

Figure 6: Environments from the EARL benchmark [5] used for simulated experiments. From left to right, the environments are: Peg insertion, Door closing and Tabletop organization.




Figure 7: (left) Randomized position of the cube in the grasping task. The position marked by violet boundary are within the distribution of expert demonstrations, and the rest are outside the distribution. (right) Architecture overview for MEDAL++.

542 A Appendix

543 A.1 Algorithm Overview

544 The pseudocode for training is given in Algorithm 1. First, the parameters and data buffers in \mathcal{F} and
 545 \mathcal{B} are initialized and the forward and backward demonstrations are loaded into \mathcal{D}_f^* and \mathcal{D}_b^* respec-
 546 tively. Next, we update the forward and backward goal sets, as described above. After initializing
 547 the environment, the forward policy π_f interacts with the environment and collects data, updating
 548 the networks and buffers in \mathcal{F} . The control switches over to the backward policy π_b after a fixed
 549 number of steps, and the networks and buffers in \mathcal{B} are updated. The backward policy interacts for
 550 a fixed number of steps, after which the control is switched over to the forward policy and this cycle
 551 is repeated thereon. When executing in the real world, humans are allowed to intervene and reset
 552 the environment intermittently, switching the control over to π_f after the intervention to restart the
 553 forward-backward cycle.

554 We now expand on how the networks are updated for π_f during training (also visualized in Figure 9);
 555 the updates for π_b are analogous. First, the new transition in the environment is added to \mathcal{D}_f . Next,
 556 we sample a batch of states from \mathcal{D}_f and label them 0, and sample a batch of *equal size* from \mathcal{D}_f^* and
 557 label them 1. The classifier C_f is updated using gradient descent on the combined batch to minimize
 558 the cross-entropy loss. Note, the classifier is not updated for every step collected in the environment.
 559 As stated earlier, the classification problem is easier than learning the policy, and therefore, it helps
 560 to train the classifier slower than the policy. Finally, the policy π_f , Q -value networks $\{Q_n^f, \bar{Q}_n^f\}_{n=1}^N$
 561 and the encoder \mathcal{E} are updated on a batch of transitions constructed by sampling $(1 - \rho)B$ transi-
 562 tions from \mathcal{D}_f and ρB transitions from \mathcal{D}_f^* . The Q -value networks and the encoder are updated
 563 by minimizing $\frac{1}{N} \sum_{n=1}^N \ell(Q_n, \mathcal{E})$ (Eq 1), and the target Q -networks are updated as an exponential
 564 moving average of Q -value networks. The policy π_f is updated by maximizing $\mathcal{L}(\pi)$. We update

Algorithm 1: MEDAL++

```
initialize  $\mathcal{F}, \mathcal{B}$ ; // forward, backward parameters  
 $\mathcal{F}.\mathcal{D}_f^*, \mathcal{B}.\mathcal{D}_b^* \leftarrow \text{load\_demonstrations}()$   
 $\mathcal{F}.\mathcal{G}_f \leftarrow \text{get\_states}(\mathcal{F}.\mathcal{D}_f^*, -K:)$  // last  $K$  states  
// exclude last  $K$  states from  $\mathcal{D}_f^*$ , use only the last  $K$  states from  $\mathcal{D}_b^*$   
 $\mathcal{B}.\mathcal{G}_b \leftarrow \text{get\_states}(\mathcal{F}.\mathcal{D}_f^*, :-K) \cup \text{get\_states}(\mathcal{B}.\mathcal{D}_b^*, -K:)$   
 $s \sim \rho_0; \mathcal{A} \leftarrow \mathcal{F}$ ; // initialize environment  
while not done do  
   $a \sim \mathcal{A}.\text{act}(s); s' \sim \mathcal{T}(\cdot | s, a);$   
   $\mathcal{A}.\text{update\_buffer}(\{s, a, s'\});$   
   $\mathcal{A}.\text{update\_classifier}();$   
   $\mathcal{A}.\text{update\_parameters}();$   
  // switch policy after a fixed interval  
  if switch then  
    |  $\text{switch}(\mathcal{A}, (\mathcal{F}, \mathcal{B}));$   
  // allow intermittent human interventions  
  if interrupt then  
    |  $s \sim \rho_0;$   
    |  $\mathcal{A} \leftarrow \mathcal{F};$   
  else  
    |  $s \leftarrow s';$ 
```




Figure 8: Visualizing the positive target states for forward classifier C_f and backward classifier C_b from the expert demonstrations. For forward demonstrations, last K states are used for C_f (orange) and the rest are used for C_b (pink). For backward demonstrations, last K states are used for C_b .




Figure 9: An overview of MEDAL++ training. The classifier is trained to discriminate states visited by an expert from the states visited online. The robot reinforcement learning learns on a combination of self-collected and expert transitions, and the policy learning is regularized using the behavior cloning loss.

565 the Q -value networks multiple times for every step collected in the environment, whereas the policy
566 network is updated once for every step collected in the environment [15].

567 A.2 Implementation Details and Practical Tips

568 An overview of the architecture used by the forward and backward networks is shown in Figure 7.

569 *Visual Encoder:* For the encoder, we use the same architecture as DrQ-v2 [12]: 4 convolutional
570 layers with 32 filters of size (3, 3), stride 1, followed by ReLU non-linearities. The high-dimensional
571 output from the CNN is embedded into a 50 dimensional feature using a fully-connected layer,
572 followed by LayerNorm and tanh non-linearity (to output the features normalized to $[-1, 1]$). For
573 real-robot experiments, the first person and third person views are concatenated channel wise before
574 being passed into the encoder. The output of the encoder is fused with proprioceptive information,
575 in this case, the end-effector position, before being passed to actor and critic networks.

576 *Actor and Critic Networks:* Both actor and critic networks are parameterized as 4 layer fully-
577 connected networks with 1024 ReLU hidden units for every layer. The actor parameterizes a Gaus-

578 sian distribution over the actions, where a tanh non-linearity on the output restricts the actions to
579 $[-1, 1]$. We use an ensemble size of 10 critics.

580 *Discriminators:* The discriminator for the forward and backward policies use a similar visual en-
581 coder but with 2 layers instead of 4. The visual embedding is passed to a fully connected network
582 with 2 hidden layers with 256 ReLU units. When training the network, we use mixup and spectral
583 norm regularization [53, 52] for the entire network.

584 *Training Hyperparameters:* For all our experiments, $K = 20$, i.e. the number of frames used as goal
585 frames. The forward policy interacts with the environment for 200 steps, then the backward policy
586 interacts for 200 steps. In real world experiments, we also reset the arm every 1000 steps to avoid
587 hitting singular positions. Note, this reset does not require any human intervention as the controller
588 just resets the arm to a fixed joint position. We use a batch size of 256 to train the policy and critic
589 networks, out of which 64 transitions are sampled from the demonstrations (*oversampling*). We use
590 a batch size of 512 to train the discriminators, 256 of the states come from expert data and the other
591 256 comes from the online data. Further, the discriminators are updated every 1000 steps collected
592 in the environment. The update-to-data ratio, that is the number of gradient updates per transition
593 collected in the environment is 3 for simulated environments and 1 for the real-robot experiments.
594 We use a linearly decaying schedule for behavior cloning regularization from 1 to 0.1 over the first
595 50000 steps which remains fixed at 0.1 onwards throughout training.

596 For real world experiments, we use a wrist camera to improve the overall performance [58], and
597 provide **only** the wrist-camera view to both discriminators. We find that this further regularizes
598 the discriminator. Finally, we provide no proprioceptive information for the VICE discriminator,
599 but we give MEDAL discriminator the proprioceptive information, as it needs a stronger notion
600 of the robot’s localization to adequately reset to a varied number of initial positions for improved
601 robustness.

602 *Teleoperation:* To collect our demonstrations on the real robot, we use an Xbox controller that
603 manipulates the end-effector position, orientation and the gripper state. Two salient notes: (1) The
604 forward and backward demonstrations are collected together, one after the other and (2) the initial
605 position for demonstrations is randomized to cover as large a state-space as feasible. The increased
606 coverage helps with exploration during autonomous training.

607 A.3 EARL Environments, Training and Evaluation

608 **Environments.** We consider three sparse-reward continuous-control environments from EARL
609 benchmark [5], shown in Appendix, Fig 6). *Tabletop organization* is a simplified manipulation
610 environment where a gripper is tasked to move the mug to one of the four coasters from a wide set
611 of initial states, *sawyer door closing* task requires a sawyer robot arm to learn how to close a door
612 starting from various positions, and finally the *sawyer peg insertion* task requires the sawyer robot
613 arm to grasp the peg and insert it into a goal. Not only does the robot have to learn how to do the task
614 (i.e. close the door or insert the peg), but it has to learn how to undo the task (i.e. open the door or
615 remove the peg) to try task repeatedly in the non-episodic training environment. The sparse reward
616 function is given by $r(s, a) = \mathbb{1}(\|s - g\| \leq \epsilon)$, where g denotes the goal, and ϵ is the tolerance for
617 the task to be considered completed.

618 **Training and Evaluation.** The environments are setup to return 84×84 RGB images as obser-
619 vations with a 3-dimensional action space for the *tabletop organization* (2D end-effector deltas in
620 the XY plane and 1D for gripper) and a 4-dimensional action space for *sawyer* environments (3D
621 end-effector delta control + 1D gripper). The training environment is reset to $s_0 \sim \rho_0$ every 25,000
622 steps of interaction with the environment. This is extremely infrequent compared to episodic set-
623 tings where the environment is reset to the initial state distribution every 200-1000 steps. EARL
624 comes with 5-15 forward and backward demonstrations for every environment to help with explo-
625 ration in these sparse reward environments. We evaluate the forward policy every 10,000 training
626 steps, where the evaluation approximates $\mathbb{E}_{s_0 \sim \rho_0} [\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)]$ by averaging the return of the




Figure 10: An overview of MEDAL++ on the task of inserting the peg into the goal location. (*top*) Starting with a set of expert trajectories, MEDAL++ learns a forward policy to *insert the peg* by matching the goal states and a backward policy to *remove and randomize the peg position* by matching the rest of the states visited by an expert. (*bottom*) Chaining the rollouts of forward and backward policies allows the robot to practice the task autonomously. The rewards indicate the similarity to their respective target states, output by a discriminator trained to classify online states from expert states.

627 policy over 10 episodes starting from $s_0 \sim \rho_0$. These roll-outs are used only for evaluation, and not
 628 for training.

629 A.4 Real-world Experiment Analysis

630 We discuss the four manipulation tasks in detail. We recommend viewing the supplemental website
 631 for training and evaluation videos:

632 (1) *Cube Grasping*: The goal in this task is to grasp the cube from varying initial positions and con-
 633 figurations and raise it. For this task, we consider a controlled setting to isolate one potential source
 634 of improvement from autonomous reinforcement learning: robustness to the initial state distribution.
 635 Specifically, all the forward demonstrations are collected starting from a narrow set of initial states
 636 (*ID*), but, the robot is evaluated starting from both *ID* states and out-of-distribution (*OOD*) states,
 637 visualized in Appendix, Figure 7. BC policy is competent on *ID* states, but it performs poorly on
 638 states that are *OOD*. However, after autonomous self-improvement using MEDAL++, we see an
 639 improvement of 15% on *ID* performance, and a large improvement of 74% on *OOD* performance.
 640 Autonomous training allows the robot to practice the task from a diverse set of states, including
 641 states that were *OOD* relative to the demonstration data. This suggests that improvement in success
 642 rate results partly from being robust to the initial state distribution, as a small set of demonstrations
 643 is unlikely to cover all possible initial states a robot can be evaluated from.

644 (2) *Cloth on the Hook*: In this task, the robot is tasked with grasping the cloth and putting it through
 645 a fixed hook. To practice the task repeatedly, the backward policy has to remove the cloth from
 646 the hook and drop it on platform. Here, MEDAL++ improves the success rate over BC by 36%.
 647 The BC policy has several failure modes: (1) it fails to grasp the cloth, (2) it follows through with
 648 hooking because of memorization, or (3) it hits into into the hook because it drifts from the right
 649 trajectory and could not recover. Autonomous self-improvement improves upon all these issues,
 650 but particularly, it learns to re-try grasping the cloth if it fails the first time, rather than following a
 651 memorized trajectory observed in the forward demonstrations.

652 (3) *Bowl Covering with Cloth*: The goal of this task is to cover a bowl entirely using the cloth.
 653 The cloth can be a wide variety of initial states, ranging from ‘laid out flat’ to ‘scrunched up’ in
 654 varying locations. The task is challenging as the robot has to grasp the cloth at the correct location
 655 to successfully cover the entire bowl (partial coverage is counted as a failure). Here, MEDAL++
 656 improves the performance over BC by 34%. The failure modes of BC are similar to previous task,
 657 including failure to grasp, memorization and failure to re-try, and incomplete coverage due to wrong
 658 initial grasp. Autonomous self-improvement substantially helps with the grasping (including re-
 659 trying) and issues related to memorization. While it plans the grasps better than BC, there is room
 660 for improvement to reduce failures resulting from partially covering the bowl.

661 (4) *Peg Insertion*: Finally, we consider the task of inserting a peg into a goal location. The location
 662 and orientation of the peg is randomized, in service of which we use 5DoF control for this task. A
 663 successful insertion requires the toy to be perpendicular to the goal before insertion, and the error

664 margin for a successful insertion is small given the size of the peg and the goal. Additionally, the
665 peg here is a soft toy, it can be grasped while being in the wrong orientation. Here, MEDAL++
666 improves the performance by 48% over BC. In addition to failures described in the previous tasks,
667 a common cause of failure is the insertion itself where the agent takes an imprecise trajectory and
668 is unable to insert the peg. After autonomous self-improvement, the robot employs an interesting
669 strategy where it keeps retries the insertion till it succeeds. The policy is also better at grasping,
670 though the failures of insertion often result from orienting the gripper incorrectly before the grasp
671 which makes insertion infeasible.