

LazyDagger: Reducing Context Switching in Interactive Imitation Learning

Ryan Hoque, Ashwin Balakrishna, Carl Putterman, Michael Luo,
Daniel S. Brown, Daniel Seita, Brijen Thananjayan, Ellen Novoseller, Ken Goldberg

Abstract—Corrective interventions while a robot is learning to automate a task provide an intuitive method for a human supervisor to assist the robot and convey information about desired behavior. However, these interventions can impose significant burden on a human supervisor, as each intervention interrupts other work the human is doing, incurs latency with each context switch between supervisor and autonomous control, and requires time to perform. We present LazyDagger, which extends the interactive imitation learning (IL) algorithm SafeDagger to reduce context switches between supervisor and autonomous control. We find that LazyDagger improves the performance and robustness of the learned policy during both learning and execution while limiting burden on the supervisor. Simulation experiments suggest that LazyDagger can reduce context switches by an average of 60% over SafeDagger on 3 continuous control tasks while maintaining state-of-the-art policy performance. In physical fabric manipulation experiments with an ABB YuMi robot, LazyDagger reduces context switches by 60% while achieving a 60% higher success rate than SafeDagger at execution time.

I. INTRODUCTION

Imitation learning allows a robot to learn from human feedback and examples [3, 4, 33]. In particular, *interactive imitation learning* (IL) [24, 45, 53], in which a human supervisor periodically takes control of the robotic system during policy learning, has emerged as a popular imitation learning method, as interventions are a particularly intuitive form of human feedback [45]. However, a key challenge in interactive imitation learning is to reduce the burden that interventions place on the human supervisor [24, 53].

One source of this burden is the cost of *context switches* between human and robot control. Context switches incur significant time cost, as a human must interrupt the task they are currently performing, acquire control of the robot, and gain sufficient situational awareness before beginning the intervention. As an illustrative example, consider a robot performing a task for which an action takes 1 time unit and an intervention requires two context switches (one at the start and one at the end). We define *latency* L as the number of time units associated with a single context switch. For instance, $L \gg 1$ for a human supervisor who will need to pause an ongoing task and walk over to a robot that requires assistance. If the supervisor takes control 10 times for 2 actions each, she spends $20L + 20$ time units helping the robot. In contrast, if the human takes control 2 times for 10 actions each, she spends only $4L + 20$ time units. The latter significantly reduces the burden on the

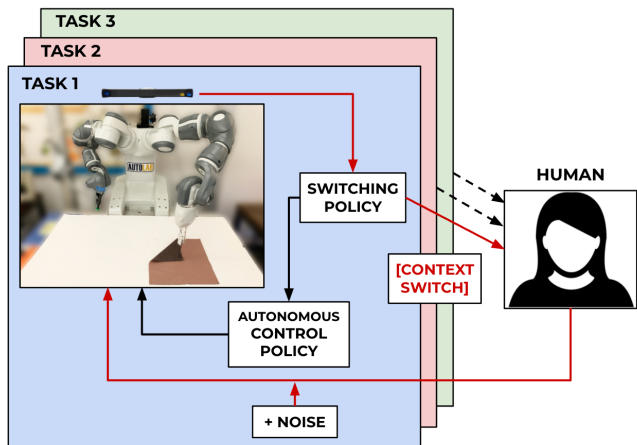


Fig. 1: LazyDagger learns to cede control to a supervisor in states in which it estimates that its actions will significantly deviate from those of the supervisor. LazyDagger reduces context switches between supervisor and autonomous control to reduce burden on a human supervisor working on multiple tasks.

supervisor. Furthermore, prior work suggests that frequent context switches can make it difficult for the supervisor to perform other tasks in parallel [46] or gain enough situational awareness to provide useful interventions [39].

We present LazyDagger (Figure 1), an algorithm which initiates useful interventions while limiting context switches. The name LazyDagger is inspired by the concept of lazy evaluation in programming language theory [29], where expressions are evaluated only when required to reduce computational burden. As in SafeDagger [53], LazyDagger learns a meta-controller which determines when to context switch based on the estimated discrepancy between the learner and supervisor. However, unlike SafeDagger, LazyDagger reduces context switching by (1) introducing asymmetric switching criteria and (2) injecting noise into the supervisor control actions to widen the distribution of visited states. One appealing property of this improved meta-controller is that even after training, LazyDagger can be applied at execution time to improve the safety and reliability of autonomous policies with minimal context switching. We find that across 3 continuous control tasks in simulation, LazyDagger achieves task performance on par with Dagger [40] with 88% fewer supervisor actions than Dagger and 60% fewer context switches than SafeDagger. In physical fabric manipulation experiments, we observe similar results, and find that at execution time, LazyDagger achieves 60% better task performance than SafeDagger with 60% fewer context switches.

II. BACKGROUND AND RELATED WORK

Challenges in learning efficiency and reward function specification have inspired significant interest in algorithms that can leverage supervisor demonstrations and feedback for policy learning.

Learning from Offline Demonstrations: Learning from demonstrations [3, 4, 33] is a popular imitation learning approach, as it requires minimal supervisor burden: the supervisor provides a batch of offline demonstrations and gives no further input during policy learning. Many methods use demonstrations directly for policy learning [20, 35, 36, 49], while others use reinforcement learning to train a policy using a reward function inferred from demonstrations [1, 9, 15, 17, 54]. Recent work has augmented demonstrations with additional offline information such as pairwise preferences [10, 11], human gaze [42], and natural language descriptions [50]. While offline demonstrations are often simple to provide, the lack of online feedback makes it difficult to address specific bottlenecks in the learning process or errors in the resulting policy due to covariate shift [40].

Learning from Online Feedback: Many policy learning algorithms’ poor performance stems from a lack of online supervisor guidance, motivating active learning methods such as DAgger, which queries the supervisor for an action in every state that the learner visits [40]. While DAgger has a number of desirable theoretical properties, labeling every state is costly in human time and can be a non-intuitive form of human feedback [27]. More generally, the idea of learning from action advice has been widely explored in imitation learning algorithms [5, 21, 22, 28]. There has also been significant recent interest in active preference queries for learning reward functions from pairwise preferences over demonstrations [7, 10, 13, 19, 34, 41]. However, many forms of human advice can be unintuitive, since the learner may visit states that are significantly far from those the human supervisor would visit, making it difficult for humans to judge what correct behavior looks like without interacting with the environment themselves [38, 45].

Learning from Supervisor Interventions: There has been significant prior work on algorithms for learning policies from interventions. Kurenkov *et al.* [26] and Xie *et al.* [52] leverage interventions from suboptimal supervisors to accelerate policy learning, but assume that the supervisors are algorithmic and thus can be queried cheaply. Nolan Wagener [32], Saunders *et al.* [43], and Thananjeyan* *et al.* [47] also leverage interventions from algorithmic policies, but for constraint satisfaction during learning. Amir *et al.* [2], Kahn *et al.* [23], Kelly *et al.* [24], Mandlekar *et al.* [30], Spencer *et al.* [45], and Wang *et al.* [51] instead consider learning from human supervisors and present learning algorithms which utilize the timing and nature of human interventions to update the learned policy. By giving the human control for multiple timesteps in a row, these algorithms show improvements over methods that only hand over control on a state-by-state basis [6]. However, the above algorithms assume that the human is continuously monitoring the system to determine

when to intervene, which may not be practical in large-scale systems or continuous learning settings [12, 14, 25, 46]. Such algorithms also assume that the human knows when to cede control to the robot, which requires guessing how the robot will behave in the future. Zhang *et al.* [53] and Menda *et al.* [31] present imitation learning algorithms SafeDAgger and EnsembleDAgger, respectively, to address these issues by learning to request interventions from a supervisor based on measures such as state novelty or estimated discrepancy between the learner and supervisor actions. These methods can still be sample inefficient, and, as we discuss later, often result in significant context switching. By contrast, LazyDAgger encourages interventions that are both easier to provide and more informative. To do this, LazyDAgger prioritizes (1) sustained interventions, which allow the supervisor to act over a small number of contiguous sequences of states rather than a large number of disconnected intervals, and (2) interventions which demonstrate supervisor actions in novel states to increase robustness to covariate shift in the learned policy.

III. PROBLEM STATEMENT

We consider a setting in which a human supervisor is training a robot to reliably perform a task. The robot may query the human for assistance, upon which the supervisor takes control and teleoperates the robot until the system determines that it no longer needs assistance. We assume that the robot and human policy have the same action space, and that it is possible to pause task execution while waiting to transfer control. We formalize these ideas in the context of prior imitation learning literature.

We model the environment as a discrete-time Markov decision process (MDP) \mathcal{M} with states $s \in \mathcal{S}$, actions $a \in \mathcal{A}$, and time horizon T [37]. The robot does not have access to the reward function or transition dynamics of \mathcal{M} but can cede control to a human supervisor, who executes some deterministic policy $\pi_H : \mathcal{S} \rightarrow \mathcal{A}$. We refer to times when the robot is in control as *autonomous mode* and those in which the supervisor is in control as *supervisor mode*. We minimize a surrogate loss function $J(\pi_R)$ to encourage the robot policy $\pi_R : \mathcal{S} \rightarrow \mathcal{A}$ to match that of the supervisor (π_H):

$$J(\pi_R) = \sum_{t=1}^T \mathbb{E}_{s_t \sim d_t^{\pi_R}} [\mathcal{L}(\pi_R(s_t), \pi_H(s_t))], \quad (\text{III.1})$$

where $\mathcal{L}(\pi_R(s), \pi_H(s))$ is an action discrepancy measure between $\pi_R(s)$ and $\pi_H(s)$ (e.g., the squared loss or 0-1 loss), and $d_t^{\pi_R}$ is the marginal state distribution at timestep t induced by executing π_R in MDP \mathcal{M} .

In interactive IL we require a meta-controller π that determines whether to query the robot policy π_R or to query for an intervention from the human supervisor policy π_H ; importantly, π consists of both (1) the high-level controller which decides whether to switch between π_R and π_H and (2) the low-level robot policy π_R . A key objective in interactive IL is to minimize some notion of supervisor burden. To this end, let $m_I(s; \pi)$ be an indicator which records whether a context switch between autonomous (π_R) and supervisor (π_H)

modes occurs at state s_t (either direction). Then, we define $C(\pi)$, the expected number of context switches in an episode under policy π , as follows: $C(\pi) = \sum_{t=1}^T \mathbb{E}_{s_t \sim d_t^\pi} [m_I(s_t; \pi)]$, where d_t^π is the marginal state distribution at timestep t induced by executing the meta-controller π in MDP \mathcal{M} . Similarly, let $m_H(s_t; \pi)$ indicate whether the system is in supervisor mode at state s_t . We then define $D(\pi)$, the expected number of supervisor actions in an episode for the policy π , as follows: $D(\pi) = \sum_{t=1}^T \mathbb{E}_{s_t \sim d_t^\pi} [m_H(s_t; \pi)]$.

We define *supervisor burden* $B(\pi)$ as the expected time cost imposed on the human supervisor. This can be expressed as the sum of the expected total number of time units spent in context switching and the expected total number of time units in which the supervisor is actually engaged in performing interventions:

$$B(\pi) = L \cdot C(\pi) + D(\pi), \quad (\text{III.2})$$

where L is context switch latency (Section I) in time units, and each time unit is the time it takes for the supervisor to execute a single action. The learning objective is to find a policy π that matches supervisor performance, π_H , while limiting supervisor burden to lie within a threshold Γ_b , set by the supervisor to an acceptable tolerance for a given task. To formalize this problem, we propose the following objective:

$$\pi = \arg \min_{\pi' \in \Pi} \{J(\pi'_R) \mid B(\pi') \leq \Gamma_b\}, \quad (\text{III.3})$$

where Π is the space of all meta-controllers, and π'_R is the low-level robot policy associated with meta-controller π' .

IV. PRELIMINARIES: SAFEDAGGER

We consider interactive IL in the context of the objective introduced in Equation (III.3): to maximize task reward while limiting supervisor burden. To do this, LazyDagger builds on SafeDagger [53], a state-of-the-art algorithm for interactive IL. SafeDagger selects between autonomous mode and supervisor mode by training a binary action discrepancy classifier f to discriminate between “safe” states which have an action discrepancy below a threshold β_H (i.e., states with $\mathcal{L}(\pi_R(s), \pi_H(s)) < \beta_H$) and “unsafe” states (i.e. states with $\mathcal{L}(\pi_R(s), \pi_H(s)) \geq \beta_H$). The classifier f is a neural network with a sigmoid output layer (i.e., $f(s) \in [0, 1]$) that is trained to minimize binary cross-entropy (BCE) loss on the datapoints $(s_t, \pi_H(s_t))$ sampled from a dataset \mathcal{D} of trajectories collected from π_H . This is written as follows:

$$\begin{aligned} \mathcal{L}_S(\pi_R(s_t), \pi_H(s_t), f) &= -f^*(\pi_R(s_t), \pi_H(s_t)) \log f(s_t) \\ &\quad - (1 - f^*(\pi_R(s_t), \pi_H(s_t))) \log(1 - f(s_t)), \end{aligned} \quad (\text{IV.1})$$

where the training labels are given by $f^*(\pi_R(s_t), \pi_H(s_t)) = \mathbb{1}\{\mathcal{L}(\pi_R(s_t), \pi_H(s_t)) \geq \beta_H\}$, and $\mathbb{1}$ denotes the indicator function. Thus, $\mathcal{L}_S(\pi_R(s_t), \pi_H(s_t), f)$ penalizes incorrectly classifying a “safe” state as “unsafe” and vice versa.

SafeDagger executes the meta-policy π which selects between π_R and π_H as follows:

$$\pi(s_t) = \begin{cases} \pi_R(s_t) & \text{if } f(s_t) < 0.5 \\ \pi_H(s_t) & \text{otherwise,} \end{cases} \quad (\text{IV.2})$$

Algorithm 1 LazyDagger

Require: Number of epochs N , time steps per epoch T , intervention thresholds β_H, β_R , supervisor policy π_H , noise σ^2

- 1: Collect $\mathcal{D}, \mathcal{D}_S$ offline with supervisor policy π_H
- 2: $\pi_R \leftarrow \arg \min_{\pi_R} \mathbb{E}_{(s_t, \pi_H(s_t)) \sim \mathcal{D}} [\mathcal{L}(\pi_R(s_t), \pi_H(s_t))] \quad \triangleright \text{Eq. (III.1)}$
- 3: $f \leftarrow \arg \min_f \mathbb{E}_{(s_t, \pi_H(s_t)) \sim \mathcal{D} \cup \mathcal{D}_S} [\mathcal{L}_S(\pi_R(s_t), \pi_H(s_t), f)] \quad \triangleright \text{Eq. (IV.1)}$
- 4: **for** $i \in \{1, \dots, N\}$ **do**
- 5: Initialize s_0 , Mode \leftarrow Autonomous
- 6: **for** $t \in \{1, \dots, T\}$ **do**
- 7: $a_t \sim \pi_R(s_t)$
- 8: **if** Mode = Supervisor or $f(s_t) \geq 0.5$ **then**
- 9: $a_t^H = \pi_H(s_t)$
- 10: $\mathcal{D} \leftarrow \mathcal{D} \cup \{(s_t, a_t^H)\}$
- 11: Execute $\tilde{a}_t^H \sim \mathcal{N}(a_t^H, \sigma^2 I)$
- 12: **if** $\mathcal{L}(a_t, \tilde{a}_t^H) < \beta_R$ **then**
- 13: Mode \leftarrow Autonomous
- 14: **else**
- 15: Mode \leftarrow Supervisor
- 16: **else**
- 17: Execute a_t
- 18: $\pi_R \leftarrow \arg \min_{\pi_R} \mathbb{E}_{(s_t, \pi_H(s_t)) \sim \mathcal{D}} [\mathcal{L}(\pi_R(s_t), \pi_H(s_t))]$
- 19: $f \leftarrow \arg \min_f \mathbb{E}_{(s_t, \pi_H(s_t)) \sim \mathcal{D} \cup \mathcal{D}_S} [\mathcal{L}_S(\pi_R(s_t), \pi_H(s_t), f)]$

where $f(s_t) < 0.5$ corresponds to a prediction that $\mathcal{L}(\pi_R(s_t), \pi_H(s_t)) < \beta_H$, i.e., that s_t is “safe.” Intuitively, SafeDagger only solicits supervisor actions when f predicts that the action discrepancy between π_R and π_H exceeds the safety threshold β_H . Thus, SafeDagger provides a mechanism for querying the supervisor for interventions only when necessary. In LazyDagger, we utilize this same mechanism to query for interventions but enforce new properties once we enter these interventions to lengthen them and increase the diversity of states observed during the interventions.

V. LAZYDAGGER

We summarize LazyDagger in Algorithm 1. In the initial phase (Lines 1-3), we train π_R and safety classifier f on offline datasets collected from the supervisor policy π_H . In the interactive learning phase (Lines 4-19), we evaluate and update the robot policy for N epochs, ceding control to the supervisor when the robot predicts a high action discrepancy.

A. Action Discrepancy Prediction

SafeDagger uses the classifier f to select between π_R and π_H (Equation (IV.2)). However, in practice, this often leads to frequent context switching (Figure 3). To mitigate this, we make two observations. First, we can leverage that in supervisor mode, we directly observe the supervisor’s

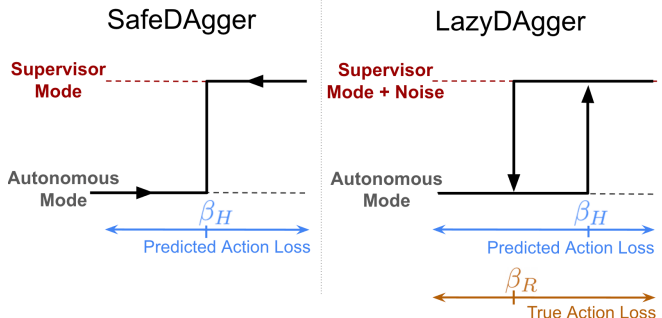


Fig. 2: **LazyDagger Switching Strategy:** SafeDagger switches between supervisor and autonomous mode if the predicted action discrepancy is above threshold β_H . In contrast, LazyDagger uses asymmetric switching criteria and switches to autonomous mode based on ground truth action discrepancy. The gap between β_R and β_H defines a hysteresis band [8].

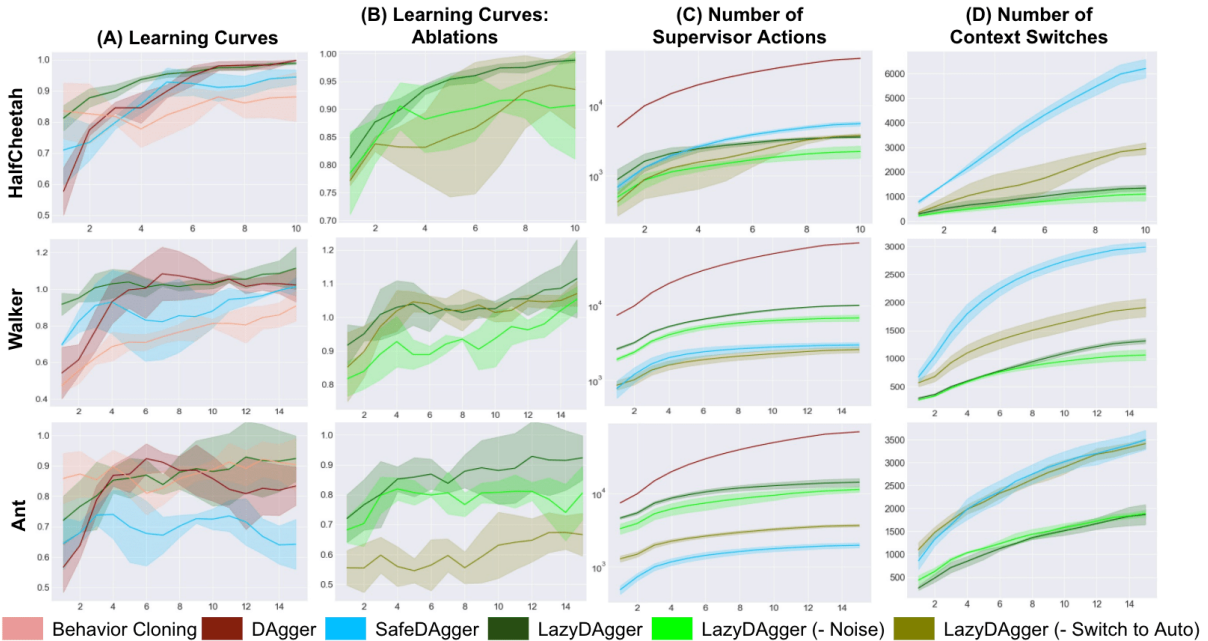


Fig. 3: **MuJoCo Simulation Results:** We study task performance (A), ablations (B), online supervisor burden (C), and total bidirectional context switches (D) for LazyDagger and baselines over 3 random seeds. For Columns (A)-(D), the x-axis for all plots shows the number of epochs over the training dataset, while the y-axes indicate normalized reward (A, B), counts of supervisor actions (C, log scale), and context switches (D) with shading for 1 standard deviation. We find that LazyDagger outperforms all baselines and ablations, indicating that encouraging lengthy, noisy interventions improves performance. Additionally, LazyDagger uses far fewer context switches than other baselines while requesting far fewer supervisor actions than DAgger.

actions. Thus, there is no need to use f , which may have approximation errors, to determine whether to remain in supervisor mode; instead, we can compute the ground-truth action discrepancy $\mathcal{L}(\pi_R(s_t), \pi_H(s_t))$ exactly for any state s_t visited in supervisor mode by comparing the supplied supervisor action $\pi_H(s_t)$ with the action proposed by the robot policy $\pi_R(s_t)$. In contrast, SafeDagger uses f to determine when to switch modes both in autonomous and supervisor mode, which can lead to very short interventions when f prematurely predicts that the agent can match the supervisor’s actions. Second, to ensure the robot has returned to the supervisor’s distribution, the robot should only switch back to autonomous mode when the action discrepancy falls below a threshold β_R , where $\beta_R < \beta_H$. As illustrated in Figure 2, LazyDagger’s asymmetric switching criteria create a hysteresis band, as is often utilized in control theory [8]. Motivated by Eq. (III.3), we adjust β_H to reduce context switches $C(\pi)$ and adjust β_R as a function of β_H to increase intervention length. We hypothesize that redistributing the supervisor actions into fewer but longer sequences in this fashion both reduces burden on the supervisor and improves the quality of the online feedback for the robot. Details on setting these hyperparameter values in practice, the settings used in our experiments, and a hyperparameter sensitivity analysis are provided in the Appendix.

B. Noise Injection

If the safety classifier is querying for interventions at state s_t , then the robot either does not have much experience in the neighborhood of s_t or has trouble matching the demonstrations at s_t . This motivates exploring novel states near s_t so that the robot can receive maximal feedback on the

correct behavior in areas of the state space where it predicts a large action discrepancy from the supervisor. Inspired by prior work that has identified noise injection as a useful tool for improving the performance of imitation learning algorithms (e.g. Laskey *et al.* [27] and Brown *et al.* [9]), we diversify the set of states visited in supervisor mode by injecting isotropic Gaussian noise into the supervisor’s actions, where the variance σ^2 is a scalar hyperparameter (Line 11 in Algorithm 1).

VI. EXPERIMENTS

We study whether LazyDagger can (1) reduce supervisor burden while (2) achieving similar or superior task performance compared to prior algorithms. Implementation details are provided in the supplementary material. In all experiments, \mathcal{L} measures Euclidean distance.

A. Simulation Experiments: MuJoCo Benchmarks

Environments: We evaluate LazyDagger and baselines on 3 continuous control environments from MuJoCo [48], a standard simulator for evaluating imitation and reinforcement learning algorithms. In particular, we evaluate on HalfCheetah-v2, Walker2D-v2 and Ant-v2.

Metrics: For LazyDagger and all baselines, we report learning curves which indicate how quickly they can make task progress in addition to metrics regarding the burden imposed on the supervisor. To study supervisor burden, we report the number of supervisor actions, the number of context switches, and the total supervisor burden (as defined in Eq. (III.2)). Additionally, we define $L^* \geq 0$ to be the latency value such that for all $L > L^*$, LazyDagger has a lower supervisor burden than SafeDagger. We report this




Fig. 4: **Fabric Smoothing Simulation Results:** We study task performance measured by final fabric coverage (A), total supervisor actions (B), and total context switches (C) for LazyDagger and baselines in the Gym-Cloth environment from [44]. The horizontal dotted line shows the success threshold for fabric smoothing. LazyDagger achieves higher final coverage than Behavior Cloning and SafeDagger with fewer context switches than SafeDagger but more supervisor actions. At execution time, we again observe that LazyDagger achieves similar coverage as SafeDagger but with fewer context switches.

L^* value, which we refer to as the *cutoff latency*, for all experiments to precisely study the types of domains in which LazyDagger is most applicable.

Baselines: We compare LazyDagger to Behavior Cloning [49], DAgger [40], and SafeDagger [53] in terms of the total supervisor burden and task performance. The Behavior Cloning and DAgger comparisons evaluate the utility of human interventions, while the comparison to SafeDagger, another interactive IL algorithm, evaluates the impact of soliciting fewer but longer interventions.

Experimental Setup: For all MuJoCo environments, we use a reinforcement learning agent trained with TD3 [16] as an algorithmic supervisor. We begin all LazyDagger, SafeDagger, and DAgger experiments by pre-training the robot policy with Behavior Cloning on 4,000 state-action pairs for 5 epochs, and similarly report results for Behavior Cloning after the 5th epoch. To ensure a fair comparison, Behavior Cloning uses additional offline data equal to the average amount of online data seen by LazyDagger during training. All results are averaged over 3 random seeds.

Results: In Figure 3, we study the performance of LazyDagger and baselines. After every epoch of training, we run the policy for 10 test rollouts *where interventions are not allowed* and report the task reward on these rollouts in Figure 3. Results suggest that LazyDagger is able to match or outperform all baselines in terms of task performance across all simulation environments (Figure 3A). Additionally, LazyDagger requires far fewer context switches compared to SafeDagger (Figure 3D), while requesting a similar number of supervisor actions across domains (Figure 3C): we observe a 79%, 56%, and 46% reduction in context switches on the HalfCheetah, Walker2D, and Ant environments respectively. LazyDagger and SafeDagger both use an order of magnitude fewer supervisor actions than DAgger. While SafeDagger requests much fewer supervisor actions than LazyDagger in the Ant environment, this limited amount of supervision is insufficient to match the task performance of LazyDagger or any of the baselines, suggesting that SafeDagger may be terminating interventions prematurely. We study the total supervisor burden of SafeDagger and LazyDagger as defined in Equation (III.2) and find that in HalfCheetah, Walker2D, and Ant, the cutoff latencies L^* are

0.0, 4.3, and 7.6 respectively, i.e. LazyDagger achieves lower supervisor burden in the HalfCheetah domain for any L as well as lower burden in Walker2D and Ant for $L > 4.3$ and $L > 7.6$ respectively. The results suggest that LazyDagger can reduce total supervisor burden compared to SafeDagger even for modest latency values, but that SafeDagger may be a better option for settings with extremely low latency.

Ablations: We study 2 key ablations for LazyDagger in simulation: (1) returning to autonomous mode with $f(\cdot)$ rather than using the ground truth discrepancy (LazyDagger (-Switch to Auto) in Figure 3), and (2) removal of noise injection (LazyDagger (-Noise)). LazyDagger outperforms both ablations on all tasks, with the exception of ablation 1 on Walker2D, which performed similarly well. We also observe that LazyDagger consistently requests more supervisor actions than either ablation. This aligns with the intuition that both using the ground truth action discrepancy to switch back to autonomous mode and injecting noise result in longer but more useful interventions that improve performance.

B. Fabric Smoothing in Simulation

Environment: We evaluate LazyDagger on the fabric smoothing task from [44] (shown in Figure 4) using the simulation environment from [44]. The task requires smoothing an initially crumpled fabric and is challenging due to the infinite-dimensional state space and complex dynamics, motivating learning from human feedback. As in prior work [44], we utilize top-down $100 \times 100 \times 3$ RGB image observations of the workspace and use actions which consist of a 2D pick point and a 2D pull vector. See [44] for further details on the fabric simulator.

Experimental Setup: We train a fabric smoothing policy in simulation using DAgger under supervision from an analytic corner-pulling policy that leverages the simulator’s state to identify fabric corners, iterate through them, and pull them towards the corners of the workspace [44]. We transfer the resulting policy for a 16×16 grid of fabric into a new simulation environment with altered fabric dynamics (i.e. lower spring constant, altered fabric colors, and a higher-fidelity 25×25 discretization) and evaluate LazyDagger and baselines on how rapidly they can adapt the initial policy to the new domain. As in [44], we terminate rollouts when

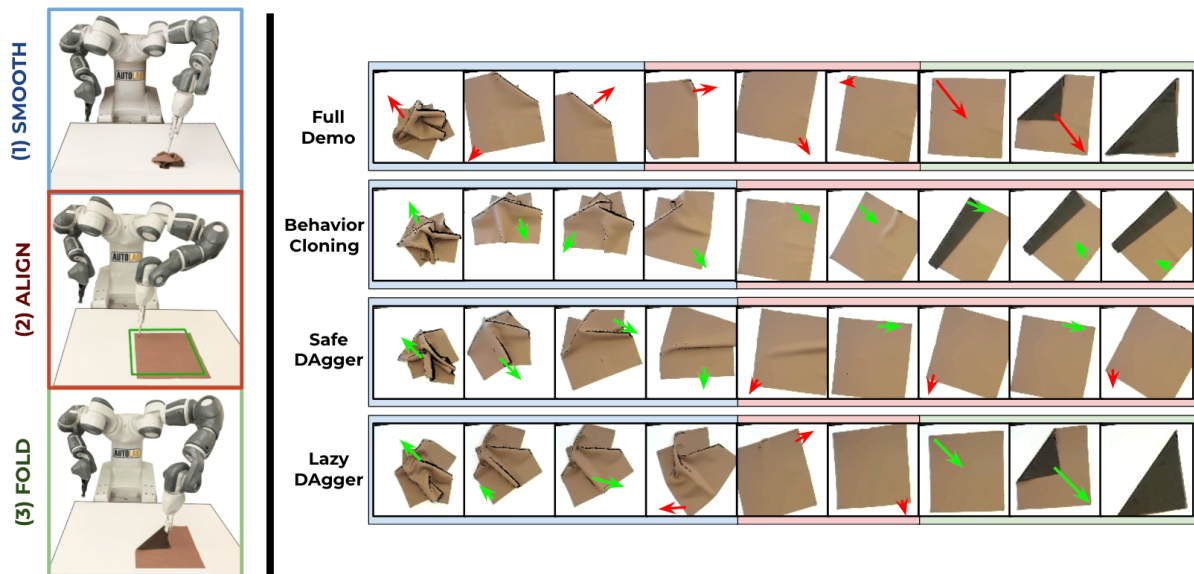


Fig. 5: **Physical Fabric Manipulation Task:** *Left:* We evaluate on a 3-stage fabric manipulation task consisting of smoothing a crumpled fabric, aligning the fabric so all corners are visible in the observations, and performing a triangular fold. *Right:* Rollouts of the fabric manipulation task, where each frame is a $100 \times 100 \times 3$ overhead image. Human supervisor actions are denoted in red while autonomous robot actions are in green. Rollouts are shaded to indicate task progress: blue for smoothing, red for alignment, and green for folding. SafeDAgger ends human intervention prematurely, resulting in poor task performance and more context switches, while LazyDAgger switches back to robot control only when confident in task completion.

Algorithm	Task Successes	Task Progress			Context Switches	Supervisor Actions	Robot Actions	Failure Modes			
		(1)	(2)	(3)				A	B	C	D
Behavior Cloning	0/10	6/10	0/10	0/10	N/A	N/A	119	2	1	7	0
SD-Execution	2/10	6/10	4/10	2/10	53	34	108	5	0	0	3
LD-Execution	8/10	10/10	10/10	8/10	21	43	47	0	0	0	2

TABLE I: **Physical Fabric Manipulation Experiments:** We evaluate LazyDAgger-Execution and baselines on a physical 3-stage fabric manipulation task and report the success rate and supervisor burden in terms of total supervisor actions and bidirectional context switches (summed across all 10 trials). Task Progress indicates how many trials completed each of the 3 stages: Smoothing, Aligning, and Folding. LazyDAgger-Execution achieves more successes with fewer context switches ($L^* = 0.28$). We observe the following failure modes (Table I): (A) action limit hit (> 15 total actions), (B) fabric is more than 50% out of bounds, (C) incorrect predicted pick point, and (D) the policy failed to request an intervention despite high ground truth action discrepancy.

we exceed 10 time steps, 92% coverage, or have moved the fabric more than 20% out of bounds. We evaluate performance based on a coverage metric, which measures the percentage of the background plane that the fabric covers (fully smooth corresponds to a coverage of 100).

Results: We report results for the fabric smoothing simulation experiments in Figure 4. Figure 4 (A) shows the performance of the SafeDAgger and LazyDAgger policies during learning. To generate this plot we periodically evaluated each policy on *test* rollouts without interventions. Figure 4 (B) and (C) show the number of supervisor actions and context switches required during learning; LazyDAgger performs fewer context switches than SafeDAgger but requires more supervisor actions as the interventions are longer. Results suggest that the cutoff latency (as defined in Section VI-A) is $L^* = 1.5$ for fabric smoothing. Despite fewer context switches, LazyDAgger achieves comparable performance to SafeDAgger, suggesting that LazyDAgger can learn complex, high-dimensional robotic control policies while reducing the number of hand-offs to a supervisor. We also evaluate LazyDAgger-Execution and SafeDAgger-Execution, in which interventions are allowed but the policy is no longer updated (see Section VI-C). We see that in this case, LazyDAgger achieves similar final coverage as

SafeDAgger with significantly fewer context switches.

C. Physical Fabric Manipulation Experiments

Environment: In physical experiments, we evaluate on a multi-stage fabric manipulation task with an ABB YuMi robot and a human supervisor (Figure 5). Starting from a crumpled initial fabric state, the task consists of 3 stages: (1) fully smooth the fabric, (2) align the fabric corners with a tight crop of the workspace, and (3) fold the fabric into a triangular fold. Stage (2) in particular requires high precision, motivating human interventions. As in the fabric simulation experiments, we use top-down $100 \times 100 \times 3$ RGB image observations of the workspace and have 4D actions consisting of a pick point and pull vector. The actions are converted to workspace coordinates with a standard calibration procedure and analytically mapped to the nearest point on the fabric. Human supervisor actions are provided through a point-and-click interface for specifying pick-and-place actions. See the supplement for further details.

Experimental Setup: Here we study how interventions can be leveraged to improve the final task performance even at *execution time*, in which policies are no longer being updated. We collect 20 offline task demonstrations and train an initial policy with behavior cloning. To prevent

overfitting to a small amount of real data, we use standard data augmentation techniques such as rotating, scaling, changing brightness, and adding noise to create 10 times as many training examples. We then evaluate the behavior cloning agent (Behavior Cloning) and agents which use the SafeDagger and LazyDagger intervention criteria but do not update the policy with new experience or inject noise (SafeDagger-Execution and LazyDagger-Execution respectively). We terminate rollouts if the fabric has successfully reached the goal state of the final stage (i.e. forms a perfect or near-perfect dark brown right triangle as in Hoque *et al.* [18]; see Figure 5), more than 50% of the fabric mask is out of view in the current observation, the predicted pick point misses the fabric mask by approximately 50% of the plane or more, or 15 total actions have been executed (either autonomous or supervisor).

Results: We perform 10 physical trials of each technique. In Table I, we report both the overall task success rate and success rates for each of the three stages of the task: (1) Smoothing, (2) Alignment, and (3) Folding. We also report the total number of context switches, supervisor actions, and autonomous robot actions summed across all 10 trials for each algorithm (Behavior Cloning, SafeDagger-Execution, LazyDagger-Execution). In Figure 5 we provide representative rollouts for each algorithm. Results suggest that Behavior Cloning is insufficient for successfully completing the alignment stage with the required level of precision. SafeDagger-Execution does not improve the task success rate significantly due to its inability to collect interventions long enough to navigate bottleneck regions in the task (Figure 5). LazyDagger-Execution, however, achieves a much higher success rate than SafeDagger-Execution and Behavior Cloning with far fewer context switches than SafeDagger-Execution: LazyDagger-Execution requests 2.1 context switches on average per trial (i.e. 1.05 interventions) as opposed to 5.3 switches (i.e. 2.65 interventions). LazyDagger-Execution trials also make far more task progress than the baselines, as all 10 trials reach the folding stage. LazyDagger-Execution does request more supervisor actions than SafeDagger-Execution, as in the simulation environments. LazyDagger-Execution also requests more supervisor actions relative to the total amount of actions due to the more conservative switching criteria and the fact that successful episodes are shorter than unsuccessful episodes on average. Nevertheless, results suggest that for this task, LazyDagger-Execution reduces supervisor burden for any $L > L^* = 0.28$, a very low cutoff latency that includes all settings in which a context switch is at least as time-consuming as an individual action (i.e. $L \geq 1$).

In experiments, we find that SafeDagger-Execution’s short interventions lead to many instances of Failure Mode A (see Table I), as the policy is making task progress, but not quickly enough to perform the task. We observe that Failure Mode C is often due to the fabric reaching a highly irregular configuration that is not within the training data distribution, making it difficult for the robot policy to make progress. We find that SafeDagger and LazyDagger experience Failure

Mode D at a similar rate as they use the same criteria to solicit interventions (but different termination criteria). However, we find that all of LazyDagger’s failures are due to Failure Mode D, while SafeDagger also fails in Mode A due to premature termination of interventions.

VII. DISCUSSION AND FUTURE WORK

We propose context switching between robot and human control as a metric for supervisor burden in interactive imitation learning and present LazyDagger, an algorithm which can be used to efficiently learn tasks while reducing this switching. We evaluate LazyDagger on 3 continuous control benchmark environments in MuJoCo, a fabric smoothing environment in simulation, and a fabric manipulation task with an ABB YuMi robot and find that LazyDagger is able to improve task performance while reducing context switching between the learner and robot by up to 79% over SafeDagger. In subsequent work, we investigate more intervention criteria and apply robot-gated interventions to controlling a fleet of robots, where context switching can negatively impact task throughput.

VIII. ACKNOWLEDGMENTS

This research was performed at the AUTOLAB at UC Berkeley in affiliation with the Berkeley AI Research (BAIR) Lab, and the CITRIS “People and Robots” (CPAR) Initiative. This research was supported in part by the Scalable Collaborative Human-Robot Learning (SCHooL) Project, NSF National Robotics Initiative Award 1734633. The authors were supported in part by donations from Google, Siemens, Toyota Research Institute, Autodesk, Honda, Intel, and Hewlett-Packard and by equipment grants from PhotoNeo, Nvidia, and Intuitive Surgical. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the sponsors. We thank colleagues Lawrence Chen, Jennifer Grannen, and Vincent Lim for providing helpful feedback and suggestions.

REFERENCES

- [1] P. Abbeel and A. Y. Ng, “Apprenticeship learning via inverse reinforcement learning”, in *Proceedings of the twenty-first international conference on Machine learning*, 2004, p. 1.
- [2] O. Amir, E. Kamar, A. Kolobov, and B. J. Grosz, “Interactive teaching strategies for agent training”, in *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.
- [3] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, “A survey of robot learning from demonstration”, *Robotics and autonomous systems*, vol. 57, no. 5, pp. 469–483, 2009.
- [4] S. Arora and P. Doshi, “A survey of inverse reinforcement learning: Challenges, methods and progress”, *arXiv preprint arXiv:1806.06877*, 2018.
- [5] A. Balakrishna*, B. Thananjeyan*, J. Lee, F. Li, A. Zahed, J. E. Gonzalez, and K. Goldberg, “On-policy robot imitation learning from a converging supervisor”, in *Conference on Robot Learning (CoRL)*, PMLR, 2019.
- [6] E. Båvenstrand and J. Berggren, “Performance evaluation of imitation learning algorithms with human experts”, in *Technical Report KTH Royal Institute of Technology Sweden*, 2019.
- [7] E. Bıyık, M. Palan, N. C. Landolfi, D. P. Losey, and D. Sadigh, “Asking easy questions: A user-friendly approach to active reward learning”, *arXiv preprint arXiv:1910.04365*, 2019.
- [8] B. K. Bose, “An adaptive hysteresis-band current control technique of a voltage-fed pwm inverter for machine drive system”, *IEEE Transactions on Industrial Electronics*, vol. 37, no. 5, 1990.
- [9] D. S. Brown, W. Goo, and S. Niekum, “Better-than-demonstrator imitation learning via automatically-ranked demonstrations”, in *Conference on Robot Learning (CoRL)*, 2019.
- [10] D. S. Brown, S. Niekum, R. Coleman, and R. Srinivasan, “Safe imitation learning via fast bayesian reward inference from preferences”, in *International Conference on Machine Learning*, 2020.

- [11] D. S. Brown, W. Goo, N. Prabhat, and S. Niekum, "Extrapolating beyond suboptimal demonstrations via inverse reinforcement learning from observations", in *Proceedings of the 36th International Conference on Machine Learning, ICML 2019*, 2019.
- [12] J. Y. Chen and M. J. Barnes, "Human-agent teaming for multirobot control: A review of human factors issues", *IEEE Transactions on Human-Machine Systems*, vol. 44, no. 1, pp. 13–29, 2014.
- [13] P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei, "Deep reinforcement learning from human preferences", in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [14] J. W. Crandall, M. A. Goodrich, D. R. Olsen, and C. W. Nielsen, "Validating human-robot interaction schemes in multitasking environments", *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 35, no. 4, pp. 438–449, 2005.
- [15] J. Fu, K. Luo, and S. Levine, "Learning robust rewards with adversarial inverse reinforcement learning", *arXiv preprint arXiv:1710.11248*, 2017.
- [16] S. Fujimoto, H. van Hoof, and D. Meger, "Addressing Function Approximation Error in Actor-Critic Methods", in *Proc. Int. Conf. on Machine Learning*, 2018.
- [17] J. Ho and S. Ermon, "Generative adversarial imitation learning", in *Advances in Neural Information Processing Systems*, 2016.
- [18] R. Hoque, D. Seita, A. Balakrishna, A. Ganapathi, A. Tanwani, N. Jamali, K. Yamane, S. Iba, and K. Goldberg, "VisuoSpatial Foresight for Multi-Step, Multi-Task Fabric Manipulation", in *Proc. Robotics: Science and Systems (RSS)*, 2020.
- [19] B. Ibarz, J. Leike, T. Pohlen, G. Irving, S. Legg, and D. Amodei, "Reward learning from human preferences and demonstrations in atari", in *Advances in Neural Information Processing Systems*, 2018.
- [20] A. J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal, "Dynamical movement primitives: Learning attractor models for motor behaviors", *Neural computation*, vol. 25, no. 2, 2013.
- [21] S. Jauhari, C. Celemin, and J. Kober, "Interactive imitation learning in state-space", *arXiv preprint arXiv:2008.00524*, 2020.
- [22] K. Judah, A. Fern, and T. Dietterich, "Active imitation learning via state queries", in *Proceedings of the icml workshop on combining learning strategies to reduce label cost*, Citeseer, 2011.
- [23] G. Kahn, P. Abbeel, and S. Levine, "LaND: Learning to Navigate from Disengagements", in *arXiv preprint arXiv:2010.04689*, 2020.
- [24] M. Kelly, C. Sidrane, K. Driggs-Campbell, and M. J. Kochenderfer, "Hg-dagger: Interactive imitation learning with human experts", in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2019.
- [25] T. Kessler Faulkner, R. A. Gutierrez, E. S. Short, G. Hoffman, and A. L. Thomaz, "Active attention-modified policy shaping: Socially interactive agents track", in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 2019.
- [26] A. Kurenkov, A. Mandlekar, R. Martin-Martin, S. Savarese, and A. Garg, "Ac-teach: A bayesian actor-critic method for policy learning with an ensemble of suboptimal teachers", in *Conf. on Robot Learning (CoRL)*, 2019.
- [27] M. Laskey, J. Lee, R. Fox, A. Dragan, and K. Goldberg, "Dart: Noise injection for robust imitation learning", in *Conf. on Robot Learning (CoRL)*, 2017.
- [28] M. Laskey, S. Staszak, W. Hsieh, J. Mahler, F. Pokorný, A. Dragan, and K. Goldberg, "Shiv: Reducing supervisor burden using support vectors for efficient learning from demonstrations in high dimensional state spaces", in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2016.
- [29] J. Launchbury, "A natural semantics for lazy evaluation", ser. POPL '93, Charleston, South Carolina, USA: Association for Computing Machinery, 1993, pp. 144–154.
- [30] A. Mandlekar, D. Xu, R. Martín-Martín, Y. Zhu, L. Fei-Fei, and S. Savarese, *Human-in-the-loop imitation learning using remote teleoperation*, 2020. arXiv: 2012.06733 [cs.LG].
- [31] K. Menda, K. Driggs-Campbell, and M. J. Kochenderfer, "EnsembleDagger: A Bayesian Approach to Safe Imitation Learning", in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2019.
- [32] C.-A. C. Nolan Wagener Byron Boots, "Safe reinforcement learning using advantage-based intervention", in *International Conference on Machine Learning (ICML)*, 2021.
- [33] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, and J. Peters, "An algorithmic perspective on imitation learning", *arXiv preprint arXiv:1811.06711*, 2018.
- [34] M. Palan, N. C. Landolfi, G. Shevchuk, and D. Sadigh, "Learning reward functions by integrating human demonstrations and preferences", in *Proceedings of Robotics: Science and Systems (RSS)*, 2019.
- [35] A. Paraschos, C. Daniel, J. R. Peters, and G. Neumann, "Probabilistic movement primitives", in *Advances in neural information processing systems*, 2013, pp. 2616–2624.
- [36] D. A. Pomerleau, "Efficient training of artificial neural networks for autonomous navigation", *Neural Computation*, vol. 3, no. 1, 1991.
- [37] M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [38] S. Reddy, A. D. Dragan, and S. Levine, "Shared autonomy via deep reinforcement learning", *Proc. Robotics: Science and Systems (RSS)*, 2018.
- [39] —, "Where do you think you're going?: Inferring beliefs about dynamics from behavior", in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [40] S. Ross, G. J. Gordon, and J. A. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning", in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- [41] D. Sadigh, A. D. Dragan, S. S. Sastry, and S. A. Seshia, "Active preference-based learning of reward functions", in *Proceedings of Robotics: Science and Systems (RSS)*, 2017.
- [42] A. Saran, E. S. Short, A. Thomaz, and S. Niekum, "Understanding teacher gaze patterns for robot learning", L. P. Kaelbling, D. Kragic, and K. Sugiura, Eds., ser. Proceedings of Machine Learning Research, vol. 100, PMLR, Oct. 2020, pp. 1247–1258.
- [43] W. Saunders, G. Sastry, A. Stuhlmüller, and O. Evans, "Trial without error: Towards safe rl with human intervention", *17th International Conference on Autonomous Agents and MultiAgent Systems*, 2018.
- [44] D. Seita, A. Ganapathi, R. Hoque, M. Hwang, E. Cen, A. K. Tanwani, A. Balakrishna, B. Thananjeyan, J. Ichnowski, N. Jamali, K. Yamane, S. Iba, J. Canny, and K. Goldberg, "Deep Imitation Learning of Sequential Fabric Smoothing From an Algorithmic Supervisor", in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2020.
- [45] J. Spencer, S. Choudhury, M. Barnes, M. Schmittle, M. Chiang, P. Ramadge, and S. Srinivasa, "Learning from interventions: Human-robot interaction as both explicit and implicit feedback", in *Proc. Robotics: Science and Systems (RSS)*, 2020.
- [46] G. Swamy, S. Reddy, S. Levine, and A. D. Dragan, "Scaled autonomy: Enabling human operators to control robot fleets", in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2020, pp. 5942–5948.
- [47] B. Thananjeyan*, A. Balakrishna*, S. Nair, M. Luo, K. Srinivasan, M. Hwang, J. E. Gonzalez, J. Ibarz, C. Finn, and K. Goldberg, "Recovery rl: Safe reinforcement learning with learned recovery zones", in *Robotics and Automation Letters (RAL)*, IEEE, 2021.
- [48] E. Todorov, T. Erez, and Y. Tassa, "MuJoCo: A Physics Engine for Model-Based Control", in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2012.
- [49] F. Torabi, G. Warnell, and P. Stone, "Behavioral cloning from observation", in *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, Stockholm, Sweden, Jul. 2018.
- [50] H.-Y. Tung, A. W. Harley, L.-K. Huang, and K. Fragkiadaki, "Reward learning from narrated demonstrations", in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7004–7013.
- [51] F. Wang, B. Zhou, K. Chen, T. Fan, X. Zhang, J. Li, H. Tian, and J. Pan, "Intervention aided reinforcement learning for safe and practical policy optimization in navigation", in *Conf. on Robot Learning (CoRL)*, 2018.
- [52] L. Xie, S. Wang, S. Rosa, A. Markham, and N. Trigoni, "Learning with training wheels: Speeding up training with a simple controller for deep reinforcement learning", in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2018.
- [53] J. Zhang and K. Cho, "Query-efficient imitation learning for end-to-end autonomous driving", in *Association for the Advancement of Artificial Intelligence (AAAI)*, 2017.
- [54] 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.