

# Out-of-Distribution Recovery with Object-Centric Keypoint Inverse Policy For Visuomotor Imitation Learning

George Jiayuan Gao, Tianyu Li, Nadia Figueroa  
University of Pennsylvania  
{gegao,tianyuli,nadiafig}@seas.upenn.edu

**Abstract:** We propose an object-centric recovery policy framework to address the challenges of out-of-distribution (OOD) scenarios in visuomotor policy learning. Previous behavior cloning (BC) methods rely heavily on a large amount of labeled data coverage, failing in unfamiliar spatial states. Without relying on extra data collection, our approach learns a recovery policy constructed by an inverse policy inferred from object keypoint manifold gradient in the original training data. The recovery policy serves as a simple add-on to any base visuomotor BC policy, agnostic to a specific method, guiding the system back towards the training distribution to ensure task success even in OOD situations. We demonstrate the effectiveness of our object-centric framework in both simulation and real robot experiments, achieving an improvement of **77.7%** over the base policy in OOD. Project Website: <https://sites.google.com/view/ocr-penn>

**Keywords:** Visuomotor Policy, OOD Recovery, Object-Centric




Figure 1: **Object-Centric Recovery (OCR) on Bottle Pick and Place Task.** The base visuomotor policy (**Right**), trained on the bottle’s initial pose within the green-shaded region of the table, exhibits limited generalization when the bottle is initialized in the red-shaded region, which is considered out-of-distribution (OOD). (**Left**) showed the recovery policy using our OCR framework to recover from the red-shaded OOD region, returning the system to a region of high confidence for the base visuomotor policy, where it resumes control.

## 1 Introduction

Robot learning has achieved significant success in deploying Imitation Learning (IL) methods on real-world robotic systems [1]. One widely studied approach within IL is Behavior Cloning (BC), which has been explored extensively in recent work [2, 3, 4, 5, 6, 7]. BC methods enable learning control policies directly from demonstrations without the need for explicit environmental modeling, making the process relatively straightforward. However, despite producing promising results, BC is well-known for its susceptibility to the covariate shift problem [1]. This issue arises because traditional BC approaches depend heavily on large quantities of labeled data, which are often obtained through labor-intensive methods such as teleoperation or kinesthetic teaching. Consequently, BC may struggle to perform reliably in out-of-distribution (OOD) scenarios, where data is sparse or

noisy—reflecting a broader challenge faced in supervised learning. Addressing this issue typically requires either returning to laborious data collection or utilizing corrective mechanisms, such as guidance from human operators or reinforcement learning (RL) agents [8, 9, 10, 11], both of which impose additional deployment efforts on robotic systems.

To enjoy the benefits of strong performing BC policies in distribution (ID) settings while not requiring the human effort of collecting more data or the compute effort of running an RL step when OOD, in this work, we propose a recovery policy framework that brings the system back to the training distribution to ensure task success even when OOD. In particular, we focus on the key challenges of visuomotor policy learning by integrating a recovery policy constructed from the gradient of the training data manifold with a base visuomotor BC policy (e.g., a diffusion policy [2]). Inspired by the “Back to the Manifold” approach [12] the recovery policy guides the system back towards the training manifold, at which point the base policy resumes control. However, unlike [12], which focuses on recovering from OOD scenarios related to the robot’s state, our approach takes an object-centric perspective, specifically addressing OOD situations for task-relevant object states. We believe this object-centric approach significantly enhances the OOD recovery capabilities of visuomotor policies, leading to more robust learning for object manipulation tasks. Furthermore, our recovery framework is designed to be agnostic to the choice of base policy, allowing it to be seamlessly integrated with various BC implementations. This flexibility makes our method adaptable for future developments in imitation learning (IL). In this paper, we make the assumption that we have access to relevant object models. Also, we mainly consider the OOD cases in which the relevant object enters unfamiliar spatial regions.

The rest of the paper is organized as follows: Section 2 presents an overview of the existing works. Section 3 describes the problem formulation. Section 4 presents the object-centric recovery policy framework in detail, including its construction of the training data manifold and the keypoint inverse policy. In section 5, we demonstrate the effectiveness of our approach on several benchmarks, including both simulation and real robot experiments, showing that our recovery policy improves the performance when entering unfamiliar states. We also show that our method has the desired property for lifelong learning of visuomotor policies, improving the performance of OOD while not diminishing the in-distribution performance. Section 6 discusses the limitations and future directions.

## 2 Related Work

When deploying to the real world, vision-based IL could easily be initialized or moved to OOD situations, possibly due to bias in data collection and compounding errors. Deploying BC methods OOD could lead to unknown behavior in the low-data region. To address this, a well-known family of approaches is Data Aggregation, which gathers extra data from expert policies (usually provided by humans) through online interaction [13, 8, 10, 11]. However, performing such an online data collection procedure is an additional burden to the human when building a system. Our method tries to avoid additional online interaction and cumbersome data collection by squeezing as much information from the existing training data. The OOD problem has received much attention from the offline RL community. Offline RL suffers less compounding error than BC methods as it optimizes for long-term outcomes [14]. However, it could still struggle with distribution shifts like extrapolation errors due to limited data. To tackle the OOD problem, methods like [15, 16, 17] try to penalize actions that are far from the data. Moreover, several works [18, 19] also propose to recover back to the training data region, which indirectly shares a similar idea as our work. The paradigm of BC+RL has also been a popular choice for addressing the OOD problem [20, 9, 21]. Our approach takes on a similar direction for training a recovery policy, but instead, we use object-centric BC as the add-on for the base BC policy. Closely related to our work is [12], which introduces a vision-based OOD recovery policy by 1) learning an equivariant map that encodes visual observations into a latent space whose gradient corresponds to robot end-effector positions, and 2) following a gradient learned by a Mixture Density Network (MDN) [22]. Rather than recovering the robot action, our work focuses on recovering task-relevant objects and inducing the robot action. When dealing with objects, it could

be more general to utilize object-centric representation. Many works [6, 23, 24, 25, 26] demonstrate the success of object-centric representation policy learning. In our work, we use keypoints deriving from pose estimation [27] as the object representation for the recovery policy. There are existing works on using keypoint as the representation for policy learning [28, 29]. However, it is difficult for these methods to provide consistent 3D keypoints on the object across training and inference, which our method heavily relies on. We would like to address this for future work. Recent works also explore using Large Language Models to determine failure and perform recovery [30, 31].

### 3 Recovery Problem Formulation

Distribution shift for learning models is commonly quantified by measuring the Kullback-Leibler (KL) divergence between the distribution of observations during training and the distribution of observations encountered at test time [32, 33, 34]. This divergence reflects how much the test-time observations deviate from those seen during training, providing a metric to measure if a scenario we encountered is out-of-distribution. Formally, if the probability distribution of the training and testing observations is  $P(O)$  and  $Q(O)$  respectively, with  $O$  representing the set of observations, we say that the testing observation will be considered out-of-distribution (OOD) if,

$$\mathcal{D}_{KL}(P(O)||Q(O)) > \epsilon \quad (1)$$

is asserted to be true with some threshold  $\epsilon > 0$ .

We formulate our visuomotor policy interaction with the environment as a Partially Observable Markov Decision Model (POMDP) [35]. We describe this POMDP by the tuple  $(S, A, O, T, E)$ , where  $s \in S$  is the set of environmental states, which are directly observable,  $a \in A$  is the set of robot actions, and  $o \in O$  is the set of visual observations. The transition function  $T : S \times A \rightarrow S$  dictates how the unobservable state changes when robot actions are performed, and the emission function  $E : S \rightarrow O$  is a surjective function that determines the visual observations given states.

Given this formulation, we can reformulate the KL Divergence OOD metric as follows,

$$\mathcal{D}_{KL}(P(E(S))||Q(E(S))) > \epsilon. \quad (2)$$

Hence, fundamentally, given an observation-level out-of-distribution scenario, if the environmental state variables are recovered back into the training distribution, the observations will also be recovered back into distribution. However, the recovery of *all* state variables is difficult to tackle all at once under the imitation learning framework, which typically has access to only task-relevant demonstrations. Therefore, for this work, we specifically focus on the recovery of task-relevant objects in manipulation tasks. Unlike previous data aggregation or reinforcement learning approaches, we aim for our recovery framework to exclusively leverage the training demonstrations of the base policy and not require any additional policy-related data collection.

## 4 Method

We present our approach in augmenting a *base* policy trained via Behavior Cloning (BC) by incorporating an *object-centric* recovery strategy, which enables task-relevant objects to return to its Euclidean training manifold where the base BC policy functions at its best. For our work, we will assume task-relevant objects in the scene are rigid and non-deformable.

To achieve this, we introduce the Object-Centric Recovery (OCR) framework, as illustrated in Figure 2. We first explicitly model the distribution of objects keypoints in the training dataset with a Gaussian Mixture Model (GMM) [36]  $p(\rho_{k,t}^{(i)}|\theta_k) = \sum_{m=1}^M \lambda_{k,m} \mathcal{N}_{\theta_k}(\rho_{k,t}^{(i)}|\mu_{k,m}, \Sigma_{k,m})$  where  $\rho_{k,t}^{(i)}$  are keypoints in the dataset and  $\theta_k = \{(\lambda_{k,m}, \mu_{k,m}, \Sigma_{k,m})\}_{m=1}^M$  parameters of the GMM (Section 4.2.1, 4.2.2). At test time, we evaluate the gradient  $\nabla p(\rho_k^{test}|\theta_k)$  to obtain the object-recovery vectors, which we use to plan for an object-recovery trajectory  $\zeta_{rec}^L$  (Section 4.2.4). We then translate this trajectory into robot actions via a Keypoint Inverse Policy  $\pi_{inv}$  that is trained using the base dataset (Section 4.2.3). Lastly, Section 4.2.5 describes how the base policy interacts with the recovery policy to become the OCR *joint policy*.

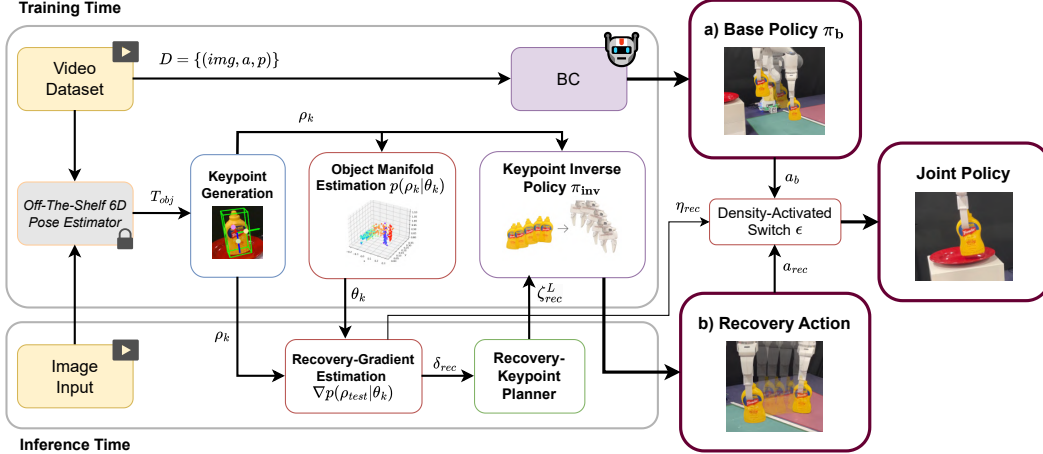


Figure 2: **Object Centric Recovery (OCR) Framework.** The OCR Framework augments a base policy  $\pi_b$ , trained via BC, by returning task-relevant objects to their training manifold, where the base policy takes over. First, we model the distribution of object keypoints in the training data using a Gaussian Mixture Model (GMM). At test time, we compute the gradient of the GMM to derive object-recovery vectors, which are used to plan a recovery trajectory. This trajectory is then converted into robot actions through a Keypoint Inverse Policy  $\pi_{inv}$ , trained *solely* on the base dataset. Finally, the base policy and the recovery policy are combined into a *joint policy*, allowing seamless interaction between recovery and task execution.

#### 4.1 Base BC Policy

Our formulation considers a generic visuomotor policy that outputs future actions based on past visual observations as the base BC policy. We consider such a liberal formulation to demonstrate that our framework can work alongside any variations of BC policy. Formally, we define a typical visuomotor policy training dataset as  $\mathbf{D}_b = \{\mathbf{d}_b^{(i)}\}_{i=1}^N$ , where each episode  $\mathbf{d}_b^{(i)} = \{(\mathbf{o}_t^{(i)}, \mathbf{a}_t^{(i)}, \mathbf{p}_t^{(i)})\}_{t=1}^T$  consists of the observations  $\mathbf{o}_t^{(i)}$ , robot actions  $\mathbf{a}_t^{(i)}$ , and robot proprioception  $\mathbf{p}_t^{(i)}$  at time step  $t$ . Then, under the imitation learning framework, a base visuomotor policy  $\pi_b$  that is parameterized by  $\phi_b$  is learning by optimizing the following behavior cloning objective:

$$\pi_b^* = \arg \min_{\theta_b} E_{(\mathbf{o}, \mathbf{a}, \mathbf{p}) \sim \mathbf{D}_b} [\mathcal{L}(\pi_b(\mathbf{o}, \mathbf{p}), \mathbf{a})] \quad (3)$$

Where the loss function  $\mathcal{L}$  is typically Cross-Entropy Loss or Mean-Squared Error.

#### 4.2 Object-Centric Recovery Policy

##### 4.2.1 Keypoint Generation

We choose to use artificial object keypoints to represent object poses for studying object-centric recovery, as keypoints allow us to tightly couple the position and orientation of the object, facilitating a more accurate estimation of its distribution during training.

We consider the same visuomotor policy training dataset formulation  $\mathbf{D}_b = \{\mathbf{d}_b^{(i)}\}_{i=1}^N$ , where each episode  $\mathbf{d}_b^{(i)} = \{(\mathbf{o}_t^{(i)}, \mathbf{a}_t^{(i)}, \mathbf{p}_t^{(i)})\}_{t=1}^T$  consists of the observations  $\mathbf{o}_t^{(i)}$ , robot actions  $\mathbf{a}_t^{(i)}$ , and robot proprioception  $\mathbf{p}_t^{(i)}$  at time step  $t$ . To extract object poses from these visuomotor datasets, we employ off-the-shelf object pose estimators [37, 38, 27] to transform each observation frame  $\mathbf{o}_t^{(i)}$  into the object pose  $\mathbf{T}_{obj,t}^{(i)}$ . Next, we define an arbitrary set of keypoints  $\mathbf{P} = \{p_k\}_{k=1}^n$ , where each  $p_k \in \mathbb{R}^d$ . For each keypoint  $k$  at time step  $t$  in demonstration  $i$ , we compute the transformed keypoints  $\rho_{k,t}^{(i)} = h^{-1}(\mathbf{T}_{obj,t}^{(i)} h(p_k))$ , where  $h$  represents the function that converts points into homogeneous coordinates. The transformed keypoint  $\rho_t^{(i)} = \{\rho_{k,t}^{(i)}\}_{k=1}^n$  then serves as the keypoint representation of the object's current pose. Thus, using  $\mathbf{D}_b$ , we create a new dataset that will be used for recovery  $\mathbf{D}_{rec} = \{\mathbf{d}_{rec}^{(i)}\}_{i=1}^N$ , where each episode  $\mathbf{d}_{rec}^{(i)} = \{(\rho_t^{(i)}, \mathbf{T}_{obj,t}^{(i)}, \mathbf{a}_t^{(i)}, \mathbf{p}_t^{(i)})\}_{t=1}^T$  consists of the keypoints, object poses, robot actions, and robot proprioception at each time step.

#### 4.2.2 Object Manifold Estimation

To estimate the manifold of the object distribution in the training dataset, we fit a Gaussian Mixture Model (GMM) [36] on each keypoint using its positions across every time step in every demonstration. Specifically, given dataset  $\mathbf{D}_{kp,k} = \left\{ \left\{ (\rho_{k,t}^{(i)}) \right\}_{t=1}^T \right\}_{i=1}^N$  consisting of one object keypoint  $k$  across all time step  $t$  in every demonstration  $i$ , we model the probability of each  $\rho_{k,t}^{(i)}$  as a weighted sum of  $M$  Gaussian distributions:

$$p(\rho_{k,t}^{(i)}|\theta_k) = \sum_{m=1}^M \lambda_{k,m} \mathcal{N}_{\theta_k}(\rho_{k,t}^{(i)}|\mu_{k,m}, \Sigma_{k,m}), \quad (4)$$

where  $\lambda_{k,m}$  is the mixing coefficient of keypoint  $k$  for the  $m$ -th Gaussian,  $\mathcal{N}(\rho_{k,t}^{(i)}|\mu_{k,m}, \Sigma_{k,m})$  is the Gaussian probability density function of keypoint  $k$  for the  $m$ -th component with mean  $\mu_{k,m}$  and covariance  $\Sigma_{k,m}$ , and  $\theta_k = \{(\lambda_{k,m}, \mu_{k,m}, \Sigma_{k,m})\}_{m=1}^M$  are the parameters of the model that estimates the distribution of keypoint  $k$ . To fit this GMM, we used Expectation-Maximization [36] to maximize the likelihood estimation of the model on the data. Computing a GMM for all  $n$  keypoints would result in parameters  $\Theta = \{\theta_k\}_{k=1}^n$  that collectively estimate the probability distribution of the keypoints of the object.

#### 4.2.3 Keypoint Inverse Policy

To facilitate object manipulation for recovery, we propose the use of an **Keypoint Inverse Policy**  $\pi_{inv}$ , which is designed to translate a sequence of object-keypoint trajectory along with the robot’s current state, into the corresponding robot actions necessary to execute those object motions effectively. Formally, if we define  $K$  to be the set of object keypoints observed, and  $P \subseteq S$  to be the set of robot proprioception states,  $\pi_{inv} : K^L \times P \rightarrow A^L$ , where  $L$  is the observation length. We already have access to the dataset of object keypoints, pose, and action tuples  $\mathbf{D}_{rec} = \left\{ \left\{ (\varrho_t^{(i)}, \mathbf{T}_{obj,t}^{(i)}, \mathbf{a}_t^{(i)}, \mathbf{p}_t^{(i)}) \right\}_{t=1}^T \right\}_{i=1}^N$  that we obtained in Section 4.2.1, which we could utilize to directly train  $\pi_{inv}$  with the learning objective:

$$\pi_{inv}^* = \arg \min_{\theta_{inv}} E_{\left( \left\{ \varrho_t^{(i)} \right\}_{t=j}^{j+L}, \left\{ \mathbf{a}_t^{(i)} \right\}_{t=j}^{j+L}, \mathbf{p}_j^{(i)} \right) \sim \mathbf{D}_{rec}} \left[ \mathcal{L} \left( \pi_{inv} \left( \left\{ \varrho_t^{(i)} \right\}_{t=j}^{j+L}, \mathbf{p}_j^{(i)} \right), \left\{ \mathbf{a}_t^{(i)} \right\}_{t=j}^{j+L} \right) \right] \quad (5)$$

where sequences of length  $L$  are pulled from the keypoint and action datasets to form  $\{\varrho_t^{(i)}\}_{t=j}^{j+L}$  and  $\{\mathbf{a}_t^{(i)}\}_{t=j}^{j+L}$ , and the initial proprioception of the sequence  $\mathbf{p}_j^{(i)}$ . However, by training this way directly, we will still run into the same issue of distribution shift, having no keypoints-to-action coverage on the OOD regions to generate properly useful manipulation outputs. To alleviate this, we propose the use of the initial object pose  $\mathbf{T}_{obj,t}^{(i)}$  to “zero-out” the data sequence. Specifically, instead of using the original sequence, we use the initial object pose of each sequence  $\mathbf{T}_{obj,j}^{(i)}$  to modify the sequence into  $\{(\mathbf{T}_{obj,j}^{(i)})^{-1} \varrho_t^{(i)}\}_{t=j}^{j+L}$ ,  $\{(\mathbf{T}_{obj,j}^{(i)})^{-1} \mathbf{a}_t^{(i)}\}_{t=j}^{j+L}$ , and  $(\mathbf{T}_{obj,j}^{(i)})^{-1} \mathbf{p}_j^{(i)}$ . For simplicity, we will name these quantities  $\varrho_{zero}^L$ ,  $\mathbf{a}_{zero}^L$ ,  $\mathbf{p}_{zero}$  respectively. Hence, the learning objective becomes:

$$\pi_{inv}^* = \arg \min_{\theta_{inv}} E_{\left( \varrho_{zero}^L, \mathbf{a}_{zero}^L, \mathbf{p}_{zero} \right) \sim \mathbf{D}_{rec}} \left[ \mathcal{L} \left( \pi_{inv} \left( \varrho_{zero}^L, \mathbf{p}_{zero} \right), \mathbf{a}_{zero}^L \right) \right] \quad (6)$$

In other words, the keypoint inverse policy only needs to learn to output robot actions from object keypoint trajectories that initialize from the identity frame. At test time, to carry out an object motion, we input the desired keypoint trajectory in the current object frame, and the policy outputs the corresponding robot action in that same frame. To execute the robot action, we then transform the action from the object frame to the robot frame. This way, regardless of the object and robot end-effector’s true pose in Euclidean space, OOD or ID, as long as we have access to the current object pose, we can output robot actions that are useful for manipulation. In addition, we can think of this as a way of “compressing” the input domain of the keypoint inverse policy based on the information available to us (object pose), making the learning problem extremely data efficient.

#### 4.2.4 Desired Object-Recovery Motion

**Object-Recovery Vectors.** At test time, we obtain the explicit current pose of the recovery object via pose estimation and generate keypoints using the same methodology as described in Section

3.2.1, obtaining  $\varrho^{test} = \{\rho_k^{test}\}_{k=1}^n$ . For each keypoint, we build a computation graph of the probability density function (pdf) of the GMM with parameters  $\theta_k$  to output the probability density  $\eta_k^{test}$  with respect to  $\rho_k^{test}$ . With this, we used automatic differentiation to output the gradient vector  $\delta_k^{test} = \nabla p(\rho_k^{test}|\theta_k)$ . However, the norm of this gradient vector  $\|\delta_k^{test}\|$  is strictly non-negative and increases as  $\rho_k^{test}$  approaches regions of increasingly higher density parameterized by the GMM with parameters  $\theta_k$ , which is in contrast with how we want recovery to take place - to approach recovery faster when the object is further away, and slower when the object is closer. To solve this, we modify the magnitude of  $\delta_k^{test}$  by a monotonically decreasing function, the *parameterized negative exponential* function, which we define as  $q(x) = e^{-\frac{\phi-x}{\eta}}$ . Thus, the modified recovery gradient is

$$\delta_k^{mod} = q(\|\delta_k^{test}\|) \frac{\delta_k^{test}}{\|\delta_k^{test}\|} \quad (7)$$

Since the recovery gradient might differ for each keypoint  $k$ , we will use the mean recovery gradient as our final object recovery, given by  $\delta_{rec} = \sum_{k=1}^n \frac{\delta_k^{mod}}{n}$ . Similarly, we take the mean of density to gain a single value for the current object density  $\eta_{rec} = \sum_{k=1}^n \frac{\eta_k^{test}}{n}$ . Hence, at each time step at test time, we output the object recovery tuple  $(\delta_{rec}, \eta_{rec})$ .

**Recovery Keypoint Planner.** From the object recovery vector  $\delta_{rec}$ , we can generate a naive recovery keypoint trajectory like so:

$$[\{t\alpha\delta_{rec} + \rho_k^{test}\}_{k=1}^n]_{t=1}^L \quad (8)$$

where  $\alpha$  is a scaling hyperparameter that we can tune at test time to optimize for the trajectory step size. However, this formulation does not take into account the feasibility of executing such a trajectory, which is paramount in ensuring the quality of the recovery. To this end, we propose a heuristic planner, using the distance of the position between the robot end-effector and the object pose as a heuristic for how much “delay” is added to the object trajectory before it starts moving, thus providing the robot with enough time to approach the object for manipulation. Specifically, we will define a maximum and a minimum distance where the object can be effectively manipulated which we denote as  $d_{max}$  and  $d_{min}$ , which can be tuned easily at test time. Then, we simply fit a linear function between points  $(d_{min}, L)$  and  $(d_{max}, 0)$ , and clip the range between  $[L, 0]$ . Formally, if we denote the norm between the position of the end-effector and object as  $d_{pos}$ , then the delay function is written as  $df(d_{pos}) = \min(\max(0, \lfloor \frac{-L}{d_{max}-d_{min}} * (d_{pos} - d_{min}) + L \rfloor), L)$ . Thus, our proposed keypoint recovery trajectory is expressed as:

$$\zeta_{rec}^L = [\{\max(0, t - df(d_{pos}))\alpha\delta_{rec} + \rho_k^{test}\}_{k=1}^n]_{t=1}^L \quad (9)$$

#### 4.2.5 Final Joint Policy

We join our base policy and recovery action as a *Joint Policy* through a density activation switch. Specifically, after computing the mean keypoint density  $\eta_{rec}$ , we use a tunable hyperparameter  $\epsilon_{rec}$  to define the threshold for distinguishing between out-of-distribution and in-distribution scenarios. If the scenario is classified as OOD, the recovery pipeline is activated; otherwise, if classified as ID, the base policy will proceed with the standard BC process. Formally, our joint policy is:

$$\pi_{\text{joint}} = \mathbb{I}_{\{\eta_{rec} \geq \epsilon_{rec}\}} \pi_b(\mathbf{O}_t, p_t) + \mathbb{I}_{\{\eta_{rec} < \epsilon_{rec}\}} \pi_{inv}(\zeta_{rec}^L, p_t) \quad (10)$$

We also describe our Joint Policy algorithmically in Algorithm 1 in Appendix A.

## 5 Evaluation

We first evaluated the Object-Centric Recovery framework’s performance using two simulation tasks, covering both prehensile and non-prehensile manipulation. These tasks were selected to demonstrate the framework’s versatility and robustness in handling a range of manipulation scenarios. Additionally, we evaluated a prehensile task in a real-world robot setting to showcase the framework’s effectiveness in practical applications. Given the absence of existing methods for object-centric recovery in visuomotor policies, to the best of our knowledge, we benchmarked our results against the out-of-distribution performance of the base BC policy. To ensure consistency across all tasks, we employed the vision-input U-Net-based diffusion policy [2] as our base policy,

and the low-dimensional diffusion policy was utilized as the architecture for our keypoint inverse policy. Our results show that, compared to the base policy, the Object-Centric Recovery framework consistently achieved a high task-completion rate in object OOD scenarios, with an average success rate of 81.0% across the three evaluated tasks, which is an improvement of **77.7%** over the base policy in OOD.

In addition, we show that for a life-long continuous learning scenario, we can employ the OCR framework to collect *recovery demonstrations* that can be augmented alongside the original base policy training dataset to imbue the *base policy* with the ability to recover at a high success rate without diminishing its performance in the original ID regions.

### 5.1 Experimental Setups

*Experiment 1) Non-Prehensile, Sim, Push-T Task [39]:* This 2D simulated task involves pushing a T-shaped block (gray) toward a fixed target using a circular end-effector agent (blue). At each reset, both the initial pose of the T block and the initial position of the end-effector are randomized. The task is particularly challenging due to the requirement for discontinuous, non-linear end-effector actions. To highlight the OCR framework’s capability to handle such complexity, we provided demonstrations initialized exclusively on the *left* side of the screen, as shown by the dashed line separating the screen in Figure 3. This setup designates the left side as in-distribution (ID) and the right side as out-of-distribution. For PushT, we recorded 100 demonstrations in the ID region.

*Experiment 2) Prehensile, Sim, Robomimic Square Task [40]:* This simulated task requires the robot to pick up a notched square-shaped object with a hole in the middle, transport it, and drop it through a fixed square peg. The initial pose of the square object is randomized within the  $SE(2)$  space on the table, and the initial position of the end-effector is also randomized. We used Robomimic’s PH dataset, which is 200 demonstrations initialized exclusively on the *right* side of the table, as shown by the red-shaded region in Figure 3. Hence, the left side of the table, or the green-shaded region is considered in-distribution

*Experiment 3) Prehensile, Real, Bottle Task:* The objective of the bottle task is for the robot to grasp a yellow bottle, transport it, and place it onto an elevated red plate. The initial pose of the bottle is randomized within the  $SE(2)$  space on the table, and the initial position of the end-effector is also randomized. As illustrated in Figure 4, the green-shaded region represents the demonstrated (in-distribution, ID) region, while the red-shaded area indicates the OOD region. We provided 115 demonstrations within the green-shaded region for policy training. We used a Franka Panda robot as the policy agent, and Polymetis [41] as the controller interface. For our 6D object pose estimation, we used Foundation Pose [27], and Grounded-SAM [42] to obtain the object mask.




Figure 3: **(Left)** shows the Push-T Task’s ID and OOD regions divided by a dashed line. **(Right)** shows the Square Task’s ID and OOD regions drawn out by the green and red regions respectively.

	Base Policy		Joint Policy (Ours)
	ID	OOD	OOD
Push-T	0.90	0.10	<b>0.93</b>
Square	0.87	0.00	<b>0.80</b>

Table 1: Simulated task success rate of the base policy vs. joint policy in OOD scenarios, with ID scenario baseline as the base policy.




Figure 4: **(Right)** shows our Franka recovering from OOD (red-shaded) to ID (green-shaded) in the Bottle Task.

	Base Policy		Joint Policy (Ours)
	ID	OOD	OOD
Bottle	0.60	0.00	<b>0.70</b>

Table 2: Real task success rate of the base policy vs. joint policy in OOD scenarios, with ID scenario baseline the base policy.

	Org Base Policy		Aug Base Policy	
	ID	OOD	ID	OOD
Push-T	0.90	0.10	0.97	<b>0.80</b>

Table 3: Push-T task success rate of the original base policy vs. augmented base policy (with OCR generated data) in ID and OOD scenarios.

*Experiment 4) Continuous Learning on Push-T:* To demonstrate the OCR framework’s effectiveness in lifelong continuous learning, we again applied Object-Centric Recovery to OOD-initialized Push-T scenarios. This time, however, the policy rolled out actions until the predefined  $\epsilon_{rec}$  threshold for reaching ID was met, while recording observations and proprioceptions. From 100 OOD initializations, we collected the dataset  $D_{aug}$ , which we appended to  $D_b$  directly to resume training on the base policy checkpoint. We then evaluated the resulting *augmented base policy* in both the original ID and OOD scenarios.

## 5.2 Experimental Results & Analysis

1) **Push-T.** Shown in Table 1, the Push-T base policy generalized to the OOD region object initialization in only 10% of cases, mainly due to random initialization placing the end-effector on the right, allowing it to accidentally push the block back into the in-distribution region. In contrast, with the OCR framework, recovery actions guide the system to a feasible interaction region (e.g., positioning on the right to push the block left) and gently return the block to the ID region, where the base policy completes the task. Across 30 random OOD initializations, the OCR framework achieved a 93% task success rate, significantly outperforming the base policy.

2) **Square.** In OOD object initialization scenarios, the Square base policy consistently attempted to move the end-effector toward the direction of the object but failed to reach it in all cases. In contrast, as shown in Table 1, the OCR framework successfully manipulated the square object for recovery, achieving an 80% task success rate in OOD scenarios, a substantial improvement over the base policy.

3) **Bottle.** In OOD object initialization scenarios, the Bottle base policy failed similarly to the Square base policy, as both tasks involve prehensile manipulation. However, as shown in Table 2, the OCR framework effectively handled the recovery for both tasks, achieving an 70% task success rate in OOD scenarios, significantly outperforming the base policy. Interestingly, we observed that the OCR joint policy’s OOD success rate is significantly higher than the base policy’s ID success rate for the bottle task. We hypothesize that this can be attributed to the OCR framework’s ability to move objects toward regions of high training density regardless of where the objects are initialized.

4) **Continuous Learning on Push-T.** By resuming training on the base policy using both the original dataset  $D_b$  and the augmented dataset  $D_{aug}$  - autonomously collected via the OCR framework - we enabled the base policy to recover independently, achieving 80% task completion in the original OOD regions, as shown in Table 3,. This improvement did not come at the cost of performance in the original ID scenarios; in fact, the augmented policy showed enhanced performance in ID as well, improving from 90% to 97%. We hypothesize this is due to the OCR’s augmented dataset consistently providing robot actions that move the object toward regions of higher density, even on the ID side. In other words, OCR demonstrations likely offer actions that **converge** the object to the base policy, complementing it rather than replacing it. We believe this showcases the OCR framework’s ability to provide valuable data for continuous learning.

## 6 Conclusion

In this work, we proposed the Object-Centric Recovery policy framework designed to address out-of-distribution challenges in visuomotor policy learning, by recovering task-relevant objects into distribution without requiring additional data collection. When our framework was tested against various manipulation tasks, it demonstrated considerable improvement in performance in OOD regions. Furthermore, the framework’s capacity for continuous learning highlights its potential to autonomously enhance policy behavior over time. However, a key limitation of our approach is the reliance on explicit object poses, which restricts its applicability to articulated and deformable objects. We aim to extend our framework by incorporating more flexible scene representations to recover from a broader range of OOD scenarios. Despite this limitation, we believe our framework represents a step towards improving the robustness of visuomotor policies in real-world settings.



## Acknowledgments

We thank Shubhdeep Shiv Aditya for assistance in setting up the real robot vision pipeline, and Dr. Dinesh Jayaraman for the valuable discussions. We also appreciate the thoughtful feedback and comments provided by the reviewers.

## References

- [1] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, J. Peters, et al. An algorithmic perspective on imitation learning. *Foundations and Trends® in Robotics*, 7(1-2):1–179, 2018.
- [2] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song. Diffusion policy: Visuomotor policy learning via action diffusion. *arXiv preprint arXiv:2303.04137*, 2023.
- [3] J. Pari, N. Muhammad, S. P. Arunachalam, and L. Pinto. The surprising effectiveness of representation learning for visual imitation. *arXiv preprint arXiv:2112.01511*, 2021.
- [4] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.
- [5] Z. Fu, T. Z. Zhao, and C. Finn. Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. *arXiv preprint arXiv:2401.02117*, 2024.
- [6] Y. Zhu, A. Joshi, P. Stone, and Y. Zhu. Viola: Imitation learning for vision-based manipulation with object proposal priors. In *Conference on Robot Learning*, pages 1199–1210. PMLR, 2023.
- [7] L. X. Shi, A. Sharma, T. Z. Zhao, and C. Finn. Waypoint-based imitation learning for robotic manipulation. *arXiv preprint arXiv:2307.14326*, 2023.
- [8] S. Ross, G. Gordon, and D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.
- [9] S. Haldar, J. Pari, A. Rai, and L. Pinto. Teach a robot to fish: Versatile imitation from one minute of demonstrations. *arXiv preprint arXiv:2303.01497*, 2023.
- [10] M. Kelly, C. Sidrane, K. Driggs-Campbell, and M. J. Kochenderfer. Hg-dagger: Interactive imitation learning with human experts. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8077–8083. IEEE, 2019.
- [11] J. Spencer, S. Choudhury, M. Barnes, M. Schmittle, M. Chiang, P. Ramadge, and S. Srinivasa. Learning from interventions. In *Robotics: Science and Systems (RSS)*, 2020.
- [12] A. Reichlin, G. L. Marchetti, H. Yin, A. Ghadirzadeh, and D. Kragic. Back to the manifold: Recovering from out-of-distribution states. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8660–8666. IEEE, 2022.
- [13] S. Ross and D. Bagnell. Efficient reductions for imitation learning. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 661–668. JMLR Workshop and Conference Proceedings, 2010.
- [14] S. Levine, A. Kumar, G. Tucker, and J. Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- [15] A. Kumar, A. Zhou, G. Tucker, and S. Levine. Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 33:1179–1191, 2020.

- [16] T. Yu, A. Kumar, R. Rafailov, A. Rajeswaran, S. Levine, and C. Finn. Combo: Conservative offline model-based policy optimization. *Advances in neural information processing systems*, 34:28954–28967, 2021.
- [17] I. Kostrikov, A. Nair, and S. Levine. Offline reinforcement learning with implicit q-learning. *arXiv preprint arXiv:2110.06169*, 2021.
- [18] K. Jiang, J.-y. Yao, and X. Tan. Recovering from out-of-sample states via inverse dynamics in offline reinforcement learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- [19] H. Zhang, J. Shao, Y. Jiang, S. He, G. Zhang, and X. Ji. State deviation correction for offline reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 9022–9030, 2022.
- [20] S. Fujimoto and S. S. Gu. A minimalist approach to offline reinforcement learning. *Advances in neural information processing systems*, 34:20132–20145, 2021.
- [21] L. Ankile, A. Simeonov, I. Shenfeld, M. Torne, and P. Agrawal. From imitation to refinement—residual rl for precise visual assembly. *arXiv preprint arXiv:2407.16677*, 2024.
- [22] C. M. Bishop. Mixture density networks. 1994.
- [23] Y. Chen, C. Wang, Y. Yang, and K. Liu. Object-centric dexterous manipulation from human motion data. In *8th Annual Conference on Robot Learning*, 2024. URL <https://openreview.net/forum?id=KAzku0Uyh1>.
- [24] J. Gao, Z. Tao, N. Jaquier, and T. Asfour. K-vil: Keypoints-based visual imitation learning. *IEEE Transactions on Robotics*, 2023.
- [25] P. Mandikal and K. Grauman. Learning dexterous grasping with object-centric visual affordances. In *2021 IEEE international conference on robotics and automation (ICRA)*, pages 6169–6176. IEEE, 2021.
- [26] C. Devin, P. Abbeel, T. Darrell, and S. Levine. Deep object-centric representations for generalizable robot learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7111–7118. IEEE, 2018.
- [27] B. Wen, W. Yang, J. Kautz, and S. Birchfield. Foundationpose: Unified 6d pose estimation and tracking of novel objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17868–17879, 2024.
- [28] C. Wen, X. Lin, J. So, K. Chen, Q. Dou, Y. Gao, and P. Abbeel. Any-point trajectory modeling for policy learning. *arXiv preprint arXiv:2401.00025*, 2023.
- [29] M. Xu, Z. Xu, Y. Xu, C. Chi, G. Wetzstein, M. Veloso, and S. Song. Flow as the cross-domain manipulation interface. *arXiv preprint arXiv:2407.15208*, 2024.
- [30] Y. Dai, J. Lee, N. Fazeli, and J. Chai. Racer: Rich language-guided failure recovery policies for imitation learning. *arXiv preprint arXiv:2409.14674*, 2024.
- [31] C. Cornelio and M. Diab. Recover: A neuro-symbolic framework for failure detection and recovery. *arXiv preprint arXiv:2404.00756*, 2024.
- [32] P. Li, Z. Li, H. Zhang, and J. Bian. On the generalization properties of diffusion models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=hCUG1MCFk5>.
- [33] D. Hendrycks, N. Mu, E. D. Cubuk, B. Zoph, J. Gilmer, and B. Lakshminarayanan. AugMix: A simple data processing method to improve robustness and uncertainty. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020.

- [34] M. Federici, R. Tomioka, and P. Forré. An information-theoretic approach to distribution shifts. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=GrZmKDYCp6H>.
- [35] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1):99–134, 1998. ISSN 0004-3702. doi:[https://doi.org/10.1016/S0004-3702\(98\)00023-X](https://doi.org/10.1016/S0004-3702(98)00023-X). URL <https://www.sciencedirect.com/science/article/pii/S000437029800023X>.
- [36] C. M. Bishop and N. M. Nasrabadi. *Pattern recognition and machine learning*, volume 4. Springer, 2006.
- [37] Y. Hai, R. Song, J. Li, and Y. Hu. Shape-constraint recurrent flow for 6d object pose estimation. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4831–4840, 2023. doi:[10.1109/CVPR52729.2023.00468](https://doi.org/10.1109/CVPR52729.2023.00468).
- [38] P. Ausserlechner, D. Habegger, S. Thalhammer, J.-B. Weibel, and M. Vincze. Zs6d: Zero-shot 6d object pose estimation using vision transformers. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 463–469, 2024. doi:[10.1109/ICRA57147.2024.10611464](https://doi.org/10.1109/ICRA57147.2024.10611464).
- [39] P. R. Florence, C. Lynch, A. Zeng, O. Ramirez, A. Wahid, L. Downs, A. S. Wong, J. Lee, I. Mordatch, and J. Tompson. Implicit behavioral cloning. *ArXiv*, abs/2109.00137, 2021. URL <https://api.semanticscholar.org/CorpusID:237346088>.
- [40] A. Mandlekar, D. Xu, J. Wong, S. Nasiriany, C. Wang, R. Kulkarni, L. Fei-Fei, S. Savarese, Y. Zhu, and R. Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. In *arXiv preprint arXiv:2108.03298*, 2021.
- [41] Y. Lin, A. S. Wang, G. Sutanto, A. Rai, and F. Meier. Polymetis. <https://facebookresearch.github.io/fairo/polymetis/>, 2021.
- [42] T. Ren, S. Liu, A. Zeng, J. Lin, K. Li, H. Cao, J. Chen, X. Huang, Y. Chen, F. Yan, Z. Zeng, H. Zhang, F. Li, J. Yang, H. Li, Q. Jiang, and L. Zhang. Grounded sam: Assembling open-world models for diverse visual tasks, 2024.

## A Final Joint Policy Algorithm

---

**Algorithm 1** Joint Policy Algorithm

---

```
1: Initialize  $\pi_b, \pi_{inv}$ , GMMs with parameters  $\Theta$ 
2: while Task not done do
3:   Collect observation  $\mathbf{o}_t$ , proprioception  $p_t$ . Compute object pose  $\mathbf{T}_{obj,t}$ , keypoints  $\rho_t$ .
4:   Evaluate mean keypoint recovery vector  $\delta_{rec}$  and mean keypoint density  $\eta_{rec}$ .
5:   if  $\eta_{rec} < \epsilon_{rec}$  then
6:     Compute keypoint recovery trajectory  $\zeta_{rec}^L$ 
7:     Compute recovery action trajectory  $\mathbf{a}_{out}^L = \pi_{inv}(\zeta_{rec}^L, p_t)$ 
8:   else
9:     Compute base action trajectory  $\mathbf{a}_{out}^L = \pi_b(\mathbf{o}_t, p_t)$ 
10:  end if
11:  for  $a$  in  $\mathbf{a}_{out}^L$  do
12:    Execute action  $a$ 
13:  end for
14: end while
```

---