

P3-PO: Prescriptive Point Priors for Visuo-Spatial Generalization of Robot Policies

Mara Levy¹

Siddhant Haldar²

Lerrel Pinto²

Abhinav Shirivastava¹

¹University of Maryland

²New York University

point-priors.github.io

Abstract—Developing generalizable robot policies that can robustly handle varied environmental conditions and object instances remains a fundamental challenge in robot learning. While considerable efforts have focused on collecting large robot datasets and developing policy architectures to learn from such data, naïvely learning from visual inputs often results in brittle policies that fail to transfer beyond the training data. This work presents Prescriptive Point Priors for Policies or P3-PO, a novel framework that constructs a unique state representation of the environment leveraging recent advances in computer vision and robot learning to achieve improved out-of-distribution generalization for robot manipulation. This representation is obtained through two steps. First, a human annotator prescribes a set of semantically meaningful points on a single demonstration frame. These points are then propagated through the dataset using off-the-shelf vision models. The derived points serve as an input to state-of-the-art policy architectures for policy learning. Our experiments across four real-world tasks demonstrate an overall 43% absolute improvement over prior methods when evaluated in identical settings as training. Further, P3-PO exhibits 58% and 80% gains across tasks for new object instances and more cluttered environments respectively. Videos illustrating the robot’s performance are best viewed at point-priors.github.io.

I. INTRODUCTION

A long standing goal in robotics has been to develop robot policies that are robust to environmental changes and can operate across variations in spatial configurations and object instances. While significant advances have been made in this direction for computer vision [1, 2] and natural language processing [3, 4, 5], the majority of robot policies remain confined to controlled laboratory environments with carefully designed settings. Robotic policies struggle to generalize to real-world scenarios because of the challenges and high costs associated with gathering diverse, high-quality robotic data.

Recent efforts aim to address this data problem by either aggregating existing robot datasets under a common framework [6] or collecting extensive real-world datasets through easy-to-use teleoperation tools [7, 8, 9, 10]. However, aggregated datasets suffer from inconsistencies across actions recorded for different projects [6], while large teleoperated collections, though useful, are often specific to a single robot type [8, 7] and it is unclear whether these approaches would scale for different robot morphologies. Consequently, developing generalized robot models still largely depends on collecting more expert demonstrations.

Correspondence to: mlevy@umd.edu



Fig. 1: A human prescribes key points one time for one instance of an object and those points are transferable to all other instances of the same object.

One way to get around this data problem is to use strong representation priors that transfer across scenarios and feed these representations as input into existing policy architectures. While priors such as object proposals [11, 12, 13] and pose estimation [14, 15] have been used in prior work, they often lose information which makes policy learning harder or require accurate modeling of the object poses to make the policy work. In this work, we explore if there exists a representation that is flexible, serves as a strong prior, and can provide the object-centric abstraction of a scene to enhance generalization. Compared to segmentation and object models, a point-based representation retains fine-grained spatial information without requiring accurate modeling of object boundaries or poses. By representing objects and scenes as a set of unstructured points, these representations extract only the essential geometric relationships between relevant elements in the scene. This allows the policy to focus exclusively on the key spatial interactions.

We present Prescriptive Point Priors for Policies or P3-PO, a novel framework that leverages the generalization capabilities of state-of-the-art computer vision models alongside state-of-the-art robot policy architectures. Through P3-PO, we demonstrate improved spatial generalization, the ability to

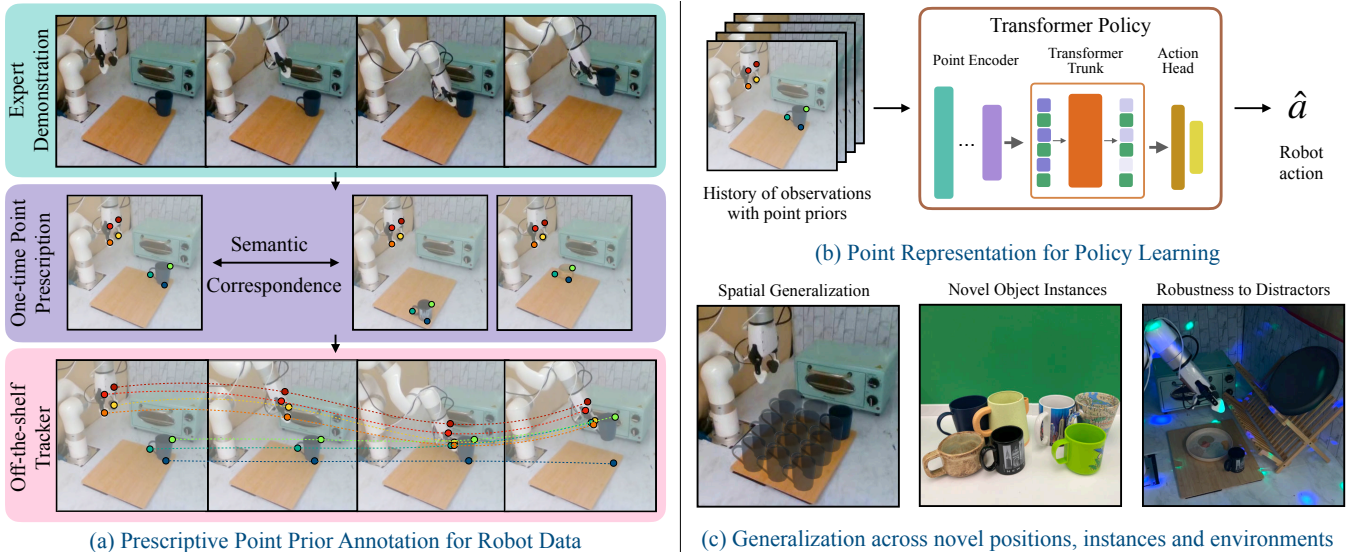


Fig. 2: Overview of the Prescriptive Point Priors for Policies (P3-PO) framework. (a) A human annotator prescribes a set of semantically meaningful key points on a **single demonstration frame**, typically in under 5 seconds. Off-the-shelf computer vision models are then used to automatically propagate these key points throughout the entire dataset without further human input. (b) The derived key points are leveraged by a transformer policy to predict the action. (c) P3-PO enables learning policies with improved generalization capabilities, including spatial generalization (i.e. generalization to new locations), generalization to novel object instances, and robustness to background distractors. P3-PO combines the strengths of vision and policy prediction methods through simple yet effective human-prescribed semantic guidance.

generalize to novel object instances, and robustness to large environmental changes. P3-PO is built on three key ideas. First, a human annotator prescribes a set of semantically meaningful points on a single demonstration frame, a process that often requires less than 5 seconds. Second, a diffusion-based visual correspondence model [16] and a state-of-the-art point tracker [17] are used to seamlessly transfer these points to the entire dataset without further human involvement. Third, the derived points are combined with state-of-the-art policy architectures like BAKU [18] for learning robust robot policies. The novelty of P3-PO lies in strategically combining the strengths of vision models and policy methods, yielding a simple yet effective approach for generalizable robot learning.

We demonstrate the effectiveness of P3-PO through experiments on four real-world tasks in an xArm Kitchen environment. Our main findings are summarized below:

- 1) P3-PO exhibits an overall 43% improvement over prior state-of-the-art policy learning algorithms across 4 real world tasks (Section V-E).
- 2) P3-PO generalizes to novel object instances, exhibiting a 58% improvement on a set of held-out objects as compared to prior work (Section V-F).
- 3) Policies trained with P3-PO are robust to the presence of background distractors (Section V-G) and work with both true depth and predicted metric depth from state-of-the-art models like Depth Anything [19, 20] (Section V-H).

All of our datasets, and training and evaluation code are publicly available. Videos of our trained policies can be seen here: point-priors.github.io.

II. RELATED WORKS

A. Imitation Learning (IL)

Imitation Learning (IL) [21] refers to training policies with expert demonstrations, without requiring a predefined reward function. In the context of reinforcement learning (RL), this is often referred to as inverse RL [22, 23], where the reward function is derived from the demonstrations and used to train a policy [24, 25, 26, 27, 28]. While these methods reduce the need for extensive human demonstrations, they still suffer from significant sample inefficiency. As a result of this inefficiency in deploying RL policies in the real world, behavior cloning (BC) [29, 30, 31, 32] has become increasingly popular in robotics. Recent advances in BC have demonstrated success in learning policies for both long-horizon tasks [33, 34, 35] and multi-task scenarios [18, 36, 6, 37]. However, most of these approaches rely on image-based representations [38, 18, 39, 36, 6, 40], which limits their ability to generalize to new objects and function effectively outside of controlled lab environments. In this work, we propose P3-PO, which attempts to address this reliance on image representations by directly using points priors as an input to the policy instead of raw images. Through extensive experiments, we observe that such an abstraction helps learn robust policies that generalize across varying scenarios.

B. Object-centric Representation Learning

Object-centric representation learning aims to create structured representations for individual components within a scene, rather than treating the scene as a whole. Common techniques in this area include segmenting scenes into

bounding boxes [41, 34, 42, 43, 12] and estimating object poses [44, 45]. While bounding boxes show promise, they share similar limitations with non object-centric image-based models, such as overfitting to specific object instances. Pose estimation, although less prone to overfitting, requires separate models for each object in a task. Another popular method involves using point clouds [13, 46], but their high dimensionality necessitates specialized models, making it difficult to accurately capture spatial relationships. In contrast, P3-PO leverages point prescription, eliminating the need to learn a representation because it is predefined by a human. This enables zero-shot generalization to both new objects and new spatial configurations. Similar to our method prior work [47] uses correspondence to identify where to interact with an object, however, this method relies on anygrasp [43], which is limited to a certain set of objects. Additionally, this method requires learning the affordance of an object, which can introduce errors that are less likely with point prescription.

III. BACKGROUND

A. Behavior Cloning

Behavior cloning [48, 49] aims to learn a behavior policy π^b given access to either the expert policy π^e or trajectories derived from the expert policy \mathcal{T}^e . This work operates in the setting where the agent only has access to observation-based trajectories, i.e. $\mathcal{T}^e \equiv \{(o_t, a_t)_{t=0}^T\}_{n=0}^N$. Here N and T denote the number of demonstrations and episode timesteps respectively. We choose this specific setting since obtaining observations and actions from expert or near-expert demonstrators is feasible in real-world settings [7, 50] and falls in line with recent work in this area [18, 51, 52, 7, 39].

B. Semantic Correspondence

Finding corresponding points across multiple images of the same scene is a well-established problem in computer vision [53, 54]. Correspondence is essential for solving a range of larger challenges, including 3D reconstruction [55, 56], motion tracking [17, 57, 58, 59], image registration [54], and object recognition [60]. In contrast, semantic correspondence focuses on matching points between a source image and an image of a different scene (e.g., identifying the left eye of a cat in relation to the left eye of a dog). Traditional correspondence methods [54, 53] often struggle with semantic correspondence due to the substantial differences in features between the images. Recent advancements in semantic correspondence utilize deep learning and dense correspondence techniques to enhance robustness [61, 62, 63] across variations in background, lighting, and camera perspectives. In this work, we adopt a diffusion-based point correspondence model, DIFT [16], to establish correspondences between a reference and an observed image, which is illustrated in Figure 3.

C. Point Tracking

Point tracking across videos is a problem in computer vision, where a set of reference points are given in the first frame of the video, and the task is to track these points




Fig. 3: Results of the correspondence model when used on the pick mug and plate off rack tasks. On the left is the frame that is annotated by a human. On the right we show that semantic correspondence [16] is able to identify the same points across a variety of instances of each object.

across multiple frames of the video sequence. Point tracking has proven crucial for many applications, including motion analysis [64], object tracking [65], and visual odometry [66]. The goal is to establish reliable correspondences between points in one frame and their counterparts in subsequent frames, despite challenges such as changes in illumination, occlusions, and camera motion. While traditional point tracking methods rely on detecting local features in images, more recent advancements leverage deep learning and dense correspondence methods to improve robustness and accuracy [17, 57, 58]. In this work, we use Co-Tracker [17] to track a set of reference points defined in the first frame of a robot’s trajectory. These points tracked through the entire trajectory are then used to train generalizable robot policies for the real world. This can be visualized in Figure 2.a.

IV. PRESCRIPTIVE POINT PRIORS FOR POLICIES (P3-PO)

Given demonstrations for robot manipulation tasks that cover a small set of possible object configurations and types, we seek to learn a generalizable robot policy that is robust to significant environmental variations and applicable to diverse object locations and types. To achieve this, we introduce P3-PO, an algorithm that decouples perception and planning to promote generalization. P3-PO operates in two phases. First, given a small set of robot demonstrations, the user annotates a single task frame with a set of semantically meaningful points. These reference points are propagated to the rest of the dataset using a combination of semantic correspondence and point tracking. The points obtained are fed into a transformer-based policy model for action prediction. An overview of our method has been provided in Figure 2. Below, we describe each component in detail.

A. One-Time Point Prescriptions

Our method begins by collecting robot demonstrations for a task through robot teleoperation [50]. The user then randomly selects one demonstration and annotates semantically meaningful points on the first frame that are relevant to performing the task, such as points on the robot and the




Fig. 4: Illustration of spatial variation used in our experiments.

objects being manipulated. This process often requires less than 5 seconds. These user-annotated points serve as priors for the rest of the data generation process. Using an off-the-shelf semantic correspondence model called DIFT [16], we transfer the points on the first frame to the corresponding locations on the first frames of all other demonstrations in the dataset. This allows us to initialize the key points across the entire dataset without additional human effort. For each demonstration, we then employ an off-the-shelf point tracking algorithm, Co-Tracker [17], to automatically track the initialized key points through the entire trajectory. In this way, by leveraging existing vision models for correspondence and tracking, we can efficiently compute the key points on every frame in the dataset while requiring the user to only annotate a single frame. This process, visualized in Figure 1, takes advantage of large, internet-scale pre-training of the vision models to generalize to new object instances and scenes without additional training. We prefer point tracking over correspondence at every frame due to its faster inference speed and ability to handle occlusions by continuing to track points. During inference, DIFT is used to mark the corresponding key points on the first frame, followed by Co-Tracker tracking the points during execution.

B. Policy Architecture

We employ the BAKU [18] architecture for policy learning. Instead of feeding in raw images, we use points derived from the previous section as input to the policy. Each 2D image point is first back-projected to 3D using the depth information from the camera. The points are flattened into a single vector in the order in which they are annotated. This point representation is then encoded using a multilayer perceptron (MLP) encoder. Given the noise in real-world depth sensing, we aggregate a history of observation features and feed them as separate tokens into the BAKU causal transformer policy. The policy predicts the action corresponding to each historical token using a deterministic action head and action chunking with exponential temporal averaging [7].



Fig. 5: Illustration of objects used in our experiments. In each image, the left pile depicts the in-domain objects while on the right are novel objects used in our generalization experiments.

V. EXPERIMENTS

Our experiments are designed to answer the following questions: (1) How well does P3-PO work for policy learning? (2) How well does P3-PO work for novel object instances? (3) Can P3-PO handle background distractors? (4) How does P3-PO perform with estimated depth? (5) Can P3-PO be improved with stronger priors?

A. Experimental Setup

Our experiments are performed on a Ufactory xArm 7 robot with an xArm Gripper in a kitchen environment. The policies are trained with RGB-D images from a third-person camera view and robot proprioception as input. The action space is comprised of the robot end effector pose and the gripper state. We collect a total of 160 demonstrations across 4 real-world tasks with varied object positions and types. The demonstrations are collected using a VR-based teleoperation system [50] at a 30Hz frequency, which are then subsampled to 5Hz. The learned policies are deployed at 5Hz.

B. Task Descriptions

We experiment with four manipulation tasks that exhibit significant variability in object position, type, and background context. Figure 6 provides rollouts of the tasks performed in our real-world setup. For each task, we collect expert demonstrations across a variety of object sizes and appearances. We refer to objects and environments seen in our collected data as in-domain. During evaluations, we add novel object instances that are unseen in the training data. The variations in positions and object instances for each task are depicted in Figure 4 and Figure 5 respectively. We provide a brief description of each task below.

a) Pick mug: The robot arm picks up a mug placed on the kitchen counter. The position of the mug is varied for each evaluation. We collect 15 demonstrations for 4 different mugs, resulting in a total of 60 demonstrations for the task. During evaluation, we introduce 3 novel mugs.

TABLE I: In-domain policy performance

Method	Pick mug	Lift plate	Mug on plate	Take out bottle
RGB [18]	13/40	22/30	8/40	5/20
RGB-D	19/40	22/30	7/40	6/20
GROOT [13]	7/40	7/30	0/40	0/20
P3-PO	39/40	29/30	32/40	13/20

b) *Lift plate*: The robot arm lifts a plate placed on the upper level of a rack. We collect 8 demonstrations for 3 different plates, resulting in a total of 24 demonstrations. During evaluation we introduce 2 novel plates.

c) *Mug on plate*: The robot arm picks up a mug placed on the kitchen counter and places it on a plate. We collect 15 demonstrations for 4 different mugs placed on the same plate, resulting in a total of 60 demonstrations for the task. During evaluation we use 1 novel mug and 1 novel plate.

d) *Take out bottle*: The robot arm takes a bottle out from the lower level of a rack. We collect 8 demonstrations for 2 different bottles, leading to a total of 16 demonstrations. During evaluation we introduce 2 novel bottles.

For all tasks, the xArm is initialized at its home position, while the object locations are varied across trials. During evaluation, the objects are placed in a held-out set of positions to keep comparisons fair across baselines.

C. Baselines

We compare P3-PO with three primary baselines.

a) *Full RGB Representation*: This method utilizes the BAKU transformer architecture [18], which takes the full RGB image of the scene and robot proprioception as input.

b) *RGB-D Representation*: This is a depth-based extension of BAKU that separately processes the depth image using an encoder and appends a depth token to the policy input for action prediction.

c) *GROOT [13]*: GROOT is a transformer-based imitation learning algorithm that constructs an object-centric 3D representation using Segment Anything [60] and a clustered point cloud which makes it robust to background distractors and novel objects. We refer the reader to the paper [13] for more details about GROOT.

D. Considerations for policy learning

P3-PO generates a point-based representation from 512x512 pixel images. For correspondence, P3-PO leverages DIFT [16], using the first layer of the hundredth time step with an ensemble size of 4. Point tracking in P3-PO is performed by a modified version of online Co-Tracker that enables tracking one frame at a time, rather than in chunks. P3-PO and GROOT utilize observation history while the RGB and RGB-D baselines do not [18]. Additionally, the RGB and RGB-D baselines incorporate robot proprioception as an input, while we follow GROOT and do not use robot proprioception as an input to P3-PO.

E. How well does P3-PO perform for policy learning?

We first evaluate P3-PO’s performance on the in-domain environment configurations using in-domain objects. We

TABLE II: Policy performance on novel object instances

Method	Pick mug	Lift plate	Mug on plate	Take out bottle
RGB	6/30	5/20	0/20	2/20
RGB-D	3/30	10/20	1/20	4/20
GROOT	4/30	5/20	0/20	0/20
P3-PO	29/30	18/20	16/20	10/20

TABLE III: Policy performance with background distractors

Method	Pick mug		Lift plate	
	In-domain	Novel object	In-domain	Novel object
RGB	0/5	0/5	4/5	0/5
RGB-D	1/5	0/5	2/5	0/5
GROOT	0/5	1/5	1/5	1/5
P3-PO	5/5	5/5	5/5	5/5

conduct 10 trials per object per task resulting in a variable number of total trials per task. The results have been reported in Table I. We observe that P3-PO outperforms the strongest baseline on average by 43% across our four real-world tasks. It must be noted that we only use 8-15 demonstrations per object per task, which is much smaller than prior works studying spatial generalization in robot policy learning [9, 7, 8]. For the full RGB and RGB-D baselines, most failures stem from minor errors, suggesting these policies could be improved with additional demonstrations. In the case of GROOT, having introduced larger variations in both the rotations and spatial locations of objects than the original paper [13], we observe that the policy is unable to generalize. We believe this limitation arises because GROOT normalizes each object-specific point cloud by its centroid, losing information about its position in space. To address this, GROOT adds the positional embedding of the centroid to the processed representation. However, this may not optimally reinforce the positional information. Videos on our website provide examples supporting these hypotheses.

F. How well does P3-PO work for novel objects?

Table II compares the performance of P3-PO with the baselines when tested on novel objects. These objects can be seen in Figure 5. We conduct 10 trials for each novel object for each task resulting in a variable number of total trials per task. We observe that P3-PO’s visual representation allows it to effectively generalize to novel object instances, outperforming the strongest baseline by 58% across all tasks. The RGB and RGB-D baselines exhibit reduced performance on novel objects due to their reliance on visual features learned from the training set. While designed for generalization, GROOT struggles with spatial variations resulting in lower accuracy. These results suggest that P3-PO, with its point-based non-image specific representation, is better equipped to generalize to novel objects than prior methods.

G. Can P3-PO handle background distractors?

We evaluate the performance of P3-PO in the presence of distractors in the task background. An illustration of the distractors has been included in Figure 2(c). We study this on two tasks - *pick mug* and *lift plate*. For each, we evaluate




Fig. 6: Real-world rollouts showing P3-PO’s capability on (a) in-domain objects and (b) novel objects.

TABLE IV: Effect of camera vs. predicted depth on P3-PO

P3-PO	Pick mug		Lift plate	
	In-domain	Novel object	In-domain	Novel object
Camera Depth	5/5	5/5	5/5	5/5
Depth Anything	5/5	5/5	5/5	5/5

the performance using 5 trials each on an in-domain object and a novel object. Table III provides these results. We observe that P3-PO outperforms the strongest baseline by 80%. The image-based baselines exhibit low accuracy due to their reliance on visual features while GROOT struggles with spatial generalization. These results reinforce that P3-PO’s point representation, decoupled from raw pixel values, enables policies that are robust to environmental perturbations.

H. How does P3-PO without ground truth depth?

Given recent advances in monocular depth prediction [19, 20], we investigate the importance of true depth values for the performance of P3-PO. To evaluate this, we compare P3-PO when using true depth from an RGB-D camera versus predicted depth from an off-the-shelf monocular depth estimation model, Depth Anything 2 [20]. As shown in Table IV, we observe that P3-PO achieves equivalent performance on two tasks, with one in-domain object and one novel object, regardless of whether true or predicted depth was provided. This is an interesting result, as it implies that P3-PO may be applicable to large-scale robot datasets [6, 67] which might not always include depth data.

I. Can P3-PO be improved with stronger priors?

Prior work has shown that encoding relational structure between inputs can improve policy learning generalization [68, 69, 70, 71]. In this section, we investigate whether encoding the spatial relationships between key points as a graph prior could further enhance P3-PO’s performance. Specifically, we represent the key points as a fully connected graph, with edges encoding the 3D distance between each pair of points. Leveraging the annotation order, we flatten this graph into a vector representation encoding the spatial relations between all point pairs. Policies are then trained on this graph-structured input. As shown in Table V, encoding the key points as a graph prior results in similar performance to directly using the key points as input. While additional structure did not provide clear benefits in this case, future work could explore more sophisticated relational encodings or combining our approach with other structural priors.

J. Can P3-PO complete more complex tasks?

In this section, we demonstrate that in addition to the tasks shown above, P3-PO excels at tasks that are more complex and dexterous than simple pick-and-place operations. First, we present results for a sweeping task where the robot picks up a broom and sweeps a nearby cutting board (Fig 7). This task is challenging for two reasons - (1) The broom’s handle is rounded, requiring precise handling to prevent slipping. This demonstrates that the point-based context provides sufficient environmental understanding for precise manipulation. (2) The task is long-horizon, consisting of two stages: lifting the broom and sweeping the board. P3-PO is able to understand its place in the sequence and act accordingly.

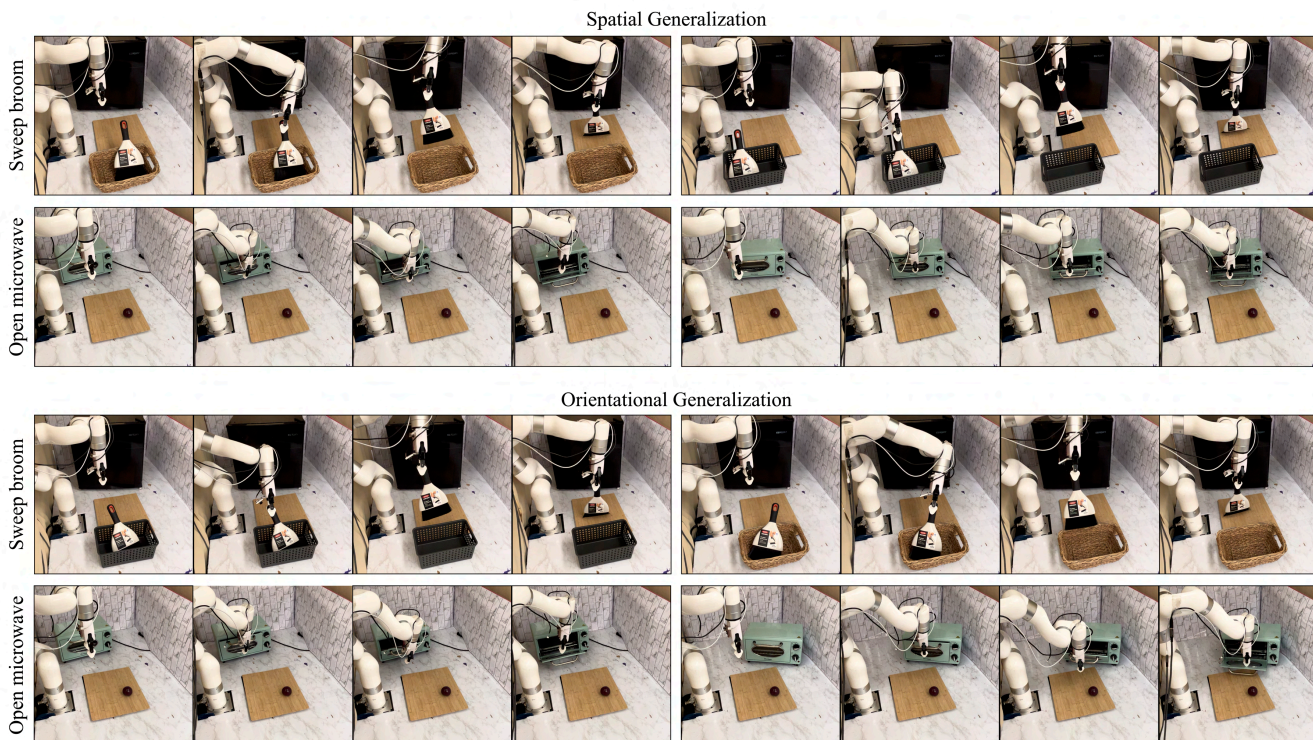


Fig. 7: Rollouts of the two complex tasks. On the top we show that for both of these tasks P3-PO can generalize to the object being in a different location. On the bottom we show that P3-PO can also generalize to different orientations of the object.

To evaluate P3-PO’s adaptability, we tested the task with two different baskets, showcasing its ability to handle varying vertical and horizontal angles. This is demonstrated in Fig 7. The model achieves an **80%** success rate when trained on 30 demonstrations and evaluated on 10 trials.

Next, we present results for an open-microwave task, which requires the robot to navigate into the thin opening between the handle and the microwave. Despite relying solely on point-based input, the robot is able to succeed with a high level of precision. This task is further challenging due to the wide range of orientations and variations in the microwave’s initial placement. The results demonstrate P3-PO’s ability to generalize to new locations and accurately interpret object orientations. P3-PO achieves an **80%** success rate on this task when evaluated on 10 trials and trained on 22 expert demonstrations. Variations in microwave locations and task execution are shown in Figure 7.

K. What do failures look like in P3-PO?

In Figure 8, we present examples of both successful and failed episodes across all six tasks. These examples highlight that while P3-PO does not always succeed, its failures come close to achieving the task objectives. For instance, in the "open microwave" task, the failures occur when the robot cannot open the microwave door far enough for the door to remain open. Similarly, tasks like "pick mug" and "lift plate" likely fail due to noise in the depth measurements. Notable these failures are infrequent and we believe they can be addressed in future iterations of this work.

TABLE V: Effect of stronger priors on P3-PO

Input	Lift plate		Take out bottle	
	In-domain	Novel object	In-domain	Novel object
Point	5/5	4/5	3/5	3/5
Graph	5/5	5/5	4/5	3/5

VI. CONCLUSION AND LIMITATIONS

In this work, we presented Prescriptive Point Priors for Policies (P3-PO), a simple yet effective framework that leverages human-provided semantic key points to enable more robust policy learning. P3-PO demonstrates improved generalization to spatial variations, novel objects, and distracting backgrounds compared to prior state-of-the-art methods. We recognize a few limitations in this work: (a) P3-PO’s reliance on existing vision models makes it susceptible to their failures. For instance, point tracking failures under occlusion hurt policy performance. However, we believe that continued advances in computer vision will serve to further strengthen performance of P3-PO. (b) While point abstractions facilitate better generalization, they lose information about scene context that could be important for navigation amid obstacles or clutter. Future work developing algorithms to retain sparse contextual cues while maintaining P3-PO’s object-centric representation may help address this. (c) In this work, we primarily study the single task performance of point prior policies. Extending the framework to multitask learning would be an interesting research direction. Overall, we believe P3-PO takes an important step toward developing general, data-




Fig. 8: Demonstrations of success and failure for each task. On the left we show successful demonstrations and on the right demonstrations of episodes that fail.

efficient robot policies suitable for real-world deployment by grounding them in human point priors.

VII. ACKNOWLEDGMENTS

We would like to thank Raunaq Bhirangi and Venkatesh Pattabiraman for their valuable feedback and discussions. This work was supported by grants from Honda, Google, NSF award 2339096, and ONR awards N00014-21-1-2758 and N00014-22-1-2773. LP is supported by the Packard Fellowship. AS is supported by NSF CAREER Award (#2238769). The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of NSF or the U.S. Government.

REFERENCES

- [1] J. Betker, G. Goh, L. Jing, T. Brooks, J. Wang, L. Li, L. Ouyang, J. Zhuang, J. Lee, Y. Guo, *et al.*, “Improving image generation with better captions,” *Computer Science*. <https://cdn.openai.com/papers/dall-e-3.pdf>, vol. 2, no. 3, p. 8, 2023.
- [2] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. L. Denton, K. Ghasemipour, R. Gontijo Lopes, B. Karagol Ayan, T. Salimans, *et al.*, “Photorealistic text-to-image diffusion models with deep language understanding,” *Advances in neural information processing systems*, vol. 35, pp. 36479–36494, 2022.
- [3] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, *et al.*, “Gpt-4 technical report,” *arXiv preprint arXiv:2303.08774*, 2023.
- [4] M. Reid, N. Savinov, D. Teplyashin, D. Lepikhin, T. Lillicrap, J.-b. Alayrac, R. Soricut, A. Lazaridou, O. Firat, J. Schrittwieser, *et al.*, “Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context,” *arXiv preprint arXiv:2403.05530*, 2024.
- [5] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, *et al.*, “Llama: Open and efficient foundation language models,” *arXiv preprint arXiv:2302.13971*, 2023.
- [6] A. Padalkar, A. Pooley, A. Jain, A. Bewley, A. Herzog, A. Irpan, A. Khazatsky, A. Rai, A. Singh, A. Brohan, *et al.*, “Open x-embodiment: Robotic learning datasets and rt-x models,” *arXiv*

- preprint *arXiv:2310.08864*, 2023.
- [7] 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.
 - [8] N. M. M. Shafiullah, A. Rai, H. Etukuru, Y. Liu, I. Misra, S. Chintala, and L. Pinto, "On bringing robots home," *arXiv preprint arXiv:2311.16098*, 2023.
 - [9] C. Chi, Z. Xu, C. Pan, E. Cousineau, B. Burchfiel, S. Feng, R. Tedrake, and S. Song, "Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots," *arXiv preprint arXiv:2402.10329*, 2024.
 - [10] H. Etukuru, N. Naka, Z. Hu, S. Lee, J. Mehu, A. Edsinger, C. Paxton, S. Chintala, L. Pinto, and N. M. M. Shafiullah, "Robot utility models: General policies for zero-shot deployment in new environments," *arXiv preprint arXiv:2409.05865*, 2024.
 - [11] M. Sieb, Z. Xian, A. Huang, O. Kroemer, and K. Fragkiadaki, "Graph-structured visual imitation," in *Conference on Robot Learning*. PMLR, 2020, pp. 979–989.
 - [12] 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*. PMLR, 2023, pp. 1199–1210.
 - [13] Y. Zhu, Z. Jiang, P. Stone, and Y. Zhu, "Learning generalizable manipulation policies with object-centric 3d representations," in *7th Annual Conference on Robot Learning*, 2023.
 - [14] C. Pan, B. Okorn, H. Zhang, B. Eisner, and D. Held, "Tax-pose: Task-specific cross-pose estimation for robot manipulation," in *Conference on Robot Learning*. PMLR, 2023, pp. 1783–1792.
 - [15] P. Vitiello, K. Dreczkowski, and E. Johns, "One-shot imitation learning: A pose estimation perspective," *arXiv preprint arXiv:2310.12077*, 2023.
 - [16] L. Tang, M. Jia, Q. Wang, C. P. Phoo, and B. Hariharan, "Emergent correspondence from image diffusion," in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
 - [17] N. Karaev, I. Rocco, B. Graham, N. Neverova, A. Vedaldi, and C. Rupprecht, "Cotracker: It is better to track together," 2023.
 - [18] S. Haldar, Z. Peng, and L. Pinto, "Baku: An efficient transformer for multi-task policy learning," *arXiv preprint arXiv:2406.07539*, 2024.
 - [19] L. Yang, B. Kang, Z. Huang, X. Xu, J. Feng, and H. Zhao, "Depth anything: Unleashing the power of large-scale unlabeled data," in *CVPR*, 2024.
 - [20] L. Yang, B. Kang, Z. Huang, Z. Zhao, X. Xu, J. Feng, and H. Zhao, "Depth anything v2," *arXiv:2406.09414*, 2024.
 - [21] A. Hussein, M. M. Gaber, E. Elyan, and C. Jayne, "Imitation learning: A survey of learning methods," *ACM Comput. Surv.*, vol. 50, no. 2, apr 2017.
 - [22] A. Y. Ng and S. J. Russell, "Algorithms for inverse reinforcement learning," in *Proceedings of the Seventeenth International Conference on Machine Learning*, ser. ICML '00. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2000, p. 663–670.
 - [23] P. Abbeel and A. Y. Ng, "Apprenticeship learning via inverse reinforcement learning," in *Proceedings of the Twenty-First International Conference on Machine Learning*, ser. ICML '04. New York, NY, USA: Association for Computing Machinery, 2004, p. 1.
 - [24] M. Levy, N. Saini, and A. Shrivastava, "Wayex: Waypoint exploration using a single demonstration," 2024.
 - [25] S. Haldar, V. Mathur, D. Yarats, and L. Pinto, "Watch and match: Supercharging imitation with regularized optimal transport," in *Conference on Robot Learning*. PMLR, 2023, pp. 32–43.
 - [26] S. Haldar, J. Pari, A. Rai, and L. Pinto, "Teach a robot to fish: Versatile imitation from one minute of demonstrations," 2023.
 - [27] J. Ho and S. Ermon, "Generative adversarial imitation learning," *CoRR*, vol. abs/1606.03476, 2016.
 - [28] A. Nair, A. Gupta, M. Dalal, and S. Levine, "Awac: Accelerating online reinforcement learning with offline datasets," *arXiv preprint arXiv:2006.09359*, 2020.
 - [29] D. Pomerleau, "Alvinn: An autonomous land vehicle in a neural network," in *Proceedings of (NeurIPS) Neural Information Processing Systems*, D. Touretzky, Ed. Morgan Kaufmann, December 1989, pp. 305 – 313.
 - [30] F. Torabi, G. Warnell, and P. Stone, "Recent advances in imitation learning from observation," *arXiv preprint arXiv:1905.13566*, 2019.
 - [31] S. Schaal, "Learning from demonstration," *Advances in neural information processing systems*, vol. 9, 1996.
 - [32] 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*. JMLR Workshop and Conference Proceedings, 2011, pp. 627–635.
 - [33] Y. Chen, C. Wang, L. Fei-Fei, and C. K. Liu, "Sequential dexterity: Chaining dexterous policies for long-horizon manipulation," 2023.
 - [34] A. Mandelkar, D. Xu, R. Martín-Martín, S. Savarese, and L. Fei-Fei, "Learning to generalize across long-horizon tasks from human demonstrations," *CoRR*, vol. abs/2003.06085, 2020.
 - [35] M. Shridhar, L. Manuelli, and D. Fox, "Cliport: What and where pathways for robotic manipulation," *CoRR*, vol. abs/2109.12098, 2021.
 - [36] H. Bharadhwaj, J. Vakil, M. Sharma, A. Gupta, S. Tulsiani, and V. Kumar, "Roboagent: Generalization and efficiency in robot manipulation via semantic augmentations and action chunking," *arXiv preprint arXiv:2309.01918*, 2023.
 - [37] H. Bharadhwaj, R. Mottaghi, A. Gupta, and S. Tulsiani, "Track2act: Predicting point tracks from internet videos enables diverse zero-shot robot manipulation," *arXiv preprint arXiv:2405.01527*, 2024.
 - [38] T. Zhang, Z. McCarthy, O. Jow, D. Lee, X. Chen, K. Goldberg, and P. Abbeel, "Deep imitation learning for complex manipulation tasks from virtual reality teleoperation," 2018.
 - [39] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," in *Proceedings of Robotics: Science and Systems (RSS)*, 2023.
 - [40] E. Jang, A. Irpan, M. Khansari, D. Kappler, F. Ebert, C. Lynch, S. Levine, and C. Finn, "Bc-z: Zero-shot task generalization with robotic imitation learning," in *Conference on Robot Learning*. PMLR, 2022, pp. 991–1002.
 - [41] C. Devin, P. Abbeel, T. Darrell, and S. Levine, "Deep object-centric representations for generalizable robot learning," *CoRR*, vol. abs/1708.04225, 2017.
 - [42] Y. Duan, M. Andrychowicz, B. C. Stadie, J. Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba, "One-shot imitation learning," *CoRR*, vol. abs/1703.07326, 2017.
 - [43] H.-S. Fang, C. Wang, H. Fang, M. Gou, J. Liu, H. Yan, W. Liu, Y. Xie, and C. Lu, "Anygrasp: Robust and efficient grasp perception in spatial and temporal domains," *IEEE Transactions on Robotics (T-RO)*, 2023.
 - [44] J. Tremblay, T. To, B. Sundaralingam, Y. Xiang, D. Fox, and S. Birchfield, "Deep object pose estimation for semantic robotic grasping of household objects," *CoRR*, vol. abs/1809.10790, 2018.
 - [45] S. Tyree, J. Tremblay, T. To, J. Cheng, T. Mosier, J. Smith, and S. Birchfield, "6-dof pose estimation of household objects for robotic manipulation: An accessible dataset and benchmark," in *International Conference on Intelligent Robots and Systems (IROS)*, 2022.
 - [46] D. Bauer, T. Patten, and M. Vincze, "Reagent: Point cloud registration using imitation and reinforcement learning," 2021.
 - [47] Y. Ju, K. Hu, G. Zhang, G. Zhang, M. Jiang, and H. Xu,

- “Robo-abc: Affordance generalization beyond categories via semantic correspondence for robot manipulation,” in *European Conference on Computer Vision*. Springer, 2025, pp. 222–239.
- [48] D. Pomerleau, “An autonomous land vehicle in a neural network,” *Advances in Neural Information Processing Systems*, vol. 1, 1998.
- [49] N. M. M. Shafiullah, S. Feng, L. Pinto, and R. Tedrake, “Supervised policy learning for real robots,” July 2024, tutorial presented at the Robotics: Science and Systems (RSS), Delft. [Online]. Available: <https://supervised-robot-learning.github.io>
- [50] A. Iyer, Z. Peng, Y. Dai, I. Guzey, S. Haldar, S. Chintala, and L. Pinto, “Open teach: A versatile teleoperation system for robotic manipulation,” *arXiv preprint arXiv:2403.07870*, 2024.
- [51] Z. J. Cui, Y. Wang, N. M. M. Shafiullah, and L. Pinto, “From play to policy: Conditional behavior generation from uncurated robot data,” *arXiv preprint arXiv:2210.10047*, 2022.
- [52] S. Lee, Y. Wang, H. Etukuru, H. J. Kim, N. M. M. Shafiullah, and L. Pinto, “Behavior generation with latent actions,” *arXiv preprint arXiv:2403.03181*, 2024.
- [53] T. Lindeberg, *Scale Invariant Feature Transform*, 05 2012, vol. 7.
- [54] B. Zitova and J. Flusser, “Image registration methods: a survey,” *Image and vision computing*, vol. 21, no. 11, pp. 977–1000, 2003.
- [55] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “Nerf: Representing scenes as neural radiance fields for view synthesis,” *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [56] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, “3d gaussian splatting for real-time radiance field rendering,” *ACM Trans. Graph.*, vol. 42, no. 4, pp. 139–1, 2023.
- [57] A. W. Harley, Z. Fang, and K. Fragkiadaki, “Particle video revisited: Tracking through occlusions using point trajectories,” in *ECCV*, 2022.
- [58] Y. Zheng, A. W. Harley, B. Shen, G. Wetzstein, and L. J. Guibas, “Pointodyssey: A large-scale synthetic dataset for long-term point tracking,” in *ICCV*, 2023.
- [59] C. Doersch, Y. Yang, M. Vecerik, D. Gokay, A. Gupta, Y. Aytar, J. Carreira, and A. Zisserman, “Tapir: Tracking any point with per-frame initialization and temporal refinement,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 10 061–10 072.
- [60] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, “Segment anything,” 2023.
- [61] Y. Fu, Y. Lei, T. Wang, W. J. Curran, T. Liu, and X. Yang, “Deep learning in medical image registration: a review,” *Physics in Medicine & Biology*, vol. 65, no. 20, p. 20TR01, 2020.
- [62] S. Huang, L. Yang, B. He, S. Zhang, X. He, and A. Shrivastava, “Learning semantic correspondence with sparse annotations,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022.
- [63] K. Gupta, V. Jampani, C. Esteves, A. Shrivastava, A. Makadia, N. Snaveley, and A. Kar, “Asic: Aligning sparse in-the-wild image collections,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 4134–4145.
- [64] J. K. Aggarwal and Q. Cai, “Human motion analysis: A review,” *Computer vision and image understanding*, vol. 73, no. 3, pp. 428–440, 1999.
- [65] A. Yilmaz, O. Javed, and M. Shah, “Object tracking: A survey,” *Acm computing surveys (CSUR)*, vol. 38, no. 4, pp. 13–es, 2006.
- [66] D. Nistér, O. Naroditsky, and J. Bergen, “Visual odometry,” in *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, vol. 1. Ieee, 2004, pp. I–I.
- [67] A. Khazatsky, K. Pertsch, S. Nair, A. Balakrishna, S. Dasari, S. Karamcheti, S. Nasiriany, M. K. Srirama, L. Y. Chen, K. Ellis, *et al.*, “Droid: A large-scale in-the-wild robot manipulation dataset,” *arXiv preprint arXiv:2403.12945*, 2024.
- [68] S. Kumar, J. Zamora, N. Hansen, R. Jangir, and X. Wang, “Graph inverse reinforcement learning from diverse videos,” in *Conference on Robot Learning*. PMLR, 2023, pp. 55–66.
- [69] M. Sieb, X. Zhou, A. Huang, O. Kroemer, and K. Fragkiadaki, “Graph-structured visual imitation,” in *Conference on Robot Learning*, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:196470945>
- [70] F. Pistilli and G. Averta, “Graph learning in robotics: A survey,” *IEEE Access*, vol. PP, pp. 1–1, 01 2023.
- [71] Z.-H. Yin and P. Abbeel, “Offline imitation learning through graph search and retrieval,” *Robotics: Science and Systems*, 2024.