

Zero-Shot Transfer of Haptics-Based Object Insertion Policies

Samarth Brahmhatt¹, Ankur Deka¹, Andrew Spielberg², and Matthias Müller¹




Fig. 1: Top-left to bottom-right: sequential frames from a plate insertion episode. The robot is controlled by our plate insertion policy trained entirely in simulation. The policy successfully deals with multiple collisions caused by noise in the target slot location, and an unseen object (a cup) blocking it, and inserts the plate into the neighbouring free slot. Zoom in to see collision details.

Abstract—Humans naturally exploit haptic feedback during contact-rich tasks like loading a dishwasher or stocking a bookshelf. Current robotic systems focus on avoiding unexpected contact, often relying on strategically placed environment sensors. Recently, contact-exploiting manipulation policies have been trained in simulation and deployed on real robots. However, they require some form of real-world adaptation to bridge the sim-to-real gap, which might not be feasible in all scenarios. In this paper we train a contact-exploiting manipulation policy in simulation for the contact-rich household task of loading plates into a slotted holder, which transfers without any fine-tuning to the real robot. We investigate various factors necessary for this zero-shot transfer, like time delay modeling, memory representation, and domain randomization. Our policy transfers with minimal sim-to-real gap and significantly outperforms heuristic and learnt baselines. It also generalizes well to a cup and plates of different sizes and weights. The project website is <https://sites.google.com/view/compliant-object-insertion>.

I. INTRODUCTION

Common household tasks like loading a dishwasher, stocking a bookshelf, or inserting a pot into a coffee maker involve

¹ S. Brahmhatt, A. Deka, and M. Müller are with Intel Labs. E-mail: {samarth.manoj.brahmhatt, ankur.deka, matthias.mueller}@intel.com.

² A. Spielberg is with the Massachusetts Institute of Technology and contributed mostly during an internship at Intel Labs. E-mail: aespielberg@csail.mit.edu.

object placement in constrained locations. Humans are able to do these almost effortlessly. For example, when inserting a dish into an empty dishwasher slot, a single glance is enough to approximate the scene geometry and the target pose of the plate. Humans then move the plate towards the slot, handle unexpected collisions along the way, and detect success by leveraging experiential knowledge of how the haptics should feel when the plate is firmly in the slot. Importantly, they do this without access to accurate pose estimation and tracking, and often without live visual feedback. Studies in neurophysics provide evidence that specialized regions in the brain form a “haptic working memory” from a visual snapshot and use it to plan motions [1]–[3].

Robots need to perform similar object placement tasks, as they are increasingly collaborating with humans in dynamic, unseen, and un-instrumented human environments. However, several challenges currently limit this. The lack of accuracy of or prior access to object and environment 3D models leads to uncertainty in the task state. Unexpected collisions between the object and environment surfaces are therefore very likely to occur, making this a contact-rich task. Contact-exploiting robot motion policies are likely to perform better than contact-avoiding policies in these cases.

However, existing learnt contact-exploiting insertion policies require either real-world expert demonstrations [12], [13], robot training from scratch in the real-world [19],

	[6], [10]	[23]	[18]	[19]	[15], [20]	[21]	[12]	[22]	[16]	ours
Insertion task	peg	peg, USB	peg, clip	plug, USB	plug, gear	pins	peg	boxes	pegs	plates
Learnable policy	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Generalizes to different geometries	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Avoids need for unoccluded vision	✓	✓	✗	✗	✓	✓	✓	✓*	✗	✓
Avoids need for expert demonstrations	✓	✗	✗	✓	✓	✗	✗	✓	✓	✓
Avoids need for real world pre-training	✓	✗	✓	✓	✓	✓	✓	✓	✗	✓
Avoids need for real world training / finetuning	✗	✓	✗	✗	✗	✗	✗	✗	✗	✓

TABLE I: Comparison of insertion tasks, policy formulations, and training requirements of the proposed algorithm with related work. *: [22] requires a GelSlim planar tactile sensor.

[22], or real-world adaptation *i.e.* data collection by the robot in the target environment before the policy is fully trained [15], [20]. Because this is expensive and potentially unsafe, policies that are trained completely in simulation are desirable. The use of high-dimensional and difficult-to-simulate signals like live images for contact and environment state feedback can also contribute to the need for real-world training or adaptation [16], [18].

In this paper we focus on a common household task: plate insertion into a slotted plate holder (see Figure 1) when the slot location is noisy, and unseen objects block the target slot. Differently from other works, we investigate how a successful real-world contact-exploiting insertion policy can be trained completely in simulation. The policy is trained using reinforcement learning (RL) considering the delayed-reward nature of the task. The predicted action is passed as a relative 6-DOF motion target to a low-level operational space controller [32]. The observation and reward design encourage contact-awareness under state uncertainty, while the combination of impedance built into the operational space controller and a soft gripper on the real robot result in compliant policy execution. We do not assume access to strategically placed cameras providing live occlusion-free visual feedback. Instead, our policy relies only an approximate target location sensed once *before* the interaction begins. During the interaction, proprioceptive and low-dimensional haptics *i.e.* end-effector wrench features are used to construct policy observations. We also do not assume access to accurate 3D models of the plate or the slotted plate holder.

Real robot experiments demonstrate that our policy successfully inserts the plate into the slot in spite of multiple unexpected collisions, outperforms baselines, and generalizes to plates of different sizes and masses, and even a cup. Through an ablation study, we find that proper memory representation in the policy observation, and incorporating delay between observation capture time and action execution time, significantly increases the transferability of the policy to the real robot. To the best of our knowledge, this is the first non-hardcoded insertion policy trained completely in simulation *i.e.* without pre-training, training, or finetuning with real-world data or demonstrations. To summarize, our contributions are:

- a haptics-based policy for contact-aware and compliant insertion of plates into a slotted plate holder,
- demonstration of zero-shot policy transfer to a real

robot, comparison to baselines, generalization to insertion of different plates and even a cup, and

- an ablation study validating design choices and analyzing their impact on sim-to-real transfer.

II. RELATED WORK

Object insertion has been widely studied in robotics in the context of automated industrial part assembly through peg-in-hole insertion [4], [5] problems. Approaches range from handcrafted impedance control [6]–[8], force control [9], [10], and vision-impedance hybrid control [11] strategies to learning from human demonstration [12], [13] and reinforcement learning (RL) with visual, force, and proprioceptive features [14]–[16]. Recently, RL has also been used for learning key [17], clip [18], USB and electrical plug [19], [20] insertion policies. We summarize the drawbacks of existing approaches in Table I and elaborate below.

Current RL-based approaches require some form of training or finetuning with the real robot in the target environment e.g. [16] train a multimodal representation from real-world data collected with a heuristic policy, [18] and [21] collect human demonstrations, and [22], [19], and [15] train the RL policy on the real robot. [20] reduce the amount of real robot training through meta-RL. However, real-world training is expensive to monitor, potentially unsafe, and requires access to the target environment prior to deployment. In contrast, we learn the policy only in simulation and transfer it directly to the real robot.

Another common requirement is live visual feedback *e.g.* [16] and [19] use image observations from a camera which is situated such that it can observe the entire interaction without robot body or gripper occlusion. This might not be feasible for a robot operating in an uninstrumented environment. [23] develop wrist camera image-based oriented keypoint detectors for generating object pose observations, which are used to train an imitation learning peg-hole insertion policy [11]. In contrast, we rely exclusively on end-effector wrench and proprioceptive features as policy observations. Vision is used to sense the approximate target location only once before the interaction begins, when occlusion is unlikely. Our work is complementary to visual policy methods because our conclusions about increasing the zero-shot transferability of haptics-based policies can be used to augment these methods in cases where strategic camera placement in the environment is acceptable or the grasped object can be tracked under occlusion [24], [25].




Fig. 3: Policy execution architecture.

Our real robot setup, inspired from household robotics, is also different: the target slot is already occupied by an unseen blocking object (see Figure 4 and Section IV), so the policy has to find an empty slot. In contrast, existing approaches *e.g.* [12], [13], [15], [16], [19], [20] focus on inserting the object in a single empty slot.

III. METHOD

We assume that the object has already been grasped before the policy execution begins. The policy is divided into two stages: an open-loop planner from MoveIt [26], [27] is first used to move the end-effector 12 cm above its approximate target pose (see Section IV for details). We assume no obstacles are present in this path. The policy then switches to the second stage *i.e.* a learnt controller that handles the contact-rich segment of the task. This is the main contribution of this paper. We describe its details below.

Problem formulation: We model the task as a finite-horizon Markov Decision Process (MDP) with a state space \mathcal{S} , continuous action space \mathcal{A} , transition dynamics $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$, discount factor $\gamma \in [0, 1)$, and reward function $R(s_t, \mathbf{a}_t, \mathbf{a}_{t-1})$. Given a horizon T , we aim to learn a policy $\pi : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})$ mapping observations to action probabilities $P \in \mathcal{P}$ that maximize the discounted return $J(\pi) = \mathbb{E}_\pi \left[\sum_{t=0}^T \gamma^t R(s_t, \mathbf{a}_t, \mathbf{a}_{t-1}) \right]$. \mathcal{P} is the family of unimodal Gaussian PDFs. π is parameterized with a 2-layer (256, 256) MLP and is learnt using the Soft Actor-Critic RL algorithm [28] containing a 2-layer (256, 256) MLP critic.

Observation: In the following, aT_b denotes the 6-DOF pose of coordinate system of body b w.r.t. coordinate system of body a . The robot base, target, grasped object, and end-effector are denoted by b , t , o , and e respectively, and the $\bar{}$ superscript indicates that the body is in its final intended state. The end-effector pose that would place the grasped object at the intended target is thus ${}^bT_{\bar{e}} = {}^bT_t {}^tT_{\bar{o}} {}^oT_e$. However, as mentioned in Section I, ${}^bT_{\bar{e}}$ and the target pose bT_t are known only approximately in practice. We model this uncertainty by adding 6-DOF noise sampled from a uniform distribution $\mathcal{U}[-\epsilon, \epsilon]$ to the perfect bT_t from simulation, denoting the noisy version by $\tilde{}$: ${}^bT_{\bar{e}} = {}^bT_{\tilde{t}} {}^tT_{\bar{o}} {}^oT_e$. ${}^tT_{\bar{o}}$ and oT_e are assumed constant, even though multiple contacts during the episode keep changing oT_e . This results from the requirement of no object tracking, and means the policy should be robust to in-hand object motion. The observation (see Figure 3) contains the current end-effector pose w.r.t. the noisy end-effector target pose ${}^eT_{\bar{e}} = {}^bT_{\bar{e}}^{-1} {}^bT_e$ (represented with 3 numbers for translation and 3 for axis-angle rotation)

and the wrench \mathcal{W} exerted by the end-effector, expressed in the b frame. Gravity, coriolis, and acceleration components are removed from this wrench (see supplementary for details), enabling it to capture object weight and contact forces exclusively. Thus the observation depends on the object mass. **History:** Several relevant quantities like end-effector velocity, contact history, current contact duration etc. require some form of memory. Hence, we stack eT_e and \mathcal{W} for $H = 8$ steps in the policy observation to increase observability of the true state [29]. We evaluate other memory representations like recurrent policies [30] in Section IV. See supp. for architecture details.

Action: The policy outputs an action vector $\mathbf{a}_t \in [-1, 1]^6$. The first 3 elements are scaled by 45 cm to get a translation vector, while the last 3 elements are scaled by 45° to get an axis-angle rotation vector. The resulting 6-DOF pose is the residual motion target for the end-effector w.r.t. its current pose. Like [31], it is composed with a similarly scaled base motion target that moves the end-effector straight to ${}^bT_{\bar{e}}$.

Low-level controller: As shown in Figure 3, the composite motion target is passed to a low-level operational space controller [32] running at 1 kHz. A new state vector is captured from the robot after every 50 low-level control steps. The operational space formulation and the end-effector wrench formulation described above decouple the simulated robot dynamics from the real robot dynamics. Because robot manufacturers often have closed-source dynamics APIs and simulators only approximate the dynamics parameters, this is important to reduce the simulation-reality gap. We do not randomize robot dynamics.

Reward: The reward function $R(s_t, \mathbf{a}_t, \mathbf{a}_{t-1})$ sums the following components: per-step time penalty $R_{time} = -1/T$, object drop penalty (also ends the episode) $R_{drop} = -1.1$, success reward (also ends the episode) $R_{success} = +0.5$, distance to true target penalty $R_{dist} = -\sum_{m \in \{trans, rot\}} K_{dist}^{(m)} \min(\lambda_{dist}^{(m)}, \|{}^eT_e^{(m)}\|)$, and action change penalty for smoothness $R_{\Delta a} = -\sum_{m \in \{trans, rot\}} K_{\Delta a}^{(m)} \min(\lambda_{\Delta a}^{(m)}, \|{}^{a_{t-1}}T_{a_t}^{(m)}\|)$. $K_{dist}^{(m)}$ and $K_{\Delta a}^{(m)}$ are scale factors, while $\lambda_{dist}^{(m)}$ and $\lambda_{\Delta a}^{(m)}$ are cutoff values. In contrast to some other works like [16], [19], we do not need a staged reward function for different phases.

Encouraging sim-to-real transfer: bT_t noise: In initial experiments without noise added to bT_t we found that the policy overfits by remembering the obstacle geometry and learning a ${}^bT_{\bar{e}}$ trajectory that avoids all collisions. However, such a policy fails on the real robot because the observations are noisy and collisions become inevitable. Adding 6-DOF noise $\mathcal{U}[-\epsilon, \epsilon]$ to bT_t as described above forces collisions, because the true target pose is no longer observable by the policy. This allows us to achieve the objective of learning collision-friendly policies. We increase the bT_t noise parameter ϵ in a piecewise linear curriculum. Partial insertion initialization: But this also significantly reduces the probability of the object reaching its intended ${}^bT_{\bar{e}}$ target pose with purely random exploration. Hence, following reverse curriculum genera-

tion [33], we initialize the robot in 50% of the episodes in poses such that the object is already partially inserted. **Policy inference delay:** In a simulator, time stops during policy inference and round-trip communication of state and action from the environment to the policy. This is not true for the real world. Hence we randomize this delay in the range of [7, 13] (based on delay in the real robot setup) low-level control steps and found this to be significant for sim-to-real transfer (see Table II).

Policy learning: Contact simulation and the randomizations described above increase the simulation and environment reset complexity. Hence we learn the policy with the off-policy Soft Actor Critic (SAC) [28] algorithm, because it is more sample efficient than on-policy algorithms [34]. We seed the replay buffer with 500 episodes of transitions collected with a random policy to encourage exploration. ‘Success’ is detected when the object pose is below a height threshold and the tabletop in-plane coordinates (X and Y) are within $d_{success}$ distance of any slot center. To dampen in-place oscillations once the plate is inserted in the slot, we increased the number of policy steps in the ‘success’ state required to end the episode from 1 to 10, and $d_{success}$ to cover the entire slot width.

Implementation details: We use the simulation environment and operational space controller implementations in Robosuite [35], which uses the MuJoCo physics engine [36]. Training is performed using tf-agents [37] and Reverb [38]. While the Franka Emika Panda robot [39] is used for both simulation and real-world experiments, the former use the Panda rigid gripper and the latter use a Soft Robotics Inc. mGrip gripper [40] (see Fig. 4). We found that in practice the mGrip’s high friction and balance of firmness and compliance allows it to hold thin objects like plates much better than the Panda rigid gripper. As our experiments validate, the policy trained in simulation transfers well to the real world in spite of this difference. Modeling the mGrip gripper in simulation for potentially more accurate wrench observations is out of the scope of this paper. See the supplementary material for values of all hyperparameters mentioned in this section. We trained the learnable policies on a cluster with Intel(R) Xeon(R) Platinum 8180 CPU @ 2.50 GHz, 112 cores and 1 TB RAM. Training one policy took 14 hours.

IV. EXPERIMENTS

Real robot experiments evaluation protocol: The episode starts with the end-effector at a random pose in a cuboidal region of dimensions $20 \times 40 \times 15$ cm above the central slot of the plate holder. We assume that the object has been grasped before the episode starts. In simulation, this is done by applying gravity compensation to the object held at the appropriate graspable pose, while in the real world an operator holds it between the gripper fingers while the grip closes. The operator holds an ArUco [41], [42] marker above the target slot, estimating the end effector location for the plate to be fully inserted. The marker pose is detected by the shoulder camera and treated as the noisy target end effector

pose w.r.t. robot base bT_e used to calculate observations (see Section III). An episode succeeds when the plate is inserted in any slot (determined by thresholding the end-effector height and medial axis error w.r.t. bT_e) and is kept there for 3 seconds. We release the grasp and raise the arm upon detecting success. It fails when the plate is kept stationary at a location other than the slots (including jamming with high force), or the robot hits torque or joint angle limits. The policy trained in simulation is deployed directly on the real robot without adaptation. As seen in Figure 4, the plate holder has 6 slots. We perform 4 trials for each slot and compute the success rate separately across slots for the first, second, third, and fourth trial to get 4 success rate values. We report the mean and standard deviation of these for each method.




Fig. 4: (a): Real robot policy execution setup. (b): Detailed view of the main plate, plate holder slots, and the unseen blocking objects - cup and charger.

Blocking objects: As explained in Section III, noise is added in the observations during simulation training to force the policy to learn collision-friendly behaviours. However, the real plate holder slots are much more ergonomically designed than the cuboidal geometries used in simulation (see Figure 4). Even though the clearance between the plate and slot widths is 2.5 mm in the real setup, a policy that centers above the slot and moves downward without drift is likely to succeed trivially. Hence, to rigorously test the learnt policies’ collision behaviour we place a blocking object e.g. a cup or a phone charger in the target slot. This tests the policy’s ability to handle unexpected collisions, and better reflects real-world application conditions like a mobile manipulator attempting to load a partially loaded dishwasher.

Differences from simulation: Figure 5 shows how the simulated plate and slotted holder geometries are very rough approximations of their real counterparts. Other approximations include the ideal plate pose in its target state w.r.t. the slotted holder tT_o , the pose of the end-effector w.r.t. the grasped plate oT_e , soft gripper simulated as a rigid gripper with spring-loaded joints, and blocking the target slot with a cup or charger. Our proposed algorithm yields high-performance policies despite these large differences.

Baselines: We compare our policy to the following baselines also implemented at the low-level through the same operational space controller as ours:

- `straight-down`: The end-effector moves down-




Fig. 5: Characteristics of the simulated plate used for training, and plates and cup used for real robot experiments, showing the highly approximate nature of simulation geometry. The ablation study (Figure 8) is done with the “regular” plate. Note also the gripper differences.

wards indefinitely, implemented by setting the operational space controller’s target to $[0, 0, -0.05]$ cm constantly.

- `random-search` [43]: The end-effector moves to a point randomly sampled from a horizontal square of side length 5 cm, and then downwards till a contact force of 3 N is experienced. If the episode does not succeed in this pose, it moves upwards, samples another point from the square, and repeats.
- `[16] no vision`: While the full vision-based policy from [16] is not comparable to this work, its “no vision” variant provides a useful baseline for network architecture. It has separate branches for proprioception and wrench observations, and uses temporal convolution to represent history. We train the policy neural network with all the randomizations described in Section III directly on our RL objective instead of the auxiliary pre-training objectives from [16] to avoid training with the real robot. The original architecture has a large number of learnable parameters and does not perform well in simulation ($17.0 \pm 13.1\%$ success rate). Hence we compare to a version with smaller 16-d wrench and proprioception feature vectors.

Comparison to baselines: Table II compares our full policy with the three baseline policies in terms of real-world success rate. Our full policy significantly outperforms them and exhibits the ability to deal with multiple collisions and finish the task (see Figures 1, 6). Qualitatively, the `straight-down` baseline is sensitive to noise in the target pose, while the `random-search` baseline is sensitive to yaw error in the target pose, which causes the plate to contact the slot sides when it is only partially inserted, triggering an upward motion. The `no vision` policy [16] performs better in simulation with smaller feature vectors but overfits, because the real robot performance is significantly lower. Qualitatively, it often keeps retrying at the same location after recovering from collisions.

Generalization: All policies are trained in simulation with a perfectly cylindrical plate, while the real robot experiments are done with three plates of varying sizes, shapes, and weights, all different from the simulated plate (see Figure 5).

Method	Sim. success rate (%) mean over 4 trials \pm std. dev.	Real success rate (%) mean over 4 trials \pm std. dev.
<code>straight-down</code>	40.0 ± 15.5	33.3 ± 27.2
<code>[16] no-vision</code>	89.0 ± 13.4	41.7 ± 9.62
<code>random-search</code>	50.0 ± 19.5	45.8 ± 8.33
ours	84.0 ± 15.0	83.3 ± 13.6

TABLE II: Plate insertion performance in simulation and real-world. Our full algorithm significantly outperforms other methods in real-world performance, and has the smallest sim-to-real gap.

As shown in Figure 7, our policy consistently outperforms baselines across plates. It also generalizes with almost the same performance from simulation to the “regular” and “medium” plates. We also used the same policy to insert the “small” plate and the cup shown in Figure 5. Even though their size and geometry differs significantly from the training simulated plate, we found that our policy performs reasonably well after adding a constant offset to the Z translation component of eT_e observations to compensate for the height difference. Figure 6 shows a multiple object insertion application with our policy, where the cup is inserted first and the plate, targeted to the slot covered by the cup, is then inserted in the next free slot.

Ablation study: Figure 8 shows the simulation and real robot performance of various ablated versions of our full policies. All policies are evaluated under the same ‘full policy’ conditions with the “regular” plate i.e. full-scale bT_t noise in observations and non-zero policy inference delay. History representation: The experiments with the real robot show that observation stacking, incorporating recurrent units in the policy network, and including gripper velocity w.r.t. goal in observations all improve true state observability compared to no history, because they result in higher success rate. Observation stacking performs best. However, the policies with $H = 16$ and explicit gripper velocity significantly underperform our main policy ($H = 8$), probably because a larger observation vector linearly increases the number of learnable parameters. Sim-to-real transfer: Training with a history of stacked observations, bT_t noise in observations, and a non-zero policy inference delay are crucial design choices for ensuring the policy’s transferability to the real robot. Training: In addition (not shown in Fig. 8), the policy completely failed in simulation training without partial insertion initialization and the residual action formulation.

Qualitative behaviour: The policy trained without the policy inference delay has very limited responsiveness to collisions, while that trained without observation history drifts and tends to converge to a set of fixed joint positions. Finally, the policy trained without noise cannot differentiate between collisions with slot tops and bottoms and hence lifts the plate out even after successfully inserting it sometimes. While our full policy achieves almost identical performance in simulation and on the real robot, the sim-to-real gap increases significantly for the ablated models, highlighting the importance of each component. Multiple policy executions are shown in




Fig. 6: Left to right: Sequential video frames showing multiple object insertion and the generalization ability of our policy. Our policy, which is trained in simulation to insert a plate (Figure 5) into a slot whose location is only approximately known, is deployed on the real robot without any real-world adaptation. (a)-(d): It first inserts a cup, a completely unseen object. (e): When success is detected, the grasp is released, arm raised, and the “regular” plate is grasped. The same policy is commanded to insert the plate in the slot blocked by the cup. (f)-(h): It inserts the plate in the next free slot. Both episodes exhibit the collision-exploiting behaviour learnt by our policy. Zoom in to see collision details. Full video in the supplementary material.




Fig. 7: **Generalization:** Real robot performance for different plates shown in Figure 5 (error bars show 20% std. dev.). Simulation performance is also shown for reference. Our algorithm consistently outperforms baselines.



Fig. 8: **Ablation study:** Impact of various design choices. Bar height indicates real robot success rate, while simulation success rates are mentioned inside the bars (error bars show 20% std. dev.). Our full algorithm, which represents memory with $H = 8$ stacked observations, performs best. The importance of training with SAC, policy inference delay, and bT_t noise in observations is also shown.




Fig. 9: Left to right: Sequential video frames showing a failure case of our policy where it gets jammed in a collision and cannot recover. Zoom in to see collision details. Full video in the supplementary material.

the supplementary video.

Limitations: Currently the policy does not perform well when there is a large horizontal gap between the end-effector pose at the start of stage 2 and the target pose bT_e , and when the base motor is rotated far from its center, even though it is trained in simulation under these conditions. We hypothesize that this is caused by imperfectly identified inertial parameters used in the operational space controller.

V. CONCLUSION

We presented a robot learning system for contact-rich object insertion that uses end-effector wrench and proprioception feedback, and a noisy insertion target pose. In contrast to other works, the policy is trained exclusively in simulation and transfers directly to the real robot. We demonstrate the capability of our approach on the task of inserting a plate into a slotted plate holder and generalizing to plates of different sizes and masses, through an ablation study and comparisons to other baseline methods. We further show that our policy can also insert a cup, an object with a completely different geometry than the training simulation plate. The ability of this system to deal with uncertainty in the insertion target pose can potentially be useful in mobile manipulation use-cases, where robot base motion is a significant source of manipulation uncertainty. We hope to integrate this system with a mobile manipulator in the future.

Acknowledgements: We would like to thank Benoit Landry for help with real robot controller implementation. We would also like to thank Vladen Koltun for important contributions to the problem statement definition, solution methodology, and feedback on real robot experiments.

REFERENCES

- [1] F. Maule, G. Barchiesi, T. Brochier, and L. Cattaneo, "Haptic working memory for grasping: the role of the parietal operculum," *Cerebral Cortex*, vol. 25, no. 2, pp. 528–537, 2015.
- [2] K. Sathian, "Analysis of haptic information in the cerebral cortex," *Journal of Neurophysiology*, vol. 116, no. 4, pp. 1795–1806, 2016.
- [3] A. Murata, V. Gallese, M. Kaseda, and H. Sakata, "Parietal neurons related to memory-guided hand manipulation," *Journal of Neurophysiology*, vol. 75, no. 5, pp. 2180–2186, 1996.
- [4] D. E. Whitney, "Historical perspective and state of the art in robot force control," *The International Journal of Robotics Research*, vol. 6, no. 1, pp. 3–14, 1987.
- [5] H. McCallion, G. Johnson, and D. Pham, "A compliant device for inserting a peg in a hole," *Industrial Robot: An International Journal*, 1979.
- [6] J. F. Broenink and M. L. Tiernego, "Peg-in-hole assembly using impedance control with a 6 dof robot," in *Proceedings of the 8th European Simulation Symposium*. Citeseer, 1996, pp. 504–508.
- [7] T. Tsumugiwa, R. Yokogawa, and K. Hara, "Variable impedance control with virtual stiffness for human-robot cooperative peg-in-hole task," in *IEEE/RSJ international conference on intelligent robots and systems*, vol. 2. IEEE, 2002, pp. 1075–1081.
- [8] H. Park, J. Park, D.-H. Lee, J.-H. Park, M.-H. Baeg, and J.-H. Bae, "Compliance-based robotic peg-in-hole assembly strategy without force feedback," *IEEE Transactions on Industrial Electronics*, vol. 64, no. 8, pp. 6299–6309, 2017.
- [9] D. Stokic, M. vnkobratovic, and D. Hristic, "Implementation of force feedback in manipulation robots," *The International journal of robotics research*, vol. 5, no. 1, pp. 66–76, 1986.
- [10] L. Johannsmeier, M. Gerchow, and S. Haddadin, "A framework for robot manipulation: Skill formalism, meta learning and adaptive control," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 5844–5850.
- [11] G. Morel, E. Malis, and S. Boudet, "Impedance based combination of visual and force control," in *Proceedings. 1998 IEEE International Conference on Robotics and Automation (Cat. No.98CH36146)*, vol. 2, 1998, pp. 1743–1748 vol.2.
- [12] T. Tang, H.-C. Lin, and M. Tomizuka, "A learning-based framework for robot peg-hole-insertion," in *Dynamic Systems and Control Conference*, vol. 57250. American Society of Mechanical Engineers, 2015, p. V002T27A002.
- [13] S. Gubbi, S. Kolathaya, and B. Amrutur, "Imitation learning for high precision peg-in-hole tasks," in *2020 6th International Conference on Control, Automation and Robotics (ICCAR)*, 2020, pp. 368–372.
- [14] T. Inoue, G. De Magistris, A. Munawar, T. Yokoya, and R. Tachibana, "Deep reinforcement learning for high precision assembly tasks," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 819–825.
- [15] C. C. Beltran-Hernandez, D. Petit, I. G. Ramirez-Alpizar, and K. Harada, "Variable compliance control for robotic peg-in-hole assembly: A deep-reinforcement-learning approach," *Applied Sciences*, vol. 10, no. 19, p. 6923, 2020.
- [16] M. A. Lee, Y. Zhu, P. Zachares, M. Tan, K. Srinivasan, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg, "Making sense of vision and touch: Learning multimodal representations for contact-rich tasks," *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 582–596, 2020.
- [17] F. Voigt, L. Johannsmeier, and S. Haddadin, "Multi-level structure vs. end-to-end-learning in high-performance tactile robotic manipulation," in *Proceedings of the 2020 Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, J. Kober, F. Ramos, and C. Tomlin, Eds., vol. 155. PMLR, 16–18 Nov 2021, pp. 2306–2316. [Online]. Available: <https://proceedings.mlr.press/v155/voigt21a.html>
- [18] M. Vecerik, O. Sushkov, D. Barker, T. Rothörl, T. Hester, and J. Scholz, "A practical approach to insertion with variable socket position using deep reinforcement learning," in *2019 international conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 754–760.
- [19] G. Schoettler, A. Nair, J. Luo, S. Bahl, J. A. Ojea, E. Solowjow, and S. Levine, "Deep reinforcement learning for industrial insertion tasks with visual inputs and natural rewards," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 5548–5555.
- [20] G. Schoettler, A. Nair, J. A. Ojea, S. Levine, and E. Solowjow, "Meta-reinforcement learning for robotic industrial insertion tasks," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 9728–9735.
- [21] Y. Ma, D. Xu, and F. Qin, "Efficient insertion control for precision assembly based on demonstration learning and reinforcement learning," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 7, pp. 4492–4502, 2020.
- [22] S. Dong, D. K. Jha, D. Romeres, S. Kim, D. Nikovski, and A. Rodriguez, "Tactile-rl for insertion: Generalization to objects of unknown geometry," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 6437–6443.
- [23] W. Gao and R. Tedrake, "kpm 2.0: Feedback control for category-level robotic manipulation," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2962–2969, 2021.
- [24] S. Brahmhbhatt, H. B. Amor, and H. Christensen, "Occlusion-aware object localization, segmentation and pose estimation," in *Proceedings of the British Machine Vision Conference (BMVC)*, X. Xie, M. W. Jones, and G. K. L. Tam, Eds. BMVA Press, September 2015, pp. 80.1–80.13. [Online]. Available: <https://dx.doi.org/10.5244/C.29.80>
- [25] P. Grady, C. Tang, C. D. Twigg, M. Vo, S. Brahmhbhatt, and C. C. Kemp, "ContactOpt: Optimizing contact to improve grasps," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [26] D. Coleman, I. Sukan, S. Chitta, and N. Correll, "Reducing the barrier to entry of complex robotic software: a moveit! case study," *arXiv preprint arXiv:1404.3785*, 2014.
- [27] I. A. Sukan and S. Chitta, "MoveIt," <https://moveit.ros.org>, 2014, [Online; accessed 2022-06-11].
- [28] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, and S. Levine, "Soft actor-critic algorithms and applications," *CoRR*, vol. abs/1812.05905, 2018. [Online]. Available: <http://arxiv.org/abs/1812.05905>
- [29] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al., "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [30] M. Hausknecht and P. Stone, "Deep recurrent Q-learning for partially observable MDPs," in *2015 AAAI fall symposium series*, 2015.
- [31] T. Johannink, S. Bahl, A. Nair, J. Luo, A. Kumar, M. Loskyll, J. A. Ojea, E. Solowjow, and S. Levine, "Residual reinforcement learning for robot control," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 6023–6029.
- [32] O. Khatib, "A unified approach for motion and force control of robot manipulators: The operational space formulation," *IEEE Journal on Robotics and Automation*, vol. 3, no. 1, pp. 43–53, 1987.
- [33] C. Florensa, D. Held, M. Wulfmeier, M. Zhang, and P. Abbeel, "Reverse curriculum generation for reinforcement learning," in *Conference on robot learning*. PMLR, 2017, pp. 482–495.
- [34] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [35] Y. Zhu, J. Wong, A. Mandlekar, and R. Martín-Martín, "robosuite: A modular simulation framework and benchmark for robot learning," in *arXiv preprint arXiv:2009.12293*, 2020.
- [36] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 5026–5033.
- [37] S. Guadarrama, A. Korattikara, O. Ramirez, P. Castro, E. Holly, S. Fishman, K. Wang, E. Gonina, N. Wu, E. Kokiopoulou, L. Sbaiz, J. Smith, G. Bartók, J. Berent, C. Harris, V. Vanhoucke, and E. Brevdo, "TF-Agents: A library for reinforcement learning in tensorflow," <https://github.com/tensorflow/agents>, 2018, [Online; accessed 2022-06-12].
- [38] A. Cassirer, G. Barth-Maron, E. Brevdo, S. Ramos, T. Boyd, T. Sottiaux, and M. Kroiss, "Reverb: A framework for experience replay," 2021.
- [39] Franka Emika, "Panda research v2," <https://www.franka.de/research,2022>, [Online; accessed 2022-06-11].
- [40] Soft Robotics Inc., "mGrip soft gripper and controller system," <https://www.softroboticsinc.com/products/mgrip-modular-gripping-solution-for-food-automation>, 2022, [Online; accessed 2022-06-11].
- [41] F. J. Romero-Ramirez, R. Muñoz-Salinas, and R. Medina-Carnicer, "Speeded up detection of squared fiducial markers," *Image and Vision Computing*, vol. 76, pp. 38–47, 2018.
- [42] S. Garrido-Jurado, R. Muñoz-Salinas, F. J. Madrid-Cuevas, and R. Medina-Carnicer, "Generation of fiducial marker dictionaries using

mixed integer linear programming,” *Pattern recognition*, vol. 51, pp. 481–491, 2016.

[43] J. A. Marvel, R. Bostelman, and J. Falco, “Multi-robot assembly

strategies and metrics,” *ACM Computing Surveys (CSUR)*, vol. 51, no. 1, pp. 1–32, 2018.