

MILE: Model-based Intervention Learning

Yigit Korkmaz, Erdem Bıyık

Abstract—Imitation learning techniques have been shown to be highly effective in real-world control scenarios, such as robotics. However, these approaches not only suffer from compounding error issues but also require human experts to provide complete trajectories. Although there exist interactive methods where an expert oversees the robot and intervenes if needed, these extensions usually only utilize the data collected during intervention periods and ignore the feedback signal hidden in non-intervention timesteps. In this work, we create a model to formulate how the interventions occur in such cases, and show that it is possible to learn a policy with just a handful of expert interventions. Our key insight is that it is possible to get crucial information about the quality of the current state and the optimality of the chosen action from expert feedback, regardless of the presence or the absence of intervention. We evaluate our method on various discrete and continuous simulation environments, a real-world robotic manipulation task, as well as a human subject study. Videos and the code can be found at <https://liralab.usc.edu/mile>.

I. INTRODUCTION

Imagine training a household robot to help users place the dishes in the dishwasher. One way to do this is to use reinforcement learning (RL) that has been proven successful in several areas ranging from gaming to dialogue systems and autonomous driving [1]–[4]. However, its need for lots of online interactions with the environment as well as a well-defined reward function make it unsuitable in a real-world situation like this. An alternative is to use imitation learning (IL) where an expert provides demonstrations of how to place the dishes. IL requires fewer interactions in the world than RL and does not require a reward function. A common drawback of this approach is the compounding distributional shift, which results from the accumulation of errors when deploying a learned policy [5]: the small inaccuracies in the robot’s learned policy will move it to an out-of-distribution state where the policy may fail more significantly, which may result in the robot breaking the dishes.

Interactive learning methods try to overcome this compounding errors problem by iteratively querying the expert with system states, and fine-tuning the policy based on the expert actions [6]–[11]. Most of these methods do not let the human intervene at will but transfer the control to the human according to some criterion [8], [9]. In others, the human can take over the control at any timestep [7], [12]–[14].

Going back to the running example, assume the robot had a mediocre policy in the beginning, thanks to some initial training done in the factory setup. In this interactive scheme, we would control the robot only when we think it is doing or about to do something wrong. This is clearly

more convenient than operating the robot for long periods to generate full demonstrations. But what about the intervals where we did *not* input any actions? Did we not provide any information to the robot? The answer is no: the fact that we *chose* to not intervene means the robot’s actions were already good enough in those intervals. This is what the existing interactive learning techniques are missing: even though they try to improve the robot’s policy based on the states where the expert intervenes, they do not utilize any structure of how or when those interventions occur, ignoring an important feedback signal that is leaking through the states where the expert does not intervene.

To efficiently use the information in both the states with and without human interventions, we argue one must understand and utilize the structure behind how interventions occur. To this end, we make the following contributions:

- 1) We propose a novel model that is fully differentiable to formulate how and when an expert intervenes.
- 2) We utilize this intervention model to fine-tune a weak policy and evaluate our method in various simulated and real-world environments to prove its effectiveness.
- 3) We compare our method against the state-of-the-art baselines which utilize interventions to show its higher sample-efficiency and performance.

II. RELATED WORK

Interactive Imitation Learning. In imitation learning, expert data is used to train a policy in a supervised way [15]–[20]. Interactive imitation learning methods try to overcome its compounding errors problem [5], [6] by querying the expert on the learned policy’s rollouts [6]–[9], [15], [21]. While the expert relabels either the whole trajectory of the policy with actions, or some parts of it that are automatically selected by estimating various task performance quantities [6], [9], there are also works where the expert has the freedom to intervene at will [7], [12], [13], [22]. Some of these methods consider the implicit feedback coming from the states where the expert chooses to not intervene. They attempt to incorporate the information leaking from non-intervention intervals either by enforcing those state-action pairs to be constrained in the action-value cost functions [14] or utilize weighted behavioral cloning (BC) to incorporate the signal from non-interventions, with different heuristics used for assigning weights to on-policy robot samples, and expert human interventions [12], [13], [22]. However, none of these algorithms uses a model to understand and learn from why the expert chooses to not intervene. In our method, we propose a model that attempts to capture how the interventions occur, and how satisfied the user is with

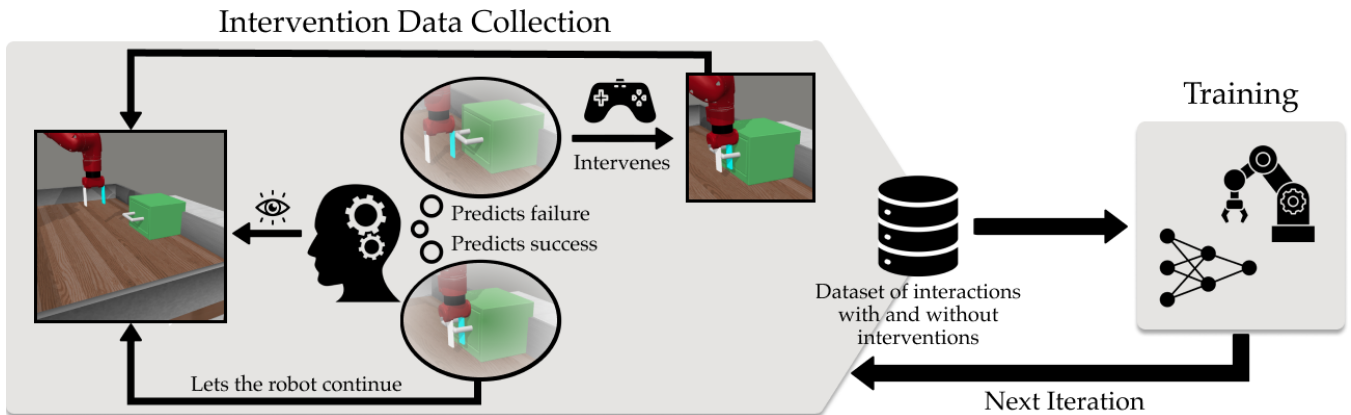


Fig. 1. **Overall MILE System.** A human operator oversees the robot during task execution and may decide to take over control at any timestep. The human makes the decision to intervene or not based on their prediction of the robot’s potential failure, without observing the robot’s action in that particular state. During the data collection phase, all interactions with the environment are recorded, both with and without interventions. The policy is then trained on this dataset using an iterative process that incorporates our novel intervention model.

the current action in the case of not intervening. Moreover, we investigate how we can efficiently use the information coming from that model for training. To our knowledge, this is the first work that models how human interventions happen in robotics while utilizing it for better policy learning.

Reinforcement Learning with Expert Data. Many studies focus on integrating expert demonstrations into RL. This is often achieved by populating the replay buffer to warm-start the initial policy and guide its exploration in a favorable direction [23]–[26]. There are also approaches that incorporate interventions into RL. One approach introduces an extra cost term in the learning objective to minimize interventions [27]. Another recent method relabels the rewards based on interventions, assigning negative rewards at intervention states, while giving a zero reward to all other states [28]. However, most of these methods assume access to some type of reward information, either through explicit reward functions or through Q-functions of the task. Our method does not require these assumptions. Not to mention, these methods still require a lot of online interactions with the environment and RL algorithms that can handle both on-policy and off-policy data coming from the intervening agent.

Modeling Interventions. There have been efforts to provide formal definitions and qualitative analyses of interventions in the human-robot interaction research [29], along with developing metrics to evaluate operator workload [30], [31]. Within computational modeling, various studies have addressed the modeling of trust in human-robot collaborative settings, which include both binary and continuous measures [32]. There are also works exploring the temporal dynamics of trust conditioned on task performance [33], [34]. One such work models trust as a latent variable in a partially observable Markov decision process, utilizing this model to maximize task performance [34]. To our knowledge, no prior work has developed a similar computational model to explain the occurrence of interventions in human-robot teams.

III. PROBLEM DEFINITION

We formulate the problem as a discrete time Markov decision process (MDP) with the standard $\langle S, A, \rho, f, R, T, \gamma \rangle$

notation. The robot does not know the reward function R or the transition dynamics f of the environment. It starts with an initial policy $\pi_\theta : S \rightarrow A$ parameterized by θ .

In addition, a potentially noisy expert human can intervene (take over the control) with any action $a_h \in A$ at any timestep. However, in contrast to some of the existing works [35], the human has to make the decision about intervening or not without observing the robot’s action $a_r \in A$ in that particular state. This is a more viable setting in robotics where it is not realistic to expect the robot to check each and every action with the human.

We let the robot and the human interact in this setup, and we record (s, a_r, a_h, s') from every timestep, with the possibility that a_h is undefined if the human did not intervene in that timestep. Our objective is to find a policy π_{θ^*} that maximizes the expected cumulative discounted reward:

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_{t=0}^{T-1} \gamma^t R(s_t, a_t) \right]$$

where $s_0 \sim \rho(\cdot)$, and (s_t, a_t) pairs are sampled according to the policy π_θ and transition function f .

We assume the human interacts with the robot for N iterations, each of which includes k episodes. After each iteration, we train the robot policy using the data collected so far. In the next section, we describe how we do this training.

IV. MODEL-BASED INTERVENTION LEARNING (MILE)

Our method to solve this problem is based on a computational model of when and how the human may decide to intervene the robot’s operation. We will introduce this intervention model in the first subsection. Subsequently, we will leverage this model to create a framework where we update the robot’s policy π_θ based on its interactions with the human and the environment.

Intervention Model. Consider our running example: would you intervene the robot’s operation when it is quickly moving toward a glass jar that stands at the edge of the countertop? How about if the jar was made of plastic? This example highlights the important factors that affect when humans intervene the robot. Firstly, consider the case of glass jar. The human is likely to intervene because they predict the

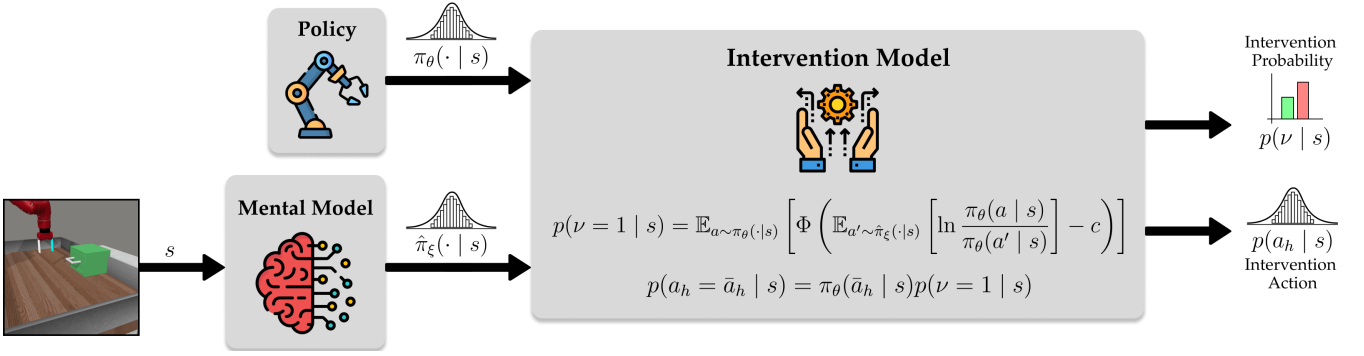


Fig. 2. **Framework for learning from interventions.** Starting from an initial policy π_θ , we jointly train the mental model, and the policy using our intervention model.

robot may hit and break the jar. So the human’s belief about what the robot will do, i.e., their mental model of the robot, affects their decision to intervene, as this was previously shown in human-robot interaction research [34].

Next, let’s compare the cases with glass and plastic jars. The human is less likely to intervene in the latter, because even if the robot takes bad actions, the outcome is not going to be as bad as in the case of glass jar. This points out that the human’s interventions also depend on how good/bad the potential outcomes are.

Considering these factors, we propose an intervention model based on the probit model from discrete decision theory [36]. For the ease of the presentation, we will start with discrete action spaces and then extend our formulation to continuous domains.

Let ν be a binary random variable that indicates whether the human intervenes ($\nu = 1$) or not ($\nu = 0$), and \bar{a}_h denote the nominal human action, i.e., the action human would take provided they decide to intervene. Mathematically, $a_h = \bar{a}_h$ if and only if $\nu = 1$. Otherwise, a_h is not defined in that state. Finally, let $\hat{\pi}$ denote the human’s mental model of the robot, i.e., what the human believes the robot will do in a given state. This prediction is needed as the human has to intervene before seeing the robot’s action. While some learning from intervention works assumes the human first sees the system’s proposed action [35], ours is a more realistic setting, especially in robotics where conveying an action without taking it is both difficult and time-consuming.

We start with modeling the probability that the human will intervene at a given state s . We use σ for the softmax operation and Φ for the cdf of a standard normal distribution.

$$\begin{aligned}
 p(\nu = 1 | s) &= \sum_{a \in A} p(\bar{a}_h = a, \nu = 1 | s) \\
 &= \sum_{a \in A} p(\bar{a}_h = a | s) p(\nu = 1 | \bar{a}_h = a, s) \quad (1)
 \end{aligned}$$

Here, the first term inside the summation is just the probability that the human would take action a at state s . Since we assume the human is a (noisy) expert, we use a Boltzmann policy under the true reward function of the environment to model this probability, similar to the prior work [37]–[41]:

$$p(\bar{a}_h = a | s) = \pi_h(a | s) := \frac{\exp(Q(s, a))}{\sum_{a' \in A} \exp(Q(s, a'))}$$

The second term in Eq. (1) denotes the probability that

the human will intervene conditioned on the state and their nominal action. This is where we bring the probit model into play: the human will intervene only if their nominal action is considerably better than what they expect the robot to do:

$$p(\nu = 1 | \bar{a}_h = a, s) = \Phi \left(Q(s, a) - \mathbb{E}_{a' \sim \hat{\pi}(\cdot | s)} [Q(s, a')] - c \right),$$

where c is a scalar hyperparameter. In this probit-based model, the human is more likely to intervene if the value of their nominal action is much higher than the expected value of the robot’s action. The scalar c depends on the effort the human needs to put to intervene the robot. If it is difficult to intervene, c is going to be high and the human will only intervene if the difference in action values is extremely high. On the other extreme, if interventions are free, then c will be low and the human may continually intervene.

Now that we modeled the probability of human interventions, we continue with *how* they intervene when they do:

$$\begin{aligned}
 p(a_h = \bar{a}_h | s) &= p(a_h = \bar{a}_h | \nu = 0, s) p(\nu = 0 | s) \\
 &\quad + p(a_h = \bar{a}_h | \nu = 1, s) p(\nu = 1 | s) \\
 &= \sigma(Q(s, \bar{a}_h)) p(\nu = 1 | s),
 \end{aligned}$$

where $p(a_h = \bar{a}_h | \nu = 0, s) = 0$ by the definition of a_h .

We note the following relation due to the Boltzmann policy formulation we use for the human:

$$\begin{aligned}
 \ln \pi_h(a | s) - \ln \pi_h(a' | s) &= \ln \left(\frac{\exp Q(s, a)}{\sum_{a'' \in A} \exp Q(s, a'')} \right) \\
 &\quad - \ln \left(\frac{\exp Q(s, a')}{\sum_{a'' \in A} \exp Q(s, a'')} \right) \\
 &= Q(s, a) - Q(s, a').
 \end{aligned}$$

This allows us to rewrite $p(\nu = 1 | s)$ and $p(a_h = \bar{a}_h | s)$ only dependent on the policy and not the Q -function:

$$\begin{aligned}
 p(\nu = 1 | s) &= \sum_{a \in A} \pi_h(a | s) \Phi \left(\mathbb{E}_{a' \sim \hat{\pi}(\cdot | s)} [\ln \pi_h(a | s) - \ln \pi_h(a' | s)] - c \right) \\
 &= \mathbb{E}_{a \sim \pi_h(\cdot | s)} \left[\Phi \left(\mathbb{E}_{a' \sim \hat{\pi}(\cdot | s)} [\ln \pi_h(a | s) - \ln \pi_h(a' | s)] - c \right) \right] \\
 p(a_h = \bar{a}_h | s) &= \pi_h(\bar{a}_h | s) p(\nu = 1 | s) \quad (2)
 \end{aligned}$$

This change of variables allows us to readily use this intervention model in continuous domains. The main difference is that, in the discrete case, a_h has a categorical distribution over the action space as well as the no-intervention ($\nu = 0$). In continuous case, ν acts like a gate controlling which continuous policy (π_h or π_θ) will be active. We will now

present how we use this intervention model to learn from humans’ intervention feedback.

Learning from Interventions. The robot misses information about two critical components in the intervention model: $\hat{\pi}$ that models what the human thinks the robot will do, and the human’s policy π_h . In our learning algorithm, we model both of these policies with neural networks, $\hat{\pi}_\xi$ and π_θ , respectively (see “mental model” and “policy” in Fig. 2). Since the intervention model is differentiable, we conveniently utilize the gradients coming from it to jointly train these networks using the dataset of (s, a_r, a_h, s') tuples. During inference time, we only employ the trained policy π_θ .

In discrete domains, we minimize the cross entropy loss between the ground truth actions (a_h if $\nu = 1$ and a_r otherwise) in the dataset and estimated final action probabilities, where no-intervention ($\nu = 0$) is an extra class in the action space. Mathematically, the loss function is:

$$J(\theta, \xi) = \text{CE}(\hat{\mathbf{P}}(\mathbf{S}; \theta, \xi), \mathbf{a}_h) = -\frac{1}{N} \sum_{i=1}^N \log(\hat{p}_{i, a_h^i}(\mathbf{s}_i; \theta, \xi)),$$

where \mathbf{S} is the set of states in the dataset, $\hat{\mathbf{P}}(\mathbf{S}; \theta, \xi)$ is the matrix of probabilities with each row $\hat{\mathbf{p}}_i(\mathbf{s}_i; \theta, \xi)$ being the predicted probability vector for sample i , \mathbf{a}_h is the vector of true intervention action labels including no-intervention, and $\hat{p}_{i, a_h^i}(\mathbf{s}_i; \theta, \xi)$ is the predicted probability for the true class of sample i .

In continuous domains, we use a combination of discrete and continuous loss functions as we want our model to learn both when and how the human intervenes. In order to accomplish the former, we use binary cross entropy loss between ground truth intervention signals ν and the estimated probability of intervention from Eq. (2), where π_θ replaces π_h and $\hat{\pi}_\xi$ replaces $\hat{\pi}$:

$$J_1(\theta, \xi) = \text{BCE}(\hat{\nu}(\mathbf{S}; \theta, \xi), \nu) \\ = -\frac{1}{N} \sum_{i=1}^N [\nu_i \log(\hat{\nu}(\mathbf{s}_i; \theta, \xi)) + (1 - \nu_i) \log(1 - \hat{\nu}(\mathbf{s}_i; \theta, \xi))],$$

where $\hat{\nu}(\mathbf{S}; \theta, \xi)$ is the predicted probability of whether an intervention will happen. For the latter, we minimize the negative log-likelihood of the human actions a_h (when $\nu = 1$) under the estimated action distribution in that state:

$$J_2(\theta) = \text{NLL}(\pi_\theta(\mathbf{a}_h | \mathbf{S})) = -\frac{1}{N} \sum_{i=1}^N \log(\pi_\theta(a_h^i | \mathbf{s}_i)) \quad (4)$$

J_1 (intervention loss) updates both the policy and the mental model, while J_2 (policy loss) updates only the policy π_θ , both terms contributing to policy learning:

$$J(\theta, \xi) = \lambda J_1(\theta, \xi) + (1 - \lambda) J_2(\theta) \quad (5)$$

In our experiments, we set $\lambda = 0.5$. We show the complete pipeline in Fig. 2, and the pseudocode on the project website.

V. SIMULATION EXPERIMENTS

Experiment Setup. We tested our method on four different simulation tasks, one with a discrete and three with continuous action spaces. The discrete action space environment is Lunarlander from Gymnasium [42]. The remaining three

tasks—Drawer-Open, Peg-Insertion, and Button-Press—are from the Metaworld suite [43] where the primary goal is to control a 6 DoF robot arm with a gripper to accomplish various tasks that require precise joint control and are prone to distributional shifts and local minima issues during RL optimization. For the Metaworld experiments, we use true world states as observations. To retain temporal information, we concatenate the states of the previous three timesteps with the current timestep. The action space is composed of 4 DoF: the Cartesian positions of the end effector, which always points down, and +1 dimension for the gripper. Since our model does not rely on any assumptions about the environment’s reward structure, and our primary comparisons are with other interactive learning methods, we follow the prior work [44] to evaluate the methods using success rates. Lunarlander is the only exception where we used the environment rewards due to the simplicity of the task.

We use simulated humans to mimic human interventions in our simulation experiments. The goal of these experiments is to scale up the tasks and to show that it is possible to finetune an initial policy in a sample-efficient way using our framework. The data collection phase involves two distinct policies: a suboptimal agent for generating rollouts and an expert agent for interventions guided by our intervention model. In place of the simulated human’s mental model of the robot (π_ζ), we train a BC policy on the rollouts of the suboptimal policy. During data collection, we rollout the suboptimal policy and intervene with the simulated human. In other words, the policy π_h of the human in the intervention model is replaced by π^* during the data collection using simulated humans. We see that these simulated interventions increase the success rate by comparing the dashed and the dotted lines in Fig. 3. During training, we jointly train the suboptimal agent and the mental model as we described in Section IV. At test time, we only use the trained policy and discard the mental model.

Baselines. We compare our methods against (i) BC intervention where the model estimates human’s interventions with only a neural network, (ii) HG-Dagger [7], a state-of-the-art algorithm that iteratively finetunes the policy based on the interventions, (iii) RLIF [28], a recent work which models the interactive imitation learning as an RL problem and uses interventions as the reward signals to train the policy in an online fashion, (iv) IWR [13] and (v) Sirius [12], two methods utilizing weighted BC with different heuristics to learn from both intervention and non-intervention timesteps. We start with the same initial suboptimal policy for all methods, and assume there is no access to any expert dataset, apart from the interventions provided.

To be consistent with the realistic settings, we keep the intervention ratio in the rollouts to be less than 30% by tuning the hyperparameter c . For our initial results, we train our method with $N = 1$ iteration and $k = 15$ episodes. For fairness, we train the other baselines until they have the same number of interventions in their datasets as ours.

Results. We present the results in Fig. 3. Our method achieves the best results across all environments (only tied

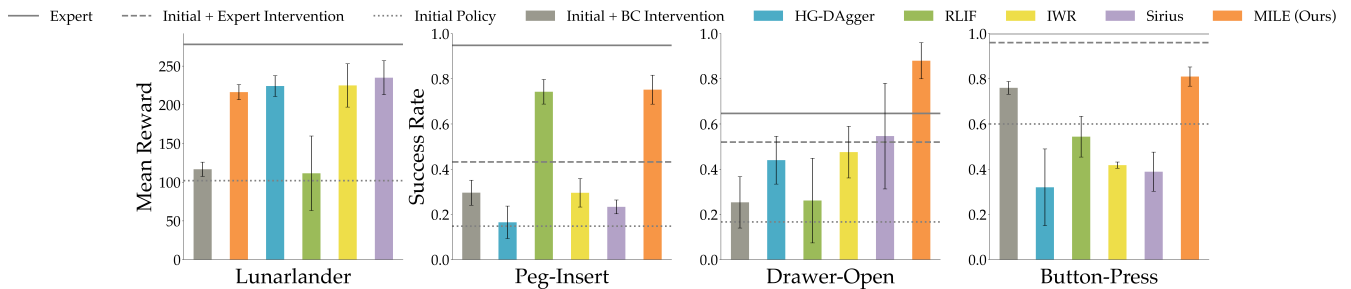


Fig. 3. **Success rates for single iteration training (mean±std).** MILE was trained with $N = 1$ iteration and $k = 15$ episodes. To make it a fair comparison, we train the other baselines until they have the same number of interventions in their datasets as ours.

with HG-Dagger, IWR and Sirius in Lunarlander, and with RLIF in Peg-Insert), proving its sample-efficiency. Due to the small amount of data which mostly consist of states with no interventions, IL methods suffer from overfitting as well as compounding errors. On the other hand, RLIF fails to produce reliable results, possibly due to the number of samples being too few to learn a successful policy.

To align more with the online nature of the baselines, we also perform iterative training with $N = 20$ iterations and $k = 1$ in Drawer-Open and Peg-Insert environments, as Button-Press was a relatively easier task. We run HG-Dagger, IWR, and Sirius in the same setting. For RLIF, we set the number of samples in its replay buffer approximately the same as MILE and other methods. The results are shown in Fig. 4. We observe that our method significantly outperforms the baselines. In just 10 iterations, it is able to achieve the success rate of the expert policy which provided the interventions, following with an even greater performance while the baselines mostly fall short.

Offline Demonstration Ablation. All of the baselines we compare against normally assume access to a set of offline demonstrations. Although MILE does not require such access, we run an ablation study in Drawer-Open, comparing our method against the baselines when they have access to expert demonstrations. Results are shown in Fig. 5. Our method performs better than RLIF even when it has access to 5 expert demonstrations. IL-based methods’ performance increases with the number of demonstrations in the dataset and becomes competitive with MILE but plateaus around the success rate of the expert agent. We believe this is due to




Fig. 5. **Success rates for the demo ablation study (mean±s.e.).**

overfitting and compounding errors. In both IWR and Sirius, robot actions are used as direct supervision signals in cases of non-intervention, albeit with different weights. In contrast, MILE uses the computational model for human interventions that we developed to construct probability distributions over the action space which are then used to train the models, thereby increasing the information bandwidth of the human feedback by computationally modeling them.

Overall, simulation results show our model for interventions improves learning, and increases adaptability, especially in low data regime where other interactive learning methods fail. With a handful of intervention trajectories, we are able to achieve near-optimal success levels in various tasks.

VI. REAL ROBOT EXPERIMENT

In order to show the effectiveness of our method in real world settings, we designed an experiment using a WidowX

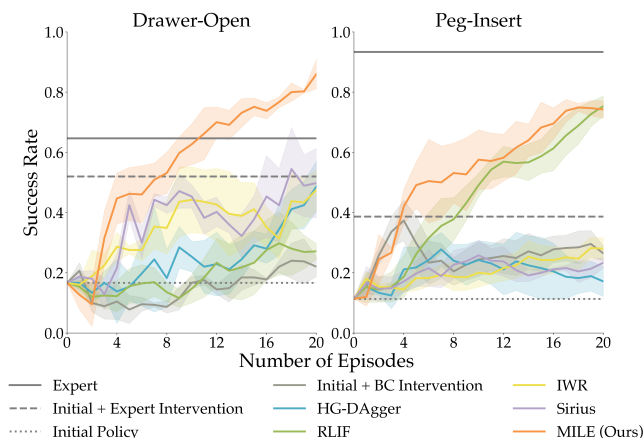


Fig. 4. **Success rates for iterative training (mean±s.e.).**

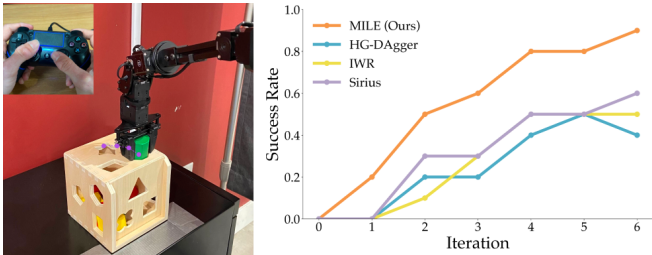


Fig. 6. **Real robot experiment setup and Success rates for MILE and other IL-based methods.** All methods were trained for 6 iterations with collecting 3 trajectories each iteration.

6-DoF robot arm and with a real human. The task is to put the octagonal block into the wooden box through the correct hole as shown in Fig. 6. The sensitivity of this task is very high, as there exists a very small margin between the block and the hole, which makes it a suitable candidate when the small interventions provided by the humans can have substantial effect on the robot’s success. We used image observations along with the robot’s end effector position.

The robot is initialized with a mediocre policy that is unable to complete the task. During the execution, the human can take over the controls of the robot at any time via a PS3 controller. Similarly, they can also give the controls back to the robot. We run the experiment with $N = 6$ iterations, with $k = 3$ intervention trajectories collected in each iteration.

We compare our method against IL-based methods, which have been somewhat more successful than RLIF in the simulation studies. All methods use the same data collection interface. Similar to the simulation experiments, we do not assume access to any offline expert demonstrations and only use the intervention trajectories. The results are shown in Fig. 6. Our method is able to achieve 80% success rate just after 4 iterations, while other baselines struggle to improve the policy. This result showcases that our method is applicable to real world settings as well.

VII. HUMAN SUBJECT STUDY ON REAL ROBOT

Finally, we conducted a user study to analyze how accurately our model estimates the humans’ interventions, as the success of our method relies on the success of the intervention model capturing when and how humans intervene the robot. The study is approved by the IRB office of the University of Southern California. We used the same task setting and experiment setup as in the previous section.

Procedure. We recruited 10 participants (3 female, 7 male, ages 20-28). The study started with a training period where subjects had a chance to get familiar with the robot controls.

We had 2 experiment settings. At the beginning of each experiment, the subject was first shown 3 trajectories from the initial policy to make them aware of the robot’s capabilities. In the first experiment, the user was asked to collect 5 intervention trajectories ($N = 1$, $k = 5$), where they can take over the initial policy at any time. These trajectories were then used to fine-tune the initial policy using our method, and the final policy was shown to the user. We asked the users to rate the initial and the final policies.

The second experiment differed from the first by only

using $N = 5$ and $k = 2$. The users could intervene and take control at any time during data collection: we did not provide any instructions on how often to intervene. We showed the finetuned policy at the end of each iteration and asked the user to rate its performance. At the end of the study, we did a post-study survey to collect participants’ feedback on the method’s effectiveness, improvement and their satisfaction on the final result.

Results. We compared our intervention model with the majority estimator that always predicts no intervention (the intervention rate of humans is around 30-40% throughout the procedure) and a neural network based estimator trained with the data from previous iterations to classify interventions in the current iteration. As shown in Fig. 7, our model outperforms both methods in predicting when an intervention will occur. Additionally, most users rated our model’s effectiveness and adaptation speed as adequate, demonstrating it successfully mirrors real human interventions.




Fig. 7. **User Study Results (mean±s.e.).** Legend: S=Satisfaction, I=Improvement, E=Effectiveness

VIII. CONCLUSION

Summary. In this work, we presented MILE, a novel way of integrating interventions in policy learning. We developed a fully-differentiable intervention model to estimate when and how the interventions occur, and a policy training framework that uses the intervention model.

Limitations and Future Work. Our intervention model considers only the current state to estimate interventions. This does not fully capture the temporal analysis humans make before intervening, e.g., waiting for robot to fail in a recoverable way before taking control. Future work should explore ways of integrating this temporal aspect of intervention estimation, possibly with maintaining a belief over robot actions and using Bayesian updates. Another limitation is our assumption that humans have an accurate mental model of the robot policy. Investigating how incorrect mental models affect the performance is an important direction for future research. Recently, some works showed that the interventions can be used for data generation and allow better sim-to-real generalization [45], [46]. Testing our method’s capabilities in these domains and exploring whether using a structure behind interventions enables further improvements in generalization are interesting future directions. Finally, learning from natural language-based corrections [47], [48] is a recent exciting research area that may benefit from our intervention model.

REFERENCES

- [1] D. Kalashnikov, A. Irpan, P. Pastor, *et al.*, “Scalable deep reinforcement learning for vision-based robotic manipulation,” in *Conference on Robot Learning (CoRL)*, 2018.
- [2] S. Levine, C. Finn, T. Darrell, and P. Abbeel, “End-to-end training of deep visuomotor policies,” *The Journal of Machine Learning Research*, vol. 17, no. 1, 2016.
- [3] L. Ouyang, J. Wu, X. Jiang, *et al.*, “Training language models to follow instructions with human feedback,” in *Conference on Neural Information Processing Systems (NeurIPS)*, 2022.
- [4] D. Silver, J. Schrittwieser, K. Simonyan, *et al.*, “Mastering the game of go without human knowledge,” *nature*, vol. 550, no. 7676, 2017.
- [5] S. Ross and D. Bagnell, “Efficient reductions for imitation learning,” in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, 2010.
- [6] S. Ross, G. Gordon, and D. Bagnell, “A reduction of imitation learning and structured prediction to no-regret online learning,” in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, 2011.
- [7] M. Kelly, C. Sidrane, K. Driggs-Campbell, and M. J. Kochenderfer, “Hg-dagger: Interactive imitation learning with human experts,” in *International Conference on Robotics and Automation (ICRA)*, 2019.
- [8] K. Menda, K. Driggs-Campbell, and M. J. Kochenderfer, “Ensembledagger: A bayesian approach to safe imitation learning,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [9] R. Hoque, A. Balakrishna, E. Novoseller, A. Wilcox, D. S. Brown, and K. Goldberg, “Thriftydagger: Budget-aware novelty and risk gating for interactive imitation learning,” in *Conference on Robot Learning (CoRL)*, 2022.
- [10] R. Hoque, L. Y. Chen, S. Sharma, *et al.*, “Fleet-dagger: Interactive robot fleet learning with scalable human supervision,” in *Conference on Robot Learning (CoRL)*, 2023.
- [11] K. Baraka, I. Idrees, T. Kessler Faulkner, *et al.*, “Human-interactive robot learning: Definition, challenges, and recommendations,” 2025.
- [12] H. Liu, S. Nasiriany, L. Zhang, Z. Bao, and Y. Zhu, “Robot learning on the job: Human-in-the-loop autonomy and learning during deployment,” in *Robotics: Science and Systems (RSS)*, 2023.
- [13] A. Mandlekar, D. Xu, R. Martín-Martín, Y. Zhu, L. Fei-Fei, and S. Savarese, “Human-in-the-loop imitation learning using remote teleoperation,” *arXiv preprint arXiv:2012.06733*, 2020.
- [14] J. Spencer, S. Choudhury, M. Barnes, *et al.*, “Learning from interventions: Human-robot interaction as both explicit and implicit feedback,” in *Robotics: Science and Systems (RSS)*, 2020.
- [15] C. Celemin, R. Pérez-Dattari, E. Chisari, *et al.*, “Interactive imitation learning in robotics: A survey,” *Foundations and Trends in Robotics*, vol. 10, no. 1-2, 2022.
- [16] 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, 2009.
- [17] J. Ho and S. Ermon, “Generative adversarial imitation learning,” in *Conference on Neural Information Processing Systems (NeurIPS)*, 2016.
- [18] M. Laskey, J. Lee, R. Fox, A. Dragan, and K. Goldberg, “Dart: Noise injection for robust imitation learning,” in *Conference on Robot Learning (CoRL)*, 2017.
- [19] P. Abbeel, A. Coates, and A. Y. Ng, “Autonomous helicopter aerobatics through apprenticeship learning,” *The International Journal of Robotics Research*, vol. 29, no. 13, 2010.
- [20] D. A. Pomerleau, “Alvin: An autonomous land vehicle in a neural network,” in *Conference on Neural Information Processing Systems (NeurIPS)*, 1988.
- [21] V. Myers, E. Biyik, and D. Sadigh, “Active reward learning from online preferences,” in *International Conference on Robotics and Automation (ICRA)*, 2023.
- [22] E. Chisari, T. Welschhold, J. Boedecker, W. Burgard, and A. Valada, “Correct me if i am wrong: Interactive learning for robotic manipulation,” *arXiv preprint arXiv:2110.03316*, 2021.
- [23] T. Hester, M. Vecerik, O. Pietquin, *et al.*, “Deep q-learning from demonstrations,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.
- [24] P. J. Ball, L. Smith, I. Kostrikov, and S. Levine, “Efficient online reinforcement learning with offline data,” in *International Conference on Machine Learning (ICML)*, 2023.
- [25] A. Nair, B. McGrew, M. Andrychowicz, W. Zaremba, and P. Abbeel, “Overcoming exploration in reinforcement learning with demonstrations,” in *International Conference on Robotics and Automation (ICRA)*, 2018.
- [26] T. Xie, N. Jiang, H. Wang, C. Xiong, and Y. Bai, “Policy finetuning: Bridging sample-efficient offline and online reinforcement learning,” in *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- [27] Q. Li, Z. Peng, and B. Zhou, “Efficient learning of safe driving policy via human-ai copilot optimization,” *arXiv preprint arXiv:2202.10341*, 2022.
- [28] J. Luo, P. Dong, Y. Zhai, Y. Ma, and S. Levine, “Rlif: Interactive imitation learning as reinforcement learning,” in *International Conference on Learning Representations (ICLR)*, 2023.
- [29] S. Wang, A. Wang, S. Goncharova, B. Scassellati, and T. Fitzgerald, “Effects of robot competency and motion legibility on human correction feedback,” in *International Conference on Human-Robot Interaction (HRI)*, 2025.
- [30] J. Scholtz, “Theory and evaluation of human robot interactions,” in *International Conference on System Sciences*, 2003.
- [31] A. Steinfeld, T. Fong, D. Kaber, *et al.*, “Common metrics for human-robot interaction,” in *International Conference on Human-robot Interaction (HRI)*, 2006.
- [32] L. Rodriguez Rodriguez, C. E. Bustamante Orellana, E. K. Chiou, L. Huang, N. Cooke, and Y. Kang, “A review of mathematical models of human trust in automation,” *Frontiers in Neuroergonomics*, vol. 4, 2023.
- [33] Y. Guo and X. J. Yang, “Modeling and predicting trust dynamics in human-robot teaming: A bayesian inference approach,” *International Journal of Social Robotics*, vol. 13, no. 8, 2021.
- [34] M. Chen, S. Nikolaidis, H. Soh, D. Hsu, and S. Srinivasa, “Planning with trust for human-robot collaboration,” in *International Conference on Human-robot Interaction (HRI)*, 2018.
- [35] W. Saunders, G. Sastry, A. Stuhlmüller, and O. Evans, “Trial without error: Towards safe reinforcement learning via human intervention,” in *International Conference on Autonomous Agents and Multiagent Systems*, 2018.
- [36] C. I. Bliss, “The method of probits,” *Science*, vol. 79, no. 2037, 1934.
- [37] A. D. Dragan and S. S. Srinivasa, “A policy-blending formalism for shared control,” *The International Journal of Robotics Research*, vol. 32, no. 7, 2013.
- [38] S. Javdani, S. S. Srinivasa, and J. A. Bagnell, “Shared autonomy via hindsight optimization,” in *Robotics: Science and Systems (RSS)*, 2015.
- [39] M. Kwon, E. Biyik, A. Talati, K. Bhasin, D. P. Losey, and D. Sadigh, “When humans aren’t optimal: Robots that collaborate with risk-aware humans,” in *International Conference on Human-Robot Interaction (HRI)*, 2020.
- [40] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey, “Maximum entropy inverse reinforcement learning,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2008.
- [41] C. Laidlaw and A. Dragan, “The boltzmann policy distribution: Accounting for systematic suboptimality in human models,” in *International Conference on Learning Representations (ICLR)*, 2022.
- [42] M. Towers, J. K. Terry, A. Kwiatkowski, *et al.*, *Gymnasium*, 2023. DOI: [10.5281/zenodo.8127026](https://doi.org/10.5281/zenodo.8127026).
- [43] T. Yu, D. Quillen, Z. He, *et al.*, “Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning,” in *Conference on Robot Learning (CoRL)*, 2019.
- [44] S. A. Sontakke, J. Zhang, S. M. R. Arnold, *et al.*, “Roboclip: One demonstration is enough to learn robot policies,” in *Conference on Neural Information Processing Systems (NeurIPS)*, 2023.
- [45] Y. Jiang, C. Wang, R. Zhang, J. Wu, and L. Fei-Fei, “Transic: Sim-to-real policy transfer by learning from online correction,” in *Conference on Robot Learning (CoRL)*, 2024.
- [46] R. Hoque, A. Mandlekar, C. Garrett, K. Goldberg, and D. Fox, “Intervengen: Interventional data generation for robust and data-efficient robot imitation learning,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- [47] Z. Yang, M. Jun, J. Tien, S. J. Russell, A. Dragan, and E. Biyik, “Trajectory improvement and reward learning from comparative language feedback,” in *Conference on Robot Learning (CoRL)*, 2024.
- [48] L. X. Shi, Z. Hu, T. Z. Zhao, *et al.*, “Yell at your robot: Improving on-the-fly from language corrections,” in *Robotics: Science and Systems (RSS)*, 2024.