

Robot Learning on the Job: Human-in-the-Loop Autonomy and Learning During Deployment

Journal Title
XX(X):1–15
©The Author(s) 2024
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/

SAGE

Huihan Liu¹, Soroush Nasiriany¹, Lance Zhang¹, Zhiyao Bao¹, Yuke Zhu¹

Abstract

With the rapid growth of computing powers and recent advances in deep learning, we have witnessed impressive demonstrations of novel robot capabilities in research settings. Nonetheless, these learning systems exhibit brittle generalization and require excessive training data for practical tasks. To harness the capabilities of state-of-the-art robot learning models while embracing their imperfections, we present Sirius, a principled framework for humans and robots to collaborate through a division of work. In this framework, partially autonomous robots are tasked with handling a major portion of decision-making where they work reliably; meanwhile, human operators monitor the process and intervene in challenging situations. Such a human-robot team ensures safe deployments in complex tasks. Further, we introduce a new learning algorithm to improve the policy's performance on the data collected from the task executions. The core idea is re-weighting training samples with approximated human trust and optimizing the policies with weighted behavioral cloning. We evaluate Sirius in simulation and on real hardware, showing that Sirius consistently outperforms baselines over a collection of contact-rich manipulation tasks, achieving an 8% boost in simulation and 27% on real hardware than the state-of-the-art methods in policy success rate, with twice faster convergence and 85% memory size reduction. Videos and more details are available at <https://ut-austin-rpl.github.io/sirius/>

Keywords

Human-in-the-loop Learning, Robot Learning, Imitation Learning

1 Introduction

Recent years have witnessed great strides in deep learning techniques for robotics. In contrast to the traditional form of robot automation, which heavily relies on human engineering, these data-driven approaches show great promise in building robot autonomy that is difficult to design manually. While learning-powered robotics systems have achieved impressive demonstrations in research settings (Andrychowicz et al. 2018; Kalashnikov et al. 2018; Lee et al. 2020), the state-of-the-art robot learning algorithms still fall short of generalization and robustness for widespread deployment in real-world tasks. The dichotomy between rapid research progress and the absence of real-world application stems from the lack of performance guarantees in today's learning systems, especially when using black-box neural networks. It remains opaque to the potential practitioners of these learning systems: how often they fail, in what circumstances the failures occur, and how they can be continually enhanced to address them.

To harness the power of modern robot learning algorithms while embracing their imperfections, a burgeoning body of research has investigated new mechanisms to enable effective human-robot collaborations. Specifically, *shared autonomy* methods (Javdani et al. 2015; Reddy et al. 2018) aim at combining human input and semi-autonomous robot control to achieve a common task goal. These methods typically use a pre-built robot controller rather than seeking to improve robot autonomy over time. Meanwhile, recent advances in *interactive imitation learning* (Kelly et al. 2019; Mandlkar et al. 2020c; Ross et al. 2011; Celemin et al.

2022) have aimed to learn policies from human feedback in the learning loop. Although these learning algorithms can improve the overall efficacy of autonomous policies, these policies still fail to meet the performance requirements for real-world deployment.

This work aims at developing a human-in-the-loop learning framework for human-robot collaboration and continual policy learning in deployed environments. We expect our framework to satisfy two key requirements: 1) it ensures task execution to be consistently successful through human-robot teaming, and 2) it allows the learning models to improve continually, such that human workload is reduced as the level of robot autonomy increases. To build such a framework, this idea of *robot learning on the job* resembles the Continuous Integration, Continuous Deployment (CI/CD) principles in software engineering (Shahin et al. 2017). Realizing this idea for learning-based manipulation invites fundamental challenges.

The foremost challenge is developing the infrastructure for human-robot collaborative manipulation. We develop a system that allows a human operator to monitor and intervene the robot's policy execution (see Fig. 1). The human can take over control when necessary and handle challenging situations to ensure safe and reliable task

¹The University of Texas at Austin

Corresponding author:

Huihan Liu, 2317 Speedway, Austin, Texas, USA.
Email: huihanl@utexas.edu

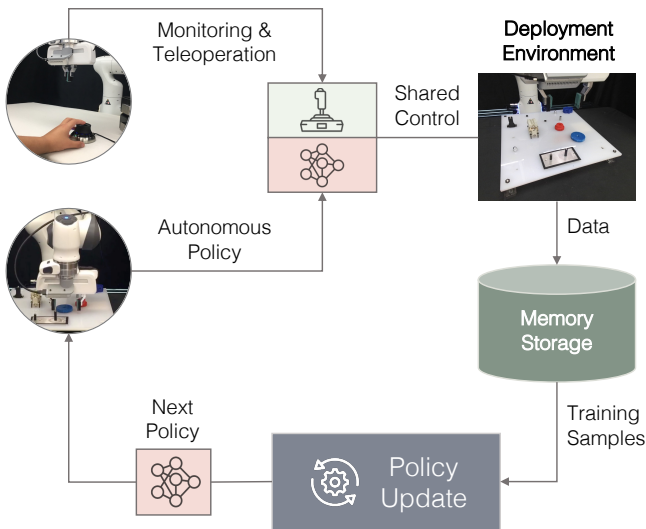


Figure 1. Overview of Sirius, our human-in-the-loop learning and deployment framework. Sirius enables a human and a robot to collaborate on manipulation tasks through shared control. The human monitors the robot’s autonomous execution and intervenes to provide corrections through teleoperation. Data from deployments will be used by our algorithm to improve the robot’s policy in consecutive rounds of policy learning.

execution. Meanwhile, human interventions implicitly reveal the *task structure* and the level of *human trust* in the robot. As recent work (Kelly et al. 2019; Mandlekar et al. 2020c; Hoque et al. 2021) indicates, human interventions inform *when* the human lacks trust in the robot, *where* the risk-sensitive task states are, and *how* to traverse these states. We can thus take advantage of the occurrences of human interventions during deployments as informative signals for policy learning.

The subsequent challenge is updating policies on an ever-growing dataset of shifting distributions. As our framework runs over time, the policy would adapt its behaviors through learning, and the human would adjust their intervention patterns accordingly. Deployment data from human-robot teams can be multimodal and suboptimal. Learning from such deployment data requires us to selectively use them for policy updates. We want the robot to learn from good behaviors to reinforce them and also to recover from mistakes and deal with novel situations. At the same time, we want to prevent the robot from copying bad actions that would lead to failure. Our key insight is that we can assess the importance of varying training data based on human interventions for policy learning.

To this end, we develop a simple yet effective learning algorithm that uses the occurrences of human intervention to re-weight training data. We consider the robot rollouts right before an intervention as “low-quality” (as the human believes the robot is about to fail) and both human demonstrations and interventions as “high-quality” for policy training. We label training samples with different weights and train policies on these samples using weighted behavioral cloning, the state-of-the-art algorithm for imitation learning (Sasaki and Yamashina 2021; Zolna et al. 2020; Xu et al. 2022) and offline reinforcement learning (Wang et al. 2020; Nair et al. 2021; Kostrikov et al. 2021). This supervised learning algorithm lends itself to the

efficiency and stability of policy optimization on our large-scale and growing dataset.

Furthermore, deploying our system in long-term missions leads to two practical considerations: 1) it incurs a heavy burden of memory storage to store all past experiences over a long duration, and 2) a large number of similar experiences may inundate the small subset of truly valuable data for policy training. We thus examine different memory management strategies, aiming at adaptively adding and removing data samples from the memory storage of fixed size. Our results show that even with 15% of the full memory size, we can retain the same level of performance or achieve even better performance than keeping all data, and moreover enables three times faster convergence for rapid model updates between consecutive rounds.

We name our framework Sirius, the star symbolizing our human-robot team with its binary star system. We evaluate Sirius in two simulated and two real-world tasks requiring contact-rich manipulation with precise motor skills. Compared to the state-of-the-art methods of learning from offline data (Nair et al. 2021; Kostrikov et al. 2021; Mandlekar et al. 2021a) and interactive imitation learning (Mandlekar et al. 2020c), Sirius achieves higher policy performance and reduced human workload. Sirius reports an 8% boost in policy performance in simulation and 27% on real hardware over the state-of-the-art methods.

2 Related Work

Human-in-the-loop Learning. A human-in-the-loop learning agent utilizes interactive human feedback signals to improve its performance (Zhang et al. 2019; Cruz and Igarashi 2020; Cui et al. 2021). Human feedback can serve as a rich source of supervision, as humans often have a priori domain information and can interactively guide the agent with respect to its learning progress. Many forms of human feedback exist, such as interventions (Kelly et al. 2019; Spencer et al. 2020; Mandlekar et al. 2020c), preferences (Christiano et al. 2017; Bıyık et al. 2022; Lee et al. 2021; Wang et al. 2022), rankings (Brown et al. 2019), scalar-valued feedback (MacGlashan et al. 2017; Warnell et al. 2018), and human gaze (Zhang et al. 2020). These feedback forms can be integrated into the learning loop through learning techniques such as policy shaping (Knox and Stone 2009; Griffith et al. 2013) and reward modeling (Daniel et al. 2014; Leike et al. 2018), enabling model updates from asynchronous policy iteration loops (Chisari et al. 2021).

Within the context of robot manipulation, one approach is to incorporate human interventions in imitation learning algorithms (Kelly et al. 2019; Spencer et al. 2020; Mandlekar et al. 2020c; Dass et al. 2022). Another approach is to employ deep reinforcement learning algorithms with learned rewards, either from preferences (Lee et al. 2021; Wang et al. 2022) or reward sketching (Cabi et al. 2019). While these methods have demonstrated higher performance compared to those without humans in the loop, they require a large amount of supervision from humans and also fail to incorporate human control feedback in deployment into the learning loop again to improve model performance. In contrast, we specifically consider the above scenarios which are critical to real-world robotic systems.

Shared Autonomy. Human-robot collaborative control is often necessary for real-world tasks when we do not have full robot autonomy while full human teleoperation control is burdensome. In shared autonomy (Dragan and Srinivasa 2013; Javdani et al. 2015; Gopinath et al. 2017; Reddy et al. 2018), the control of a system is shared by a human and a robot to accomplish a common goal (Tan et al. 2021). The existing literature on shared autonomy focuses on efficient collaborative control from human intent prediction (Dragan and Srinivasa 2012; Muelling et al. 2015; Perez-D’Arpino and Shah 2015). However, they do not attempt to learn from human intervention feedback, so there is no policy improvement. We examine a context similar to that of shared autonomy where a human is involved during the actual deployment of the robot system; however, we also put human control in the feedback loop and use them to improve the learning itself.

Learning from Offline Data. An alternative to the human-in-the-loop paradigm is to learn from fixed robot datasets via imitation learning (Pomerleau 1989; Zhang et al. 2018; Mandlkar et al. 2020b; Florence et al. 2021) or offline reinforcement learning (offline RL) (Levine et al. 2020; Fujimoto et al. 2019; Kumar et al. 2020; Kidambi et al. 2020; Yu et al. 2020, 2021; Mandlkar et al. 2020a; Kostrikov et al. 2021). Offline RL algorithms, particularly, have demonstrated promise when trained on large diverse datasets with suboptimal behaviors (Singh et al. 2020; Kumar et al. 2022; Ajay et al. 2021). Among a number of different methods, advantage-weighted regression methods (Wang et al. 2020; Nair et al. 2021; Kostrikov et al. 2021) have recently emerged as a popular approach to offline RL. These methods use a weighted behavior cloning objective to learn the policy, using learned advantage estimates as the weight. In this work, we also use weighted behavior cloning; however, we explicitly leverage human intervention signals from our online human-in-the-loop setting to obtain weights rather than using task rewards to learn advantage-based weights. We show that this leads to superior empirical performance for our manipulation tasks.

3 Background and Overview

3.1 Problem Formulation

We formulate a robot manipulation task as a Markov Decision Process $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P}, p_0, \gamma)$ representing the state space, action space, reward function, transition probability, initial state distribution, and discount factor. In this work, we adopt an intervention-based learning framework in which the human can choose to intervene and take control of the robot. Given the current state $s_t \in \mathcal{S}$, the robot action $a_t^R \in \mathcal{A}$ is drawn from the policy $\pi_R(\cdot | s_t)$, and the human can override this action with a human action $a_t^H \in \mathcal{A}$. The policy π for the human-robot team can thus be formulated as:

$$\pi(\cdot | s_t) = I_H(s_t)\pi_H(\cdot | s_t) + (1 - I_H(s_t))\pi_R(\cdot | s_t),$$

where I_H is a binary indicator function of human interventions and π_H is the implicit human policy. Our learning objective is two-fold: 1) we want to improve the level of robot autonomy by finding the

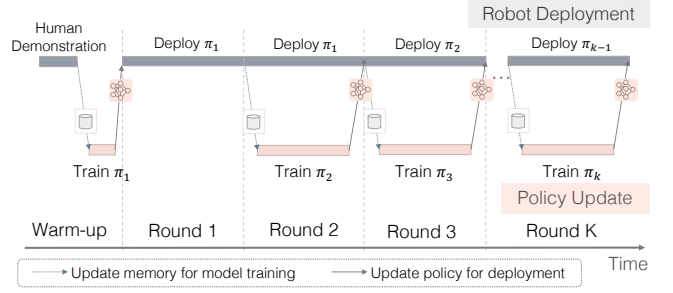


Figure 2. Illustration of the workflow in Sirius. Robot deployment and policy update co-occur in two parallel threads. Deployment data are passed to policy training, while a newly trained policy is deployed to the target environment for task execution.

autonomous policy π_R that maximizes the cumulative rewards $\mathbb{E}_{\pi_R} [\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, s_{t+1})]$, and 2) we want to minimize the human’s workload in the system, *i.e.*, the expectation of interventions $\mathbb{E}_{\pi} [I_H(s_t)]$ under the state distribution induced by the team policy π .

3.2 Weighted Behavioral Cloning Methods

We aim to learn a robot policy π_R with the deployment data to enhance robot autonomy and reduce human costs in human-robot collaboration. Weighted Behavioral Cloning (BC) has recently become one promising approach to learning policies from multimodal and suboptimal data. In standard BC methods, we train a model to mimic the action for each state in the dataset. The objective is to learn a policy π_R parameterized by θ that maximizes the log-likelihood of actions a conditioned on the states s :

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{(s,a) \sim \mathcal{D}} [\log \pi_{\theta}(a | s)], \quad (1)$$

where (s, a) are samples from the dataset \mathcal{D} . For weighted BC, the log-likelihood term of each (s, a) pair is scaled by a weight function $w(s, a)$, which assigns different importance scores to different samples:

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{(s,a) \sim \mathcal{D}} [w(s, a) \log \pi_{\theta}(a | s)]. \quad (2)$$

The weighted BC framework lays the foundation of several state-of-the-art methods for offline reinforcement learning (RL) (Nair et al. 2021; Kostrikov et al. 2021; Wang et al. 2020). Different weight assignments differentiate high-quality samples from low-quality ones, such that the algorithm prioritizes high-quality samples for learning. In particular, advantage-based offline RL algorithms calculate weights as $w(s, a) = f(Q^{\pi}(s, a))$, where $f(\cdot)$ is a non-negative scalar function related to the learned advantage estimates $A^{\pi}(s, a)$. High-advantage samples indicate that their actions likely contribute to higher future returns and, therefore, should be weighted more. Through the sample-weighting scheme, these methods filter out low-advantage samples and focus on learning from the higher-quality ones in the dataset. Nonetheless, effectively learning value estimates can be challenging in practice, especially when the dataset does not cover a sufficiently wide distribution of states and actions—a challenge highlighted by prior work (Gulcehre et al. 2020; Fu et al. 2020). In the

deployment setting, the data only constitute successful trajectories that complete the task eventually. Empirically, we find in Section 5 that the nature of our deployment data makes today’s offline RL methods struggle to learn values.

In contrast to the value learning framework, some prior works (Mandlekar et al. 2020c; Gandhi et al. 2022a; Chisari et al. 2021) have developed weighted BC approaches that are specialized for the human-in-the-loop setting. In particular, Mandlekar et al. (2020c) proposes Intervention-weighted Regression (IWR) which designs weights based on whether a sample is a human intervention. Inspired by these prior works, we introduce a simple yet practical weighting scheme that harnesses the unique properties of deployment data to learn performant agents. We elaborate on our weighting scheme in the following section.

4 Sirius: Human-in-the-loop Learning and Deployment

We present Sirius, our human-in-the-loop framework that learns and deploys continually improving policies from human and robot deployment data. First, we define the human-in-the-loop deployment setting and give an overview of our system design. Next, we describe our weighting scheme, which can learn effective policies from mixed, multi-modal data throughout deployment. Finally, we introduce memory management strategies that reduce the computational complexities of policy learning and improve the efficiency of the system.

4.1 Human-in-the-loop Deployment Framework

Our human-in-the-loop system aims to constantly learn from the deployment experience and human corrective feedback so as to obtain a high-performing robot policy and reduce human workload over time. It consists of two components that happen simultaneously: Robot Deployment and Policy Update. In Robot Deployment (top thread in Fig. 2), the robot performs task executions with human monitoring; in Policy Update (bottom thread), the system improves the policy with the deployment data for the next round of task execution.

The system starts with an initial policy in the warm-up phase, where we bootstrap a robot policy π_1 trained on a small number of human demonstrations. Initially, the memory buffer comprises a set of human demonstration trajectories $\mathcal{D}^0 = \{\tau_j\}$, where each trajectory $\tau_j = \{s_t, a_t, r_t, c_t = \text{demo}\}$ consists of the states, actions, task rewards, and the data class type flag c_t indicating whether these trajectories are human demonstrations.

Upon training the initial policy π_1 , we deploy the robot to perform the task, and in the process, we collect a set of trajectories to improve the policy. A human operator who continuously monitors the robot’s execution will intervene based on whether the robot has performed or will perform suboptimal behaviors. Note that we adapt human-gated control (Kelly et al. 2019) rather than robot-gated control (Hoque et al. 2021) to guarantee task execution success and trustworthiness of the system for real-world deployment. Through this process, we obtain a new dataset

\mathcal{D}' of trajectories $\tau_j = \{s_t, a_t, r_t, c_t\}$, where c_t either indicates the transition is a *robot action* ($c_t = \text{robot}$) or a *human intervention* ($c_t = \text{intv}$). We append this data to the existing memory buffer collected so far $\mathcal{D}^1 \leftarrow \mathcal{D}^0 \cup \mathcal{D}'$, and train a new policy π_2 on this new dataset.

In subsequent rounds, we deploy the robot to collect new data while simultaneously updating the policy. We define “Round” as the interval for policy update and deployment: It consists of the completion of training for one policy, and at the same time, the collection of one set of deployment data. In Round i , we train policy π_i using all previously collected data. Maintaining the previous rounds of collected data allows us to retain a diverse coverage of the state distribution, which has the potential benefit of regularizing the policy and keeping the policy robust (Mandlekar et al. 2018; Hoque et al. 2021). Meanwhile, the robot is continuously being deployed using the current latest policy π_{i-1} , and gathered deployment data \mathcal{D}' . At the end of round i we append this data to the existing memory buffer collected so far $\mathcal{D}^i \leftarrow \mathcal{D}^{i-1} \cup \mathcal{D}'$ and train a new policy π_{i+1} on this aggregated dataset.

Our system aggregates data from deployment environments over long-term deployments. This presents a unique set of challenges: first, the generated data comes from **mixed distributions** consisting of robot policy actions, human interventions, and human demonstrations; also, the system produces data that is **constantly growing in size**, imposing memory burden and computational inefficiency for learning algorithms. We address these challenges in the following sections.

4.2 Human-in-the-loop Policy Learning

We present a simple yet effective learning method that takes advantage of the unique characteristics of deployment data to learn effective policies. We have a critical insight that human interventions provide informative signals of human trust and human judgement of the robot executions, which we will use to guide the design of our algorithm. The core idea of our approach is to harness the structure of the human correction feedback to re-weight training samples based on an approximate quality score. With these weighted samples, we train the policy with the weighted behavioral cloning method to learn the policy on mixed-quality data. Our approach is motivated by two insights on how the human intervention structure could be used.

Our first intuition is that human intervention samples are highly important samples and should be prioritized in learning. Human-operated samples are expensive to obtain and should be optimized in general, but human intervention occurs in situations where the robot is unable to complete the task and requires help. These are risk-sensitive task states, so data in these regions are highly valuable. Therefore, these state-action pairs should be ranked high by the weighting function, and we should upweight the human intervention samples such that these samples will positively influence learning more.

Moreover, we should not only make use of *what* human samples to use, but also *when* the human samples take place. We make the critical observation that when the robot operates autonomously, it usually performs reasonable behaviors. But when it demands interventions, it is when the robot

Algorithm 1 Human-in-the-loop Learning at Deployment**Notations**

L : memory buffer maximum fixed size
 X : maximum deployment rounds
 M : number of initial human demonstration trajectories
 K : number of rollout episodes in each deployment round
 b : batch size
 n : number of gradient steps in each learning round
 α : policy learning rate

 \triangleright *warmstart phase*

Collect M human demonstrations τ_1, \dots, τ_M

$\mathcal{D}^0 \leftarrow \{\tau_1, \dots, \tau_M\}$

Initialize BC policy π_1^θ :

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{(s,a) \sim \mathcal{D}^0} [\log \pi_1^\theta(a | s)]$$

 \triangleright *initial deployment data*

$\mathcal{D}^1 \leftarrow \text{DEPLOYMENT}(\pi_1^\theta, \mathcal{D}^0)$

 \triangleright *deployment-learning loop*

for $i \leftarrow 1$ to X **do**

 Run in parallel:

$\mathcal{D}^{i+1} \leftarrow \text{DEPLOYMENT}(\pi_i^\theta, \mathcal{D}^i)$

$\pi_{i+1}^\theta \leftarrow \text{LEARNING}(\mathcal{D}^i)$

 \triangleright *deployment thread*

function DEPLOYMENT(π_θ, \mathcal{D})

 Collect rollout episodes $\tau_1, \dots, \tau_K \sim p_{\pi_\theta}(\tau)$

$\mathcal{D}^+ \leftarrow \mathcal{D} \cup \{\tau_1, \dots, \tau_K\}$

if $|\mathcal{D}^+| > L$ **then**

 Discard trajectories in \mathcal{D}^+ s.t. $|\mathcal{D}^+| \leq L$

 with a memory management strategy (in 4.3)

return \mathcal{D}^+

 \triangleright *learning thread*

function LEARNING(\mathcal{D})

 Initialize π_θ

for each class c **do**

$\mathcal{D}_c \leftarrow \{(s, a, c') \in \mathcal{D} \mid c' = c\}$

$P(c) \leftarrow |\mathcal{D}_c|/|\mathcal{D}|$

 Obtain $P^*(c)$ (see 4.4)

for n gradient steps **do**

 Sample mini-batch $(s^i, a^i, c^i)_{i=1}^b \sim \mathcal{D}$

 Compute $w(s^i, a^i, c^i) \leftarrow \frac{P^*(c^i)}{P(c^i)}$ for the mini-batch

$\mathcal{L}_\pi(\theta) = -\frac{1}{b} \sum_i [w(s^i, a^i, c^i) \cdot \log \pi_\theta(a^i | s^i)]$

$\theta \leftarrow \theta - \alpha \nabla_\theta \mathcal{L}_\pi(\theta)$

return π_θ

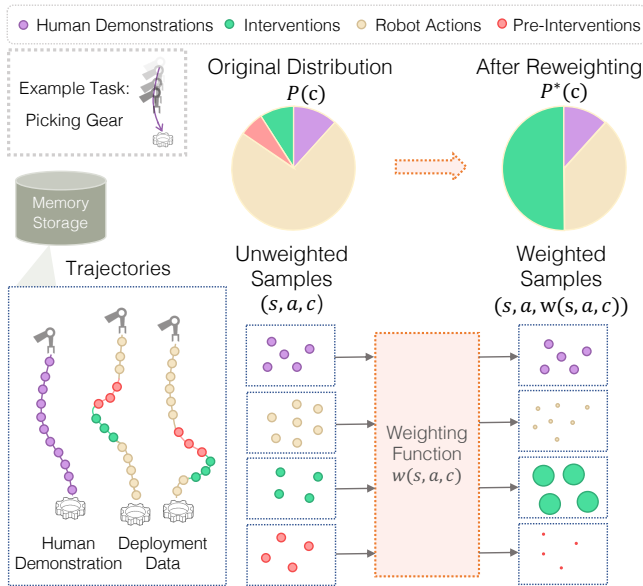


Figure 3. Overview of our human-in-the-loop learning model. We maintain an ever-growing database of diverse experiences spanning four categories: human demonstrations, autonomous robot data, human interventions, and transitions preceding interventions which we call pre-interventions. We set weights according to these four categories, with a high weight given to interventions over other categories. We use these weighted samples to continually learn vision-based manipulation policies during deployment.

has made mistakes or has performed suboptimal behaviors. Therefore, human interventions implicitly signify human value judgment of the robot behavior—the samples before

human interventions are less desirable and of lower quality. We aim to minimize their impact on learning.

With these insights, we devise a weighting scheme according to intervention-guided data class types. Recall that each sample (s, a, r, c) in our dataset contains a data class type c , indicating whether the sample denotes a human demonstration action, robot action, or human intervention action. To incorporate the timing of human interventions, we distinguish and penalize the samples taken prior to each human intervention. We define the segment preceding each human intervention as a separate class, pre-intervention (`preintv`) (see Fig. 3). This classification is based on the implicit human evaluation from the human partner, thresholding the robot samples into either normal robot samples or suboptimal `preintv` samples. Overall, this yields four class types $c \in \{\text{demo}, \text{intv}, \text{robot}, \text{preintv}\}$. We choose to keep the robot action data and learn from them because adding robot action data to the training set regularizes the training process, as shown by Mandlekar et al. (2020c). Removing the robot action data reduces the method to be HG-Dagger (Kelly et al. 2019), which is known to perform worse in several findings (Mandlekar et al. 2018; Hoque et al. 2021; Li et al. 2022).

We derive the weight for each individual sample according to its corresponding class type c . Suppose the dataset \mathcal{D} has total number of samples N , and n_c is the number of samples that is class c . We use \mathcal{D}_c to represent the collection of samples of class c in \mathcal{D} . The original class distribution is $P(c) = n_c/N$ for class c , and the unweighted BC objective

under this distribution is:

$$\begin{aligned} & \arg \max_{\theta} \mathbb{E}_{(s,a) \sim \mathcal{D}} [\log \pi_{\theta}(a | s)] \\ &= \arg \max_{\theta} \mathbb{E}_{P(c)} \mathbb{E}_{(s,a) \sim \mathcal{D}_c} [\log \pi_{\theta}(a | s)]. \end{aligned} \quad (3)$$

In a long-term deployment setting, most data will be robot actions, and human interventions usually constitute a small ratio of the dataset samples since interventions only happen at critical regions in a trajectory; the pre-intervention samples constitute a small but non-negligible proportion which can have detrimental effects (see Fig. 3, left pie chart). We will now change the class distribution to a new distribution $P^*(c)$, in which we increase the ratio of human intervention samples and decrease the ratio of the pre-intervention samples (see Fig. 3, right pie chart). Under this new distribution, the weight $w(s, a, c)$ of the training samples in each individual class c can be equivalently set as $w(s, a, c) = P^*(c)/P(c)$ by the rule of importance sampling. We outline the details of our specific distribution $P^*(c)$ in Sec. 4.4. This way, we obtain the sample weights for weighted BC, leveraging the inherent structure of human-robot team data.

4.3 Memory Management

As the deployment continues and the dataset increases, large data slows down training convergence and takes up excessive memory space. We hypothesize that forgetting (routinely discarding samples from memory) helps prioritize important and useful experiences for learning, speeding up convergence and even further improving policy. Moreover, the right kind of forgetting matters, since we want to preserve the data that is most beneficial to learning. Therefore, we would like to investigate the following question—with limited data storage and a never-ending deployment data flow, how do we absorb the most useful data and preserve more valuable information for learning?

We assume that we have a fixed-size memory buffer that replaces existing samples with new ones when full. We consider five strategies for managing the memory buffer of deployment data. Each strategy tests out a different hypothesis listed below:

1. **LFI** (Least-Frequently-Intervened): first reject samples from trajectories with the least interventions. *(Preserving the most human intervened trajectories keeps the most valuable human and critical state examples, which helps learning the most.)*
2. **MFI** (Most-Frequently-Intervened): first reject samples from trajectories with the most interventions. *(Successful, unintervened robot trajectories yield higher quality data for learning compared to those that require intervention.)*
3. **FIFO** (First-In-First-Out): reject samples in the order that they were added to the buffer. *(More recent data from a higher performing policy are higher quality data for learning.)*
4. **FILO** (First-In-Last-Out): reject the most recently added samples first. *(Initial data from a worse performing policy have greater state coverage and data diversity for learning.)*

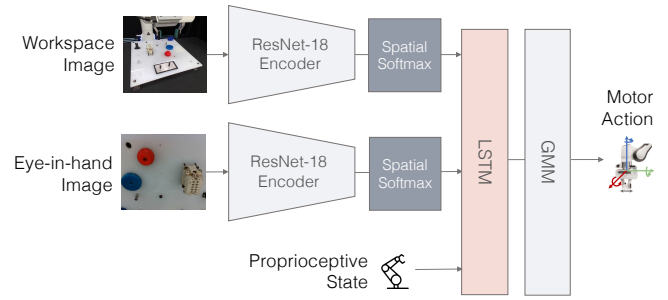


Figure 4. Policy Architecture. Our vision-based policy uses BC-RNN as our policy backbone. Our inputs are workspace camera image and eye-in-hand camera image, as well as robot proprioceptive states.

5. **Uniform**: reject samples uniformly at random.

(Uniformly selecting trajectories can yield a balanced mix of diverse samples, aiding in the learning process.)

With the intervention-guided weighting scheme for policy update and memory management strategies, we present the overall workflow of human-in-the-loop learning in deployment in Algorithm 1.

4.4 Implementation Details

For the robot policy (see Fig. 4), we adopt BC-RNN (Mandlekar et al. 2021a), the state-of-the-art behavioral cloning algorithm, as our model backbone. We use ResNet-18 encoders (He et al. 2016) to encode third person and eye-in-hand images (Mandlekar et al. 2021a, 2020b). We concatenate image features with robot proprioceptive state as input to the policy. The network outputs a Gaussian Mixture Model (GMM) distribution over actions. GMM is used to handle the multimodality of human actions: GMM generates different modes and creates a probability distribution of the different modes the actions can take, and therefore is more flexible and effective than deterministic actions, or a single gaussian distribution of actions (Chernova and Veloso 2007; Mandlekar et al. 2021b).

For our intervention-guided weighting scheme, we set $P^*(\text{intv}) = \frac{1}{2}$. The 50% ratio is adapted from prior work (Mandlekar et al. 2020c) that increases the weight of intervention to a reasonable level. We conduct an ablation study in Section 5 how changing $P^*(\text{intv})$ affects the policy performance. We set $P^*(\text{preintv}) = 0$, essentially nullifying the impact of pre-intervention samples. The demo weight maintains the true ratio of demonstration samples in the dataset: $P^*(\text{demo}) = P(\text{demo})$. Finally, $P^*(\text{robot})$ adjusts itself accordingly. Under this new distribution, we implicitly decrease the proportion of the `robot` class (see Fig. 3) due to increasing the proportion of the `intv` class. Note that the ratio of the demonstration remains unchanged as they are still important and useful samples to learn from, especially during initial rounds of updates when the robot generates lower-quality data. This is in contrast to IWR by Mandlekar et al. (Mandlekar et al. 2020c), which treats all non-intervention samples as a single class, thus lowering the contribution of demonstrations from their unweighted ratio. The weight for each individual sample is $w(s, a, c) = P^*(c)/P(c)$, as discussed in Section 4.2.




Figure 5. Human Robot Teaming. Left: The robot executes the task by default while a human supervises the execution. Right: When the human detects undesirable robot behavior, the human intervenes.

We set a segment of length ℓ before each human intervention as the class `preinterv`. The optimal choice on the hyperparameter ℓ depends on the *human reaction time*, which quantifies how fast the human operator reacted to the robot’s undesired behavior. Prior works (Stiber et al. 2022; Spencer et al. 2020) indicate that a response delay exists between the time the robot starts to perform mistakes and the time human actually perform corrective interventions. Our empirical observation based on our human operator shows an average reaction time of 2 seconds, roughly corresponding to the time of 15 robot actions. We thus set $\ell = 15$.

5 Experiments

In our experiments, we seek to answer the following research questions: 1) How effective is Sirius in improving autonomous robot policy performance over time? 2) Can this system reduce human workload over time? 3) How do the individual design choices in our learning algorithm affect overall performance? and 4) Which memory management strategy is most effective for learning with constrained memory storage?

5.1 Human-Robot Teaming

We illustrate the actual human-robot teaming process during human-in-the-loop deployment in Figure 5. The robot executes a task (e.g., gear insertion) by default while a human supervises the execution. In this gear insertion scenario, the expected robot behavior is to pick up the gear and insert it down the gear shaft. When the human detects undesirable robot behavior (e.g., gear getting stuck), the human intervenes by taking over control of the robot. The human directly passes in action commands to perform the desired behavior. When the human judges that the robot can continue the task, the human passes control back to the robot.

To enable effective shared human control of the robot, we seek a teleoperation interface that (1) enables humans to control the robot effectively and intuitively and (2) switches between robot and human control immediately once the human decides to intervene or pass the control back to the robot. To this end, we employ SpaceMouse¹ control. The human operator controls a 6-DoF SpaceMouse and passes the position and orientation of the SpaceMouse as action commands. The user can pause when monitoring the computer screen by pressing a button, exert control until the robot is back to an acceptable state, and pass the control back to the robot by stopping the motion on the SpaceMouse.

5.2 Tasks

We design a set of simulated and real-world tasks that resemble common industrial tasks in manufacturing and logistics. We consider long-horizon tasks that require precise contact-rich manipulation, necessitating human guidance. For all tasks, we use a Franka Emika Panda robot arm equipped with a parallel jaw gripper. Both the agent and human control the robot in task space. We use a SpaceMouse as the human interface device to intervene.

We systematically evaluate the performance of our method and baselines in the robosuite simulator (Zhu et al. 2020). We choose the two most challenging contact-rich manipulation tasks in the robomimic benchmark (Mandlekar et al. 2021a):

Nut Assembly. The robot picks up a square nut from the table and inserts the nut into a column.

Tool Hang. The robot picks up a hook piece and inserts it into a very small hole, then hangs a wrench on the hook. As noted in robomimic (Mandlekar et al. 2021a), this is a difficult task requiring precise and dexterous control.

In the real world, we design two tasks representative of industrial assembly and food packaging applications:

Gear Insertion. The robot picks up two gears on the NIST board and inserts each of them onto the gear shafts.

Coffee Pod Packing. The robot opens a drawer, places a coffee pod into the pod holder, and closes the drawer.

5.3 Baselines

We compare our method with the state-of-the-art human-in-the-loop learning method for robot manipulation, Intervention Weighted Regression (**IWR**) (Mandlekar et al. 2020c). Furthermore, to ablate the impacts of algorithms versus data distributions, we compare the state-of-the-art imitation learning algorithm **BC-RNN** (Mandlekar et al. 2021a) and offline RL algorithm Implicit Q-Learning (**IQL**) (Kostrikov et al. 2021). We run these two latter baselines on the deployment data generated by our method for a fair comparison.

Our codebase is based on robomimic (Mandlekar et al. 2021a), a recent open-source project that benchmarks a range of learning algorithms on offline data. We standardize all methods with the same state-of-the-art policy architectures and hyperparameters from robomimic. The architectural design includes ResNet-18 image encoders, random cropping for image augmentation, GMM head, and the same training procedures. The list of hyperparameter choices is presented in Table 1. For all BC-related methods, including Ours, IWR, and BC-RNN, we use the same BC-RNN architecture specified in Table 2.

For IQL (Kostrikov et al. 2021), we reimplemented the method in our robomimic-based codebase to keep the policy backbone and common architecture the same across all methods. Our implementation is based on the publicly available PyTorch implementation of IQL².

We follow the paper’s original design with some slight modifications. In particular, the original IQL uses the sparse reward setting where the reward is based on task success. We add a denser reward for IQL to incorporate information on human intervention. To mimic the intervention-guided weights for IQL, we use the following rewards: $r = 1.0$ upon task success, $r = 0.25$ for intervention states, $r = -0.25$ for

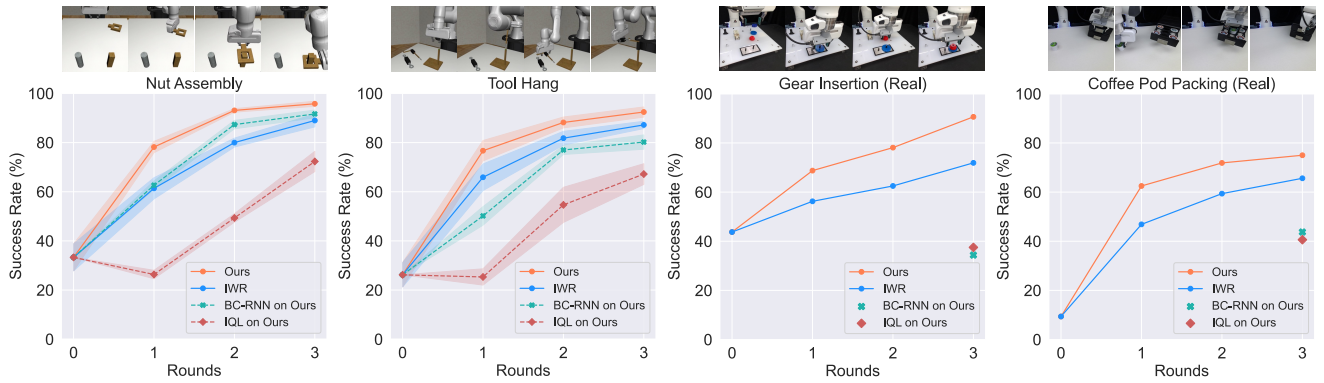


Figure 6. Quantitative evaluations. We compare our method with human-in-the-loop learning, imitation learning, and offline reinforcement learning baselines. Our results in simulated and real-world tasks show steady performance improvements of the autonomous policies over rounds. Our model reports the highest performance in all four tasks after three rounds of deployments and policy updates. Solid line: human-in-the-loop; dashed line: offline learning on data from our method.

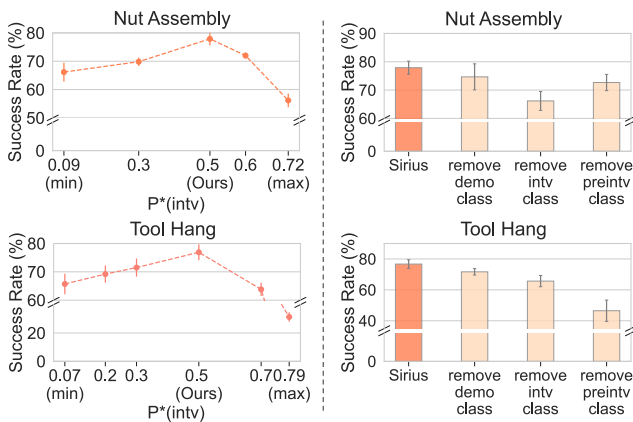


Figure 7. (Left) Ablation on intervention ratio weight. We show how policy performance first increase then decrease as $P^*(intv)$ increases, peaking at $P^*(intv) = 0.5$. **(Right) Ablation on weight function design.** Our results show that removing each class label hurts model performance.

Table 1. Common hyperparameters

Hyperparameter	Value
GMM number of modes	5
Image encoder	ResNet-18
Random crop ratio	90% of image height
Optimizer	Adam
Batch size	16
# Training steps per epoch	500
# Total training epochs	1000
Evaluation interval (in epoch)	50

pre-intervention states, and $r = 0$ for all other states. We found that this version of IQL outperforms the default sparse reward setting. Note that in contrast to our method, IQL requires additional information on task rewards, which may be expensive to obtain in real-world settings. We list the hyperparameters for IQL baseline in Table 3.

Table 2. BC backbone hyperparameters

Hyperparameter	Value
RNN hidden dim	1000
RNN sequence length	10
# of LSTM layers	2
Learning rate	$1e-4$

Table 3. IQL hyperparameters

Hyperparameter	Value
Reward scale	1.0
Termination	false
Discount factor r	0.99
Beta β	1.0
Adv filter	exponential
V function quantile	0.75
Actor lr	$1e-4$
Actor lr decay factor	0.1
Actor mlp layers	[1024, 1024]
Critic lr	$1e-4$
Critic lr decay factor	0.1
Critic mlp layers	[1024, 1024]

5.4 Evaluation Protocol

To provide a fair comparison with existing human-in-the-loop methods, we follow the round update protocol established by prior work (Mandlekar et al. 2020c; Kelly et al. 2019): three rounds of policy learning and deployment, where each round of deployment runs until the number of intervention samples reaches one third of the initial human demonstration samples. The motivation for this protocol is to ensure that the total amount of intervention data across all rounds matches the demonstration data, allowing us to evaluate the ability to learn from the same amount of human

interventions across different human-in-the-loop baselines. We benchmark human-in-the-loop deployment systems in two aspects: 1) **Policy Performance**. Our human-robot team achieves a reliable task success of 100%. Here we evaluate the success rate of the autonomous policy after each round of model update; and 2) **Human Workload**. We measure human workload as the percentage of intervention in the trajectories in each round. While we acknowledge that human workload is a complex domain that can benefit from qualitative metrics such as the NASA-TLX (Task Load Index) (Hart and Staveland 1988), our current evaluation follows the convention of prior human-in-the-loop literature by focusing on quantitative measures (Hoque et al. 2021; Li et al. 2022; Hoque et al. 2022). A comprehensive study using these advanced methodologies could be done in future research. We perform rigorous evaluations of policy performance as follows:

- *Simulation experiments*: We evaluate the success rate of each method across 3 seeds. For each seed, we evaluate the success rate at a set of regularly spaced training checkpoints and record the average over the top three performing checkpoints to avoid outliers. For each checkpoint, we evaluate whether the agent successfully completed the task over 100 trials.
- *Real-world experiments*: We evaluate each method for one seed due to the high time cost for real robot evaluation. Since real robot evaluations are subject to noise and variation across checkpoints, we first perform an initial evaluation of different checkpoints (5 checkpoints) for each method, evaluating each of them for a small number of trials (5 trials). For the checkpoint that gives the best initial quantitative behavior, we perform 32 trials and report the success rate over them.

5.5 Experiment Results

Quantitative Results. We show in Fig. 6 that our method significantly outperforms the baselines on our evaluation tasks. Our method consistently outperforms IWR over the rounds. We attribute this difference to our fine-grained weighting scheme, enabling the method to better differentiate high-quality and suboptimal samples. This advantage over IWR cascades across the rounds, as we obtain a better policy, which in turn yields better deployment data.

We also show that our method significantly outperforms the BC-RNN and IQL baselines under the same dataset distribution. This highlights the importance of our weighting scheme — BC-RNN performs poorly due to copying the suboptimal behaviors in the dataset, while IQL fails to learn values as weights that yield effective policy performance.

Ablation Studies. We perform an ablation study to examine the contribution of each component in our weighting scheme in Fig. 7 (Right). We study how removing each class, *i.e.*, treating each class as the robot action class (and thus removing the special weight for that class), affects the policy performance:

- **remove demo class**: not preserving the true ratio of demo class, which lowers its contribution (see 4.4).
- **remove intv class**: not upweighting the `intv` class, which is equivalent to (`min`) in Fig. 7 (Left).




Figure 8. Ablation on memory management strategies. We study the five different strategies introduced in Section 4.3. LFI (discarding least frequently intervened trajectories) matches and even yields better performance over keeping all data samples (Base) while taking much less memory storage.




Figure 9. Human Intervention Sample Ratio. We evaluate the human intervention sample ratio for the four tasks. The human intervention sample ratio decreases over deployment round updates. Our methods have a larger reduction in human intervention ratio as compared with IWR.

- **remove preintv class**: not downweighting the `preintv` class but treating it as `robot` class.

We run each ablated version of our method on Round 1 data for the simulation tasks. We choose Round 1 data for this study because they are generated from the initial BC-RNN policy rather than biased toward data generated from our method. As shown in Fig. 7 (Right), removing any class weight hurts the policy performance. This shows the effectiveness of our fine-grained weighting scheme, where each class contributes differently to the learning of the deployment data.

We also conduct an in-depth study on the influence of human intervention reweighting ratio $P^*(\text{intv})$. In the unweighted distribution, the human intervention samples take up a small proportion of the dataset size, which we denote as the *minimum ratio*; the maximum ratio it can take is to nullify the proportion of `robot` samples altogether (so that the dataset only constitutes human demonstrations and human interventions). We run our method with a different ratio ranging from minimum to maximum using Round 1 data on both simulation tasks. The specific range for Nut Assembly and Tool Hang can be found in Fig. 7 (Left). The overall trend is that the policy performance peaks at $P^*(\text{intv}) = 0.5$, and is worse when $P^*(\text{intv})$ gets larger or smaller. Our intuition is that if the intervention ratio is too small, we are not making the best use of the intervention samples; if it is too large, it will limit the diversity of training data. Either way has an adversarial effect.

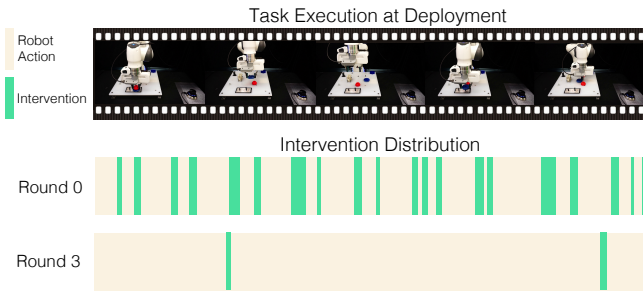


Figure 10. Human Intervention Distribution. The two color bars represent the time duration over 10 consecutive trajectories and whether each step is autonomous robot action (yellow) or human intervention (green). In Round 1, much human intervention is needed to handle difficult situations. In Round 3, the policy needs very little human intervention, and the robot can run autonomously most of the time.

Analysis on Memory Management. We compare the effectiveness of Memory Management strategies in Section 4.3 at deployment. Fig. 8 shows the result of memory size reduction on the two simulation tasks in Round 3, where the Nut Assembly accumulated 3000+ trajectories and the Tool Hang task 1600+ trajectories. By capping our memory buffer size at 500 trajectories, we manage to reduce memory size to a much small proportion of the original dataset size (15% for Nut Assembly and 30% for Tool Hang).

Among all of the strategies, LFI (discarding least frequently intervened trajectories) is the only strategy that matches and even yields better performance over keeping all data samples (Base). In addition to minimizing storage requirements, LFI also improves learning efficiency. Under LFI, the policy converged twice as fast as Base for both tasks (where we define convergence as the number of epochs to reach 90% success rate). The faster convergence speed, in turn, yields faster model iterations in real-world deployments.

There are a number of potential explanations for the superior performance of LFI. First, note that among all of the strategies, LFI preserves the largest number of human intervention samples. This suggests that human interventions have high intrinsic value to our learning algorithm, as they help to ensure robust policy execution under suboptimal scenarios. Another perspective is that LFI preserves the more frequently intervened trajectories, which exhibit wider state coverage and a diverse array of events. This facilitates the trained policies to operate effectively under rare and unexpected scenarios. MFI (discarding most intervened trajectories) has the opposite effect, favoring trajectories that require less human supervision and often exhibit less diverse behaviors. The results on FIFO and FILO suggest that managing samples according to deployment time is not the most effective strategy, as valuable training data can be collected all throughout the deployment of the system. Finally, the naïve Uniform strategy is ineffective as it does not incorporate any distinguishing characteristics of samples to manage the memory.

Human Workload Reduction. Lastly, we highlight the effectiveness of our method in reducing human workload. In Fig. 9, we plot the human intervention sample ratio for every round, *i.e.*, the percentage of intervention samples in




Figure 11. Learning from Single-Human Demonstration. There are large variations in policy performance when the dataset is the demonstrations from a single human source.




Figure 12. Policy performance from the mixed multi-human demonstration dataset (Round 0), across different dataset sizes.

all samples per round. We compare the results for the HITL methods, Ours and IWR. We see that the human intervention ratio decreases over rounds for both methods, as policy performance increases over time. Furthermore, we see that this reduction in human workload is greater for our method compared to IWR.

Qualitatively, we visualize how the division of work of the human-robot team evolves in Figure 10. For the Gear Insertion task, we do 10 trials of task execution in sequence for our method in Round 0 and Round 3, respectively, and record the time duration for human intervention needed during the deployment. Comparing Round 0 and Round 3, the policy in Round 3 needs very little human intervention, and the intervention duration is also much shorter. This serves as a qualitative illustration of the changing human-robot dynamics within our framework, visualizing the changing nature of human-in-the-loop deployment.

6 Multi-Human Sirius

Real-world deployment scenarios often require multiple human operators to manage a fleet of robots. Different humans vary in their individual skills, familiarity with the system, and level of risk tolerance, which could potentially influence the intervention behavior. To address this variability, we study Sirius’s human-in-the-loop deployment in the *multi-human* setting. To this end, we conduct a comprehensive human study involving 12 participants, who engaged in both human demonstrations and intervention data collection under the Sirius framework. The research goal is to analyze the distribution and characteristics of multi-human data, and to assess its impact on the learning algorithm’s performance in diverse multi-human real-world scenarios.

Participants and Procedures. For our user studies, we selected a diverse group of 12 university students, aiming to encompass a range of experiences and backgrounds. This group included 8 males and 4 females, aged between 19-26 years, with an average age of 23.3. Half of the participants were PhD students, and 7 out of the 12 worked in fields

related to robotics. To ensure a broad representation of skill levels, only 4 of the 12 participants had prior experience in teleoperating robots. We adhered to the IRB protocol approved by the university for all our human studies.

We asked each participant to perform 50 human demonstration and 150 intervention rollouts in the simulated Nut Assembly environment. Each participant was informed about the task goal, practices teleoperation for around 10 minutes before starting the actual experiment, regardless of their prior teleoperation background.

6.1 Learning from Initial Human Demonstrations

To better understand multi-human intervention in Sirius, we start by analyzing human demonstrations without robot interaction. This approach allows us to observe how different levels of operator expertise influence the behavioral learning process. In this section, we seek to answer the following research questions:

- How does the different individual human expertise affect the policy performance for learning from demonstrations?
- How does data diversity (in terms of human operator skill expertise) affect policy performance?

Experiment Results. First, we present the multi-human data quality distribution, measured by the policy performance of learning from the human demonstration dataset (50 trajectories) of each single person. We show the distribution of policy success rate in Figure 11. We see that this group has large variance of demonstration performance, with maximum success rate = 42.9% and minimum success rate = 14.0%. The considerable variation in performance highlights the need for human-in-the-loop algorithms to be robust, effectively accommodating and learning from a diverse array of human behaviors and decision-making patterns.

Table 4. Statistics of the *Demonstration* Dataset: A summary of key measures for each individual’s average demonstration trajectory length.

Statistic	Average Trajectory Length
Mean	201.23
Standard Deviation	52.69
Max	302.48
Min	128.60

We also calculate the average trajectory lengths of each human operator’s demonstration dataset in Table 4. The group demonstrates large variance in the average trajectory lengths, showing high variance in skills and demonstration habits. The disparity in trajectory lengths can be attributed to several factors. Operators with more expertise or familiarity with the task may complete trajectories more efficiently, resulting in shorter average lengths. On the other hand, less experienced operators may follow longer routes, exhibit more indecision in their actions, and make more errors with subsequent corrections, resulting in longer average trajectory lengths.




Figure 13. Sirius outperforms BC-RNN and IWR baselines for multi-human deployment dataset over one round (Round 0 + Round 1).




Figure 14. Learning from Single-Human Intervention. There are large variations in policy performance when the dataset is one round of deployment data from a single human source.

We also show results of performing behavioral cloning on the *mixed* multi-human dataset. We show results of using 50, 200, 400 and 600 demonstrations (sampling around 8%, 33%, 67% and 100% trajectories from each individual human dataset respectively) in Figure 12. First, we demonstrate that the multi-human dataset is effective for the purpose of BC, with the learning performance progressively improving as the volume of data increases. Additionally, for the case with 50 demonstrations, we observe that the mixed dataset yields a moderate policy performance in success rate (29.2%) when compared to the performance of policies trained on each individual’s demonstrations (shown in Figure 11). Learning from mixed quality dataset is worse than learning from highest quality dataset from the most skilled human operator, which could potentially be attributed to two factors. Firstly, although the heterogeneity offers more diverse approaches and strategies for task completion, potentially enhancing generalization, the diversity of the data leads to multimodality, which has been known to hurt BC performance (Shafiullah et al. 2022; Gandhi et al. 2022b). Second, a more diverse dataset also incorporates demonstrations of lower quality. These suboptimal demonstrations potentially add noise and introduce less effective strategies into the learning process.

6.2 Human-in-the-loop Deployment

In this section, we evaluate the multi-human Sirius setting where multiple humans interact with robots, providing interventions. We seek to answer the following research questions:

- Is the Sirius deployment framework and learning algorithm still effective for a multi-human setting?
- How does different individual human expertise affect the policy performance for human-in-the-loop deployment?

- How does data diversity (in terms of human operator skill expertise) affect policy performance?

Experiment Results. We combine the demonstration data from Round 0 and intervention data from Round 1 as the multi-human Sirius dataset. We run three algorithms: **Sirius**, **IWR** and **BC-RNN** on the multi-human dataset to evaluate the **Sirius** algorithm’s effectiveness. Figure 13 shows the policy performance at Round 1 of deployment, where Sirius outperforms the other two baselines. The results suggest that Sirius is well-suited to multi-human settings, where the data is more diverse, and appears less prone to overfitting to a single human operator. The reason why Sirius performs better, as we hypothesised, could be attributed to its intervention-based reweighting scheme. This approach allows us to capture the common principles of human-robot interaction, independent of the variability in specific human behaviors. By leveraging this scheme, Sirius can generalize better across different human operators, making it robust in more diverse and realistic multi-human environments.

Influence of Varied Human Expertise on Learning from Interventions: Does the level of skill and experience of individual human operators play a critical role in the policy learning process when it involves human interventions? This inquiry seeks to understand how the distinct abilities of each operator affect the performance of collaborative human-robot learning systems. To study this problem, we evaluate the policy performance of learning from each *individual* intervention dataset from Round 0 + Round 1 (50 demonstration trajectories + 150 intervention trajectories), each run from a single human operator. Figure 14 shows the variance in policy performance for each single human operator, with maximum success rate = 72.7% and minimum success rate = 44.8%, showing that the difference in individual skills contributes significantly to policy learning. We also show the statistics of the intervention dataset in Table 5. In contrast to the demonstration trajectories, the intervention trajectories exhibit a lower mean and smaller standard deviation. This difference is likely due to the presence of a robot policy that handles the majority of tasks, with human variation primarily occurring in the decision of when and how long to intervene.

Additionally, we utilize the Intervention Sample Ratio in Section 5 to measure the fraction of timesteps completed by humans in comparison to the total number of timesteps. This metric shows the various extent of human involvement in the process. We also define and include two other key metrics: the Average Intervention Length and the Average Number of Interventions per Trajectory. The Average Intervention Length quantifies the average duration of each human intervention, showing how long humans typically engage in the task during an intervention before giving control back to the robot. Meanwhile, the Average Number of Interventions per Trajectory indicates how frequently humans intervene in a single trajectory. As shown in Table 5, the considerable variability in the three metrics highlights that different individuals exhibit unique patterns of intervention, both in terms of frequency and duration.

Impact of Data Diversity on Policy Performance: This study aims to determine if policy learning outcomes significantly vary between a dataset sourced from a larger,




Figure 15. Impact of Data Diversity on Policy Performance. We compare the performance of the most diverse sampled dataset against the three more homogeneous subgroups with 33.3% of the data.

more diverse pool of individuals and a dataset that is less diverse, sourced from fewer individuals and thus more consistent. Participants are divided into 3 sub-groups based on their individual policy performance, creating 3 sub-groups of datasets: Best, Medium and Worse. Additionally, we create another dataset that is more diverse by collecting samples from every participant. We ensure that the combined number of trajectories matches those in the 3 sub-groups of datasets. We compare the policy performance of learning from this more diverse dataset with the 3 homogeneous sub-groups. Figure 15 illustrates that when utilizing the most diverse dataset (Max Diversity), there is a noticeable decline in policy performance compared to the results achieved with the best consistent group (Best). Nonetheless, the enhancement in data diversity does not lead to performance dropping below that of the Worse group. This indicates that the quality of individual datasets plays a more crucial role than diversity alone in determining overall effectiveness.

7 Conclusion

We introduce Sirius, a framework for human-in-the-loop robot manipulation and learning at deployment that both guarantees reliable task execution and also improves autonomous policy performance over time. We utilize the properties and assumptions of human-robot collaboration to develop an intervention-based weighted behavioral cloning method for effectively using deployment data. We also design a practical system that trains and deploys new models continuously under memory constraints. For future work, we would like to improve the flexibility and adaptability of the human-robot shared autonomy, including more intuitive control interfaces and faster policy learning from human feedback. Another direction for future research is to alleviate the human cognitive burden of monitoring and teleoperating the system. To ensure trustworthy execution, our current system still requires the human to constantly monitor the robot. Developing deployment monitoring mechanism would allow the system to automatically detect robot errors without constant human supervision. Lastly, to study human workload reduction, we employed a simple way of measuring human workload based on the intervention percentage. Conducting qualitative human studies to measure human mental workload would provide deeper insights.

Table 5. Statistics of the *Intervention* Dataset: A summary of key measures for each individual's average intervention rollout trajectory length, intervention ratio, average intervention length and average number of interventions per trajectory.

Statistic	Average Trajectory Length	Intervention Sample Ratio	Average Intervention Length	Average Number of Interventions per Trajectory
Mean	185.84	0.32	15.28	2.80
Standard Deviation	19.34	0.11	4.70	1.03
Max	216.19	0.57	23.67	5.15
Min	161.88	0.19	7.09	1.79

Acknowledgements

We thank Ajay Mandlekar for having multiple insightful discussions, and for sharing well-designed simulation task environments and codebases during development of the project. We thank Yifeng Zhu for valuable advice and system infrastructure development for real robot experiments. We would like to thank Tian Gao, Jake Grigsby, Zhenyu Jiang, Ajay Mandlekar, Braham Snyder, and Yifeng Zhu for providing helpful feedback for this manuscript.

Declaration of Conflicting Interests

The Authors declare that there is no conflict of interest.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: We acknowledge the support of the National Science Foundation [grant numbers 1955523, 2145283]; the Office of Naval Research [grant number N00014-22-1-2204]; and Amazon.

Notes

1. <https://3dconnexion.com/us/spacemouse/>
2. <https://github.com/rail-berkeley/rkit/tree/master/examples/iql>

References

- Ajay A, Kumar A, Agrawal P, Levine S and Nachum O (2021) Opal: Offline primitive discovery for accelerating offline reinforcement learning. In: *ICLR*.
- Andrychowicz M, Baker B, Chociej M, Józefowicz R, McGrew B, Pachocki JW, Petron A, Plappert M, Powell G, Ray A, Schneider J, Sidor S, Tobin J, Welinder P, Weng L and Zaremba W (2018) Learning dexterous in-hand manipulation. pp. 20–3.
- Biyik E, Losey DP, Palan M, Landolfi NC, Shevchuk G and Sadigh D (2022) Learning reward functions from diverse sources of human feedback: Optimally integrating demonstrations and preferences. In: *IJRR*, volume 41. SAGE Publications Sage UK: London, England, pp. 45–67.
- Brown DS, Goo W, Nagarajan P and Niekum S (2019) Extrapolating beyond suboptimal demonstrations via inverse reinforcement learning from observations. In: *ICML*.
- Cabi S, Colmenarejo SG, Novikov A, Konyushkova K, Reed S, Jeong R, Zolna K, Aytar Y, Budden D, Vecerik M, Sushkov O, Barker D, Scholz J, Denil M, de Freitas N and Wang Z (2019) Scaling data-driven robotics with reward sketching and batch reinforcement learning. In: *arXiv preprint arXiv:1909.12200*.
- Celemin C, Pérez-Dattari R, Chisari E, Franzese G, de Souza Rosa L, Prakash R, Ajanović Z, Ferraz M, Valada A and Kober J (2022) Interactive imitation learning in robotics: A survey. pp. 1–197.
- Chernova S and Veloso M (2007) Confidence-based policy learning from demonstration using gaussian mixture models. In: *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS '07*. New York, NY, USA: Association for Computing Machinery. ISBN 9788190426275. DOI:10.1145/1329125.1329407. URL <https://doi.org/10.1145/1329125.1329407>.
- Chisari E, Welschehold T, Boedecker J, Burgard W and Valada A (2021) Correct me if i am wrong: Interactive learning for robotic manipulation. In: *RAL*, volume 7. pp. 3695–3702.
- Christiano P, Leike J, Brown TB, Martic M, Legg S and Amodei D (2017) Deep reinforcement learning from human preferences.
- Cruz CA and Igarashi T (2020) A survey on interactive reinforcement learning: Design principles and open challenges. In: *DIS*.
- Cui Y, Koppol P, Admoni H, Simmons RG, Steinfeld A and Fitzgerald T (2021) Understanding the relationship between interactions and outcomes in human-in-the-loop machine learning. In: *IJCAI*.
- Daniel C, Viering M, Metz J, Kroemer O and Peters J (2014) Active reward learning. In: *RSS*.
- Dass S, Pertsch K, Zhang H, Lee Y, Lim JJ and Nikolaidis S (2022) Pato: Policy assisted teleoperation for scalable robot data collection.
- Dragan A and Srinivasa S (2012) Formalizing assistive teleoperation. In: *RSS*.
- Dragan A and Srinivasa S (2013) A policy-blending formalism for shared control. SAGE Publications Sage UK: London, England, pp. 790–805.
- Florence P, Lynch C, Zeng A, Ramirez O, Wahid A, Downs L, Wong A, Lee J, Mordatch I and Tompson J (2021) Implicit behavioral cloning. In: *CoRL*.
- Fu J, Kumar A, Nachum O, Tucker G and Levine S (2020) D4rl: Datasets for deep data-driven reinforcement learning. In: *arXiv preprint arXiv:1802.01744*.
- Fujimoto S, Meger D and Precup D (2019) Off-policy deep reinforcement learning without exploration. In: *ICML*.
- Gandhi K, Karamcheti S, Liao M and Sadigh D (2022a) Eliciting compatible demonstrations for multi-human imitation learning. In: *CoRL*.

- Gandhi K, Karamcheti S, Liao M and Sadigh D (2022b) Eliciting compatible demonstrations for multi-human imitation learning.
- Gopinath D, Jain S and Argall BD (2017) Human-in-the-loop optimization of shared autonomy in assistive robotics. In: *RAL*.
- Griffith S, Subramanian K, Scholz J, Isbell CL and Thomaz AL (2013) Policy shaping: Integrating human feedback with reinforcement learning. In: *NeurIPS*.
- Gulcehre C, Wang Z, Novikov A, Paine TL, Colmenarejo SG, Zolna K, Agarwal R, Merel J, Mankowitz DJ, Paduraru C, Dulac-Arnold G, Li JZ, Norouzi M, Hoffman MW, Nachum O, Tucker G, Heess NMO and de Freitas N (2020) RL unplugged: Benchmarks for offline reinforcement learning. In: *NeurIPS*.
- Hart SG and Staveland LE (1988) Development of nasa-tlx (task load index): Results of empirical and theoretical research. In: Hancock PA and Meshkati N (eds.) *Human Mental Workload, Advances in Psychology*, volume 52. North-Holland, pp. 139–183. DOI:[https://doi.org/10.1016/S0166-4115\(08\)62386-9](https://doi.org/10.1016/S0166-4115(08)62386-9). URL <https://www.sciencedirect.com/science/article/pii/S0166411508623869>.
- He K, Zhang X, Ren S and Sun J (2016) Deep residual learning for image recognition. In: *CVPR*. pp. 770–778.
- Hoque R, Balakrishna A, Novoseller ER, Wilcox A, Brown DS and Goldberg K (2021) Thriftydagger: Budget-aware novelty and risk gating for interactive imitation learning. In: *CoRL*.
- Hoque R, Chen LY, Sharma S, Dharmarajan K, Thananjeyan B, Abbeel P and Goldberg K (2022) Fleet-dagger: Interactive robot fleet learning with scalable human supervision.
- Javdani S, Srinivasa SS and Bagnell JA (2015) Shared autonomy via hindsight optimization. In: *RSS*, volume 2015.
- Kalashnikov D, Irpan A, Pastor P, Ibarz J, Herzog A, Jang E, Quillen D, Holly E, Kalakrishnan M, Vanhoucke V and Levine S (2018) QT-Opt: Scalable deep reinforcement learning for vision-based robotic manipulation. In: *CoRL*.
- Kelly M, Sidrane C, Driggs-Campbell K and Kochenderfer MJ (2019) HG-DAGger: Interactive imitation learning with human experts. In: *ICRA*. pp. 8077–8083.
- Kidambi R, Rajeswaran A, Netrapalli P and Joachims T (2020) Morel: Model-based offline reinforcement learning. In: *NeurIPS*.
- Knox WB and Stone P (2009) Interactively shaping agents via human reinforcement: The tamer framework. In: *K-CAP*.
- Kostrikov I, Nair A and Levine S (2021) Offline reinforcement learning with implicit q-learning. In: *ICLR*.
- Kumar A, Hong J, Singh A and Levine S (2022) When should we prefer offline reinforcement learning over behavioral cloning? In: *ICLR*.
- Kumar A, Zhou A, Tucker G and Levine S (2020) Conservative q-learning for offline reinforcement learning. In: *NeurIPS*.
- Lee J, Hwangbo J, Wellhausen L, Koltun V and Hutter M (2020) Learning quadrupedal locomotion over challenging terrain. *Science Robotics*.
- Lee K, Smith L and Abbeel P (2021) Pebble: Feedback-efficient interactive reinforcement learning via relabeling experience and unsupervised pre-training. In: *ICML*.
- Leike J, Krueger D, Everitt T, Martic M, Maini V and Legg S (2018) Scalable agent alignment via reward modeling: A research direction. In: *arXiv preprint arXiv:1811.07871*.
- Levine S, Kumar A, Tucker G and Fu J (2020) Offline reinforcement learning: Tutorial, review, and perspectives on open problems.
- Li Q, Peng Z and Zhou B (2022) Efficient learning of safe driving policy via human-ai copilot optimization.
- MacGlashan J, Ho MK, Loftin R, Peng B, Wang G, Roberts DL, Taylor ME and Littman ML (2017) Interactive learning from policy-dependent human feedback. In: *ICML*. pp. 2285–2294.
- Mandlekar A, Ramos F, Boots B, Savarese S, Fei-Fei L, Garg A and Fox D (2020a) Iris: Implicit reinforcement without interaction at scale for learning control from offline robot manipulation data. In: *ICRA*.
- Mandlekar A, Xu D, Martín-Martín R, Savarese S and Fei-Fei L (2020b) Learning to generalize across long-horizon tasks from human demonstrations. In: *RSS*.
- Mandlekar A, Xu D, Martín-Martín R, Zhu Y, Fei-Fei L and Savarese S (2020c) Human-in-the-loop imitation learning using remote teleoperation. In: *arXiv preprint arXiv:2012.06733*.
- Mandlekar A, Xu D, Wong J, Nasiriany S, Wang C, Kulkarni R, Fei-Fei L, Savarese S, Zhu Y and Martín-Martín R (2021a) What matters in learning from offline human demonstrations for robot manipulation. In: *CoRL*.
- Mandlekar A, Xu D, Wong J, Nasiriany S, Wang C, Kulkarni R, Fei-Fei L, Savarese S, Zhu Y and Martín-Martín R (2021b) What matters in learning from offline human demonstrations for robot manipulation.
- Mandlekar A, Zhu Y, Garg A, Booher J, Spero M, Tung A, Gao J, Emmons J, Gupta A, Orbay E, Savarese S and Fei-Fei L (2018) Roboturk: A crowdsourcing platform for robotic skill learning through imitation. In: *CoRL*. pp. 879–893.
- Muelling K, Venkatraman A, Valois JS, Downey JE, Weiss JM, Javdani S, Hebert M, Schwartz AB, Collinger JL and Bagnell JA (2015) Autonomy infused teleoperation with application to bci manipulation.
- Nair A, Gupta A, Dalal M and Levine S (2021) Awac: Accelerating online reinforcement learning with offline datasets. In: *arXiv preprint arXiv:2006.09359*.
- Perez-D’Arpino C and Shah JA (2015) Fast target prediction of human reaching motion for cooperative human-robot manipulation tasks using time series classification. In: *ICRA*. pp. 6175–6182.
- Pomerleau DA (1989) Alvin: An autonomous land vehicle in a neural network. In: *NeurIPS*.
- Reddy S, Levine S and Dragan AD (2018) Shared autonomy via deep reinforcement learning.
- Ross S, Gordon G and Bagnell D (2011) A reduction of imitation learning and structured prediction to no-regret online learning. In: *AISTATS*. pp. 627–635.
- Sasaki F and Yamashina R (2021) Behavioral cloning from noisy demonstrations. In: *ICLR*. URL <https://openreview.net/forum?id=zrT3HcsWSAt>.
- Shafiqullah NMM, Cui ZJ, Altanzaya A and Pinto L (2022) Behavior transformers: Cloning k modes with one stone.
- Shahin M, Ali Babar M and Zhu L (2017) Continuous integration, delivery and deployment: A systematic review on approaches, tools, challenges and practices. pp. 3909–3943.
- Singh A, Yu A, Yang J, Zhang J, Kumar A and Levine S (2020) Cog: Connecting new skills to past experience with offline reinforcement learning. In: *CoRL*.

- Spencer J, Choudhury S, Barnes M, Schmittle M, Chiang M, Ramadge P and Srinivasa S (2020) Learning from interventions: Human-robot interaction as both explicit and implicit feedback. In: *RSS*.
- Stiber M, Taylor R and Huang CM (2022) Modeling human response to robot errors for timely error detection. In: *IROS*. pp. 676–683.
- Tan W, Koleczek D, Pradhan S, Perello N, Chettiar V, Ma N, Rajaram A, Rohra V, Srinivasan S, Hossain S and Chandak Y (2021) Intervention aware shared autonomy.
- Wang X, Lee K, Hakhamaneshi K, Abbeel P and Laskin M (2022) Skill preferences: Learning to extract and execute robotic skills from human feedback. In: *CoRL*. PMLR, pp. 1259–1268.
- Wang Z, Novikov A, Zolna K, Springenberg JT, Reed S, Shahriari B, Siegel N, Merel J, Gulcehre C, Heess N and de Freitas N (2020) Critic regularized regression. pp. 7768–7778.
- Warnell G, Waytowich N, Lawhern V and Stone P (2018) Deep tamer: Interactive agent shaping in high-dimensional state spaces. In: *AAAI*, volume 32.
- Xu H, Zhan X, Yin H and Qin H (2022) Discriminator-weighted offline imitation learning from suboptimal demonstrations. In: *ICML*. PMLR, pp. 24725–24742.
- Yu T, Kumar A, Rafailov R, Rajeswaran A, Levine S and Finn C (2021) Combo: Conservative offline model-based policy optimization. In: *NeurIPS*.
- Yu T, Thomas G, Yu L, Ermon S, Zou J, Levine S, Finn C and Ma T (2020) Mopo: Model-based offline policy optimization. In: *NeurIPS*.
- Zhang R, Saran A, Liu B, Zhu Y, Guo S, Niekum S, Ballard D and Hayhoe M (2020) Human gaze assisted artificial intelligence: A review. In: *IJCAI*. pp. 4951–4958. Survey track.
- Zhang R, Torabi F, Guan L, Ballard DH and Stone P (2019) Leveraging human guidance for deep reinforcement learning tasks. In: *IJCAI*.
- Zhang T, McCarthy Z, Jow O, Lee D, Chen X, Goldberg K and Abbeel P (2018) Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. In: *ICRA*.
- Zhu Y, Wong J, Mandlekar A and Martín-Martín R (2020) robosuite: A modular simulation framework and benchmark for robot learning. In: *arXiv preprint arXiv:2009.12293*.
- Zolna K, Novikov A, Konyushkova K, Gülçehre Ç, Wang Z, Aytar Y, Denil M, de Freitas N and Reed SE (2020) Offline learning from demonstrations and unlabeled experience. In: *CoRR*.