2026 IEEE/SICE International Symposium on System Integration (SII)
January 11-14, 2026. Cancún, México

An Object Placement Optimization System for Efficient and Unbiased Imitation Learning Data Collection

Hiromasa Yamaguchi[1], Yuga Yano[1] and Hakaru Tamukoh[12]

Abstract— **Training data diversity is an important factor in improving the performance of imitation learning. However, object placement diversity and systems that support object placement during data collection have not been sufficiently investigated. In this paper, we analyze the effectiveness of object placement diversity in imitation learning. We evaluate how different placement conditions affect task success using a simulator. Based on the results, we design optimal placement conditions and propose the object placement support system. The proposed system enabled more effective and efficient object placement than human-judged placement, and achieved comparable performance with less data.**

I. Introduction

In recent years, social demand for robots has increased as a solution to labor shortages and hazardous tasks caused by a declining birthrate and an aging population [1]. To meet this demand, technological developments in areas such as perception, manipulation, and mobility have been actively pursued [2], [3], [4], [5], [6]. Among these, imitation learning has recently attracted significant attention [7], [8]. Imitation learning is a method in which a robot learns a policy directly from demonstrations of tasks performed by a human operator. This approach can generate flexible behavior without manually designing complex perception, decision-making, or control processes. Recent methods such as diffusion policy [9], Action Chunking with Transformers (ACT)[10], and π_0[11] have demonstrated high-accuracy behavior generation through imitation learning.

Previous work has established that diversity in object placement within imitation learning datasets is a key factor influencing policy generalization and enhancing task performance [12], [13], [14]. However, most studies have focused on increasing diversity by collecting large amounts of data, while the efficiency of data collection has received less attention. In this study, we analyze how different placement conditions affect task success through preliminary experiments in a simulator. Based on the identified placement requirements, we propose an object placement support system that assists data collection in physical environments. The system aims to reduce placement bias during data collection and enable the acquisition of diverse and effective training data.

This paper is based on results obtained from project JPNP16007 commissioned by the New Energy and Industrial Technology Development Organization (NEDO). This work was supported by JSPS KAKENHI Grant Numbers 23H03468 and 23K18495. This work was supported JST ALCANext Grant Number JPMJAN23F3. This work was supported by JST SPRING, Japan Grant Number JPMJSP2154.

[1]All authors are with Kyushu Institute of Technology, Fukuoka, Japan. {yamaguchi.hiromasa512, yano.yuuga158,}@mail.kyutech.jp and tamukoh@brain.kyutech.ac.jp

[2]Hakaru Tamukoh is also affiliated with the Research Center for Neuromorphic AI Hardware, Fukuoka, Japan.

II. Related works

A. Imitation learning models

Various imitation learning models have been proposed in recent years [9], [10], [11], [15], [16]. For example, Diffusion Policy [9] uses diffusion models to generate diverse behaviors. ACT [10] predicts fixed-length action chunks to capture long-term structure and produce smooth, consistent behaviors. π_0 [11] is a pretrained imitation learning model designed for a wide range of robot tasks, enabling rapid adaptation to new tasks with only a few training data. Because it is publicly available and pretrained, π_0 reduces variability caused by initialization and training dynamics, and thus provides more stable performance than ACT or Diffusion Policy, which require training from scratch.

B. Importance of object placement in imitation learning

Saxena et al. [12] demonstrated that diversity in camera viewpoints and in the spatial arrangement of objects within large-scale imitation learning datasets has a significant impact on downstream task performance. In particular, a lack of diversity in object placements constrains the generalization ability of learned behaviors and decreases the task success rate. However, clear guidelines, such as the required level of diversity and the desirable distribution of object placements, are still insufficient. To address this issue, this study systematically evaluates the effect of object placement diversity through both simulation-based experiments and physical world validation.

III. Preliminary experiments

We conduct preliminary experiments to analyze how placement bias affects task success, with the goal of improving the efficiency of object placement. In these experiments, we chose a **transfer cube** task, in which a robot grasps a cube with one arm and transfers it to the opposite arm. This task is based on the ACT [10]. Figure 1 shows the object size and placement area used in this experiment. We collect RGB images and robot joint angles as training data and fine-tune the π_0 [11] model. The trained model predicts target robot joint angles from the current RGB images, along with robot joint angles. We obtain the pretrained π_0 model from hugging face [17]. We trained the model for 50,000 steps with a batch size of 8. Both training and evaluation were conducted on a server equipped with an NVIDIA A100 GPU. The success rate of each table is defined as the proportion of successful trials

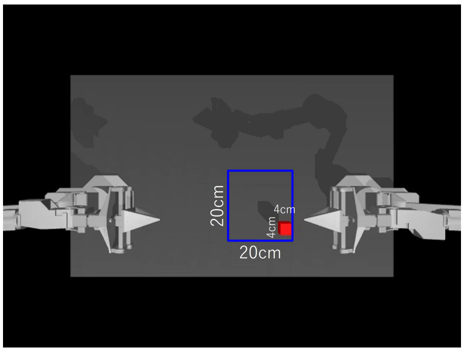

Fig. 1. Object size and placement area

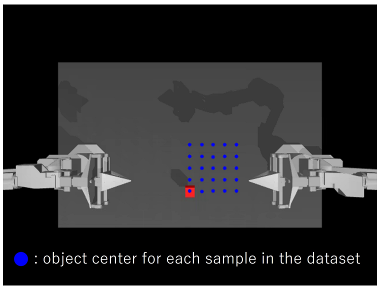

Fig. 2. Visualization of object center positions as blue dots for the Uniform25

out of 100, where a trial is considered successful if the cube is successfully handed over. A higher success rate indicates that the method consistently achieves the desired outcomes. To reduce the influence of randomness in a single trial, we trained models using three different random seeds. Each seed generated a different shuffling of the training data and resulted in different batch partitions. For each shuffled dataset, we trained a separate model and conducted 100 evaluation trials. The final results are reported as the average success rate over all runs.

A. Data collection conditions

To analyze the impact of object placement from multiple perspectives, we designed several conditions based on three criteria. For each condition, we collected training data and fine-tuned π_0 model.

1) Effect of placement uniformity: We compared a dataset in which 25 objects were placed randomly (`Random25`) with one in which the objects uniformly covered the entire placement area (`Uniform25`). Figure 2 shows the object placement of `Uniform25`.

2) Effect of data quantity and bias: We compared `Uniform25` with `Uniform25+Random25`, in which 25 randomly placed objects were added to `Uniform25`.

3) Sufficiency of coverage rate: We created datasets with coverage rates restricted to 90%, 80%, and 70%, and compared the task success rate for each fine-tuned model. Figure 3 shows the object center positions for each coverage rate. Each condition was constructed by randomly removing a specified proportion of samples from the `Uniform25` dataset.

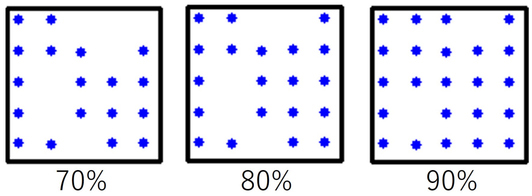

Fig. 3. Object center positions visualized as blue dots in each coverage rate

TABLE I
SUCCESS RATES FOR UNIFORM AND RANDOM PLACEMENTS
(MEAN $\pm$ STD. OVER 3 SEEDS)

Placement	Success rate [%]
Random25	73.0 ± 5.3
Uniform25	**88.6** ± **4.5**

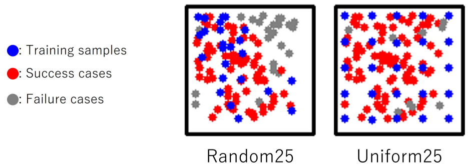

Fig. 4. Relationship between object placements in training and test data for the placement uniformity experiment

The 80% condition was created by removing 10% of the samples from the 90% dataset, and the 70% condition by removing 10% from the 80% dataset.

B. Results and discussion

1) Effect of placement uniformity: Table I summarizes the success rates for uniform and random object placements. Figure 4 shows the relationship between training samples and test samples under different placement conditions. `Uniform25` had an 15.6-point higher success rate than `Random25`. This result indicates that uniform placement with full coverage is more effective for imitation learning than random placement.

2) Effect of data quantity and bias: Table II summarizes the success rates for different data quantities and biases. Figure 5 shows the relationship between training samples and test samples under different data quantity and bias conditions. `Uniform25` achieved a success rate of 88.6%, whereas `Uniform25+Random25` dropped to 84.3%. This comparison shows that simply increasing the dataset size does not necessarily improve performance, and that biased data can negatively affect policy generalization. These results suggest that balanced and uniform object placement is more important for improving imitation learning performance than relying solely on dataset size.

3) Sufficiency of coverage rate: Table III summarizes the success rates for different coverage rates. Figure 7 shows the relationship between training samples and test samples under each coverage condition. All conditions with restricted coverage resulted in lower success rates compared to `Uniform25` with 100% coverage (88.6%), which can be attributed to

TABLE II

Success rate as a function of data quantity and placement bias (mean ± std. over 3 seeds)

Placement	Success rate [%]
Uniform25 + Random25	84.3 ± 2.1
Uniform25	**88.6** ± **4.5**

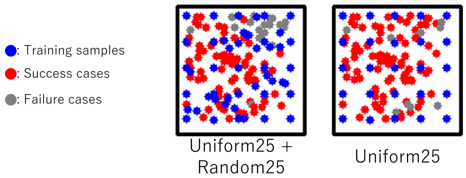

Fig. 5. Relationship between object placements in training and test data for the data quantity and placement bias experiment

TABLE III

Success rate for different coverage rates (mean ± std. over 3 seeds)

Coverage rate	Success rate [%]
70% Coverage	80.0 ± 1.7
80% Coverage	81.0 ± 4.0
90% Coverage	85.0 ± 7.0
100% Coverage (Uniform25)	**88.6** ± **4.5**

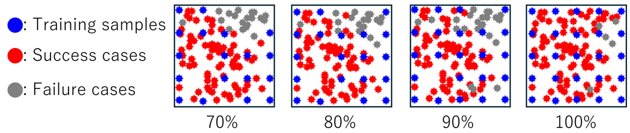

Fig. 7. Relationship between object placements in training and test data for the coverage experiment

uncovered areas and spatial bias. In all cases except 100%, the task success rate dropped significantly when the block was placed at the top-right. We think that the failure is due to missing training data in the top-right region. These results suggest that avoiding uncovered regions and spatial bias is important for stable performance in imitation learning.

Based on these results, effective data collection for imitation learning requires object placement that (1) ensures uniform coverage of the entire placement area, and (2) avoids adding biased or random samples.

IV. Proposed system

Preliminary experiments revealed that spatial uniformity in object placement affects the generalization performance of policies in imitation learning. However, in data collection in a physical world, an operator decides object placement, which may result in biased and redundant position selection. To address this, we propose an object placement support system designed to facilitate the collection of highly representative training data that contributes to improved imitation learning performance.

The proposed system utilizes a spatial density heatmap generated from the history of past object placements. Based on this heatmap, the system automatically estimates uncovered regions and recommends the next placement positions to the user. This approach enables efficient collection of spatially uniform data with minimal bias, without relying on human judgment. Figure 6 shows the proposed system, which consists of five steps described below.

1) **ROI specification**: A user specifies the region of interest (ROI) for object placement on the graphical user interface (GUI).
2) **Object position detection**: The system detects objects of a specified color within the ROI in the RGB image and computes their center coordinates.
3) **Heatmap generation**: For each detected placement point, we apply a rectangular 2D Gaussian kernel centered on it to assign values within the kernel area. An object is represented as a Gaussian patch centered at its location. The kernel size is determined based on the dimensions of the bounding box, and each patch is normalized to have a peak value of 1. All patches are superimposed to generate the spatial density heatmap.
4) **Placement recommendation**: The system recommends the next placement position by selecting the region with the lowest accumulated density within the area corresponding to the object's size.
5) **Redundancy avoidance**: In the generated heatmap, the proportion of pixels with values exceeding zero within the ROI is calculated. If this proportion surpasses the threshold α, the area is regarded as sufficiently covered, and data collection is terminated.

We designed the system to satisfy the requirements identified in the preliminary experiments: (1) ensure uniform coverage of the entire placement area, and (2) avoid biased or random samples.

V. Experimental settings

In this study, we conducted two types of experiments to verify the effectiveness of the proposed system. We used the ALOHA2 [18] in a physical world, the same type of robot as in the preliminary experiments. The horizontal and vertical standard deviations of the Gaussian patches are set to one-fourth of the object's width and height, respectively. This setting ensures that the spread of each distribution roughly matches the object size.

A. Experiment 1: Appropriateness and efficiency in object placement

This experiment, which involved only placing objects, evaluated whether the proposed system can achieve more uniform and efficient placement compared with human-judged placement. We asked three participants to place objects as uniformly as possible until a complete fill. In the proposed system, we set the coverage rate threshold α to 0.8, since

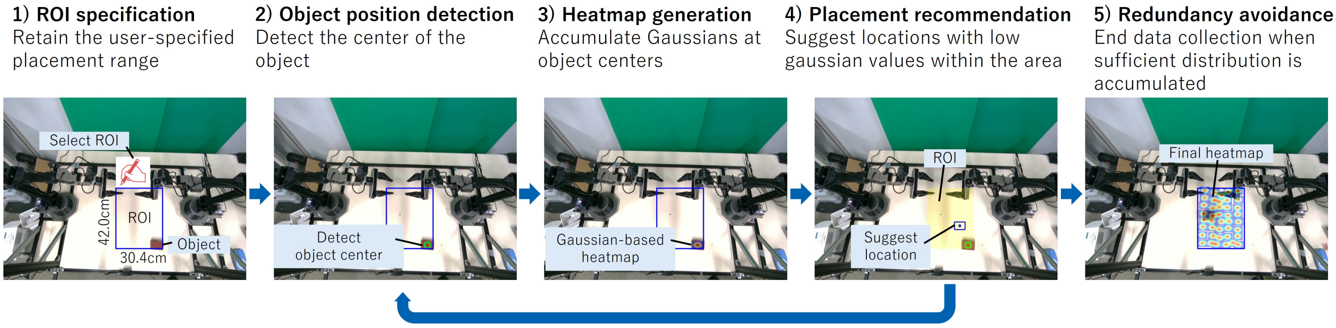

Fig. 6. Proposed object placement support system

TABLE IV
COMPARISON OF HUMAN-JUDGED AND SYSTEM-GENERATED PLACEMENT RESULTS

Method / Participant	# of placements ↓	coverage rate [%] ↑	overlap[%] ↓
Participant A	**26**	36.90	23.64
Participant B	44	79.60	22.29
Participant C	31	64.59	**9.54**
Proposed system	42	**80.30**	17.12

TABLE V
SUCCESS RATE COMPARISON FOR DIFFERENT OBJECT PLACEMENT METHODS

Condition	Training samples ↓	Success rate[%] ↑
Proposed system	42	86.67
Human-judged (Participant D)	52	**90.00**
Human-judged (Participant E, original)	**32**	53.33
Human-judged (Participant E, matched to system count)	42	50.00

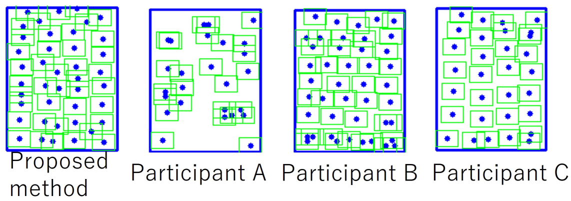

Fig. 8. Visualization of object center positions

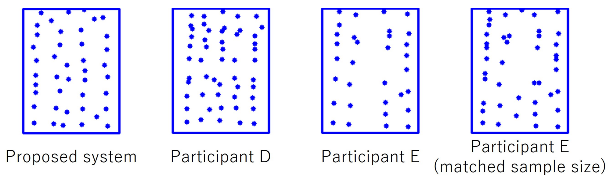

Fig. 9. Object center positions in the training dataset

objects are placed by hand and small misalignments and gaps are unavoidable, making exact 100% coverage impractical. The evaluation metrics are as follows.

- **coverage rate**: The ratio of the ROI area where objects are placed.
- **Overlap rate**: The ratio of the overlapping area to the covered area:

$$\text{Overlap rate} = \frac{A_{\text{overlap}}}{A_{\text{covered}}}$$

  Here, A_{covered} is the region of the placement area covered at least once, and A_{overlap} is the region covered at least twice.
- **Number of placements**: The total number of object placements required to complete the placement process.

Experimental conditions: The placement area was 30.4 cm× 42.0 cm, and the target object size was 7.5 cm × 5.1 cm × 4.9 cm.

B. Experiment 2: Imitation learning performance

This experiment examined the effect of different data collection methods on imitation learning performance. The target task involved the robot grasping an object and transferring it to the opposite arm. The same operator collected all training data, following a consistent motion policy throughout. We prepared four datasets for comparison. The first was collected using the proposed system. The second and third were based on human-judged placement by Participant D and Participant E, respectively. If the number of placements by a participant was smaller than that of the proposed system, we collected additional samples by asking the participant to place more objects until the sample count matched. We fine-tuned the pretrained π_0 model with each dataset. After training, we executed the same task 30 times under the same environment for each condition and evaluated the task success rate. The training conditions were identical to those in the preliminary experiments. We trained on a server with an NVIDIA H200 GPU and evaluated on a desktop PC with an NVIDIA GeForce RTX 4080 GPU.

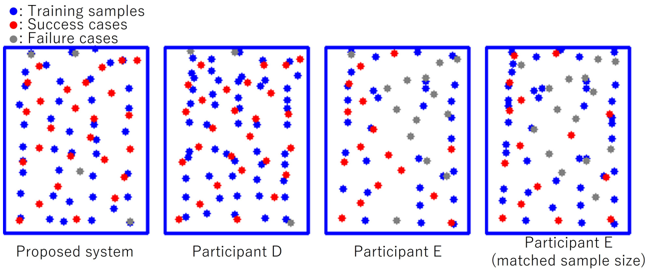

Fig. 10. Object center positions in the training and test samples

VI. Results

A. Experiment 1: Appropriateness and efficiency in object placement

Figure 8 shows the visualization of object center positions, object shapes, and placement area. Table IV summarizes the evaluation of human-judged and system-supported object placement. The results show that the proposed system outperformed human-judged placement in terms of coverage rate (80.30%). For the overlap rate (17.12%), the proposed system is worse than Participant C (9.54%) but better than Participant A (23.64%) and Participant B (22.29%). Considering these two metrics, the proposed system achieved the best balance between high coverage and low overlap, resulting in efficient placement with minimal redundancy.

B. Experiment 2: Imitation learning performance

Figure 9 shows the distribution of object center positions in the training dataset. The first dataset, collected using the proposed system, contained 42 samples. The second and third datasets, collected through human-judged placement by Participant D and Participant E, contained 52 and 32 samples, respectively. The fourth dataset was based on additional placements by Participant E to match the sample count of the proposed system, resulting in 42 samples.

Table V summarizes the success rates under different object placement conditions. Figure 10 shows the distribution of success and failure cases. The proposed system had a success rate of 86.67% (26/30). Human-judged placement by Participant D had a success rate of 90.00% (27/30). Human-judged placement by Participant E had a success rate of 53.33% (16/30). Additional placements by Participant E, matching the sample count of the proposed system, had a success rate of 50.00% (15/30).

VII. Discussion

To cover the entire placement area, participants B and C placed objects starting from one edge, incrementally shifting the position each time. Participant A placed objects in dispersed locations rather than following a fixed sequence. Incrementally shifting placement positions can achieve relatively broad coverage of the entire placement area. This strategy is effective when verifying only object placement, as in Experiment 1. However, when collecting training data, the operator needs to manipulate the object using the robot. As a result, the exact position of the object before data collection is unknown, and it is difficult to place the object very close to its previous location. In Experiment 2, we collected training data using the robot while adopting a strategy of incrementally shifting placement positions. As shown in Fig. 9, Participant D placed objects in a generally sequential manner, but with frequent overlaps. In contrast, Participant E tended to avoid the central region, leaving it largely uncovered. This outcome was due to the difficulty of tracking exact past positions during data collection. These results indicate that the proposed system can achieve uniform coverage of the entire placement area, which is difficult to accomplish when collecting data with a robot through intuitive human placement.

In experiment 2, placements of participant D resulted in a slightly higher success rate than the proposed system. However, ten additional training samples were needed to obtain only one more successful trial. This result suggests that, without the proposed system, humans may collect more data than necessary, resulting in redundant samples. On the other hand, the significantly lower success rate with placements by Participant E indicates that inconsistent or biased placement can hinder learning performance. Even when the number of samples was increased to match that of the proposed system, the success rate did not improve, highlighting the importance of placement diversity and coverage rather than quantity alone.

VIII. Conclusion

This study examined how specific object placement conditions influence imitation learning performance. Based on these results, we proposed the object placement support system that facilitates effective and unbiased placement for physical world data collection. The proposed system achieved efficient placement with minimal redundancy through a balance of high coverage and low overlap. Imitation learning models trained on data collected with this system achieved task success rates comparable to or higher than those trained on human-judged placement data.

However, this study focused on a single task and a limited range of placement scenarios, which constrains the generality of the findings. To overcome this limitation, we should explore multi-task settings and dynamic placement environments in future research that better reflect real-world diversity. In addition, extending the proposed system to fully three-dimensional spatial arrangements will be essential for applications involving aerial or stacked object configurations, thereby enhancing its effectiveness in complex physical environments.

References

[1] D. Macis, S. Perilli, and C. Gena, "Employing socially assistive robots in elderly care," in *Adjunct proceedings of the 30th ACM conference on user modeling, adaptation and personalization*, 2022, pp. 130–138.

[2] T. Ono, D. Kanaoka, T. Shiba, S. Tokuno, Y. Yano, A. Mizutani, I. Matsumoto, H. Amano, and H. Tamukoh, "Solution of world robot challenge 2020 partner robot challenge (real space)," *Advanced Robotics*, vol. 36, no. 17-18, pp. 870–889, 2022.

[3] Y. Yano, K. Isomoto, T. Ono, and H. Tamukoh, "Autonomous waiter robot system for recognizing customers, taking orders, and serving food," in *Robot World Cup*. Springer, 2023, pp. 252–261.

[4] K. Yamao, D. Kanaoka, K. Isomoto, and H. Tamukoh, “Object positions interpretation system for service robots through targeted object marking,” in *Proceedings of the 2025 IEEE/SICE International Symposium on System Integration (SII)*, 2025, pp. 998–1003.

[5] Y. Yano, A. Mizutani, Y. Fukuda, D. Kanaoka, T. Ono, and H. Tamukoh, “Unified understanding of environment, task, and human for human-robot interaction in real-world environments,” in *2024 33rd IEEE International Conference on Robot and Human Interactive Communication (ROMAN)*, 2024, pp. 224–230.

[6] A. Mizutani, Y. Tanaka, H. Tamukoh, O. Nomura, K. Tateno, and T. Morie, “A hippocampus-inspired environment-specific knowledge acquisition system utilizing common knowledge with contextual information,” in *Proceedings of the 2024 International Joint Conference on Neural Networks (IJCNN)*, 2024.

[7] B. Fang, S. Jia, D. Guo, M. Xu, S. Wen, and F. Sun, “Survey of imitation learning for robotic manipulation,” *International Journal of Intelligent Robotics and Applications*, vol. 3, no. 4, pp. 362–369, 2019.

[8] M. Zare, P. M. Kebria, A. Khosravi, and S. Nahavandi, “A survey of imitation learning: Algorithms, recent developments, and challenges,” *IEEE Transactions on Cybernetics*, vol. 54, no. 12, pp. 7173–7186, 2024.

[9] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, “Diffusion Policy: Visuomotor policy learning via action diffusion,” *Robotics: Science and Systems XIX*, 2023.

[10] T. Zhao, V. Kumar, S. Levine, and C. Finn, “Learning fine-grained bimanual manipulation with low-cost hardware,” *Robotics: Science and Systems XIX*, 2023.

[11] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter *et al.*, “π_0: A vision-language-action flow model for general robot control,” *arXiv preprint arXiv:2410.24164*, 2024.

[12] V. Saxena, M. Bronars, N. R. Arachchige, K. Wang, W. C. Shin, S. Nasiriany, A. Mandlekar, and D. Xu, “What matters in learning from large-scale datasets for robot manipulation,” *arXiv preprint arXiv:2506.13536*, 2025.

[13] F. Lin, Y. Hu, P. Sheng, C. Wen, J. You, and Y. Gao, “Data scaling laws in imitation learning for robotic manipulation,” in *Proceedings of The Thirteenth International Conference on Learning Representations*, 2025.

[14] J. Gao, A. Xie, T. Xiao, C. Finn, and D. Sadigh, “Efficient data collection for robotic manipulation via compositional generalization,” *arXiv preprint arXiv:2403.05110*, 2024.

[15] V. Vosylius and E. Johns, “Instant policy: In-context imitation learning via graph diffusion,” in *The Thirteenth International Conference on Learning Representations*, 2025.

[16] A. Xie, O. Rybkin, D. Sadigh, and C. Finn, “Latent diffusion planning for imitation learning,” in *Forty-second International Conference on Machine Learning*, 2025.

[17] Hugging Face. lerobot/pi0. Accessed: Aug. 15, 2025. [Online]. Available: https://huggingface.co/lerobot/pi0

[18] J. Aldaco, T. Armstrong, R. Baruch, J. Bingham, S. Chan, K. Draper, D. Dwibedi, C. Finn, P. Florence, S. Goodrich *et al.*, “Aloha 2: An enhanced low-cost hardware for bimanual teleoperation,” *arXiv preprint arXiv:2405.02292*, 2024.