

Demonstration-Guided Reinforcement Learning with Efficient Exploration for Task Automation of Surgical Robot

Tao Huang¹, Kai Chen¹, Bin Li², Yun-Hui Liu², Qi Dou¹

Abstract—Task automation of surgical robot has the potentials to improve surgical efficiency. Recent reinforcement learning (RL) based approaches provide scalable solutions to surgical automation, but typically require extensive data collection to solve a task if no prior knowledge is given. This issue is known as the exploration challenge, which can be alleviated by providing expert demonstrations to an RL agent. Yet, how to make effective use of demonstration data to improve exploration efficiency still remains an open challenge. In this work, we introduce Demonstration-guided EXploration (DEX), an efficient reinforcement learning algorithm that aims to overcome the exploration problem with expert demonstrations for surgical automation. To effectively exploit demonstrations, our method estimates expert-like behaviors with higher values to facilitate productive interactions, and adopts non-parametric regression to enable such guidance at states unobserved in demonstration data. Extensive experiments on 10 surgical manipulation tasks from SurRoL, a comprehensive surgical simulation platform, demonstrate significant improvements in the exploration efficiency and task success rates of our method. Moreover, we also deploy the learned policies to the da Vinci Research Kit (dVRK) platform to show the effectiveness on the real robot. Code is available at <https://github.com/med-air/DEX>.

I. INTRODUCTION

Surgical robots nowadays assist surgeons to conduct minimally invasive interventions in clinical routine [1]. Automating surgical tasks has been increasingly desired to improve surgical efficiency, with many efforts being made on different tasks such as suturing [2–4], endoscope control [5–7], tissue manipulation [8–10] and pattern cutting [11–13]. Recently, reinforcement learning (RL) approaches have exhibited high scalability to learn diverse control policies and yielded promising performance in surgical automation [14–20], but typically require extensive data collection to solve a task if no prior knowledge is given. This issue, known as the exploration challenge, gives rise to the idea of providing expert knowledge from demonstration data to an RL agent [21]. Yet, how to make full effective of demonstrations to improve exploration efficiency still remains an open challenge.

One classical approach is to give higher sampling priority to demonstrated data over self-collected data in a simulator [22–24]. While straightforward, the expert

This research work was supported in part by CUHK T Stone Robotics Institute, Hong Kong Innovation and Technology Commission Project No. ITS/237/21FP, Hong Kong Research Grants Council TRS Project No.T42-409/18-R, and InnoHK Multi-Scale Medical Robotics Center.

¹T. Huang, K. Chen, and Q. Dou are with the Department of Computer Science and Engineering, The Chinese University of Hong Kong.

²B. Li, and Y. H. Liu are with the Department of Mechanical and Automation Engineering, The Chinese University of Hong Kong.

Corresponding author: Qi Dou (qidou@cuhk.edu.hk).

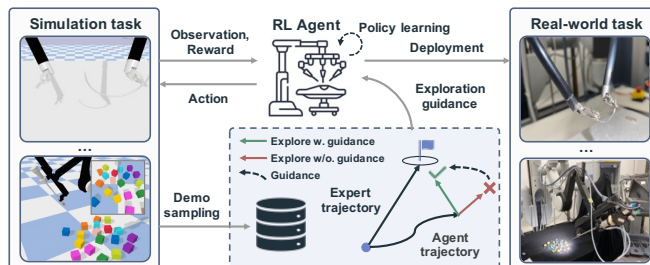


Fig. 1: Our method DEX leverages demonstration data to efficiently guide the RL exploration in simulation tasks. The learned policies are successfully deployed to the real-world surgery subtasks on the dVRK platform.

knowledge from demonstrations is not explicitly exploited during the exploration, making these methods still inefficient in robotic tasks. To provide explicit exploration guidance, one other approach offers RL agent with an additional reward function learned from demonstrations, which encourages the agent to stay close to expert samples [25–27] by discerning expert-like and -unlike behaviors. However, these methods are not easily applicable to surgical tasks because they not only require careful task-specific design of environment reward function [28], but also introduces the risk of local optima during policy learning.

To overcome such limitations, a promising approach is to directly regularize the robot policy to mimic the expert policy within the actor-critic framework [29–33]. Such guidance is typically operationalized by regularizing the actor loss with a penalty measuring the behavioral dissimilarity between the robot and the expert. Effective as these methods are, their exploration efficiency is still unsatisfactory due to two limitations. Firstly, these methods enforce little regularization on the critic, making the critic suffers from the notorious overestimation issue [34, 35]. Without an accurate critic, the overestimated values of expert-unlike actions hinder the robot from exploring expert-like actions. Secondly, these methods provide effective guidance only when the current state is close to the ones in the demonstrations. This impedes their capability to guide exploration early in the learning process, as the robot with a mediocre policy is likely to visit states far from demonstrated ones. A stopgap is to cover as many possible demonstrated states as possible by collecting a large number of demonstrations [36]. Unfortunately, collecting many surgical expert data is often unaffordable due to the resources and ethical constraints.

In this work, we aim to improve the exploration efficiency of RL given a modest set of demonstration data for task automation of surgical robot. To tackle the overestimation

issue of the critic, we propose a novel algorithmic framework, named expert-guided actor-critic, that regularizes both actor and critic with the action dissimilarity between agent and expert. As illustrated in Figure 1, at a high level, with this regularization the critic lowers the value estimates of expert-unlike actions (colored with red), thus further encouraging the exploration on expert-like actions (colored with green). To guide exploration at states unobserved in demonstrations, we adopt a non-parametric regression model to robustly propagate the guidance (dashed arrow) from demonstration data to these states. Consequently, our method, named Demonstration-guided EXploration (DEX), is of high exploration efficiency with a limited amount of demonstration data. The empirical results on surgical automation tasks from SurRoL [37] show that DEX significantly outperforms prior learning-based methods. We also validate the transferability of learned policies on da Vinci surgical robots. The contributions of this work are summarized as follows:

- We propose a novel actor-critic framework to mitigate the overestimation issue of the critic for encouraging explorations on expert-like actions in RL;
- We adopt non-parametric guidance propagation to guide exploration at demonstration-unobserved states;
- We empirically demonstrate that our method significantly outperforms prior RL-based approaches on the surgical robot learning tasks from SurRoL;
- We successfully deploy our trained policies on the real dVRK, indicating its great potential for task automation of surgical robot in the real world.

II. RELATED WORK

Surgical Robot Task Automation. Considerable efforts have been made by researchers to design skill-oriented controllers for surgical task automation, including surgical suturing [2, 38], endoscope control [5, 6], tissue manipulation [8, 39] and pattern cutting [40, 41]. However, these works follow a traditional paradigm that requires domain knowledge for designing control strategies and suffers from low generality, i.e., a controller designed for one task is often hard to generalize [42]. In contrast, learning-based approaches, representatively RL, enable robots to flexibly learn to perform tasks from collected data. These methods thus do not require task-specific control strategies and have shown improved generalization capabilities in automating complex surgical tasks [14–20]. Promising as it is, running existing RL methods in surgical tasks is still inefficient due to the exploration challenge, which is often alleviated through substantial reward engineering for each task. To this end, we propose a novel RL algorithm to solve diverse tasks efficiently by exploiting expert demonstrations for facilitating task automation in robotic surgery.

Learning from Demonstrations. Imitation learning (IL), also known as learning from demonstrations [43], is a typical learning-based framework that leverages demonstrations for policy learning. A well-known approach is behavior

cloning (BC, [44]) which supervises the agent to imitate expert behavior. However, it suffers from the distribution shift problem and thus is of poor generalization performance [45, 46]. Some IL methods overcome this issue by using generative adversarial networks to match the trajectory distribution of demonstrations [47, 48] or provide the agent with inferred reward function [49, 50]. But they require a great many online samples and demonstrations to ensure fair generalization capability and suffer from poor training stability [51]. In contrast, our method goes beyond such limitations by utilizing the expert knowledge from demonstration data to expedite RL exploration, which combines advances of IL and RL and alleviates their respective drawbacks.

Demonstration-Guided RL. A number of demonstration-guided RL approaches have been proposed to mitigate the exploration challenges of RL, including prioritizing expert samples over self-collected samples [22–24] or learning extra reward functions [25–27]. However, these methods either make no explicit use of demonstrations or still need careful reward engineering in surgical tasks. To overcome these issues, a promising approach is to leverage demonstrations to directly regularize the RL policy to explore expert-like behaviors [29–33]. It penalizes the behavioral dissimilarity between agent and expert with state-action pairs from demonstrations. Nonetheless, this approach often requires a large number of demonstrations to ensure high exploration efficiency due to the impediment of overestimated values and the incapability of offering guidance at states far from the demonstrated ones. Our method addresses these limitations through critic regularization in the proposed actor-critic framework and robust guidance propagation based on non-parametric regression.

III. METHOD

We develop Demonstration-guided EXploration (DEX), a novel exploration-efficient demonstration-guided RL algorithm for surgical subtask automation with limited demonstrations. Our method addresses the potential overestimation issue in existing methods based on our proposed actor-critic framework in Section III-A. To offer exploration guidance at states unobserved in demonstration data, our method propagates the guidance from demonstrations through a non-parametric module in Section III-B. Figure 2 illustrates the overall framework of the proposed method DEX.

Problem Formulation: We consider an off-policy RL agent that interacts within an environment formulated by a Markov decision process. At time step t , the agent takes an action a_t based on the current state s_t and its deterministic policy π . The environment then rewards the agent with $r_t := r(s_t, a_t)$ and transits it to the successor state s_{t+1} . After each transition, the agent stores the experience (s_t, a_t, r_t, s_{t+1}) into a replay buffer \mathcal{D}_A . Meanwhile, it also maintains a separate demonstration buffer \mathcal{D}_E , where experiences are collected through an unavailable expert policy π^e in advance.

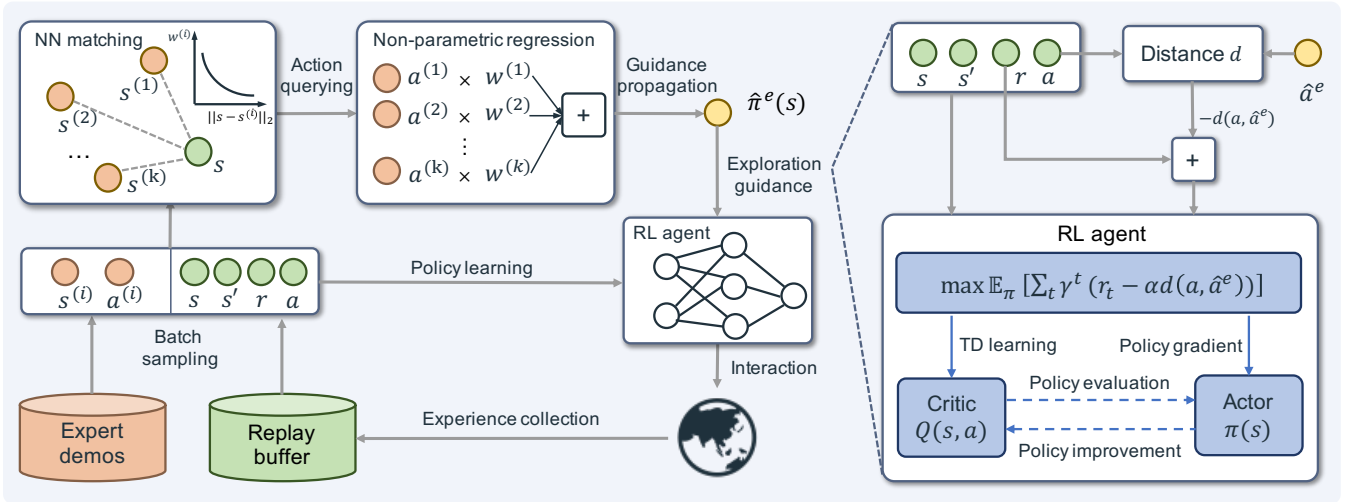


Fig. 2: RL with demonstrations-guided exploration. We present the overall illustration of our proposed method DEX. It consists of two parts, a novel actor-critic based policy learning module that efficiently leverages demonstration data to guide the RL exploration (right), and a non-parametric module based on nearest-neighbor matching and locally weighted regression for robust guidance propagation at states far from demonstrated ones (upper left).

A. Expert-Guided Actor-Critic Framework

Many existing actor-critic based methods learn optimal policy by training an actor that maximizes the expected return $\mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t r_t]$, which is estimated by a critic that approximates the Q-value function, where $\gamma \in (0, 1]$ is a discount factor. Their common way to guide exploration is to jointly maximize the Q-value and the behavioral similarity between the agent and expert. Effective as it is, the inaccurate Q-value estimates of expert-unlike actions may dominate the guidance embodied as behavioral dissimilarity, and thus hamper the exploration, reminiscent of the regularization issue in offline RL literature [52, 53]. In other words, the agent is likely to take unproductive actions due to value misestimate, which hinders it from exploring expert-like actions.

To remedy this issue, we introduce a novel RL objective that augments the environment reward with the behavior gap between agent policy and expert policy:

$$\max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t (r_t - \alpha d(a_t, a_t^e)) \right], \quad a_t^e := \pi^e(s_t), \quad (1)$$

where α is an exploration coefficient and the function $d(\cdot, \cdot)$ is a distance metric that measures the behavioral similarity between agent action and expert action. Correspondingly, we define the regularized Q-value function as $Q^\pi(s, a) := \mathbb{E}_{\pi} [\sum_{t=0}^{\infty} \gamma^t (r_t - \gamma \alpha d(a_{t+1}, a_{t+1}^e)) | s, a]$. It assigns higher value estimates to actions similar to the expert action at any given state. It thus mitigates the effect of the inaccurate estimates of undesired actions by lowering its value through a dissimilarity regularization. Meanwhile, echoed with SAC [54], an appropriate choice of α steers the exploration-exploitation balance and benefits policy learning.

Based on the formulation, we introduce our proposed expert-guided actor-critic frameworks. The regularized value function (critic) Q_θ , parameterized by θ , minimizes the

squared Bellman residual:

$$\mathcal{L}_Q(\theta) = \mathbb{E}_{\mathcal{D}_A \cup \mathcal{D}_E} \left[(r_t + \gamma V_{\bar{\theta}}(s_{t+1}) - Q_\theta(s_t, a_t))^2 \right], \quad (2)$$

where $V_{\bar{\theta}}$ is state value function served as Q-target whose parameters $\bar{\theta}$ are the exponential moving average of θ that stabilize the training process. The actor, parameterized by ϕ , is trained by maximizing following the objective:

$$\mathcal{L}_\pi(\phi) = \mathbb{E}_{\mathcal{D}_A \cup \mathcal{D}_E} [Q_\theta(s_t, \pi_\phi(s_t)) - \alpha d(\pi_\phi(s_t), a_t^e)]. \quad (3)$$

Our framework bears some resemblances to behavior regularized actor-critic in offline RL domain [55, 56]. At the online fine-tuning stage, their formulas target conservative policy updates by constraining the behavior gap between old and new policies. Different from their focus, our framework treats expert policy as a static reference policy and uses the behavior gap between agent and expert as a regularization term to accelerate the RL exploration more efficiently.

B. Guidance Propagation from Limited Demonstrations

Prior methods provide the actor with exploration guidance only at states observed in demonstration data, since the expert action a_t^e is available only therein. However, the agent is prone to visit regions uncovered by demonstrations at the initial stage of training, where the demonstrations are incapable of supervising the actor exploration. On the other hand, such an issue also limits the feasibility and efficacy of critic regularization. Only regularizing the demonstration-covered Q-value introduces an underestimation bias on the value estimate to these state-action pairs due to the negativity of the regularization term. This negates the original purpose of enforcing critic regularization. Propagating guidance from limited demonstrations is thus of high necessity to efficiently realize the proposed actor-critic framework.

One natural solution is to learn a parametric expert policy approximator $\hat{\pi}^e$ from demonstrations in a parametric manner, e.g., behavior cloning. Although using a parametric

model already outperforms existing methods, which demonstrates the effectiveness of our framework, we empirically observe that the parametric propagated guidance, i.e., $\hat{\pi}(s)$, fairly differs from the ground-truth expert action when states are far from the demonstration ones. We thus resort to a non-parametric regression model that empirically propagates more robust guidance from limited demonstrations.

Specifically, we first sample a minibatch of states and actions from demonstration buffer \mathcal{D}_E . Then, given any state s , we search states in the minibatch and find k nearest neighbors of s based on the Euclidean distance. The selected k states and their associated actions are denoted by $\{(s^{(i)}, a^{(i)})\}_{i=1}^k$, a set of transitions from \mathcal{D}_E . Subsequently, we approximate the expert policy through locally weighted regression [57] with exponential kernel function:

$$\hat{\pi}^e(s) = \frac{\sum_{i=1}^k \exp(-\|s - s^{(i)}\|_2) \cdot a^{(i)}}{\sum_{i=1}^k \exp(-\|s - s^{(i)}\|_2)}. \quad (4)$$

The assumption here is that similar states share similar optimal actions, whose effectiveness and robustness have been demonstrated in some robotic learning tasks [58–60]. Unlike these IL methods, we aim to use the limited demonstration data to accelerate the exploration of RL and increase the generalization ability of policy.

C. Implementation Details

We implement our method with deep deterministic policy gradient (DDPG, [61]) from OpenAI Baselines [62] and adopt their main hyperparameter settings. We choose DDPG to avoid the potential pathology of density-based methods [63], though our implementation can be naturally extended to the stochastic setting. Specifically, both actor and critic are parameterized as four fully-connected layers of hidden dimension 256 interpolated with ReLU activations, where actions are scaled to the range $[-1, 1]$ by a Tanh activation in the actor network. They are trained with ADAM [64]. The exploration noise is set as Gaussian noise with a scale factor of 0.1. For goal-conditioned environments, we adopt hindsight experience replay [65] with future sampling strategy on 80% examples for both replay buffer and demonstration buffer. The distance metric is set as L2-norm for all tasks. The coefficient of exploration term α is set as 5 for all tasks. The number of nearest neighbors k used for expert policy estimation is set to 5. More analyses on α and k would be given in the ablation studies.

IV. EXPERIMENTS

In this section, we first conduct experiments on SurRoL [37], a comprehensive simulation environment for task automation of surgical robots. Extensive evaluation and ablation analysis on 10 challenging surgical manipulation tasks demonstrate the effectiveness of our method DEX for surgical robot learning. Moreover, we transfer the learned policies to a dVRK platform for real-world robotic deployment.




Fig. 3: Task description of SurRoL. We choose ten surgical tasks from SurRoL and divide them into three domains according to the type of manipulators. The arrows in each task represent the task flow as well as the execution process of the scripted expert policy provided by the platform.

A. Experimental Setup

Different Tasks on Surgical Manipulation. To evaluate whether our method DEX can effectively handle different surgical manipulation scenarios, we conduct experiments on 10 challenging tasks from SurRoL. These 10 tasks widely cover three different domains according to the type of manipulator: (1) Single-handed patient-sided manipulator (**PSM**); (2) Bimanual PSM (**Bi-PSM**); and (3) Endoscopic camera manipulator (**ECM**). Figure 3 illustrates the involved 10 manipulation tasks. All tasks are goal-conditioned with sparse reward functions that indicate task success, except for *ActiveTrack*, whose reward function is dense that depends on the precision of object tracking. We use low-dimensional state representation that consists of object state (3D Cartesian positions and 6D pose) and robot proprioceptive state (jaw status and end-effector position), and Cartesian-space control as action space. Please refer to [37] for more details on the different surgical manipulation tasks. Besides, we sample 100 successful episodes of demonstration data for each task through the scripted expert policy provided by SurRoL.

Evaluation Metrics. Similar to [68, 69], we measure the manipulation performance of each task after a uniform 100k environment steps. The manipulation score for each task is linearly re-scaled to be within $[0, 1]$ for comparisons across tasks as well as better visualization. Following [70], we adopt interquartile mean (IQM) and 95% interval estimates via stratified bootstrap confidence intervals (CIs) to measure the aggregate performance over each domain robustly.

Baselines. First of all, we compare our method with two state-of-the-art pure RL algorithms without using demonstrations: (1) **SAC** [54] and (2) **DDPG** [61]. Besides, we compare our method with three representative IL methods: (3) **BC** [44] that supervises the agent to imitate expert actions from demonstrations. (4) **SQL** [67] that reduces IL to RL by assigning positive rewards to demonstrations, a representative of the state-of-the-art IL algorithms based on SAC. (5) **VINN** [58] that directly uses the estimate of expert policy in Equation (4) as the agent policy. Moreover, we compare our method with four existing demonstration-guided RL algorithms: (6) **DDPGBC** [29]

TABLE I: Performance after training 100k environment steps. We present mean scores with standard deviations of each task and aggregate IQMs with 95% stratified bootstrap CIs of each domain (grey-colored cells). Results are over 10 runs with different seeds, where each run averages 20 evaluations. We abbreviate sparse and dense reward to S and D , respectively. Our method achieves prominent manipulation performance on complex tasks and domains.

Task	$S/A/r$	SAC [66]	DDPG [61]	BC [44]	SQL [67]	VINN [58]	DDPGBC [29]	AMP [27]	CoL [33]	AWAC [55]	DEX	
ECM	Aggregate	–	0.99 (± 0.03)	0.99 (± 0.02)	1.00 (± 0.00)	0.24 (± 0.06)	0.58 (± 0.06)	1.00 (± 0.00)	1.00 (± 0.01)	1.00 (± 0.00)	0.99 (± 0.01)	1.00 (± 0.00)
	ECMReach	$\mathbb{R}^{12}/\mathbb{R}^3/S$	1.00 (± 0.06)	1.00 (± 0.00)	1.00 (± 0.00)	0.07 (± 0.04)	0.49 (± 0.10)	1.00 (± 0.00)	0.99 (± 0.02)	1.00 (± 0.00)	1.00 (± 0.00)	1.00 (± 0.00)
	StaticTrack	$\mathbb{R}^{16}/\mathbb{R}^3/S$	0.92 (± 0.14)	0.98 (± 0.05)	1.00 (± 0.00)	0.43 (± 0.26)	0.56 (± 0.10)	1.00 (± 0.00)	0.97 (± 0.03)	1.00 (± 0.00)	1.00 (± 0.00)	1.00 (± 0.00)
	MisOrient	$\mathbb{R}^{11}/\mathbb{R}^1/S$	1.00 (± 0.00)	1.00 (± 0.00)	1.00 (± 0.00)	0.56 (± 0.10)	0.50 (± 0.11)	0.99 (± 0.02)	0.98 (± 0.02)	0.99 (± 0.02)	0.98 (± 0.03)	0.99 (± 0.02)
	ActiveTrack	$\mathbb{R}^{10}/\mathbb{R}^3/D$	0.79 (± 0.08)	0.67 (± 0.08)	0.95 (± 0.01)	0.07 (± 0.06)	0.92 (± 0.06)	0.81 (± 0.05)	0.94 (± 0.01)	0.96 (± 0.01)	0.51 (± 0.12)	0.94 (± 0.01)
PSM	Aggregate	–	0.0 (± 0.00)	0.00 (± 0.00)	0.40 (± 0.05)	0.00 (± 0.00)	0.02 (± 0.02)	0.80 (± 0.04)	0.00 (± 0.00)	0.85 (± 0.06)	0.46 (± 0.19)	0.89 (± 0.03)
	NeedleReach	$\mathbb{R}^{13}/\mathbb{R}^5/S$	1.00 (± 0.00)	1.00 (± 0.00)	1.00 (± 0.00)	0.07 (± 0.09)	0.89 (± 0.06)	1.00 (± 0.00)	0.99 (± 0.02)	1.00 (± 0.00)	0.94 (± 0.20)	1.00 (± 0.00)
	GauzeRetrieve	$\mathbb{R}^{25}/\mathbb{R}^5/S$	0.00 (± 0.00)	0.00 (± 0.00)	0.07 (± 0.05)	0.00 (± 0.00)	0.01 (± 0.02)	0.63 (± 0.11)	0.00 (± 0.00)	0.71 (± 0.16)	0.43 (± 0.43)	0.73 (± 0.12)
	NeedlePick	$\mathbb{R}^{25}/\mathbb{R}^5/S$	0.00 (± 0.00)	0.00 (± 0.00)	0.21 (± 0.06)	0.00 (± 0.00)	0.02 (± 0.02)	0.91 (± 0.05)	0.00 (± 0.00)	0.96 (± 0.05)	0.26 (± 0.33)	0.94 (± 0.05)
	PegTransfer	$\mathbb{R}^{25}/\mathbb{R}^5/S$	0.00 (± 0.00)	0.00 (± 0.00)	0.56 (± 0.11)	0.02 (± 0.05)	0.05 (± 0.04)	0.48 (± 0.22)	0.00 (± 0.00)	0.58 (± 0.23)	0.31 (± 0.32)	0.73 (± 0.20)
Bi-PSM	Aggregate	–	0.00 (± 0.00)	0.00 (± 0.00)	0.08 (± 0.04)	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)	0.39 (± 0.11)
	NeedleRegrasp	$\mathbb{R}^{41}/\mathbb{R}^{10}/S$	0.00 (± 0.00)	0.00 (± 0.00)	0.09 (± 0.03)	0.01 (± 0.00)	0.01 (± 0.02)	0.05 (± 0.08)	0.00 (± 0.00)	0.04 (± 0.07)	0.00 (± 0.00)	0.63 (± 0.19)
	BiPegTransfer	$\mathbb{R}^{41}/\mathbb{R}^{10}/S$	0.00 (± 0.00)	0.00 (± 0.00)	0.09 (± 0.05)	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)	0.01 (± 0.02)	0.00 (± 0.00)	0.18 (± 0.14)
Overall	–	0.46 (± 0.03)	0.45 (± 0.01)	0.68 (± 0.02)	0.02 (± 0.02)	0.24 (± 0.03)	0.83 (± 0.05)	0.48 (± 0.01)	0.87 (± 0.03)	0.58 (± 0.08)	0.92 (± 0.02)	

TABLE II: Ablation on demonstration amount. We offer our method and baselines different amounts of demonstrations. Results are IQMs with 95% stratified bootstrap CIs of 5 runs over PSM and Bi-PSM domains.

Method	Number of episodes in demonstrations				
	10 epi.	25 epi.	50 epi.	75 epi.	100 epi.
BC [44]	0.00 (± 0.00)	0.05 (± 0.02)	0.15 (± 0.03)	0.19 (± 0.03)	0.22 (± 0.03)
SQL [67]	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)
VINN [58]	0.03 (± 0.03)	0.02 (± 0.03)	0.02 (± 0.02)	0.01 (± 0.02)	0.01 (± 0.01)
DDPGBC [29]	0.00 (± 0.03)	0.20 (± 0.11)	0.35 (± 0.05)	0.47 (± 0.10)	0.45 (± 0.08)
AMP [27]	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)
CoL [33]	0.18 (± 0.07)	0.47 (± 0.09)	0.46 (± 0.09)	0.49 (± 0.05)	0.43 (± 0.10)
AWAC [55]	0.00 (± 0.05)	0.05 (± 0.14)	0.07 (± 0.18)	0.40 (± 0.15)	0.39 (± 0.13)
DEX (ours)	0.24 (± 0.02)	0.50 (± 0.09)	0.68 (± 0.07)	0.78 (± 0.07)	0.80 (± 0.06)

that guides DDPG actor with Q-filtered BC loss. (7) AMP [27] that augments environment with GAIL [47] reward based on SAC. (8) CoL [33] that initializes the policy with BC offline and integrates BC with DDPG in the online training stage. (9) AWAC [55] that trains an offline RL with demonstrations and online fine-tunes it with a conservative constraint between current and old policies.

B. Main Results on SurRoL Platform

We present the aggregate and distributed performance of each method in Table I. The results show that two pure RL methods achieve great performance in ECM domain, but poor performance in PSM and Bi-PSM. This implies that their exploration strategies are not applicable in complex surgical manipulation tasks, while our demonstration-depend exploration is able to handle these tasks. Similar to pure RL, three IL methods perform well in ECM domain, while their performance decreases dramatically in the PSM and Bi-PSM domain, which has a higher dimension of state and action space and requires more complex manipulation skills. This indicates that, given a limited number of demonstrations, the generalization ability of IL is severely deteriorated by the complexity of the task. Compared to IL, our method not only perform similarly great in ECM domain, but achieves **2.2x** IQM in PSM domain and, especially, **4.9x** in Bi-PSM domain. This demonstrates the advances of our method that leverages demonstrations to accelerate RL exploration.

Meanwhile, we also observe that three demonstration-guided RL methods, except for AMP which suffers from




Fig. 4: Ablation studies. In (a), we propose two variants of DEX that partially regularize the actor-critic to investigate the effectiveness of our proposed critic regularization. In (b), we ablate the choice of exploration coefficient α . In (c), we test the influence of the number of nearest neighbors k in non-parametric propagation. All results are IQMs with 95% stratified bootstrap CIs (error bar) over 5 runs on PSM and Bi-PSM domains.

the instability issue of adversarial training [46], perform favorably in ECM and PSM domains. This may attribute to their utilization of demonstration for RL exploration. However, these methods still get stuck in the complex PSM tasks, such as *PegTransfer*, let alone the Bi-PSM domain that additionally demands coordination skills between two manipulators. Compared with them, our method fully leverages demonstrations to mitigate the guidance impediment of critic and guide the exploration efficiently at demonstration-uncovered states, leading to a prominent performance improvement of **0.39** in the Bi-PSM domain.

C. Ablation Studies

Effect of Demonstration Amount. We further study the manipulation performance of each method with different amounts of demonstrations. We train agents with 10%, 25%, 50%, 75% and 100% of demonstrations, respectively. Table II presents the aggregate results over PSM and Bi-PSM domains. Our proposed method DEX consistently outperforms other competing methods with different amounts of demonstrations. This result indicates that with the same amount of demonstrations, our DEX can utilize demonstrations more efficiently to achieve a higher manipulation performance. More importantly, our proposed DEX exhibits a lower demand for the number of demonstrations for task automation. Even with merely 25% of demonstrations, DEX has outperformed most competing methods that are trained with 100% of demonstrations. This characteristic demonstrates that DEX is well suited for the challenging surgical scenarios where collecting large amounts of demonstrations is unaffordable.

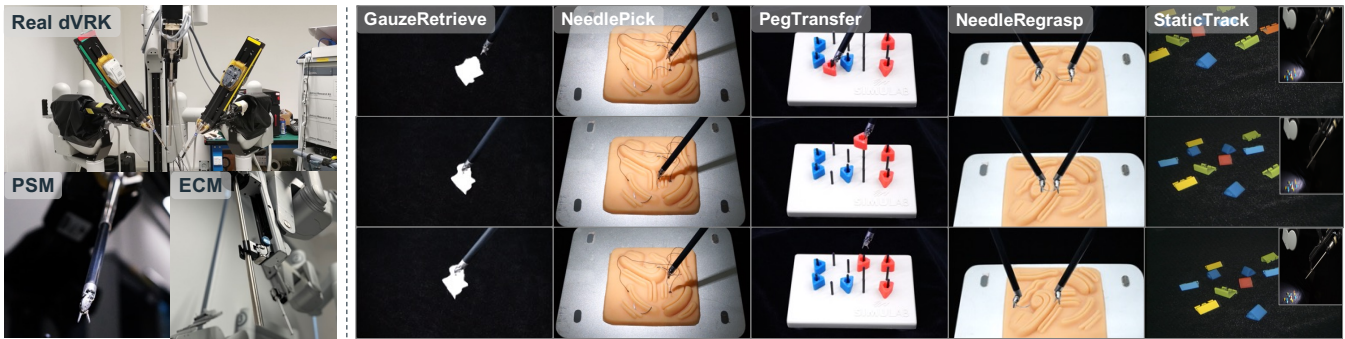


Fig. 5: Deployment on real dVRK. We deploy the learned policies on the real dVRK platform (left) and present the generated trajectories of five representative tasks (right). The successful deployment demonstrates the transfer ability of our method from simulation to real-world surgical scenarios.

Effect of Critic Regularization. While some existing methods only impose the guidance on actor [29, 33, 36], our method guide both actor and critic to remedy the guidance impediments of the ordinary critic. In order to further investigate the effects of critic regularization on the manipulation performance, we propose two variants of DEX that partially regularize the actor-critic modules in Figure 4(a). The results show that, compared with the ordinary no-guidance actor-critic (DEX-AC), the variant that only regularizes actor (DEX-RA) brings an improvement of 0.20 w.r.t. IQM, and the proposed one which regularizes both actor and critic (DEX) brings 0.39. It demonstrates the impediment of inaccurate value estimates of ordinary critic and the effectiveness of our proposed critic regularization to alleviate this issue.

Effect of Exploration Coefficient α . The exploration coefficient α determines the relative importance of the exploration term against environment reward. We investigate the effect of this parameter by gradually increasing the value of α from 0 to 20 and repeatedly train the agent with the remaining settings identical. Figure 4(b) present the experiment results. It shows that α around 5 yields the best performance, indicating that an intermediate choice balances exploration and exploitation well. We also note that an adaptive coefficient (e.g., decay α) may further balance the trade-off to avoid local minima and leave it as future work.

Effect of Guidance Propagation. To investigate the effect of guidance propagation, we propose a parametric variant (DEX-BC) that pre-trains a BC model to approximate the expert action at demonstration-uncovered states. The results show that, compared with no-propagation methods, DEX-BC still achieves a performance improvement of 0.3 IQM over PSM and Bi-PSM domains, though is 0.05 lower than DEX. It demonstrates not only the robustness of non-parametric guidance propagation but the insensitivity to the choice of propagation method in our actor-critic framework.

Effect of k for Guidance Propagation. We investigate the influence of the number of nearest neighbors k by gradually increasing its value from 1 to 9. The results in Figure 4(c) show that too large value of k leads to a slight performance drop and an intermediate value around 5 works well, indicating that our method is not sensitive to the choice of k within an appropriate interval.

TABLE III: Evaluation on real dVRK. We test success rate of the learned policies with 20 trials for each task and compare with two baselines.

Method	GauzeRetrieve	NeedlePick	PegTransfer	NeedleRegrasp	StaticTrack
BC [44]	0.00	0.85	0.00	0.40	1.00
DDPGBC [29]	0.75	0.95	0.35	0.65	1.00
DEX (ours)	0.90	0.95	0.75	0.90	1.00

D. Deployment on Real Robot Platform of dVRK

To validate the transfer ability of our method, we train RL policies in SurRoL and deploy them on dVRK systems. We conduct real-world manipulation experiments on five representative tasks, namely *GauzeRetrieve*, *NeedlePick* and *PegTransfer* in PSM domain, *NeedleRegrasp* in Bi-PSM domain, and *StaticTrack* in ECM domain. These tasks are selected to be diverse and comprehensive, requiring different levels of surgical manipulation skills. We train policies with 300k steps within SurRoL and select checkpoints with the best performance for deployment. We present experiment snapshots in Figure 5. In three PSM tasks, we observe that agents successfully learn skills of reaching, picking, and placing to accomplish the task without collisions. In *NeedleRegrasp*, the policy additionally acquires coordination skills between two manipulators compared with policies in the PSM domain. In *StaticTrack*, the trained policy successfully learns endoscopic control with a high success rate. Meanwhile, we also compare our method with one imitation learning method BC and one competitive demonstration-guided RL method DDPGBC. The results in Table III show that our method achieves the highest success rates at all tasks with 20 trials, which demonstrates the great potential of our method for real-world surgical task automation.

V. CONCLUSION

We present DEX, a novel demonstration-guided RL method that narrows down exploration space by encouraging expert-like behaviors and enabling robust guidance when confronting states unobserved in demonstrations, to improve exploration efficiency with a modest set of demonstrations. We first demonstrate its performance improvement on 10 challenging surgical manipulation tasks compared with state-of-the-art methods on the SurRoL platform. We also deploy the trained policy to the dVRK platform to show its potential for transferring to real-world surgical automation scenarios.

REFERENCES

- [1] C. D’Ettorre, A. Mariani, A. Stilli, F. R. y Baena, P. Valdastrì, A. Deguet, P. Kazanzides, R. H. Taylor, G. S. Fischer, S. P. DiMaio, *et al.*, “Accelerating surgical robotics research: A review of 10 years with the da vinci research kit,” *IEEE Robotics & Automation Magazine*, 2021.
- [2] K. L. Schwaner, D. Dall’Alba, P. T. Jensen, P. Fiorini, and T. R. Savarimuthu, “Autonomous needle manipulation for robotic surgical suturing based on skills learned from demonstration,” in *IEEE International Conference on Automation Science and Engineering (CASE)*, 2021.
- [3] V. M. Varier, D. K. Rajamani, N. Goldfarb, F. Tavakkolmoghaddam, A. Munawar, and G. S. Fischer, “Collaborative suturing: A reinforcement learning approach to automate hand-off task in suturing for surgical robots,” *IEEE International Conference on Robot and Human Interactive Communication*, 2020.
- [4] Y. Barnoy, M. O’Brien, W. Wang, and G. Hager, “Robotic surgery with lean reinforcement learning,” *arXiv preprint arXiv:2105.01006*, 2021.
- [5] H. Gao, W. Fan, L. Qiu, X. Yang, Z. Li, X. Zuo, Y. Li, M. Q.-H. Meng, and H. Ren, “Savanet: Surgical action-driven visual attention network for autonomous endoscope control,” *IEEE Transactions on Automation Science and Engineering (TASE)*, 2022.
- [6] J. J. Ji, S. Krishnan, V. Patel, D. Fer, and K. Goldberg, “Learning 2d surgical camera motion from demonstrations,” in *IEEE International Conference on Automation Science and Engineering (CASE)*, 2018.
- [7] A. Pandya, L. A. Reisner, B. King, N. Lucas, A. Composto, M. Klein, and R. D. Ellis, “A review of camera viewpoint automation in robotic and laparoscopic surgery,” *Robotics*, 2014.
- [8] H. Saeidi, J. D. Opfermann, M. Kam, S. Wei, S. Léonard, M. H. Hsieh, J. U. Kang, and A. Krieger, “Autonomous robotic laparoscopic surgery for intestinal anastomosis,” *Science Robotics*, 2022.
- [9] E. Tagliabue, A. Pore, D. Dall’Alba, E. Magnabosco, M. Piccinelli, and P. Fiorini, “Soft tissue simulation environment to learn manipulation tasks in autonomous robotic surgery,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [10] A. Pore, D. Corsi, E. Marchesini, D. Dall’Alba, A. Casals, A. Farinelli, and P. Fiorini, “Safe reinforcement learning using formal verification for tissue retraction in autonomous robotic-assisted surgery,” *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [11] B. Thananjeyan, A. Garg, S. Krishnan, C. Chen, L. Miller, and K. Goldberg, “Multilateral surgical pattern cutting in 2d orthotropic gauze with deep reinforcement learning policies for tensioning,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
- [12] T. Nguyen, N. D. Nguyen, F. Bello, and S. Nahavandi, “A new tensioning method using deep reinforcement learning for surgical pattern cutting,” in *2019 IEEE international conference on industrial technology (ICIT)*, 2019.
- [13] N. D. Nguyen, T. Nguyen, S. Nahavandi, A. Bhatti, and G. Guest, “Manipulating soft tissues by deep reinforcement learning for autonomous robotic surgery,” in *2019 IEEE International Systems Conference (SysCon)*. IEEE, 2019, pp. 1–7.
- [14] H. Su, Y. Hu, Z. Li, A. Knoll, G. Ferrigno, and E. De Momi, “Reinforcement learning based manipulation skill transferring for robot-assisted minimally invasive surgery,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [15] W. Chi, G. Dagnino, T. M. Y. Kwok, A. Nguyen, D. Kundrat, M. E. M. K. Abdelaziz, C. V. Riga, C. D. Bicknell, and G.-Z. Yang, “Collaborative robot-assisted endovascular catheterization with generative adversarial imitation learning,” *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [16] P. M. Scheickl, B. Gyenes, T. Davitashvili, R. Younis, A. Schulze, B. P. Müller-Stich, G. Neumann, M. Wagner, and F. Mathis-Ullrich, “Cooperative assistance in robotic surgery through multi-agent reinforcement learning,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [17] A. Segato, M. Di Marzo, S. Zucchelli, S. Galvan, R. Secoli, and E. De Momi, “Inverse reinforcement learning intra-operative path planning for steerable needle,” *IEEE Transactions on Biomedical Engineering*, 2021.
- [18] X. Tan, C.-B. Chng, Y. Su, K.-B. Lim, and C.-K. Chui, “Robot-assisted training in laparoscopy using deep reinforcement learning,” *Robotics and Automation Letters (RA-L)*, 2019.
- [19] B. Keller, M. Draelos, K. Zhou, R. Qian, A. N. Kuo, G. Konidakis, K. Hauser, and J. A. Izatt, “Optical coherence tomography-guided robotic ophthalmic microsurgery via reinforcement learning from demonstration,” *IEEE Transactions on Robotics*, 2020.
- [20] C. D’Ettorre, S. Zirino, N. N. Dei, A. Stilli, E. De Momi, and D. Stoyanov, “Learning intraoperative organ manipulation with context-based reinforcement learning,” *International Journal of Computer Assisted Radiology and Surgery*, 2022.
- [21] J. Ramírez, W. Yu, and A. Perrusquía, “Model-free reinforcement learning from expert demonstrations: a survey,” *Artificial Intelligence Review*, 2022.
- [22] S. Levine and V. Koltun, “Guided policy search,” in *International Conference on Machine Learning (ICML)*, 2013.
- [23] T. Hester, M. Vecerík, O. Pietquin, M. Lanctot, T. Schaul, B. Piot, D. Horgan, J. Quan, A. Sendonaris, I. Osband, G. Dulac-Arnold, J. P. Agapiou, J. Z. Leibo, and A. Gruslys, “Deep q-learning from demonstrations,” in *Association for the Advancement of Artificial Intelligence (AAAI)*, 2018.
- [24] M. Vecerík, T. Hester, J. Scholz, F. Wang, O. Pietquin, B. Piot, N. M. O. Heess, T. Rothörl, T. Lampe, and M. A. Riedmiller, “Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards,” *arXiv preprint arXiv:1707.08817*, 2017.
- [25] Y. Zhu, Z. Wang, J. Merel, A. A. Rusu, T. Erez, S. Cabi, S. Tunyasuvunakool, J. Kramár, R. Hadsell, N. de Freitas, and N. M. O. Heess, “Reinforcement and imitation learning for diverse visuomotor skills,” in *Robotics: Science and Systems (RSS)*, 2018.
- [26] X. B. Peng, P. Abbeel, S. Levine, and M. van de Panne, “Deepmimic: Example-guided deep reinforcement learning of physics-based character skills,” *ACM Transactions on Graphics (ToG)*, 2018.
- [27] X. B. Peng, Z. Ma, P. Abbeel, S. Levine, and A. Kanazawa, “Amp: Adversarial motion priors for stylized physics-based character control,” *ACM Transactions on Graphics (ToG)*, 2021.
- [28] A. Pore, E. Tagliabue, M. Piccinelli, D. Dall’Alba, A. Casals, and P. Fiorini, “Learning from demonstrations for autonomous soft-tissue retraction,” in *2021 International Symposium on Medical Robotics (ISMR)*, 2021.
- [29] A. Nair, B. McGrew, M. Andrychowicz, W. Zaremba, and P. Abbeel, “Overcoming exploration in reinforcement learning with demonstrations,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [30] A. Rajeswaran, V. Kumar, A. Gupta, J. Schulman, E. Todorov, and S. Levine, “Learning complex dexterous manipulation with deep reinforcement learning and demonstrations,” in *Robotics: Science and Systems (RSS)*, 2018.
- [31] H. Zhu, A. Gupta, A. Rajeswaran, S. Levine, and V. Kumar, “Dexterous manipulation with deep reinforcement learning: Efficient, general, and low-cost,” in *International Conference on Robotics and Automation (ICRA)*, 2019.
- [32] R. M. Shah and V. Kumar, “Rrl: Resnet as representation for reinforcement learning,” in *International Conference on Machine Learning (ICML)*, 2021.
- [33] V. G. Goecks, G. M. Gremillion, V. J. Lawhern, J. Valasek, and N. R. Waytowich, “Integrating behavior cloning and reinforcement learning for improved performance in dense and sparse reward environments,” in *International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2020, pp. 465–473.
- [34] S. Fujimoto, H. Hoof, and D. Meger, “Addressing function approximation error in actor-critic methods,” in *International Conference on Machine Learning (ICML)*, 2018.
- [35] H. V. Hasselt, A. Guez, and D. Silver, “Deep reinforcement learning with double q-learning,” in *Association for the Advancement of Artificial Intelligence (AAAI)*, 2016.
- [36] Z.-Y. Chiu, F. Richter, E. K. Funk, R. K. Orosco, and M. C. Yip, “Bimanual regrasping for suture needles using reinforcement learning for rapid motion planning,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [37] J. Xu, B. Li, B. Lu, Y.-H. Liu, Q. Dou, and P.-A. Heng, “Surrol: An open-source reinforcement learning centered and dvrk compatible platform for surgical robot learning,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [38] S. Leonard, K. L. Wu, Y. Kim, A. Krieger, and P. C. Kim, “Smart tissue anastomosis robot (star): A vision-guided robotics system for laparoscopic suturing,” *IEEE Transactions on Biomedical Engineering*, 2014.
- [39] A. Attanasio, B. Scaglioni, M. Leonetti, A. F. Frangi, W. Cross, C. S. Biyani, and P. Valdastrì, “Autonomous tissue retraction in

- robotic assisted minimally invasive surgery—a feasibility study,” *IEEE Robotics and Automation Letters (RA-L)*, 2020.
- [40] V. Patel, S. Krishnan, A. Goncalves, C. Chen, W. D. Boyd, and K. Goldberg, “Using intermittent synchronization to compensate for rhythmic body motion during autonomous surgical cutting and debriement,” in *International Symposium on Medical Robotics (ISMR)*, 2018.
- [41] A. Murali, S. Sen, B. Kehoe, A. Garg, S. McFarland, S. Patil, W. D. Boyd, S. Lim, P. Abbeel, and K. Goldberg, “Learning by observation for surgical subtasks: Multilateral cutting of 3d viscoelastic and 2d orthotropic tissue phantoms,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015.
- [42] M. Yip and N. Das, “Robot autonomy for surgery,” in *The Encyclopedia of MEDICAL ROBOTICS: Volume 1 Minimally Invasive Surgical Robotics*, 2019.
- [43] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, “A survey of robot learning from demonstration,” *Robotics and Autonomous Systems*, 2009.
- [44] M. Bain and C. Sammut, “A framework for behavioural cloning,” in *Machine Intelligence*, 1999.
- [45] J. Ibarz, J. Tan, C. Finn, M. Kalakrishnan, P. Pastor, and S. Levine, “How to train your robot with deep reinforcement learning: lessons we have learned,” *The International Journal of Robotics Research (IJRR)*, 2021.
- [46] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, J. Peters, *et al.*, “An algorithmic perspective on imitation learning,” *Foundations and Trends® in Robotics*, 2018.
- [47] J. Ho and S. Ermon, “Generative adversarial imitation learning,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [48] Y. Ding, C. Florensa, P. Abbeel, and M. Phielipp, “Goal-conditioned imitation learning,” *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [49] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, “Maximum entropy inverse reinforcement learning,” in *Association for the Advancement of Artificial Intelligence (AAAI)*, 2008.
- [50] J. Fu, K. Luo, and S. Levine, “Learning robust rewards with adversarial inverse reinforcement learning,” in *International Conference on Learning Representations (ICLR)*, 2018.
- [51] I. Kostrikov, K. K. Agrawal, D. Dwibedi, S. Levine, and J. Tompson, “Discriminator-actor-critic: Addressing sample inefficiency and reward bias in adversarial imitation learning,” in *International Conference on Learning Representations (ICLR)*, 2019.
- [52] I. Kostrikov, R. Fergus, J. Tompson, and O. Nachum, “Offline reinforcement learning with fisher divergence critic regularization,” in *International Conference on Machine Learning (ICML)*, 2021.
- [53] S. Fujimoto, D. Meger, and D. Precup, “Off-policy deep reinforcement learning without exploration,” in *International Conference on Machine Learning (ICML)*, 2019.
- [54] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in *International Conference on Machine Learning (ICML)*, 2018.
- [55] A. Nair, A. Gupta, M. Dalal, and S. Levine, “Awac: Accelerating online reinforcement learning with offline datasets,” *arXiv preprint arXiv:2006.09359*, 2020.
- [56] Y. Wu, G. Tucker, and O. Nachum, “Behavior regularized offline reinforcement learning,” *arXiv preprint arXiv:1911.11361*, 2019.
- [57] C. G. Atkeson, A. W. Moore, and S. Schaal, “Locally weighted learning,” *Lazy learning*, 1997.
- [58] J. Pari, N. M. M. Shafiullah, S. P. Arunachalam, and L. Pinto, “The surprising effectiveness of representation learning for visual imitation,” in *Robotics: Science and Systems (RSS)*, 2022.
- [59] K. Hakhamaneshi, R. Zhao, A. Zhan, P. Abbeel, and M. Laskin, “Hierarchical few-shot imitation with skill transition models,” in *International Conference on Learning Representations (ICLR)*, 2022.
- [60] S. P. Arunachalam, S. Silwal, B. Evans, and L. Pinto, “Dexterous imitation made easy: A learning-based framework for efficient dexterous manipulation,” *arXiv preprint arXiv:2203.13251*, 2022.
- [61] 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 (ICLR)*, 2016.
- [62] P. Dhariwal, C. Hesse, O. Klimov, A. Nichol, M. Plappert, A. Radford, J. Schulman, S. Sidor, Y. Wu, and P. Zhokhov, “Openai baselines,” <https://github.com/openai/baselines>, 2017.
- [63] T. G. Rudner, C. Lu, M. A. Osborne, Y. Gal, and Y. Teh, “On pathologies in kl-regularized reinforcement learning from expert demonstrations,” *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [64] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *International Conference on Learning Representations (ICLR)*, 2015.
- [65] M. Andrychowicz, D. Crow, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, P. Abbeel, and W. Zaremba, “Hindsight experience replay,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [66] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, and S. Levine, “Soft actor-critic algorithms and applications,” *arXiv preprint arXiv:1812.05905*, 2018.
- [67] S. Reddy, A. D. Dragan, and S. Levine, “Sqil: Imitation learning via reinforcement learning with sparse rewards,” in *International Conference on Learning Representations (ICLR)*, 2019.
- [68] M. Laskin, A. Srinivas, and P. Abbeel, “Curl: Contrastive unsupervised representations for reinforcement learning,” in *International Conference on Machine Learning (ICML)*, 2020.
- [69] D. Yarats, I. Kostrikov, and R. Fergus, “Image augmentation is all you need: Regularizing deep reinforcement learning from pixels,” in *International Conference on Learning Representations (ICLR)*, 2020.
- [70] R. Agarwal, M. Schwarzer, P. S. Castro, A. C. Courville, and M. Bellemare, “Deep reinforcement learning at the edge of the statistical precipice,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.