

# Minimizing Human Assistance: Augmenting a Single Demonstration for Deep Reinforcement Learning

Abraham George<sup>1</sup>, Alison Bartsch<sup>1</sup>, and Amir Barati Farimani<sup>1</sup>

**Abstract**—The use of human demonstrations in reinforcement learning has proven to significantly improve agent performance. However, any requirement for a human to manually ‘teach’ the model is somewhat antithetical to the goals of reinforcement learning. This paper attempts to minimize human involvement in the learning process while retaining the performance advantages by using a single human example collected through a simple-to-use virtual reality simulation to assist with RL training. Our method augments a single demonstration to generate numerous human-like demonstrations that, when combined with Deep Deterministic Policy Gradients and Hindsight Experience Replay (DDPG + HER) significantly improve training time on simple tasks and allows the agent to solve a complex task (block stacking) that DDPG + HER alone cannot solve. The model achieves this significant training advantage using a single human example, requiring less than a minute of human input. Moreover, despite learning from a human example, the agent is not constrained to human-level performance, often learning a policy that is significantly different from the human demonstration.

## I. INTRODUCTION

Reinforcement learning (RL) is widely used to allow robots to autonomously learn how to complete tasks, such as grasping and manipulation [1]–[6], walking [7]–[10], aerial robotic flight [11], [12], and many others. However, the training process is often both time and resource intensive, even for relatively simple tasks. A recent survey paper highlighting the challenges in the field of deep reinforcement learning (DRL), [13] noted sample efficiency, reliable and stable learning, and side-stepping exploration challenges all as current bottlenecks. As such, finding ways to improve the stability and learning speed of reinforcement learning models is vital to their widespread use. One such way is to provide human examples of the task being completed to help guide the reinforcement learning algorithm [14]. Although human examples have proven to be effective at increasing the training speed and accuracy of DRL algorithms [14]–[17], they are somewhat counter-productive, as one of the major advantages of DRL is that it can complete tasks without requiring an expert to spend time implementing a controller, or generating numerous demonstrations. Additionally, training human experts to use demonstration methods that provide rich information, such as teleoperation, can be time-consuming. However, this challenge can be addressed with virtual reality based teleoperation, because in the virtual space a human can intuitively perform actions without much difficulty, while still generating rich trajectory information.

<sup>1</sup>With the Department of Mechanical Engineering, Carnegie Mellon University {aigeorge, abartsch, afarimani}@andrew.cmu.edu

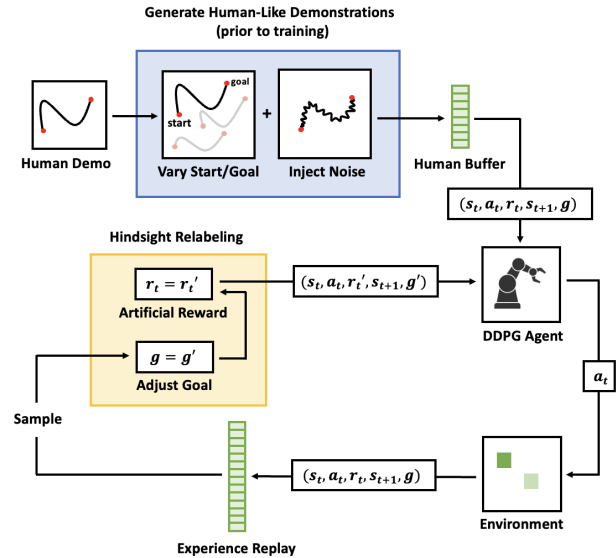


Fig. 1. A diagram outlining the components of our method. It consists of a standard DDPG + HER architecture that is augmented with a secondary buffer populated with a variety of human-like trajectories, which are generated from a single collected human demonstration.

In this paper, we propose a middle ground solution to these challenges: a method that uses a single human example, collected through a simple-to-use virtual reality (VR) simulation, to assist with DRL training. By limiting the human input as much as possible, this paper endeavors to gain significant training advantages while requiring less than a minute of human input. The key aspects of our method are

- A demonstration trajectory collection method that is easy and intuitive to use, leveraging VR technology.
- A trajectory generation method that augments the single human demonstration by varying the start and goal positions and injecting noise. It is used to populate the demonstration buffer with many varying demonstration trajectories created from a single human example.
- A DDPG + HER architecture which is combined with the demonstration buffer to train the agent with no additional reward engineering necessary.

A schematic diagram outlining this architecture can be seen in Fig. 1. The primary technical contribution of this work is the development of an augmentation method that, given a single human demonstration, can generate a diverse set of trajectories to train a DRL agent. Our system converges to successful policies on standard manipulation tasks faster than the baseline of Deep Deterministic Policy Gradients

(DDPG) [18] with Hindsight Experience Replay (HER) [19], while requiring minimal additional effort to collect a single demonstration. Additionally, we demonstrate that our method can achieve success on more complicated and longer time-horizon tasks, such as block stacking, that standard DDPG + HER is unable to learn.

## II. RELATED WORKS

1) *Imitation Learning*: Imitation learning uses demonstrations to determine an agent’s actions by having the agent replicate expert examples. This task can be accomplished through machine learning using a supervised learning approach such as classification [20]. These forms of behavior cloning have proven effective on complex tasks such as biped locomotion [21], but they require large data sets and do not function well in environments outside of the range of examples they trained on. Data set Aggregation (DAGger) addresses some of these issues by augmenting the learned policy with expert demonstrations collected throughout the training process, interconnecting the learned and expert policies [22]. However, DAGger can be difficult to implement because it requires the use of expert examples collected throughout the duration of training.

2) *Human Demonstrations in Reinforcement Learning*: Human demonstrations have proven a powerful aid in augmenting reinforcement learning. By combining human demonstrations with reinforcement learning, classical control tasks such as the pendulum swing up task can be learned [23]. Human demonstrations have been combined with Deep Q-learning to significantly accelerate learning, even with relatively small amounts of data [17]. This Deep Q-learning approach has also been extended to the continuous action space, where demonstrations have been shown to aid DDPG in learning to drive a car in a simulated environment [16]. Additionally, demonstration learning has proven effective in assisting with robotic sparse reward tasks, such as peg insertion [15]. [24] found moderate success pre-training with non-expert human demonstrations, first the authors used supervised learning to learn task features from human demonstrations, then combined this pre-trained feature network with asynchronous advantage actor-critic to learn a policy for the Atari domain. Notably, this method still improved over the baseline with low-quality demonstrations. In [25], researchers present an algorithm that transfers 20 human demonstrations into a baseline policy that is refined with reinforcement learning, which produces small but statistically significant improvement in performance. They found the best performance when combining the best demonstrations from two teachers, highlighting the need for high quality demonstrations for this technique. However, neither [24] nor [25] significantly improved over the baselines, while both required multiple human demonstrations. A different vein of work focuses on using human preferences and feedback to shape learned policies. [26] presents a Bayesian approach to maximize the information gain from the human feedback, converting this feedback into a policy, and combining this policy with a RL policy. In [27], the goal of the agent

is defined through human preferences between trajectories. They demonstrate the method was able to solve a suite of simulated robotics tasks without a reward function. However, each task required an hour of human time.

In this work, our experiments closely match those done by [14], where a data set of 100 human demonstrations was used to significantly increase the performance of deep reinforcement learning agents on robot manipulation tasks, including pick-and-place and block stacking. The human examples were added to the replay buffer, the loss function was adjusted with a behavioral cloning component to incentivize the agent to match the human demonstrations, and the agent was initialized in demonstration states to assist with the long-horizon sparse reward tasks, in this case block stacking. The primary limitation of [14] is the necessity of collecting 100 human demonstrations for each task. Our work seeks to replicate the successes while only using a single human example and not requiring any modifications to the loss function.

## III. METHODS

### A. Human Example Collection and Augmentation

A Meta Oculus Quest 2 virtual reality (VR) headset was used to collect human examples of the reinforcement learning (RL) tasks being completed. The VR environment shows the user a first person view of the task they needed to complete. For example, when demonstrating the pick-and-place task, the user is shown the block that needs to be manipulated, along with a highlighted goal region where the block is to be placed. A view of this environment can be found in Fig. 2. To manipulate the blocks, the user controls a Franka Emika Panda [28] end effector, which is mapped onto the right hand Oculus controller, and whose gripper is opened and closed with the controller joy-stick. The VR environment records the user’s actions when completing the desired task, creating a trajectory file that can then be used to generate a replay buffer of human actions.

Although only a single human example is sampled, a wide range of demonstrations with different start and goal conditions, along with different approaches to solving the task, is ideal for training the reinforcement learning algorithm. To help increase the diversity of our human samples, a method to generate human-like examples from the single recorded human example was developed. First, random object start and object goal locations are chosen. The recorded trajectory is then linearly scaled such that the recorded start and goal locations match the new locations. For the stack task, which has multiple object start and goal locations, the trajectory is split into two sub-trajectories (move block one, then move block two) and each sub-trajectory is scaled independently. The robot then traverses the task environment using the generated trajectory and a proportional controller, and the resultant actions, rewards, and states are recorded. Although this scaling allows for the generation of new trajectories, it does not add much variety. To decrease the trajectory similarity, Ornstein–Uhlenbeck noise [29] is added to the robot’s actions. The augmented trajectories that successfully

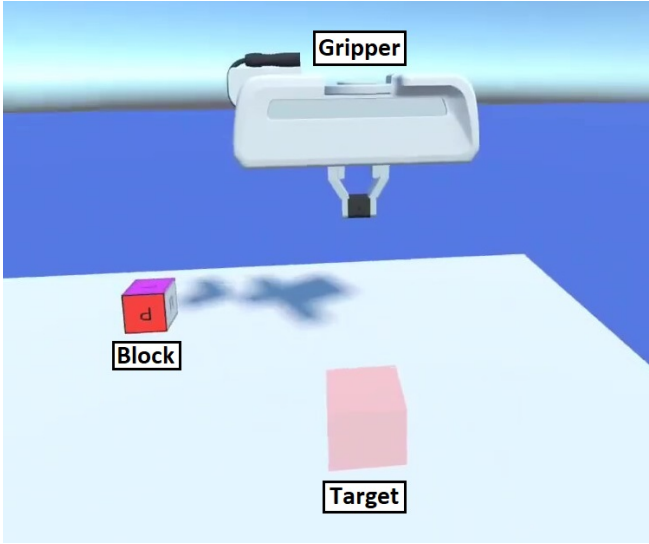


Fig. 2. View from the Oculus VR headset of the pick-and-place task. The user moves the gripper with a hand-held controller, and is tasked with moving the block to the target location.

complete the task are used to populate a secondary replay buffer (referred to as the human buffer) which can then be used to train the reinforcement learning algorithm. The process for generating human-like actions is shown below in equations 1-4.

$$a = \begin{cases} \frac{p_{gen,goal} - p_{gen,start}}{p_{rec,goal} - p_{rec,start}}, & \text{if } p_{rec,goal} \neq p_{rec,start} \\ 1, & \text{otherwise} \end{cases} \quad (1)$$

$$b = p_{gen,start} - a \cdot p_{rec,start} \quad (2)$$

$$p_{gen,i} = a \cdot p_{rec,i} + b \quad (3)$$

$$action_{p,i} = K(p_{gen,i} - p_{measured,i}) + N_{OU} \quad (4)$$

Where  $p$  is a trajectory coordinate (the process is run for  $p = x, y,$  and  $z$ ),  $rec$  refers to the recorded trajectory,  $gen$  refers to the trajectory that is to be generated,  $goal$  refers to the goal block location,  $start$  refers to the start block location,  $i$  is the time step,  $K$  is the proportionality constant for the controller,  $p_{measured,i}$  is the measured location of the end effector at time step  $i$ , and  $N_{OU}$  is Ornstein-Uhlenbeck noise.

### B. Learning Architecture

Our human aided reinforcement learning method follows the general DDPG + HER DRL structure, except with a secondary replay buffer storing the augmented human demonstrations. During training, 25% of the transitions are sampled from this secondary human buffer, while the remaining 75% are sampled from the standard experience replay buffer. This method, which is similar to the replay buffer spiking introduced by Lipton et. al, [30] exposes the RL agent to the human example without tying it to the human

developed policy. Although previous work has shown that directly linking human demonstrations to the training of the agent through methods such as using behavior cloning loss can be effective, [14] we fear that the lack of diversity in our human samples may cause this type of approach to lead to overfitting.

Additionally, a demonstration based pre-play method was implemented for use in tasks with a long time horizon and a sparse reward, in this case block stacking. This type of task is particularly challenging because the agent needs to learn a more complex sequence of actions before it receives a reward. The pre-play method addresses this issue by initializing the training environment using actions based on the human demonstration. In 50% of the simulation episodes, the initial conditions of the environment are used to generate a human-like trajectory, which the robot follows for the first  $n$  time steps, where  $n$  is randomly generated from a uniform distribution between 0 and the total number of simulation steps. After the  $n$  steps, the robot's actions are determined by the DDPG agent. This approach provides the agent with simplified versions of the task by having the demonstration based pre-play partially solve the task. By exposing the agent to the reward early in the training process, we hope this method will function similarly to curriculum learning and help improve the agent's performance. Additionally, since a method to generate human-like examples for a given start and goal location has been developed, the human-like trajectory used to determine the action policy before step  $n$  can be tailored to the specific environment conditions, allowing this method to be used when training on hardware.

## IV. EXPERIMENTAL EVALUATION

### A. Simulation Environment

The RL model was evaluated on three Panda-Gym tasks created by Gallouedec et al.: push, pick-and-place, and stack, as shown in Fig. 3 [31]. The panda-gym tasks, which run on the PyBullet physics engine, feature a 7-degree-of-freedom Franka Emika Panda robot arm which is used to manipulate a 4 cm cube block. During the RL experiments, the orientation of the gripper is locked, making the end effector a 4-degree-of-freedom system (position and gripper width). A sparse reward was used during training, returning a reward of  $-1$  when the task is not complete and  $0$  when the task is complete. Task completion is determined by taking the distance from the current cube location (opaque cube in Fig. 3) to the goal cube location (transparent cube in Fig. 3), and checking if that distance is less than the cutoff distance (5 cm in the push and pick-and-place tasks and 4 cm for both cubes in the stack task). Two modifications were made to Gallouedec et al.'s Panda-Gym tasks. In the push task, the fingers were unlocked so that they could be used to manipulate the block, and in the stack task, the completion criteria was changed from both cubes having less than 5cm of error to both cubes having less than 4cm of error to remove the possibility of false positive rewards. To evaluate our method's performance, all experiments were conducted across five or more random seeds.




Fig. 3. Panda-Gym environments used for RL model evaluation. The block(s) to be stacked are shown as opaque cubes, while the goal locations are shown as transparent cubes.

From the Panda-Gym simulation, the DRL agent observes the following information: the end effector position and velocity, the gripper width, and the block(s) position, rotation, and velocity (both translational and rotational). Given this state information and the goal, the DDPG agent selects an action (the output of the actor network). This action has 4 parameters, the  $x,y,z$  displacement to apply to the end effector and the displacement to apply to the width between the gripper fingers. When training the DDPG agent, each epoch included 1280 simulations (50 time-steps each for push and pick-and-place, and 100 time-steps for stack) and 3200 network updates (batch size of 256).

## B. Results

1) *Effect of the Human Buffer:* To examine the effectiveness of the human-aided model, we compared the success rate over epoch of the baseline DDPG + HER with our human-aided DDPG + HER implementation for the push and the pick-and-place tasks. The results of these tests can be seen in Fig. 4 and 6. Our experiments show that the addition of the human buffer led to a significant increase in learning speed, especially in the more complex pick-and-place task. We found that regardless of whether the human buffer was used, the agent learned the same policy. For example, in the push task both the baseline DDPG + HER agent and the DDPG + HER agent with a human buffer learned to close the gripper, place it on top of the block, and drag it to the goal location. The agent with a human buffer learned the same policy even though the provided human example used a significantly different approach: grasp the block using the gripper, pick it up, and then place it down in the goal location (see Fig. 5). With this disparity in learning technique in mind, we repeated the pick-and-place experiment, but replaced the original human example with an example where the human drags the block using a closed gripper (the learned policy). When shown this demonstration, the DDPG + HER with human buffer agent significantly increased its performance (see Fig. 4). From these experiments, we draw the conclusion that the human buffer provides a significant assistance to DDPG + HER, even when the human policy is non-ideal.

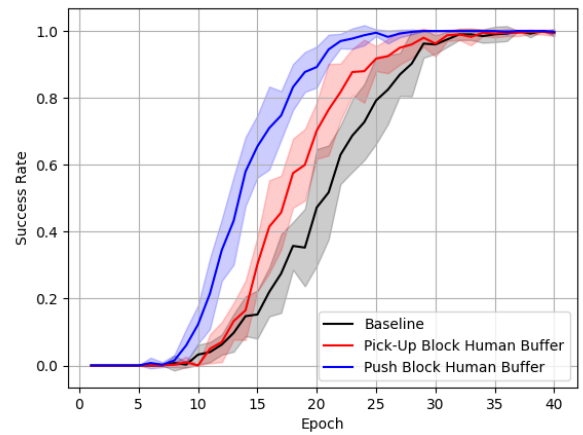


Fig. 4. Success rate on the push task for baseline DDPG + HER, DDPG + HER with the human buffer created from a pick-up demonstration, and DDPG + HER with the human buffer created from a block pushing demonstration. The shaded region shows one standard deviation from the mean. The introduction of the human buffer causes the model to train faster, even when the demonstration is significantly different from the converged policy.




Fig. 5. Push and Stack agent policies. Left - Push task: The human demonstrations (top three images) feature picking up and then placing the block, while the policy the agent converges upon (bottom three images) is to drag the block to the goal using the gripper. Right - Stack Task: The human example (top four images) first picks and places the bottom block, then picks and places the top block. The trained agent (bottom three images) picks up the top block, sets it onto the bottom block, then drags the stack to the goal location.

2) *Ablation Study:* To examine the impact of the human-generated buffer on the results, we performed an ablation study to analyze the key aspects of our approach: the example collection method, the augmentation step generating human-like samples, and the necessity of an initial demonstration. First, to test if the human buffer aided the agent’s learning because it consisted of transitions that were collected from a human demonstration, as opposed to those created by the RL agent, the pick-and-place experiment was re-run using a trained DDPG + HER agent to create the ‘human’ example. As shown in Fig. 6, this method performed similarly to its human counterpart, showing that the example trajectory does not need to be a human demonstration. Next, to test if the augmentation step of generating human-like samples was necessary, the pick-and-place task was run with a human buffer that only contained the original recorded trajectory. This method performed about the same as baseline DDPG + HER, showing that the augmentation step in our method is vital to the agent’s performance.

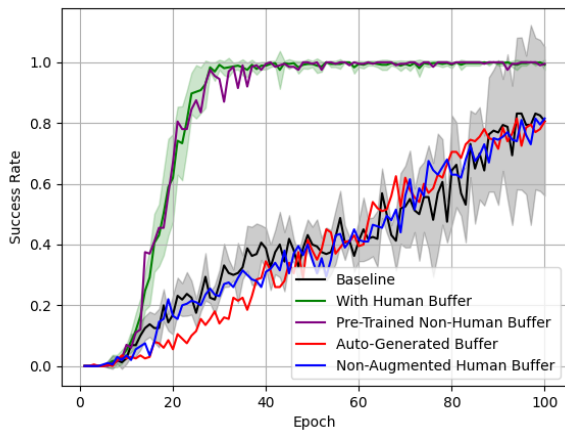


Fig. 6. Success rate on pick-and-place task vs. epoch for baseline DDPG + HER, DDPG + HER with the human buffer, DDPG + HER with a non-human buffer (created by trained agent), DDPG + HER with a generated buffer (created during training), and DDPG + HER with a human buffer without augmentation (only the single recorded trajectory). The shaded region shows one standard deviation from the mean, and was not displayed for all experiments to preserve figure readability. The results show that the augmented human demonstration method significantly outperforms DDPG + HER. Additionally, having an expert demonstration and augmenting that demonstration are vital to our method, but the expert demonstration does not have to be human generated.

Finally, to test if the increase in performance was due to the generation of new trajectories (augmentation step) rather than the quality of the original successful example, we replaced the human buffer with an auto-generated previous successes buffer that used the agent’s successful runs and our augmentation algorithm to generate example trajectories during training. The concept of leveraging the agent’s successes is similar to that of [32], except in this case these successful trajectories are augmented using our generation method. This approach had a similar success rate to baseline DDPG + HER, suggesting that the expert demonstration itself is also a vital part of the method’s success.

3) *Human-Aided Block Stacking*: After analyzing the impact of the single-example human buffer on the push and pick-and-place tasks, we experimented on the stack task. This task, which requires the robot to stack two blocks on top of each other, is more complex than the previous two tasks. Our results can be seen in Fig. 7. In this trial, the baseline DDPG + HER approach failed, with a constant 0% success rate. However, when given a human buffer the agent was successful, converging to above a 90% success rate in 270 epochs. We also tested DDPG + HER with the demonstration based pre-play method, which managed to reach 90% success in 230 epochs. By combining the pre-play method with the human buffer, we were able to further increase the agent’s performance, reaching over a 90% average success rate in only 170 epochs.

As with the push task, the agent learned a policy for completing the stack task that was significantly different from the human example. The human example solves the task by first placing the bottom block in place, then retrieving the top block and stacking it onto the bottom block. However,

the RL agent learned to first grab the top block, place the top block onto the bottom block, then use the still grasped top block to drag the stack to the goal location (shown in Fig. 5). This discrepancy shows that the demonstration based pre-play approach assists training without restricting the agent to the demonstration policy.

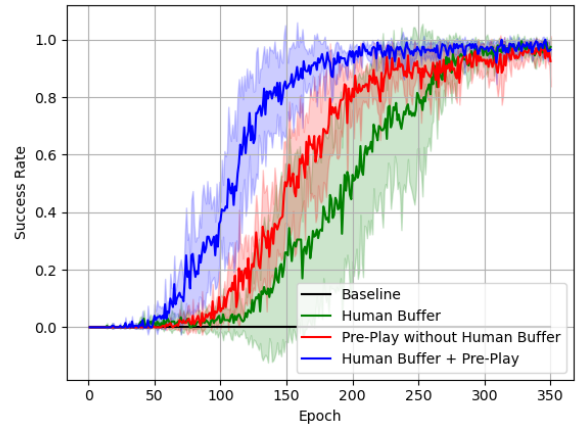


Fig. 7. Success rate on the stack task for baseline DDPG + HER, DDPG + HER with the human buffer, DDPG + HER without the human buffer but with demonstration based pre-play, and DDPG + HER with the human buffer and demonstration based pre-play. The shaded region shows one standard deviation from the mean. The results show that both the human buffer and the demonstration based pre-play improve the agent’s performance.

4) *Effect of HER with Human Buffer*: Hindsight Experience Replay (HER) is a commonly used technique that assists sparse reward reinforcement learning by creating artificially successful experiences to train on [19]. These memories are vital because they present the agent with some examples of positive rewards, which in long horizon tasks may be difficult for the agent to reach organically [19]. However, this role is also fulfilled by both the generated human buffer and the demonstration based pre-play implementation. To determine if the agent still needs HER when using the human buffer, we used DDPG with the human buffer (without HER) to solve the push and pick-and-place tasks, and DDPG with the human buffer and demonstration based pre-play (without HER) to solve the stack tasks (Fig. 8 and 9). Comparing these results to their HER counterparts, we can see that removing HER appears to have slight positive effect on the pick-and-place and push tasks and significantly improves the performance of the stack task. The disparity between our method with and without HER appears to be correlated with the task complexity.

### C. Hardware Validation

In order to test the viability of our trained policies for hardware control, we implemented push, pick-and-place, and stack policies, which we had trained in simulation, on a Franka-Emika Panda arm without hardware fine-tuning. A wrist-mounted Intel RealSense D415 RGB-D depth camera was used to monitor the position of the blocks, which were labeled using AprilTags. An image of the hardware setup can be found in Fig. 10. For each task, we tested three




Fig. 8. Effect of hindsight experience replay (HER) on the push and pick-and-place tasks. The shaded region shows one standard deviation from the mean. HER does not seem to have a significant effect on the agent’s learning.




Fig. 9. Effect of hindsight experience replay (HER) on the stack task. The shaded region shows one standard deviation from the mean. When HER is removed, the agent performs significantly better.

separate trained policies and averaged the results. The pick-and-place task performed the best, with an 87% success rate, while the push task fared worse, with only a 53% success rate. This discrepancy is likely due to the push task relying on more complicated contact dynamics (sliding the block along the table), which are significantly different between simulation and hardware. Finally, the stack task’s learned policy failed on hardware, achieving a 0% success rate, due to the learned policy’s reliance on sliding the stack of blocks along the table, as described in Fig 5. When executing on hardware, the agent is able to stack the blocks, but when it tries to slide the stack to the goal location, the stack doesn’t slide, instead either toppling or remaining stationary. Overall, these results are promising, especially considering they were achieved with zero-shot sim-to-real transfer. Additionally, we think that the push and stack tasks’ performances could be significantly improved with the application of sim-to-real techniques such as hardware tuning and domain randomization during training [33].




Fig. 10. Hardware validation setup. A Franka Emika Panda arm was used to execute the trained policies. A wrist-mounted RealSense camera was used to track the position of the block(s) using AprilTags. 3D printed finger-tip extensions were used to keep the fingers and blocks centered in the camera frame.

## V. CONCLUSIONS

Our results show that a single human example can be used to significantly improve DDPG + HER, with the addition of a generated human buffer decreasing the training time (time to 80% accuracy) of the pick-and-place task to under one fourth of plain DDPG + HER. Additionally, the use of a demonstration based pre-play approach along with the human buffer enabled the block stacking task to be solved. Moreover, because the human example was collected by directly mapping a user’s hand motions in VR, the single demonstration can be completed quickly and without much human effort. Furthermore, our method has been shown to work when the human demonstration is non-ideal, without restricting the agent from finding a more optimal policy. These two aspects combine to allow a non-expert with minimum training to provide all of the necessary human input with a single short demonstration, allowing for a significant increase in model performance from under a minute of human input.

Our augmentation method is in theory algorithm-agnostic, as it can be paired with any DRL algorithm that employs a replay buffer, such as Soft Actor Critic [34] and Twin Delayed DDPG [35]. Thus, moving forward we hope to evaluate the impact of the human-like buffer when combined with these other DRL algorithms to determine if the improvement from the augmentation method holds regardless of algorithm.

Additionally, we would like to improve our model’s generalizability, specifically by developing alternative, more robust methods for generating the human-like trajectories. Although the current generation method works well, it requires the task to be composed of sub-tasks that can be linearly interpolated. By developing a more general method for human-like trajectory generation, we could significantly increase the range of tasks to which the single human demonstration augmentation method can be applied.

## REFERENCES

- [1] A. Zeng, S. Song, S. Welker, J. Lee, A. Rodriguez, and T. Funkhouser, "Learning synergies between pushing and grasping with self-supervised deep reinforcement learning," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 4238–4245.
- [2] D. Quillen, E. Jang, O. Nachum, C. Finn, J. Ibarz, and S. Levine, "Deep reinforcement learning for vision-based robotic grasping: A simulated comparative evaluation of off-policy methods," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2018, pp. 6284–6291.
- [3] 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.
- [4] F. Stulp, E. A. Theodorou, and S. Schaal, "Reinforcement learning with sequences of motion primitives for robust manipulation," *IEEE Transactions on robotics*, vol. 28, no. 6, pp. 1360–1370, 2012.
- [5] I. Popov, N. Heess, T. Lillicrap, R. Hafner, G. Barth-Maron, M. Vecerik, T. Lampe, Y. Tassa, T. Erez, and M. Riedmiller, "Data-efficient deep reinforcement learning for dexterous manipulation," *arXiv preprint arXiv:1704.03073*, 2017.
- [6] J. Matas, S. James, and A. J. Davison, "Sim-to-real reinforcement learning for deformable object manipulation," in *Conference on Robot Learning*. PMLR, 2018, pp. 734–743.
- [7] J. Morimoto, G. Cheng, C. G. Atkeson, and G. Zeglin, "A simple reinforcement learning algorithm for biped walking," in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, vol. 3. IEEE, 2004, pp. 3030–3035.
- [8] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, "Learning to walk in minutes using massively parallel deep reinforcement learning," in *Conference on Robot Learning*. PMLR, 2022, pp. 91–100.
- [9] T. Haarnoja, S. Ha, A. Zhou, J. Tan, G. Tucker, and S. Levine, "Learning to walk via deep reinforcement learning," *arXiv preprint arXiv:1812.11103*, 2018.
- [10] Y. Yang, K. Caluwaerts, A. Iscen, T. Zhang, J. Tan, and V. Sindhwani, "Data efficient reinforcement learning for legged robots," in *Conference on Robot Learning*. PMLR, 2020, pp. 1–10.
- [11] C. Sampedro, H. Bavle, A. Rodriguez-Ramos, P. de La Puente, and P. Campoy, "Laser-based reactive navigation for multirotor aerial robots using deep reinforcement learning," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1024–1031.
- [12] W. Koch, R. Mancuso, R. West, and A. Bestavros, "Reinforcement learning for uav attitude control," *ACM Transactions on Cyber-Physical Systems*, vol. 3, no. 2, pp. 1–21, 2019.
- [13] J. Ibarz, J. Tan, C. Finn, M. Kalakrishnan, P. Pastor, and S. Levine, "How to train your robot with deep reinforcement learning: lessons we have learned," *The International Journal of Robotics Research*, vol. 40, no. 4-5, pp. 698–721, 2021.
- [14] A. Nair, B. McGrew, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Overcoming exploration in reinforcement learning with demonstrations," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 6292–6299.
- [15] M. Vecerik, T. Hester, J. Scholz, F. Wang, O. Pietquin, B. Piot, N. Heess, T. Rothörl, T. Lampe, and M. Riedmiller, "Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards," *arXiv preprint arXiv:1707.08817*, 2017.
- [16] S. Zuo, Z. Wang, X. Zhu, and Y. Ou, "Continuous reinforcement learning from human demonstrations with integrated experience replay for autonomous driving," in *2017 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2017, pp. 2450–2455.
- [17] T. Hester, M. Vecerik, O. Pietquin, M. Lanctot, T. Schaul, B. Piot, D. Horgan, J. Quan, A. Sendonaris, I. Osband *et al.*, "Deep q-learning from demonstrations," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
- [18] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.
- [19] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, O. Pieter Abbeel, and W. Zaremba, "Hindsight experience replay," *Advances in neural information processing systems*, vol. 30, 2017.
- [20] N. Ratliff, J. A. Bagnell, and S. S. Srinivasa, "Imitation learning for locomotion and manipulation," in *2007 7th IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2007, pp. 392–397.
- [21] J. Nakanishi, J. Morimoto, G. Endo, G. Cheng, S. Schaal, and M. Kawato, "Learning from demonstration and adaptation of biped locomotion," *Robotics and autonomous systems*, vol. 47, no. 2-3, pp. 79–91, 2004.
- [22] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proceedings of the fourteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, 2011, pp. 627–635.
- [23] C. G. Atkeson and S. Schaal, "Robot learning from demonstration," in *ICML*, vol. 97, 1997, pp. 12–20.
- [24] G. V. de la Cruz, Y. Du, and M. E. Taylor, "Pre-training with non-expert human demonstration for deep reinforcement learning," *The Knowledge Engineering Review*, vol. 34, p. e10, 2019.
- [25] M. E. Taylor, H. B. Suay, and S. Chernova, "Integrating reinforcement learning with human demonstrations of varying ability," in *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, 2011, pp. 617–624.
- [26] S. Griffith, K. Subramanian, J. Scholz, C. L. Isbell, and A. L. Thomaz, "Policy shaping: Integrating human feedback with reinforcement learning," in *Advances in Neural Information Processing Systems*, C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger, Eds., vol. 26. Curran Associates, Inc., 2013. [Online]. Available: <https://proceedings.neurips.cc/paper/2013/file/e034fb6b66aacc1d48f445ddfb08da98-Paper.pdf>
- [27] P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei, "Deep reinforcement learning from human preferences," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/d5e2c0adad503e91f91df240d0cd4e49-Paper.pdf>
- [28] *Franka Emika Robot's Instruction Handbook*. Franka Emika GmbH, 2021.
- [29] J. Nauta, Y. Khaluf, and P. Simoens, "Using the ornstein-uhlenbeck process for random exploration," in *4th International Conference on Complexity, Future Information Systems and Risk (COMPLEXIS2019)*. SciTePress, 2019, pp. 59–66.
- [30] Z. C. Lipton, J. Gao, L. Li, X. Li, F. Ahmed, and L. Deng, "Efficient exploration for dialog policy learning with deep bbq networks & replay buffer spiking," *CoRR abs/1608.05081*, 2016.
- [31] Q. Gallouédec, N. Cazin, E. Dellandréa, and L. Chen, "Multi-goal reinforcement learning environments for simulated franka emika panda robot," *arXiv preprint arXiv:2106.13687*, 2021.
- [32] J. Luo and H. Li, "Dynamic experience replay," in *Conference on robot learning*. PMLR, 2020, pp. 1191–1200.
- [33] B. Mehta, M. Diaz, F. Golemo, C. J. Pal, and L. Paull, "Active domain randomization," in *Proceedings of the Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, L. P. Kaelbling, D. Kracic, and K. Sugiura, Eds., vol. 100. PMLR, 30 Oct–01 Nov 2020, pp. 1162–1176. [Online]. Available: <https://proceedings.mlr.press/v100/mehta20a.html>
- [34] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *International conference on machine learning*. PMLR, 2018, pp. 1861–1870.
- [35] S. Fujimoto, H. Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," in *International conference on machine learning*. PMLR, 2018, pp. 1587–1596.