

Hierarchical-Task-Aware Multi-modal Mixture of Incremental LoRA Experts for Embodied Continual Learning

Ziqi Jia^{1,2,*}, Anmin Wang^{1,3,*}, Xiaoyang Qu^{1,†}, Xiaowen Yang², Jianzong Wang^{1,†}

¹Ping An Technology (Shenzhen) Co., Ltd., Shenzhen, China,

²Tsinghua Shenzhen International Graduate School, Tsinghua University, Shenzhen, China,

³Huazhong University of Science and Technology, Wuhan, China,

Correspondence: jzwang@188.com, quxiaoy@gmail.com

Abstract

Previous continual learning setups for embodied intelligence focused on executing low-level actions based on human commands, neglecting the ability to learn high-level planning and multi-level knowledge. To address these issues, we propose the Hierarchical Embodied Continual Learning Setups (HEC) that divide the agent’s continual learning process into two layers: high-level instructions and low-level actions, and define five embodied continual learning sub-setups. Building on these setups, we introduce the Task-aware Mixture of Incremental LoRA Experts (Task-aware MoILE) method. This approach achieves task recognition by clustering visual-text embeddings and uses both a task-level router and a token-level router to select the appropriate LoRA experts. To effectively address the issue of catastrophic forgetting, we apply Singular Value Decomposition (SVD) to the LoRA parameters obtained from prior tasks, preserving key components while orthogonally training the remaining parts. The experimental results show that our method stands out in reducing the forgetting of old tasks compared to other methods, effectively supporting agents in retaining prior knowledge while continuously learning new tasks.

1 Introduction

Recently, embodied intelligence has witnessed remarkable progress (Shi et al., 2024; Fu et al., 2024; Shridhar et al., 2020). Embodied intelligence integrates key capabilities such as perception, cognition, decision-making, and action, aiming to enable agents to perform complex tasks like household chores (Agrawal et al., 2024; Yang et al., 2024), navigation (Singh et al., 2023; Zheng et al., 2024; Zhang et al., 2024), and object manipulation (Parekh et al., 2024; Li et al., 2024b) within physical environments. Unlike the static environment

assumptions common in traditional embodied intelligence research, embodied agents are highly likely to encounter novel behavioral patterns or environments after deployment. Therefore, continual learning is crucial for their adaptation to the dynamically changing real world. Kim et al. (Kim et al., 2024) proposed a continual learning setup for embodied intelligence based on instruction following. However, with the emergence of large language models (LLMs), embodied agents demonstrate enhanced autonomous decision-making and environmental understanding (Szot et al., 2024; Mandi et al., 2024; Huang et al., 2022), capable of decomposing abstract human instructions into detailed task plans. Yet, previous continual learning setups did not fully address high-level task planning needs or leverage the full power of LLMs.

To address these challenges, we build upon existing embodied intelligence continual learning setups, specifically considering the capabilities of LLMs, further decouples the learning process into two levels: high-level instructions and low-level actions. We propose the Hierarchical Embodied Continual Learning Setups (HEC), which encompasses five specific continual learning setups: High-level Instruction Behavior Incremental Learning, High-level Instruction Environment Incremental Learning, Low-level Action Behavior Incremental Learning, Low-level Action Environment Incremental Learning, and Hybrid Hierarchical Incremental Learning. HEC aims to enable agents to concurrently learn specific actions and high-level task planning abilities, achieving a higher level of autonomy.

In embodied continual learning, when a model learns a new task, it often suffers from catastrophic forgetting (McCloskey and Cohen, 1989; Goodfellow et al., 2013; Gu and Feng, 2020; Ratcliff, 1990), where previously learned knowledge or skills are significantly lost. To address this issue, we propose a Task-aware Mixture of Incremental LoRA

*Equal Contribution

†Corresponding Author

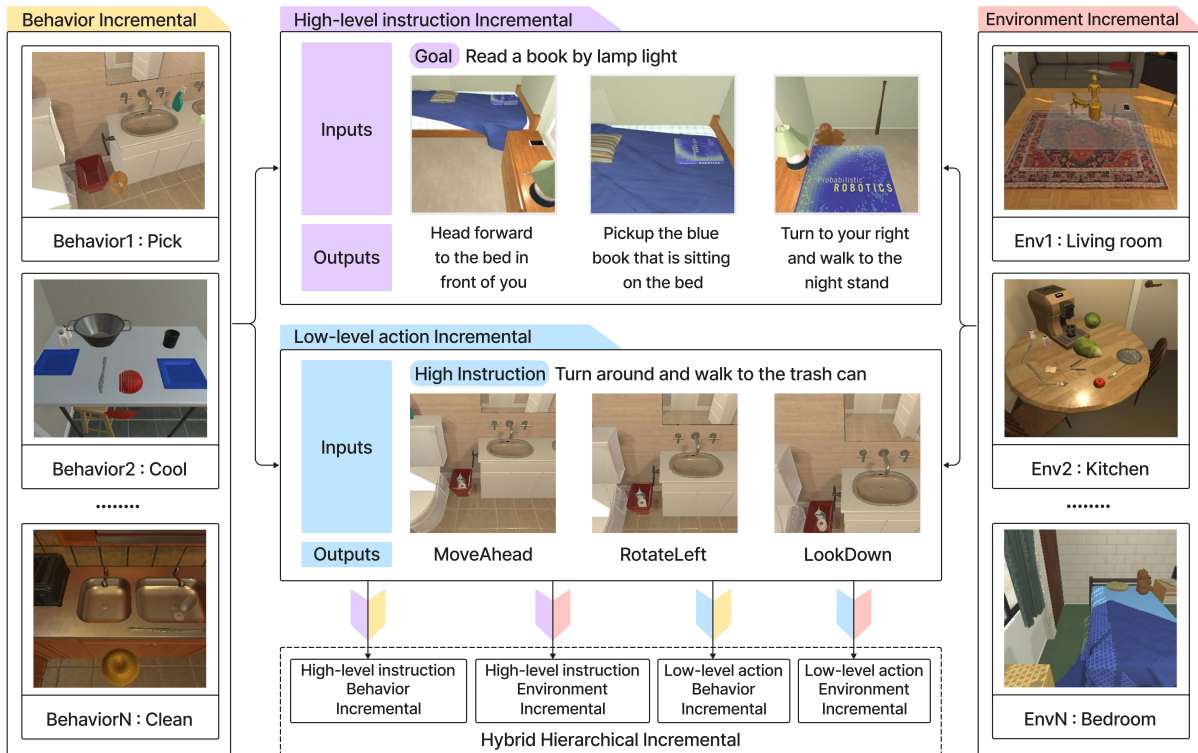


Figure 1: **Proposed Hierarchical Embodied Continual Learning Setups.** In the High-level Instruction Incremental setup, the agent learns to generate high-level instructions based on goal conditions while preserving prior knowledge. In the Low-level Action Incremental setup, the agent generates specific actions from high-level instructions without forgetting previous knowledge. Combining these with Behavior and Environment Incremental setups, we propose five incremental learning setups. Notably, Hybrid Hierarchical Incremental Learning evaluates the model’s ability to retain knowledge across different levels by treating the other four setups as independent tasks for continual learning.

Experts (Task-aware MoILE) method. Considering that models in real-world continual learning scenarios typically do not have prior knowledge of task IDs and need to make initial judgments about inputs to assign samples to LoRA experts trained on similar tasks, we first perform clustering analysis on visual-text embeddings processed by a tokenizer and the CLIP model and compute the centroid vector of the cluster to which a new input embedding belongs. This centroid vector represents the fuzzy task category of the current input and is fed into a task-level router to assist the MoE in selecting the appropriate expert. Furthermore, considering the semantic similarities between tasks in embodied continual learning (e.g., "pick" and "place") and to effectively mitigate catastrophic forgetting between different levels (e.g., high-level instructions and low-level actions), we perform singular value decomposition (SVD) on the LoRA parameters trained on previous tasks, retaining the principal components and training the remaining residual components while imposing orthogonality constraints to ensure their orthogonality with the

principal components. This approach enables continued training on LoRA experts of similar tasks under the MoE selection mechanism while effectively suppressing catastrophic forgetting and preserving previously acquired skills. Our main contributions include:

- We propose the Hierarchical Embodied Continual Learning Setups specifically designed for LLM-driven embodied agents, which decomposes learning into high-level instructions and low-level actions and defines five incremental learning setups.
- We introduce a Task-aware Mixture of Incremental LoRA Experts (Task-aware MoILE) continual learning method that effectively mitigates catastrophic forgetting and improves continual learning performance through task-aware expert selection and an SVD-based orthogonal training strategy.
- Evaluations on our HEC demonstrate that our proposed Task-aware MoILE method achieves

state-of-the-art results, attaining the highest accuracy and the lowest forgetting measure compared to previous methods.

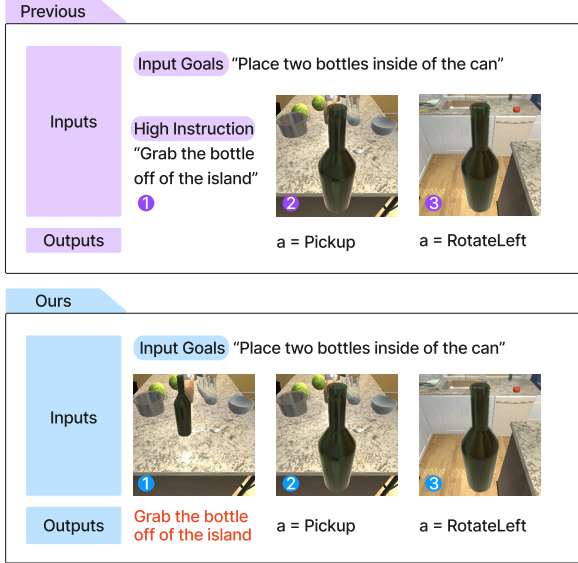


Figure 2: **Comparison of Our Setup with Previous Work.** Our setup focuses on enabling agents to learn knowledge at different levels, including instruction understanding and action execution, ultimately allowing them to autonomously generate instructions and perform actions in the environment.

2 Problem Formulation

Previous research (Kim et al., 2024) has proposed two paradigms for embodied incremental learning: Behavior Incremental Learning and Environment Incremental Learning. These approaches mainly focus on agents following and executing human instructions, while overlooking their autonomous planning and decision-making capabilities. To address this gap, and in light of the growing task planning abilities of Large Language Models and the increasing need for autonomous robotic planning, we propose the Hierarchical Embodied Continual Learning Setups. As illustrated in **Figure 1**, HEC consists of five distinct incremental learning modes: High-level Instruction Behavior Incremental Learning (HB), High-level Instruction Environment Incremental Learning (HE), Low-level Action Behavior Incremental Learning (LB), Low-level Action Environment Incremental Learning (LE), and Hybrid Hierarchical Incremental Learning (HH). As illustrated in **Figure 2**, HEC emphasizes that during actual task execution, agents not only continually learn execution skills but also continually learn task decomposition and planning.

2.1 High-level Instruction Incremental Learning

For high-level instructions, we divide this into two incremental learning modes: High-level Instruction Behavior Incremental Learning, which includes the task set $\mathcal{T}_{HB} = \{\text{EXAMINE, PICK\&PLACE, HEAT, COOL, CLEAN, PICK2\&PLACE, MOVABLE}\}$; and High-level Instruction Environment Incremental Learning, which includes the environment set $\mathcal{T}_{HE} = \{\text{KITCHENS, LIVINGROOMS, BEDROOMS, BATHROOMS}\}$. In this setup, we use high-level instructions as labels y_i^t for the i -th data point of the t -th task, with task goals \mathcal{G}_i^t and scene images \mathcal{I}_i^t as input. Based on the input $x_i^t = \{\mathcal{G}_i^t, \mathcal{I}_i^t\}$, the model outputs the high-level plan h_i^t for the corresponding task, which is a decomposed sub-task. For example, given the task goal "Read a book under the light," the model will generate different high-level instructions based on different scene images, such as "Pick up the blue book that is sitting on the bed."

2.2 Low-level Action Incremental Learning

For low-level actions, we also divide this into two incremental learning modes: Low-level Action Behavior Incremental Learning, which includes the task set $\mathcal{T}_{LB} = \{\text{EXAMINE, PICK\&PLACE, HEAT, COOL, CLEAN, PICK2\&PLACE, MOVABLE}\}$; and Low-level Action Environment Incremental Learning, which includes the task set $\mathcal{T}_{LE} = \{\text{KITCHENS, LIVINGROOMS, BEDROOMS, BATHROOMS}\}$. In this setup, we use low-level actions as labels y_i^t for the i -th data point of the t -th task, with task goals \mathcal{G}_i^t , the current high-level instruction h_i^t , and scene images \mathcal{I}_i^t as input. Based on the input $x_i^t = \{\mathcal{G}_i^t, h_i^t, \mathcal{I}_i^t\}$, the model outputs the specific action a_i^t for the current task and high-level instruction. For example, given the current high-level instruction "Go to the bookshelf," the model will output specific low-level actions based on the scene image, such as "MoveAhead" and "TurnLeft."

2.3 Hybrid Hierarchical Incremental Learning

To enable agents to continually learn both high-level instructions and low-level actions in an environment, we propose Hybrid Hierarchical Incremental Learning (HH). We consider the four incremental learning modes mentioned above (LB, HB, LE, and HE) as four independent tasks and learn

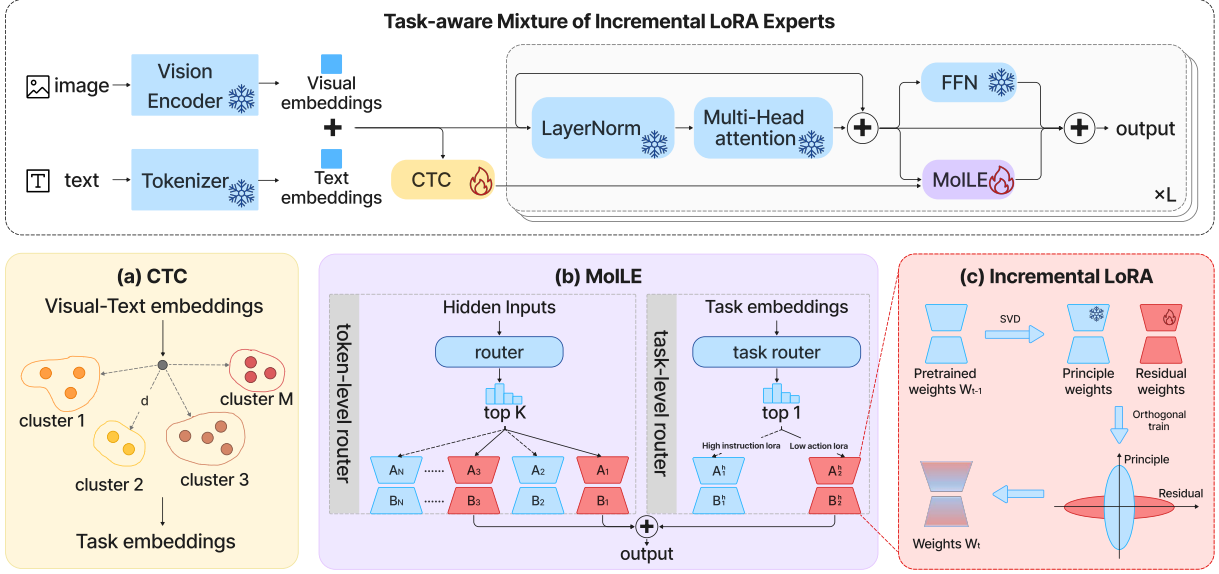


Figure 3: **The overall architecture of the proposed Task-aware MoILE.** The model transforms Visual-Text embeddings into Task embeddings through (a) CTC, and then performs token-level and task-level expert selection using (b) MoILE. The selected experts are updated via (c) Incremental LoRA, effectively preventing catastrophic forgetting.

them sequentially to evaluate the model’s forgetting performance when learning tasks at different levels.

3 Proposed Methods

To enable the agent to continually learn at both the high-instruction and low-action levels, we propose the method shown in **Figure 3**. The method first processes the input embeddings through clustering (3.1), then selects experts at the token-level and task-level using the token-level router and task-level router (3.2). We then perform SVD on the selected LoRA experts, freezing the principal components and training the residuals, ensuring the model retains old knowledge while learning new information (3.3).

3.1 Cross-modal Task Clustering

Since directly obtaining task IDs is challenging in real-world scenarios, we use cross-modal task clustering (CTC) to determine the task type of inputs. Specifically, to fully leverage the information from both modalities in the vision-language large model (VLM), we first transform the image and text into feature vectors using pre-trained visual and language encoders, respectively. These feature vectors are then fused or concatenated to generate a unified visual-text embedding vector x^m . To allow the cluster centers to dynamically adapt to changes in data distribution, the CTC module updates the

positions of the cluster centers with each new batch of samples processed. The specific update formula is as follows:

$$c_j^{new} = c_j^{old} + \frac{\alpha}{|S_j^{batch}|} \sum_{x_i^m \in S_j^{batch}} (x_i^m - c_j^{old}), \quad (1)$$

where c_j^{new} is the new cluster center, c_j^{old} is the old cluster center, α is the learning rate, S_j^{batch} is the set of samples assigned to cluster center j in the current batch, and x_i^m is the embedding vector of the i -th sample. For each input sample x_i^m , calculate its Euclidean distance from all cluster centers c_l :

$$j = \arg \min_{1 \leq l \leq M} \|x_i^m - c_l\|^2, \quad (2)$$

where x_i^m is the embedding vector of the i -th sample, c_l is the l -th cluster center, and M represents the number of cluster centers. Once it has been assigned to the nearest cluster center c_j , the CTC module outputs the corresponding cluster center vector as the task embedding e_i . This task embedding represents the task category of the sample in the multimodal space and is passed to the subsequent task router module.

3.2 Token-level Router and Task-level Router

To enable the model to better output based on task hierarchy while maintaining generalization ability in hierarchical embodied continual learning, we introduce two types of LoRA experts in

the MoE layer: token-level LoRA and task-level LoRA. Specifically, the token-level router selects the top- K token-level LoRA experts based on the hidden input x from the previous layer. Meanwhile, the task-level router selects the top-1 task-level LoRA expert based on the input task embedding e , ensuring the model chooses the appropriate high-instruction or low-action expert for the current task level. The forward propagation process can be expressed as follows:

$$f(x) = W_0x + \Delta Wx, \quad (3)$$

$$\begin{aligned} \Delta Wx &= \sum_{i=1}^{N_1} G_1(x)_i \cdot E_i(x) \\ &+ \sum_{i=1}^{N_2} G_2(e)_i \cdot E_i^h(x), \end{aligned} \quad (4)$$

where W_0 represents the pretrained weights of the base model, ΔW indicates the weight updates produced by the LoRA-enhanced experts, N_1 denotes the number of token-level LoRA instances, and N_2 represents the number of task-level LoRA instances. Therefore, the total number of LoRA experts per layer N can be expressed as $N = N_1 + N_2$. The token-level expert $E(\cdot)$ and the task-level expert $E^h(\cdot)$ compute their outputs using the following LoRA update rule:

$$\begin{cases} E_i(x) = B_i A_i x_i, \\ E_i^h(x) = B_i^h A_i^h x_i, \end{cases} \quad (5)$$

where $B_i \in \mathbb{R}^{d \times r}$, $A_i \in \mathbb{R}^{r \times k}$, $B_i^h \in \mathbb{R}^{d \times r}$, $A_i^h \in \mathbb{R}^{r \times k}$, the dimensions d and k correspond to the size of the pretrained weight matrix $W_0 \in \mathbb{R}^{d \times k}$ in large language models. The parameter r denotes the low-rank dimension, with $r \ll \min(d, k)$. The router logits $G_1(\cdot)$ denote the routing probabilities for each token-level expert, while $G_2(\cdot)$ represent the routing probabilities for each task-level expert, which can be expressed as follows:

$$\begin{cases} G_1(x) = \{i \mid p_i(x) \geq p_{(K)}(x)\}, \\ G_2(e) = \{i \mid p_i(e) \geq p_{(1)}(e)\}, \end{cases} \quad (6)$$

where $p_i(\cdot)$ denotes the probability assigned to expert i , $p_{(K)}(\cdot)$ represents the K -th highest probability, and $p_{(1)}(\cdot)$ is the highest probability.

3.3 Incremental LoRA

To effectively prevent catastrophic forgetting and promote knowledge transfer across different tasks, we propose the Incremental LoRA method. Specifically, when updating LoRA expert parameters, we

Algorithm 1 Algorithm of Task-aware MoILE

Input: Pre-trained VLM parameters, number of tasks T , task data $\{D_t\}_{t=1}^T$.

Output: Updated model parameters.

- 1: **for** each task $t \in 1 : T$ **do**
 - 2: Obtain W_p and \bar{W} (Eqs. (8),(9)).
 - 3: **for** each batch $D_{batch} \in D_t$ **do**
 - 4: Update cluster centers c_{new_j} (Eq. (1)).
 - 5: Compute task embeddings e (Eq. (2)).
 - 6: Perform MoE forward pass (Eqs. (3),(4)).
 - 7: Calculate losses L_s, L_o (Eqs. (12), (14)).
 - 8: Update parameters \bar{W} .
 - 9: **end for**
 - 10: **end for**
 - 11: **return** Updated model parameters.
-

first perform Singular Value Decomposition (SVD) on the LoRA weight matrix W , which can be expressed as follows:

$$W = U \Sigma V^\top, \quad (7)$$

where Σ is a diagonal matrix whose entries along the diagonal are the singular values, U and V are orthogonal matrices consisting of the left and right singular vectors.

We analyze the singular values after SVD and find that only a small portion of the LoRA parameters have significant values, while most have relatively low values. This indicates that a few parameters retain the core knowledge of the task. Based on this, we divide the LoRA expert parameters into principal component parameters and residual component parameters. The parameters corresponding to the top p singular values are retained as the principal components W_p , while the remaining parameters are considered residual components \bar{W} :

$$W_p = U_p \Sigma_p V_p^\top, \quad (8)$$

$$\bar{W} = W - W_p, \quad (9)$$

where we freeze the principal component parameters W_p and update only the residual parameters \bar{W} during training. After merging the two, we obtain the LoRA parameters \dot{W} :

$$\dot{W} = \bar{W} + W_p, \quad (10)$$

$$W_p = B_p A_p, \bar{W} = \bar{B} \bar{A}. \quad (11)$$

To prevent interference from new tasks from causing catastrophic forgetting, we introduce a LoRA singular value loss function L_s during training. The singular value loss function aims to minimize the difference between the top p singular

Table 1: Low-action Incremental Comparison. The best value per metric is in **bold**.

Methods	LB								LE							
	Order1		Order2		Order3		Avg		Order1		Order2		Order3		Avg	
	AA↑	FM↓	AA↑	FM↓	AA↑	FM↓	AA↑	FM↓	AA↑	FM↓	AA↑	FM↓	AA↑	FM↓	AA↑	FM↓
EWC	62.87	12.13	62.43	11.24	62.03	11.09	62.44	11.49	60.82	10.19	60.84	10.06	60.57	10.42	60.74	10.22
CAMA	62.89	11.53	62.96	11.50	62.71	12.02	62.85	11.68	61.41	10.14	60.30	10.64	60.36	10.81	60.69	10.53
MoELoRA	63.33	10.01	63.04	10.31	63.67	10.42	63.35	10.25	61.67	10.53	60.65	10.53	60.78	10.52	61.03	10.53
O-LoRA	64.61	8.85	64.53	7.70	64.68	8.62	64.61	8.39	63.06	7.91	62.20	8.05	62.75	8.14	62.67	8.03
InfLoRA	65.94	7.69	65.16	7.43	65.74	7.19	65.61	7.44	63.24	7.04	62.32	7.90	62.24	7.97	62.60	7.64
Ours	67.75	3.83	67.52	3.09	68.45	3.18	67.91	3.37	65.43	3.61	64.42	4.05	64.78	3.71	64.88	3.79

Table 2: High instruction Incremental Comparison. The best value per metric is in **bold**.

Methods	HB								HE							
	Order1		Order2		Order3		Avg		Order1		Order2		Order3		Avg	
	AA↑	FM↓	AA↑	FM↓	AA↑	FM↓	AA↑	FM↓	AA↑	FM↓	AA↑	FM↓	AA↑	FM↓	AA↑	FM↓
EWC	52.03	10.39	51.24	10.70	51.37	11.39	51.55	10.83	53.59	11.30	53.60	10.78	53.03	11.87	53.41	11.32
CAMA	51.45	10.84	51.43	10.69	51.08	11.53	51.32	11.02	53.25	10.96	53.04	10.72	53.19	10.51	53.16	10.73
MoELoRA	51.99	10.95	52.29	10.31	51.96	10.59	52.08	10.62	53.84	12.88	53.87	11.20	54.00	11.02	53.90	11.70
O-LoRA	54.57	4.96	54.22	5.69	54.28	5.47	54.36	5.37	56.48	3.68	56.41	4.09	56.82	4.89	56.57	4.22
InfLoRA	54.52	4.60	54.82	4.12	54.17	5.03	54.50	4.58	56.56	3.30	56.88	4.29	56.49	4.53	56.64	4.04
Ours	55.90	2.89	55.34	3.40	55.51	3.07	55.58	3.12	58.01	1.62	57.84	1.97	57.21	1.69	57.69	1.76

values of the LoRA parameters \hat{W} and those of the frozen principal component LoRA parameters W_p . Specifically, this loss function constrains the top p singular values of the LoRA parameters \hat{W} to be as close as possible to those of the frozen principal component LoRA parameters W_p . This ensures that the model retains knowledge from previous tasks, and the formulation is as follows:

$$L_s = \left| \sqrt{\sum_{i=1}^p \hat{\sigma}_i^2} - \sqrt{\sum_{i=1}^p \sigma_i^2} \right|, \quad (12)$$

where $\hat{\sigma}$ denotes the singular values of the LoRA parameters \hat{W} , and σ denotes the singular values of the frozen principal component LoRA parameters W_p . Additionally, we introduce an orthogonal loss L_o . Specifically, we apply an orthogonal constraint between the residual LoRA parameters \bar{W} and the principal component LoRA parameters W_p to reduce interference between them. The corresponding formula is as follows:

$$O = \bar{A}^T A_p, \quad (13)$$

$$L_o = \sum_{u,v} \|O[u, v]\|^2, \quad (14)$$

where $O[u, v]$ denotes the element at the u -th row and v -th column of the matrix O . In summary, our total loss function L_{total} is defined as follows:

$$L_{total} = L + \lambda_1 L_s + \lambda_2 L_o, \quad (15)$$

where L is the next-token prediction loss of the large language model, λ_1 and λ_2 are hyperparameters that balance the contributions of each loss term. Our pseudocode is shown in **Algorithm 1**.

4 Experiments

4.1 Experimental Setting

Metrics. We use Average Accuracy (AA) to evaluate the model’s performance on tasks:

$$AA = \frac{1}{T} \sum_{t=1}^T A_{T,t}, \quad (16)$$

where $A_{T,t}$ is the performance on t -th task after training the task T . And we employ Forgetting Measure (FM) to measure the extent of catastrophic forgetting:

$$f_{t,T} = \max_{i \in \{1, \dots, T-1\}} (A_{i,t} - A_{T,t}), \forall t < T, \quad (17)$$

$$FM_T = \frac{1}{T-1} \sum_{t=1}^{T-1} f_{t,T}. \quad (18)$$

Baseline. We compare our method with state-of-the-art techniques in the continual learning domain, including EWC (Kirkpatrick et al., 2017), CAMA (Kim et al., 2024), MoELoRA (Chen et al., 2024), O-LoRA (Wang et al., 2023), and InfLoRA (Liang and Li, 2024).

Implementation details. In our experiments, we employ the **AI2-THOR** environment alongside the ALFRED dataset (Shridhar et al., 2020). For each HEC setup, we train the agent using multiple randomly ordered sequences and evaluate it in unseen environments. For more details on these sequences, please refer to **Appendix A**. The backbone language model is LLAMA 2-7B, and the

Table 3: Hybrid Comparison. The best value per metric is in **bold**.

Methods	Order1		Order2		Order3		Avg	
	AA↑	FM↓	AA↑	FM↓	AA↑	FM↓	AA↑	FM↓
EWC	39.74	44.84	39.93	44.55	39.46	44.44	39.71	44.61
CAMA	39.59	43.51	39.88	43.90	39.76	43.55	39.74	43.65
MoELoRA	40.04	43.83	39.75	44.33	39.86	43.41	39.88	43.86
O-LoRA	43.89	37.99	44.70	38.42	43.15	37.53	43.91	37.98
InfLoRA	44.77	36.00	45.24	36.50	44.74	35.61	44.92	36.04
Ours	52.84	13.41	54.15	13.62	53.76	13.50	53.58	13.51

vision encoder is set to CLIP ViT-L/14-336 by default. For additional details on the model training parameters, please refer to **Appendix B**.

4.2 Comparison with State-of-the-art Methods

Table 1 compares our method with other state-of-the-art approaches in the Low-level Action Incremental Learning setup. The results show that, in the LB and LE setups, our method outperforms the second-best by 2.3% and 2.21% in average accuracy, while reducing the forgetting measure by 4.07% and 3.85%. **Table 2** presents the comparison in the High-level Instruction Incremental Learning setup. In the HB and HE setups, our method improves average accuracy by 1.08% and 1.05%, and reduces the forgetting measure by 1.46% and 2.28%. **Table 3** shows the comparison in the Hybrid Hierarchical Incremental Learning setup. In this more practical setup, which closely aligns with real-world embodied intelligence tasks, our method boosts average accuracy by 8.66% and reduces the forgetting measure by 22.53%.

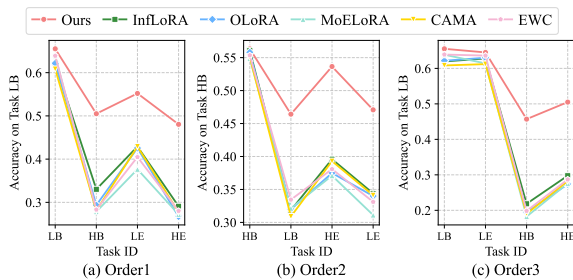


Figure 4: The performance trajectory of the first task in the continual learning process under different orders of the Hybrid Hierarchical Incremental Setup.

As shown in **Figure 4**, with continuous learning, especially when learning different levels of knowledge (e.g., low-level actions and high-level instructions), our method effectively reduces the forgetting of previously learned knowledge. In contrast, the accuracy of other methods significantly declines. Compared to O-LoRA and InfLoRA, our

method makes more accurate expert selections and better aligns with the tasks, thereby improving accuracy. When compared to MoELoRA, our method preserves existing knowledge while learning new knowledge, and, without knowing the task ID, it can choose the most suitable experts via clustering and task-level routing, thus improving accuracy and reducing forgetting.

Table 4: The Ablation Study of Different Components on the Hybrid Hierarchical Incremental Learning setup. **TR** stands for Task-level Router, **IL** represents Incremental LoRA, and **LS** and **LO** refer to Singular Value Loss and Orthogonal Loss.

Methods	AA↑	FM↓
Ours - w/oIL	45.53	38.68
Ours - w/oTR	47.87	19.45
Ours - w/oLS	49.10	21.05
Ours - w/oLO	49.31	22.72
Ours	53.58	13.51

4.3 Ablation Study




Figure 5: Routing behavior across the token-level router and task-level router for different datasets. The length of the color bar represents the router's preference for LoRA experts.

Analysis of the individual components. As shown in **Table 4**, we conducted an ablation study on the Hybrid Hierarchical Incremental Learning setups to examine the role of each component in our method. "w/oIL", "w/oTR", "w/oLS", and "w/oLO" represent versions of the method that do not use Incremental LoRA, task-level router, singular value loss, and orthogonal loss, respectively. Compared to the full method, "w/oIL" shows a

significant increase in task forgetting measure, indicating that Incremental LoRA plays a crucial role in preventing catastrophic forgetting. "w/oTR" exhibits a noticeable drop in average accuracy, highlighting the effectiveness of the task router in improving task accuracy. Both "w/oLS" and "w/oLO" show an increase in forgetting measure, suggesting L_s and L_o are important for mitigating catastrophic forgetting. These results validate the effectiveness of the individual components in our method.

Analysis of Routing Results. To further explore how our method alleviates the problem of catastrophic forgetting, we randomly selected a fixed number of samples from each dataset and visualized their LoRA block routing behaviors, as shown in **Figure 5**. This reflects the choices of LoRA experts made by the token-level router and task-level router for different tasks, where the length of the color bar represents the router’s preference for each LoRA expert. LoRA1-6 represent token-level LoRA, and LoRA7-8 represent task-level LoRA. From the visualization, we can see that the token-level router makes relatively balanced choices for token-level LoRA, although there are some preferences for different tasks. On the other hand, the task-level router effectively selects the corresponding high-level instruction or low-level action LoRA experts based on the task hierarchy. For a comparison of different hyperparameters of our method, please refer to **Appendix B**.

5 Related Work

Continual Learning. Traditional continual learning methods can be categorized into memory-based (Rebuffi et al., 2017; Yan et al., 2021), regularization-based (Wang et al., 2024), and dynamic expansion approaches (Xu et al., 2021). Recently, parameter-efficient fine-tuning (PEFT) and mixture-of-experts (MoE) approaches have outperformed traditional methods in continual learning. PEFT methods (Yan et al., 2024; Qiao and Mahdavi, 2024) introduce a few parameters to minimize modifications, preserving past knowledge with limited flexibility. MoE-based methods (Li et al., 2024a; Dou et al., 2023; Gou et al., 2023) use router mechanisms for dynamic expert selection, achieving satisfactory results but still facing catastrophic forgetting. Our approach introduces a cluster-based hierarchical MoE structure for precise expert allocation without task-id supervision, while promoting knowledge transfer. Additionally,

we propose Incremental LoRA, which retains the principal components of existing LoRA parameters and progressively learns new tasks, effectively addressing continual learning challenges.

Lifelong Learning for Embodied Agents. Embodied continual learning enables agents to continuously learn new tasks and adapt their behavior in dynamic environments (Lesort et al., 2020), while preserving previously acquired knowledge. While existing research on manipulation tasks primarily focuses on reinforcement learning (Wolczyk et al., 2021) and imitation learning (Mendez et al., 2018; Gao et al., 2021), vision-language navigation (VLN) has explored incremental learning approaches. Kim et al. (Kim et al., 2024) introduced Behavior Incremental Learning and Environment Incremental Learning for VLN, aiming to equip agents with new skills and adaptability to new environments. However, these methods mainly focus on low-level action learning, neglecting the importance of high-level planning and decision-making. To address this, we introduce Hierarchical Embodied Continual Learning Setups, proposing High-level Instruction Incremental Learning and Low-level Action Incremental Learning for VLN, enabling agents to effectively learn high-level instructions for task planning and decision-making, while retaining low-level action execution capabilities, advancing embodied intelligence continual learning.

6 Conclusion

This paper addresses the challenges of continual learning in embodied agents driven by LLMs, and proposes an innovative embodied continual learning setups and method. The main contributions include: First, we design the Hierarchical Embodied Continual Learning Setups that decomposes the learning process into high-level instructions and low-level actions, and define five incremental learning setups to enhance the agent’s ability for autonomous continual learning. Second, the proposed continual learning method, Task-aware MoILE, combines cross-modal task clustering, a hierarchical task routing mechanism, and incremental LoRA, enabling the agent to efficiently select task-specific experts without prior knowledge of task IDs, while ensuring compatibility and minimal interference between new and old tasks. Through detailed comparative experiments and ablation studies, we validate the effectiveness of each compo-

ment in preventing catastrophic forgetting and improving task execution accuracy.

7 Limitations

Despite our study making significant strides by introducing Hierarchical Embodied Continual Learning Setups and the Task-aware MoILE method to address catastrophic forgetting in cross-level learning for agents, there are several key areas that warrant further exploration and improvement in future work: Firstly, our research has only been experimentally validated on the ALFRED dataset. To more comprehensively assess the effectiveness and generalizability of our approach, future studies should aim to extend its application to more complex and varied real-world scenarios. Secondly, the current experiments rely on existing high-quality datasets for training and testing. However, in the practical application scenarios of embodied intelligence, it is often challenging to obtain large amounts or high-quality data. Therefore, an important direction for future research is to explore how to apply our method in few-shot or semi-supervised continual learning environments. This would not only broaden the applicability of the method but also enhance the adaptability and learning efficiency of agents in data-scarce conditions.

8 Acknowledgements

This work was supported by the National Key Research and Development Program of China (Youth Scientist Project) under Grant No. 2024YFB4504300 and the Shenzhen-Hong Kong Joint Funding Project (Category A) under Grant No. SGDX20240115103359001.

References

- Ayush Agrawal, Raghav Prabhakar, Anirudh Goyal, and Dianbo Liu. 2024. [Physical reasoning and object planning for household embodied agents](#). *Trans. Mach. Learn. Res.*, 2024.
- Cheng Chen, Junchen Zhu, Xu Luo, Hengtao Shen, Jingkuan Song, and Lianli Gao. 2024. Coin: A benchmark of continual instruction tuning for multi-model large language models. In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Shihan Dou, Enyu Zhou, Yan Liu, Songyang Gao, Jun Zhao, Wei Shen, Yuhao Zhou, Zhiheng Xi, Xiao Wang, Xiaoran Fan, et al. 2023. Loramoe: Revolutionizing mixture of experts for maintaining world knowledge in language model alignment. *arXiv preprint arXiv:2312.09979*, 4(7).
- Dayuan Fu, Biqing Qi, Yihuai Gao, Che Jiang, Guanting Dong, and Bowen Zhou. 2024. [Msi-agent: Incorporating multi-scale insight into embodied agents for superior planning and decision-making](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 643–659. Association for Computational Linguistics.
- Chongkai Gao, Haichuan Gao, Shangqi Guo, Tianren Zhang, and Feng Chen. 2021. [CRIL: continual robot imitation learning via generative and prediction model](#). In *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2021, Prague, Czech Republic, September 27 - Oct. 1, 2021*, pages 6747–6754. IEEE.
- Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. 2013. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*.
- Yunhao Gou, Zhili Liu, Kai Chen, Lanqing Hong, Hang Xu, Aoxue Li, Dit-Yan Yeung, James T. Kwok, and Yu Zhang. 2023. [Mixture of cluster-conditional lora experts for vision-language instruction tuning](#). *CoRR*, abs/2312.12379.
- Shuhao Gu and Yang Feng. 2020. [Investigating catastrophic forgetting during continual training for neural machine translation](#). In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 4315–4326. International Committee on Computational Linguistics.
- Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, Pierre Sermanet, Tomas Jackson, Noah Brown, Linda Luu, Sergey Levine, Karol Hausman, and Brian Ichter. 2022. [Inner monologue: Embodied reasoning through planning with language models](#). In *Conference on Robot Learning, CoRL 2022, 14-18 December 2022, Auckland, New Zealand*, volume 205 of *Proceedings of Machine Learning Research*, pages 1769–1782. PMLR.
- Byeonghwi Kim, Minhyuk Seo, and Jonghyun Choi. 2024. Online continual learning for interactive instruction following agents. *arXiv preprint arXiv:2403.07548*.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526.

- Timothée Lesort, Vincenzo Lomonaco, Andrei Stoian, Davide Maltoni, David Filliat, and Natalia Díaz Rodríguez. 2020. [Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges](#). *Inf. Fusion*, 58:52–68.
- Hongbo Li, Sen Lin, Lingjie Duan, Yingbin Liang, and Ness B Shroff. 2024a. Theory on mixture-of-experts in continual learning. *arXiv preprint arXiv:2406.16437*.
- Zhaoyi Li, Kelin Yu, Shuo Cheng, and Danfei Xu. 2024b. League++: Empowering continual robot learning through guided skill acquisition with large language models. In *ICLR 2024 Workshop on Large Language Model (LLM) Agents*.
- Yan-Shuo Liang and Wu-Jun Li. 2024. Inflora: Interference-free low-rank adaptation for continual learning. *CVPR 2024, seattle, wa, usa, june 16-22, 2024*. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 23638–23647.
- Zhao Mandi, Shreeya Jain, and Shuran Song. 2024. [Roco: Dialectic multi-robot collaboration with large language models](#). In *IEEE International Conference on Robotics and Automation, ICRA 2024, Yokohama, Japan, May 13-17, 2024*, pages 286–299. IEEE.
- Michael McCloskey and Neal J Cohen. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier.
- Jorge A. Mendez, Shashank Shivkumar, and Eric Eaton. 2018. [Lifelong inverse reinforcement learning](#). In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 4507–4518.
- Amit Parekh, Nikolas Vitsakis, Alessandro Suglia, and Ioannis Konstantas. 2024. [Investigating the role of instruction variety and task difficulty in robotic manipulation tasks](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pages 19389–19424. Association for Computational Linguistics.
- Fuli Qiao and Mehrdad Mahdavi. 2024. [Learn more, but bother less: parameter efficient continual learning](#). In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Roger Ratcliff. 1990. Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychological review*, 97(2):285.
- Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. 2017. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010.
- Haochen Shi, Zhiyuan Sun, Xingdi Yuan, Marc-Alexandre Côté, and Bang Liu. 2024. [Opex: A component-wise analysis of llm-centric agents in embodied instruction following](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 622–636. Association for Computational Linguistics.
- Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. 2020. Alfred: A benchmark for interpreting grounded instructions for everyday tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10740–10749.
- Kunal Pratap Singh, Jordi Salvador, Luca Weihs, and Aniruddha Kembhavi. 2023. Scene graph contrastive learning for embodied navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 10884–10894.
- Andrew Szot, Max Schwarzer, Harsh Agrawal, Bogdan Mazouze, Rin Metcalf, Walter Talbott, Natalie Mackraz, R. Devon Hjelm, and Alexander T. Toshev. 2024. [Large language models as generalizable policies for embodied tasks](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Xiao Wang, Tianze Chen, Qiming Ge, Han Xia, Rong Bao, Rui Zheng, Qi Zhang, Tao Gui, and Xuanjing Huang. 2023. [Orthogonal subspace learning for language model continual learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 10658–10671. Association for Computational Linguistics.
- Yifan Wang, Yafei Liu, Chufan Shi, Haoling Li, Chen Chen, Haonan Lu, and Yujiu Yang. 2024. [Insl: A data-efficient continual learning paradigm for fine-tuning large language models with instructions](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), NAACL 2024, Mexico City, Mexico, June 16-21, 2024*, pages 663–677. Association for Computational Linguistics.
- Maciej Wolczyk, Michal Zajac, Razvan Pascanu, Lukasz Kucinski, and Piotr Milos. 2021. [Continual world: A robotic benchmark for continual reinforcement learning](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 28496–28510.

Ju Xu, Jin Ma, Xuesong Gao, and Zhanxing Zhu. 2021. Adaptive progressive continual learning. *IEEE transactions on pattern analysis and machine intelligence*, 44(10):6715–6728.

Shipeng Yan, Jiangwei Xie, and Xuming He. 2021. Der: Dynamically expandable representation for class incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, CVPR2021, virtual, June 19-25, 2021*, pages 3014–3023.

Zhiyuan Yan, Jiangming Wang, Zhendong Wang, Peng Jin, Ke-Yue Zhang, Shen Chen, Taiping Yao, Shouhong Ding, Baoyuan Wu, and Li Yuan. 2024. Effort: Efficient orthogonal modeling for generalizable ai-generated image detection. *arXiv preprint arXiv:2411.15633*.

Jingkang Yang, Yuhao Dong, Shuai Liu, Bo Li, Ziyue Wang, Haoran Tan, Chencheng Jiang, Jiamu Kang, Yuanhan Zhang, Kaiyang Zhou, et al. 2024. Octopus: Embodied vision-language programmer from environmental feedback. In *European Conference on Computer Vision, ECCV 2024, Milan, Italy, September 29-October 4, 2024*, pages 20–38. Springer.

Hongxin Zhang, Weihua Du, Jiaming Shan, Qinhong Zhou, Yilun Du, Joshua B. Tenenbaum, Tianmin Shu, and Chuang Gan. 2024. Building cooperative embodied agents modularly with large language models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.

Duo Zheng, Shijia Huang, Lin Zhao, Yiwu Zhong, and Liwei Wang. 2024. Towards learning a generalist model for embodied navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 13624–13634.

A Details of Hierarchical Embodied Continual Learning Setups

A.1 High-level Instruction Behavior Incremental Learning

In our setups, there are seven distinct behavior categories: EXAMINE, PICK&PLACE, HEAT, COOL, CLEAN, PICK2&PLACE, and MOVABLE. In the high-level Instruction Behavior Incremental Learning setup, the model is required to learn high-level planning under different behavior types. The total dataset for this setup comprises 187,103 instances. We randomly ordered these behavior categories to generate the following three sequences for model training.

1. EXAMINE → HEAT → PICK2&PLACE → COOL → PICK&PLACE → CLEAN → MOVABLE

2. PICK2&PLACE → CLEAN → MOVABLE → PICK&PLACE → HEAT → EXAMINE → COOL

3. MOVABLE → COOL → PICK&PLACE → HEAT → CLEAN → EXAMINE → PICK2&PLACE

A.2 High-level Instruction Environment Incremental Learning

In the dataset, four unique environments are included: KITCHENS, LIVINGROOMS, BEDROOMS, and BATHROOMS. In the high-level Instruction Environment Incremental Learning setup, the model is tasked with acquiring the ability to perform high-level planning across varying environmental contexts. The total dataset for for this setup comprises 81048 instances. We also randomly ordered these environments to generate the following three sequences for model training.

1. BEDROOMS → BATHROOMS → LIVINGROOMS → KITCHENS

2. KITCHENS → BEDROOMS → LIVINGROOMS → BATHROOMS

3. BATHROOMS → LIVINGROOMS → KITCHENS → BEDROOMS

A.3 Low-level Action Behavior Incremental Learning

In the Low-level Action Behavior Incremental Learning setup, the model is required to acquire low-level control capabilities across diverse behavior types. We trained the model using the same sequences as described above in A.1.

1. EXAMINE → HEAT → PICK2&PLACE → COOL → PICK&PLACE → CLEAN → MOVABLE

2. PICK2&PLACE → CLEAN → MOVABLE → PICK&PLACE → HEAT → EXAMINE → COOL

3. MOVABLE → COOL → PICK&PLACE → HEAT → CLEAN → EXAMINE → PICK2&PLACE

A.4 Low-level Action Environment Incremental Learning

In the Low-level Action Environment Incremental Learning setup, the model is required to acquire

low-level control capabilities across diverse environmental contexts. We trained the model using the same sequences as described above in A.2.

1. BEDROOMS \rightarrow BATHROOMS \rightarrow LIVINGROOMS \rightarrow KITCHENS
2. KITCHENS \rightarrow BEDROOMS \rightarrow LIVINGROOMS \rightarrow BATHROOMS
3. BATHROOMS \rightarrow LIVINGROOMS \rightarrow KITCHENS \rightarrow BEDROOMS

A.5 Hybrid Hierarchical Incremental Learning

In our experimental setup, we also designed three distinct sequences comprising four incremental learning modes, resulting in the following orders for model training.

1. Low-level Action Behavior \rightarrow High-level Instruction Behavior \rightarrow Low-level Action Environment \rightarrow High-level Instruction Environment
2. High-level Instruction Behavior \rightarrow Low-level Action Behavior \rightarrow High-level Instruction Environment \rightarrow Low-level Action Environment
3. Low-level Action Behavior \rightarrow Low-level Action Environment \rightarrow High-level Instruction Behavior \rightarrow High-level Instruction Environment

B Details of Model Training

In this section, we provide the detailed hyperparameter settings used for training the proposed model. The hyperparameters are set as follows. We conducted the training using two NVIDIA RTX 4090 D GPUs. A complete training cycle of the Hybrid Hierarchical Incremental Learning setup typically requires approximately 8 hours.

- In Equation 2, M is set to 4, where M denotes the number of clusters.
- In Equation 4, N_1 is set to 6, and N_2 is set to 2, where N_1 represents the number of token-level LoRA experts, and N_2 represents the number of task-level LoRA experts.
- In Equation 6, K is set to 2, where K denotes the number of selected token-level LoRA experts in the top- K gating mechanism of the MoE layer.

- In Equation 15, λ_1 is set to 1.0, and λ_2 is set to 0.5, where λ_1 is the weight for the LoRA singular value loss, and λ_2 is the weight for the orthogonal loss.

Additionally, we conducted experiments with different hyperparameter settings to evaluate the model’s performance.

Table 5: Model performance (AA and FM) under different hyperparameter settings.

N_1	N_2	K	λ_1	λ_2	M	AA \uparrow	FM \downarrow
2	2	1	1.0	0.5	4	48.10	18.56
6	2	2	0.5	1.0	4	53.10	14.27
6	2	2	1.0	0.5	8	52.30	15.84
6	2	2	1.0	0.5	4	53.58	13.51

The results summarized in Table 5 indicate that the current hyperparameter setting achieves optimal performance. Reducing token-level experts to $N_1 = 2$ with $K = 1$ degrades AA by 5.48%, suggesting insufficient capacity for token-level adaptation. Setting the LoRA singular value loss weight ($\lambda_1 = 0.5$) smaller than the orthogonal loss weight ($\lambda_2 = 1.0$) results in slightly degraded performance. This suggests that a stronger singular value constraint is more effective in balancing parameter diversity and adaptation stability. Setting cluster number M to 8 increases FM by 2.33% and degrades AA by 1.28%, implying over-clustering introduces noisy task-feature disentanglement. These results demonstrate that our final hyperparameter setting balances these factors.

Table 6: Comprehensive analysis of computational overheads.

Method	Memory (GB)	Training Time (h)
EWC	15.37	5.2
CAMA	15.38	5.8
MoELoRA	16.32	6.8
O-LoRA	19.26	7.9
InfLoRA	19.58	7.2
Ours	16.06	7.6

Table 7: Performance comparison under different numbers of clusters.

Number of Clusters	AA \uparrow	FM \downarrow
2	53.27	14.03
4	53.58	13.51
8	52.30	15.84

As demonstrated in Table 6, our proposed method achieves the highest task accuracy and the

lowest knowledge forgetting (see Tables 1, 2, and 3 in the paper for detailed comparisons) while maintaining memory efficiency and inference efficiency comparable to the baseline method MoELoRA. Notably, this performance advantage is accompanied by a moderate increase in training time.

In designing the cross-modal clustering module, our primary objective was to enable the model to effectively distinguish between high-level instructions and low-level actions. To this end, we employed a fixed number of clusters. **Table 7** illustrates the impact of varying cluster numbers on model performance under consistent experimental conditions. As the results show, the number of clusters exhibits minimal influence on overall performance. Based on these findings, we adopted a fixed number of clusters in practical applications. Specifically, setting the number of clusters to 2 is sufficient to differentiate between high-level instructions and low-level actions, ensuring both simplicity and effectiveness in real-world scenarios.

C Analysis of LoRA Parameters and Visual-Text Embeddings




Figure 6: The singular values of the LoRA parameters after SVD.

As shown in **Figure 6**, we analyze the singular values after SVD and find that only a small portion of the LoRA parameters have significant values, while most have relatively low values. This indicates that a few parameters retain the core knowledge of the task. Based on this, we divide the LoRA expert parameters into principal component parameters and residual component parameters to prioritize critical information during model updates.

As shown in **Figure 7**, we use t-SNE to visual-

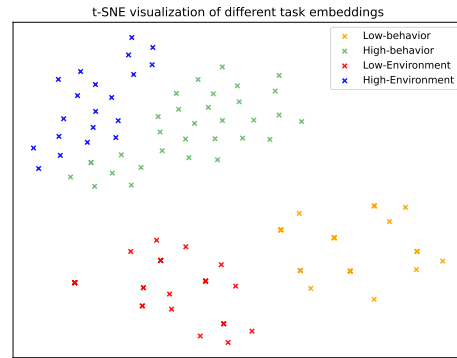


Figure 7: **T-SNE visualization of visual-text embeddings.** Different colors indicating different hierarchical continual learning datasets.

ize the visual-text embeddings of different datasets. It is evident that tasks of the same type are clustered together in the feature space, while tasks with similar hierarchical levels, such as High-behavior and High-Environment under High-level Instruction, are more closely distributed. This geometric arrangement provides insights for hierarchical knowledge preservation.