

Offline Imitation Learning Through Graph Search and Retrieval

Zhao-Heng Yin and Pieter Abbeel
UC Berkeley EECS




Fig. 1: Collecting high-quality human demonstrations for imitation learning can be very difficult. Consider the problem of using a spoon, tying a rubber band, and picking up tiny objects with a tweezer. These tasks require very precise and accurate manipulator movement. A human operator can exhibit suboptimal behaviors and retry multiple times during a single demonstration. We propose a novel algorithm called GSR to learn proficient behavior from such suboptimal data through graph search and retrieval.

Abstract—Imitation learning is a powerful machine learning algorithm for a robot to acquire manipulation skills. Nevertheless, many real-world manipulation tasks involve precise and dexterous robot-object interactions, which make it difficult for humans to collect high-quality expert demonstrations. As a result, a robot has to learn skills from suboptimal demonstrations and unstructured interactions, which remains a key challenge. Existing works typically use offline deep reinforcement learning (RL) to solve this challenge, but in practice these algorithms are unstable and fragile due to the deadly triad issue. To overcome this problem, we propose GSR, a simple yet effective algorithm that learns from suboptimal demonstrations through Graph Search and Retrieval. We first use pretrained representation to organize the interaction experience into a graph and perform a graph search to calculate the values of different behaviors. Then, we apply a retrieval-based procedure to identify the best behavior (actions) on each state and use behavior cloning to learn that behavior. We evaluate our method in both simulation and real-world robotic manipulation

tasks with complex visual inputs, covering various precise and dexterous manipulation skills with objects of different physical properties. GSR can achieve a 10% to 30% higher success rate and over 30% higher proficiency compared to baselines. Our project page is at <https://zhaohengyin.github.io/gsr>.

I. INTRODUCTION

Imitation learning [36, 22] is a powerful approach to learning robots that has achieved great success in robotic manipulation in recent years [38, 2, 3, 48, 8]. It involves training robots to mimic human demonstrations, allowing them to acquire manipulation skills in complex environments. Despite the remarkable strides made so far, we notice that most existing works usually assume expert-level task demonstrations, while many real-world robotic manipulation tasks involve precise, dexterous manipulation with complex objects which are hard

for humans to provide high-quality demonstrations (Figure 1). For example, consider the task of making a robot manipulator grasp a tweezer to manipulate a tiny object, which is a common application scenario in a science laboratory. This task involves very fine-grained, precise control of gripping force and gripper movement. A human teleoperator can easily fail to provide successful demonstrations. Moreover, there usually exist suboptimal behaviors within a successful demonstration, such as retrying to grip the item if the first attempt fails. Naively using Behavior Cloning (BC) to imitate the collected suboptimal demonstration can result in incapable control policies that can reproduce failure behavior [28].

We need an algorithm that can learn from suboptimal human demonstrations to solve this problem. Typically, a suboptimal human demonstration is composed of multiple segments of different expertise levels, and the whole demonstration dataset contains a set of useful temporal segments that combined together can approximate the behavior of a good policy. The desired learning algorithm should be able to pick up those good segments and recover this desired policy, which is also known as *trajectory stitching* in reinforcement learning (RL) research [28]. In principle, it can be achieved by an offline deep-RL algorithm. However, in the application, previous research suggests that existing offline RL algorithms can usually struggle in realistic, long-horizon, pixel-based tasks [30, 40]. The simplest imitation learning algorithm like BC still performs the best among offline policy learning algorithms in many scenarios [30] and continues to be the standard practice in recent works.

To overcome the learning difficulties associated with typical offline RL, we propose GSR, a simple yet effective algorithm that learns from suboptimal demonstrations through Graph Search and Retrieval. Our key insight is to identify and stitch good behaviors directly through the use of graph representation and retrieval without bootstrapping deep RL networks. Specifically, we leverage pretrained representations to build a graph representation for the collected interaction data. This graph serves as a nonparametric world model, with which we can reason about whether a temporal segment is proficient and can lead to desired future outcomes effectively through a graph search. Then, for each state in the dataset, we can retrieve the most proficient behavior on that state and use it for imitation learning (Figure 2). Our algorithm can be viewed as a filtering procedure to preprocess the dataset for imitation learning and can be used with any kind of deep imitation learning models (architectures). As a direct approach that uses graph search rather than deep RL, our method enjoys high time efficiency. The whole preprocessing procedure typically consumes 10 to 30 minutes (depend on the dataset) overhead compared to BC, making it practical for real-world applications.

In the experiments, we test our method in both simulation and real-world robotic manipulation tasks of various visual and physical complexities, involving high-resolution, multiview camera observation, and precise robot-object interactions. Despite numerous challenges in both perception and action, our method can consistently improve baselines’ success rate

by 10% to 30% and proficiency by over 30%. We also provide various quantitative and qualitative analyses to show that our method is capable of identifying good behaviors in the dataset.

II. RELATED WORK

Learning from Suboptimal and Unstructured Data How to extract and learn good behaviors from suboptimal datasets is a central challenge in policy learning. When reward labels are available (e.g. task completion reward in goal-conditioned setup in this paper), offline deep Reinforcement Learning [27, 28, 16, 7, 23, 24, 32, 42, 26, 35] is one prevailing approach to address this problem. Most of the existing offline deep RL methods take the form of value-based learning. They learn value functions to estimate the goodness of actions with recursive Bellman updates and train policy to select good actions that lead to high values. However, deep RL is highly unstable and fragile due to deadly triad [45, 40], and people usually find that such (offline) RL algorithms struggle in application scenarios involving high-dimensional visual inputs and complex physical dynamics, and fail to outperform simple imitation learning algorithms like BC [30]. Our method is also based on value-based learning, but different from the existing methods, we directly use graph search and retrieval instead of prevailing deep RL updates to evaluate the value functions and perform policy optimization. This direct approach sidesteps the difficulty of bootstrapping deep value functions in complex environments [40].

When reward labels are not available, another line of work focuses on identifying proficient behaviors through other forms of feedback such as optimality labels, ranking, and preference [5, 4, 9, 43, 6, 46, 1, 31, 47, 19, 25]. This is an orthogonal research direction to the offline RL and our method. We believe our method can also benefit from these approaches by incorporating this additional information into the graph structure. This paper mainly considers a goal-conditioned setup and the proficiency can be defined through graph distances.

Graph and Retrieval in Policy Learning Our method is built upon graph search and retrieval, and there are many policy learning works incorporating these techniques. Some works propose to incorporate a graph into the policy for goal-conditioned problems [37, 14, 13, 44, 49]. During inference, these methods involve locating the agent on the generated graph and running a graph search to decide the next subgoal. In other words, the graph can be considered as a high-level map-based planner to guide a low-level goal-conditioned policy. Our method differs from these works in the following ways. First, we only use the graph to estimate the goodness of behavior in the offline phase and do not use it as a high-level planner during inference which can fail to generalize and act robustly (pick up inappropriate next goal in deployment) [13]. Besides, we relate graphs to extracting better policies in an offline robotic dataset. Finally, we study the use of pretrained representations in building the graph in application scenarios for the first time, while existing works build the graph by training a specialized reachability network on the collected data [37].

Recently, people have found retrieval a useful tool for policy learning, especially in the field of robotics. For example,

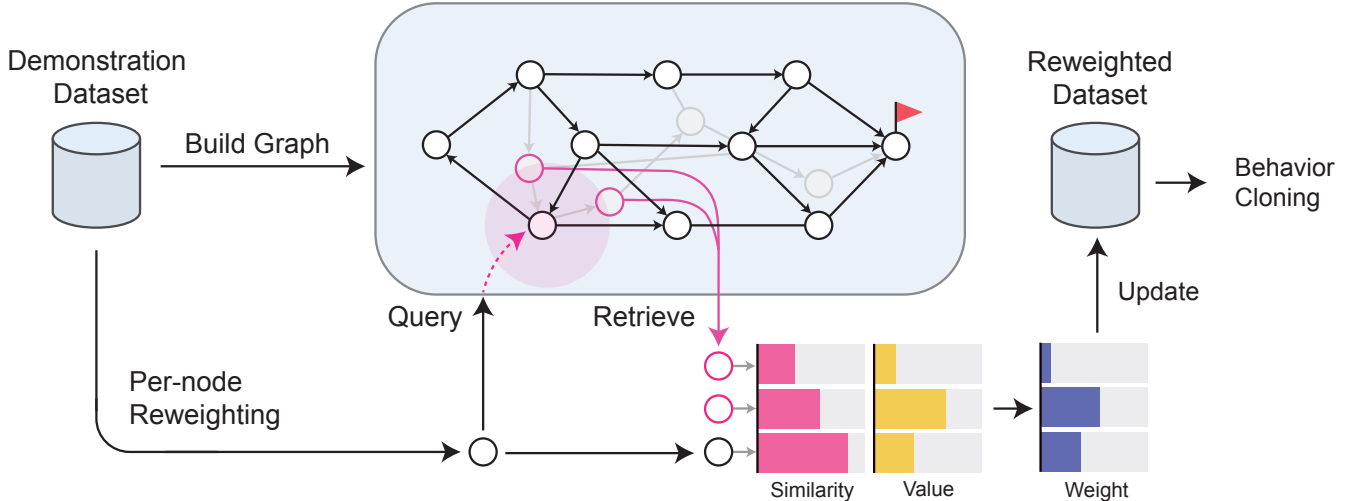


Fig. 2: Overview of our algorithm. We first build a graph to represent the demonstration dataset and run a graph search to evaluate the goodness (value) of each node. Then, we use a retrieval process to reassign weights to the transitions associated with each node. The weight reassignment is based on the similarity and goodness (value) score. The entire process will give more weight to the optimal behaviors around each node, resulting in a reweighted dataset. We use this reweighted dataset for behavior cloning.

VINN [34] finds that a K-Nearest Neighbor policy with BYOL representation [17] makes a power robotic manipulation policy. Lately, Palo et al. [11] use a retrieval process to find and execute relevant skill segments in the collected dataset to solve a robotic task. These works consider retrieval as part of the learned BC policy, while our method uses retrieval for sample weighting in offline policy learning. Perhaps the most related approach is Behavior Retrieval [12], which uses retrieval to find task-relevant transitions in a task-agnostic play dataset for imitation learning. We step further by using retrieval as a policy improvement operator, which is not studied by their work.

III. PRELIMINARIES

A. Problem Formulation

In this paper, we study an offline policy learning setup. We provide the robot agent with a human interaction dataset \mathcal{D} containing several task-relevant interaction trajectories. Each trajectory τ is a sequence of observations $o_{0:T}$ and corresponding actions $a_{0:T}$, i.e., $\tau = (o_0, a_0, o_1, a_1, \dots, o_T, a_T)$. For each trajectory $\tau \in \mathcal{D}$, the human teleoperator controls the robot to solve the task and achieves success in the end, i.e., the last observation of the trajectory is a task success example. However, the demonstrated behavior can be highly suboptimal, including numerous errors and misses in the execution process. The setup favors the data collection process: It allows the data generated through human practice to serve as a demonstration, provided that the task is ultimately accomplished. We expect the robot to learn the task-solving skill by learning from \mathcal{D} without using any online interaction samples.

B. Offline Policy Learning

Many offline policy learning algorithms usually involve training a policy network $\pi_\theta(a|o)$ to maximize the following weighted log-likelihood objective function:

$$\mathcal{L}(\theta) = \mathbb{E}_{(o,a) \sim \mathcal{D}} [w(o,a) \log \pi_\theta(a|o)], \quad (1)$$

where $w(o,a)$ is a predefined or learned weight function. For example, if we define $w(o,a) = 1$ for all $(o,a) \sim \mathcal{D}$, the objective corresponds to the behavior cloning. If we define $w(o,a) = \exp(A(o,a))$ where A is the advantage of taking action a at observation o , this corresponds to the policy extraction objective used in Advantage-Weighted Regression (AWR) [35] and Implicit Q-Learning (IQL) [24].

IV. POLICY LEARNING

BY GRAPH SEARCH AND RETRIEVAL

GSR is a simple method that computes each weighting coefficient $w(o,a)$ using graph search and retrieval. Compared to the previous offline RL-based approach, it does not involve advantage computation with an extra Q-learning procedure, which is known to be unstable for high-dimensional pixel-based problems. For clarity, we first derive our idea step-by-step from a tabular case example in section IV-A. We introduce the implementation details in the remaining sections.

A. Overview

For simplicity, let us first consider the problem of defining $w(o,a)$ in a tabular case where we want the agent to reach a certain goal observation o_g efficiently at the fewest cost after learning. Here, the cost function is the number of steps taken to reach o_g . Obviously, for any observation o , the agent should pick an action a whose resulting next observation (more

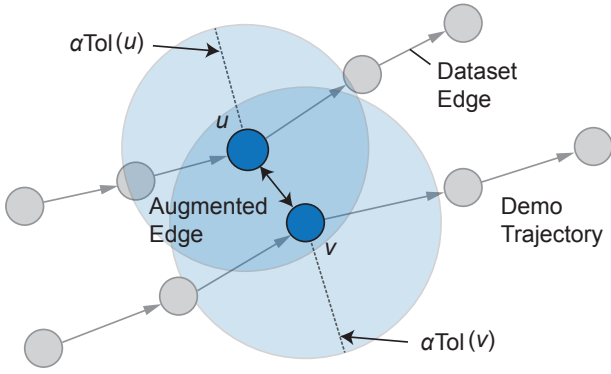


Fig. 3: Identifying connectivity. Augmented edge: We add a bidirectional edge between two nodes u and v if they both lie in the tolerance range of each other in the pretrained representation space. Dataset edge: It represents ground truth transitions on each demonstration trajectory.

precisely, underlying state) o_a has the shortest distance to the goal o_g . We denote the optimal action as

$$a^* = \arg \min_{a \in \mathcal{A}_o = \{a \in \mathcal{A} | (o, a) \in \mathcal{D}\}} d(o_a, o_g). \quad (2)$$

The distance function d is directly computable via a shortest path search. Then, we can define $w(o, a) = \mathbf{1}[a = a^*]$. In other words, we drop all the suboptimal goal-reaching actions from the dataset and only fit the remaining optimal ones. One may also use a softer version using $w(o, a) = \exp(-d(o_a, o_g))/Z$ to introduce stochasticity into the policy for future online fine-tuning, where $Z = \sum_{a \in \mathcal{A}_o} \exp(-d(o_a, o_g))$ is a normalizer. Nevertheless, the above formula is not useful in practice. For the real-world continuous control problem, each encountered observation is unique in the human demonstration dataset \mathcal{D} , and we have the following two problems: (1) Similar states on different data trajectories are disconnected, and as a result, the distance estimate can be very inaccurate. (2) The argmin operator has only one candidate action ($|\mathcal{A}_{o_i}| = |\{a_i\}| = 1$) and degenerates ($w(o_i, a_i) = 1$ everywhere). In this case, no policy improvement can be performed.

To address the first problem, we propose to identify and connect similar states in the dataset to form a better distance estimate in section IV-B. For the second problem, our insight is to create an approximated candidate set with nearest neighbor retrieval, and bias the weight towards relevant transitions that has higher value. We provide an overview of the proposed framework in Figure 2.

B. Graph Construction

The recent advance in large-scale representation pretraining provides us with powerful tools to identify similarities between different images. We can compare experiences from different trajectories in the representation space, and connect similar states to build a graph representation of the world. Concretely, we construct a graph $G(\mathcal{V}, \mathcal{E})$ to represent the whole dataset as follows.

Algorithm 1: GSR

- 1 [Optional] Finetune pretrained f_θ on \mathcal{D} ;
 - 2 Build graph $G(\mathcal{V}, \mathcal{E})$ using procedure in Section IV-B;
 - 3 Set $w[v] = 0$ for each $v \in \mathcal{V}$;
 - 4 Compute feature embedding f_v for each $v \in \mathcal{V}$;
 - 5 **for** $v \in \mathcal{V}$ **do**
 - 6 $\mathcal{N}(v) \leftarrow \{v\} \cup K$ -Nearest neighbors of v ;
 - 7 **for** $u \in \mathcal{N}(v)$ **do**
 - 8 $w[u] \leftarrow w[u] + \frac{\exp[\beta_1 S(u, v) + \beta_2 \bar{Q}(u)]}{\sum_{u' \in \mathcal{N}(v)} \exp[\beta_1 S(u', v) + \beta_2 \bar{Q}(u')]}$;
 - 9 Set data weight $w(o, a)$ using w :
 $w(o_{in:(i+1)n}, a_{in:(i+1)n}) \leftarrow w[v_i]$ for each $v_i \in \mathcal{V}$;
 - 10 Train policy π_θ by maximizing Equation 1;
-

1) *Building Vertex Set \mathcal{V}* : We define the vertex set \mathcal{V} by subsampling each trajectory, resulting in a subset of the dataset observations. Specifically, for each trajectory $o_{0:T}$, we add o_0, o_n, o_{2n}, \dots to the vertex set, and we also refer to them as v_0, v_1, \dots . The subsampling is simply used to make adjacent nodes on the data trajectory have enough visual difference (adjacent nodes can look identical otherwise). To specify the task objective conveniently, we also add a “virtual” goal vertex g to \mathcal{V} .

2) *Building Edge Set \mathcal{E}* : The *directed* edge set \mathcal{E} contains two kinds of edges: dataset edge and augmented edge. The dataset edge represents the transitions on each demo trajectory, serving as an approximation of the world dynamics. We connect each vertex to its next (future) vertex on the same trajectory (i.e. $(v_0, v_1), (v_1, v_2), \dots$). The last vertex of each trajectory corresponding to task success is connected to the goal vertex g . We set the weight of each dataset edge to 1.

So far, the vertices belonging to different trajectories are still disconnected from each other. We define augmented edges to bridge similar states across trajectories. To identify similar states, we use the off-shelf pretrained vision models to compute features for similarity computation. We first encode each vertex (observation) v_i with some pretrained encoder f_θ , giving us the representation $f_{v_i} = f_\theta(v_i)$. Then, we determine a quantity called neighborhood tolerance for each node v_i to decide connectivity, which we denote as $\text{Tol}(v_i)$. We compute this using the following formulas:

$$\text{Tol}_{\text{raw}}(v_i) = \min\{\text{sim}(f_{v_i}, f_{v_{i+1}}), \text{sim}(f_{v_i}, f_{v_{i-1}})\}, \quad (3)$$

$$\text{Tol}(v_i) = \frac{1}{M} \sum_{u \in \mathcal{N}_M(v_i)} \text{Tol}_{\text{raw}}(u). \quad (4)$$

Here, sim is a metric to measure the similarity between features. In our case, it is negative l_2 distance $\text{sim}(x, y) = -\|x - y\|_2$. In the first equation, we use the distance of each node to its one-step temporal neighbors as a rough neighborhood threshold. However, instead of using this single sample estimate, we average this using its M nearest neighbors in the feature space, denoted as $\mathcal{N}_M(v_i)$. This nearest-neighbor

set is computed using all the trajectories. Then, we connect two node v_i and v_j (bi-directionally) if

$$\text{sim}(f_{v_i}, f_{v_j}) > \alpha \max\{\text{Tol}(v_i), \text{Tol}(v_j)\}. \quad (5)$$

Hence, we require each node to fall within each other’s neighborhood range, multiplied by a hyperparameter α . The procedure is shown in Figure 3. In practice, α is set to a value around 1: the connected vertices should be roughly within a 1-step reach from each other’s view.

C. Policy Improvement with Retrieval

With a constructed graph, we can calculate the value (goodness) of each dataset transition from v_i to v_{i+1} (i.e. $(o_{in:(i+1)n}, a_{in:(i+1)n})$) using the distance on the graph. We define this as $Q(v_i) = Q(v_i, v_i \rightarrow v_{i+1}) = -1 - d(v_{i+1}, g)$. With this function, we bias the action toward the better ones **for each** v . To do this, we first retrieve K nearest neighbors of v in the feature space, denoted $\mathcal{N}_K(v)$. Then, we reallocate the default dataset weight (which is 1) assigned to v (which is 1) to v and these retrieved K nearest neighbors. Intuitively, we should put high weight to a retrieved vertex u only if (1) u is similar to v and (2) u has a high value. More formally, inspired by the idea of locally-weighted regression [10], the weight allocated to each vertex $u \in \mathcal{N}(v) = \{v\} \cup \mathcal{N}_K(v)$ should be proportional to $\text{Similarity}(u, v) \times \exp Q(u)$. Defining the first similarity term using an exponential form, we design the weight reallocation criterion at the vertex v as

$$w_{alloc}(u)|_v = \frac{\exp[\beta_1 S(u, v) + \beta_2 \tilde{Q}(u)]}{\sum_{u' \in \mathcal{N}(v)} \exp[\beta_1 S(u', v) + \beta_2 \tilde{Q}(u')]} \quad (6)$$

where

$$S(u, v) = \frac{1}{2} \left(\frac{1}{|\text{Tol}(u)|} + \frac{1}{|\text{Tol}(v)|} \right) \text{sim}(f_u, f_v) \leq 0 \quad (7)$$

is a normalized feature similarity function (with respect to 1-step neighborhood distance computed at u and v) to evaluate similarity between u and v . β_1 and β_2 are two temperature hyperparameters, controlling the strength of the similarity term and value term, respectively. In particular, when $\beta_1 \rightarrow \infty$ and $\beta_2 \rightarrow 0$, the weight to distribute concentrates on v and vanishes at the retrieved neighbors, and the policy learning reduces to the case of BC. Therefore, **our method smoothly interpolates between aggressive policy improvement and naive BC (no policy improvement)**.

D. Implementation and Time Complexity

We use R3M [33] as pretrained feature since it is pretrained with a contrastive objective, which we find can represent fine-grained movement well. In the task that involves multiple camera images, we calculate the feature embedding by concatenating the embedding of each image. We can also finetune the pretrained representation on the downstream task demo dataset [29] to make the retrieval more robust. We perform finetuning with a time-contrastive learning objective in simulated experiments and find that it improves the robustness to the selection of hyperparameters. We leave the discussion to




Fig. 4: The real world robot manipulation setup. We conduct experiments on a UR5 robot arm with Robotiq gripper. We use 3 workspace cameras with 256×256 RGB observation, highlighting the challenge in perception especially for our considered precise manipulation tasks.

the appendix. The pseudo-code of our method is summarized in Algorithm 1. The whole process is mainly composed of a graph search and single-round retrieval procedure, followed by BC. In particular, when we do not finetune the pretrained representation, the overhead computation compared to BC (Line 2-9) in general takes less than 10 minutes in our experiments, making our method convenient to use. The time complexity of this procedure is $O(N \log N + |E|)$, where N is the size of the dataset, and $|E|$ is the number of graph edges. Full implementation is in the appendix.

V. EXPERIMENTS

In experiments, we evaluate our method in both simulation and real world to validate its effectiveness in various robotic manipulation scenarios. We first study how much performance gain our method can achieve compared to the state-of-the-art imitation learning baseline. Then, through both quantitative and qualitative analysis, we show that our method identify and chain useful behaviors in the dataset to learn a robust policy. Finally, we present hyperparameter analysis to study design choices of our method.

A. Experiment Setup

1) *Simulation Experiments*: We use the image-based Robomimic benchmark [30] as our testbed in simulation, which provides several robotic manipulation tasks with FrankaPanda robot. Specifically, we use three tasks that contain human demonstrations of diverse qualities (Worse, Okay, Better):

- **Can Pick-and-Place** In this task, the robot is required to pick up a can on the table and place it into a slot. We use the following two dataset setups: Worse-Okay20

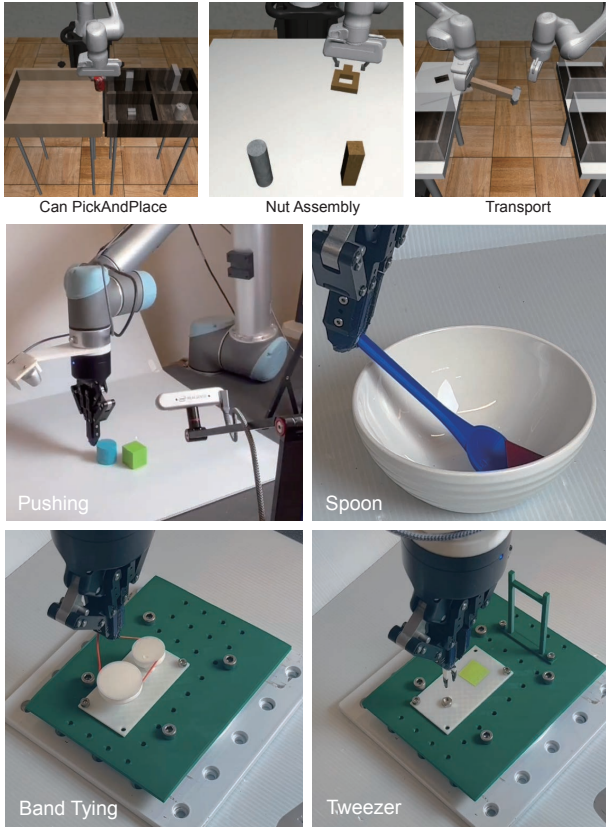


Fig. 5: Illustration of the used tasks. Above: Tasks from the robomimic benchmark. Bottom: Our real-world tasks.

and Worse-Better20 (the whole worse-human dataset with 20% data of the better-human dataset).

- **Nut Assembly** In this task, the robot is required to pick up a square hole and plug it into a square peg. This task requires precise and accurate manipulation. We use the following two dataset setup: Worse-Okay50 and Worse-Better50.
- **Transport** This is a long-horizon dual arm manipulation task. One robot arm is required to pick up a hammer in a case and transport it to another arm. Then, the other arm should place the hammer on the target location. We use the default MH setup provided by the benchmark.

The number of Okay/Better demos is based on the difficulty of the task. The result of using different demos is in the appendix.

2) *Real-world Experiments*: In the real-world, we design tasks of various complexities to highlight various aspects of our method. The robot workspace setup is shown in Figure 4.

- **Pushing** This is the simplest task in our task suite. The robot is required to push a blue cylinder toward a green cube on the table.
- **Spoon Scooping** In this task, the robot is required to pick up a spoon and use it to scoop a small object inside a slippery bowl. This task requires precise control of the spoon motion.
- **Tying a Rubber Band** This task is inspired by the NIST

board challenge. The robot should pick up a thin rubber band and tie it onto a pair of wheels. This task highlights the challenge of robust perception against partial occlusion and fine-grained manipulation.

- **Tweezer Manipulation** In this task, the robot needs to first pick up a tweezer that hangs on the shelf and then use the tweezer to pick up a nut and place it on a target platform. This task features both precise manipulation and long-horizon control.

We use the GELLO system [41] to collect human demonstrations, and the number of demonstrations and example rollouts can be found in the appendix. The Pushing and Spoon Scooping tasks use 2 cameras. The Rubber Band and Tweezer Manipulation tasks use 3 cameras. The camera images are all 256×256 RGB images. The policy also uses the robot arm’s proprioception information (joint positions) as input.

3) *Evaluation Metric*: To evaluate the performance of a trained policy, we use the following metrics.

- **Success rate (SR)** is defined as the number of task successes divided by the number of trials in evaluation.
- **Time-to-success (TTS)** is the average time a method takes to complete a task, which is computed only among the successful trials. A good method is expected to achieve a high success rate with low execution time.
- **Normalized Proficiency (NP)** is a metric we use for simulation. We calculate it by normalizing TTS with respect to the best human operator’s mean TTS ($NP=1.0$) and the worst human operator’s mean TTS ($NP=0.0$).

B. Baselines

In the experiments, we mainly compare our method to the following baselines:

- **Diffusion Policy (DP)**. DP [8] is a state-of-the-art behavior cloning baseline based on the diffusion model [21]. It is effective in overcoming the behavior multi-modality challenge in learning from human demonstrations in robotic manipulation.
- **Implicit-Q-Learning (IQL)**. IQL [24] is a strong offline RL baseline. We use R3M feature during Q-learning phase as in [20] and implement its policy model as the diffusion policy above.

Our method also uses DP as the policy model. For the policy network architecture, we use separate ResNet-18 [18] to process each input image, concatenate the output features with the proprioception information, and feed them into a Transformer-based denoiser network to generate the action. The whole policy network is trained from scratch. The denoising procedure is based on Denoising Diffusion Implicit Models (DDIM) [39] for fast deployment time inference. Specifically, in the simulation experiment, the last average pooling layer of ResNet-18 is replaced by a spatial softmax [15] as in previous works [30, 8].

C. Simulated Results

Table I shows our results in simulation. First, we observe that our method can reduce the average time to success across

TABLE I: Performance of evaluated methods on the **image-based** Robomimic benchmark. The results are averaged on 3 different seeds. For each seed, the evaluation is averaged over the last 5 checkpoints in training. Each evaluation takes 50 trials. We provide sufficient time for each policy to succeed (according to the recommendation of the benchmark). Our proposed method can outperform the baselines in both success rate (SR) and normalized proficiency (NP) across different tasks.

Task	Can (WO)		Can (WB)		NutAssembly (WO)		NutAssembly (WB)		Transport	
	SR(\uparrow)	NP(\uparrow)	SR(\uparrow)	NP(\uparrow)	SR(\uparrow)	NP(\uparrow)	SR(\uparrow)	NP(\uparrow)	SR(\uparrow)	NP(\uparrow)
DP	0.77 \pm 0.02	0.46 \pm 0.06	0.80 \pm 0.02	0.47 \pm 0.06	0.44 \pm 0.03	0.48 \pm 0.05	0.48 \pm 0.02	0.62 \pm 0.01	0.53 \pm 0.02	0.54 \pm 0.02
IQL-DP	0.77 \pm 0.01	0.57 \pm 0.04	0.81 \pm 0.01	0.52 \pm 0.03	0.40 \pm 0.02	0.48 \pm 0.08	0.46 \pm 0.02	0.33 \pm 0.03	0.53 \pm 0.02	0.56 \pm 0.02
Ours	0.85\pm0.02	0.86\pm0.02	0.87\pm0.01	0.76\pm0.03	0.53\pm0.01	1.16\pm0.03	0.57\pm0.04	0.88\pm0.03	0.67\pm0.03	0.63\pm0.02

TABLE II: Performance of evaluated methods on the real-world tasks. The results are averaged on the last checkpoint of 3 seeds. Each checkpoint evaluation takes 30 trials. We provide sufficient time for each policy to succeed. The time-to-success (TTS) metric is measured in seconds.

Task	Pushing		Spoon Scooping		Band Tying		Tweezer	
	SR(\uparrow)	TTS(\downarrow)	SR(\uparrow)	TTS(\downarrow)	SR(\uparrow)	TTS(\downarrow)	SR(\uparrow)	TTS(\downarrow)
DP	0.58 \pm 0.14	12.0 \pm 1.6	0.63 \pm 0.08	30.2 \pm 1.8	0.57 \pm 0.14	34.8 \pm 2.6	0.48 \pm 0.04	41.2 \pm 2.8
IQL-DP	0.82 \pm 0.07	9.4 \pm 0.8	0.67 \pm 0.08	29.4 \pm 1.3	0.57 \pm 0.06	30.6 \pm 2.4	0.46 \pm 0.12	37.0 \pm 2.0
Ours	0.95\pm0.03	7.8\pm0.4	0.81\pm0.05	24.0\pm1.9	0.86\pm0.05	22.5\pm1.4	0.76\pm0.06	29.7\pm2.5




Fig. 6: Dataset weight distribution plot (cumulative distribution function using sample statistics) v.s. different temperatures. As β_1 decreases or β_2 increases, we can see a larger contrast within samples (more samples away from weight 1). The plot is based on Can-WB20.

different tasks. Compared to previous methods, we can improve NP by 30%-50% on Can PickAndPlace and Nut Assembly, and can reach the best human operator’s time efficiency in some cases. Note that the best human operator’s proficiency is already high enough in these tasks, so we do not expect our algorithm’s NP to exceed the human operator’s performance ($>100\%$). The second finding is that our algorithm can also improve the success rate of state-of-the-art algorithms. The effect is also task-dependent. Interestingly, the improvement is the most significant (around 14%) in the most difficult long-horizon bimanual transport task, while for the other simpler tasks like NutAssembly, the improvement is milder (5% to 10%). We conjecture that this is because all the used demonstrations eventually lead to successes and the suboptimal behaviors in those datasets are mainly caused by slow motion or minor corrections. In this case, following the suboptimal trajectories will not harm the success rate much.

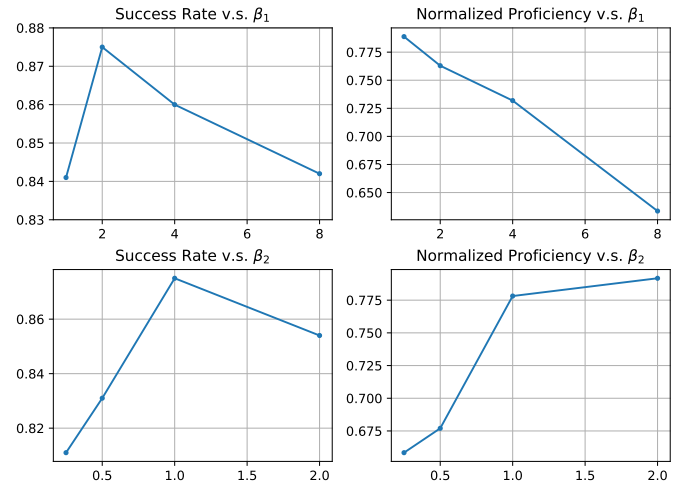


Fig. 7: Success rate and Normalized Proficiency metric under different β_1 and β_2 . The results are computed on the Can WB20. In general, we observe that decreasing β_1 and increasing β_2 can lead to improved success rate and proficiency for a wide range of β .

D. Real-world Results

As shown in Table II, our results in simulation also transferred to the real-world experiments well. We find that our method can achieve a success rate greater than 80% in the considered task and outperform all baselines in execution time. We observe that the baselines appear to repeat the failed human attempts during demo collection more frequently compared to our method. In some cases, they may eventually recover from failed behavior modes (by repeating subsequent retrying behavior demonstrated by a human) and complete the task.

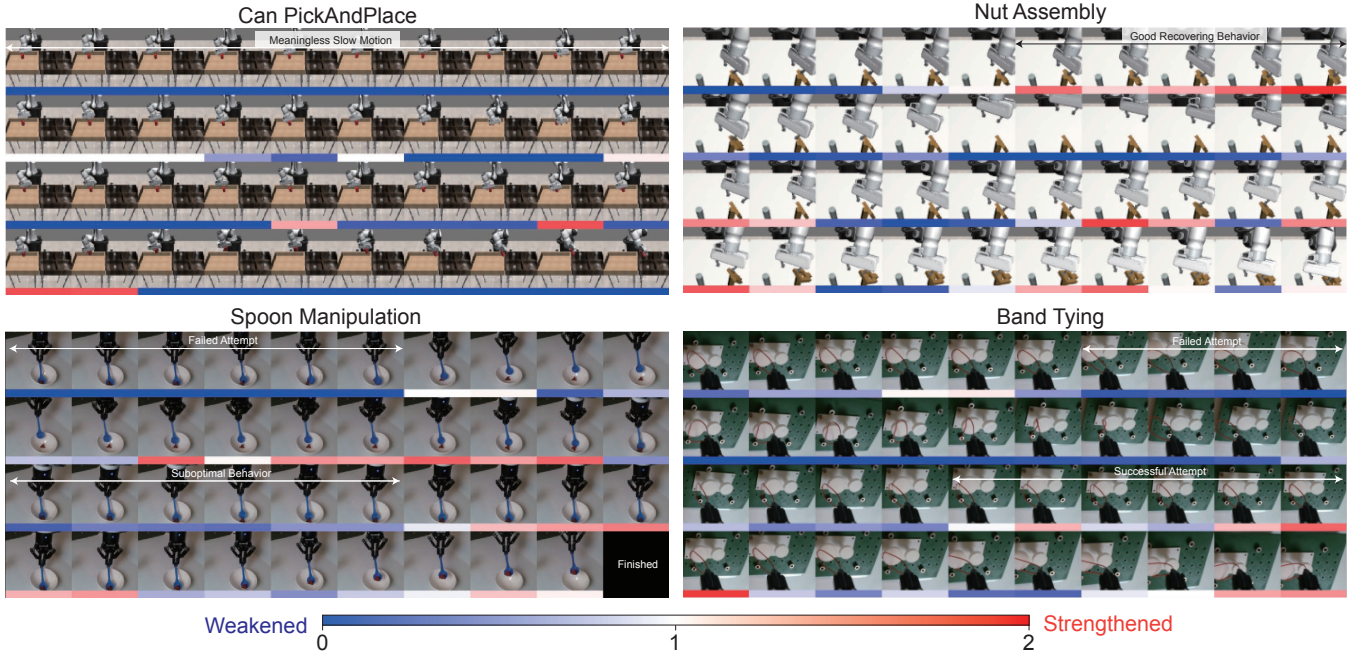


Fig. 8: Visualization of calculated weight w for transitions on a demonstration trajectory. We select 4 trajectories (part) from 4 different tasks. We use red to indicate strengthened transitions and blue to indicate weakened transitions. Frame Order: Left to Right, Top to Bottom.




Fig. 9: Relative Success rate and Normalized Proficiency improvement compared to baseline BC (DP) under different α . The results are averaged on the Can and NutAssembly.

However, in many cases, they will get stuck or go out of distribution, leading to a complete failure. In contrast, our method can identify the good behavior in the dataset, and avoid learning these failure behavior patterns (see next sections).

E. Qualitative Results: What is our method learning?

To understand what kind of behavior our policy is trying to learn, we perform qualitative analysis by visualizing our weight function w . We conduct the study in both simulation and real-world experiments. The result is shown in Figure 8. We use red and blue colors to indicate high (strengthened behavior) and low (weakened behavior) w respectively. Interestingly, we have the following findings: (1) All the temporal segments that lead to the failures are weakened and have low weights. (2). However, the segments that demonstrate failure-recovering behavior are not weakened because the best option for these states is to follow the dataset transitions. (3). The proficient temporal segments are usually strengthened as desired, and

the segments that correspond to slow movement are slightly weakened. We provide more visualization results in the appendix.

F. Hyperparameter Analysis

Having known that our method indeed strengthened desired behavior, in this section, we further study the effect of the main hyperparameters in our algorithm:

a) β_1 and β_2 : These temperatures control the policy improvement strength as discussed in section IV-C. We would like to understand how they affect weight distribution over the dataset and the resulting learning performance. We visualize the weight distribution in Figure 6. We can see that as β_1 decreases or β_2 increases, more data has lower weight and is rejected (zero weight), and we can observe a stronger contrast in the weight distribution, with fewer samples' weight center around 1. Then we plot the change in success rate and normalized proficiency in Figure 7. We find that both success rate and proficiency increase as we decrease β_1 or increase β_2 (β_1 decrease: using more dissimilar samples; β_2 increase: biasing more towards proficient samples) for a wide range of β s. However, when β_1 becomes too small or β_2 becomes too large, we find the success rate can drop. Intuitively, this is because less data is used for training. Although the policy can become more proficient in observed (narrower) regions, it may fail to solve the task once it deviates from the data distribution.

b) Tolerance α : α controls the range of tolerance and decide connections. We sweep it over different values to see how it impacts learning outcome (Figure 9). We find that both low and high α will affect the success rate and proficiency. A low alpha will kill most of the augmented edges, and can

provide an inaccurate distance estimate. On the other hand, a high α can connect nodes that are actually far away from each other and also harm distance estimate. Therefore, one needs to specify α in an appropriate range for best performance. We find that α can be tuned conveniently through some visualizations: we simply need to set α to different values examine whether the generated edges contain obvious errors (such as connecting two very dissimilar states), and select a large α whose error ratio is below a certain threshold.

VI. CONCLUSION

In this paper, we have presented GSR, a simple offline imitation learning method that can learn from suboptimal datasets through graph search and retrieval. We carried out experiments in both challenging simulation and real-world environments to demonstrate its effectiveness. In the future, there are several open problems to study. For example, how do we improve the quality of the representation to make our algorithm more powerful? In this paper, we considered demonstrations in the same context (i.e., the same workspace). However, it would be more desirable to have a task-invariant representation that can bridge experience from different contexts (backgrounds). Besides, in the context of large-scale pretraining, it would be interesting to use pretrained forward prediction models to generate more long-horizon connections, rather than a 1-step connection in this paper. This can leverage general knowledge from other relevant tasks for policy optimization. We hope that our work can inspire more powerful policy optimization methods.

ACKNOWLEDGEMENT

This work is supported by InnoHK Centre for Logistics Robotics and ONR MURI N00014-22-1-2773. The authors would like to thank Philipp Wu, Yide Shentu, and Toru Lin for the hardware support.

REFERENCES

- [1] Mark Beliaev, Andy Shih, Stefano Ermon, Dorsa Sadigh, and Ramtin Pedarsani. Imitation learning by estimating expertise of demonstrators. In *International Conference on Machine Learning*, pages 1732–1748. PMLR, 2022.
- [2] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022.
- [3] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023.
- [4] Daniel Brown, Wonjoon Goo, Prabhat Nagarajan, and Scott Niekum. Extrapolating beyond suboptimal demonstrations via inverse reinforcement learning from observations. In *International conference on machine learning*, pages 783–792. PMLR, 2019.
- [5] Benjamin Burchfiel, Carlo Tomasi, and Ronald Parr. Distance minimization for reward learning from scored trajectories. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.
- [6] Letian Chen, Rohan Paleja, and Matthew Gombolay. Learning from suboptimal demonstration via self-supervised reward regression. In *Conference on robot learning*, pages 1262–1277. PMLR, 2021.
- [7] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021.
- [8] Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *Robotics: Science and Systems*, 2023.
- [9] Sungjoon Choi, Kyungjae Lee, and Songhwai Oh. Robust learning from demonstrations with mixed qualities using leveraged gaussian processes. *IEEE Transactions on Robotics*, 35(3):564–576, 2019.
- [10] William S Cleveland and Susan J Devlin. Locally weighted regression: an approach to regression analysis by local fitting. *Journal of the American statistical association*, 83(403):596–610, 1988.
- [11] Norman Di Palo and Edward Johns. On the effectiveness of retrieval, alignment, and replay in manipulation. *IEEE Robotics and Automation Letters*, 2024.
- [12] Maximilian Du, Suraj Nair, Dorsa Sadigh, and Chelsea Finn. Behavior retrieval: Few-shot imitation learning by querying unlabeled datasets. *Robotics: Science and Systems*, 2023.
- [13] Scott Emmons, Ajay Jain, Misha Laskin, Thanard Kurutach, Pieter Abbeel, and Deepak Pathak. Sparse graphical memory for robust planning. *Advances in Neural Information Processing Systems*, 33:5251–5262, 2020.
- [14] Ben Eysenbach, Russ R Salakhutdinov, and Sergey Levine. Search on the replay buffer: Bridging planning and reinforcement learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- [15] Chelsea Finn, Xin Yu Tan, Yan Duan, Trevor Darrell, Sergey Levine, and Pieter Abbeel. Deep spatial autoencoders for visuomotor learning. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 512–519. IEEE, 2016.
- [16] Scott Fujimoto and Shixiang Shane Gu. A minimalist approach to offline reinforcement learning. *Advances in neural information processing systems*, 34:20132–20145, 2021.
- [17] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Dohersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural*

- information processing systems*, 33:21271–21284, 2020.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [19] Donald Joseph Hejna III and Dorsa Sadigh. Few-shot preference learning for human-in-the-loop rl. In *Conference on Robot Learning*, pages 2014–2025. PMLR, 2023.
- [20] Minh Heo, Youngwoon Lee, Doohyun Lee, and Joseph J Lim. Furniturebench: Reproducible real-world benchmark for long-horizon complex manipulation. *arXiv preprint arXiv:2305.12821*, 2023.
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [22] Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35, 2017.
- [23] Michael Janner, Qiyang Li, and Sergey Levine. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34:1273–1286, 2021.
- [24] Ilya Kostrikov, Ashvin Nair, and Sergey Levine. Offline reinforcement learning with implicit q-learning. *arXiv preprint arXiv:2110.06169*, 2021.
- [25] Sachit Kuhar, Shuo Cheng, Shivang Chopra, Matthew Bronars, and Danfei Xu. Learning to discern: Imitating heterogeneous human demonstrations with preference and representation learning. In *Conference on Robot Learning*, pages 1437–1449. PMLR, 2023.
- [26] Aviral Kumar, Justin Fu, Matthew Soh, George Tucker, and Sergey Levine. Stabilizing off-policy q-learning via bootstrapping error reduction. *Advances in Neural Information Processing Systems*, 32, 2019.
- [27] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 33:1179–1191, 2020.
- [28] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- [29] Arjun Majumdar, Karmesh Yadav, Sergio Arnaud, Jason Ma, Claire Chen, Sneha Silwal, Aryan Jain, Vincent-Pierre Berges, Tingfan Wu, Jay Vakil, et al. Where are we in the search for an artificial visual cortex for embodied intelligence? *Advances in Neural Information Processing Systems*, 36, 2024.
- [30] 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. *arXiv preprint arXiv:2108.03298*, 2021.
- [31] Vivek Myers, Erdem Biyik, Nima Anari, and Dorsa Sadigh. Learning multimodal rewards from rankings. In *Conference on Robot Learning*, pages 342–352. PMLR, 2022.
- [32] Ashvin Nair, Abhishek Gupta, Murtaza Dalal, and Sergey Levine. Awac: Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.
- [33] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3m: A universal visual representation for robot manipulation. *arXiv preprint arXiv:2203.12601*, 2022.
- [34] Jyothish Pari, Nur Muhammad Shafiullah, Sridhar Pandian Arunachalam, and Lerrel Pinto. The surprising effectiveness of representation learning for visual imitation. *arXiv preprint arXiv:2112.01511*, 2021.
- [35] Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*, 2019.
- [36] Dean A Pomerleau. Alvin: An autonomous land vehicle in a neural network. *Advances in neural information processing systems*, 1, 1988.
- [37] Nikolay Savinov, Alexey Dosovitskiy, and Vladlen Koltun. Semi-parametric topological memory for navigation. *arXiv preprint arXiv:1803.00653*, 2018.
- [38] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Perceiver-actor: A multi-task transformer for robotic manipulation. In *Conference on Robot Learning*, pages 785–799. PMLR, 2023.
- [39] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [40] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [41] Philipp Wu, Yide Shentu, Zhongke Yi, Xingyu Lin, and Pieter Abbeel. Gello: A general, low-cost, and intuitive teleoperation framework for robot manipulators. *arXiv preprint arXiv:2309.13037*, 2023.
- [42] Yifan Wu, George Tucker, and Ofir Nachum. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.
- [43] Haoran Xu, Xianyuan Zhan, Honglei Yin, and Huiling Qin. Discriminator-weighted offline imitation learning from suboptimal demonstrations. In *International Conference on Machine Learning*, pages 24725–24742. PMLR, 2022.
- [44] Lunjun Zhang, Ge Yang, and Bradley C Stadie. World model as a graph: Learning latent landmarks for planning. In *International Conference on Machine Learning*, pages 12611–12620. PMLR, 2021.
- [45] Shangdong Zhang, Hengshuai Yao, and Shimon Whiteson. Breaking the deadly triad with a target network. In *International Conference on Machine Learning*, pages 12621–12631. PMLR, 2021.
- [46] Songyuan Zhang, Zhangjie Cao, Dorsa Sadigh, and Yanan Sui. Confidence-aware imitation learning from

demonstrations with varying optimality. *Advances in Neural Information Processing Systems*, 34:12340–12350, 2021.

- [47] Wenjia Zhang, Haoran Xu, Haoyi Niu, Peng Cheng, Ming Li, Heming Zhang, Guyue Zhou, and Xianyuan Zhan. Discriminator-guided model-based offline imitation learning. In *Conference on Robot Learning*, pages 1266–1276. PMLR, 2023.
- [48] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.
- [49] Deyao Zhu, Li Erran Li, and Mohamed Elhoseiny. Value memory graph: A graph-structured world model for offline reinforcement learning. *arXiv preprint arXiv:2206.04384*, 2022.