

# Behavior Retrieval: Few-Shot Imitation Learning by Querying Unlabeled Datasets

Maximilian Du, Suraj Nair, Dorsa Sadigh, Chelsea Finn  
Stanford University  
{maxjdu, surajn}@stanford.edu

**Abstract**—Enabling robots to learn novel visuomotor skills in a data-efficient manner remains an unsolved problem with myriad challenges. A popular paradigm for tackling this problem is through leveraging large unlabeled datasets that have many behaviors in them and then adapting a policy to a specific task using a small amount of task-specific human supervision (i.e. interventions or demonstrations). However, how best to leverage the narrow task-specific supervision and balance it with offline data remains an open question. Our key insight in this work is that task-specific data not only provides new data for an agent to train on but can also *inform the type of prior data the agent should use for learning*. Concretely, we propose a simple approach that uses a small amount of downstream expert data to selectively query relevant behaviors from an offline, unlabeled dataset (including many sub-optimal behaviors). The agent is then jointly trained on the expert and queried data. We observe that our method learns to query only the relevant transitions to the task, filtering out sub-optimal or task-irrelevant data. By doing so, it is able to learn more effectively from the mix of task-specific and offline data compared to naively mixing the data or only using the task-specific data. Furthermore, we find that our simple querying approach outperforms more complex goal-conditioned methods by 20% across simulated and real robotic manipulation tasks from images. See <https://sites.google.com/view/behaviorretrieval> for videos and code.

## I. INTRODUCTION

One of the promises of deep learning applied to robotics is the ability to learn control from sensor observations in a data-efficient manner, a prerequisite to deploying robots in rich real-world settings like homes. Towards accomplishing this goal, a number of prior methods have explored leveraging broad, unlabeled data to capture general knowledge and a small amount of task-specific data to adapt to a particular task. A natural form for this task-specific data to take is human expert data (e.g. demonstrations or online interventions), which can be imitated to efficiently adapt the agent to the target task. While this data is convenient to learn from, it comes at a cost: expensive human-in-the-loop supervision for every target task. Thus, even with pre-training data providing general knowledge, learning a performant task-specific policy without substantial human effort remains an open challenge.

Unlike the many works that study different techniques for offline pre-training from sub-optimal data, from visual pre-training [32, 30, 42] to multi-task imitation [20, 9, 39], in this work we focus instead on how to best use the sub-optimal offline data in fine-tuning a policy from small amount of expert data. Critically, most prior work do not make use of the offline data during fine-tuning, and those that do [9] tend to use

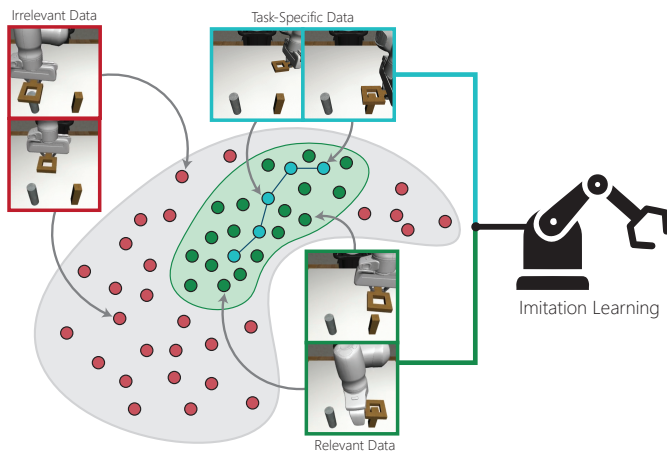


Fig. 1: Using Task-Specific Data to Query Offline Datasets. Using a small amount of task-specific human expert feedback (i.e. interventions or demonstrations) (blue), our method learns to select the relevant portions (green) of a broad, unlabeled offline dataset (red) to efficiently learn the target task. In this example the task is to have the robot place the square on the right peg (shown with the initial and final frame in blue). While the broader dataset might include irrelevant data (initial and final frame shown in red) where the robot is placing the square on the left peg, it includes useful data providing diversity in terms of how the square needs to be placed on the right peg (shown with initial and final frame in green). Our algorithm identifies these relevant data from the broader dataset and learns from them while ignoring the irrelevant data.

simple heuristics that often fail to effectively balance offline and task-specific expert data, or focus on using the offline data for skills instead of low-level control [31]. The interesting question to ask is *why balancing these data sources tend to be difficult*. On one hand, fine-tuning the policy only on a small amount of downstream expert data can be unstable, diverging significantly from the pre-trained policy and losing much of the knowledge gained during pre-training. On the other hand, jointly training the agent on the online expert data and the offline pre-training data can be highly sub-optimal, as much of the pre-training data may be irrelevant or damaging to learning the target task. We study precisely this tradeoff, and seek to understand how best to balance task-specific human expert data and broad unlabeled data in learning control policies.

Our key insight in this work is that the small amount of task-specific human expert data not only provides new data for an agent to train on but also **informs what prior data the agent should retrieve for learning**. That is, when a human provides

supervision to the agent through for example an intervention (a state-action tuple), the agent should not solely fine-tune on the tuple. Rather, the agent can learn to retrieve many relevant state-action tuples from a broad offline dataset and instead fine-tune on the larger combined data.

Concretely, we propose behavior retrieval, a method for training a policy on a small amount of human expert data (interventions or demonstrations), while simultaneously leveraging a broad, unlabeled dataset that may include sub-optimal, task-irrelevant, and task-adversarial behaviors. Our method begins by using the unlabeled dataset to pre-train a state-action similarity metric measuring the similarity between any two state-action pairs. Then, when adapting the policy with a small amount of human expert data, we retrieve relevant experience from the unlabeled dataset using our learned similarity metric, and jointly fine-tune the policy on the expert and queried data (See Figure 1). In a set of simulated and real-world robotic manipulation tasks from both low-dimensional states and image observations, we find that selectively sampling from offline data for joint fine-tuning with our method leads to more stable and performant learning than only finetuning on the human expert data or using hand-designed heuristics for balancing offline and task-specific data. Moreover, we find our proposed approach outperforms multi-task pre-training and finetuning by over 20% by using its simple retrieval strategy.

## II. RELATED WORK

Towards the goal of data-efficient policy learning, there are many relevant works spanning few-shot learning from experts which studies learning with just a few demos, offline pre-training which aims to learn priors from broad datasets, and meta-learning which learns policies that can adapt efficiently.

**Few-shot Learning from Experts.** Learning behaviors from human expert data through imitation learning is a long-studied problem in robotics [34, 37, 36, 22, 33, 1, 5, 15]. More recently, the emphasis for such imitation learning methods has been on learning efficiently, minimizing the effort placed on the human expert. For example, many works have explored how a policy might be pre-trained on demonstrations then fine-tuned with online expert feedback [36, 24, 28, 22, 41, 19]. Similarly, many works have explored how to learn from broad data with many tasks and fine-tune on a new task given a small amount task-specific data [9, 27]. Rather than simply fine-tuning on expert interventions or demos, our method leverages *retrieval*, filtering only the relevant prior experience and jointly training on the combined data. In our experiments, we compare to all of these prior methods and find that our approach behavior retrieval outperforms them by selectively retrieving and training on offline data.

Additionally, several works have explored designing action spaces or planning trajectories in a way that enables more efficient learning [44, 21, 17, 16]. Behavior retrieval instead proposes a method for sampling from offline data, and in principle can be combined with any such action spaces in a straightforward manner. Recent approaches also attempt to be transparent about what type of data the imitation learning

algorithm requires and further influence the demonstrator through interactive data collection to provide informative, in-distribution data [12, 38]. Such methods are complementary to our approach, and the two could be naturally combined to influence the demonstrator for better interventions which then enables even more effective retrieval from offline data using behavior retrieval. Finally, our work is not the first to consider retrieval in imitation learning. Nasiriany et al. has explored learning skills from offline data and retrieving relevant skills for long-horizon tasks. Unlike this work, we learn low-level control policies, and focus on retrieving individual state-action tuples rather than high-level skills [31]. This allows us to leverage offline data to learn more performant low-level policies, as opposed to being tied to a high-level skill space. Other works like Goyal et al. [13] also explore the use of retrieval, however focus on sample efficient RL using retrieval from the replay buffer, rather than imitation from unstructured datasets like we consider in this work.

**Robotic Pre-training.** While our approach focuses on retrieving behaviors as a way of leveraging offline data, there exist many alternative methods for pre-training from offline data. For example, one line of work studies pre-training visual representations [32, 30, 42, 35, 26] that can then be used frozen or finetuned for more efficient downstream imitation or reinforcement learning. This approach is particularly effective at using data that does not contain action labels, like videos of humans. Another line of work focuses on using robot data with actions, and pre-training multi-task (e.g. goal-conditioned or language conditioned) policies through imitation learning [6, 25, 3, 4, 9, 20] or offline reinforcement learning [23]. Like our work, some prior work also explores selectively sampling or weighting expert data [45, 2]. However, the focus of these works is on learning from a mix of sub-optimal experts for a particular task, while we instead focus on retrieving from a broad offline dataset during task-specific fine-tuning. In total, all of these works are complementary to the direction we study in this work, as any of the above pre-training schemes can leverage the behavior retrieval approach to efficiently finetune.

**Meta-Imitation Learning.** Towards the goal of few-shot learning of control policies for new tasks and environments, another well-studied approach is that of meta-learning. Meta-imitation learning methods [11, 7, 18, 43, 46] train to produce effective policies using a few demonstrations of many tasks, such that given just a few demos of a new task they can perform well. Critically all of these methods require a large dataset of well-defined, labeled tasks for meta-training, which can be expensive. Our approach is instead able to learn from unlabeled “play” data without annotated task boundaries and use it for few-shot imitation learning.

## III. BEHAVIOR RETRIEVAL: IMITATION LEARNING BY QUERYING UNLABELED DATASETS

### A. Problem Setup

In this work, we consider the problem setting of learning a target task from a small amount of task-specific expert data and

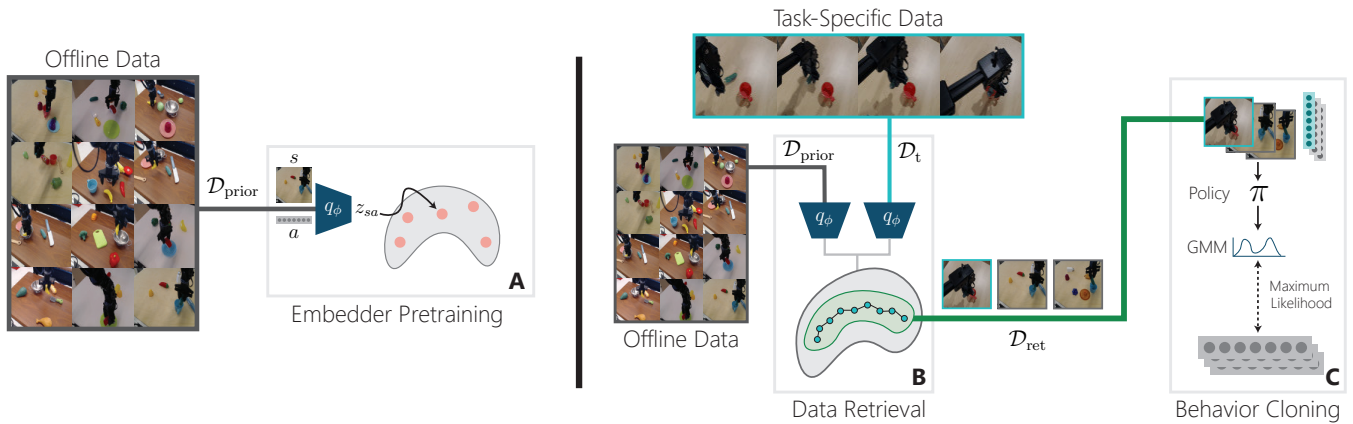


Fig. 2: **The Behavior Retrieval Method.** Our approach has 3 main steps. (A) Using the unlabeled offline data  $\mathcal{D}_{\text{prior}}$  we pre-train a state-action embedding. (B) We use the pre-trained embedding to look up similar transitions in the offline data  $\mathcal{D}_{\text{prior}}$  that are relevant to the task data  $\mathcal{D}_t$ . (C) We then train a policy with behavior cloning on the mix of the task-specific and retrieved data.

a large amount of sub-optimal, unlabeled data, where all the data shares a state space  $\mathcal{S}$  and an action space  $\mathcal{A}$ . Concretely, for a particular target task  $t$  in the space of all tasks  $\mathcal{T}$ , we assume access to:

- 1)  $\mathcal{D}_t$ : a dataset of expert  $(s, a)$  tuples for the target task  $t$ , which may come either from online interventions or small set ( $\sim 10$ ) of demonstrations.
- 2)  $\mathcal{D}_{\text{prior}}$ : which contains an unsegmented mix of data from tasks in the space of  $\mathcal{T}$ , which may include everything from random data to near-expert data for task  $t$  to sub-optimal data for any task other than  $t$ .

In practice, we expect the prior data to have some data that is relevant to task  $t$  and some that is not relevant. Given these data sources, the final goal is to learn a policy  $\pi_t : \mathcal{S} \rightarrow \mathcal{A}$  that is performant on task  $t$ , and to do so in a way that leverages the prior data to get better performance. In the extreme case where none of the prior data is relevant, the optimal method should ideally not be harmed by the prior data and learn a policy as good purely learning from  $\mathcal{D}_t$ .

### B. Overview

In behavior retrieval, we explore going beyond the standard approach of pre-training a policy on the prior dataset  $\mathcal{D}_{\text{prior}}$  and fine-tuning on the task-specific dataset  $\mathcal{D}_t$ . Our key insight is that  $\mathcal{D}_t$  not only provides data to train on but also informs what data from  $\mathcal{D}_{\text{prior}}$  is relevant for learning. Moreover, we can learn the metric for measuring this relevance from  $\mathcal{D}_{\text{prior}}$ . Concretely, our method, behavior retrieval, first learns a  $(s, a)$ -similarity metric from  $\mathcal{D}_{\text{prior}}$ , then given  $\mathcal{D}_t$  uses the similarity metric to retrieve relevant  $(s, a)$  tuples from  $\mathcal{D}_{\text{prior}}$ , and finally trains  $\pi_t$  on the union of  $\mathcal{D}_t$  and the retrieved data with imitation learning (See Figure 2). We cover the similarity metric training in Section III-C, the filtering process in Section III-D, and the policy training in Section III-E. See Algorithm 1 for a summary of the behavior retrieval method.

### C. Similarity Metric Pre-training

To effectively retrieve behaviors from  $\mathcal{D}_{\text{prior}}$  relevant to  $\mathcal{D}_t$ , we need some way of measuring the similarity between two  $(s, a)$  tuples, a non-trivial problem for high-dimensional




Fig. 3: **Training the State-Action Embedding.** We train a variational auto-encoder jointly on states and actions to produce our state-action embedding  $z_{sa}$ .

state spaces like images. Critically, because  $\mathcal{D}_{\text{prior}}$  may have behaviors from completely different tasks than the target task, it's not sufficient to simply pull transitions with similar states, as the actions may be poorly-suited to the target task. Therefore, we want to learn a function  $\mathcal{F} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  which takes as input two state-action pairs and predicts a scalar score measuring their functional similarity. The first component of our method is learning  $\mathcal{F}$  using  $\mathcal{D}_{\text{prior}}$ .

At a high level, we do so by learning a low dimensional  $(s, a)$  embedding using  $\mathcal{D}_{\text{prior}}$ , and instantiate  $\mathcal{F}$  as the distance between the embeddings of  $(s_1, a_1) \sim \mathcal{D}_{\text{prior}}$  and  $(s_2, a_2) \sim \mathcal{D}_t$  tuples. There are many possible approaches to learn these  $(s, a)$  embeddings, ranging from contrastive learning to compression via an auto-encoder. In our experiments we find that compressing states and actions with a VAE provides a simple approach with the best performance (See Figure 3).

Concretely, we simply learn a joint variational autoencoder for the states and actions, which are trained to maximize the evidence lower bound (ELBO)

$$\mathbb{E}_{q_\phi(z_{sa}|s,a)}[\log p_\theta(s, a|z_{sa})] - \beta \mathcal{D}_{KL}[q_\phi(z_{sa}|s, a)||p(z_{sa})] \quad (1)$$

on state-action pairs sampled from  $\mathcal{D}_{\text{prior}}$ , where  $\beta$  is a tuned hyperparameter. The state and action are encoded separately then combined through a multi-layer perceptron (MLP) to form the state-action embedding. The action is additionally concatenated to this state-action embedding as  $z = [z_{sa}, a]$ , which we found to help with behavior separability.

Finally given the two tuples  $(s_1, a_1)$  and  $(s_2, a_2)$  we pro-

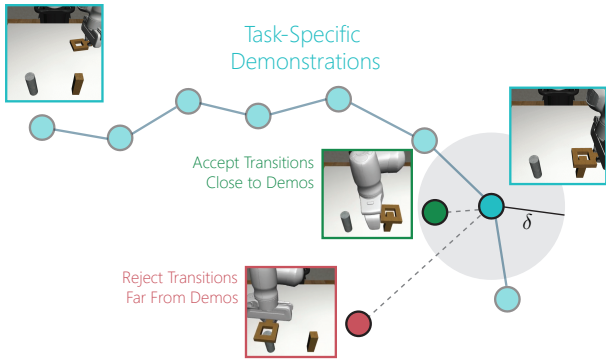


Fig. 4: **Retrieving from the Unlabeled Dataset.** Using our pre-trained state-action embedder, we compute the embeddings for the offline dataset  $\mathcal{D}_{\text{prior}}$  and the small number of task-specific demos  $\mathcal{D}_t$ . Then, we select transitions in the offline dataset within a certain distance of the task-specific demos, in embedding space.

duce  $z_1$  and  $z_2$ . The distance function,  $\mathcal{F}$ , is then simply the negative L2 distance between the two embeddings:

$$\mathcal{F}(s_1, a_1, s_2, a_2) = -\|z_1 - z_2\|_2 \quad (2)$$

#### D. Filtering Relevant Data

Given the pre-trained  $\mathcal{F}$ , how might we retrieve relevant data from  $\mathcal{D}_{\text{prior}}$  using  $\mathcal{D}_t$ ? While there are many valid strategies from upweighing losses to importance sampling, we hypothesize that transitions in  $\mathcal{D}_{\text{prior}}$  where the expert is taking similar actions in similar states will be useful to imitate for the target task, and thus should be leveraged. Therefore, if a transition in  $\mathcal{D}_{\text{prior}}$  is relevant to **any** transition in  $\mathcal{D}_t$ , we choose to retrieve it and learn from it with behavior cloning.

Concretely, let the upper bound on similarity  $\mathcal{F}^+ = \max_{(s^*, a^*) \in \mathcal{D}_{\text{prior}}} \max_{(s, a) \in \mathcal{D}_t} \mathcal{F}(s^*, a^*, s, a)$  and the lower bound  $\mathcal{F}^- = \min_{(s^*, a^*) \in \mathcal{D}_{\text{prior}}} \max_{(s, a) \in \mathcal{D}_t} \mathcal{F}(s^*, a^*, s, a)$ . Then, for any transition  $(s^*, a^*) \in \mathcal{D}_{\text{prior}}$ , we retrieve it if:

$$\frac{\max_{(s, a) \in \mathcal{D}_t} \mathcal{F}(s^*, a^*, s, a) - \mathcal{F}^-}{\mathcal{F}^+ - \mathcal{F}^-} > \delta \quad (3)$$

We take this minmax normalization across  $\mathcal{F}$  to account for varying magnitudes of embeddings. In our experiments, we use  $\delta \in [0.6, 0.8]$ , although we show in section IV-D that our method is robust to changes in  $\delta$ . Since we retrieve the top  $1 - \delta\%$  of  $\mathcal{D}_{\text{prior}}$  that we think is most relevant, when using behavior retrieval, a user should choose  $\delta$  to be one minus the percent of  $\mathcal{D}_{\text{prior}}$  likely to be relevant. For instance, if almost none of  $\mathcal{D}_{\text{prior}}$  is expected to be relevant,  $\delta$  should be set close to 1, and if all of  $\mathcal{D}_{\text{prior}}$  is relevant  $\delta$  should be set close to 0. See Figure 4 for visual illustration of the filtering process.

#### E. Training with Retrieved Data

Finally, given the retrieved data,  $\mathcal{D}_{\text{ret}}$ , we then train the agent to jointly imitate  $\mathcal{D}_{\text{ret}}$  and  $\mathcal{D}_t$  with a behavior cloning loss. That is, we train  $\pi_\psi$  with the objective:

$$\min_{\psi} \mathbb{E}_{(s, a) \sim \mathcal{D}_t} [-\log \pi_\psi(a|s)] + \mathbb{E}_{(s, a) \sim \mathcal{D}_{\text{ret}}} [-\log \pi_\psi(a|s)] \quad (4)$$

#### Algorithm 1 Behavior Retrieval

---

```

1: Input: Prior dataset  $\mathcal{D}_{\text{prior}} = [(s_1, a_1), \dots, (s_T, a_T)]_{1:N}$ 
2: Input: Task dataset  $\mathcal{D}_t = [(s_1, a_1), \dots, (s_T, a_T)]_{1:10}$ 
3: Initialize  $\psi, \phi, \theta$  randomly,  $\mathcal{D}_{\text{ret}} \leftarrow \emptyset$ . Batch size  $B$ .
4: /* Train VAE embedding space */
5: while VAE not converged do
6:    $s_t, a_t \sim \mathcal{D}_{\text{prior}}$ 
7:   Update  $\theta, \phi$  to maximize Eq. 1
8: /*Retrieve task-relevant data*/
9: for  $(s, a) \in \mathcal{D}_{\text{prior}}$  do
10:  if Eq. 3 is True for  $(s, a)$  then
11:     $\mathcal{D}_{\text{ret}} \leftarrow \mathcal{D}_{\text{ret}} \cup (s, a)$ 
12: /* Fit policy to task data and retrieved data */
13: while  $\pi_\psi$  not converged do
14:   $[(s, a)]_t^{1:B} \sim \mathcal{D}_t$ 
15:   $[(s, a)]_{\text{ret}}^{1:B} \sim \mathcal{D}_{\text{ret}}$ 
16:  Update  $\psi$  to maximize  $\log \pi_\psi(a|s)$ 
    on  $[(s, a)]_t^{1:B}$  and  $[(s, a)]_{\text{ret}}^{1:B}$ 

```

---

We encode the observation modalities separately before combining them into a shared state embedding for the policy. For the policy, we use an LSTM [14] architecture, which allows for object permanence amidst occlusion and the observation of features like object velocity. The policy outputs the parameters to a Gaussian mixture model (GMM), and the final actions are chosen by sampling from the GMM, which allows for better modeling of multi-modal behavior often present in a large multi-task datasets. See Appendix A for more details on the policy architecture and training. Algorithm 1 summarizes the method, with lines 5-7 describing the embedding pre-training, lines 9-11 summarizing the filtering process, and lines 13-16 describing the policy training phase.

#### IV. EXPERIMENTS

The main goal of our experiments is to understand the empirical performance of behavior retrieval, both in terms of how it works and how well it works. After describing the experimental set-up (Section IV-A), we will first examine what data is retrieved by our method in both simulated and real robotic manipulation tasks (Section IV-B). We will then analyze the performance of behavior retrieval in comparison to pre-training with finetuning (Section IV-C1), in comparison to other retrieval strategies (Section IV-C2), and in evaluations on a physical robot (Section IV-C3) using a heterogenous large-scale dataset collected in prior work [9]. Beyond final performance, we also study the robustness of behavior retrieval (Section IV-D) and ablate design choices and hyperparameters (Section IV-E). In the appendix, we additionally include experiments that study behavior retrieval’s ability to stitch together behaviors from short-horizon tasks into a sequential task (Appendix D-A), its ability to be combined with other policy pretraining methods [25] (Appendix D-B), its robustness to larger and more diverse  $\mathcal{D}_{\text{prior}}$  (Appendix D-C), and an analysis of failure modes of behavior retrieval and comparisons (Appendix D-D).




Fig. 5: **Simulation Environments.** We consider three simulated domains. In each domain we highlight the downstream task data  $\mathcal{D}_t$  in blue, relevant offline data from  $\mathcal{D}_{\text{prior}}$  in green, and irrelevant downstream data from  $\mathcal{D}_{\text{prior}}$  in red. In RoboSuite Can Pick and Place (left), the agent must pick a place a can into the bin, and irrelevant data involves throwing the can randomly. In RoboSuite Nut Assembly (middle), the agent must insert a square into the correct peg, and irrelevant data involves putting it onto the wrong peg. In PyBullet WidowX Office Cleanup (right), the agent must pick and place an eraser into a specified tray, where irrelevant data involved many actions with other objects in the scene.

#### A. Experimental Domains

We study behavior retrieval and comparisons in three simulated domains with 7DOF end-effector control and three real robot domains with 4DOF end-effector control, visualized in Figures 5 and 6 respectively. The task data  $\mathcal{D}_t$  for all environments contain 10 expert demos for the target task.

**Sim: CanPick.** In this RoboSuite environment [47, 29], the task is for a simulated Franka robot to pick and place a coke can from one bin to another.  $\mathcal{D}_{\text{prior}}$  contains a mix of 400 human-collected demos [29] where half complete the task and half fail by randomly throwing the can out of the bin.

**Sim: NutAssembly.** The modified RoboSuite [47, 29] task is for a simulated Franka robot to pick and insert a square into the right peg.  $\mathcal{D}_{\text{prior}}$  contains a mix of 400 machine-generated demos where half complete the task and, in the other half, the robot puts the square onto the wrong (left) peg.

**Sim: Office.** The third task is for a simulated Widowx robot in an office environment [40] to pick an eraser and place it into a specified tray.  $\mathcal{D}_{\text{prior}}$  contains 1200 machine-generated demos of 8 separate pick-place behaviors. Out of these,  $\mathcal{D}_{\text{prior}}$  contains 150 demos of the target eraser-tray task.

**Real: NutAssembly** Analogous to the simulated nut assembly task, a real WidowX robot needs to pick up a blue square hole and insert it onto the red peg. As constructed, the hole and peg have a narrow tolerance, requiring precise control for successful insertion.  $\mathcal{D}_{\text{prior}}$  contains a mix of 160 human-collected demos where half complete the task and half fail by inserting the peg onto the green peg.

**Real: Pickle** The second real-world task is for a real WidowX robot to pick up a pickle from the table and place it into

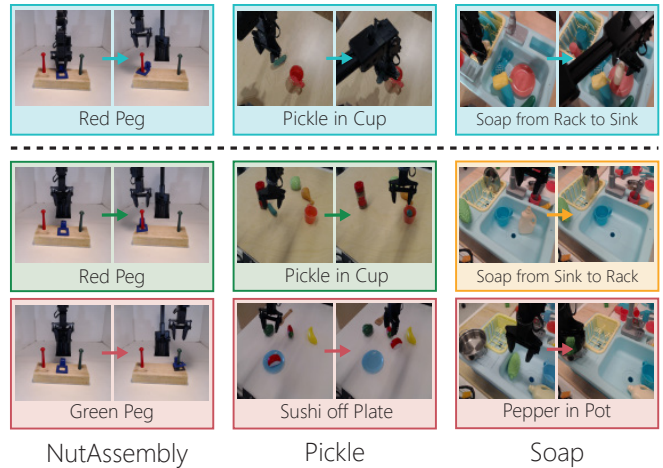


Fig. 6: **WidowX Robot Environments.** We consider three robot domains. In each domain we highlight the downstream task data  $\mathcal{D}_t$  in blue, relevant offline data from  $\mathcal{D}_{\text{prior}}$  in green, partially relevant offline data from  $\mathcal{D}_{\text{prior}}$  in orange, and irrelevant offline data from  $\mathcal{D}_{\text{prior}}$  in red. In Nut Assembly (left), the agent must insert a square into the correct peg, and irrelevant data involves putting it onto the wrong peg (like the task in simulation). In Pickle in Cup (middle), the agent must pick and place a pickle in a cup, and  $\mathcal{D}_{\text{prior}}$  is a large subset of the Bridge Dataset collected in a similar tabletop setting [9]. In Soap in Sink (right), the agent must pick and place a soap in sink, a totally novel task, and again  $\mathcal{D}_{\text{prior}}$  is a large subset of the Bridge Dataset collected in a similar sink setting. [9].

a cup randomly placed within the robot’s visual field.  $\mathcal{D}_{\text{prior}}$  is a subset of the Bridge Data [9] that contains tabletop manipulations, totalling to 285 demos. Out of these demos, 45 correspond to an analogous pickle-cup task, although there are significant shifts in lighting, table color, and distractor objects.

**Real: Soap** The final task is for a real Widowx robot to pick up a detergent container from the dish rack of a toy kitchen sink and place it, standing up, on a red plate in the sink.  $\mathcal{D}_{\text{prior}}$  is a subset of the Bridge Dataset [8] that contains all the 485 trajectories collected on a copy of this sink. There is no analogous task in the Bridge Dataset, although there are trajectories involving the detergent bottle, as well as sub-trajectories that pull other objects from the dish rack or place other objects into the sink.

#### B. What Data Does Behavior Retrieval Retrieve?

In our first experiment, we qualitatively analyze what type of behaviors are retrieved in the Sim: NutAssembly and Real: Pickle tasks. In Figure 7, we show example states from  $\mathcal{D}_t$ , and three randomly sampled states from the top-25 highest ranked transitions in  $\mathcal{D}_{\text{prior}}$  according to behavior retrieval and three randomly sampled states from the lowest 25 transitions in  $\mathcal{D}_{\text{prior}}$  according to behavior retrieval. We observe that indeed the states ranked highly by behavior retrieval are relevant to the target task in  $\mathcal{D}_t$ , and the lowest ranked transitions are visibly from different or adversarial tasks (e.g. moving the nut to the wrong peg).

In our next experiment, we quantitatively evaluate how well behavior retrieval separates task relevant and irrelevant data, particularly looking at the NutAssembly task in both




Fig. 7: **Visualizing the Retrieved Data:** For each target task in  $\mathcal{D}_t$  (top), we highlight the 3 random states and their scores from the top-25 highest ranked transition by behavior retrieval (middle), as well as 3 random states and their scores from the lowest 25 ranked transition by behavior retrieval (bottom). We see that behavior retrieval is indeed ranking highly transitions that visually look relevant to the target task in  $\mathcal{D}_t$ , and rejecting states that are clearly from other tasks.

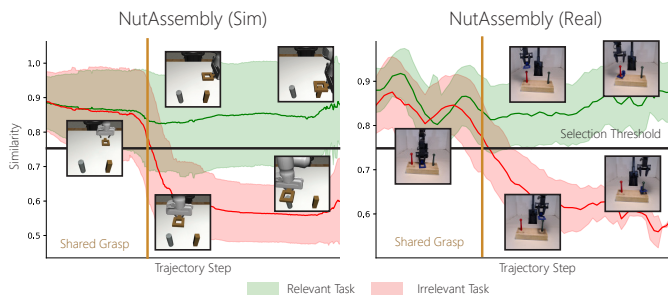


Fig. 8: **Demonstration of Embedding Separation.** We measure how behavior retrieval scores data from  $\mathcal{D}_{\text{prior}}$  for the target task. Green is the ground truth relevant data, while red is the ground truth irrelevant demos. We see that behavior retrieval learns to make use of the part of the irrelevant demos that are useful to the target task (grasping), but only highly scores the relevant demos afterwards. This highlights that behavior retrieval is able to make use of all relevant transitions, from both relevant and irrelevant tasks, but effectively discards transitions that aren’t useful to the target task.

simulation and the real robot (other environments in the Appendix C). Concretely, we plot the average embedding similarity for the ground truth relevant data (green) and for the ground truth irrelevant data (red) (See Figure 8). For example, in the nut assembly task data from  $\mathcal{D}_{\text{prior}}$  that puts the square on the correct peg is green, and on the incorrect peg is red. While the agent is grasping the square, behavior retrieval retrieves data from both sources as the grasp is always relevant to the downstream task. However, once the agent has grasped, we see the score for the task-irrelevant data drop significantly. This highlights behavior retrieval’s ability to make use of all transitions in  $\mathcal{D}_{\text{prior}}$  that are useful for the target task, while effectively filtering out those that are not useful, a valuable property for being able to leverage large, heterogeneous datasets as  $\mathcal{D}_{\text{prior}}$ .

	GCBC $\mathcal{D}_{\text{prior}}$	GCBC $\mathcal{D}_{\text{prior}}$ + FT $\mathcal{D}_t$	Ours
CanPick (Image)	68 ± 14%	62 ± 2%	<b>74 ± 1%</b>
CanPick (State)	33 ± 1%	45 ± 1%	<b>73 ± 1%</b>
NutAssembly (Image)	22 ± 4%	10 ± 1%	<b>58 ± 9%</b>
NutAssembly (State)	15 ± 1%	10 ± 1%	<b>75 ± 5%</b>
Office (Image)	0 ± 0%	10 ± 4%	<b>64 ± 3%</b>
Office (State)	76 ± 5%	76 ± 1%	<b>85 ± 4%</b>
<b>Average</b>	<b>35 ± 4%</b>	<b>36 ± 2%</b>	<b>72 ± 2%</b>

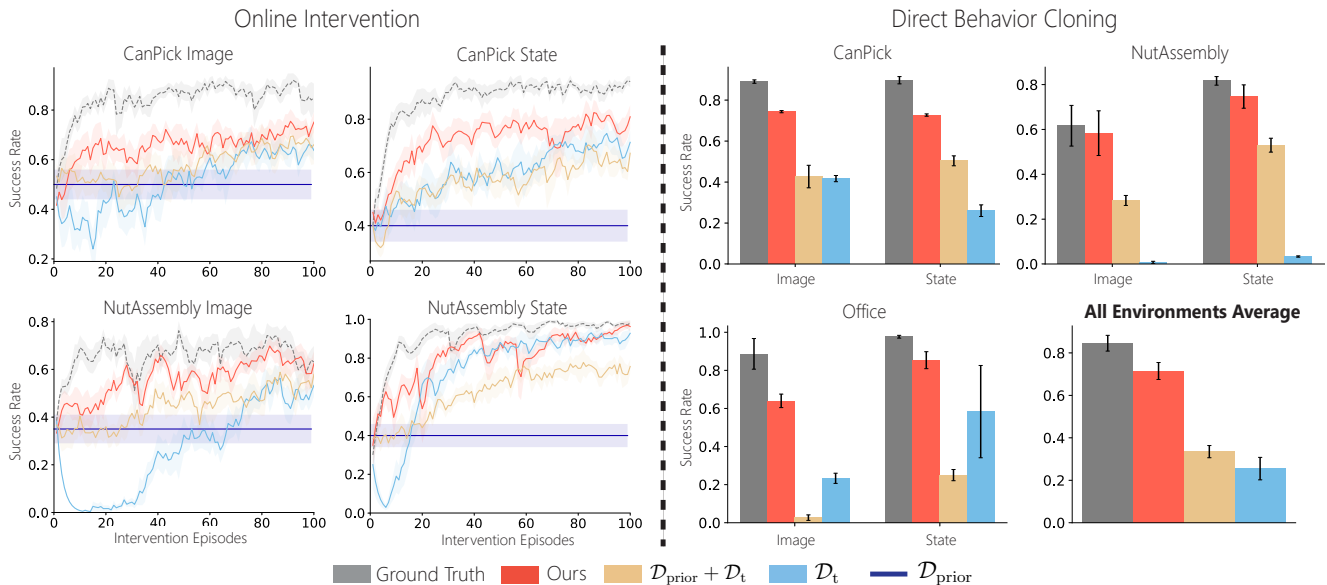
TABLE I: **Comparing Behavior Retrieval to Pre-training + Finetuning in Simulation.** We report the success rate on our three simulated tasks across 100 trials (averaged over 3 random seeds with standard error shown) of behavior retrieval compared to goal-conditioned behavior cloning (GCBC) on  $\mathcal{D}_{\text{prior}}$  and with finetuning on  $\mathcal{D}_t$ . We see that both with and without fine-tuning multi-task pre-training struggles and behavior retrieval performs best across all settings, improving performance by over 30% on average.

### C. Does Behavior Retrieval Improve Performance ...

1) *Over Other Pre-training+Finetuning Schemes?:* First, we aim to understand how behavior retrieval compares to other techniques that use pre-training and adapt the policy on a target dataset. A conventional approach is to perform multi-task pre-training of  $\mathcal{D}_{\text{prior}}$ , followed by fine-tuning on the target task data  $\mathcal{D}_t$ . We begin by comparing behavior retrieval (**Ours**) to such methods in simulation: pre-training with goal-conditioned behavior cloning on  $\mathcal{D}_{\text{prior}}$  (**GCBC  $\mathcal{D}_{\text{prior}}$** ), as well as with additional finetuning of the goal-conditioned policy on balanced batches of  $\mathcal{D}_t$  and  $\mathcal{D}_{\text{prior}}$  (**GCBC  $\mathcal{D}_{\text{prior}}$  +FT  $\mathcal{D}_t$** ). See Appendix A for more details on the implementation of GC and GC+Finetune. In Table I we observe that across all three simulated tasks, behavior retrieval improves over goal-conditioned BC, with and without finetuning, improving the success rate by 30% on average.

We hypothesize that purely offline goal-conditioned BC struggles with underfitting, particularly in cases where goal-states may over-specify the task (i.e. many goal-states can represent the same task). Similarly, adding fine-tuning can be unstable – for example if training only on  $\mathcal{D}_t$  the fine-tuned policy can overfit, while if training on balanced batches with the offline data, the underfitting issue remains. This is particularly the case in settings where the desired change in state only makes up a small portion of the goal-states, like in Office (Image). Ultimately behavior retrieval is able to pull only the relevant data from  $\mathcal{D}_{\text{prior}}$ , and train on all of it without overfitting to  $\mathcal{D}_t$ , resulting in the best performance.

2) *Over Other Retrieval Strategies?:* We also compare behavior retrieval (**Ours**) to other retrieval strategies across all of our simulated environments. This includes not doing any retrieval ( $\mathcal{D}_t$ ), training on balanced batches from  $\mathcal{D}_{\text{prior}}$  and  $\mathcal{D}_t$  ( $\mathcal{D}_{\text{prior}} + \mathcal{D}_t$ ), and using the ground-truth task label to filter only task-relevant demos from  $\mathcal{D}_{\text{prior}}$  (**Ground Truth**). This last approach can be thought of as a privileged Oracle, that does not use our learned filtering metric, but instead uses the ground truth label to filter. This will retrieve demos from  $\mathcal{D}_{\text{prior}}$  that are known to be the same task, however will not



**Fig. 9: Comparing Behavior Retrieval to Other Retrieval Strategies.** In simulation, we compare the success rate of behavior retrieval (**Ours**) to other strategies for retrieval including only training on  $\mathcal{D}_t$  ( $\mathcal{D}_t$ ), balanced batches of  $\mathcal{D}_{\text{prior}}$  and  $\mathcal{D}_t$  ( $\mathcal{D}_{\text{prior}} + \mathcal{D}_t$ ), only ( $\mathcal{D}_{\text{prior}}$ ) (interventions only), and using **Ground Truth** separation which uses privileged information. We consider both online interventions (top) and 10 demos (bottom) as  $\mathcal{D}_t$ . In both cases we see behavior retrieval approach the ground truth oracle performance, while significantly outperforming the comparisons, which either overfit to only using  $\mathcal{D}_t$ , or train too much on sub-optimal data from  $\mathcal{D}_{\text{prior}}$ . All results are computed over 100 trials (or 20 trials every episode in the intervention setting), all averaged over 3 seeds with standard error shading.

retrieve transitions from a different task that may be relevant. We consider both the setting where downstream expert data comes in the form of online interventions (Figure 9 (left)) and when it comes in the form of 10 demos (Figure 9 (right)).

We observe in Figure 9 that despite the expert data coming in very different forms, whole demos vs. online interventions, in both settings behavior retrieval significantly improves success rate over other retrieval strategies, by over 20% given a small set of demos, and by around 10% given online interventions. In the online intervention setting, the human supervision is strong (on-policy corrections), therefore we expect  $\mathcal{D}_t$  to be particularly useful and using  $\mathcal{D}_{\text{prior}}$  to be less important, hence why we see smaller gains from behavior retrieval. In the setting where the agent is given a small set of demonstrations (Figure 9 (right)), behavior retrieval nearly matches using the oracle retrieval from  $\mathcal{D}_{\text{prior}}$ , suggesting that it is effectively capturing the relevant data to  $\mathcal{D}_t$ .

In the case of online interventions,  $\mathcal{D}_t$  struggles to learn from the small amount of online expert data, and thus can overfit significantly or experience sharp drops in performance (Figure 9 (left), NutAssembly and CanPick (Image)). Even in the case of a handful of target demos,  $\mathcal{D}_t$  has larger variance (Figure 9 (right), Office (State)). Naïvely using balanced batches of  $\mathcal{D}_{\text{prior}}$  and  $\mathcal{D}_t$  can lead to training with task-irrelevant data, leading to worse task-specific performance.

Comparing online interventions overall (Figure 9 (left)) to full expert demos (Figure 9 (right)) for  $\mathcal{D}_t$ , we observe that on average performance with online expert data is slightly better, an expected result as the human supervision is in-the-loop and on the states the agent most needs feedback. Additionally, we see on average that learning from low-dim state information

tends to perform slightly better than from images, with some variation across tasks, again matching our intuition that lower-dimensional state space of object and robot positions is easier to learn from.

3) *On a Real Robot?:* While behavior retrieval performs effectively, do the benefits translate to a physical robot? In this experiment, we compare behavior retrieval to all of the above comparisons from Sections IV-C1 and IV-C2 on a real WidowX robot. We test on three real robot tasks, NutAssembly, Pickle, and Soap, as described in Section IV-A, where  $\mathcal{D}_{\text{prior}}$  for the Pickle and Soap tasks involve using a subset the pre-existing Bridge Dataset [9].

In our results (Figure 10), we again see that behavior retrieval improves over both other retrieval strategies and pre-training + finetuning by more than 20% in success rate. In fact, we see that because it is able to use useful transitions even from irrelevant tasks in the offline data, behavior retrieval exceeds using only prior data from the same task (Ground Truth), which uses privileged information in the form of task labels.

Additionally, we note that like in our simulation experiments, only using  $\mathcal{D}_t$  tends to be unstable, completely failing on some tasks like NutAssembly. Unlike our simulation experiments, we see that adding finetuning to the goal-conditioned pre-training has a much more significant impact on performance, improving over GCBC by over 30%. We suspect this is the case because there is a large shift from the Bridge Data  $\mathcal{D}_{\text{prior}}$  to the downstream task specific environment, so finetuning is much more important.

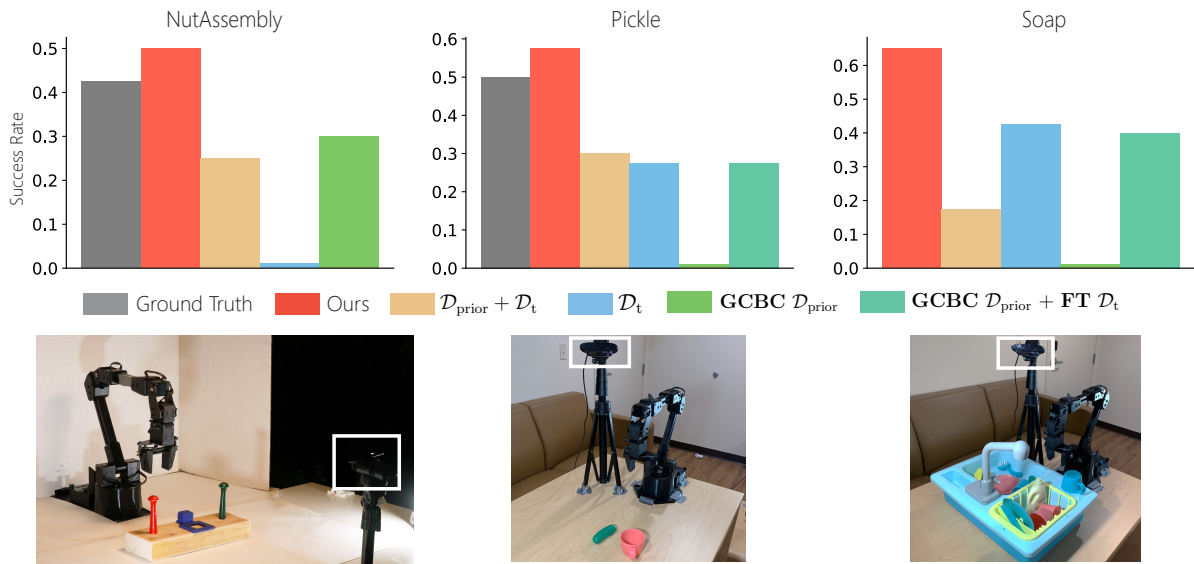


Fig. 10: **Real Robot Comparisons.** Across our three robot tasks, we compare the success rate given 10 target demos across 40 evaluations of behavior retrieval, other retrieval methods, and pre-training and fine-tuning (**top**). We again see that behavior retrieval improves over the comparisons by over 20% in success rate. Because it is able to leverage the useful transitions from irrelevant tasks in the offline data, behavior retrieval exceeds using the ground truth filtering which uses only data from  $\mathcal{D}_{\text{prior}}$  from the same exact task. We show the robot setup for each task (**bottom**), as well as the camera capturing observations highlighted in the white box.

	$\mathcal{D}_t$	Ours
Different Cup Color	1/10	<b>6/10</b>
Real Pickle	0/10	<b>4/10</b>
Distractors	0/10	<b>4/10</b>
Physical Perturbation	3/10	<b>4/10</b>

TABLE II: **Robustness to Distribution Shifts.** We test behavior retrieval’s robustness to distribution shifts at test time, including changing the cup color. We find effectively using  $\mathcal{D}_{\text{prior}}$  to improve robustness, maintaining performance close to a 50% success rate while using only task-specific data ( $\mathcal{D}_t$ ) drops considerably to around a 10% success rate.

#### D. How Robust is Behavior Retrieval?

In our next experiment, we study how robust behavior retrieval is with respect to distribution shifts, embedding spaces, and hyperparameters.

First on the real robot Pickle in Cup task, we study testing on a different colored cup than training, using a real pickle instead of a toy pickle, visual distractors, and physical perturbations. For visual distractors, we place four toy kitchen items in the visual field of the robot and shuffle them for each trial. For physical perturbations, we allow the robot to begin the alignment process with the pickle in its grasp, and then we move the cup to a different location.

In Table II we find that across all distribution shifts, behavior retrieval is able to maintain strong performance (nearly a 50% success rate), while the model using only task-specific data completely collapses to around a 10% success rate. We hypothesize that retrieving broader data that is still useful to the task makes the trained policy robust to such variation. Moreover, behavior retrieval is able to make use of the Bridge Data [9], which contains a WidowX250, while our evaluations are carried out on a WidowX200 robot, further highlighting its

	Ours w/ Contrastive	Ours w/ VAE (default)
CanPick (Image)	47 ± 1%	<b>74 ± 1%</b>
CanPick (State)	55 ± 3%	<b>73 ± 1%</b>
NutAssembly (Image)	30 ± 3%	<b>58 ± 9%</b>
NutAssembly (State)	19 ± 3%	<b>75 ± 5%</b>
Office (Image)	N/A	<b>64 ± 3%</b>
Office (State)	39 ± 1%	<b>85 ± 4%</b>
Average	38 ± 2%	<b>72 ± 2%</b>

TABLE III: **Comparing State-Action Embedding Spaces for Retrieval.** In our simulation tasks we compare the success rate over 100 trials (and 3 random seeds) of our method with different choices for learning the state-action embedding space. We find that using other techniques like contrastive RL to learn the state-action embedding can be unstable, leading to worse performance than the simple VAE in all tasks, and diverging in the case of Office (Image).

robustness to distribution shifts.

#### E. Ablating Components of Behavior Retrieval

In this section, we ablate different design choices and hyperparameters in behavior retrieval, specifically the choice of (1)  $(s, a)$  embedding space and (2) filtering threshold  $\delta$ .

**Ablating  $(s, a)$  Embedding.** We begin by comparing our VAE embedding space to alternative methods for learning the  $(s, a)$  embedding. Concretely, we also explore using contrastive RL [10], which learns an  $(s, a)$  embedding that captures the future distribution of visited states with contrastive learning. In principle, this embedding should capture a functional embedding of  $(s, a)$ . We then explore using behavior retrieval swapping in a contrastive RL embedding instead of our default VAE  $(s, a)$  embeddings.

In Table III, we observe that the VAE embedding performs better across all simulated tasks. While in principle contrastive



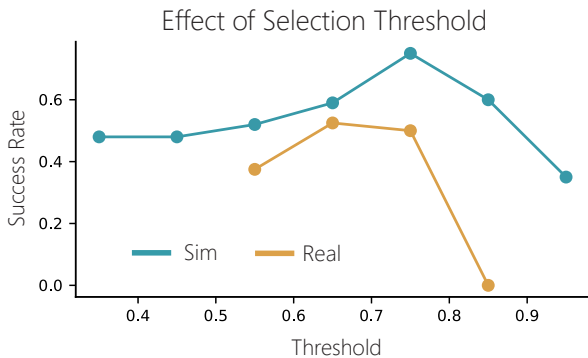


Fig. 11: **Impact of Selection Threshold.** On the NutAssembly task in both simulation and the real world, we find that behavior retrieval is generally robust to the choice of  $\delta$  with the best performance in both sim and real with  $\delta$  around 0.75. Interestingly, we see performance scales gracefully as  $\delta$  decreases, but performance drops significantly as  $\delta$  increases (filtering out all of  $\mathcal{D}_{\text{prior}}$ ).

RL should capture more “functional” state-action embeddings, we found it to be unstable, even diverging and failing to train the representation effectively in the Office environment from image observations.

**Ablating  $\delta$ .** Additionally, we study the method’s robustness to hyper-parameters, specifically the embedding distance threshold  $\delta$  for when to accept or reject samples from  $\mathcal{D}_{\text{prior}}$ . In Figure 11, we evaluate different thresholds in simulation and on the real robot. First we see that, across domains, a threshold of around  $\delta = 0.75$  works best. Interestingly, we find that the method is robust to lower values of  $\delta$ , and performance degrades gracefully. Intuitively, this means that the method is sampling more task-irrelevant data, but is still capturing all the task-relevant data. On the other hand, performance drops sharply with too high of a  $\delta$ , which again makes sense: too high of a  $\delta$  will throw out all of  $\mathcal{D}_{\text{prior}}$  including useful data.

## V. DISCUSSION

**Summary.** We have presented behavior retrieval, a method for leveraging large, unlabeled offline datasets to efficiently learn target tasks from a small amount of expert data. Our key insight is that rather than simply pre-training and finetuning, we can learn to selectively retrieve only the useful data from the prior dataset to improve task-specific training. behavior retrieval presents a simple technique for tackling this problem, and in our experimental evaluations outperforms both multi-task pre-training and finetuning approaches as well as other retrieval strategies. Moreover, we see qualitatively that behavior retrieval indeed retrieves task-relevant transitions while discarding those that are irrelevant.

**Limitations and Future Work.** While we are excited about the potential of behavior retrieval, a number of important limitations remain. First, a major limitation is that our current lookup relies on similarity in a compressed embedding space of states and actions. While this has worked effectively in our experiments, it focuses on general visual similarity and may not capture the task-relevant components of a state and action

in all scenarios, e.g. in cases where  $\mathcal{D}_{\text{prior}}$  is visually very different from  $\mathcal{D}_t$ . Behavior retrieval also relies on access to the prior dataset when learning the new task, which may prevent its applicability to memory-limited scenarios where using only a pre-trained model is feasible. Finally, while we suspect that behavior retrieval can be naturally combined with many other directions of research towards more efficient imitation learning (like pre-trained representations, restricted action spaces, and large-scale pre-training), we have yet to experimentally verify that this is the case. Combining behavior retrieval with these methods and seeing if the benefits compound is an exciting direction for future research.

## ACKNOWLEDGMENTS

The authors thank Jonathan Yang and Sasha Khazatsky for their assistance with the real robot setup, Suneel Belkhale for assistance in implementing the LMP baseline, and the authors of the Bridge Dataset and RoboMimic suite for their help throughout the project. This work is funded by ONR grants N00014-22-1-2293 and N00014-22-1-2621.

## REFERENCES

- [1] Brenna D. Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009. ISSN 0921-8890. doi: <https://doi.org/10.1016/j.robot.2008.10.024>. URL <https://www.sciencedirect.com/science/article/pii/S0921889008001772>.
- [2] Mark Beliaev, Andy Shih, Stefano Ermon, Dorsa Sadigh, and Ramtin Pedarsani. Imitation learning by estimating expertise of demonstrators. In *39th International Conference on Machine Learning (ICML)*, 2022.
- [3] Suneel Belkhale and Dorsa Sadigh. Plato: Predicting latent affordances through object-centric play. *ArXiv*, abs/2203.05630, 2022.
- [4] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alexander Herzog, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Tomas Jackson, Sally Jesmonth, Nikhil J. Joshi, Ryan C. Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Kuang-Huei Lee, Sergey Levine, Yao Lu, Utsav Malla, Deeksha Manjunath, Igor Mordatch, Ofir Nachum, Carolina Parada, Jodilyn Peralta, Emily Perez, Karl Pertsch, Jornell Quiambao, Kanishka Rao, Michael S. Ryoo, Grecia Salazar, Pannag R. Sanketi, Kevin Sayed, Jaspiar Singh, Sumedh Anand Sontakke, Austin Stone, Clayton Tan, Huang Tran, Vincent Vanhoucke, Steve Vega, Quan Ho Vuong, F. Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and Brianna Zitkovich. Rt-1: Robotics transformer for real-world control at scale. *ArXiv*, abs/2212.06817, 2022.
- [5] Carlos Celemin, Rodrigo Pérez-Dattari, Eugenio Chisari, Giovanni Franzese, Leandro de Souza Rosa, Ravi Prakash, Zlatan Ajanović, Marta Ferraz, Abhinav Valada, and Jens Kober. Interactive imitation learning

- in robotics: A survey. *Foundations and Trends® in Robotics*, 10(1-2):1–197, 2022. ISSN 1935-8253. doi: 10.1561/23000000072. URL <http://dx.doi.org/10.1561/23000000072>.
- [6] Yiming Ding, Carlos Florensa, Pieter Abbeel, and Mariano Phielipp. Goal-conditioned imitation learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/c8d3a760ebab631565f8509d84b3b3f1-Paper.pdf>.
- [7] Yan Duan, Marcin Andrychowicz, Bradly Stadie, OpenAI Jonathan Ho, Jonas Schneider, Ilya Sutskever, Pieter Abbeel, and Wojciech Zaremba. One-shot imitation learning. *Advances in neural information processing systems*, 30, 2017.
- [8] Frederik Ebert, Chelsea Finn, Sudeep Dasari, Annie Xie, Alex Lee, and Sergey Levine. Visual foresight: Model-based deep reinforcement learning for vision-based robotic control. *arXiv:1812.00568*, 2018.
- [9] Frederik Ebert, Yanlai Yang, Karl Schmeckpeper, Bernadette Bucher, Georgios Georgakis, Kostas Daniilidis, Chelsea Finn, and Sergey Levine. Bridge data: Boosting generalization of robotic skills with cross-domain datasets. *ArXiv*, abs/2109.13396, 2021.
- [10] Benjamin Eysenbach, Tianjun Zhang, Ruslan Salakhutdinov, and Sergey Levine. Contrastive learning as goal-conditioned reinforcement learning. *ArXiv*, abs/2206.07568, 2022.
- [11] Chelsea Finn, Tianhe Yu, T. Zhang, P. Abbeel, and Sergey Levine. One-shot visual imitation learning via meta-learning. In *CoRL*, 2017.
- [12] Kanishk Gandhi, Siddharth Karamcheti, Madeline Liao, and Dorsa Sadigh. Eliciting compatible demonstrations for multi-human imitation learning. *ArXiv*, abs/2210.08073, 2022.
- [13] Anirudh Goyal, Abram L. Friesen, Andrea Banino, Théophane Weber, Nan Rosemary Ke, Adrià Puigdomènech Badia, Arthur Guez, Mehdi Mirza, Ksenia Konyushkova, Michal Valko, Simon Osindero, Timothy P. Lillicrap, Nicolas Manfred Otto Heess, and Charles Blundell. Retrieval-augmented reinforcement learning. *ArXiv*, abs/2202.08417, 2022.
- [14] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [15] Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. Imitation learning: A survey of learning methods. *ACM Comput. Surv.*, 50(2), apr 2017. ISSN 0360-0300. doi: 10.1145/3054912. URL <https://doi.org/10.1145/3054912>.
- [16] Stephen James and P. Abbeel. Coarse-to-fine q-attention with tree expansion. *ArXiv*, abs/2204.12471, 2022.
- [17] Stephen James and Andrew J. Davison. Q-attention: Enabling efficient learning for vision-based robotic manipulation. *IEEE Robotics and Automation Letters*, 7(2): 1612–1619, 2022. doi: 10.1109/LRA.2022.3140817.
- [18] Stephen James, Michael Bloesch, and Andrew J Davison. Task-embedded control networks for few-shot imitation learning. In *Conference on robot learning*, pages 783–795. PMLR, 2018.
- [19] Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey Levine, and Chelsea Finn. Bc-z: Zero-shot task generalization with robotic imitation learning. In *Conference on Robot Learning*, 2022.
- [20] Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey Levine, and Chelsea Finn. Bc-z: Zero-shot task generalization with robotic imitation learning. In Aleksandra Faust, David Hsu, and Gerhard Neumann, editors, *Proceedings of the 5th Conference on Robot Learning*, volume 164 of *Proceedings of Machine Learning Research*, pages 991–1002. PMLR, 08–11 Nov 2022. URL <https://proceedings.mlr.press/v164/jang22a.html>.
- [21] Edward Johns. Coarse-to-fine imitation learning: Robot manipulation from a single demonstration. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [22] Michael Kelly, Chelsea Sidrane, K. Driggs-Campbell, and Mykel J. Kochenderfer. Hg-dagger: Interactive imitation learning with human experts. *2019 International Conference on Robotics and Automation (ICRA)*, pages 8077–8083, 2018.
- [23] Aviral Kumar, Anika Singh, Frederik Ebert, Yanlai Yang, Chelsea Finn, and Sergey Levine. Pre-training for robots: Offline rl enables learning new tasks from a handful of trials. *ArXiv*, abs/2210.05178, 2022.
- [24] Michael Laskey, Jonathan Lee, Roy Fox, Anca D. Dragan, and Ken Goldberg. Dart: Noise injection for robust imitation learning. In *Conference on Robot Learning*, 2017.
- [25] Corey Lynch, Mohi Khansari, Ted Xiao, Vikash Kumar, Jonathan Tompson, Sergey Levine, and Pierre Sermanet. Learning latent plans from play. In *Conference on Robot Learning*, 2019.
- [26] Yecheng Jason Ma, Shagun Sodhani, Dinesh Jayaraman, Osbert Bastani, Vikash Kumar, and Amy Zhang. Vip: Towards universal visual reward and representation via value-implicit pre-training. *ArXiv*, abs/2210.00030, 2022.
- [27] Zhao Mandi, Fangchen Liu, Kimin Lee, and P. Abbeel. Towards more generalizable one-shot visual imitation learning. *2022 International Conference on Robotics and Automation (ICRA)*, pages 2434–2444, 2021.
- [28] Ajay Mandlekar, Danfei Xu, Roberto Mart'in-Mart'in, Yuke Zhu, Li Fei-Fei, and Silvio Savarese. Human-in-the-loop imitation learning using remote teleoperation. *ArXiv*, abs/2012.06733, 2020.
- [29] Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín.

- What matters in learning from offline human demonstrations for robot manipulation. In *arXiv preprint arXiv:2108.03298*, 2021.
- [30] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhi Gupta. R3m: A universal visual representation for robot manipulation. *ArXiv*, abs/2203.12601, 2022.
- [31] Soroush Nasiriany, Tian Gao, Ajay Mandlekar, and Yuke Zhu. Learning and retrieval from prior data for skill-based imitation learning. *ArXiv*, abs/2210.11435, 2022.
- [32] Simone Parisi, Aravind Rajeswaran, Senthil Purushwalkam, and Abhinav Kumar Gupta. The unsurprising effectiveness of pre-trained vision models for control. 2022.
- [33] Rom Parnichkun, Matthew N. Dailey, and Atsushi Yamashita. Reil: A framework for reinforced intervention-based imitation learning, 2022. URL <https://arxiv.org/abs/2203.15390>.
- [34] Dean Pomerleau. Alvin: An autonomous land vehicle in a neural network. In D.S. Touretzky, editor, *Proceedings of (NeurIPS) Neural Information Processing Systems*, pages 305 – 313. Morgan Kaufmann, December 1989.
- [35] Ilija Radosavovic, Tete Xiao, Stephen James, Pieter Abbeel, Jitendra Malik, and Trevor Darrell. Real-world robot learning with masked visual pre-training. *CoRL*, 2022.
- [36] Stéphane Ross, Geoffrey J. Gordon, and J. Andrew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *AISTATS*, 2011.
- [37] Stefan Schaal, Jan Peters, Jun Nakanishi, and Auke Ijspeert. Learning movement primitives. In Paolo Dario and Raja Chatila, editors, *Robotics Research. The Eleventh International Symposium*, pages 561–572. Berlin, Heidelberg, 2005. Springer Berlin Heidelberg. ISBN 978-3-540-31508-7.
- [38] Mariah L Schrum, Erin Hedlund-Botti, and Matthew Gombolay. Reciprocal MIND MELD: Improving learning from demonstration via personalized, reciprocal teaching. In *6th Annual Conference on Robot Learning*, 2022. URL [https://openreview.net/forum?id=f\\_XmiyZcsjL](https://openreview.net/forum?id=f_XmiyZcsjL).
- [39] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Cliport: What and where pathways for robotic manipulation. In *CoRL*, 2021.
- [40] Avi Singh, Huihan Liu, Gaoyue Zhou, Albert Yu, Nicholas Rhinehart, and Sergey Levine. Parrot: Data-driven behavioral priors for reinforcement learning. In *International Conference on Learning Representations*, 2021.
- [41] Jonathan Spencer, Sanjiban Choudhury, Matt Barnes, Matthew Schmitt, Mung Chiang, Peter Ramadge, and Siddhartha Srinivasa. Learning from interventions: Human-robot interaction as both explicit and implicit feedback. 07 2020. doi: 10.15607/RSS.2020.XVI.055.
- [42] Tete Xiao, Ilija Radosavovic, Trevor Darrell, and Jitendra Malik. Masked visual pre-training for motor control. 2022.
- [43] Tianhe Yu, Pieter Abbeel, Sergey Levine, and Chelsea Finn. One-shot hierarchical imitation learning of compound visuomotor tasks. *arXiv preprint arXiv:1810.11043*, 2018.
- [44] Andy Zeng, Pete Florence, Jonathan Tompson, Stefan Welker, Jonathan Chien, Maria Attarian, Travis Armstrong, Ivan Krasin, Dan Duong, Vikas Sindhwani, and Johnny Lee. Transporter networks: Rearranging the visual world for robotic manipulation. *Conference on Robot Learning (CoRL)*, 2020.
- [45] Songyuan Zhang, Zhangjie Cao, Dorsa Sadigh, and Yanan Sui. Confidence-aware imitation learning from demonstrations with varying optimality. In *Neural Information Processing Systems*, 2021.
- [46] Allan Zhou, Eric Jang, Daniel Kappler, Alex Herzog, Mohi Khansari, Paul Wohlhart, Yunfei Bai, Mrinal Kalakrishnan, Sergey Levine, and Chelsea Finn. Watch, try, learn: Meta-learning from demonstrations and reward. *arXiv preprint arXiv:1906.03352*, 2019.
- [47] Yuke Zhu, Josiah Wong, Ajay Mandlekar, Roberto Martín-Martín, Abhishek Joshi, Soroush Nasiriany, and Yifeng Zhu. robosuite: A modular simulation framework and benchmark for robot learning. In *arXiv preprint arXiv:2009.12293*, 2020.

## APPENDIX A MODEL DETAILS

### A. Data Modalities

We include a third-person camera view (84 x 84 x 3), an eye-in-hand camera view (84 x 84 x 3), end-effector position (3-dim), end-effector orientation (4-dim), and gripper state (2-dim). For state-only experiments in simulation, we use the provided object states and the same proprioception.

For real robot experiments, we use the end-effector position and a third-person camera view, either mounted facing the front of the robot (nut assembly task) or looking directly down at robot’s work area (tasks from the Bridge data)

### B. Policy Architecture

We used a memory-augmented GMM policy. Images are processed with independent ResNet-18 encoders with a modified spatial softmax pooling and last layer to give a 64-vector embedding for each image. These embeddings are then concatenated with the proprioceptive lowdim data before being encoded in a two-layer LSTM with a horizon of 10 and a hidden dimension size of 1000.

To compute the action, the LSTM-encoded state representation is passed through an MLP to yield the parameters of a GMM with 5 modes. The model is trained through log-likelihood. During inference, we sample from the GMM with a scaled-down variance. The policy is trained end-to-end, including the visual embedders.

### C. VAE Embedding Architecture

Similar to the policy modality encoder, the VAE embedder uses a ResNet-18 to encode all images, and concatenates the lowdim proprioceptive and/or object states to the image embeddings. In addition, the action is treated as a lowdim modality and concatenated to the image embeddings.

From this modality embedding, the VAE runs an MLP encoder consisting of two hidden layers (300, 400) and output a latent 128-vector. We use a fixed gaussian prior, and we use a  $\beta$  weight of 0.0001 on the KL loss.

During training, we reconstruct the input through the VAE embedding. Low-dimensional states are recreated directly using MLPs, and the images are recreated using a ResNet-18 structure with transposed convolutions and upsampling.

For the real robot tasks on Bridge data, we preloaded R3M [30] weights on the ResNet-18 encoders, which were helpful in getting the model to adapt to the visually diverse data.

### D. Tuning and Training

For the VAE Images, we train for 500k steps with batch size 50. For the VAE lowdim, we train for 500k steps with batch size 100. We can monitor the progress of VAE training by examining the quality of its image reconstruction. In practice, this was sufficient to tune the main hyperparameter  $\beta$ , although we also made use of plots like those in Figure 8.

The plots like Figure 8 also allowed us to estimate the best  $\delta$ , although in practice, the exact  $\delta$  is not critical, as long as it is not too high.

For policies on images, we train for 300k steps with batch size 16. We use random crop augmentation on the sim environments, and we also use color jitter for the real environments. For policies on state, we train for 600k steps with batch size 100.

For our method, all data, and goal-conditioned fine-tuned, we sample from two datasets (offline, task-specific) and so the batch sizes are doubled.

### E. Goal Conditioned

We construct the goal-conditioned baseline by treating the last state in the demonstration as an input to the policy. We train only on the offline data, and during inference, we take the last state from a trajectory sampled from  $\mathcal{D}_t$ , which we assume access to during test-time.

### F. Goal Conditioned Fine Tune

To fine-tune the goal-conditioned policy, we load the pre-trained goal-conditioned model and then train on a balanced batch of offline and task-specific data.

### G. Data Selection

We precompute the embeddings for the offline and task dataset, and then compute the batch negative L2 distance.

Because we use a memory-augmented policy, we implicitly assume that the past 10 steps of a selected transition are also relevant. This is not a very strong assumption, as any sequence of state-actions that get us to a task-relevant state should be considered by our task policy.

### H. Our Method on Interventions

The core of our method stays the same with the HG-Dagger experiments, except that the task-relevant data comes from the last 1000 transitions where the expert has intervened.

The expert policy is either a machine policy or a policy trained on a larger dataset of task data. Every step, we take the L2 distance of the apprentice policy action and the expert action. If the L2 distance is above a certain threshold, we take the expert action and mark the transition as an intervention. Otherwise, we take the apprentice policy action.

We pre-train the policy with the whole offline dataset. While this yields a policy that does not know which location to place the peg or the soda can, the pre-training gave the policy the ability to grasp objects and also gave it features that made it adapt relatively quickly to the interventions.

## APPENDIX B ENVIRONMENT DETAILS

### A. Generating the Offline and Expert Data in Sim

For the NutAssembly and Office simulation tasks, machine policies were used to generate the offline dataset. For NutAssembly, we use a hand-coded algorithm to grasp the Nut piece, rotate it to align the peg, and then plunge onto the peg. We add gaussian noise to the actions, and we randomly select the peg to insert onto.

For Office, we use the provided Roboverse machine policy [40]. We used four objects and two bins, totaling to eight separate pick-place tasks.

## B. Real Robot Setup

We used a Widowx 200 with 3DoF delta control running at 3Hz. The images come from a Logitech C920 camera secured in place during collection and testing. The testing location had consistent lighting during all trials.

Task-relevant data and the offline data (for the nut assembly environment) were collected using a PS3 controller. For consistency, only one experienced human demonstrator was used. The 10 task-relevant data took around 10 minutes to collect.

## C. Real NutAssembly Task Setup

This setup was constructed to mimic the simulation task. Two wooden coathanger pegs were secured into a wooden base and painted red and green. The nut was constructed from MDF Plywood and painted blue. An additional piece of wood was glued to the top of the plywood to allow for easier grasping. The hole has a  $\sim 4$ mm tolerance on the peg, and the peg has a larger head that makes it likely that a poorly aligned hole will fall to the side.

## D. Pickle Task Setup

This setup was constructed to mimic the pickle cup task found in the bridge data. The same red cup and a similar pickle was used in the task demonstration. The red cup was randomly placed in 20 cm x 8 cm region for every demonstration and evaluation. The pickle was placed in 10 cm x 4 cm region for every demonstration and evaluation. While the bridge data contains distractors and multiple colors of cups, the task demonstrations only contained the pickle and the red cup.

For a run to be counted as a success, the pickle must be completely inside the cup. We exclude the rare cases where the pickle rests on top of the cup.

## E. Real Soap Task Setup

This setup was constructed to mimic the sink tasks found in the bridge data. The blue sink is placed at roughly the same position as that of the Bridge Data, with a collection of toy kitchen items like cups, plates, and toy foods. The soap is placed in the corner of the dish rack, and the red plate target is placed randomly in the sink for every demonstration and evaluation. The distractor items are randomized for every trial as well.

For a run to be counted as a success, the soap must be on the red plate and standing up by itself. In the rare cases where the plastic gets stuck to the robot gripper, we redo the trial.

## F. Using the Bridge Data

For the pickle-cup task, we selected the tabletop manipulation trajectories from the bridge data, which consisted of manipulations on three different colors of tables with various objects. For the soap in sink task, we selected all of the trajectories that are performed on this particular sink.

For both tasks, we set up our environment to match that of the bridge data as closely as possible. This included approximating the position of the fixed camera, and we also applied an approximate transformation in the proprioceptive

data to match our setup. Because we were using only 3DoF control, we discarded the orientation in the proprioception and actions of the bridge data.

## APPENDIX C ADDITIONAL VISUALIZATIONS

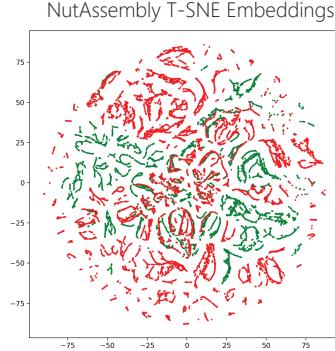


Fig. 12: **Visualizing the Embedding Space:** Using a T-SNE projection and access to ground-truth task labels (which determine the color of the point), we are able to see how trajectories follow paths in the projection, and how the different task types find natural clusterings.

Figure 12 shows a projection of the embeddings for 100 NutAssembly simulation trajectories. We see curves in the embedding plot which consist of individual or grouped transitions. Furthermore, we see how there are distinct clusterings of trajectory types, and we also see overlaps of trajectories, which we understand to be the shared sub-trajectories.

Figure 13 shows all of the separation plot for the six simulation configurations and the two real robot tasks with ground truth task labels. As can be seen, our VAE embeddings create a separation between relevant and irrelevant tasks. For the CanPick and NutAssembly (both sim and real), we see an overlapping of similarities where there is a shared grasping of the target object.

Figure 14 visualizes some of the selected and rejected data while running behavior retrieval on all of our six environments. As can be seen, the most similar transitions have correspondences to the desired task, and the least similar transitions typically contain irrelevant states. We note that for most environments, the grasping process is found in the top 25 transitions. This is desirable, as it shows that our method can make use of shared sub-trajectories.

## APPENDIX D ADDITIONAL EXPERIMENTS

### A. Can Behavior Retrieval Perform Task Stitching?

In our next experiment, we seek to understand if behavior retrieval is capable of performing task-stitching in the simulated WidowX office environment. That is, we consider the case where  $\mathcal{D}_{\text{prior}}$  only contains short-horizon behaviors (e.g. picking a single object and placing it in a container), but test on a target task  $\mathcal{D}_t$  which requires picking and placing two objects in sequence. We see in Table IV that in this setting behavior retrieval is the only method able to achieve a

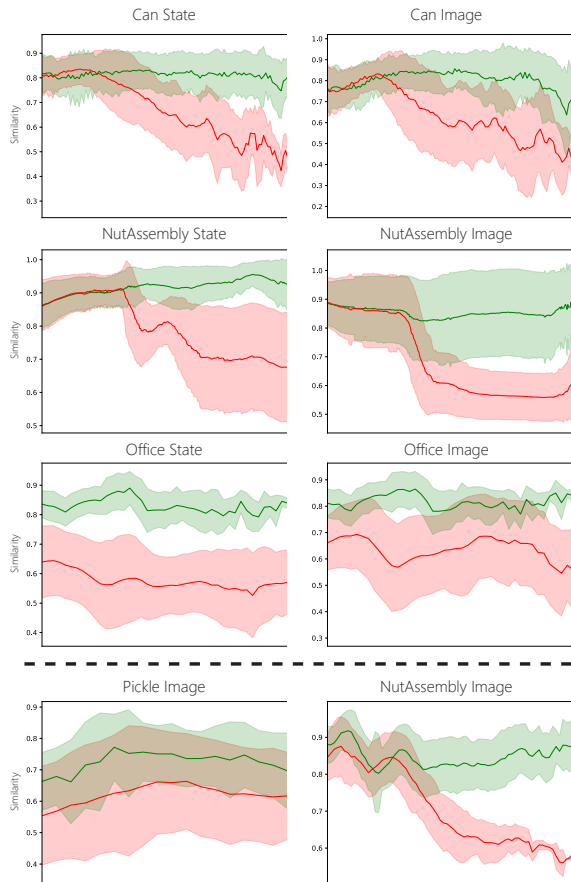


Fig. 13: **Embedding Separation Plots:** We show the embedding separation plots for all of the environments with ground truth task labels.

success rate above 10%. behavior retrieval is able to pull the useful transitions from each of the short horizon trajectories in the offline data, and jointly finetune on it with  $\mathcal{D}_t$  to more effectively tackle the sequential task.

$\mathcal{D}_t$	$(\mathcal{D}_{\text{prior}} + \mathcal{D}_t)$	GCBC $\mathcal{D}_{\text{prior}}$	GCBC $\mathcal{D}_{\text{prior}} + \text{FT } \mathcal{D}_t$	Ours
$5 \pm 3\%$	0	0	$3 \pm 2\%$	<b><math>20 \pm 2\%</math></b>

TABLE IV: **Robustness to Task Horizon.** We study if behavior retrieval can perform task-stitching, that is, leveraging short horizon trajectories in  $\mathcal{D}_{\text{prior}}$  to learn a sequential task in  $\mathcal{D}_t$ . We see that behavior retrieval is the only method to get a success rate that is not near zero, due to it’s ability to pull the relevant parts of short-horizon trajectories help learn the multi-step task.

### B. Can Behavior Retrieval Be Combined with Other Pre-Training Methods?

In this experiment, we investigate if behavior retrieval can be combined with other methods than vanilla behavior cloning for pre-training policies. Concretely, we combine behavior retrieval finetuning with the LMP approach proposed in Lynch et al. [25], which learns a skill embedding space by encoding trajectories which the policy is then conditioned on. In Table V we look at success rates in the Robosuite CanPick




Fig. 14: **Visualizing Retrieved Data for All Environments:** We show samples from the top 25 and the bottom 25 transitions when we apply our method to the offline data collected in each environment.

environment, comparing LMP trained on  $\mathcal{D}_{\text{prior}}$ , LMP trained on  $\mathcal{D}_{\text{prior}}$  and naively finetuned on  $\mathcal{D}_t$ , and LMP trained on  $\mathcal{D}_{\text{prior}}$  and finetuned with behavior retrieval on  $\mathcal{D}_{\text{prior}}$ . We observe that naively finetuning the LMP policy on  $\mathcal{D}_t$  helps performance slightly, but using behavior retrieval finetuning leads to more than a 10% improvement in success rate. This suggests that behavior retrieval as a finetuning method can be naturally combined with different policy training approaches.

LMP ( $\mathcal{D}_{\text{prior}}$ )	LMP ( $\mathcal{D}_{\text{prior}}$ ) + FT ( $\mathcal{D}_t$ )	LMP ( $\mathcal{D}_{\text{prior}}$ ) + FT (Ours)
$55 \pm 2\%$	$56 \pm 2\%$	<b><math>67 \pm 4\%</math></b>

TABLE V: **Combining Behavior Retrieval with LMP.** We study if behavior retrieval can be combined with other policy pre-trained methods like LMP [25]. We find that it can, producing over a 10% improvement in success rate.

### C. Using the Full Bridge Dataset.

In the original Pickle and Soap real robot tasks above, we set  $\mathcal{D}_{\text{prior}}$  to only be a subset of the full Bridge Dataset that was captured in visually similar scenes to the target environment. As a result,  $\mathcal{D}_{\text{prior}}$  only consisted of 10-20% of the full Bridge Dataset. To test if behavior retrieval can handle the full, more visually diverse dataset, we ran an additional experiment on the task of picking and placing a cup, where  $\mathcal{D}_{\text{prior}}$  consisted of the full Bridge Dataset, including mostly data from different scenes. We found that out of 20 trials, behavior retrieval had a **55%** success rate at the task, while  $\mathcal{D}_t$  only has a 30% success rate, and  $\mathcal{D}_{\text{prior}} + \mathcal{D}_t$  has a 35% success rate. Thus, even when only a small portion of a visually diverse  $\mathcal{D}_{\text{prior}}$  is

relevant to the task, with the correct choice of  $\delta$  (0.9 in this case), behavior retrieval can still provide significant benefits to performance.

#### D. Failure Mode Analysis

	Grasp	Intermediate	Alignment / Place
Ground Truth	52%	0%	58%
Ours	54%	2%	44%
Task-Specific Only	60%	9%	31%
All Data	22%	31%	47%
GC	83%	5%	12%
GC + Finetuned	15%	0%	85%

TABLE VI: **Failure Mode Analysis.** Averaged across all of our real robot tasks (40 trials each), we measure the percentage of each type of failure for each of the methods including our own. We observe that the distribution of failures is similar to Ground Truth, suggesting that with effective retrieval, methods tend to only fail at the challenging bottleneck regions of the task (grasping and placing).

For the real robot experiments, we kept track of the types and frequencies of failures on the tasks. We grouped these failures into three categories relevant to the three tasks we consider: grasp, intermediate, and alignment / place. Intermediate failures include anything between the grasping of the object and the final alignment process, like the premature release of the object.

We plot the percentage of each failure type for each method in Table VI, averaged across the three tasks and 40 trials.

Averaged across the three tasks, the failure distribution of behavior retrieval is comparable to the ground truth. Namely, most of the failures are at the bottleneck regions of grasping and placing. When all the data is included, the robot experiences more intermediate failures, as the offline dataset contains irrelevant and even adversarial demonstrations, like in the case of the two pegs in Nut Assembly.

Finally, with the goal-conditioned models, the significant distribution shift between the offline and task-specific demonstrations often yielded policies that failed to successfully grasp the object. The grasping is fixed with fine-tuning, although perhaps due to the complexity of the goal state, the robot struggles with the higher precision alignment part of the task.