

Learning from Physical Human Feedback: An Object-Centric One-Shot Adaptation Method

Alvin Shek¹, Bo Ying Su¹, Rui Chen¹, Changliu Liu¹

Abstract—For robots to be effectively deployed in novel environments and tasks, they must be able to understand the feedback expressed by humans during intervention. This can either correct undesirable behavior or indicate additional preferences. Existing methods either require repeated episodes of interactions or assume prior known reward features, which is data-inefficient and can hardly transfer to new tasks. We relax these assumptions by describing human tasks in terms of object-centric sub-tasks and interpreting physical interventions in relation to specific objects. Our method, Object Preference Adaptation (OPA), is composed of two key stages: 1) pre-training a base policy to produce a wide variety of behaviors, and 2) online-updating according to human feedback. The key to our fast, yet simple adaptation is that general interaction dynamics between agents and objects are fixed, and only object-specific preferences are updated. Our adaptation occurs online, requires only one human intervention (one-shot), and produces new behaviors never seen during training. Trained on cheap synthetic data instead of expensive human demonstrations, our policy correctly adapts to human perturbations on realistic tasks on a physical 7DOF robot. Videos, code, and supplementary material: <https://alvinosaur.github.io/AboutMe/projects/opa>.

I. INTRODUCTION

Robots are useful in many real world tasks, e.g., where there are health risks or where tedious efforts are required. Ways to achieve desired behavior range from traditional planning and control to reinforcement or imitation learning. Cost functions, rewards, and pre-recorded demonstrations that are representative of the target scenarios are often fixed beforehand. However, during actual deployment, unpredictable factors can always arise such as varying task specifications and human preferences on how the tasks should be carried out. Such changes can be easily and naturally handled if a human can provide corrective physical feedback in real time. In such physical human-robot interactions (pHRI), robots need to understand this feedback and adapt their behavior accordingly for future tasks. In this paper, we study a special yet prevalent setting where the online human feedback is *object-centric*, as opposed to non-object-centric settings such as those encoding temporal preferences. That is, human interventions can be interpreted as leading the robot to either avoid or visit certain objects and areas in the environment with possibly preferred orientations. Hence, by leveraging this assumption, the robot can infer which objects are relevant and better understand human feedback.

Illustrating Example: To better understand this insight, consider a factory setting where a robot has been trained to carry printed items from a 3D printer to a bin for cleaning.

¹Carnegie Mellon University, Pittsburgh, PA. Contact: {ashek, boyings, ruic, cliu6}@andrew.cmu.edu



Fig. 1: OPA interprets human feedback in relation to specific objects, enabling fast online adaptation. In this scene, the default behavior (left) is to carry a can upright to a dropoff zone. After a perturbation in orientation (middle), the robot learns to hold the can correctly, allowing the human inspector to scan its label properly (right).

Now, the manufacturing process suddenly changes: the robot should additionally place items upright in a scanner before dropping them in the cleaning bin. As a human operator, a natural and quick way to communicate this task modification would be a physical correction. After the robot picks up an item, the human would drag and guide the arm towards the scanner and have the end-effector hold the item upright for a few seconds before guiding the arm towards the bin. This detour could be simply described as an additional object-centric sub-task, which requires the item to be moved with an upright orientation *relative to* the downward-facing scanner. In this work, we take advantage of this object-centric interpretation of human intentions to enable fast, yet simple online adaptation of robot behaviors.

Existing Approaches: Several approaches have been proposed to adapt robot behaviors using human feedback. [1]–[4] [3] assumes human intentions can be modelled as a linear combination of predefined feature functions. However, these pre-defined basis functions might not cover changing human preferences, especially when robots are deployed to new tasks, such as scanning labels under a newly-installed scanner. More recent works attempt to learn such features with deep neural networks [5]–[7], but effective learning then requires sufficient data and time and still may not generalize to new test scenarios [8]. [9]–[11] attempt to learn small MLP’s using more structured human feedback, but such simple MLP’s may not generalize to new scenarios.

Our Approach: To address the above challenges, we propose Object Preference Adaptation (OPA) to 1) explicitly consider pairwise relations between each object and the robot and 2) interpret human feedback in terms of these relations. We first pre-train a policy to reproduce randomly-generated trajectories that follow a specific pattern: reaching a goal while interacting with nearby objects. The policy parameters are divided into two groups: the core weights encoding the general interaction dynamics between objects, and the object-specific features. We assume there are only four

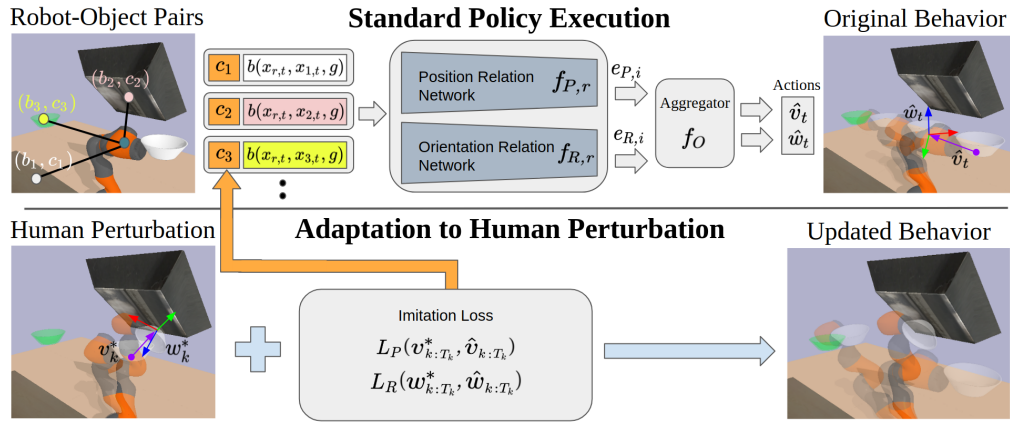


Fig. 2: Illustration of adaptation with OPA where the desired behavior is to have the bowl’s bottom side face the large black screen. OPA consists of a position relation network $f_{P,r}$ and an orientation relation network $f_{R,r}$ whose outputs are aggregated by f_O . **Top:** For each object $i \in [N]$ in a scene (including the green goal), a state-based feature $\mathbf{b}_R(\mathbf{x}_{r,t}, \mathbf{x}_{i,t}, \mathbf{g})$ is computed from both object state $\mathbf{x}_{i,t}$ and robot state $\mathbf{x}_{r,t}$. This is passed in along with a learned, preference feature \mathbf{c}_i into each relation network to produce edge features $e_{P,i}$ and $e_{R,i}$. Together, all edge features are fed into an aggregator function f_O that finally computes predicted translation $\hat{\mathbf{v}}_t$ shown in purple and orientation $\hat{\mathbf{w}}_t$ shown as Cartesian axes in the top right. **Bottom:** In the bottom left, humans can perturb both position and orientation to produce $\mathbf{v}_{k:T_k}^*$ and $\mathbf{w}_{k:T_k}^*$, which can be treated as ground truth for an imitation loss. Imitation loss includes position loss L_P and rotation loss L_R , defined as 6 and 7 respectively. Only object preference features \mathbf{c}_i are updated, allowing for fast and effective adaptation of behavior such as scanning the bottom of a bowl as shown in the final image to the right.

possible classes of object interactions during data generation: either moving closer or farther away in position, and either ignoring or matching the orientation of an object with a fixed rotational offset. Real-world examples of this include moving above and horizontally tilting to pour water into a cup, or staying away from an open flame. As the policy learns to imitate these trajectories, it also learns two key properties: how objects generally influence a trajectory (the interaction dynamics), and how specific types of object relations can be represented in a continuous latent space of finite dimensions.

When adapting to physical perturbations, the core policy weights are frozen because the general dynamics of object interactions should not change. Rather, *object-specific features* need to be adapted to capture various behaviors. In the illustrating example, the scanner should always be able to attract or repel the robot (the dynamics), but which specific behavior and to what extent are object-specific. We argue that such separation can be naturally captured by a graph with the agent and objects as nodes. The interaction dynamics are shared among all edges, while how each node pair influences the robot behaviors also depends on the actual objects. Following this idea, we compose our robot policy based on Graph Neural Network (GNN), which further allows end-to-end learning given human feedback. As a result, during online adaptation, we constrain the search to a compact latent space of object-specific features, which can be done with only a few gradient steps. Although our pre-training data contained only four classes of object relations, optimization over a continuous latent space allows the policy to express new, unseen types of interactions with objects, such as ignoring objects for position, or learning arbitrary 3D orientations relative to an object’s orientation.

To evaluate our model’s adaptability, we compared our method with three other state-of-the-art methods in a real-life task where a human inspector must perturb the robot to move closer and present cans with a desired orientation as

shown in Fig. 1.

Overall, we summarize our contributions as follows:

- 1) We interpret human physical feedback as object-specific interactions and model these with a graph representation, enabling fast adaptation.
- 2) We show that synthesized data alone can train a policy that handles realistic tasks and unstructured human perturbations.
- 3) We experimentally show adaptability to human perturbation on a real robot task.

II. UNDERSTANDING OBJECT-CENTRIC HUMAN FEEDBACK

At time t , let $\mathbf{x}_{r,t} \in \mathbb{R}^d$ represent the current state of the robot r , and $\mathbf{x}_{i,t} \in \mathbb{R}^d$ represent the states of other objects $i \in [N]$ in a scene. This assumes we can detect, uniquely identify, and extract the 6D pose of each object in a scene. We leave unstructured environments for future work.

The robot’s state updates according to dynamics $\mathbf{x}_{r,t+1} = \mathbf{x}_{r,t} + \alpha \mathbf{u}_t$ where α is a constant step size and $\mathbf{u}_t \in \mathbb{R}^d$ is the action either applied externally by the human if available, or otherwise generated by the policy. We denote the human intervention as \mathbf{u}_t^* and the robot policy as $\hat{\mathbf{u}}_t = f_O(\mathbf{x}_{r,t}, \{\mathbf{x}_{i,t}\}_{i \in [N]}, \mathbf{g})$ where the goal \mathbf{g} is assumed given. The general learning objective is to match the human’s actions over T steps:

$$\min_{f_O} \sum_{t \in [T]^*} \|f_O(\mathbf{x}_{r,t}, \{\mathbf{x}_{i,t}\}_{i \in [N]}, \mathbf{g}) - \mathbf{u}_t^*\|, \quad (1)$$

where $[T]^*$ is the set of time indices where human interventions are available.

To infer object-centric meanings from human actions, we represent the scene as a graph, where nodes are the robot and all other objects, and directed edges pointing only from each object to the robot. A graph naturally lets us model each object and its relation to the robot separately; graphs can

also capture potential object-centric preferences from human feedback. Note that our approach can be easily extended to consider relations between objects with additional edges. In those cases, our approach shows even more advantage in providing a computationally efficient way to capture complex interactions between objects.

In our graph, each of the agent-object relations is composed of two elements: relative state features $\mathbf{b}_{i,t}$ and learned “preference” features \mathbf{c}_i . State features are computed from robot, object, and goal state $\mathbf{b}_{i,t} = \mathbf{b}(\mathbf{x}_{r,t}, \mathbf{x}_{i,t}, g)$ and help the policy reason about spatial proximity and direction of each object relative to the robot and its goal. Preference features \mathbf{c}_i capture the specific way that each object should interact with the agent. Together, these can be processed by a *relation network* f_r to generate a latent edge features $\mathbf{e}_{i,t} = f_r(\mathbf{b}_{i,t}, \mathbf{c}_i)$, as shown in the top of Fig. 2. We can rewrite the robot’s overall policy as an “aggregator” f_O of edge features $f_O(\{\mathbf{e}_{i,t}\}_{i \in [N]})$ and its objective as

$$\min_{\{\mathbf{c}_i, f_r\}} \sum_{t \in [T]^*} \|f_O(\{\mathbf{e}_{i,t}\}_{i \in [N]}) - \mathbf{u}_t^*\|. \quad (2)$$

Note that we learn preference features \mathbf{c}_i while keeping the manually designed state features $\mathbf{b}_{i,t}$ fixed. This naturally isolates the objective information of the scene from subjective preferences and helps the model to focus on object-specific preference features.

III. OBJECT-CENTRIC REPRESENTATION

We now consider the specific domain of full rigid body motion in 2D or 3D for OPA. Time index t is omitted for clarity. We assume object-specific tasks can generally be represented as reach and avoid, and only focus on this setting without considering the complexity of actual grasping. We leave this for future work. The objects and robot are represented as spheres with radii of influence $s_{i \in [N]}$ and s_r , respectively. Tricks to handle non-spherical objects are discussed in task 1 of the experiments.

Actions \mathbf{u} are composed of translation \mathbf{v} and rotation \mathbf{w} actions. \mathbf{v} is a unit vector describing the direction of translation. \mathbf{w} describes the desired orientation at the next timestep. For 2D, $\mathbf{w} = [\cos(\theta), \sin(\theta)]$ for angle θ on the x-y plane. For 3D, $\mathbf{w} = [R_x^T, R_y^T] \in \mathbb{R}^6$ where $R_x, R_y \in \mathbb{R}^3$ are the x and y axes of an orientation expressed by a rotation matrix $R = [R_x, R_y, R_z] \in SO(3)$. This is shown in the top right of Fig. 2, where $\hat{\mathbf{v}}_t$ points from the end-effector’s current to future position, whereas $\hat{\mathbf{w}}_t$ is the future orientation.

States \mathbf{x}_i are separated by position $\mathbf{x}_{P,i}$ and orientation $\mathbf{x}_{R,i}$. State features \mathbf{b}_i are separated by position $\mathbf{b}_{P,i} = \mathbf{b}_P(\mathbf{x}_r, \mathbf{x}_i, g)$ and orientation $\mathbf{b}_{R,i} = \mathbf{b}_R(\mathbf{x}_r, \mathbf{x}_i, g)$. Preference features \mathbf{c}_i are composed of position component $\mathbf{c}_{P,i}$ and orientation component $\mathbf{c}_{R,i}$. $\mathbf{c}_{P,i}$ is purely latent whereas $\mathbf{c}_{R,i} = [\mathbf{c}_{R,i}^{\text{latent}}, \mathbf{c}_{R,i}^{\Delta}]$ is composed of a latent part $\mathbf{c}_{R,i}^{\text{latent}}$ and a learned rotational offset $\mathbf{c}_{R,i}^{\Delta}$. This learned offset is applied directly to the input orientation of an object and allows the model to learn relative orientations. For example, scanning the barcode underneath a bowl in Fig. 2 requires

the barcode to face the scanner, and this relative orientation can be represented by a rotational offset.

A. Graph Representation Overview

To process a graph of arbitrarily many object nodes, we choose to use a GNN due to its desirable property of invariance to both the number and ordering of nodes. Also, since GNN’s apply the same network to process every edge in a graph, they enforce a strong inductive bias that general edge computations remain the same. Only each pair of node’s relative features distinguishes specific behavior. This is the key for GNN’s well-known generalization to unseen scenarios.

To process each robot-object node pair, two different networks compute latent edge features: $\mathbf{e}_{P,i} = f_{P,r}(\mathbf{b}_{P,i}, \mathbf{c}_{P,i})$ for position, and $\mathbf{e}_{R,i} = f_{R,r}(\mathbf{b}_{R,i}, \mathbf{c}_{R,i})$ for orientation. Finally, the output actions can be computed using the same function f_O but with different sets of edge features: $\hat{\mathbf{v}} = f_O(\{\mathbf{e}_{P,i}\}_{i \in [N]})$ and $\hat{\mathbf{w}} = f_O(\{\mathbf{e}_{R,i}\}_{i \in [N]})$, as shown in Fig. 2. In the following sections, we will explain the computation of these features for each object and how these are aggregated to produce a single output action.

B. Relation Network and Aggregator

In this section, we first discuss our general relation network f_r , which takes as input pairs of state and preference features $(\mathbf{b}_i, \mathbf{c}_i)$.

Typically, neural network inputs are raw states, hand-crafted state features, or outputs from a separate network. Our approach, however, treats the object-specific features \mathbf{c}_i as updatable weights to be learned with the rest of the network through gradient descent. As our policy trains on the same set of object types, these low-dimensional preference features gradually discriminate to represent very different behaviors. In fact, as will be described in later sections, these object feature spaces can be as small as 1D and still capture diverse behaviors. We note that during training, the object types are provided as input in the form of indices. This enables the model to associate the correct feature \mathbf{c}_i with each object, ensuring that they are trained consistently.

Pairs of state-based and preference features are fed into a multi-layer perceptron (MLP) f_r^1 to produce intermediate features $\tilde{\mathbf{e}}_i$. A second MLP f_r^2 computes unnormalized attention-like weights $\tilde{\alpha}_i$ for these features, and overall edge features $\mathbf{e}_i = \tilde{\alpha}_i \tilde{\mathbf{e}}_i$ take the form of actual actions. Given this set of weighted actions, the aggregator f_O simply performs a summation followed by a normalization.

For position action $\hat{\mathbf{v}}$, we normalize to a magnitude of 1 to produce a valid unit vector:

$$\hat{\mathbf{v}} = \frac{\sum_{i \in [N]} \mathbf{e}_{P,i}}{\|\sum_{i \in [N]} \mathbf{e}_{P,i}\|} \quad (3)$$

Each object essentially contributes a “force” pushing or pulling on the robot.

For orientation \hat{w}_t , we sum up all edge features and normalize for each specific axis:

$$\tilde{w} = \sum_{i \in [N]} e_{R,i} \text{ then } \tilde{w} = \begin{bmatrix} \frac{\tilde{w}_{1:3}}{\|\tilde{w}_{1:3}\|} & \frac{\tilde{w}_{4:6}}{\|\tilde{w}_{4:6}\|} \end{bmatrix}. \quad (4)$$

For the 2D case, $\hat{w} = \tilde{w} \in \mathbb{R}^2$. For the 3D case, however, since our above calculation is essentially a normalized weighted sum of rotation matrices, the resulting $\tilde{w} = [\tilde{R}_x^T, \tilde{R}_y^T]$ is not a valid rotation. \tilde{R}_x and \tilde{R}_y must be orthogonal, and to enforce this, we apply Gram-Schmidt orthogonalization to remove the component of \tilde{R}_y parallel to \tilde{R}_x while keeping $R_x = \tilde{R}_x$:

$$\tilde{R}'_y = \tilde{R}_y - \frac{\langle \tilde{R}_x, \tilde{R}_y \rangle}{\langle \tilde{R}_x, \tilde{R}_x \rangle} \tilde{R}_x \text{ and } R_y = \frac{\tilde{R}'_y}{\|\tilde{R}'_y\|}. \quad (5)$$

The resulting two axes $\hat{w} = [R_x^T, R_y^T]$ alone are sufficient to represent a full rotation in $SO(3)$, where the R_z axis is simply the cross product of the R_x and R_y axes.

C. Position Relation Network

In this section, we describe state-based and preference features for our position relation network. State-based features $b_{P,i}$ are the following:

- 1) **Size-relative distance:** Distance between robot and object divided by the sum of their radii, which helps the policy reason about when to interact with an object.
- 2) **Direction:** Unit-vector pointing from agent to each object, which helps determine the direction of ‘‘force’’ applied on agent.
- 3) **Goal-relative direction:** Inner-product between each agent-object vector and the agent-goal vector. A positive value indicates that the object lies in the same direction as towards the goal and should be considered. A negative value indicates the object is ‘‘behind’’ the agent and can be ignored.

Position preference features $c_{P,i}$ intuitively capture the magnitude of the output force, or how attracted or repelled the robot should be from each object.

The position relation network overall outputs a push-pull force on the agent. Potential field methods [12] also use this approach, but constrain this force vector to be parallel to each agent-object vector. This may seem like an intuitive way to enforce structure in the network and reduce complexity, but this constraint fails during ‘‘singularities’’ where no orthogonal component is available to avoid an obstacle lying in the same direction as the goal.

During training, we measure misalignment between ground-truth v_t^* and predicted \hat{v}_t translation directions using the following loss across B batch samples:

$$L_P = \frac{1}{B} \sum_{b=1}^B 1 - \langle v_b^*, \hat{v}_b \rangle \quad (6)$$




Fig. 3: Example scenario to train the orientation relation network in 2D. Expert path in red is annotated with an orientation vector at each waypoint. Each node (start, grey object, and goal) has both its observed orientation $x_{R,i}$ shown as a black arrow along with its desired, relative orientation $x_{R,i} \cdot R(c_{R,i}^{\Delta*})$ shown in green. The only node with nonzero rotational offset is the grey object: $c_{R,0}^{\Delta*} = -90^\circ$, which is why the green arrows are not visible for start and goal.

D. Orientation Relation Network

In this section, we describe state-based and learned features for our orientation relation network. State-based features $b_{R,i}$ are the following:

- 1) **Size-relative distance:** identical to that of the position relation network.
- 2) **Modified Orientation:** Orientation of each object, but rotated by a learned rotational offset. This helps the policy output the correct orientation relative to an object’s.

The learned rotational offset $c_{R,i}^{\Delta}$ is represented as a scalar $\Delta\theta$ for 2D and as a unit-normalized quaternion for 3D. $c_{R,i}^{\Delta}$ is applied to the x and y axes of the object’s rotation matrix. We choose to manually apply a learned rotational offset to reduce burden on the network of needing to learn how to apply valid rotations. An additional learned feature $c_{R,i}$ intuitively determines whether an object’s orientation should be ignored or not. This is important for modeling real-world tasks where a robot should move closer to an object without changing its orientation, such as handing a glass of water to a person.

Loss can be calculated by measuring the summed misalignment between ground-truth and predicted x and y axes:

$$L_R = \frac{1}{B} \sum_{b=1}^B 2 - \langle R_{x,b}^*, \hat{R}_{x,b} \rangle - \langle R_{y,b}^*, \hat{R}_{y,b} \rangle. \quad (7)$$

IV. TRAINING AND ONLINE ADAPTATION

So far, we have discussed the implementation and intuition behind our policy. We now discuss how to train and adapt such a policy at test time.

A. Training with Synthetic Data

It requires much effort to collect demonstrations for imitation learning [13]. Hence, we generate synthetic data for training. For position and orientation specifically, notice how their computation is completely independent. This allows us optimize their losses L_P and L_R separately, bypassing the issue of different units [14]. This also allows us to train both networks with different sets of data, which is necessary since

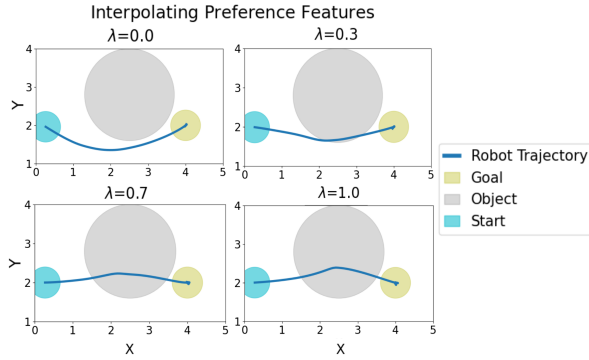


Fig. 4: Resulting change in robot’s behavior as the grey object’s preference feature shifts between “repel” $c_{P,0}$ and “attract” $c_{P,1}$ according to $(1 - \lambda)c_{P,0} + \lambda c_{P,1}$.

the behavior of the position and orientation networks should not be correlated.

For the position network specifically, interactions with nearby objects should take the form of attractions and repulsions; Elastic Bands [15] naturally model this. For the orientation network, trajectories involve interactions with a single object of two possible relations: *ignore* and *consider*. Ignored objects have no influence on the agent’s orientation. Considered objects force the agent to match their original orientation $\mathbf{x}_{R,0}$ relative to an offset $c_{R,0}^{\Delta*}$ shared among all considered objects. Fig. 3 shows an example where the expert orientation of a waypoint must match that of nearby objects. The expert initially matches the start orientation with zero offset, but switches to the grey object’s orientation in green with non-zero offset, and finally converges to the goal orientation with zero offset. Our policy would need to predict this relative orientation given only the original black orientation as input. Since observed orientations $\mathbf{x}_{R,0}$ are randomized while $c_{R,0}^{\Delta*}$ is fixed, our orientation network is forced to properly learn offset $c_{R,0}^{\Delta*}$ to predict orientations $\mathbf{x}_{R,0} \cdot R(c_{R,0}^{\Delta*})$ (note the right-multiply of offset).

Our method relies on objects to influence behavior, but we also may desire the robot to fix its orientation throughout a trajectory, even with no objects nearby. A cup of water, for example, should be carried upright. To train this soft constraint, start and goal are treated as imaginary objects whose orientation must also be considered. Both share a true rotational offset of the identity, or zero: $c_{R,g}^{\Delta*} = R(0)$.

B. Online Adaptation

In this section, we discuss our assumptions on the form of human perturbation as well as how the policy adapts to this to infer human preferences. [16] argues that people may intervene only when recent behavior has been unacceptable. This implies that the overall trajectory is not an example of a good trajectory, but rather a *transition* from bad to good. This thus provides information for *how the robot’s behavior should change* rather than the absolute trajectory that should be naively imitated.

Following that intuition, we focus only on the human perturbation trajectory ($\mathbf{x}_{k:T_k}^*, \mathbf{u}_{k:T_k}^*$) and treat this as an expert trajectory to imitate. A key assumption is that humans

only care about the final pose rather than the actual trajectory taken. Hence, we only take the start and end pose of the perturbation trajectory and linearly interpolate between the two. On the policy side, rather than calculating the output action for each individual expert state $\mathbf{x}_{k:T_k}^*$, we rollout the policy in an open-loop fashion. This allows gradients to propagate throughout the entire rollout, and early mistakes will be penalized for future errors.

Fine-tuning all weights of a neural network is well-known to lead to inefficient adaptation due to the large number of parameters. We bypass this issue by taking advantage of our graph-based architecture: updating only learned object features $c_{P,i}$ and $c_{R,i}$ while keeping the core relation network weights frozen. This allows us to drastically change the policy’s behavior with only a few steps of standard gradient descent. In Fig. 4, we visualize this flexibility in behavior. As c_P of the grey object interpolates from learned “repel” $c_{P,0}$ to “attract” $c_{P,1}$, the robot clearly moves closer to the object. The indices 0 and 1 refer to an arbitrary repel object and an arbitrary attract object respectively. At test-time, any new objects can be approximately initialized as “ignore” with value $(c_{P,0} + c_{P,1})/2$.

This works surprisingly well, even for the learned rotational offset features $c_{R,i}^{\Delta*}$. Recall that our model only had to learn two rotational offsets: $c_{R,0}^{\Delta*}$ and $c_{R,g}^{\Delta*}$. At test-time, however, our model can quickly adapt to reproduce *any arbitrary rotation* in $SO(3)$. This optimization is still non-convex, however, so in practice we run multiple optimizations with random initial values and pick the lowest cost solution.

V. REAL WORLD EXPERIMENTS

Now, we wish to compare our approach to three state-of-the-art baselines: **Online** [4], **FERL** [9], and **Unified** [11]. **Online** learns a linear combination of predefined feature functions. **FERL** and **Unified** can learn them from scratch as simple MLPs, compared to our GNN. In our experiment, a robot carries cans from one location to another and must learn to also present them *close to a human* with a certain *orientation relative to the human* as shown in Fig. 1. Although **Ours** and **Online** needed only one perturbation, **Unified** and **FERL** were given 10 perturbations so their neural networks could be fairly and properly trained. Perturbations were provided in one training scene, and each method was tested in 5 additional scenes.

We test two versions of **Online**, first **Online*** which is given exactly the right features of distance to human as well as error from human orientation $\mathbf{x}_{R,H}$ with offset ΔR^{H*} , defined as $\mathbf{x}_{R,H} \cdot \Delta R^{H*}$. **Online** is still given distance to the human but also to 5 other irrelevant objects. It isn’t given the exact rotational offset relative to the human, so instead 30 random offsets are sampled $\Delta \hat{R}^H \sim SO(3)$ and applied to the human pose in every scene along with the 5 objects, giving a total of 180 rotation features. This highlights a dilemma of sufficiently covering $SO(3)$ while trying to reduce the complexity of the predefined feature space.




Fig. 5: Comparing position loss in meters and orientation loss in radians of the methods. “Original” denotes the same scenario that each method received human feedback and trained in. “New” denotes the 5 new test scenarios.




Fig. 6: Comparing combined position and rotation loss of each method when given different amount of time for adaptation.

A. Real World Results

To evaluate the methods’ position performance, we measure their minimum achieved distance to the human. For rotation, we measure the minimum angle achieved from the desired, human-relative rotation $x_{R,H} \cdot \Delta R^{H*}$.

Fig. 5 compares performance between the “Original” training and 5 “New” test scenarios. We can first see that **Online*** achieves the lowest error in both position and orientation for both original and new scenes. This is expected as it optimized over the same loss functions that we use in evaluation and serves as a lower bound. Next, we can see that the more realistic **Online** achieves slightly higher position error but significantly higher orientation error, indicating that it learned to prioritize human distance but failed to learn the correct rotation feature from the 180 different features, highlighting the challenge of learning rotational offsets in $SO(3)$. **Ours** achieves low error in both position and orientation and notably generalizes well to the test scenarios. During execution, our policy would occasionally be attracted to the human and fail to reach the goal. That is due to the use of attention without hard constraints. Hence, we run constrained trajectory optimization to enforce goal-reaching with the minimum deviation from the nominal trajectory generated by our policy. **Unified** performs just as well for both position and orientation in the original scenarios, but does not generalize well to new test scenarios. This indicates that their small MLP overfitted to the training scenario. **FERL** does not perform well in both training and test scenarios, which could be caused by too short training time of 120 seconds, the max time we thought reasonable for “online” adaptation.

Fig. 6 compares performance against different adaptation times with position and rotation loss combined. **Online*** and **Ours** perform well even with only 1 second of adaptation. **Online*** similarly achieves its best performance given only 1 second. **Unified** requires at least 30 seconds to reach our

performance, and though **FERL** performs the worst, more adaptation time clearly leads to better performance.

VI. RELATED WORK

A. Online Adaptation from pHRI

Inverse Reinforcement Learning (IRL) methods attempt to model the preferences θ of an agent by observing its behavior [17]. The agent’s actions presumably maximize a reward function parameterized by θ , and the goal is that inferring correct θ will help model the agent’s behavior. Observations are often limited and noisy, which creates much ambiguity: many reward functions could represent these demonstrations [18]. IRL methods commonly address this by constraining the space of rewards to be composed of pre-specified basis functions $\phi(x)$ with unknown weights θ : $r(x) = \theta^T \phi(x)$ [1]–[4]. Although optimization of a convex loss with respect to these weights θ is convex, the chosen features $\phi(x)$ need to be specified a priori by an expert. In certain applications, these features are indeed known and fixed, such as holding a coffee cup upright, but in rapidly changing environments like the household, this assumption may not hold. Optimizing over the wrong features can then fail to capture the human’s desired behavior [19], [20], leading to incorrect robot behavior.

B. Learning Features Directly from State Space

Recent works have attempted to address this by indirectly learning such features through deep neural networks [5]–[7]. These deep IRL methods take in the raw state space as input and rely on the networks to learn a good representation that transfers to test time scenarios. Given a limited set of demonstrations and a high-dimensional state space, however, deep neural networks may not generalize to new scenarios [8]. The higher dimensional the state, the more diverse the data needs to be reduce ambiguity in the humans’ true preferences.

[9], [10] tackle this challenge by adding more structural constraints on how humans should provide feedback and train new feature functions as neural networks. To reduce ambiguity in state space, however, the human needs to provide multiple different trajectories to teach new behavior. [11] follows a similar strategy of learning new neural network-based reward functions from scratch. All three methods train small MLPs to reduce training time and needed data.

VII. CONCLUSION

We presented Object Preference Adaptation (OPA), a method for inferring human intentions for objects from physical corrections. OPA represents the environment as a graph with nodes as objects and edges between each object and the agent. After pre-training a base policy, OPA only needs to optimize object-specific features in a compact latent space. These aspects allow OPA to quickly and effectively adapt to real human feedback in a factory inspection task. OPA still has several limitations that could be addressed, such as learning more complex behavior and guaranteeing goal convergence.

REFERENCES

- [1] P. Abbeel and A. Y. Ng, "Apprenticeship learning via inverse reinforcement learning," ser. ICML '04. New York, NY, USA: Association for Computing Machinery, 2004, p. 1. [Online]. Available: <https://doi.org/10.1145/1015330.1015430>
- [2] A. Jain, S. Sharma, T. Joachims, and A. Saxena, "Learning preferences for manipulation tasks from online coactive feedback," 2016.
- [3] D. P. Losey and M. K. O'Malley, "Learning the correct robot trajectory in real-time from physical human interactions," *J. Hum.-Robot Interact.*, vol. 9, no. 1, Dec. 2019.
- [4] A. Bajcsy, D. P. Losey, M. K. O'Malley, and A. D. Dragan, "Learning robot objectives from physical human interaction," in *Proceedings of the 1st Annual Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, S. Levine, V. Vanhoucke, and K. Goldberg, Eds., vol. 78. PMLR, 13–15 Nov 2017, pp. 217–226.
- [5] D. S. Brown, R. Coleman, R. Srinivasan, and S. Niekum, "Safe imitation learning via fast bayesian reward inference from preferences," 2020. [Online]. Available: <https://arxiv.org/abs/2002.09089>
- [6] C. Finn, S. Levine, and P. Abbeel, "Guided cost learning: Deep inverse optimal control via policy optimization," 2016. [Online]. Available: <https://arxiv.org/abs/1603.00448>
- [7] M. Wulfmeier, D. Z. Wang, and I. Posner, "Watch this: Scalable cost-function learning for path planning in urban environments," 2016. [Online]. Available: <https://arxiv.org/abs/1607.02329>
- [8] J. Fu, K. Luo, and S. Levine, "Learning robust rewards with adversarial inverse reinforcement learning," 2017. [Online]. Available: <https://arxiv.org/abs/1710.11248>
- [9] A. Bobu, M. Wiggert, C. Tomlin, and A. D. Dragan, "Feature expansive reward learning," *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, Mar 2021.
- [10] —, "Inducing structure in reward learning by learning features," *The International Journal of Robotics Research*, April 2022.
- [11] S. A. Mehta and D. P. Losey, "Unified learning from demonstrations, corrections, and preferences during physical human-robot interaction," 2022. [Online]. Available: <https://arxiv.org/abs/2207.03395>
- [12] J.-C. Latombe, *Potential Field Methods*. Boston, MA: Springer US, 1991, pp. 295–355.
- [13] A. Mandekar, J. Booher, M. Spero, A. Tung, A. Gupta, Y. Zhu, A. Garg, S. Savarese, and L. Fei-Fei, "Scaling robot supervision to hundreds of hours with roboturk: Robotic manipulation dataset through human reasoning and dexterity," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 1048–1055.
- [14] R. Groenendijk, S. Karaoglu, T. Gevers, and T. Mensink, "Multi-loss weighting with coefficient of variations," 2020.
- [15] S. Quinlan and O. Khatib, "Elastic bands: connecting path planning and control," in *[1993] Proceedings IEEE International Conference on Robotics and Automation*, 1993, pp. 802–807 vol.2.
- [16] J. Spencer, S. Choudhury, M. Barnes, M. Schmittle, M. Chiang, P. Ramadge, and S. Srinivasa, "Learning from interventions: Human-robot interaction as both explicit and implicit feedback," 07 2020.
- [17] S. Arora and P. Doshi, "A survey of inverse reinforcement learning: Challenges, methods and progress," 2020.
- [18] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning," pp. 1433–1438, 2008.
- [19] A. Bobu, A. Bajcsy, J. F. Fisac, and A. D. Dragan, "Learning under misspecified objective spaces," 2018. [Online]. Available: <https://arxiv.org/abs/1810.05157>
- [20] L. Haug, S. Tschitschek, and A. Singla, "Teaching inverse reinforcement learners via features and demonstrations," 2018. [Online]. Available: <https://arxiv.org/abs/1810.08926>