GRANULARITY MATTERS: TASK-LEVEL PRESERVATION AND DIRECTION-LEVEL TRANSFER FOR CONTINUAL MULTIMODAL INSTRUCTION TUNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Continual multimodal instruction tuning (CMIT) requires multimodal large language models to acquire new knowledge without forgetting previous tasks. Task-level preservation addresses this challenge by protecting the expert learned for each task, while leaving historical knowledge confined within individual experts. Historical expert transfer can overcome this confinement, but assigning a single strength to the complete expert suppresses useful directions when the value is low and amplifies irrelevant or conflicting directions when it is high. Effective transfer therefore requires direction-level transfer without weakening task-level preservation. We propose PHAROS, a novel framework that decouples task-level preservation from direction-level transfer. PHAROS preserves each task with a private LoRA expert and its associated multimodal projector. To enable adaptive and selective transfer, PHAROS derives functional directions from the layer-wise updates of each historical private expert. For each new task, PHAROS learns an independent signed gate for every historical direction, allowing the sign and strength of the same direction to adapt across target tasks. The gated directions form a task-dependent historical transfer branch alongside the new private expert, enabling selective historical transfer while preserving its dedicated trainable capacity. Extensive experiments on CoIN and UCIT show that PHAROS achieves average accuracies of 70.11% and 77.60%, outperforming the best prior results by 4.77% and 3.58%, respectively. Further experiments reveal that task-level preservation and direction-level transfer provide complementary gains, while the learned gates assign different signs and strengths to functional directions from the same historical expert across target tasks. These results demonstrate the value of preserving task knowledge as complete experts while transferring historical knowledge through finer functional directions.

1 INTRODUCTION

Continual multimodal instruction tuning (CMIT) enables multimodal large language models (Alayrac et al., 2022; Li et al., 2023; Dai et al., 2023; Liu et al., 2023) to adapt sequentially to heterogeneous multimodal tasks. Sequential optimization is vulnerable to interference because updates for a new task can overwrite parameters that support earlier instruction-following and task-specific behavior, leading to catastrophic forgetting (Li & Hoiem, 2017; Kirkpatrick et al., 2017; Chen et al., 2024). Therefore, an effective CMIT method must learn each new task without overwriting previously acquired knowledge. Existing methods mitigate such interference in several ways. Orthogonality-based approaches constrain new updates to avoid optimization directions used by previous tasks (Wang et al., 2023; Qiao et al., 2024a). Parameter-protection approaches preserve or repair weights that are important to earlier tasks (Zhu et al., 2024; Chen et al., 2025). Some adapter-based approaches isolate task-specific updates in dedicated experts and freeze each expert after training (He et al., 2026a; Zhang et al., 2025; Guo et al., 2025). This preservation therefore keeps the task-specific updates stored in historical experts unchanged during subsequent training. Under strict isolation, each historical expert becomes a knowledge silo: its parameters remain protected, but its knowledge cannot support the adaptation of later tasks.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

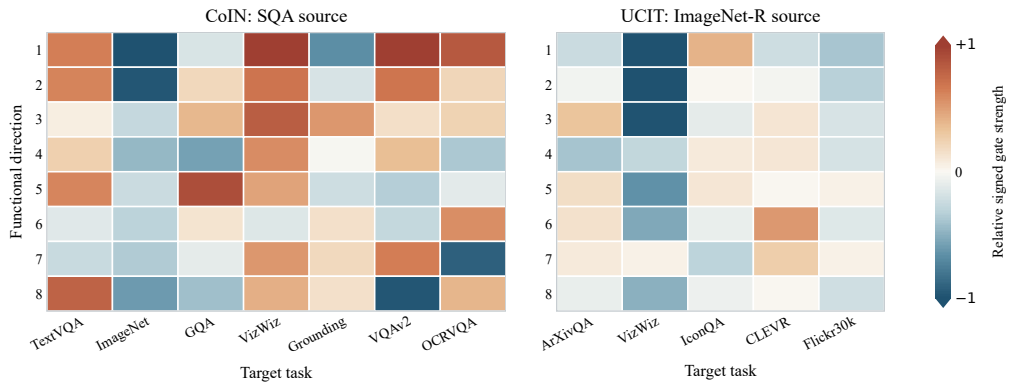


Figure 1: **Historical direction transfer on CoIN and UCIT.** The mixed signs and strengths show that functional directions from the same historical expert contribute differently across target tasks, supporting adaptive direction-level transfer rather than uniform expert-level transfer.

CMIT therefore also requires historical knowledge transfer, allowing later tasks to draw on earlier learning by incorporating historical experts into new-task adaptation (He et al., 2026a; Chen et al., 2024; Guo et al., 2025). However, treating a complete historical expert as a single transfer unit couples all functional directions within that expert. Because the relevance of these directions can vary independently for a new task, assigning a single mixing weight to the complete expert makes it difficult to strengthen useful directions while suppressing irrelevant or conflicting ones. As shown in Figure 1, functional directions derived from the same historical expert receive different signs and strengths across target tasks, a variation that a single expert-level coefficient cannot represent. The key issue is therefore not only which historical expert should provide knowledge, but also the granularity at which that knowledge should be transferred.

Motivated by this observation, we propose PHAROS, which decouples task-level preservation from direction-level transfer. For task-level preservation, PHAROS freezes the private LoRA expert learned for each task together with its associated multimodal projector, because subsequent changes to a shared projector would alter the visual representation received by the frozen expert (Liu et al., 2023; Guo et al., 2025). For direction-level transfer, PHAROS derives functional directions from each historical private expert through its layer-wise LoRA updates and stores them in a historical bank. Each functional direction is an input-to-output transformation derived from the historical expert rather than an arbitrary coordinate of its LoRA factorization. For each new task, PHAROS jointly learns the current private expert and a separate signed gate for every historical direction. Because these gates are optimized on the current-task data, the same historical direction can receive different signs and strengths across target tasks. The gated directions form a task-specific historical transfer branch, allowing historical knowledge to support new-task adaptation while the current private expert provides dedicated trainable capacity for the task.

We conduct extensive experiments on two complementary CMIT benchmarks, CoIN and UCIT, comprising fourteen heterogeneous tasks across visual question answering, classification, reasoning, grounding, OCR, and captioning. PHAROS achieves an Avg.ACC of 70.11% on CoIN and 77.60% on UCIT, surpassing the best previous results by 4.77% and 3.58%, respectively. Further experiments reveal that private-expert preservation and direction-level transfer provide complementary gains, while the learned gates assign different signs and strengths to functional directions from the same historical expert across target tasks.

Our contributions are summarized as follows:

- We identify that task knowledge in CMIT should be preserved in complete experts, while historical knowledge should be transferred through finer functional directions.
- We propose PHAROS, which preserves each task's private expert and associated projector for task-level preservation, while using functional directions controlled by adaptive signed gates for selective historical transfer.
- Extensive experiments on established benchmarks show that PHAROS achieves superior accuracy, surpassing existing methods and demonstrating its strong performance in CMIT.

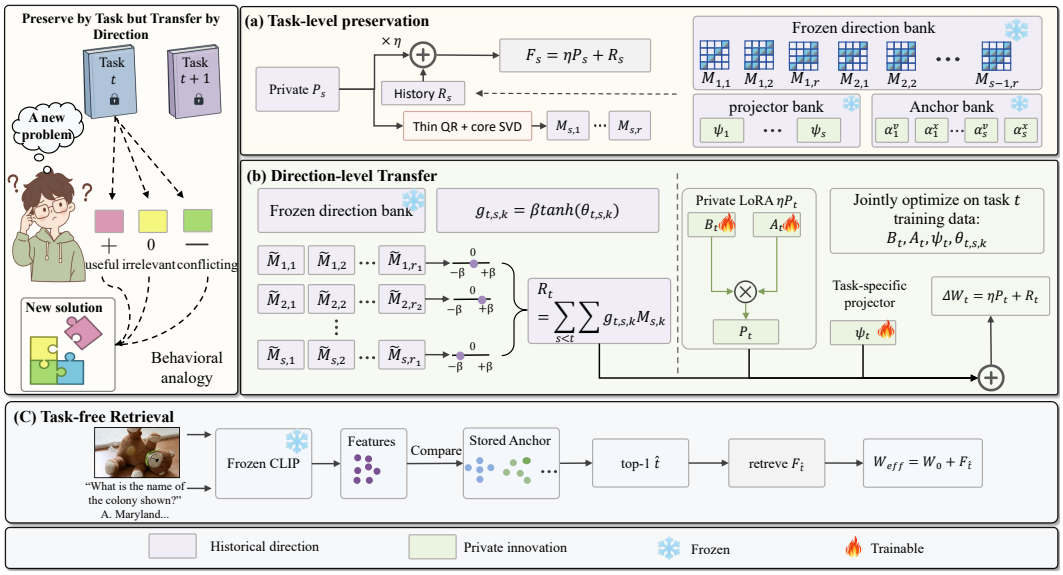


Figure 2: **Overview of PHAROS.** (a) After learning a task, PHAROS preserves its private expert and derives functional directions as fine-grained units for future transfer. (b) For a new task, task-specific signed gates adaptively reinforce, suppress, or reverse individual historical directions according to their effects on the current task. This direction-level control enables selective historical transfer while the new private expert provides dedicated trainable capacity for the current task. (c) At inference time, when task IDs are unavailable, multimodal anchor-based routing selects a task and retrieves its projector, private expert, and signed-gate policy as a unified task state.

2 Related Work

2.1 Knowledge Preservation in Continual Multimodal Learning

CMIT requires multimodal large language models to acquire new knowledge without forgetting previous tasks (Parisi et al., 2019; De Lange et al., 2021; He et al., 2026a; Chen et al., 2024; Guo et al., 2025). Because tasks arrive sequentially, optimization on a new task can overwrite parameters that support knowledge acquired from previous tasks. Knowledge distillation (Li & Hoiem, 2017), parameter regularization (Kirkpatrick et al., 2017), constrained updates (Wang et al., 2023; Qiao et al., 2024a), and selective parameter protection or repair (Zhu et al., 2024; Qiao et al., 2024b; Chen et al., 2025; Li et al., 2026; Liu et al., 2026) mitigate this interference while continuing to optimize a shared parameter space. These approaches reduce forgetting, but stronger restrictions protect previous tasks at the expense of acquiring new knowledge, whereas weaker restrictions improve new-task adaptation while leaving historical knowledge more vulnerable to overwriting. Parameter-isolation methods instead reserve or mask parameters for previous tasks (Mallya & Lazebnik, 2018; Serra et al., 2018; Mallya et al., 2018), providing stronger protection against subsequent updates. However, progressively partitioning a fixed parameter space leaves less trainable capacity available for later tasks as the task sequence grows. Rather than constraining or partitioning a shared parameter space, PHAROS assigns each task a private LoRA expert that remains fixed during subsequent training. This design directly protects task knowledge while providing the current private expert with dedicated trainable capacity for new-task adaptation. However, when a later task requires knowledge already encoded by a historical expert, this isolation forces the model to relearn that knowledge from current-task data.

2.2 Knowledge Transfer across Task-Specific Modules

Historical expert transfer allows later tasks to build on knowledge already encoded by historical experts rather than relearn it from current-task data. Modular continual-learning methods retain task-specific adapters or experts and retrieve or expand them as tasks arrive (Aljundi et al., 2017; Houlsby et al., 2019; Hu et al., 2022; He et al., 2026a; Chen et al., 2024; Guo et al., 2025; Huai et al., 2025; Yu et al., 2025; Zhang et al., 2025; Che et al., 2026; Hou et al., 2026). Beyond selecting

individual experts, AdapterFusion combines the outputs of frozen adapters (Pfeiffer et al., 2021), LoRAHub constructs weighted mixtures of complete task-specific LoRA modules (Huang et al., 2023), and task arithmetic combines complete task updates (Ilharco et al., 2023; Yang et al., 2024). TIES operates at a finer parameter-coordinate level by trimming parameter changes and resolving sign conflicts before merging (Yadav et al., 2023). These approaches make previously acquired knowledge accessible to later tasks, but their transfer granularity still limits selective control: Module-level combinations couple directions within each source module, whereas coordinate-level rules operate on parameters that do not represent coherent functional directions, preventing target-task-specific control over their contributions. PHAROS instead preserves complete experts at the task level and uses task-specific signed gates to control the sign and strength of each historical direction independently. This direction-level transfer allows useful historical knowledge to support the current task while limiting irrelevant or conflicting contributions, resulting in more effective new-task adaptation and stronger continual performance without weakening task-level preservation.

2.3 Fine-Grained Structures for Continual Adaptation

Between complete adapters and individual parameter coordinates, existing methods have explored finer structures within low-rank updates. AdaLoRA allocates rank budgets according to parameter importance (Zhang et al., 2023), while SMoRA and related vision-language methods organize rank components as independently activated experts (Zhao et al., 2026; Fa et al., 2026). In continual learning, historical updates have been organized through orthogonal or residual subspaces (Wang et al., 2023; Luo et al., 2026), separated into shared and task-specific subspaces (He et al., 2026b; Kaushik et al., 2026), compressed progressively (Qiao et al., 2026), or represented using shared bases with task-specific coefficients (Chen et al., 2026). EBLoRA further addresses forgetting by regulating spectral imbalance in low-rank updates (Gu et al., 2026). These methods demonstrate that fine-grained structures can improve parameter efficiency, knowledge sharing, or interference control. However, these structures primarily determine how the current or accumulated adaptation is learned, constrained, allocated, or stored. PHAROS assigns fine-grained structure a different role: each completed private expert remains unchanged, while functional directions derived from that expert form a separate representation for subsequent transfer. This separation enables fine-grained historical transfer without altering the task adaptation preserved in the complete expert.

3 Method

3.1 Preliminaries

We consider CMIT over an ordered sequence of T tasks, $\mathcal{T}_{1:T} = (\mathcal{T}_1, \ldots, \mathcal{T}_T)$. At training stage t, only the current-task dataset $\mathcal{D}_t = \{(v_i, x_i, y_i)\}_{i=1}^{n_t}$ is available for optimization, where n_t is the number of training examples, and v_i, x_i, and y_i are the image, instruction, and target response of the i-th example, respectively. Because training proceeds one task at a time, the current task is known during training, whereas the model must infer the task identity from each input at inference time. Under this protocol, the model must acquire new knowledge from each incoming task without forgetting previous tasks. To address this objective, PHAROS keeps the pretrained backbone frozen and represents the knowledge acquired from each task using a private LoRA expert. Specifically, let $\mathcal{L}$ denote the set of linear layers augmented with LoRA, and let $W^\ell \in \mathbb{R}^{d_{\text{out}}^\ell \times d_{\text{in}}^\ell}$ denote the frozen weight matrix of layer $\ell \in \mathcal{L}$, where d_{in}^ℓ and d_{out}^ℓ are its input and output dimensions, respectively. For task $\mathcal{T}_t$, PHAROS parameterizes the private update using LoRA (Hu et al., 2022) as

$$\Delta W_{t,\text{priv}}^\ell = \eta_\ell B_t^\ell A_t^\ell, \qquad A_t^\ell \in \mathbb{R}^{r \times d_{\text{in}}^\ell}, \quad B_t^\ell \in \mathbb{R}^{d_{\text{out}}^\ell \times r}, \tag{1}$$

where A_t^ℓ and B_t^ℓ are task-specific trainable matrices, r is the prescribed LoRA rank, and η_ℓ is the fixed scaling factor for layer ℓ. This task-specific parameterization provides a distinct private parameter space for each task, forming the basis for task-level knowledge preservation.

3.2 Task-Level Knowledge Preservation

PHAROS preserves the knowledge learned from each task by freezing its private expert after training on that task is complete. For task $\mathcal{T}_s$, let

$$\mathcal{E}_s = \left\{ \left(A_s^\ell, B_s^\ell\right) \right\}_{\ell \in \mathcal{L}} \tag{2}$$

denote its private expert. After training on $\mathcal{T}_s$ is complete, PHAROS freezes $\mathcal{E}_s$ to preserve the knowledge learned from this task. At stage t, all historical private experts $\{\mathcal{E}_s \mid 1 \leq s \leq t-1\}$ remain frozen, whereas the current private expert $\mathcal{E}_t$ is trainable. Freezing the historical private experts prevents optimization on $\mathcal{D}_t$ from overwriting the knowledge they encode. Meanwhile, task $\mathcal{T}_t$ receives a dedicated rank-r LoRA parameterization for acquiring new knowledge. Because preserving the private expert alone cannot retain its task-specific visual mapping, PHAROS pairs $\mathcal{E}_s$ with a task-specific multimodal projector Ψ_s, which maps visual features into the language-model embedding space. After training on $\mathcal{T}_s$, Ψ_s is frozen with $\mathcal{E}_s$, preventing subsequent stages from changing the visual representation supplied to the preserved expert. The frozen private experts then provide stable sources of historical knowledge for transfer to subsequent tasks.

3.3 Direction-Level Historical Transfer

A direct form of historical expert transfer treats each historical private expert as a single transfer unit. Under this strategy, a single strength controls all functional directions encoded by the expert, even though their relevance may vary for the current task. Limiting this strength to suppress irrelevant or conflicting directions consequently weakens the transfer of useful directions. PHAROS resolves this coupling by decomposing each historical private expert into functional directions whose contributions can be controlled independently.

Functional-direction construction. The functional directions must describe the transformation encoded by an expert rather than the arbitrary coordinates of its LoRA factors. For a historical task $\mathcal{T}_s$ and layer $\ell \in \mathcal{L}$, let

$$P_s^\ell = B_s^\ell A_s^\ell, \tag{3}$$

denote the unscaled layer-wise operator of its private expert. This operator remains unchanged under the equivalent reparameterization

$$(B_s^\ell G)(G^{-1}A_s^\ell) = B_s^\ell A_s^\ell = P_s^\ell, \tag{4}$$

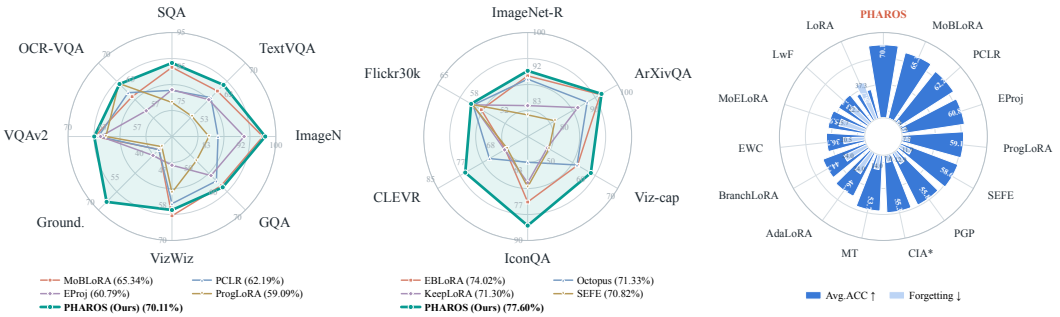
where $G \in \mathbb{R}^{r\times r}$ is any invertible matrix and G^{-1} is its inverse. Because this transformation changes the factor coordinates while preserving the encoded function, attaching transfer strengths directly to those coordinates would make transfer depend on an arbitrary parameterization. PHAROS instead derives functional directions from the compact singular value decomposition (SVD) of P_s^ℓ. Let $\rho_s^\ell = \mathrm{rank}(P_s^\ell) \leq r$ denote its number of nonzero singular values. Then

$$P_s^\ell = U_s^\ell \Sigma_s^\ell (V_s^\ell)^\top = \sum_{k=1}^{\rho_s^\ell} M_{s,k}^\ell, \qquad M_{s,k}^\ell = \sigma_{s,k}^\ell u_{s,k}^\ell (v_{s,k}^\ell)^\top, \tag{5}$$

where U_s^ℓ and V_s^ℓ contain the left and right singular vectors, $\Sigma_s^\ell = \mathrm{diag}(\sigma_{s,1}^\ell, \ldots, \sigma_{s,\rho_s^\ell}^\ell)$ contains the positive singular values, and $u_{s,k}^\ell$ and $v_{s,k}^\ell$ are the k-th columns of U_s^ℓ and V_s^ℓ, respectively. Thus, $M_{s,k}^\ell \in \mathbb{R}^{d_{\mathrm{out}}^\ell \times d_{\mathrm{in}}^\ell}$ is the k-th input-to-output transformation, and their sum reconstructs the complete layer-wise operator. This operator-based construction is invariant to equivalent LoRA reparameterizations, up to the usual basis choice for repeated singular values. In practice, PHAROS obtains the same compact decomposition from a core matrix of size at most $r \times r$ using thin QR decompositions of A_s^ℓ and B_s^ℓ, avoiding materialization of the dense operator. Although this construction removes dependence on arbitrary factor coordinates, the functional directions can have different source scales. Without normalization, a learned strength would reflect both current-task relevance and the scale inherited from the historical expert. PHAROS separates these factors by defining

$$\widetilde{M}_{s,k}^\ell = \frac{\eta_\ell M_{s,k}^\ell}{\|\eta_\ell M_{s,k}^\ell\|_F}, \qquad \|\widetilde{M}_{s,k}^\ell\|_F = 1, \tag{6}$$

where $\|\cdot\|_F$ denotes the Frobenius norm and $\widetilde{M}_{s,k}^\ell$ is the normalized functional direction. After task $\mathcal{T}_s$ is completed, PHAROS adds the normalized directions derived from its private expert to the historical direction bank, so each newly acquired direction enters the bank once.



(a) CoIN benchmark. (b) UCIT benchmark. (c) Radial method comparison.

Figure 3: **Task-wise and aggregate performance on CoIN and UCIT.** PHAROS delivers broad gains across heterogeneous tasks together with near-zero forgetting, demonstrating that strong continual performance and historical-knowledge preservation are achieved simultaneously.

Adaptive signed direction transfer. The normalized direction bank allows PHAROS to learn comparable contributions from different historical experts. For a new task $\mathcal{T}_t$, PHAROS assigns a task-specific signed gate to each direction from every historical task $\mathcal{T}_s$, where $s < t$:

$$g^{\ell}_{t,s,k} = \beta \tanh\left(\theta^{\ell}_{t,s,k}\right), \qquad g^{\ell}_{t,s,k} \in [-\beta, \beta], \tag{7}$$

where $\theta^{\ell}_{t,s,k}$ is the trainable gate logit and $\beta > 0$ is the maximum absolute transfer strength. A positive gate transfers $\widetilde{M}^{\ell}_{s,k}$ in its original orientation, a near-zero gate suppresses its contribution, and a negative gate reverses its effect. Because a separate gate is learned for every target task and historical direction, the same direction can receive different signs and strengths across subsequent tasks. The gates therefore adapt historical transfer to the current task at the direction level rather than assigning one strength to the complete expert. Summing the gated directions produces the historical transfer branch

$$\Delta W^{\ell}_{t,\text{hist}} = \sum_{s=1}^{t-1} \sum_{k=1}^{\rho^{\ell}_s} g^{\ell}_{t,s,k} \widetilde{M}^{\ell}_{s,k}, \tag{8}$$

where $\Delta W^{\ell}_{t,\text{hist}}$ is the historical update transferred to task $\mathcal{T}_t$ at layer ℓ. Combining this branch with the current private expert gives

$$\Delta W^{\ell}_{t} = \Delta W^{\ell}_{t,\text{priv}} + \Delta W^{\ell}_{t,\text{hist}}. \tag{9}$$

This branchwise construction decouples task-level preservation from direction-level transfer: historical task knowledge remains preserved in complete private experts, while the historical branch selectively transfers that knowledge through independently gated functional directions.

3.4 Multimodal Task-Free Retrieval

After learning each task, PHAROS stores its private expert, signed gates, and multimodal projector as a preserved task-specific configuration. Because no task identity is provided at inference time, PHAROS must select the appropriate preserved configuration from the input. To support this selection, it records an image anchor and a text anchor for every task. Let ϕ_v and ϕ_x denote frozen CLIP image and text encoders, respectively. During stage t, the anchors are computed as

$$a^{v}_{t} = \frac{1}{n_t} \sum_{i=1}^{n_t} \phi_v(v_i), \qquad a^{x}_{t} = \frac{1}{n_t} \sum_{i=1}^{n_t} \phi_x(x_i), \tag{10}$$

where a^{v}_{t} and a^{x}_{t} summarize the image and instruction features of task $\mathcal{T}_t$. Once computed, these anchors remain fixed for subsequent retrieval. For an inference input (v, x), the same encoders produce guide features $z^{v} = \phi_v(v)$ and $z^{x} = \phi_x(x)$. Each learned task $\mathcal{T}_j$, $1 \le j \le t$, receives the score

$$q_j(v, x) = \frac{1}{2}\left[\cos(z^{v}, a^{v}_{j}) + \cos(z^{x}, a^{x}_{j})\right], \tag{11}$$

Table 1: **Performance (accuracy in %) on COIN benchmark.** We report final accuracy, average accuracy (Avg.ACC), forgetting, and new-task accuracy (New.ACC). Best results are boldfaced.

Method	SQA	TextVQA	ImageNet	GQA	VizWiz	Ground.	VQAv2	OCR-VQA	Avg.ACC↑	Forgetting.↓	New.ACC↑
Zero-shot	49.91	2.88	0.33	2.08	0.90	0.00	0.68	0.17	7.12	–	–
Multi-task	56.77	49.35	95.55	56.65	53.90	30.09	59.50	55.65	57.18	–	–
PerTaskFT	82.48	62.23	96.55	60.67	60.64	32.24	66.28	61.59	65.41	–	–
LoRA (Hu et al., 2022)	21.26	28.74	10.25	36.78	32.45	0.83	42.50	57.08	28.74	37.29	61.36
LwF (Li & Hoiem, 2017)	63.14	39.60	8.90	34.83	14.53	2.48	40.67	62.35	33.31	22.32	52.58
EWC (Kirkpatrick et al., 2017)	67.41	40.41	8.18	35.05	37.88	2.67	41.27	61.02	36.74	20.51	54.68
MT (Zhu et al., 2024)	79.63	55.47	35.64	58.70	44.37	32.20	62.21	61.59	53.73	14.03	66.00
PGP (Qiao et al., 2024a)	85.17	56.85	32.26	**61.74**	49.43	32.74	65.74	62.20	55.77	12.94	67.09
CIA* (Qiao et al., 2024b)	75.63	54.47	43.64	60.70	43.37	36.00	65.21	63.59	55.33	7.04	61.49
SEFE (Chen et al., 2025)	75.35	58.66	83.10	54.25	48.85	16.75	65.35	**66.25**	58.57	11.94	69.02
MoELoRA (Chen et al., 2024)	58.92	38.59	8.85	37.10	44.25	2.45	41.40	55.35	35.86	25.71	58.36
AdaLoRA (Zhang et al., 2023)	73.40	51.29	35.47	44.53	46.75	0.93	55.86	62.03	46.28	23.99	63.27
EProj (He et al., 2026a)	78.51	57.53	92.35	55.93	44.67	36.59	63.74	57.00	60.79	5.42	65.54
BranchLoRA (Zhang et al., 2025)	68.24	40.18	24.60	41.40	49.83	15.94	51.23	62.14	44.20	23.98	65.18
ProgLoRA (Yu et al., 2025)	74.84	51.83	83.90	49.93	53.87	31.19	62.71	64.44	59.09	7.53	65.68
PCLR (Qiao et al., 2026)	78.33	58.24	86.08	58.14	57.61	33.04	64.17	61.92	62.19	3.39	65.16
MoBLoRA (Chen et al., 2026)	84.96	60.46	96.87	60.07	**61.70**	32.95	64.82	60.85	65.34	**0.00**	65.34
PHAROS (Ours)	**86.23**	**62.54**	**97.43**	60.71	59.74	**65.00**	**64.95**	64.27	**70.11**	0.32	**70.39**

where $\cos(\cdot,\cdot)$ denotes cosine similarity and $q_j(v,x)$ measures the agreement between the input and the two anchors of task $\mathcal{T}_j$. For a text-only input, PHAROS uses only the text similarity. It then selects

$$\hat{t} = \arg\max_{1\le j\le t} q_j(v,x), \tag{12}$$

where $\hat{t}$ is the retrieved task index. The same index retrieves $\Psi_{\hat{t}}$, the private expert $\mathcal{E}_{\hat{t}}$, and the signed gates $\{g^{\ell}_{\hat{t},s,k} \mid \ell \in \mathcal{L}, 1 \le s < \hat{t}, 1 \le k \le \rho^{\ell}_s\}$, preserving their training-time association without requiring a task label.

4 Experiments

4.1 Experimental Setup

We evaluate on the six-task UCIT stream (Guo et al., 2025) and the eight-task CoIN stream (Chen et al., 2024), using their standard task orders and official task-specific evaluators. UCIT reports trajectory-level Avg. and final-checkpoint Last, whereas CoIN reports Avg.ACC, forgetting, and New.ACC. All methods use LLaVA-v1.5-7B with CLIP ViT-L/14-336 encoders. PHAROS is trained without replay and performs task-free inference. Exact task orders, metric definitions, optimization settings, baseline provenance, and checkpoint audits are provided in Appendix A.2.

4.2 Main Results

To evaluate PHAROS across complementary CMIT conditions, we first measure continual performance and new-task acquisition on the longer eight-task CoIN stream, then examine trajectory-level performance and final retention on UCIT, and finally vary the UCIT task order to test whether the method depends on a favorable sequence. On CoIN, Table 1 shows that PHAROS achieves 70.11 Avg.ACC, exceeding the strongest baseline by 4.77 points, while its New.ACC of 70.39 surpasses the previous best of 69.02. PHAROS also obtains the best final accuracy on four of the eight tasks, with the largest improvement on Grounding, from 36.59 to 65.00. Together, these results show that PHAROS improves aggregate continual performance while retaining strong new-task acquisition across distinct task types. On UCIT, Table 2 shows that PHAROS obtains 77.60 Avg. and 77.60 Last, improving over EBLoRA by 3.58 and 4.80 points, respectively. It achieves the best final score on five of the six tasks and remains only 0.78 points below MoELoRA on Flickr30k, indicating that the improvement extends across recognition, knowledge-intensive question answering, visual reasoning, and captioning. Because historical knowledge available for transfer changes with task order, Table 3 further compares the reversed, and Random-42 sequences. The resulting 0.10-point range indicates that PHAROS maintains stable continual performance under substantially different historical-task compositions. Figure 3 complements these aggregate results by showing that the gains are distributed across tasks and by jointly visualizing CoIN accuracy and forgetting. To examine whether PHAROS remains effective with a larger language backbone, we additionally evaluate it

Table 2: **Performance (accuracy in %) on UCIT benchmark.** PHAROS is compared with representative continual-learning methods under the same task sequence. Best results are boldfaced.

	Method	ImageNet-R	ArXivQA	Viz-cap	IconQA	CLEVR	Flickr30k	Mean↑
	Zero-shot	16.27	53.73	38.39	19.20	20.63	41.88	31.68
	Multi-task	90.63	91.30	61.81	73.90	73.60	57.45	74.78
Avg	LwF (Li & Hoiem, 2017)	55.60	79.86	53.23	54.87	56.51	56.34	59.40
Avg	EWC (Kirkpatrick et al., 2017)	54.23	80.13	53.14	55.06	57.52	55.94	59.34
Avg	L2P (Wang et al., 2022)	41.52	82.32	51.98	52.21	43.16	52.77	53.99
Avg	O-LoRA (Wang et al., 2023)	75.26	86.73	55.86	58.47	57.38	53.52	64.54
Avg	MoELoRA (Chen et al., 2024)	64.49	82.42	49.54	56.87	56.35	**58.34**	61.33
Avg	HiDe-LLaVA (Guo et al., 2025)	85.70	92.70	54.10	66.87	59.12	55.15	68.94
Avg	Octopus† (Liu et al., 2026)	89.69	91.34	61.04	60.18	68.09	57.61	71.33
Avg	Octopus (Liu et al., 2026)	88.41	93.30	**61.65**	59.91	66.32	56.87	71.08
Avg	CL-MoE‡ (Huai et al., 2025)	79.17	82.12	50.98	53.30	63.70	57.30	64.43
Avg	SEFE‡ (Chen et al., 2025)	85.48	83.58	55.02	73.20	70.35	57.30	70.82
Avg	KeepLoRA‡ (Luo et al., 2026)	86.50	89.78	53.83	71.00	69.50	57.20	71.30
Avg	EBLoRA‡ (Gu et al., 2026)	90.02	94.28	61.08	76.90	66.55	55.30	74.02
Avg	**PHAROS (Ours)**	**90.80**	**94.64**	61.03	**84.27**	**77.33**	57.56	**77.60**
Last	LwF	40.27	75.93	42.76	44.38	37.43	56.34	49.52
Last	EWC	39.05	77.88	43.24	45.33	39.72	55.94	50.20
Last	L2P	32.73	80.41	43.72	42.16	39.25	52.77	48.51
Last	O-LoRA	69.36	82.42	48.64	53.66	42.53	53.52	58.36
Last	MoELoRA	49.87	77.63	43.65	46.40	36.47	**58.34**	52.06
Last	HiDe-LLaVA	80.50	89.83	48.78	62.90	47.97	55.15	64.19
Last	Octopus†	88.80	89.83	56.44	59.90	70.60	57.11	70.45
Last	Octopus	88.20	93.03	58.46	60.50	69.00	56.87	71.01
Last	CL-MoE‡	61.20	75.80	44.40	52.60	54.40	57.30	57.62
Last	SEFE‡	80.20	79.10	47.10	69.40	65.70	57.30	66.47
Last	KeepLoRA‡	82.43	86.70	46.54	67.80	66.40	57.18	67.84
Last	EBLoRA‡	89.60	94.20	56.60	75.20	66.20	55.30	72.80
Last	**PHAROS (Ours)**	**90.80**	**94.60**	**61.04**	**84.27**	**77.33**	57.56	**77.60**

† denotes the historical-task proxy variant of Octopus. ‡ denotes Avg rows reconstructed from published rounded matrices.

Table 3: **Task-order robustness on UCIT.** For each order, the task row gives the actual training sequence, and the following rows report Avg. and Last in that same sequence.

Order	Row	Task 1	Task 2	Task 3	Task 4	Task 5	Task 6	Mean↑
Reverse	Task	Flickr30k	CLEVR-Math	IconQA	Viz-cap	ArXivQA	ImageNet-R	
	Avg.	57.70	77.90	84.27	60.73	94.03	90.93	77.59
	Last	57.67	77.90	84.27	60.73	94.03	90.93	77.59
Random-42	Task	IconQA	ArXivQA	Viz-cap	CLEVR-Math	ImageNet-R	Flickr30k	
	Avg.	85.13	94.13	60.80	77.60	90.77	57.71	77.69
	Last	85.13	94.13	60.81	77.60	90.77	57.71	77.69

Table 4: **Performance (accuracy in %) of PHAROS with LLaVA-v1.5-13B on UCIT.** The performance shows that PHAROS remains effective on the larger backbone.

Method	Metric	ImageNet-R	ArXivQA	Viz-cap	IconQA	CLEVR	Flickr30k	Mean↑
PHAROS-13B	Avg.	91.17	95.02	64.51	85.37	78.63	59.51	79.03
	Last	91.17	94.97	64.51	84.70	78.63	59.51	78.92

with LLaVA-v1.5-13B under the canonical UCIT task order. Table 4 reports each task immediately after its acquisition and again after the complete six-task sequence.

4.3 Ablation Study

The main results establish the overall advantage of PHAROS but do not identify the contribution of each component. We therefore construct the method incrementally under the same CoIN task order in Table 5. Task-level preservation provides the largest gain, increasing Avg.ACC from 54.05 to 67.96 and reducing forgetting from 10.33 to 2.99. Direction-leve transfer further raises

Table 5: **Ablation study on CoIN.** Each PHAROS component provides a consistent gain.

Configuration	Avg.ACC↑	Forgetting.↓	New.ACC↑
Baseline	54.05	10.33	63.08
+ Task-level preservation	67.96	2.99	70.57
+ Direction-leve transfer	69.82	1.35	**71.01**
+ Projector (PHAROS)	**70.11**	**0.32**	70.39

Avg.ACC to 69.82 and New.ACC to 71.01 while reducing forgetting to 1.35, showing that selective historical transfer complements task-level preservation. Adding task-specific projector preservation produces the highest Avg.ACC of 70.11 and further reduces forgetting to 0.32, with New.ACC remaining above the baseline at 70.39. These incremental gains show that private experts, direction-level transfer, and task-specific projectors address complementary aspects of continual adaptation.

4.4 Forgetting Dynamics

Final aggregate metrics cannot reveal when forgetting emerges or how it accumulates as additional tasks arrive. At stage t, forgetting is computed as the mean gap between the best score previously achieved on each earlier task and its current score; Appendix A.2 provides the formal definition. Figure 4(b) shows that PHAROS maintains low forgetting as the sequence grows, whereas competing methods accumulate larger degradation. This trajectory-level evidence confirms that task-level preservation remains effective throughout continual training rather than only at the final checkpoint.

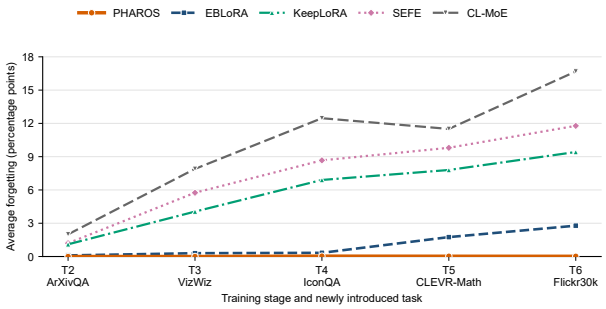


Figure 4: **Stage-wise forgetting on UCIT.** PHAROS maintains near-zero forgetting as the task sequence grows, demonstrating consistent preservation of historical knowledge throughout continual learning.

4.5 How Does PHAROS Use Historical Directions?

The ablation verifies the benefit of historical transfer but does not show whether PHAROS actually controls historical knowledge at the direction level. We therefore inspect the signed gates learned on CoIN and UCIT. Figure 1 shows how directions from a fixed historical expert vary across target tasks, while Figure 5 summarizes the positive and negative gate strengths over the complete historical bank. Both signs occur for every target task, and their strengths vary across tasks and benchmarks. These results show that PHAROS learns a target-task-specific contribution for each historical direction rather than transferring an entire historical expert with a single coefficient.

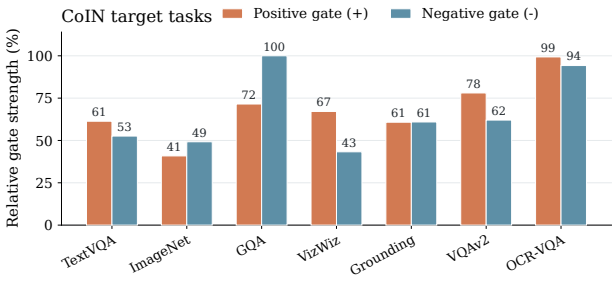


Figure 5: **Task-dependent signed historical transfer on CoIN.** The relative positive and negative gate strengths vary across target tasks, showing that PHAROS adaptively determines the sign and strength of historical-direction contributions for each new task.

5 Conclusion

This work is motivated by the observation that task-level preservation protects historical experts, but effective transfer requires finer control over their functional directions. Assigning a single strength to a complete expert can suppress useful directions when the value is low and amplify irrelevant or conflicting directions when it is high. Building on this observation, we proposed PHAROS, which preserves each task with a frozen private LoRA expert and associated multimodal projector, and transfers historical knowledge through functional directions derived from historical experts. For each new task, task-specific signed gates adaptively determine the sign and strength of every historical direction's contribution, enabling direction-level transfer without weakening task-level preservation. Experiments on CoIN and UCIT demonstrate that preserving task knowledge as complete experts while transferring historical knowledge at a finer granularity improves continual performance across diverse multimodal tasks. Maintaining a separate expert–projector pair for each task and an expanding historical transfer branch increases storage and inference costs in long task streams. Future work can improve scalability through projector sharing and selective direction activation.

REFERENCES

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736, 2022.

Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7120–7129. IEEE, 2017.

Chang Che, Ziqi Wang, Pengwan Yang, Cheems Wang, Hui Ma, and Zenglin Shi. Lora in lora: Towards parameter-efficient architecture expansion for continual visual instruction tuning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 19978–19986, 2026.

Cheng Chen, Junchen Zhu, Xu Luo, Heng T Shen, Jingkuan Song, and Lianli Gao. Coin: A benchmark of continual instruction tuning for multimodel large language models. *Advances in neural information processing systems*, 37:57817–57840, 2024.

Jinpeng Chen, Runmin Cong, Yuzhi Zhao, Hongzheng Yang, Guangneng Hu, Horace Ip, and Sam Kwong. SEFE: Superficial and essential forgetting eliminator for multimodal continual instruction tuning. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 7982–8001. PMLR, 13–19 Jul 2025.

Pei Chen, Xilai Wang, Zejian Li, Lingyun Sun, et al. From experts to bases: Orthogonal subspace mixture for continual multimodal instruction tuning. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 10545–10561, 2026.

Wenliang Dai, Junnan Li, Dongxu Li, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale N Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. *Advances in neural information processing systems*, 36:49250–49267, 2023.

Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385, 2021.

Zhan Fa, Yue Duan, Jian Zhang, Lei Qi, Wanqi Yang, and Yinghuan Shi. Decomposing and composing: Towards efficient vision-language continual learning via rank-1 expert pool in a single lora. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 20985–20993, 2026.

Hao Gu, Mao-Lin Luo, Zi-Hao Zhou, Han-Chen Zhang, Min-Ling Zhang, and Tong Wei. Spectral imbalance causes forgetting in low-rank continual adaptation. *arXiv preprint arXiv:2602.00722*, 2026.

Haiyang Guo, Fanhu Zeng, Ziwei Xiang, Fei Zhu, Da-Han Wang, Xu-Yao Zhang, and Cheng-Lin Liu. Hide-llava: Hierarchical decoupling for continual instruction tuning of multimodal large language model. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 13572–13586, 2025.

Jinghan He, Haiyun Guo, Kuan Zhu, Ming Tang, and Jinqiao Wang. Continual instruction tuning for large multimodal models. *IEEE Transactions on Image Processing*, 2026a.

Lingfeng He, De Cheng, Huaijie Wang, Xi Yang, Nannan Wang, and Xinbo Gao. Task-driven subspace decomposition for knowledge sharing and isolation in lora-based continual learning. *arXiv preprint arXiv:2603.00191*, 2026b.

Zhiyan Hou, Haiyun Guo, Haokai Ma, Yandu Sun, Yonghui Yang, and Jinqiao Wang. Pass-moe: Mitigating misaligned co-drift among router and experts via pathway activation subspaces for continual learning. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 31959–31972, 2026.

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International conference on machine learning*, pp. 2790–2799. PMLR, 2019.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022.

Tianyu Huai, Jie Zhou, Xingjiao Wu, Qin Chen, Qingchun Bai, Ze Zhou, and Liang He. Cl-moe: Enhancing multimodal large language model with dual momentum mixture-of-experts for continual visual question answering. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 19608–19617. IEEE, 2025.

Chengsong Huang, Qian Liu, Bill Yuchen Lin, Tianyu Pang, Chao Du, and Min Lin. Lorahub: Efficient cross-task generalization via dynamic lora composition, 2023.

Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing Models with Task Arithmetic. In *International Conference on Learning Representations*, 2023.

Prakhar Kaushik, Ankit Vaidya, Shravan Chaudhari, Rama Chellappa, and Alan Yuille. Shared lora subspaces for almost strict continual learning. *arXiv preprint arXiv:2602.06043*, 2026.

James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.

Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pp. 19730–19742. PmLR, 2023.

Songze Li, Mingyu Gao, Tonghua Su, Xu-Yao Zhang, and Zhongjie Wang. Multimodal continual instruction tuning with dynamic gradient guidance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10820–10829, 2026.

Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023.

Yuehao Liu, Shanyan Guan, Weijia Zhang, Xuanming Shang, Yanhao Ge, Wei Li, and Chao Ma. Octopus: History-free gradient orthogonalization for continual learning in multimodal large language models, 2026.

Mao-Lin Luo, Zi-Hao Zhou, Yi-Lin Zhang, Yuanyu Wan, Min-Ling Zhang, and Tong Wei. Keeplora: Continual learning with residual gradient adaptation. In *International Conference on Learning Representations*, volume 2026, pp. 136017–136040, 2026.

Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7765–7773. IEEE, 2018.

Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *European Conference on Computer Vision*, pp. 72–88. Springer, 2018.

German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural networks*, 113:54–71, 2019.

Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. Adapterfusion: Non-destructive task composition for transfer learning. In *Proceedings of the 16th conference of the European chapter of the association for computational linguistics: main volume*, pp. 487–503, 2021.

Jingyang Qiao, Zhizhong Zhang, Xin Tan, Chengwei Chen, Yanyun Qu, Yong Peng, and Yuan Xie. Prompt gradient projection for continual learning. In *International Conference on Learning Representations*, volume 2024, pp. 16294–16315, 2024a.

Jingyang Qiao, Zhizhong Zhang, Xin Tan, Yanyun Qu, Shouhong Ding, and Yuan Xie. Large continual instruction assistant. *arXiv preprint arXiv:2410.10868*, 2024b.

Jingyang Qiao, Zhizhong Zhang, Shaohui Liu, Yuan Xie, et al. Pclr: Progressively compressed lora for multimodal continual instruction tuning. In *International Conference on Learning Representations*, volume 2026, pp. 55054–55089, 2026.

Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International conference on machine learning*, pp. 4548–4557. PMLR, 2018.

Xiao Wang, Tianze Chen, Qiming Ge, Han Xia, Rong Bao, Rui Zheng, Qi Zhang, Tao Gui, and Xuan-Jing Huang. Orthogonal subspace learning for language model continual learning. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 10658–10671, 2023.

Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *2022 IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pp. 139–149. IEEE, 2022.

Prateek Yadav, Derek Tam, Leshem Choshen, Colin A Raffel, and Mohit Bansal. Ties-merging: Resolving interference when merging models. *Advances in neural information processing systems*, 36:7093–7115, 2023.

Enneng Yang, Zhenyi Wang, Li Shen, Shiwei Liu, Guibing Guo, Xingwei Wang, and Dacheng Tao. Adamerging: Adaptive model merging for multi-task learning. In *International Conference on Learning Representations*, volume 2024, pp. 22743–22763, 2024.

Yahan Yu, Duzhen Zhang, Yong Ren, Xuanle Zhao, Xiuyi Chen, and Chenhui Chu. Progressive lora for multimodal continual instruction tuning. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 2779–2796, 2025.

Duzhen Zhang, Yong Ren, Zhong-Zhi Li, Yahan Yu, Jiahua Dong, Chenxing Li, Zhilong Ji, and Jinfeng Bai. Enhancing multimodal continual instruction tuning with branchlora. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5743–5756, 2025.

Qingru Zhang, Minshuo Chen, Alexander Bukharin, Nikos Karampatziakis, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. Adalora: Adaptive budget allocation for parameter-efficient fine-tuning. *arXiv preprint arXiv:2303.10512*, 2023.

Ziyu Zhao, Yixiao Zhou, Xin Yu, Zhi Zhang, Didi Zhu, Tao Shen, Zexi Li, Jinluan Yang, Xuwu Wang, Jing Su, et al. Each rank could be an expert: Single-ranked mixture of experts lora for multi-task learning. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*, pp. 1998–2009, 2026.

Didi Zhu, Zhongyi Sun, Zexi Li, Tao Shen, Ke Yan, Shouhong Ding, Kun Kuang, and Chao Wu. Model tailor: Mitigating catastrophic forgetting in multi-modal large language models. *arXiv preprint arXiv:2402.12048*, 2024.

A APPENDIX

This appendix provides the task streams, metric definitions, optimization settings, and validation procedures used in our experiments. It also reports complete task trajectories, an additional task-order experiment, the full signed-direction atlases, and representative evaluation examples from both benchmarks.

AI USE STATEMENT

Generative AI tools were used to assist with refining the conceptual framing and research hypotheses, providing feedback on the methodology and experimental design, implementing and debugging parts of the research code, checking metric calculations, interpreting experimental results, and assisting with translation. They were not used to generate synthetic datasets, fabricate or alter experimental measurements, or formulate proofs. We reviewed and tested all AI-assisted code and verified the reported results against the original checkpoints and evaluation artifacts. We take full responsibility for the final content of this work.

A.1 TASK STREAMS AND EVALUATION

UCIT follows the six-task sequence ImageNet-R, ArXivQA, VizWiz captioning, IconQA, CLEVR-Math, and Flickr30k captioning (Guo et al., 2025). CoIN follows the eight-task sequence ScienceQA, TextVQA, ImageNet, GQA, VizWiz, visual grounding, VQAv2, and OCR-VQA (Chen et al., 2024). After learning task t, we evaluate the resulting checkpoint on every task observed up to that stage. Each dataset is scored with its official task-specific evaluator, and all results are reported as percentages. Generation uses greedy decoding with a temperature of zero.

A.2 METRIC DEFINITIONS

Let $a_{t,i}$ denote the score on task i after training through task t, where $1 \leq i \leq t \leq T$. The lower-triangular matrix $A = [a_{t,i}]$ records the complete continual trajectory.

UCIT reports *Last*, the macro-average over all tasks at the final checkpoint,

$$\text{Last} = \frac{1}{T}\sum_{i=1}^{T} a_{T,i}, \tag{13}$$

and *Avg.*, which first averages each task over all stages following its acquisition and then assigns equal weight to every task,

$$\text{Avg.} = \frac{1}{T}\sum_{i=1}^{T} \frac{1}{T-i+1}\sum_{t=i}^{T} a_{t,i}. \tag{14}$$

This definition prevents early tasks, which are evaluated at more stages, from receiving greater weight than later tasks.

For CoIN, average accuracy measures final performance over the complete task stream,

$$\text{Avg.ACC} = \frac{1}{T}\sum_{i=1}^{T} a_{T,i}, \tag{15}$$

whereas new-task accuracy measures acquisition performance on the diagonal,

$$\text{New.ACC} = \frac{1}{T}\sum_{i=1}^{T} a_{i,i}. \tag{16}$$

Forgetting excludes the last task because it is not followed by further training. For each earlier task, it measures the decrease from its best post-acquisition score to its final score,

$$\text{Forgetting} = \frac{1}{T-1}\sum_{i=1}^{T-1}\left(\max_{t\in\{i,\ldots,T\}} a_{t,i} - a_{T,i}\right). \tag{17}$$

The stage-wise forgetting value used in the retention curves is

$$F_t = \frac{1}{t-1} \sum_{i=1}^{t-1} \left(\max_{s \in \{i,\ldots,t\}} a_{s,i} - a_{t,i} \right), \qquad t \geq 2. \tag{18}$$

A.3 Training Configuration

All main experiments use LLaVA-v1.5-7B with CLIP ViT-L/14-336, while the language and vision backbones remain frozen. PHAROS jointly optimizes the current private expert, signed gates over historical directions, and task-specific multimodal projector before freezing the resulting task-specific parameters. Unless stated otherwise, experiments use seed 42, no replay, and task-free inference.

A.4 Baseline Provenance and Run Validation

For UCIT, per-task Last values are taken from the corresponding final checkpoints. Avg. entries marked by ‡ in Table 2 are reconstructed from the rounded lower-triangular matrices collected by Gu et al. (2026); small rounding differences may remain. CoIN baselines are taken from the matched LLaVA-v1.5-7B comparison compiled by Chen et al. (2026). These values use the same benchmark orders and backbone but are reported results rather than reruns in our codebase.

Before a checkpoint is used as the predecessor of a new task, we verify that all historical private experts, historical gate rows, routing memories, and archived projectors remain unchanged. The projector bank at stage t must contain the snapshots produced at stages 1 through t. We evaluate a run only after its checkpoint passes this causal-chain audit, and we accept an evaluation only when the merged prediction file contains the expected number of samples.

B Additional Experimental Results

Table 6: **Performance trajectories of PHAROS on UCIT.** Each sequence begins when the corresponding task is learned and continues through all subsequent stages. Final scores are boldfaced.

Dataset	Observed trajectory (%)
ImageNet-R	$90.80 \rightarrow 90.80 \rightarrow 90.80 \rightarrow 90.80 \rightarrow 90.80 \rightarrow \mathbf{90.80}$
ArXivQA	$94.70 \rightarrow 94.70 \rightarrow 94.60 \rightarrow 94.60 \rightarrow \mathbf{94.60}$
VizWiz	$61.02 \rightarrow 61.02 \rightarrow 61.02 \rightarrow \mathbf{61.04}$
IconQA	$84.27 \rightarrow 84.27 \rightarrow \mathbf{84.27}$
CLEVR-Math	$77.33 \rightarrow \mathbf{77.33}$
Flickr30k	$\mathbf{57.56}$

B.1 Complete UCIT Trajectories

Table 6 reports every observed UCIT score used to compute Avg. and Last. The trajectories are nearly flat after task acquisition: the largest change is only 0.10 point on ArXivQA, while the remaining tasks change by at most 0.02 point. These values provide the complete evidence behind the aggregate UCIT results in Table 2.

B.2 Additional Task Orders

Tables 7 and 8 extend the task-order analysis to the Random-2026 sequence on UCIT and the reversed sequence on CoIN. For each sequence, acquisition reports the score obtained when a task is first learned, whereas the final row reports its score after the complete task stream.

B.3 Complete Signed-Direction Atlases

The main paper visualizes a compact subset of the learned gates. Figures 6 and 7 provide the complete view over all eight functional directions and every valid causal source–target relation. A cell



Figure 6: **Complete signed-direction atlas on CoIN.** Each panel corresponds to one functional direction. Rows denote target tasks and columns denote earlier source tasks in continual order. Gray cells mark unavailable self-relations and future-task relations. All panels use the same symmetric CoIN color scale.

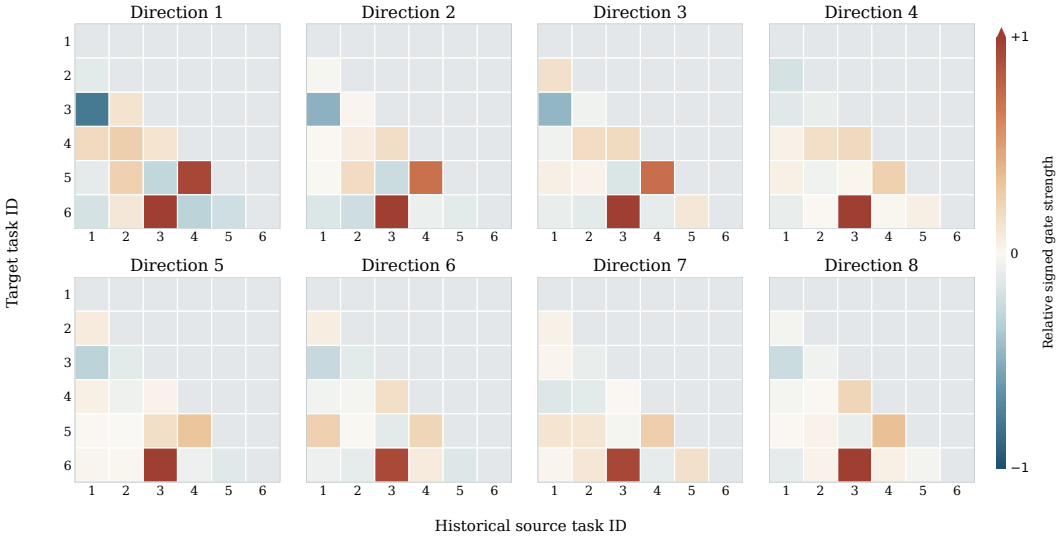


Figure 7: **Complete signed-direction atlas on UCIT.** Each panel shows one functional direction over every valid source–target relation. Gray cells mark unavailable self-relations and future-task relations. All panels use the same symmetric UCIT color scale.

is shown only when the source task precedes the target task. The coexistence of positive, negative, and near-zero gates across the atlases shows that PHAROS does not apply a uniform coefficient to a historical expert. Instead, the contribution assigned to a direction depends jointly on its source and the current target task.

B.4 Benchmark Tasks and Output Formats

Figure 8 presents one traceable evaluation example from every CoIN and UCIT task. The two streams cover classification, multiple-choice reasoning, short-form visual question answering, numerical reasoning, visual grounding, OCR, and caption generation. This diversity requires the model to preserve not only visual concepts but also task-specific output conventions, motivating evaluation with each benchmark's official task-specific scorer.

Table 7: **PHAROS under the Random-2026 UCIT order.** The Avg. and final scores show that PHAROS preserves task performance under an additional task sequence.

Method	Metric	CLEVR-Math	ArXivQA	IconQA	Flickr30k	Viz-cap	ImageNet-R	Mean↑
PHAROS	Avg.	76.50	94.43	83.40	57.33	60.78	90.67	77.18
	Last	76.50	94.40	83.40	57.31	60.78	90.67	77.18

Table 8: **PHAROS under the reversed CoIN order.** The Avg. and final scores show that PHAROS preserves task performance under the reversed eight-task sequence.

Method	Metric	OCR-VQA	VQAv2	Grounding	VizWiz	GQA	ImageNet	TextVQA	ScienceQA	Mean↑
PHAROS	Acquisition	64.38	66.11	66.34	58.91	60.92	97.45	63.66	84.81	70.32
	Last	64.02	65.42	66.34	58.60	60.84	97.45	63.63	84.81	70.14

CoIN — 8-task continual stream

Order	Task	Example	Instruction and target	Response
1	ScienceQA		What is the name of the colony shown? Target **B**	Choice
2	TextVQA		what number is on the player's jersey? Target **22**	Short answer
3	ImageNet		What is the object in the image? Target **Cabbage butterfly**	Class label
4	GQA		Is it overcast? Target **No**	Yes / No
5	VizWiz		Is there anything on the screen? Target **yes**	Yes / No
6	Grounding		Please provide the bounding box coordinate of the region this sentence describes:... Target **[0.38,0.41,0.51,0.67]**	Coordinates
7	VQAv2		Where is he looking? Target **down**	Short answer
8	OCR-VQA		Who wrote this book? Target **Amy Knapp**	Short answer

UCIT — 6-task continual stream

Order	Task	Example	Instruction and target	Response
1	ImageNet-R		What is the object in the image? Target **Centipede**	Class label
2	ArXivQA		What can be inferred about the error bars associated with the Γ_0 measurements? Target **D**	Choice
3	Viz-cap		What is happening in the image? Target **A computer screen shows a repair prompt on the screen**	Caption
4	IconQA		Bill is washing his dog one morning. His watch shows the time. What time is it? Target **1**	Choice
5	CLEVR-Math		Add 1 green rubber cubes. How many objects exist? Target **10**	Number
6	Flickr30k		What is happening in the image? Target **Several men in hard hats are operating a giant pulley system**	Caption

Figure 8: **Representative evaluation examples from CoIN and UCIT.** Tasks are arranged in their canonical continual order. Each card contains a real benchmark image, its instruction, the reference target, and the required response format.