Improved Learning of Robot Manipulation Tasks via Tactile Intrinsic Motivation

Nikola Vulin[1], Sammy Christen[1], Stefan Stevšić[1] and Otmar Hilliges[1]

Abstract—In this paper we address the challenge of exploration in deep reinforcement learning for robotic manipulation tasks. In sparse goal settings, an agent does not receive any positive feedback until randomly achieving the goal, which becomes infeasible for longer control sequences. Inspired by touch-based exploration observed in children, we formulate an intrinsic reward based on the sum of forces between a robot's force sensors and manipulation objects that encourages physical interaction. Furthermore, we introduce contact-prioritized experience replay, a sampling scheme that prioritizes contact rich episodes and transitions. We show that our solution accelerates the exploration and outperforms state-of-the-art methods on three fundamental robot manipulation benchmarks.

Index Terms—Reinforcement Learning, Deep Learning in Grasping and Manipulation, Tactile Feedback, Intrinsic Motivation

I. INTRODUCTION

MODEL-FREE deep reinforcement learning (DRL) algorithms have demonstrated great potential in solving sequential decision making problems, such as learning to play video games in Atari [1], defeating the world champion in the game of Go [2], or controlling robotic systems for locomotion [3], navigation [4] and manipulation [5]. For complex tasks and real-world scenarios, in particular in robotics, formulating a feasible reward function is inherently difficult and tedious. Hence, reward functions are most easily formulated as a sparse signal received upon reaching the final goal. This leads to inefficient exploration, because an agent receives crucial feedback rarely and in a delayed manner. Reward functions are typically based on reaching *extrinsic* states. Thus, the reward does not leverage internal robot states or other internal representations. To address inefficient exploration, an *intrinsic* reward inspired by intrinsic motivation observed in humans [6] can be added to the reward function.

Humans strongly rely on the sense of touch for exploration and use it as guidance when interacting with their surroundings. For example, humans can manipulate objects even without visual feedback, and removing tactile cues significantly reduces manipulation capabilities [7]. Furthermore, it has been shown that force feedback is used by infants to explore the

Fig. 1. Inspiration of our work comes from intrinsic motivation of infants exploring the world through physical interactions. We model tactile sensing with force sensors (red) and leverage tactile information to accelerate learning.

physical properties of objects [8]. Our insight is that force feedback can be leveraged in a general manner for learning a diverse set of manipulation tasks. Inspired by the touch-based exploration observed in small children, we propose an *intrinsic* reward based on tactile signals to enhance exploration in robotic manipulation environments. Furthermore, we introduce a sampling scheme that encourages robot-object interaction via prioritization of meaningful trajectories.

Intrinsic rewards have been introduced for discrete domains based on state novelty [9], uncertainty [10], [11] or information gain [12]. Using methods based on intrinsic motivation accelerates the training of an agent, but requires an additional model to extract a reward from state information. In DRL approaches, tactile robot-object interaction information is often ignored, although this information might be valuable during exploration. Some prior works extend the state space with tactile information to provide explicit interaction cues. This results in improvements in terms of sample efficiency and robustness for grasping [13], in-hand manipulation [14], and human-robot interaction [15]. However, these methods modify only the state space [13], [14], or focus on a set of specific tasks [13], [15].

In our method (see Figure 1), we add force sensors to the end-effector and extend the state space with force measurements, similarly to other work [13], [14], [15]. We then go beyond prior work and introduce an intrinsic reward to guide the exploration. Our intrinsic reward behaves like an intermediate, simpler to achieve, directive base. Having an intermediate base can improve exploration, since the agent returns to this good position and does not lose track due to probabilistic exploration [16]. The reward motivates the agent to touch the

object while allowing it to explore different interactions, i.e., when and how long to touch an object. We experimentally show that our reward leads to significantly better performance on benchmark manipulation tasks, especially when the task difficulty is increased.

Another way of addressing exploration is to use data augmentation techniques such as Hindsight Experience Replay (HER) [17], which replaces a task-specific goal with a state that was achieved in hindsight of a sampled transition. This will induce intermediate rewards for failed trials. However, HER randomly samples states from the replay buffer, including trajectories where the manipulation object does not move, which may slow down learning. We therefore introduce a replay buffer sampling scheme that allows an agent to prioritize meaningful trajectories. More specifically, we first modify the sampling probability of entire episodes depending on whether contacts have happened. From a batch of selected episodes, we sample virtual goals, with an increased weight on transitions that occur after initial contact. We then go back in time to find training transitions, which narrows down the search space to meaningful transitions for manipulation tasks. We show empirically that this leads to faster convergence compared to HER.

The main contributions of this paper can be summarized as follows: i) An *intrinsic* force-based reward that accelerates the exploration and improves the performance on several benchmark robotic manipulation tasks. ii) A novel sampling scheme for off-policy algorithms in contact-based tasks to increase the sample efficiency. iii) We provide a detailed evaluation of our approach and comparison with the baselines. We analyze the contributions of the method components in an ablation study and show that our method performs significantly better, especially with increased task difficulty.

II. RELATED WORK

A. Tactile Feedback

For humans, tactile sensing is an important sensory modality, besides visual perception, when manipulating objects. It is rich in information and free of external disturbances, compared to visual sensing that is influenced by occlusions and poor lighting conditions. While the majority of DRL works do not consider tactile feedback for learning robotic manipulation tasks, a few have shown the benefits and potential of considering this additional sensory modality [13], [14], [15], [18], [19].

In [13], the authors show that tactile sensing can improve the robustness of grasping under noisy position sensing, thus supporting exact position estimation. [14] demonstrates that tactile sensing provides the agent with essential information and makes the learning more sample efficient for in-hand manipulation tasks. Unlike our method, these approaches only add the force measurement to the state space.

More similar to our method, [15] and [18] extend their state space with tactile sensing and use it in the reward function. To learn human-robot interactions, [15] propose a specific reward to ensure contact between the human and the robot. [18] uses the feedback of tactile sensors as a penalty to avoid

high impacts and hence to learn gentle manipulation. However, in both methods the reward is used to trigger a task-specific behavior. In contrast, we formulate an intrinsic reward to encourage general task-independent exploration.

B. Intrinsic Motivation

Humans are efficient learners, showing intrinsically motivated behaviors when exploring the world around them [6]. Inspired by human exploration strategies, intrinsic motivation in reinforcement learning (RL) has been used to improve exploration in sparse reward settings [20]. The research mostly relies on intrinsic rewards that an agent collects for visiting novel states. This type of reward can be implemented by computing pseudo-counts through density models [9] or by estimating state novelty with neural networks [11]. Furthermore, one can model curiosity to guide the agent towards areas where the uncertainty of the agent's prediction of the following states is high [10]. These methods need to estimate additional, computationally expensive values, which slows down training. In contrast, our intrinsic reward only relies on tactile feedback present in the state space.

C. Experience Replay

Experience replay [21] is applied in off-policy DRL algorithms to store experience transitions in a buffer and reuse them during training. Popularized in DQN [1], it has become a standard procedure to mitigate sample efficiency and stabilize training in most state-of-the-art off-policy algorithms [22]. An extension includes Prioritized Experience Replay (PER) [23], where transitions are prioritized based on how valuable they are to the learning process via the temporal difference error.

For the case of learning from sparse rewards in continuous control problems, Hindsight Experience Replay (HER) [17] has proposed a data augmentation technique that achieves large performance improvements over standard approaches. The transitions stored in the buffer are used for training, with the difference that rewards are computed with respect to virtual goals. Even when the agent fails to complete the task, virtual goals provide a way to achieve intermediate rewards. HER was extended in [24] by prioritizing episodes based on the sum of kinetic and potential energy of the manipulation object. This can lead to prioritization of undesired trajectories, because the faster the object moves, the higher the priority, which may not be what is beneficial for the task at hand. In contrast, our prioritization scheme treats episodes with contact more equally and also modifies the sampling probability of training transitions within an episode.

III. PRELIMINARIES

A. Problem Definition

In this paper, we focus on using reinforcement learning to solve robotic manipulation tasks, where a robot needs to interact with an object to complete the task. For example, the robot may need to pick objects and place them in a container, build a structure, or use tools. A straightforward way to learn these tasks is to reward the agent once the task

is completed. In contrast, shaping the reward might be very difficult because there are multiple intermediate steps that the robot needs to achieve before it can reach the final goal. When sparse rewards are used, no signal directs the agent towards the goal, and hence the agent reaches it only by chance via random exploration. Thus, algorithm convergence is usually very slow, or the algorithm does not even converge if the task is too complex. We aim to improve exploration in robotic manipulation tasks and hence learn the required skills faster. In manipulation tasks, we can often resort to the goal-conditioned formulation of the reinforcement learning problem. This can help to mitigate the exploration problem by applying methods such as HER [17].

B. Goal-conditioned Deep Reinforcement Learning

To model our control problem, we use the standard formulation of a Markov Decision Process (MDP) and extend it with a set of goals $\mathcal{G}$. An MDP is defined by a tuple $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, \mathcal{G}, \mathcal{R}, \mathcal{T}, \rho_0, \gamma\}$, where $\mathcal{S}$ is the state set, $\mathcal{A}$ the action set, $\mathcal{R}_t = r(s_t, a_t, g_t)$ the reward function, $\mathcal{T} = p(s_{t+1}|s_t, a_t)$ the transition dynamics of the environment with $s_t \in \mathcal{S}$ and $a_t \in \mathcal{A}$, $\rho_0 = p(s_1)$ the initial state distribution, and γ a discount factor. Our aim is to maximize the expected sum of discounted future rewards with a policy $\pi : \mathcal{S} \times \mathcal{G} \to \mathcal{A}$, which is a mapping from states and goals to actions. We adopt the actor-critic framework and use the Q function to describe the expected future reward:

$$Q(s_t, g_t, a_t) = \mathbb{E}_{a_t \sim \pi, s_{t+1} \sim \mathcal{T}} \left[\sum_{i=t}^{T} \gamma^{i-t} \mathcal{R}_t \right] \qquad (1)$$

We approximate both the value function (critic) and the policy (actor) with neural networks and use the DDPG algorithm [22] for training. Thus, the critic network is parameterized by θ^Q and minimizes the following loss:

$$L(\theta^Q) = \mathbb{E}_{\mathcal{M}} \left[\left(Q(s_t, g_t, a_t; \theta^Q) - y_t \right)^2 \right],$$
$$y_t = r(s_t, g_t, a_t) + \gamma Q(s_{t+1}, g_{t+1}, a_{t+1}; \theta^Q). \qquad (2)$$

The policy network uses parameters θ^π and is trained to maximize the Q-value:

$$L(\theta^\pi) = \mathbb{E}_\pi \left[Q(s_t, g_t, a_t; \theta^Q) | s_t, g_t, a_t = \pi(s_t, g_t; \theta^\pi) \right] \qquad (3)$$

The algorithm is off-policy and therefore uses an experience replay buffer for training. The data is collected in episode rollouts (we use the word trajectory interchangeably in this paper) and stored as transitions $(s_t, g_t, a_t, r_t, s_{t+1})$ in the buffer. These transitions are then sampled from the replay buffer to train the neural networks that approximate the Q function and control policy π.

C. Hindsight Experience Replay (HER)

Hindsight Experience Replay (HER) [17] is a data augmentation technique that significantly improves the exploration of goal-conditioned reinforcement learning algorithms. In HER, the key idea is to learn from failed trials. Therefore, uniformly
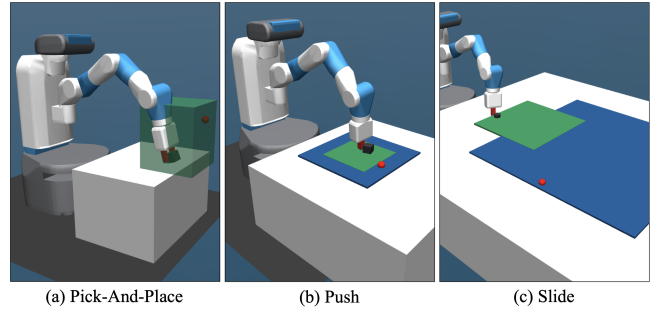

(a) Pick-And-Place   (b) Push   (c) Slide

Fig. 2. Overview of the benchmark tasks. We enlarge the benchmark's [25] original goal space (green) to a larger area (blue) to increase the task difficulty. For the *Pick-And-Place* task, we sample goals uniformly contrary to the original benchmark task, which uses 50% of targets on the table.

sampled transitions from the replay buffer are modified by replacing the goal g_t of a transition with a state achieved in a subsequent time step of the same trajectory. Accordingly, the reward r_t must be recalculated, which will consequently induce rewards for trajectories that failed to reach the original goal. From these unsuccessful trials, the goal-conditioned policy should learn to reach the original goal more efficiently via extrapolation. However, the goals used in HER are usually object-centric. Hence, it can take a significant amount of trials until the agent starts to move the object and virtual goals start to affect the learning algorithm. This is a substantial downside of HER, which we address with our method.

D. Simulation Environments

We evaluate our method on three different benchmark tasks from the robotics collection [25] of OpenAI Gym [26]. The environments are based on the physics engine MuJoCo [27]. MuJoCo allows to design tailored physical environments and computes the dynamics and contact interactions between rigid bodies. We use a 7 DoF robot manipulator (see Figure 1). Marked in red are the locations of the force sensors, which are placed on the end-effector. The robot is controlled via torque actions that are sent to its motor joints.

Figure 2 shows the three tasks considered: *Pick-And-Place*, *Push* and *Slide*. These tasks represent a diverse set of manipulation tasks. In *Pick-And-Place*, the goal is to pick up an object and move it to a target. In *Push*, the task consists of pushing an object on a surface to a target position. Contrarily, in *Slide* the robot has to kick a puck in order to reach a distant goal that is out of reach of the robot's workspace. The red sphere in Figure 2 marks the target position an object needs to reach. The original goal space [25] is illustrated in green. To increase the difficulty of the problem, we enlarge the spread of the goal space from the green to the blue area. For *Pick-And-Place*, the original problem [25] uses the simplification that 50 % of goals are on the table and not in the air. This can lead to instances where the robot does not need to learn to grasp the object in order to complete the task. Hence, we modify the task such that all goals are sampled uniformly in the marked green area.

IV. METHOD

When training robots to execute manipulation tasks via reinforcement learning, the reward function is most often
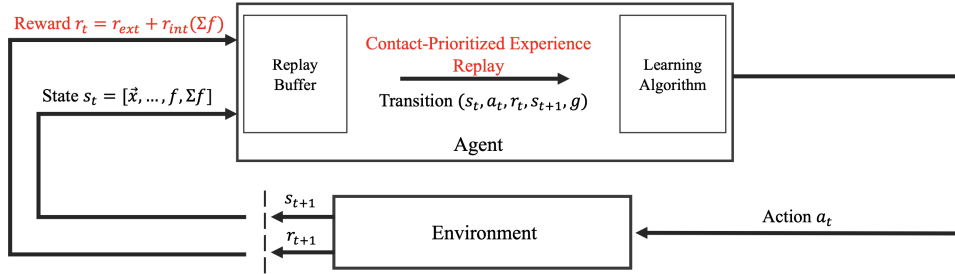
Fig. 3. The standard reinforcement learning loop and our two method extensions in red. The reward function is extended by an intrinsic reward and a sampling scheme that prioritizes contact-rich transitions is introduced.

defined as sparse. Therefore, it is essential to have an efficient exploration approach. We propose a method that improves exploration efficiency by leveraging force sensor information in the goal-conditioned reinforcement learning setting (see Figure 3).

Our main insight is that sense of touch can guide robots to take actions that result in manipulation of an object. To implement the sense of touch, we extend our state-space with force measurements from tactile sensors positioned at the end-effector [13], [14], [15]. We introduce a touch-based intrinsic reward function to direct exploration towards the states where the robot touches the object. The trajectory in the state space that leads to the task solution passes through the state where the robot touches the object. Hence, incentivizing the agent to reach such states will guide the exploration towards the final task goal. Using the intrinsic reward in combination with HER utilizes virtual goals much earlier because the robot starts to move the object significantly faster. We further exploit the touch information by introducing a Contact-Prioritized Replay Buffer (CPER). CPER uses contact clues to sample more informative transitions and virtual goals for training the Q function and control policy π, which further speeds up the learning.

A. State Space

We extend our state space with force measurements from tactile sensors positioned at the end-effector [13], [14], [15]. The original state space consists of a mixture between proprioceptive state information, i.e., joint positions and joint velocities, and relevant information about the manipulation object, such as the object's velocity and the position x_{obj}, which we will use in the reward function. We further extend the state space with the measured force f_t at the end-effector and the sum of all measured forces until the current step $\sum_{i=0}^{t} f_i$. The accumulated force is a good measure to track how much contact occurred during an episode, which is what our intrinsic reward is based on.

B. Intrinsic Force Reward (IR)

We extend an extrinsic reward function $r_{\text{ext}}(s, g)$ with our intrinsic force reward $r_{int}(s)$, which is independent of the goal:

$$r(s, g) = \omega_{\text{ext}} * r_{\text{ext}}(s, g) + \omega_{\text{int}} * r_{\text{int}}(s), \qquad (4)$$

where ω_{ext} and ω_{int} are the weights on the extrinsic and intrinsic reward function, respectively. We empirically found that the weights $\omega_{\text{ext}} = 0.75$ and $\omega_{\text{int}} = 0.25$ provide a good balance between our intrinsic reward as an auxiliary reward, while putting more emphasis on the extrinsic reward function. The extrinsic reward function is formulated sparsely, where the indicator function $\mathbb{I}$ returns a positive signal of magnitude 1 if the object's position x_{obj} is within the range ε_{pos} around the goal g:

$$r_{\text{ext}}(s, g) = \mathbb{I}\left[\left\| g - x_{\text{obj}} \right\| < \varepsilon_{\text{pos}} \right] \qquad (5)$$

In all other cases, the reward is 0. The intrinsic reward depends only on the state and returns a reward if the current sum of forces $\sum_{i=0}^{t} f_i$ is higher than a minimal amount of desired contact interaction $\varepsilon_{\text{force}}$:

$$r_{\text{int}}(s_t) = \mathbb{I}\left[\sum_{i=0}^{t} f_i > \varepsilon_{\text{force}} \right] \qquad (6)$$

The intrinsic force reward allows the agent a certain freedom about when and how long to touch the object. As soon as a minimal amount of contact occurred, i.e., when Σf is greater than a minimal threshold, the agent receives the intrinsic reward until the end of the episode and can concentrate on solving the extrinsic task. The reward encourages the agent to manipulate the object and to bring it into motion. Therefore, the value of the threshold $\varepsilon_{\text{force}}$ should be chosen large enough to avoid falsely providing the intrinsic reward due to sensor noise. We empirically set the threshold to $10N$ for the tasks *Push* and *Pick-And-Place* and $3N$ for the *Slide* task, since this environment requires only a short interaction with the object. However, we found that the threshold is robust across a range of different values.

We formulate the intrinsic reward such that it has a minimal influence on the final task. After the force threshold is reached, the agent can freely explore interaction with an object. Hence, we direct the agent to a crucial intermediate base from where the goal can be reached more quickly. Importantly, we use the same reward function across multiple tasks, demonstrating that our formulation can work in the general case. We also address the issues of HER in early exploration. Specifically, when an object is not moved during an episode, HER does not have an effect on training (see Section III-C for details). Our intrinsic reward circumvents the early exploration issue by guiding the

---

Algorithm 1 Contact-Prioritized Experience Replay (CPER)

1: **Given:** Transitions $(s_t, g_t, a_t, r_t, s_{t+1})$ of episode $e \in E$ and step $t \in T$ stored in Replay Buffer $\mathcal{R}$
2: Compute probability distribution $p_{\text{transition}}(e, t)$ ▷ Fig. 4, Eq. 7
3: Sample mini-batch $\mathcal{B}$ of episodes from $\mathcal{R}$ with marginal probability distribution $p_{\text{episode}} = \sum_{t=0}^{T} p_{\text{transition}}(e, t)$
4: **for** episode e_b in $\mathcal{B}$ **do**
5:      Sample state $s_{t'}$ from $p_{\text{transition}}(e_b, t)$
6:      Sample transition $(s_t, g_t, a_t, r_t, s_{t+1})$ back in past where $0 \leq t \leq t'$.
7:      **if** Hindsight Transition **then**   ▷ See HER [17]
8:          Replace original goal with sampled state $g_t \leftarrow s_{t'}$
9:          Recompute $r' := r(s_t, a_t, g_t)$
10:      **end if**
11: **end for**
12: Perform one step of optimization using Algorithm $\mathbb{A}$ and mini-batch $\mathcal{B}$

---

agent to interact with the object and significantly speeds up convergence as we show in Section V-A.

C. Contact-Prioritized Experience Replay (CPER)

When training the value and policy networks, transitions are sampled from a replay buffer. The standard procedure is to sample the transitions uniformly. However, not all transitions are of the same value to the learning algorithm. Therefore, prioritizing more informative transitions helps to speed up learning [23]. Furthermore, the selection of virtual goals is important as well. We leverage tactile information to alter the sampling of transitions and virtual goals.

We therefore introduce Contact-Prioritized Experience Replay (CPER), which is a sampling prioritization scheme under the assumption that contact rich trajectories include more meaningful information for the learning process. Our method relies on two important features to accelerate learning. First, we prioritize trajectories where the agent touches the object. These trajectories contain important information about the relevant motions of the agent and the object. Other trajectories include robot motions where the end-effector may be far away from the object and are therefore less relevant for training. Second, we sample for virtual goals first and then find training transitions by sampling from previous time steps. We sample virtual goals from states occurring after the contact to induce meaningful intermediate rewards. By sampling transitions from previous time steps, we increase the chance of finding relevant transitions. This reversed scheme helps in tasks like *Slide*, where most important transitions occur prior to the contact phase. We experimentally show that using CPER leads to significantly faster training compared to the case when only our intrinsic reward is used. We now explain our sampling scheme in more detail.

First, we compute the probability distribution $p_{\text{transition}}$ using $p'_{\text{transition}}$ (see Figure 4) and the following equation:

$$p'_{\text{transition}}(e, t) = \begin{cases} 1 & \text{if } \sum_{i=0}^{t} f_i < \varepsilon_{\text{force}} \\ \lambda & \text{otherwise,} \end{cases}$$

$$p_{\text{transition}}(e, t) = \frac{p'_{\text{transition}}(e, t)}{\sum_{e=0}^{E} \sum_{t=0}^{T} p'_{\text{transition}}(e, t)} \qquad (7)$$
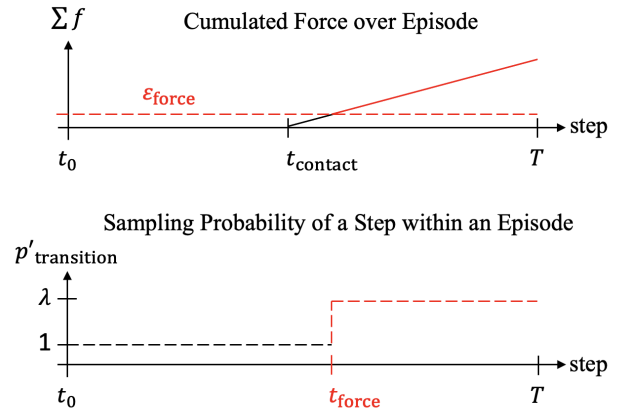

Fig. 4. Computation of $p'_{\text{transition}}$. Once the force threshold t_{force} is surpassed, the probability is multiplied by a factor λ.

Once the sum of forces reaches the threshold $\varepsilon_{\text{force}}$ at time step t_{force} within a trajectory, we increase the sampling probability of transitions for all subsequent time steps until the end of the episode T by factor λ. We empirically found that a factor of 10 leads to the best performance.

Algorithm 1 describes our method in detail and consists of two parts. First, we compute the marginal probability distribution p_{episode} from the sum of $p_{\text{transition}}$ per episode in the buffer, which overweights trajectories where contact occurred and further prioritizes based on the time step of the first contact (lines 2-3). This induces a prioritization of episodes with contact-rich information. We then sample a batch of trajectories $\mathcal{B}$ according to p_{episode} (line 3).

In the second part (lines 4-11), we sample a virtual transition at time t' within an episode from batch $\mathcal{B}$ according to our modified distribution $p_{\text{transition}}$ (line 5). We then reverse back in time to find a training transition $(s_t, g_t, a_t, r_t, s_{t+1})$ and replace the original goal g_t with the achieved state $s_{t'}$ from the virtual transition (line 6-9). Our sampling scheme allows us to find more meaningful virtual goals, narrowing down the search space to transitions that are relevant for robot manipulation tasks. In contrast to HER [17], where first transitions are sampled uniformly and then virtual goals are selected from achieved states in the near future, we search for relevant goals first and then sample back in time to find training transitions.
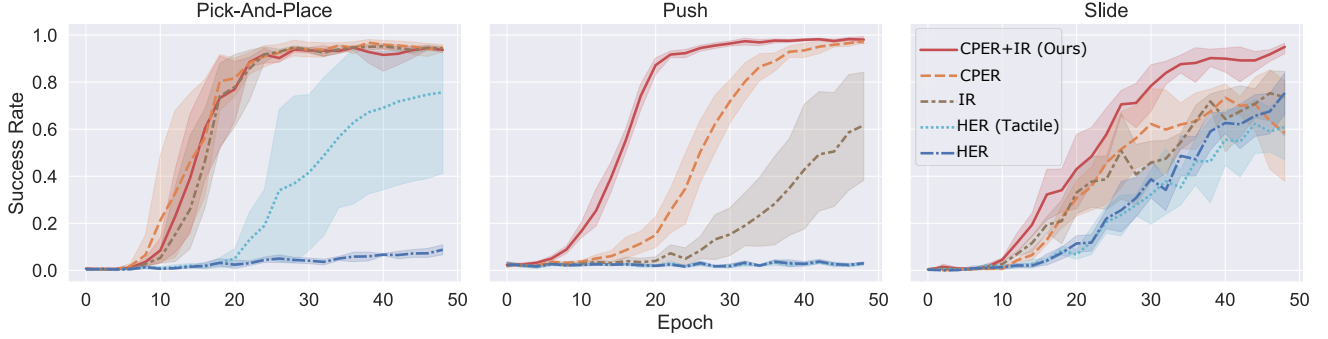
Fig. 5. Learning curves for all three tasks. We show the average success rates over 5 random seeds with the corresponding confidence interval. 1 epoch corresponds to 40 training episodes. We show our full method with contact-prioritized experience replay and the intrinsic reward (CPER+IR) (red) and two ablations: CPER (orange) and IR (brown). We compare it with two baselines: HER with tactile sensing (light blue) and the original HER method (blue).

Since we first sample goal transitions with a higher sampling probability for transitions after the first contact (see Equation 7), it ensures that states in the trajectory before the contact are not neglected. We only replace the goals for a subset of transitions (80%, the same as in [17], line 7), keeping the original goal in the other subset as the task's final goal.

V. RESULTS

A. Robot Manipulation Experiments

To evaluate the proposed approach, we choose three standard manipulation tasks from MuJoCo benchmarks: *Pick-And-Place*, *Push*, and *Slide* (see Section III-D). These three environments require significantly different interactions between the robot and the object, allowing us to evaluate the generalization abilities of our approach. The *Pick-And-Place* task requires the robot to grasp the object, the *Push* task requires one side touching, while *Slide* requires a short kick interaction. In this experiment, we use the hardest version of each task as described in Sec. III-D. Our goal is to evaluate our contributions in the most demanding settings. We use DDPG [22] with HER [17], which we call only HER for abbreviation going forward, as a baseline because it is the state-of-the-art method for training goal-conditioned control policies. In the ablative study, we test how each of the components of our method performs. First, we extend the state space with force and cumulative force sensor readings, as explained in Sec. IV-A, and use HER to train the agent (HER (Tactile)). Next, we separately add our intrinsic reward (IR) and our contact-prioritized experience replay (CPER), where we use our prioritization scheme with only the extrinsic reward. Finally, our full method uses all the components, combining the intrinsic reward and CPER (CPER+IR (Ours)).

Figure 5 shows the results of the experiments. Our approach learns to complete the task significantly faster than HER. The difference is more prominent in *Pick-And-Place* and *Push*, since HER has problems solving these tasks due to the increased task difficulty. In the original settings [25], HER manages to learn these tasks. However, our method leads to faster training in the original settings as well (see Figure 6). In all tasks, our approach converges after 20 to 30 epochs. In particular, at the beginning of training, our method finds successful actions faster. Hence, it yields a steeper learning

curve, which indicates that our method motivates the agent to manipulate the object from the beginning.

Our ablation study shows that simply adding tactile information to the state space does not automatically lead to faster training. HER usually performs the same with and without the tactile sensor readings. In *Pick-And-Place*, the tactile feedback improves the performance, but it is unstable. The faster convergence may imply that the tactile feedback gives an important signal that the object is grasped. When the intrinsic reward is used (IR), the performance is significantly improved compared to both HER and HER (Tactile). This confirms that the proposed intrinsic reward function indeed speeds up training. In *Slide*, the intrinsic reward has the smallest effect because this task demands a precise movement of the robot arm prior to the short kick of the object.

Finally, we analyze the influence of our sampling scheme CPER. While there is no significant difference between IR and CPER+IR for *Pick-And-Place*, we observe that it is crucial to choose transitions with valuable information in *Push* and *Slide* to achieve faster convergence. In *Pick-And-Place*, the agent starts to grasp the object relatively fast due to the intrinsic reward, and hence most of the trajectories are prioritized for sampling. Therefore, our sampling scheme does not have a significant effect. Furthermore, applying CPER without IR shows similar results as the full method in *Pick-and-Place*, implying that the intrinsic reward can potentially be avoided in some tasks.

B. Task Difficulty

To study the effects of increasing the task difficulty, we conducted experiments using our method and the HER baseline. The original environments use a much smaller goal space than the robot can potentially reach in *Push* and *Slide*, or use the simplification that 50% of the goals in *Pick-And-Place* are on the table, which is a curriculum that simplifies the exploration. We test both approaches in three environments where we gradually increase the difficulty. In the Simple environment, we use the original settings. For Intermediate, we use settings that are half-way between the Simple and Hard environments, i.e., we increase the range of the goal space for *Push* and *Slide*, while we decrease the percentage of goals on the table
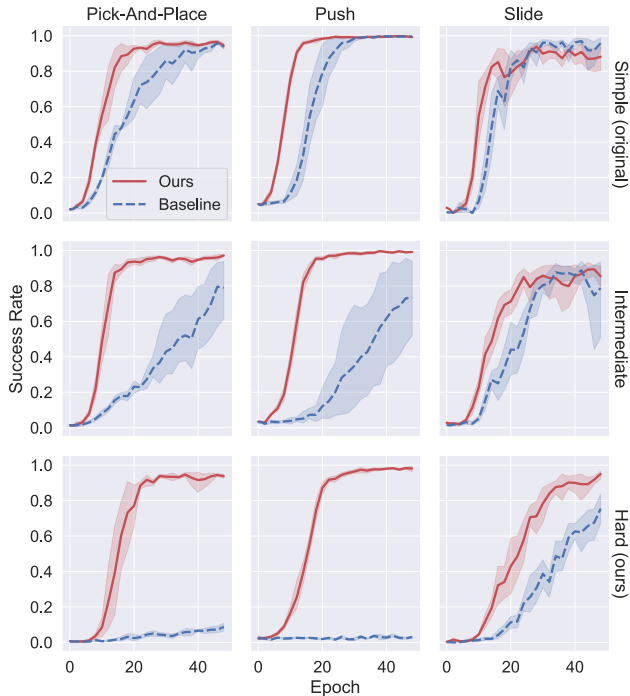
Fig. 6. Learning curves for all three tasks and difficulty levels. We show the average success rate over 5 seeds with the corresponding confidence interval. 1 epoch corresponds to 40 training episodes. We compare our full method CPER + IR (red) and the original HER method (blue).

for *Pick-And-Place*. Hard is our final environment described in Sec. III-D.

The results are illustrated in Figure 6. When the task difficulty is kept the same as in [25] (Simple, top row), both methods reach similar success rates. However, we can see that our method converges faster. Similar results are observed when the task difficulty is slightly increased (Intermediate, middle row), although the convergence of the baseline slows even more compared to our method. Lastly, if the difficulty is increased as described in Section III-D, our method shows significantly better performance both in terms of convergence and success rate. These results are in accordance with insights from [28], which demonstrates that methods often perform well due to constrained goal spaces. Because of the enhanced exploration, our method manipulates the object once discovered and can therefore more efficiently generalize to a larger goal space.

C. Sampling Ablation Study

In this experiment, we analyze our prioritization scheme CPER in more detail. CPER uses two main components: Increasing the probability of sampling transition from episodes with contact, and sampling transitions and virtual goals inside the trajectory based on the contact occurrence. We compare full CPER to a sampling scheme that uses just the first component, i.e., we prioritize episodes based on contact occurrences. Hence, we sample the virtual goal state $s_{t'}$ uniformly, instead of sampling it from $p_{\text{transition}}$ (compare line 5 in Algorithm 1). We name this sampling scheme *episode sampling*. We also compare against a more conventional baseline, where we
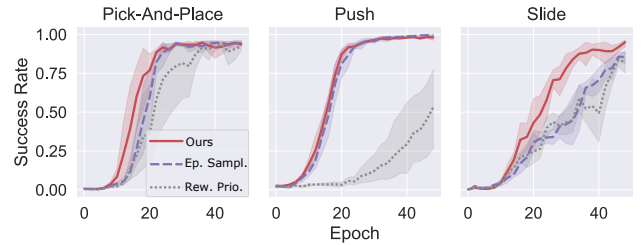

Fig. 7. Ablation study of our sampling method. We show the average success rate over 5 random seeds with the corresponding confidence interval. 1 epoch corresponds to 40 training episodes. We compare our full method CPER + IR (red) and an ablation (purple) with a classical reward prioritiza-tion scheme (grey). In our ablation we find the training episodes according to CPER, but sample the state and virtual goal uniformly (like in HER).

prioritize episodes based on the achieved reward and sample the transitions uniformly, which we call *reward prioritization*. In all three methods, we use our IR and the tactile feedback for an appropriate comparison.

Figure 7 illustrates the results. We can see that our full method is generally exploring faster. In particular, for *Slide*, where the physical interactions with the object are limited to a few steps, we can see that *episode sampling* shows slower convergence. This indicates that the full method is finding more informative transitions, which our method achieves by sampling virtual goals from $p_{\text{transition}}$ and then revert back to find a meaningful training transition. In this task, it is important to learn the swing trajectory prior to the kick and to choose virtual goals after the kick. If we choose transitions randomly, we more often sample non-informative transitions that occur after the kick or goals where the object does not move, which slows down training. The worse performance of the *reward prioritization* baseline further indicates the benefit of our prioritization scheme.

VI. DISCUSSION

Here we address the current limitations of our method and propose areas for future work. This paper focuses on the three commonly used benchmarks in RL research for robotic manipulation tasks [25]. These fundamental problems have significantly different contact patterns and a single object in the scene. This allows us to demonstrate the benefits of an intrinsic reward based on tactile information and our proposed sampling scheme. However, applying our method to more complex tasks such as the manipulation of multiple objects or a door opening task, would require adding further modules. To solve such tasks, our method could be combined with other solutions, such as using a task-specific curriculum [29] for multiple objects or learning from expert demonstrations [30] for door opening.

To study the properties of the method, we conducted experiments in simulation. The deployment of simulator policies on a real robot often fails due to the discrepancy of sensor measurements and system dynamics between the real and simulated systems [3]. To overcome this, either the accuracy of simulators needs to be improved or the policies have to be robust against imperfections in real world settings. Approaches that combine both improving the accuracy of simulators as

well as applying domain randomization have been proposed [3] and could be combined with our method. The contribution of our approach will most likely be robust for transfer to a real robot, since sensor noise will not trigger the threshold for receiving the intrinsic reward. One more aspect to consider is sensor placement, which we found crucial for avoiding exploitation of the intrinsic reward, e.g., by starting to push on the table. Here we use force sensors to get information about the contact with the object. In future work, one could infer tactile information from other sensor modalities, such as ultrasonic proximity sensors or vision [31], which would further alleviate this issue and yield a simpler solution for a real system.

While we primarily focus on manipulation tasks, where using the sense of touch as intrinsic reward is intuitive, our method could be extended to other domains. In a more general sense, if a task can be decomposed into distinct phases, demarcated by a measurable event or physical property, the intrinsic reward and the sampling scheme could be applied. For example, reaching a certain amount of vertical thrust might be useful for learning flying skills in drones.

VII. CONCLUSION

We study the challenges of exploration and efficient learning in deep reinforcement learning for robot manipulation tasks. We show in our experiments that an intrinsic force reward overcomes the difficulties of initial exploration and results in learning the intended behavior faster. We have discovered that transitions where the agent manipulates the object and brings it into motion contain valuable information for the learning process. We therefore introduce an up-sampling scheme that prioritizes exactly these transitions. We show that our prioritization scheme accelerates the learning progress even more. We find that our solution improves the performance and enhances the exploration on three fundamental manipulation tasks. Thus, we conclude that tactile feedback has the potential to advance reinforcement learning a step further.

REFERENCES

[1] 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.

[2] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton *et al.*, "Mastering the game of go without human knowledge," *Nature*, vol. 550, no. 7676, pp. 354–359, 2017.

[3] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Science Robotics*, vol. 4, no. 26, 2019.

[4] S. Christen, L. Jendele, E. Aksan, and O. Hilliges, "Learning functionally decomposed hierarchies for continuous control tasks with path planning," *IEEE Robotics and Automation Letters*, 2021.

[5] I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas *et al.*, "Solving rubik's cube with a robot hand," *arXiv preprint arXiv:1910.07113*, 2019.

[6] R. M. Ryan and E. L. Deci, "Intrinsic and extrinsic motivations: Classic definitions and new directions," *Contemporary educational psychology*, vol. 25, no. 1, pp. 54–67, 2000.

[7] L. A. Jones and S. J. Lederman, *Human hand function*. Oxford University Press, 2006.

[8] E. Kosoy, J. Collins, D. M. Chan, J. B. Hamrick, S. Huang, A. Gopnik, and J. Canny, "Exploring exploration: Comparing children with rl agents in unified environments," *arXiv preprint arXiv:2005.02880*, 2020.

[9] M. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos, "Unifying count-based exploration and intrinsic motivation," in *Advances in Neural Information Processing Systems*, 2016, pp. 1471–1479.

[10] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell, "Curiosity-driven exploration by self-supervised prediction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017, pp. 16–17.

[11] Y. Burda, H. Edwards, A. Storkey, and O. Klimov, "Exploration by random network distillation," in *International Conference on Learning Representations*, 2018.

[12] A. Baranes and P.-Y. Oudeyer, "Active learning of inverse models with intrinsically motivated goal exploration in robots," *Robotics and Autonomous Systems*, vol. 61, no. 1, pp. 49–73, 2013.

[13] H. Merzić, M. Bogdanović, D. Kappler, L. Righetti, and J. Bohg, "Leveraging contact forces for learning to grasp," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3615–3621.

[14] A. Melnik, L. Lach, M. Plappert, T. Korthals, R. Haschke, and H. Ritter, "Tactile sensing and deep reinforcement learning for in-hand manipulation tasks," *IROS Workshop on Autonomous Object Manipulation*, 2019.

[15] S. Christen, S. Stevšić, and O. Hilliges, "Guided deep reinforcement learning of control policies for dexterous human-robot interaction," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 2161–2167.

[16] A. Ecoffet, J. Huizinga, J. Lehman, K. O. Stanley, and J. Clune, "Go-explore: a new approach for hard-exploration problems," *arXiv preprint arXiv:1901.10995*, 2019.

[17] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, O. P. Abbeel, and W. Zaremba, "Hindsight Experience Replay," in *Advances in Neural Information Processing Systems*, 2017, pp. 5048–5058.

[18] S. H. Huang, M. Zambelli, J. Kay, M. F. Martins, Y. Tassa, P. M. Pilarski, and R. Hadsell, "Learning gentle object manipulation with curiosity-driven deep reinforcement learning," *CoRR*, vol. abs/1903.08542, 2019.

[19] B. Wu, I. Akinola, J. Varley, and P. K. Allen, "Mat: Multi-fingered adaptive tactile grasping via deep reinforcement learning," in *Conference on Robot Learning*, 2020, pp. 142–161.

[20] A. Aubret, L. Matignon, and S. Hassas, "A survey on intrinsic motivation in reinforcement learning," *arXiv preprint arXiv:1908.06976*, 2019.

[21] L.-J. Lin, "Self-improving reactive agents based on reinforcement learning, planning and teaching," *Machine learning*, vol. 8, no. 3-4, pp. 293–321, 1992.

[22] 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," in *International Conference on Learning Representations*, 2016.

[23] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," in *International Conference on Learning Representations*, 2016.

[24] R. Zhao and V. Tresp, "Energy-based hindsight experience prioritization," *Conference on Robot Learning*, pp. 113–122, 2018.

[25] M. Plappert, M. Andrychowicz, A. Ray, B. McGrew, B. Baker, G. Powell, J. Schneider, J. Tobin, M. Chociej, P. Welinder, V. Kumar, and W. Zaremba, "Multi-goal reinforcement learning: Challenging robotics environments and request for research," *CoRR*, vol. abs/1802.09464, 2018.

[26] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," *arXiv preprint arXiv:1606.01540*, 2016.

[27] 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*. IEEE, 2012, pp. 5026–5033.

[28] W. C. Lewis-II, M. Moll, and L. E. Kavraki, "How much do unstated problem constraints limit deep robotic reinforcement learning?" *CoRR*, vol. abs/1909.09282, 2019.

[29] S. Narvekar, B. Peng, M. Leonetti, J. Sinapov, M. E. Taylor, and P. Stone, "Curriculum learning for reinforcement learning domains: A framework and survey," *J. Mach. Learn. Res.*, vol. 21, pp. 181:1–181:50, 2020.

[30] A. Rajeswaran*, V. Kumar*, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations," in *Proceedings of Robotics: Science and Systems (RSS)*, 2018.

[31] T. Pham, A. Kheddar, A. Qammaz, and A. A. Argyros, "Towards force sensing from vision: Observing hand-object interactions to infer manipulation forces," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. IEEE Computer Society, 2015, pp. 2810–2819.