

Beyond the Majority: Long-tail Imitation Learning for Robotic Manipulation

Junhong Zhu^{1*}, Ji Zhang^{2*†}, Jingkuan Song³, Lianli Gao^{1‡}, Heng Tao Shen³

Abstract—While generalist robot policies hold significant promise for learning diverse manipulation skills through imitation, their performance is often hindered by the long-tail distribution of training demonstrations. Policies learned on such data, which is heavily skewed towards a few data-rich head tasks, frequently exhibit poor generalization when confronted with the vast number of data-scarce tail tasks. In this work, we conduct a comprehensive analysis of the pervasive long-tail challenge inherent in policy learning. Our analysis begins by demonstrating the inefficacy of conventional long-tail learning strategies (e.g., re-sampling) for improving the policy’s performance on tail tasks. We then uncover the underlying mechanism for this failure, revealing that data scarcity on tail tasks directly impairs the policy’s spatial reasoning capability. To overcome this, we introduce Approaching-Phase Augmentation (APA), a simple yet effective scheme that transfers knowledge from data-rich head tasks to data-scarce tail tasks without requiring external demonstrations. Extensive experiments in both simulation and real-world manipulation tasks demonstrate the effectiveness of APA. Our code and demos are publicly available at: <https://mldxy.github.io/Project-VLA-long-tail/>.

I. INTRODUCTION

Recent advances in robotics have been significantly driven by large-scale imitation learning, with generalist robot policies emerging as a dominant paradigm for creating general-purpose agents [1], [2], [3], [4], [5]. By training on vast, diverse datasets of human demonstrations, these policies hold the promise of enabling robots to perform a wide range of manipulation tasks from unconstrained natural language instructions. This data-driven approach has already demonstrated remarkable success in acquiring a broad spectrum of manipulation skills [6], [7], [8], paving the way for more flexible and dexterous robotic systems.

However, a critical and commonly disregarded challenge hinders the real-world applicability of these models: the naturally long-tail distribution of demonstration data. In large-scale robotics dataset, a small number of common head tasks (e.g., pick up the bowl/plate) are demonstrated frequently, while the vast majority of tail tasks (e.g., put the wine bottle on the rack) are represented by only a handful of examples [6], [9], [10]. As we will demonstrate, policies trained on such imbalanced data often perform unreliably on tail tasks [11], [12]. Fig. 1 shows that when shifting from a full (or balanced) to a long-tail (or imbalanced) dataset, the relative performance degradation

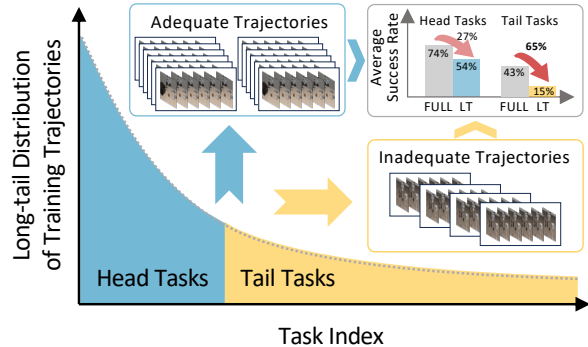


Fig. 1. **Influence of long-tail demonstrations on policy generalization.** Owing to the skewed distribution of training demonstrations, the performance degradation of generalist robot policies is particularly pronounced for tail tasks relative to those in the head.

is far more pronounced for tail tasks than for head tasks. This long-tail problem represents a fundamental challenge to developing generalist and reliable robotic agents.

A natural first step to address this issue would be to apply conventional long-tail learning strategies that have been successful in computer vision, such as class re-balancing method or data augmentation techniques [13], [12]. However, we find that these methods are largely ineffective for policy learning. Specifically, our evaluation of re-sampling strategies from [13] revealed limited efficacy, as they merely duplicate existing data without introducing the variational diversity necessary for generalization. Furthermore, augmentation techniques like mixup, while successful in classification, are incompatible with robotic control. The linear interpolation of states and actions disregards underlying dynamics, frequently resulting in kinematically infeasible or physically invalid trajectories [14], [15]. This raises the following questions:

What is the core reason for the performance drop of policies trained on long-tail demonstrations? And how can this issue be resolved without relying on external demonstration data?

To answer these questions, we construct a long-tail imitation learning benchmark based on the LIBERO environment [10], comprising 10 diverse manipulation tasks. Through an in-depth investigation, we find that policies fail on tail tasks primarily due to a degradation of their spatial reasoning capabilities. Because of training data scarcity, the model is unable

¹Junhong Zhu and Lianli Gao are with Shenzhen Institute for Advanced Study, University of Electronic Science and Technology of China. ²Ji Zhang is with Southwest Jiaotong University. ³Jingkuan Song and Heng Tao Shen are with Tongji University.

*Equal contribution, [†]Project lead, [‡]Corresponding author.

to learn the precise spatial relationships needed for tail tasks, although it has grasped the general concept from head tasks. Based on this key insight, we introduce Approaching-Phase Augmentation (APA), a simple yet effective scheme designed to improve the spatial reasoning capabilities of robot policies on tail tasks. APA is a self-contained method that leverages demonstrations from head tasks to generate new, high-quality training examples for tail tasks. Crucially, it achieves this without requiring any external demonstration data, making it a practical and efficient solution. We conduct extensive evaluations in both simulation and real-world environments, the achieved results across a total of 16 diverse manipulation tasks demonstrate the effectiveness of our method.

Our contributions in this work are threefold.

- We construct a long-tail imitation learning benchmark based on the LIBERO environment, and identify that data scarcity on tail tasks directly impairs the policy’s spatial reasoning capability.
- We propose Approaching-Phase Augmentation (APA), a simple and effective method that improves spatial reasoning by transferring knowledge from head tasks to tail tasks without requiring external demonstrations.
- Extensive experiments in both simulation and the real-world validate the effectiveness of our APA method.

II. RELATED WORKS

A. Generalist Robot Policies

Recent advances in embodied AI have been driven by robot policies, which integrate perception, language understanding, and control into an end-to-end policy. Pioneering works like RT-1 [16] introduced the tokenization of robot actions and observations for control within a transformer architecture. RT-2 [17] significantly advanced this by reframing robotic control as a vision-language problem, directly transferring web-scale knowledge to novel robotic tasks. Subsequently, Octo [18] was developed as a generalist, multi-task robot policy with a unified architecture, serving as a versatile foundation model adaptable across diverse robot embodiments. OpenVLA [6] established a new state-of-the-art by integrating a powerful vision encoder with a modern LLM, setting a new open-source benchmark for imitation learning performance. SmolVLA [19] presented an efficient Vision-Language-Action model that lowered the barrier for robotic learning by enabling training and deployment on consumer hardware using community-sourced data. Physical Intelligence company proposed π_0 [8], prioritizing high-dexterity control with continuous actions generated via flow-matching networks. π_0 -FAST variant [20] reverted to discrete actions for faster training at the expense of slower inference. Regardless of the model, task success rates are heavily dependent on the quantity and quality of training data.

B. Learning from Long-tail Distributions

Learning from long-tailed data remains a significant challenge in machine learning, where models are prone to bias towards head classes due to the severely imbalanced data distribution [21], [22], [23], [24], [25]. To address this issue: 1)

re-sampling techniques [26], [27], [28] balance the training data distribution by either oversampling tail-class instances or undersampling head-class instances, thereby reducing the frequency disparity across classes; 2) data augmentation strategies [29], [30], [31] enhance the diversity and volume of tail-class samples through generative or transformation-based methods. Both streams strive to mitigate model bias and improve generalization on tail classes, albeit through distinct mechanisms. A more elegant and popular solution is head-to-tail knowledge transfer, which leverages head class data as a reservoir of transferable structure to enhance the learning of tail classes. SMART [32] transferred the semantic covariance of head classes to enrich the feature patterns of tail classes. [33] introduced an embedding-augmentation framework that transferred intra-class angular distributions from head to tail. GLMC [34] mitigated classifier bias by applying a cumulative reweighted loss to soft labels generated from mixed head and tail samples. H2T-FAST [35] performed feature-level style transfer, injecting head-class styles such as textures and color patterns into tail-class content to synthesize augmented features. Copy-Paste [36] performed object-level augmentation by pasting tail objects onto diverse backgrounds, often derived from head images, to increase data diversity.

III. DIAGNOSING THE LONG-TAIL FAILURE MODE

To investigate the long-tail challenge in robot policies, we first formalize the problem and construct imbalanced datasets based on the LIBERO [10] benchmark. Subsequently, we demonstrate that traditional re-sampling methods are ineffective for VLAs, and then conduct a systematic analysis to identify the root causes of performance degradation.

A. Problem formulation

A robot policy π_θ is able to map multimodal observations and a natural language instruction to a sequence of executable actions for a specified task [6]. Let the observation at time t be $o_t \in O$ and the instruction be L . The observation space typically includes visual input (e.g., RGB images $I_t \in \mathbb{R}^{n \times H \times W \times 3}$, where $n \geq 1$) and may include proprioceptive states (e.g., joint angles). Commonly, the action at time t is $a_t \in A$, $a_t \in \mathbb{R}^7$ represents the end-effector pose, which uses 7-DOF for each arm [37], [38]. Each 7-DOF action includes 3-DOF relative translation offsets ($[\Delta x, \Delta y, \Delta z] \in \mathbb{R}^3$, 3-DOF rotation euler angles ($[Roll, Pitch, Yaw] \in \mathbb{R}^3$), and 1-DOF gripper state ($Open/Closed \in \mathbb{R}^1$). For each demonstration d in the dataset D , the trajectory is written as $(\mathbf{o}^{(d)}, \mathbf{a}^{(d)}, L)$ with a sequence of states $\mathbf{o}^{(d)} = (o_1^{(d)}, \dots, o_{T_i}^{(d)})$ and a corresponding sequence of actions $\mathbf{a}^{(d)} = (a_1^{(d)}, \dots, a_{T_i}^{(d)})$, where T_i is the length of the trajectory. We follow the standard language-conditioned imitation learning setting to train a policy π_θ . Its parameters θ are optimized by minimizing:

$$\min_{\theta} J(\theta) = \mathbb{E}_{(\mathbf{o}, \mathbf{a}, L) \sim D} \left[\sum_{t=1}^T \mathcal{L}_{\text{action}}(\pi_{\theta}(s_t, L), a_t) \right], \quad (1)$$

where \mathcal{L}_{act} denotes the action loss function (e.g., Mean Squared Error for continuous actions and Cross-Entropy Loss

TABLE I. The language instructions and the number of demonstrations for each task from LIBERO-Core-FULL and LIBERO-Core-LT datasets.

Task	Language Instruction	(FULL, LT)
Task 1	Pick up the black bowl next to the plate and place it on the plate	(46, 46)
Task 2	Pick up the black bowl next to the cookie box and place it on the plate	(47, 28)
Task 3	Pick up the black bowl on the cookie box and place it on the plate	(45, 19)
Task 4	Pick up the ketchup and place it in the basket	(42, 15)
Task 5	Pick up the alphabet soup and place it in the basket	(47, 11)
Task 6	Push the plate to the front of the stove	(39, 9)
Task 7	Put the bowl on top of the cabinet	(47, 8)
Task 8	Put the cream cheese in the bowl	(39, 7)
Task 9	Put the wine bottle on top of the cabinet	(45, 6)
Task 10	Put the wine bottle on the rack	(38, 5)

for discrete actions). At inference time, the model takes (o_t, L) as input and outputs $\hat{a}_t = \pi_{\theta}(o_t, L)$.

B. A Long-tail Imitation Learning Benchmark

We introduce two task sets based on the LIBERO environment [10]. First, we curate LIBERO-Core-FULL by selecting 10 representative tasks from the LIBERO-Spatial, LIBERO-Object, and LIBERO-Goal splits. Subsequently, we construct LIBERO-Core-LT by generating a long-tailed training set from these 10 tasks. We first order the 10 tasks according to the skill hierarchy of the Open X-Embodiment dataset [9]. Drawing inspiration from prior work in computer vision, such as ImageNet-LT and Places-LT [21], We then sample demonstrations for each task using a Pareto distribution [39], which allocates a varying number of demonstrations to each task. Following the proportional partitioning methodology commonly used in long-tailed recognition, we categorize tasks by their per-task demonstration counts: the top 30% (3 of 10 tasks) constitute the data-rich head, and the remaining 70% (7 of 10 tasks) form the data-scarce tail. The language instructions and the number of training demonstrations for each task are provided in Table I. A comparison of the demonstration distribution between the two datasets is shown in Figure 2.

C. Inefficiency of Re-sampling Methods

To evaluate conventional long-tail strategies, we fine-tune the miniVLA model [40], pre-trained on the LIBERO-90 dataset, on our LIBERO-Core-LT training set using various re-sampling methods. All other hyperparameters are held constant to isolate the effect of the training data distribution.

Following the methodology of [13], we employ a sampling strategy where the probability p_j of sampling a trajectory from task j is given by:

$$p_j = \frac{n_j^q}{\sum_{i=1}^C n_i^q}, \quad (2)$$

where n_j is the number of demonstrations for task j , and C is the total number of tasks. The hyperparameter $q \in [0, 1]$




Fig. 2. Comparison of training demonstration distribution between the LIBERO-Core-FULL (blue bar) and LIBERO-Core-LT (yellow bar) datasets.

TABLE II. Performance of re-sampling methods.

Method	Success Rate (%)
Baseline (Original Distribution)	26.5
Re-sampling ($q = 0.75$)	25.1
Re-sampling ($q = 0.50$)	25.1
Re-sampling ($q = 0.25$)	27.1

controls the degree of re-balancing, where a smaller value of q results in a higher sampling probability for tail tasks with scarce training data. We experiment with several values, setting $q \in \{0.75, 0.5, 0.25\}$, to progressively over-sample from tail tasks. Performance is averaged over 30 rollouts across three distinct random seeds. The results are summarized in Table II. As can be observed, these re-sampling strategies yield marginal performance improvements in only a single case. This suggests that simply increasing exposure to tail-task data is ineffective in improving the policy learning performance on tail tasks.

D. Phase-Wise Failure Analysis for Long-Tail Learning

In this section, we seek to answer the following question: What is the core reason for the performance drop of policies trained on long-tail demonstrations?

To this end, we train two separate model instances for comparison. The first is fine-tuned on our fully balanced dataset, LIBERO-Core-FULL, while the second is trained on its long-tailed variant, LIBERO-Core-LT. To ensure a fair comparison, both models use the same architecture, start from the same pre-trained checkpoint, and employ identical hyperparameters, as shown in Section III-C.

Phase Decoupling. We decouple each task trajectory into two sequential phases. 1) *Target Approaching Phase*: encompassing all actions until the robot’s end-effector reaches the immediate vicinity of the primary object. 2) *Subsequent Execution Phase*: including all actions that follow a successful approach to complete the task. Let n_{appr} denote the number of rollouts failing in target approaching phase, n_{exec} the number of rollouts that pass target approaching phase but fail in subsequent execution phase, and n_{succ} the number of fully successful rollouts. From these counts, we compute the

TABLE III. Phase-wise failure probabilities and success rates on LIBERO-Core-FULL and LIBERO-Core-LT datasets.

Task		LIBERO-Core-FULL			LIBERO-Core-LT		
		p_{appr}	p_{exec}	Success Rate (%)	p_{appr}	p_{exec}	Success Rate (%)
Head	Task 1	-	-	78.9±0.15	-	-	56.7±0.78
	Task 2	-	-	78.9±0.70	-	-	60.0±0.00
	Task 3	-	-	63.3±0.33	-	-	45.6±0.15
Tail	Task 4	0.078	0.244	67.8±0.04	0.111	0.322	56.7±0.78
	Task 5	0.022	0.289	68.9±0.04	0.322	0.533	14.4±0.48
	Task 6	0.044	0.667	28.9±0.70	0.189	0.556	25.6±0.04
	Task 7	0.100	0.378	52.2±0.04	0.933	0.067	0.0±0.00
	Task 8	0.167	0.422	41.1±0.44	0.433	0.544	2.2±0.15
	Task 9	0.422	0.278	30.0±0.04	0.711	0.244	4.4±0.26
	Task 10	0.089	0.800	11.1±0.00	0.411	0.589	0.0±0.00

unconditional failure probability for target approaching phase (p_{appr}) and the conditional failure probability for subsequent execution phase given success in target approaching phase (p_{exec}), as defined in Equation 3 and Equation 4, respectively.

$$p_{\text{appr}} = P(\text{appr} = \text{fail}) = \frac{n_{\text{appr}}}{n_{\text{appr}} + n_{\text{exec}} + n_{\text{succ}}}, \quad (3)$$

$$p_{\text{exec}} = P(\text{exec} = \text{fail} | \text{appr} = \text{succ}) = \frac{n_{\text{exec}}}{n_{\text{exec}} + n_{\text{succ}}}. \quad (4)$$

Phase-wise Analysis. To quantify the impact of the long-tailed data distribution, we employ Relative Risk (RR) [41], a concept from epidemiology and economics, to measure the increased likelihood of task failure when training on the LIBERO-Core-LT dataset (the exposed group) versus the LIBERO-Core-FULL dataset (the non-exposed group) [42]. An RR value greater than 100% indicates an increased risk, a value of 100% indicates no change, and a value less than 100% suggests a decreased risk. The RR for every phase of each task is defined in Equation 5, where i denotes the i^{th} of the tasks.

$$RR_{\text{phase}}^{(i)} = \frac{p_{\text{phase,L T}}^{(i)}}{p_{\text{phase,FULL}}^{(i)}}. \quad (5)$$

To mitigate the influence of extreme values, we compute the geometric mean of the relative risk across all tail tasks, which is well suited for averaging ratios and provides a more stable measure of central tendency. This average relative risk, RR_{phase} , is calculated over the set of tail tasks indexed from $i = m$ to n :

$$RR_{\text{phase}} = \left(\prod_{i=m}^n RR_{\text{phase}}^{(i)} \right)^{\frac{1}{n-m+1}}. \quad (6)$$

In our setting, the tail tasks are indexed from $m = 4$ to $n = 10$.

We calculated phase-wise failure probabilities, p_{appr} and p_{exec} for each tail task, with the results summarized in Table III. Based on Equation 5 and Equation 6, we found that the relative risk of failure during the target approaching phase (RR_{appr}) reached 400.89% across all tail tasks, compared to 164.34% in the execution phase (RR_{exec}). This indicates that, relative to a policy trained on the full dataset, a policy trained on long-tail data is four times more likely to fail during target approaching. Notably, the target approaching phase

depends most critically on spatial reasoning [43], [44], [45] between the robot and the object. This finding reveals that data scarcity in tail tasks directly undermines the policy’s spatial reasoning capability.

IV. APPROACHING-PHASE AUGMENTATION (APA)

Motivated by those findings, we propose the Approaching-Phase Augmentation (APA) method, which leverages demonstrations from head tasks to generate new, high-quality training examples for tail tasks. The overall process is illustrated in Fig. 3, with the detailed steps outlined below.

Step 1: Head Task Trajectory Segmentation. We begin by randomly selecting a pool of successful execution demonstrations from the data-rich head tasks. For each trajectory, we identify the segment where the robot arm is approaching the target by monitoring its proprioceptive perception (e.g., the gripper’s degree of openness or closure). This provides us with a collection of successful approaching-phase demonstrations.

Step 2: Tail to Head Object Grafting. In this step, we generate new training demonstrations by taking a trajectory segment from head tasks and replacing its original object with one from tail tasks. In the simulated LIBERO benchmark, this involves directly replacing the object asset. The new object’s initial position is inherited from the source trajectory’s original object, while its rotational orientation is specified by the target from tail tasks. The scene is then re-rendered to generate new visual data. For real-world applications, this process can be approximated by identifying the object with YOLOv8 [46] and replacing it via image inpainting. Ultimately, this recombination diversifies the training data for the target approaching phase.

Step 3: Instruction Formatting and Co-Training. The final step is to co-train the policy on a composite dataset. To ensure linguistic consistency, we use a pair of corresponding templates to format the language instructions for the original and augmented trajectories, respectively. For our newly augmented trajectories, which only contain the target approaching phase, we generate instructions using the template: approach the [target_object]. To create a consistent two-phase format across the entire dataset, we modify the instructions for all original demonstrations, instructions for the original demonstrations




Fig. 3. **Overview of the Approaching-Phase Augmentation (APA) pipeline.** Our method involves a three-step process: (1) **Head Task Trajectory Segmentation**, which isolates the target approaching phase from data-rich head tasks; (2) **Tail to Head Object Grafting**, which creates augmented trajectories using objects from tail tasks; and (3) **Instruction Formatting and Co-Training**, which formats the corresponding language instructions and then trains the policy on the combined dataset.

in source long-tail dataset are modified to approach the [target_object] then [verb_phrase_1] and [verb_phrase_2]. The model is then trained on this combined, linguistically consistent dataset.

The primary advantage of APA is its targeted efficiency. By focusing exclusively on augmenting the critical target approaching phase, our method bypasses the complexity of generating full, physically plausible manipulation sequences. Unlike general inpainting methods requiring consistent object properties (e.g., size, mass) for interaction, our approach is unconstrained by these physical requirements. This simple yet effective intervention is possible precisely because it targets the key bottleneck we identified in Section III.

V. EXPERIMENTS

In this section, we conduct extensive experiments in both simulation and real-world environments to evaluate the effectiveness of the proposed APA method.

A. Simulation Experiments

1) **Experimental Setup:** Our experiments are conducted on the LIBERO-Core-LT dataset constructed in Section III-B. We also employ the state-of-the-art miniVLA as our testbed model, which was pretrained on the LIBERO-90 dataset. To evaluate our method, we fine-tune this checkpoint on two different dataset configurations:

- **Baseline:** The model is fine-tuned directly on the original, imbalanced LIBERO-Core-LT dataset.
- **APA (Ours):** The model is fine-tuned on the LIBERO-Core-LT dataset augmented with demonstrations generated by our proposed method.

Both models were fine-tuned for 36 epochs. Each result was averaged over three random seeds from 30 evaluation rollouts.

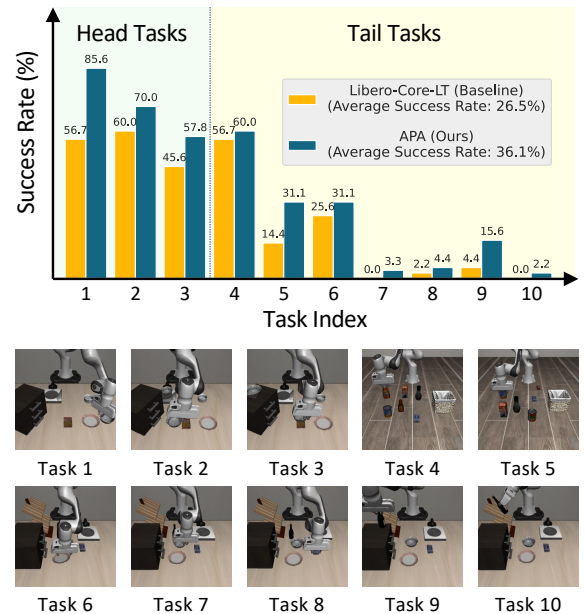


Fig. 4. Effectiveness of our proposed APA method on LIBERO-Core-LT dataset.

2) **Results and Analysis:** The results, summarized in Fig. 4, demonstrate the effectiveness of APA. Our method boosts the average success rate from 26.5% to 36.1%, a substantial relative improvement of 36.2%. Crucially, our approach improves performance on data-scarce tail tasks while simultaneously boosting results on data-rich head tasks. We attribute this positive-sum outcome to our augmentation strategy, which leverages demonstrations from the head as a source. This process appears to provide diverse data for the tail while also reinforcing the learned policies for the head tasks themselves. To confirm that this efficacy is independent

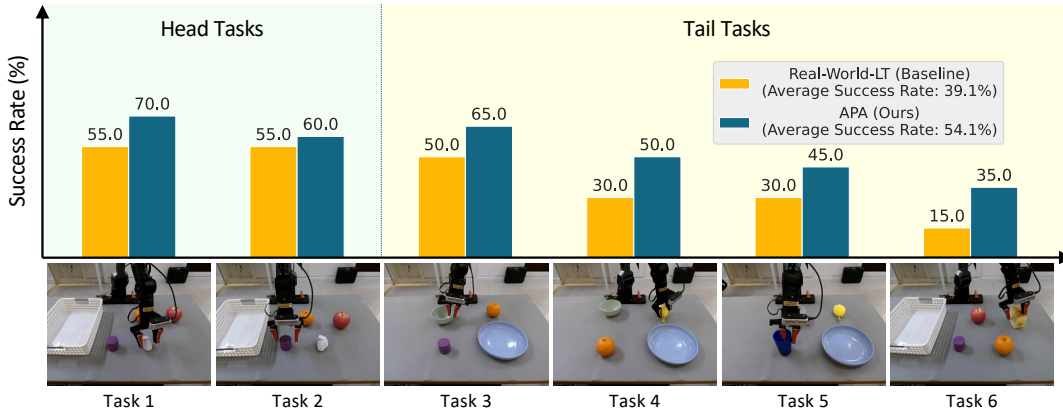


Fig. 5. Effectiveness of our proposed APA method on real-word long-tail tasks.

TABLE IV. Ablation study on the designed components.

Setting	Formatting	Augmentation	Success Rate (%)
Baseline	✗	✗	26.5
+ Formatting	✓	✗	26.0
+ Augmentation	✗	✓	26.9
Ours (Full)	✓	✓	36.1

of the task distribution, we evaluated APA on a randomly shuffled dataset—an experiment detailed in Appendix VI-A that underscores the robustness of our approach.

B. Ablation Studies

We conduct two ablation studies to analyze the key aspects of our APA framework. First, we analyze the individual impact of our core components, Cross-Task Object Grafting and Instruction Formatting. Second, we investigate the effect of varying the number of augmented demonstrations.

1) **Component Analysis:** Our component analysis reveals that the core components of our method are ineffective in isolation and rely on a strong synergy. As detailed in Table IV, ablating either the instruction formatting or the trajectory augmentation yields no significant improvement over the baseline. This is because the policy cannot ground the new language instructions without corresponding visual examples, nor can it effectively utilize the new visual data without an explicit linguistic guide. These findings therefore confirm that the instruction formatting and augmentation components are deeply intertwined and codependent, making both essential for achieving optimal performance.

2) **Effect of the Number of Augmented Demonstrations:** We analyze the impact of the number of augmented demonstrations per task added to the training set. Table V shows that while adding these demonstrations is beneficial, the performance gains are not monotonic, and the success rate begins to decline after peaking. We hypothesize that this decline occurs because an excessive number of augmented demonstrations, which only cover the initial target approaching phase, may skew the training data distribution and impact the overall success rate.

TABLE V. Impact of the num of augmented demonstrations.

Number of Augmented Demons.	0	3	6	9
Success Rate (%)	26.5	28.8	36.1	32.0

C. Real-world Experiments

1) **Experimental Setup:** To evaluate performance in a physical setting, we design a new long-tail benchmark, named `Real-World-LT`. This dataset comprises 6 real-world manipulation tasks, for which we collected demonstrations following a Pareto distribution to create a long-tail structure [21], [39]. Further details on the tasks and trajectory counts are provided in the Appendix VI-B. For the policy, we select the π_0 model, which was pretrained on the real-world data from OXE [9]. This choice parallels our simulation setup (which used the virtual-pretrained miniVLA), ensuring that we use a base model appropriately matched to its domain. We then evaluate and compare two finetuning approaches: **Baseline:** The policy is fine-tuned directly on `Real-World-LT` dataset. **APA (Ours):** The policy is fine-tuned on the same dataset augmented with APA. During evaluation, we conduct 20 rollouts for each task, randomizing the configurations and orientations of objects in each trial to ensure a robust assessment of performance. All Our real-world experiments are conducted on an AGILEX PIPER 6-DOF robot arm equipped with a 1-DOF gripper.

2) **Results and Analysis:** Our real-world experimental results, summarized in Fig. 5, validate the effectiveness of our APA method on a physical robot. Overall, APA achieves an average success rate of 54.1%, a significant improvement over the baseline’s 39.1%, which corresponds to a 38.4% relative gain. Consistent with our simulation findings, APA substantially improves the success rate of tail tasks without compromising performance on the data-rich head tasks. In fact, it enhances performance across the entire distribution, demonstrating its comprehensive benefits. Furthermore, these results provide real-world evidence for our initial hypothesis from simulation: that the performance degradation on tail tasks is disproportionately concentrated in the initial object

TABLE VI. The number of demonstrations for each task from LIBERO-Core-LT-Shuffled dataset.

Task	Language Instruction	Count
Task 1	Pick up the ketchup and place it in the basket	42
Task 2	Pick up the black bowl next to the plate and place it on the plate	29
Task 3	Pick up the alphabet soup and place it in the basket	21
Task 4	Push the plate to the front of the stove	15
Task 5	Put the wine bottle on the rack	12
Task 6	Pick up the black bowl next to the cookie box and place it on the plate	10
Task 7	Pick up the black bowl on the cookie box and place it on the plate	8
Task 8	Put the wine bottle on top of the cabinet	7
Task 9	Put the cream cheese in the bowl	6
Task 10	Put the bowl on top of the cabinet	5

approaching phase, which heavily relies on the model’s spatial reasoning capability. The success of our method confirms that by specifically targeting this bottleneck, APA offers a practical and robust solution for mitigating long-tail challenges in robotic manipulation.

VI. CONCLUSIONS

This paper presents a comprehensive analysis of the long-tail challenge in policy learning. Our in-depth analysis reveals that the primary failure mode of the policy on tail tasks is a degradation of the spatial reasoning capacity. Based on this key insight, we introduce Approaching-Phase Augmentation (APA), a novel and effective method that generates high-quality, targeted training data by grafting tail task objects onto successful approach trajectories from head tasks. Through extensive experiments on both simulated and real-world benchmarks, we show that our APA method significantly boosts the performance on tail tasks without compromising head task success.

APPENDIX

A. Robustness of APA to Different Task Distributions

To verify that the efficacy of our method is not contingent on the specific tasks designated as head in the original long-tail distribution, we evaluate its performance on a new, randomly generated task ordering. We create this new distribution by first shuffling the 10 core tasks and then re-applying Pareto sampling to construct a new long-tailed training set, which we name LIBERO-Core-LT-Shuffled. This procedure effectively reassigns which tasks are data-rich (head) and which are data-scarce (tail), with the new order detailed in Table VI. The results in Fig. 6 demonstrate that our approach still outperforms the baseline under this new distribution. This confirms that the effectiveness of APA is both robust and generalizable.

B. Task Details for the Real-World-LT Dataset

The details of the six real-world manipulation tasks in Real-World-LT dataset are presented in Table VII.




Fig. 6. Effectiveness of our proposed APA method on LIBERO-Core-LT-Shuffled dataset.

TABLE VII. The number of demonstrations for each task from Real-World-LT dataset.

Task	Language Instruction	Count
Task 1	Pick up the spitball and place it in the basket	20
Task 2	Pick up the cylinder and place it in the basket	13
Task 3	Put the bowl on the plate	9
Task 4	Put the lemon on the plate	6
Task 5	Put the cup on the plate	5
Task 6	Pick up the bread and place it in the basket	4

ACKNOWLEDGMENT

This work is supported by the Shenzhen Science and Technology Program (No. JCYJ20240813114208012), the National Natural Science Foundation of China (No. 62506310), the China Postdoctoral Science Foundation (No. 2025M781517), and the Sichuan Science and Technology Program (No. 2026NSFSC1479).

REFERENCES

- [1] Ranjan Sapkota, Yang Cao, Konstantinos I Roulmliotis, and Manoj Karkee. Vision-language-action models: Concepts, progress, applications and challenges. *arXiv preprint arXiv:2505.04769*, 2025.
- [2] Yueen Ma, Zixing Song, Yuzheng Zhuang, Jianye Hao, and Irwin King. A survey on vision-language-action models for embodied ai. *arXiv preprint arXiv:2405.14093*, 2024.
- [3] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale N Fung, and Steven Hoi. Instruct-clip: Towards general-purpose vision-language models with instruction tuning. *NeurIPS*, 36:49250–49267, 2023.
- [4] Moo Jin Kim, Chelsea Finn, and Percy Liang. Fine-tuning vision-language-action models: Optimizing speed and success. *arXiv preprint arXiv:2502.19645*, 2025.
- [5] Jiangyong Huang, Silong Yong, Xiaojian Ma, Xiongkun Linghu, Puhao Li, Yan Wang, Qing Li, Song-Chun Zhu, Baoxiang Jia, and Siyuan Huang. An embodied generalist agent in 3d world. In *ICML*, pages 20413–20451, 2024.
- [6] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, and Pannag Sanketi. Openvla: An open-source vision-language-action model. 2024.
- [7] Qingqing Zhao, Yao Lu, Moo Jin Kim, Zipeng Fu, Zhuoyang Zhang, Yecheng Wu, Zhaoshuo Li, Qianli Ma, Song Han, Chelsea Finn, et al. Cot-vla: Visual chain-of-thought reasoning for vision-language-action models. In *CVPR*, pages 1702–1713, 2025.
- [8] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. $\pi 0$: A vision-language-action flow model for general robot control. *corr*, abs/2410.24164, 2024. doi: 10.48550/arXiv.2410.24164.

- [9] Abby O’Neill, Abdul Rehman, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, et al. Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0. In *ICRA*, pages 6892–6903. IEEE, 2024.
- [10] Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. Libero: Benchmarking knowledge transfer for lifelong robot learning. *NeurIPS*, 36:44776–44791, 2023.
- [11] Qihao Zhao, Yalun Dai, Hao Li, Wei Hu, Fan Zhang, and Jun Liu. Ltgc: Long-tail recognition via leveraging llms-driven generated content. In *CVPR*, pages 19510–19520, 2024.
- [12] Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. Deep long-tailed learning: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 45(9):10795–10816, 2023.
- [13] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217*, 2019.
- [14] Beomjoon Kim and Joelle Pineau. Maximum mean discrepancy imitation learning. In *Robotics: Science and systems*, 2013.
- [15] Jonas Werner, Kun Chu, Cornelius Weber, and Stefan Wermter. Llm-based interactive imitation learning for robotic manipulation. *arXiv preprint arXiv:2504.21769*, 2025.
- [16] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Tomas Jackson, Sally Jesmonth, Nikhil Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Kuang-Huei Lee, Sergey Levine, Yao Lu, Utsav Malla, Deeksha Manjunath, Igor Mordatch, Ofir Nachum, Carolina Parada, Jodilyn Peralta, Emily Perez, Karl Pertsch, Jornell Quiambao, Kanishka Rao, Michael Ryoo, Grecia Salazar, Pannag Sanketi, Kevin Sayed, Jaspier Singh, Sumedh Sontakke, Austin Stone, Clayton Tan, Huong Tran, Vincent Vanhoucke, Steve Vega, Quan Vuong, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and Brianna Zitkovich. Rt-1: Robotics transformer for real-world control at scale. In *arXiv preprint arXiv:2212.06817*, 2022.
- [17] Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayyaan Wahid, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *CoRL*, pages 2165–2183. PMLR, 2023.
- [18] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, et al. Octo: An open-source generalist robot policy. *arXiv preprint arXiv:2405.12213*, 2024.
- [19] Mustafa Shukor, Dana Aubakirova, Francesco Capuano, Pepijn Kooijmans, Steven Palma, Adil Zouitine, Michel Aractingi, Caroline Pascal, Martino Russi, Andres Marafioti, et al. Smolvla: A vision-language-action model for affordable and efficient robotics. *arXiv preprint arXiv:2506.01844*, 2025.
- [20] Karl Pertsch, Kyle Stachowicz, Brian Ichter, Danny Driess, Suraj Nair, Quan Vuong, Oier Mees, Chelsea Finn, and Sergey Levine. Fast: Efficient action tokenization for vision-language-action models. *arXiv preprint arXiv:2501.09747*, 2025.
- [21] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *CVPR*, pages 2537–2546, 2019.
- [22] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, pages 2961–2969, 2017.
- [23] Zemin Liu, Yuan Li, Nan Chen, Qian Wang, Bryan Hooi, and Bingsheng He. A survey of imbalanced learning on graphs: Problems, techniques, and future directions. *IEEE Transactions on Knowledge and Data Engineering*, 2025.
- [24] Ji Zhang, Xu Luo, Lianli Gao, Difan Zou, Hengtao Shen, and Jingkuan Song. From channel bias to feature redundancy: Uncovering the “less is more” principle in few-shot learning. *arXiv e-prints*, pages arXiv–2310, 2023.
- [25] Ji Zhang, Jingkuan Song, Lianli Gao, Nicu Sebe, and Heng Tao Shen. Reliable few-shot learning under dual noises. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.
- [26] Jiang-Xin Shi, Tong Wei, Yuke Xiang, and Yu-Feng Li. How re-sampling helps for long-tail learning? *NeurIPS*, 36:75669–75687, 2023.
- [27] Hao Guo and Song Wang. Long-tailed multi-label visual recognition by collaborative training on uniform and re-balanced samplings. In *CVPR*, pages 15089–15098, 2021.
- [28] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *CVPR*, pages 9268–9277, 2019.
- [29] Shuang Li, Kaixiong Gong, Chi Harold Liu, Yulin Wang, Feng Qiao, and Xinjing Cheng. Metasaug: Meta semantic augmentation for long-tailed visual recognition. In *CVPR*, pages 5212–5221, 2021.
- [30] Xiaohua Chen, Yucan Zhou, Dayan Wu, Wanqian Zhang, Yu Zhou, Bo Li, and Weiping Wang. Imagine by reasoning: A reasoning-based implicit semantic data augmentation for long-tailed classification. In *AAAI*, volume 36, pages 356–364, 2022.
- [31] Jiahao Chen and Bing Su. Instance-specific semantic augmentation for long-tailed image classification. *IEEE Transactions on Image Processing*, 33:2544–2557, 2024.
- [32] Shan Zhang, Yao Ni, Jinhao Du, Yanxia Liu, and Piotr Koniusz. Semantic transfer from head to tail: Enlarging tail margin for long-tailed visual recognition. In *WACV*, pages 1350–1360, 2024.
- [33] Jialun Liu, Yifan Sun, Chuchu Han, Zhaopeng Dou, and Wenhui Li. Deep representation learning on long-tailed data: A learnable embedding augmentation perspective. In *CVPR*, pages 2970–2979, 2020.
- [34] Fei Du, Peng Yang, Qi Jia, Fengtao Nan, Xiaoting Chen, and Yun Yang. Global and local mixture consistency cumulative learning for long-tailed visual recognitions. In *CVPR*, pages 15814–15823, 2023.
- [35] Ziyao Meng, Xue Gu, Qiang Shen, Adriano Tavares, Sandro Pinto, and Hao Xu. H2t-fast: Head-to-tail feature augmentation by style transfer for long-tailed recognition. In *ECAI*, pages 1712–1719. IOS Press, 2023.
- [36] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *CVPR*, pages 2918–2928, 2021.
- [37] Jiaming Liu, Hao Chen, Pengju An, Zhuoyang Liu, Renrui Zhang, Chenyang Gu, Xiaoqi Li, Ziyu Guo, Sixiang Chen, Mengzhen Liu, et al. Hybridvla: Collaborative diffusion and autoregression in a unified vision-language-action model. *arXiv preprint arXiv:2503.10631*, 2025.
- [38] Qixiu Li, Yaobo Liang, Zeyu Wang, Lin Luo, Xi Chen, Mozheng Liao, Fanguyun Wei, Yu Deng, Sicheng Xu, Yizhong Zhang, et al. Cogact: A foundational vision-language-action model for synergizing cognition and action in robotic manipulation. *arXiv preprint arXiv:2411.19650*, 2024.
- [39] William J Reed. The pareto, zipf and other power laws. *Economics letters*, 74(1):15–19, 2001.
- [40] Suneel Belkhal and Dorsa Sadigh. Minivla: A better vla with a smaller footprint, 2024.
- [41] Stephen D Simon. Understanding the odds ratio and the relative risk. *Journal of andrology*, 22(4):533–536, 2001.
- [42] Chittaranjan Andrade. Understanding relative risk, odds ratio, and related terms: as simple as it can get. *The Journal of clinical psychiatry*, 76(7):21865, 2015.
- [43] Delin Qu, Haoming Song, Qizhi Chen, Yuanqi Yao, Xinyi Ye, Yan Ding, Zhigang Wang, JiaYuan Gu, Bin Zhao, Dong Wang, et al. Spatialvla: Exploring spatial representations for visual-language-action model. *arXiv preprint arXiv:2501.15830*, 2025.
- [44] Ji Zhang, Shihan Wu, Xu Luo, Hao Wu, Lianli Gao, Heng Tao Shen, and Jingkuan Song. Inspire: Vision-language-action models with intrinsic spatial reasoning. *arXiv preprint arXiv:2505.13888*, 2025.
- [45] Shihan Wu, Xu Luo, Ji Zhang, Junlin Xie, Jingkuan Song, Heng Tao Shen, and Lianli Gao. Policy contrastive decoding for robotic foundation models. *arXiv preprint arXiv:2505.13255*, 2025.
- [46] Rejin Varghese and M Sambath. Yolov8: A novel object detection algorithm with enhanced performance and robustness. In *ADICS*, pages 1–6. IEEE, 2024.