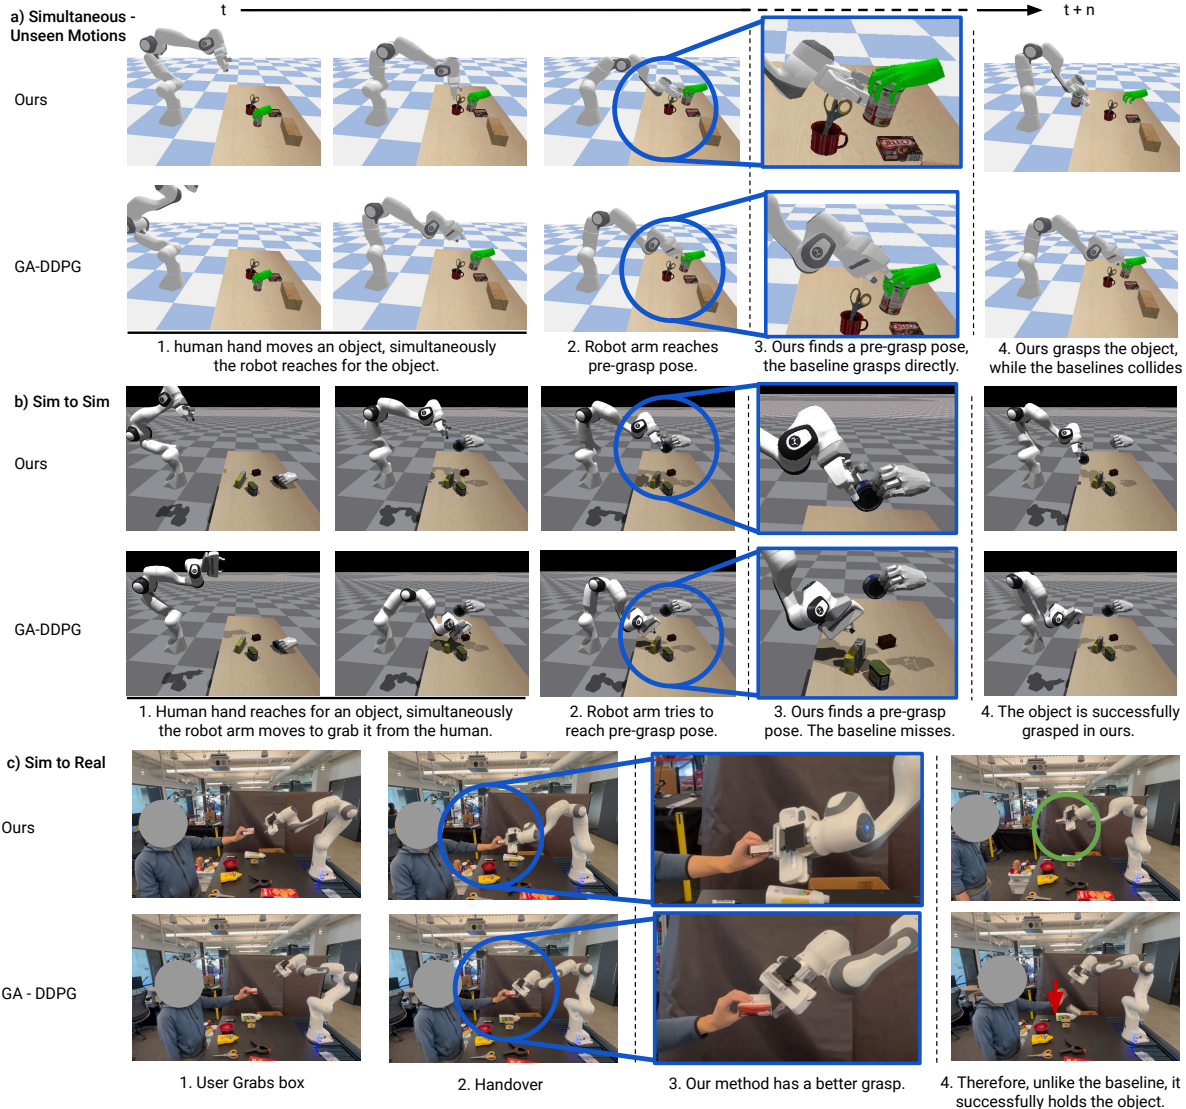


Figure 4. **Qualitative results.** We provide a comparison to show our methods’ advantages over GA-DDPG [50]. (a) Our method reacts to the moving human, while the baseline tries to go for a grasp directly, which leads to collision. (b) In the sim-to-sim transfer, we often find that the baseline does not find a grasp on the object. (c) In the sim-to-real experiment, GA-DDPG usually tries to get to a grasp directly, while our method adjusts the gripper into a stable grasping pose first. See the video in supp. material for more qualitative examples.

	Ablation Study			
	success (%)	failure (%)		
		contact	drop	timeout
w/ RGBDM + ResNet18	34.10	6.20	45.80	13.90
w/ third person view	60.42	9.95	25.69	3.94
w/o hand point cloud	59.03	24.07	11.58	5.32
w/o aux prediction	70.60	10.65	16.20	2.54
w/o standoff	52.55	7.87	36.80	2.78
w/o finetuning	73.38	9.03	13.89	3.70
Ours	75.23	9.26	13.43	2.08
w/o finetuning simult.	62.27	11.81	20.37	5.56
Ours simult.	68.75	8.8	17.82	4.63

Table 2. **Ablation.** We ablate the vision backbone, hand perception, and egocentric view. We also study the effect of finetuning, the auxiliary prediction, and splitting the task into two phases. All design choices are crucial aspects of our method with regards to overall performance. Results are averaged over 3 random seeds.

that processes the RGB and depth/segmentation (DM) images. Similar to the findings in GA-DDPG, the PointNet++ backbone performs better. Next, we train our method from third person view instead of egocentric view and without active hand segmentation (*w/o hand point cloud*), i.e., the policy only perceives the object point cloud but not the hand point cloud. We also ablate the auxiliary prediction (*w/o aux prediction*) and evaluate a variant that directly learns to approach and grasp the object instead of using the two task phases of approaching and grasping (*w/o standoff*). Lastly, we compare against our pretrained model, which was only trained in the sequential setting without finetuning (*w/o finetuning*). We find that the ablated components comprise important elements of our method. The results indicate an in-

		Sim-to-Sim			
		success (%)	failure (%)		
			contact	drop	timeout
Sequential	GA-DDPG [50]	19.44	4.86	47.22	28.47
	GA-DDPG [50] finetuned	11.81	6.25	68.75	13.19
	Ours	44.21	9.49	40.51	5.79
	Ours w/o grasp	54.40	7.87	33.34	4.40
Simult.	GA-DDPG [50]	11.11	15.97	48.61	24.31
	GA-DDPG [50] finetuned	16.67	9.72	63.89	9.72
	Ours	39.58	9.03	43.75	7.64
	Ours w/o grasp pred.	47.92	10.65	35.88	5.56

Table 3. **Sim-to-Sim Experiment.** We evaluate sim-to-sim transfer of the learning-based method to Isaac Gym [29]. Our method shows better transfer capabilities than GA-DDPG [50].

creased amount of hand collision or object drop in all ablations. A closer analysis in the simultaneous setting shows that our finetuned model outperforms the pretrained model.

5.2. Sim-to-Sim Transfer

Instead of directly transferring to the real world, we first evaluate the robustness of the models by transferring them to a different physics simulator. We re-implement the HandoverSim environment following the mechanism presented in [6] except for replacing the backend physics engine from Bullet [10] to Isaac Gym [29]. We then evaluate the models trained on the original Bullet-based environment on the test scenes powered by Isaac Gym. The results are presented in Tab. 3. We observe a significant drop for GA-DDPG on the success rates (i.e., to below 20%) in both settings. Qualitatively, we see that grasps are often either missed completely or only partially grasped (see Fig. 4 (b)). On the other hand, our method is able to retain higher success rates. Expectedly, it also suffers from a loss in performance. We analyze the influence of our grasp predictor on transfer performance and compare against a variant where we execute the grasping motion after a fixed amount of time (*Ours w/o grasp pred.*), which will leave the robot enough time to find a pre-grasp pose. Part of the performance drop is caused by the grasp predictor initiating the grasping phase at the wrong time, which can be improved upon in future work.

5.3. Sim-to-Real Transfer

Finally, we deploy the models trained in HandoverSim on a real robotic platform. We follow the perception pipeline used in [50, 52] to generate segmented hand and object point clouds for the policy, and use the output to update the end effector’s target position. We compare our method against GA-DDPG [50] with two sets of experiments: (1) a pilot study with controlled handover poses and (2) a user evaluation with free-form handovers. For experimental details and the full results, please see the supp. material.

Pilot Study We first conduct a pilot study with two subjects. The subjects are instructed to handover 10 objects from HandoverSim by grasping and presenting the objects in controlled poses. For each object, we test with 6 poses

	Subject 1		Subject 2	
	GA-DDPG [50]	Ours	GA-DDPG [50]	Ours
011_banana	3 / 6	6 / 6	6 / 6	5 / 6
037_scissors	2 / 6	5 / 6	3 / 6	5 / 6
006_mustard_bottle	1 / 6	3 / 6	2 / 6	4 / 6
024_bowl	3 / 6	4 / 6	3 / 6	3 / 6
040_large_marker	0 / 6	4 / 6	4 / 6	5 / 6
003_cracker_box	3 / 6	2 / 6	0 / 6	2 / 6
052_extra_large_clamp	1 / 6	4 / 6	5 / 6	5 / 6
008_pudding_box	3 / 6	6 / 6	4 / 6	4 / 6
010_potted_meat_can	2 / 6	2 / 6	3 / 6	4 / 6
021_bleach_cleanser	3 / 6	5 / 6	3 / 6	4 / 6
total	21 / 60	41 / 60	33 / 60	41 / 60

Table 4. **Sim-to-Real Experiment.** Success rates of the pilot study. Our method outperforms GA-DDPG [50] for both subjects.

(3 poses for each hand) with varying object orientation and varying amount of hand occlusion, resulting in 60 poses per subject. The same set of poses are used in testing both our model and GA-DDPG [50]. The success rates are shown Tab. 4. Results indicate that our method outperforms GA-DDPG [50] for both subjects on the overall success rate (i.e., 41/60 versus 21/60 for Subject 1). Qualitatively, we observe that GA-DDPG [50] tends to fail more from unstable grasping as well as hand collision. Fig. 4 (c) shows two examples of the real world handover trials.

User Evaluation We further recruited 6 users to compare the two methods and collected feedback from a questionnaire with Likert-scale and open-ended questions. In contrast to the pilot study, we asked the users to handover the 10 objects in ways that are most comfortable to them. We repeated the same experimental process for both methods, and counterbalanced the order to avoid bias. From participants’ feedback, the majority agreed that the timing of our method is more appropriate and our method can adjust between different object poses better. The interpretability of the robot’s motion was also acknowledged by their comments. Please see the supp. material for more details.

6. Conclusion

In this work, we have presented a learning-based framework for human-to-robot handovers from vision input with a simulated human-in-the-loop. We have introduced a two-stage teacher-student training procedure. In our experiments we have shown that our method outperforms baselines by a significant margin on the HandoverSim benchmark [6]. Furthermore, we have demonstrated that our approach is more robust when transferring to a different physics simulator and a real robotic system.

Acknowledgements We thank Tao Chen and Adithyavairavan Murali for laying the groundwork, Lirui Wang for the help with GA-DDPG, and Mert Albaba, Christoph Gebhardt, Thomas Langerak and Juan Zarate for their feedback on the manuscript.

References

- [1] Ireaiyo Akinola, Jingxi Xu, Shuran Song, and Peter K Allen. Dynamic grasping with reachability and motion awareness. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021. 3
- [2] Peter Anderson, Ayush Shrivastava, Joanne Truong, Arjun Majumdar, Devi Parikh, Dhruv Batra, and Stefan Lee. Sim-to-real transfer for vision-and-language navigation. In *Conference on Robot Learning (CoRL)*, 2021. 1
- [3] Antonio Bicchi and Vijay Kumar. Robotic grasping and contact: A review. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2000. 2
- [4] Jeannette Bohg, Antonio Morales, Tamim Asfour, and Danica Kragic. Data-driven grasp synthesis—a survey. *IEEE Transactions on Robotics (T-RO)*, 2013. 2
- [5] Samarth Brahmabhatt, Chengcheng Tang, Christopher D. Twigg, Charles C. Kemp, and James Hays. ContactPose: A dataset of grasps with object contact and hand pose. In *Proceedings of the European Conference on Computer Vision (ECCV)*, August 2020. 2
- [6] Yu-Wei Chao, Chris Paxton, Yu Xiang, Wei Yang, Balakumar Sundaralingam, Tao Chen, Adithyavairavan Murali, Maya Cakmak, and Dieter Fox. HandoverSim: A simulation framework and benchmark for human-to-robot object handovers. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2022. 2, 3, 6, 8
- [7] Yu-Wei Chao, Wei Yang, Yu Xiang, Pavlo Molchanov, Ankur Handa, Jonathan Tremblay, Yashraj S. Narang, Karl Van Wyk, Umar Iqbal, Stan Birchfield, Jan Kautz, and Dieter Fox. DexYCB: A benchmark for capturing hand grasping of objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 3, 4
- [8] Tao Chen, Jie Xu, and Pulkit Agrawal. A system for general in-hand object re-orientation. In *Conference on Robot Learning (CoRL)*, 2021. 3
- [9] Sammy Christen, Stefan Stevsic, and Otmar Hilliges. Demonstration-guided deep reinforcement learning of control policies for dexterous human-robot interaction. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2019. 2
- [10] Erwin Coumans and Yunfei Bai. PyBullet: a Python module for physics simulation for games, robotics and machine learning. <https://pybullet.org>, 2016–2021. 2, 8
- [11] Matt Deitke, Dhruv Batra, Yonatan Bisk, Tommaso Campari, Angel X Chang, Devendra Singh Chaplot, Changan Chen, Claudia Pérez-D’Arpino, Kiana Ehsani, Ali Farhadi, et al. Retrospectives on the embodied ai workshop. *arXiv preprint arXiv:2210.06849*, 2022. 1
- [12] Matt Deitke, Winson Han, Alvaro Herrasti, Aniruddha Kembhavi, Eric Kolve, Roozbeh Mottaghi, Jordi Salvador, Dustin Schwenk, Eli VanderBilt, Matthew Wallingford, Luca Weihs, Mark Yatskar, and Ali Farhadi. RoboTHOR: An open simulation-to-real embodied AI platform. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [13] Haonan Duan, Peng Wang, Yiming Li, Daheng Li, and Wei Wei. Learning human-to-robot dexterous handovers for anthropomorphic hand. *IEEE Transactions on Cognitive and Developmental Systems (TCDS)*, 2022. 2
- [14] Clemens Eppner, Arsalan Mousavian, and Dieter Fox. ACRONYM: A large-scale grasp dataset based on simulation. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021. 5
- [15] Clemens Eppner, Arsalan Mousavian, and Dieter Fox. A billion ways to grasp: An evaluation of grasp sampling schemes on a dense, physics-based grasp data set. In *Robotics Research: The 19th International Symposium (ISRR)*. Springer, 2022. 4
- [16] Zackory Erickson, Vamsee Gangaram, Ariel Kapusta, C. Karen Liu, and Charles C. Kemp. Assistive Gym: A physics simulation framework for assistive robotics. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2020. 2
- [17] Zicong Fan, Omid Taheri, Dimitrios Tzionas, Muhammed Kocabas, Manuel Kaufmann, Michael J. Black, and Otmar Hilliges. ARCTIC: A dataset for dexterous bimanual hand-object manipulation. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [18] Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning (ICML)*, 2018. 3
- [19] Chuang Gan, Jeremy Schwartz, Seth Alter, Damian Mrowca, Martin Schrimpf, James Traer, Julian De Freitas, Jonas Kubilius, Abhishek Bhandwaldar, Nick Haber, Megumi Sano, Kuno Kim, Elias Wang, Michael Lingelbach, Aidan Curtis, Kevin Feigelis, Daniel Bear, Dan Gutfreund, David Cox, Antonio Torralba, James J. DiCarlo, Josh Tenenbaum, Josh McDermott, and Dan Yamins. ThreeDWorld: A platform for interactive multi-modal physical simulation. In *NeurIPS Track on Datasets and Benchmarks*. 2021. 1
- [20] Guillermo Garcia-Hernando, Shanxin Yuan, Seungryul Baek, and Tae-Kyun Kim. First-person hand action benchmark with RGB-D videos and 3D hand pose annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [21] Shreyas Hampali, Mahdi Rad, Markus Oberweger, and Vincent Lepetit. HOnnotate: A method for 3D annotation of hand and object poses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [22] Yana Hasson, Gul Varol, Dimitrios Tzionas, Igor Kalevatykh, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning joint reconstruction of hands and manipulated objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [24] Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan

- Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, and Sergey Levine. QT-Opt: Scalable deep reinforcement learning for vision-based robotic manipulation. In *Conference on Robot Learning (CoRL)*, 2018. 3
- [25] Kilian Kleeberger, Richard Bormann, Werner Kraus, and Marco F. Huber. A survey on learning-based robotic grasping. *Current Robotics Reports*, 2020. 2
- [26] Kailin Li, Lixin Yang, Xinyu Zhan, Jun Lv, Wenqiang Xu, Jiefeng Li, and Cewu Lu. ArtiBoost: Boosting articulated 3D hand-object pose estimation via online exploration and synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [27] Shaowei Liu, Hanwen Jiang, Jiarui Xu, Sifei Liu, and Xiaolong Wang. Semi-supervised 3D hand-object poses estimation with interactions in time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [28] Yunze Liu, Yun Liu, Che Jiang, Kangbo Lyu, Weikang Wan, Hao Shen, Boqiang Liang, Zhoujie Fu, He Wang, and Li Yi. HOI4D: A 4D egocentric dataset for category-level human-object interaction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [29] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. Isaac Gym: High performance gpu-based physics simulation for robot learning. In *NeurIPS Track on Datasets and Benchmarks*. 2021. 2, 8
- [30] Naresh Marturi, Marek Kopicki, Alireza Rastegarpanah, Vijaykumar Rajasekaran, Maxime Adjigble, Rustam Stolkin, Aleš Leonardis, and Yasemin Bekiroglu. Dynamic grasp and trajectory planning for moving objects. *Autonomous Robots*, 2019. 2
- [31] Andrew T Miller and Peter K Allen. Graspit! a versatile simulator for robotic grasping. *IEEE Robotics & Automation Magazine (RAM)*, 2004. 2
- [32] Gyeongsik Moon, Shou-I Yu, He Wen, Takaaki Shiratori, and Kyoung Mu Lee. InterHand2.6M: A dataset and baseline for 3D interacting hand pose estimation from a single RGB image. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [33] Douglas Morrison, Peter Corke, and Jürgen Leitner. Closing the loop for robotic grasping: A real-time, generative grasp synthesis approach. In *Proceedings of Robotics: Science and Systems (RSS)*, 2018. 2
- [34] Arsalan Mousavian, Clemens Eppner, and Dieter Fox. 6-DOF GraspNet: Variational grasp generation for object manipulation. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 2
- [35] Valerio Ortenzi, Akansel Cosgun, Tommaso Pardi, Wesley P. Chan, Elizabeth Croft, and Dana Kulić. Object handovers: A review for robotics. *IEEE Transactions on Robotics (T-RO)*, 2021. 1
- [36] Yik Lung Pang, Alessio Xompero, Changjae Oh, and Andrea Cavallaro. Towards safe human-to-robot handovers of unknown containers. In *IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*, 2021. 2
- [37] Claudia Pérez-D’Arpino, Can Liu, Patrick Goebel, Roberto Martín-Martín, and Silvio Savarese. Robot navigation in constrained pedestrian environments using reinforcement learning. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021. 2
- [38] Xavier Puig, Tianmin Shu, Shuang Li, Zilin Wang, Yuan-Hong Liao, Joshua B. Tenenbaum, Sanja Fidler, and Antonio Torralba. Watch-And-Help: A challenge for social perception and human-AI collaboration. In *International Conference on Learning Representations (ICLR)*, 2021. 2
- [39] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 4
- [40] Patrick Rosenberger, Akansel Cosgun, Rhys Newbury, Jun Kwan, Valerio Ortenzi, Peter Corke, and Manfred Grafinger. Object-independent human-to-robot handovers using real time robotic vision. *IEEE Robotics and Automation Letters (RA-L)*, 2021. 2
- [41] Ricardo Sanchez-Matilla, Konstantinos Chatzilygeroudis, Apostolos Modas, Nuno Ferreira Duarte, Alessio Xompero, Pascal Frossard, Aude Billard, and Andrea Cavallaro. Benchmark for human-to-robot handovers of unseen containers with unknown filling. *IEEE Robotics and Automation Letters (RA-L)*, 2020. 2
- [42] Bokui Shen, Fei Xia, Chengshu Li, Roberto Martín-Martín, Linxi Fan, Guanzhi Wang, Claudia Pérez-D’Arpino, Shyamal Buch, Sanjana Srivastava, Lyne Tchampi, Micael Tchampi, Kent Vainio, Josiah Wong, Li Fei-Fei, and Silvio Savarese. iGibson 1.0: A simulation environment for interactive tasks in large realistic scenes. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021. 1
- [43] Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. ALFRED: A benchmark for interpreting grounded instructions for everyday tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [44] Shuran Song, Andy Zeng, Johnny Lee, and Thomas Funkhouser. Grasping in the wild: Learning 6DoF closed-loop grasping from low-cost demonstrations. *IEEE Robotics and Automation Letters (RA-L)*, 2020. 3
- [45] Sanjana Srivastava, Chengshu Li, Michael Lingelbach, Roberto Martín-Martín, Fei Xia, Kent Vainio, Zheng Lian, Cem Gokmen, Shyamal Buch, C. Karen Liu, Silvio Savarese, Hyowon Gweon, Jiajun Wu, and Li Fei-Fei. BEHAVIOR: Benchmark for everyday household activities in virtual, interactive, and ecological environments. In *Conference on Robot Learning (CoRL)*, 2021. 1
- [46] Andrew Szot, Alexander Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah Maestre, Mustafa Mukadam, Devendra Singh Chaplot, Oleksandr Maksymets, Aaron Gokaslan, Vladimír Vondruš, Sameer Dharur, Franziska Meier, Wojciech Galuba, Angel Chang, Zsolt

- Kira, Vladlen Koltun, Jitendra Malik, Manolis Savva, and Dhruv Batra. Habitat 2.0: Training home assistants to rearrange their habitat. In *Advances in Neural Information Processing Systems (NeurIPS)*. 2021. 1
- [47] Omid Taheri, Nima Ghorbani, Michael J. Black, and Dimitrios Tzionas. GRAB: A dataset of whole-body human grasping of objects. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [48] Chen Wang, Claudia Pérez-D’Arpino, Danfei Xu, Li Fei-Fei, C. Karen Liu, and Silvio Savarese. Learning diverse strategies for human-robot collaboration. In *Conference on Robot Learning (CoRL)*, 2021. 2
- [49] Lirui Wang, Yu Xiang, and Dieter Fox. Manipulation trajectory optimization with online grasp synthesis and selection. In *Proceedings of Robotics: Science and Systems (RSS)*, 2020. 2, 3, 5, 6
- [50] Lirui Wang, Yu Xiang, Wei Yang, Arsalan Mousavian, and Dieter Fox. Goal-auxiliary actor-critic for 6D robotic grasping with point clouds. In *Conference on Robot Learning (CoRL)*, 2021. 2, 3, 5, 6, 7, 8
- [51] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, Li Yi, Angel X. Chang, Leonidas J. Guibas, and Hao Su. SAPIEN: A SimulATED Part-based Interactive Environment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [52] Wei Yang, Chris Paxton, Arsalan Mousavian, Yu-Wei Chao, Maya Cakmak, and Dieter Fox. Reactive human-to-robot handovers of arbitrary objects. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021. 2, 3, 6, 8
- [53] Wei Yang, Balakumar Sundaralingam, Chris Paxton, Iretiayo Akinola, Yu-Wei Chao, Maya Cakmak, and Dieter Fox. Model predictive control for fluid human-to-robot handovers. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2022. 2
- [54] Ruolin Ye, Wenqiang Xu, Zhendong Xue, Tutian Tang, Yanfeng Wang, and Cewu Lu. H2O: A benchmark for visual human-human object handover analysis. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [55] R. Ye, W. Xu, Z. Xue, T. Tang, Y. Wang, and C. Lu. H2o: A benchmark for visual human-human object handover analysis. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2