

# DeLo: Dual Decomposed Low-Rank Experts Collaboration for Continual Missing Modality Learning

Xiwei Liu<sup>1</sup>, Yulong Li<sup>1</sup>, Feilong Tang<sup>1,2</sup>, Imran Razzak<sup>1,\*</sup>

<sup>1</sup>Mohamed bin Zayed University of Artificial Intelligence

<sup>2</sup>Monash University

{Xiwei.Liu, Imran.Razzak}@mbzuai.ac.ae

## Abstract

Adapting Large Multimodal Models (LMMs) to real-world scenarios poses the dual challenges of learning from sequential data streams while handling frequent modality incompleteness, a task known as Continual Missing Modality Learning (CMML). However, existing works on CMML have predominantly relied on prompt tuning, a technique that struggles with this task due to cross-task interference between its learnable prompts in their shared embedding space. A naive application of Low-Rank Adaptation (LoRA) with modality-shared module will also suffer modality interference from competing gradients. To this end, we propose DeLo, the first framework to leverage a novel dual-decomposed low-rank expert architecture for CMML. Specifically, this architecture resolves modality interference through decomposed LoRA expert, dynamically composing LoRA update matrix with rank-one factors from disentangled modality-specific factor pools. Embedded within a task-partitioned framework that structurally prevents catastrophic forgetting, this expert system is supported by two key mechanisms: a Cross-Modal Guided Routing strategy to handle incomplete data and a Task-Key Memory for efficient, task-agnostic inference. Extensive experiments on established CMML benchmarks demonstrate that our method significantly outperforms state-of-the-art approaches. This highlights the value of a principled, architecturally-aware LoRA design for real-world multimodal challenges.

**Code** — <https://github.com/Xiwei-web/DeLo>

## Introduction

Pre-trained multimodal models have shown great potential in many applications (Radford et al. 2021; Li et al. 2023). When fine-tuning these pre-trained models on downstream tasks, missing modality issues often occur due to equipment failure, data corruption, privacy concerns, etc. Existing methods (Zhao, Li, and Jin 2021; Ma et al. 2021; Lee et al. 2023; Jang, Wang, and Kim 2024) address missing modality issues by reconstructing missing information or aligning multimodal features. However, the sequential data in real-world applications will make these methods suffer from catastrophic forgetting (Guo et al. 2025), leading to

\*Corresponding author  
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

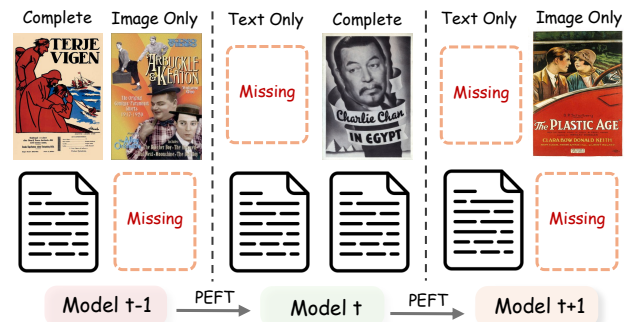


Figure 1: Illustration of continual missing modality learning.

performance degradation. To handle forgetting, an intuitive idea is to store and retrain all old data but it creates large storage overheads and potential privacy-leak issues.

In recent years, continual learning (CL) has made great progress, such as replay-based methods (Rolnick et al. 2019; Buzzega et al. 2020; Cha, Lee, and Shin 2021), regularization-based methods (Kirkpatrick et al. 2017; Zenke, Poole, and Ganguli 2017; Aljundi et al. 2018), and architecture-based methods (Serra et al. 2018; Li et al. 2019; Ebrahimi et al. 2020). Prompt-based methods (Wang et al. 2022b,a; Wang, Huang, and Hong 2022; Smith et al. 2023) have attracted much attention due to their simplicity and effectiveness. However, most of these methods are unimodal and difficult to transfer to multimodal field. Multimodal methods (Wang, Huang, and Hong 2022; Qian et al. 2023) always depend on language-image models such as CLIP (Radford et al. 2021), which makes them difficult to apply to other fields where there are more modalities. LoRA-based methods (Wistuba et al. 2023; Liang and Li 2024; Chen et al. 2024) have also emerged. These approaches aim to better resolve the interference between new and old tasks to achieve a superior stability-plasticity trade-off.

Prompt-based methods were recently extended to address a more complex challenge: Continual Missing Modality Learning (Figure 1) (Zhao et al. 2024; Yue et al. 2025; Guo et al. 2025). However, while this line of research has proven fruitful, its foundational reliance on prompt tuning as the parameter efficient fine tuning (PEFT) mechanism warrants closer examination. Recent studies (Wistuba et al. 2024) argue that this choice is often an "undefended and

Method	Missing Modality	Continual Learning	Prompt Tuning	LoRA
MAP (Lee et al. 2023)	✓	✗	✓	✗
MSP (Jang, Wang, and Kim 2024)	✓	✗	✓	✗
L2P (Wang et al. 2022b)	✗	✓	✓	✗
DualPrompt (Wang et al. 2022a)	✗	✓	✓	✗
InfLoRA (Liang and Li 2024)	✗	✓	✗	✓
RebQ (Zhao et al. 2024)	✓	✓	✓	✗
DeLo (Ours)	✓	✓	✗	✓

Table 1: Comparison of related representative works by problem area and adaptation technique.

unablated decision,” with extensive evidence demonstrating that LoRA (Hu et al. 2022) is substantially more performant than prompt tuning for CL. The core issue is the learnable prompts of prompt tuning are optimized within a shared embedding space, inherently limiting their capacity to prevent the cross-task interference that causes catastrophic forgetting. This evidence motivates a shift towards a LoRA-based approach. However, a naive application of standard LoRA for CMML is also insufficient. Such an approach would still suffer from modality interference, where conflicting updates from these heterogeneous data sources (e.g., image-only vs. text-only) would lead to the learning of entangled and sub-optimal adaptations. Table 1 highlights that there is a clear gap where no existing method effectively tailors a powerful PEFT like LoRA to the dual challenges of CMML.

Based on these findings, our work proposes DeLo, the first LoRA-based framework designed to tackle the CMML problem by leveraging a novel dual-**D**ecomposed **L**ow-rank expert architecture. DeLo adopts dynamic decomposed low-rank expert with factors selected from Modality-Specific Factor Pool through instance-based Cross-Modal Guided Routing strategy. DeLo features a task-partitioned expert architecture to physically isolate knowledge, and utilizes a Task-Key Memory to enable task-agnostic inference. We conduct extensive experiments on two multimodal datasets: UPMC-Food101-CMML and MM-IMDb-CMML. Our proposed method can consistently outperform baselines and state-of-the-art methods significantly on both two datasets.

Our main contributions can be summarized as follows:

- To the best of our knowledge, we are the first to propose a LoRA-based framework for the Continual Missing Modality Learning task, bridging a significant gap in this commonly encountered real-world scenario.
- We propose a novel Dual Decomposed Low-Rank Expert architecture that simultaneously addresses catastrophic forgetting and modality interference. The LoRA experts is composed of selectable low-rank factors from disentangled modality-specific factor pools.
- We demonstrate through extensive experiments that DeLo achieves state-of-the-art performance on CMML benchmarks, delivering superior average performance while structurally mitigating catastrophic forgetting.

## Related Work

### Multi-Modal Learning with Missing Modalities

The challenge of handling missing modalities is a significant bottleneck in multimodal learning, often causing severe performance degradation (Lee et al. 2023; Jin et al. 2023). Research to mitigate this issue has broadly followed three main directions. One line of work focuses on generative solutions (Cai et al. 2018; Zhang et al. 2024; Ke et al. 2025), such as directly synthesizing the data of the absent modality using the available information. Other approaches operate in the feature space, aiming to learn robust representations that are resilient to missing data. These include methods that align features by translating between modalities (Pham et al. 2019), predict the missing modality’s representation from available ones (Zhao, Li, and Jin 2021), or estimate latent features via Bayesian meta-learning (Ma et al. 2021).

More recently, a popular parameter-efficient strategy has emerged that leverages learnable prompt tokens within pre-trained transformers to act as proxy representation for missing information (Lee et al. 2023; Woo et al. 2023; Jang, Wang, and Kim 2024; Zhao et al. 2024; Yue et al. 2025; Guo et al. 2025). However, these prompt-based methods, much like the other strategies discussed, are predominantly designed for static fixed datasets. This shared limitation makes them ill-equipped for real-world scenarios where data arrive sequentially, a core challenge addressed by CMML.

### Continual Learning

Continual learning (CL) endeavors to empower models to learn from a sequential stream of tasks. The primary barrier to this goal is catastrophic forgetting (McCloskey and Cohen 1989), a phenomenon in which acquiring new knowledge leads to the abrupt erasure of previously learned information. Traditional CL strategies, which typically learn from scratch, can be categorized into three families: (1) Regularization-based methods (Zenke, Poole, and Ganguli 2017; Aljundi et al. 2018; Jung et al. 2020) protect prior knowledge by imposing constraints on updates to parameters deemed critical for past tasks. (2) Replay-based methods (Aljundi et al. 2019; Chrysakis and Moens 2020; Liang and Li 2023b,a) store a small buffer of representative samples from previous tasks to rehearse during new task training. (3) Architecture-based methods (Li et al. 2019; Hung et al. 2019; Ebrahimi et al. 2020) mitigate interference by dynamically expanding the network or isolating subsets of task-specific parameters. Although effective to some extent, these methods often face scalability challenges as the complexity of the model or the number of tasks increases.

The paradigm has recently shifted toward leveraging large pre-trained models (Dosovitskiy et al. 2020; He et al. 2022) for continual learning. Early attempts to fully fine-tune these large models proved to be inefficient in parameters (Boschini et al. 2022; Zheng et al. 2023). This has spurred the exploration of PEFT for CL. Prominent examples include prompt adjustment (Wang et al. 2022b,a; Khan et al. 2023; Smith et al. 2023), which has demonstrated strong performance, particularly in class-incremental scenarios, and various methods based on LoRA (Wistuba et al. 2023; Liang

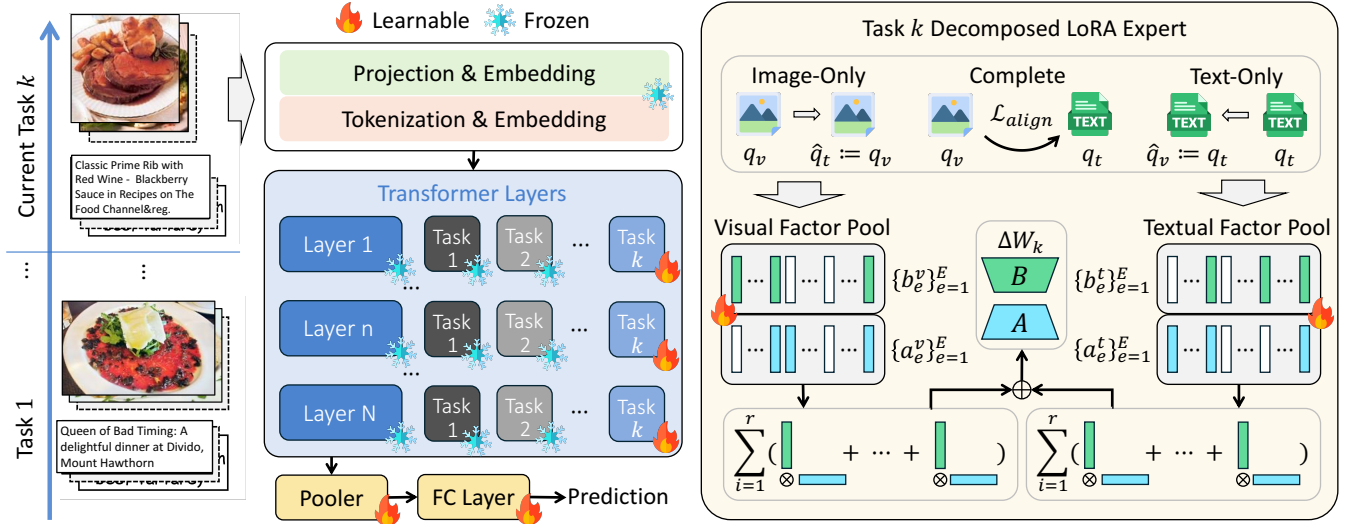


Figure 2: **Overview of our proposed DeLo framework.** The left panel illustrates the task-partitioned architecture for continual learning. To prevent catastrophic forgetting, only the expert modules and classifier for the current task  $k$  are trainable, while the main backbone and all modules for previous tasks (1 to  $k-1$ ) are frozen. The right panel details the dual decomposed LoRA expert for a given Task  $k$ , which consists of Modality-Specific Factor Pools (Visual:  $\mathcal{P}_k^v = \{(a_e^v, b_e^v)\}_{e=1}^E$  and Textual:  $\mathcal{P}_k^t = \{(a_e^t, b_e^t)\}_{e=1}^E$ ). For each input, factors are dynamically selected from these pools to compose modality-specific weight adjustments  $\Delta W^v = \sum_{i=1}^r b_i^v \otimes a_i^v$  and  $\Delta W^t = \sum_{j=1}^r b_j^t \otimes a_j^t$ , whose sum forms the final update  $\Delta W_k$ . The top of this panel shows our Cross-Modal Guided Routing for handling missing modalities, where the query from an available modality serves as a proxy for the missing one (e.g.,  $\hat{q}_t := q_v$  for image-only input) and the Alignment Loss  $\mathcal{L}_{align}$  for complete data.

and Li 2024; Chen et al. 2024), which focus on mitigating intertask interference during sequential learning. However, effective application of these powerful PEFT strategies, especially LoRA, to continual learning tasks with incomplete data remains a key open challenge.

## Proposed Method

### Problem Formulation

Our work tackles the challenging yet practical scenario of CMML, where a model must sequentially learn from tasks containing potentially incomplete multimodal data. This setting imposes three real-world constraints: (1) parameter efficiency for on-device adaptation, since full fine-tuning of large pre-trained models is infeasible; (2) privacy and storage limitations that prohibit data replay, making the model prone to catastrophic forgetting as new tasks arrive; and (3) autonomous system must perform task-agnostic inference, enabling classification without prior task IDs information.

Formally, we consider a sequence of  $T$  tasks,  $\mathcal{T} = \{\mathcal{T}_1, \dots, \mathcal{T}_T\}$ . For any given task  $\mathcal{T}_k$ , the model has access only to its corresponding training dataset  $\mathcal{D}_k$ . To simplify but without loss of generality, we consider a multimodal dataset with text (t) and visual (v) modalities, where each task's dataset  $\mathcal{D}_k$  is structured as a composite of three subsets,  $\mathcal{D} = \{\mathcal{D}^t, \mathcal{D}^v, \mathcal{D}^c\}$ , to capture the nature of modality-missing. Specifically, a set of modality-complete samples  $\mathcal{D}^c = \{(\mathbf{t}_i, \mathbf{v}_i, \mathbf{y}_i)\}_{i=1}^{N_c}$ ; a text-only set  $\mathcal{D}^t = \{(\mathbf{t}_i, \mathbf{y}_i)\}_{i=1}^{N_t}$ ; and a visual-only set  $\mathcal{D}^v = \{(\mathbf{v}_i, \mathbf{y}_i)\}_{i=1}^{N_v}$ , where  $N_c$ ,  $N_t$  and  $N_v$  are the number of samples in each respective subset.

Here,  $\mathbf{t}_i/\mathbf{v}_i$  denotes the embedding of textual/visual content and  $\mathbf{y}_i \in \mathbb{R}^C$  is the label vector where  $C$  is the number of classes. Following (Lee et al. 2023), we assign dummy inputs  $\tilde{\mathbf{t}}/\tilde{\mathbf{v}}$  to represent missing-modality data and obtain  $\tilde{\mathcal{D}}^t = \{(\tilde{\mathbf{t}}_i, \tilde{\mathbf{v}}_i, \mathbf{y}_i)\}_{i=1}^{N_t}$  and  $\tilde{\mathcal{D}}^v = \{(\tilde{\mathbf{t}}_i, \mathbf{v}_i, \mathbf{y}_i)\}_{i=1}^{N_v}$ . During the training, data from previous tasks is not available. In the inference phase, the task ID is unknown.

### Decomposed Low-Rank Expert Architecture

Inspired by the concept of Mixture-of-Experts (Shazeer et al. 2017) and its application to dynamic adapter construction (Shen et al. 2024), DeLo leverages low-rank decomposition factors as a pool of selectable experts. In standard LoRA, a frozen pre-trained weight matrix  $W \in \mathbb{R}^{d_{out} \times d_{in}}$  is augmented by injecting a parallel, trainable weight adjustment matrix  $\Delta W \in \mathbb{R}^{d_{out} \times d_{in}}$ , where  $d_{in}$  and  $d_{out}$  are the input and output dimensions respectively. This adjustment is efficiently parameterized by decomposing it into two smaller, low-rank matrices:  $\Delta W = BA$ , where  $A \in \mathbb{R}^{r \times d_{in}}$  and  $B \in \mathbb{R}^{d_{out} \times r}$ . The rank  $r$  is much smaller than the original dimensions ( $r \ll \min(d_{in}, d_{out})$ ) to ensure high parameter efficiency. The forward pass of a LoRA-adapted layer computes the input hidden state  $h$  as follows:

$$\tilde{h} = Wh + \Delta Wh = Wh + \alpha \cdot BAh, \quad (1)$$

with the  $\alpha \geq 1$  controls the influence of the weight adjustment matrices. As shown in Figure 3, DeLo departs from this static formulation by dynamically composing this weight adjustment matrix from a larger pool of rank-one factors.




Figure 3: Comparison of LoRA parameterizations. (a) Conventional LoRA, where the weight adjustment matrix is represented as a product of two low-rank matrices ( $BA$ ). (b) Decomposed LoRA, where the adjustment is expressed as a sum of  $r$  rank-one matrices, each formed by the outer product of a vector pair  $\sum_{i=1}^r b_i \otimes a_i$ .

This can be expressed via tensor decomposition as:

$$\Delta W = BA = \sum_{i=1}^r b_i \otimes a_i, \quad (2)$$

where  $\otimes$  denotes the outer product and each  $(a_i, b_i)$  is the selected vector pair with  $a_i \in \mathbb{R}^{d_{in} \times 1}$  and  $b_i \in \mathbb{R}^{d_{out} \times 1}$ . Instead of learning a single  $B$  and  $A$  matrix, DeLo maintains expanded pools of expert factors. These pools contain  $E$  total pool pairs,  $\{a_e\}_{e=1}^E, \{b_e\}_{e=1}^E, E > r$ . This allows the model to dynamically construct a unique  $\Delta W$  for specific inputs by selecting  $r$  appropriate factors. Following LoRA (Hu et al. 2022), we use a random Gaussian initialization for  $\{a_e\}_{e=1}^E$  and zero for  $\{b_e\}_{e=1}^E$ .

### Decomposed modality-specific factor pool

A central challenge in adapting multimodal models is the potential for modality interference. When a single, unified pool of experts is used to learn from distinct data streams like vision and text, its parameters are updated by competing gradients, which can lead to the learning of entangled and suboptimal representations. Inspired by the knowledge decomposition strategy in RebQ (Zhao et al. 2024), DeLo employs modality-specific factor pools to ensure the model learns specialized, disentangled adaptations for visual and textual features without interference. Specifically, the expert system is comprised of two distinct and independent sub-pools: a vision factor pool  $\mathcal{P}^v$  and a text factor pool  $\mathcal{P}^t$ . The vision factor pool  $\mathcal{P}^v$  contains  $E_v$  pairs of learnable low-rank factors dedicated to the visual modality:  $\mathcal{P}^v = \{(a_e^v, b_e^v)\}_{e=1}^{E_v}$ , where  $a_e^v$  and  $b_e^v$  are the expert factors for vision. Similarly, the text factor pool  $\mathcal{P}^t$  contains  $E_t$  pairs of factors for the textual modality:  $\mathcal{P}^t = \{(a_e^t, b_e^t)\}_{e=1}^{E_t}$ .

We adopt the instance-based routing method from MixLoRA (Shen et al. 2024) to select factors from expanded pools. For a given input hidden state  $h \in \mathbb{R}^{\text{seq} \times d_{in}}$  from the preceding layer, a query signal  $q = \text{Avg}(h)$  is first generated by averaging the hidden states across the sequence dimension. The input query  $q$  is then processed through a dense layer with weights  $W_A \in \mathbb{R}^{E \times d_{in}}$  and a softmax function to generate scores over the expert pool. The final selection vec-

tor  $g_A \in \{0, 1\}^E$  is determined by a  $\text{top}_r(\cdot)$  operation that creates a binary mask identifying the top- $r$  scoring experts:

$$g_A = \text{top}_r(\text{softmax}(W_A \cdot q)), \quad (3)$$

The selection of factors for the  $B$  matrix employs a more comprehensive process to ensure a cohesive adaptation. It is guided by two signals: the same input query  $q$ , and the composition of the  $A$  matrix that was just selected. The selection scores derived from both signals are fused via element-wise addition to determine the final probabilities for the  $B$  factors. The resulting selection vector  $g_B$ , therefore identifies a set of  $B$  factors that are not only relevant to the input but also highly compatible with the chosen  $A$  factors.

We finally select  $r$  vision experts from  $\mathcal{P}^v$  to compose the vision-specific adjustment matrix  $\Delta W^v$ . Likewise, it selects  $r$  text experts from  $\mathcal{P}^t$  to compose the text-specific adjustment matrix  $\Delta W^t$ . These are formulated as:

$$\Delta W^v = \sum_{i=1}^r b_i^v \otimes a_i^v \quad \text{and} \quad \Delta W^t = \sum_{j=1}^r b_j^t \otimes a_j^t. \quad (4)$$

The final weight adjustment matrix for the input is the collaborative sum of these two modality-specific adjustments. The complete forward pass for a layer within the model is therefore:  $\tilde{h} = Wh + (\Delta W^v + \Delta W^t)h$ .

### Cross-Modal Guided Routing

Our method presents a unique challenge: *How does the system perform modality-specific factor selection when one modality is missing?* We address this by using the query from the available modality as a proxy for the missing one, leveraging the shared semantic space of the multimodal backbone. For a complete sample with visual and textual hidden states  $h_v$  and  $h_t$ , their averaged queries  $q_v = \text{Avg}(h_v)$  and  $q_t = \text{Avg}(h_t)$  are semantically aligned ( $q_v \approx q_t$ ). Thus, when text is missing, the proxy query is simply  $\hat{q}_t := q_v$ . Expert selection for the missing modality then follows the standard routing rule:

$$g_A^t = \text{top}_r(\text{softmax}(W_A^t \cdot \hat{q}_t)). \quad (5)$$

Because both queries share the same hidden dimension  $d_{in}$ , this substitution requires no extra projection, keeping the routing mechanism simple and seamless.

### Task-Key Routing for Task-Agnostic Inference

To prevent catastrophic forgetting, DeLo employs a task-partitioned architecture. The complete set of expert pools  $\mathcal{P}$  can be defined as a collection of  $T$  task-specific pools:

$$\mathcal{P} = \{(\mathcal{P}_1^v, \mathcal{P}_1^t), (\mathcal{P}_2^v, \mathcal{P}_2^t), \dots, (\mathcal{P}_T^v, \mathcal{P}_T^t)\}, \quad (6)$$

During the training of task  $\mathcal{T}_k$ , only the parameters within its corresponding pool  $\mathcal{P}_k$  and the classifier head are updated, while all previously learned pools and heads are frozen.

This partitioned design necessitates a mechanism for task-agnostic inference to select the correct expert pool for a given input. To this end, we introduce a lightweight, non-trainable buffer designed to dynamically route inputs without explicit task IDs. This module maintains a set of representative vectors  $\{\text{key}_k\}_{k=1}^T$ , where each  $\text{key}_k \in \mathbb{R}^{d_{in}}$

$\eta$	#Image	#Text	MAP		MSP		L2P		DualPrompt		RebQ		DeLo	
			AP ( $\uparrow$ )	FG ( $\downarrow$ )	AP ( $\uparrow$ )	FG ( $\downarrow$ )	AP ( $\uparrow$ )	FG ( $\downarrow$ )	AP ( $\uparrow$ )	FG ( $\downarrow$ )	AP ( $\uparrow$ )	FG ( $\downarrow$ )	AP ( $\uparrow$ )	FG ( $\downarrow$ )
10%	100%	90%	20.66	82.50	21.45	80.03	34.09	5.80	59.56	3.38	68.67	9.50	<b>77.19</b>	10.21
	90%	100%	21.53	82.20	23.29	78.98	35.21	5.51	59.90	1.72	72.46	7.64	<b>80.23</b>	7.88
	95%	95%	22.84	80.13	22.12	79.35	34.00	6.55	58.18	6.32	71.06	8.22	<b>78.12</b>	10.32
30%	100%	70%	16.57	83.77	21.37	79.95	30.40	6.53	51.86	6.46	62.06	10.76	<b>74.08</b>	9.69
	70%	100%	24.00	77.12	22.21	76.02	34.75	4.66	52.66	5.11	71.62	6.44	<b>78.20</b>	6.34
	85%	85%	20.66	79.81	21.89	79.01	29.76	6.87	48.56	8.73	66.37	8.07	<b>73.63</b>	8.89
50%	100%	50%	18.18	79.77	17.88	77.31	28.95	6.33	47.70	6.09	55.87	12.07	<b>71.41</b>	14.77
	50%	100%	23.85	75.59	21.57	77.99	32.18	4.55	50.22	5.03	69.23	5.84	<b>76.38</b>	7.39
	75%	75%	18.66	79.80	18.10	78.04	25.30	6.17	43.67	8.52	62.40	8.24	<b>70.59</b>	8.27
70%	100%	30%	17.68	77.82	19.29	78.17	26.57	5.80	43.09	8.96	50.00	12.47	<b>70.04</b>	14.91
	30%	100%	22.48	75.21	21.22	75.21	30.43	3.99	50.28	5.19	69.41	3.73	<b>74.80</b>	5.32
	65%	65%	20.00	75.95	19.76	76.54	24.62	5.66	40.69	7.26	59.92	8.56	<b>67.42</b>	9.41
90%	100%	10%	16.92	76.60	18.62	75.14	24.94	5.60	37.26	11.46	48.15	12.76	<b>67.30</b>	15.68
	10%	100%	24.89	70.84	21.69	71.23	29.77	4.54	51.16	4.80	67.71	4.71	<b>74.21</b>	4.83
	55%	55%	18.41	75.19	20.45	74.31	24.62	5.66	40.69	7.29	54.67	8.78	<b>63.37</b>	10.97

Table 2: **Performance comparison on UPMC-Food101-CMML dataset.** We report AP and FG based on classification accuracy. Missing Modality can occur in both training and testing. The best performance is in **bold**.

serves as a unique centroid for task  $\mathcal{T}_k$ . During the training of task  $\mathcal{T}_k$ , its key is updated smoothly using the exponential moving average (EMA). At each training step  $t$ , the key is updated using the average query signal from  $q_{\text{batch}}$  the current batch as follows:

$$\text{key}_{k,t} = \beta \cdot \text{key}_{k,t-1} + (1 - \beta) \cdot q_{\text{batch},t}, \quad (7)$$

where  $\beta$  is the EMA momentum parameter, and  $\text{key}_{k,t-1}$  is the key’s state before update. At inference, given a test input with query signal  $q_{\text{test}}$ , the module predicts the task ID  $k^*$  by finding the key with the highest cosine similarity:

$$k^* = \underset{k \in \{1, \dots, T\}}{\text{argmax}} (\cos(q_{\text{test}}, \text{key}_k)). \quad (8)$$

The corresponding pool  $\mathcal{P}_{k^*}$  is then activated to process the input. This approach is effective as each Task-Key serves as a robust representation of its task’s feature cluster.

### Training Objective

In this work, we address the classification task and adopt cross-entropy (CE) loss  $\mathcal{L}_c$  (binary cross-entropy loss for the multi-label classification task) for training. To mitigate the potential modality bias inherent in our cross-modal guided routing mechanism when a modality is missing, we introduce two auxiliary losses computed exclusively on the modality-complete samples within each batch. An Alignment Loss  $\mathcal{L}_{\text{align}}$  explicitly encourages the query signals from paired visual ( $q_v$ ) and textual ( $q_t$ ) data to inhabit a shared semantic space by minimizing their cosine distance:

$$\mathcal{L}_{\text{align}} = 1 - \cos(q_v, q_t), \quad (9)$$

thereby enhancing the quality of the proxy signal. Furthermore, to ensure the proxy signal elicits a similar reasoning process as the authentic one, we employ a Consistency Loss  $\mathcal{L}_{\text{con}}$ . This term minimizes the KL divergence between the

model’s output distributions when a router is guided by the true query versus the proxy query. The final objective is:

$$\mathcal{L} = \mathcal{L}_c + \lambda_1 \mathcal{L}_{\text{align}} + \lambda_2 \mathcal{L}_{\text{con}}. \quad (10)$$

where  $\lambda_1$  and  $\lambda_2$  are hyperparameters that balance the regularization. This joint optimization forces the model not only to be accurate but also to learn robust and consistent cross-modal representations, which are crucial for effective performance in modality-missing scenarios.

## Experiments

### Datasets and Evaluation Metrics

Following the experiment protocol adopted in (Zhao et al. 2024), we adapt two curated datasets established for the CMML task: UPMC-Food101-CMML and MM-IMDb-CMML. UPMC-Food101-CMML benchmark is derived from the UPMC-Food101 (Wang et al. 2015) dataset, comprising 101 food categories and 61,142 training, 6,846 validation, and 22,716 test image-text pairs. For the CL setup, these categories are partitioned into 10 sequential sessions. MM-IMDb-CMML benchmark is built upon the MM-IMDb (Arevalo et al. 2017) dataset. This dataset consists of 15,552 training, 2,608 validation, and 7,799 test image-text pairs across 27 movie genres. Adhering to the multi-label continual learning paradigm (Dong et al. 2023), the 7 most infrequent genres are removed, and the remaining categories are structured into 5 distinct sessions, each introducing 4 new genres. For both benchmarks, the missing severity is controlled by a ratio  $\eta\%$ , representing the modality-incomplete data within each session. In cases where both modalities can be absent, this portion is typically split evenly between image-only and text-only data.

For evaluation, we employ two standard metrics from the continual learning literature (Chaudhry et al. 2018a,b; Lopez-Paz and Ranzato 2017; Loshchilov and Hutter 2017): Average Performance (AP) and Average Forgetting (FG).

$\eta$	#Image	#Text	MAP		MSP		L2P		DualPrompt		RebQ		DeLo	
			AP ( $\uparrow$ )	FG ( $\downarrow$ )	AP ( $\uparrow$ )	FG ( $\downarrow$ )	AP ( $\uparrow$ )	FG ( $\downarrow$ )	AP ( $\uparrow$ )	FG ( $\downarrow$ )	AP ( $\uparrow$ )	FG ( $\downarrow$ )	AP ( $\uparrow$ )	FG ( $\downarrow$ )
10%	100%	90%	15.77	34.67	19.49	35.11	13.67	8.96	22.79	20.42	24.19	32.46	<b>29.14</b>	32.11
	90%	100%	16.63	33.30	19.01	34.89	13.54	8.26	24.82	24.09	27.35	28.88	<b>34.77</b>	30.36
	95%	95%	17.24	31.28	15.69	32.17	13.90	8.42	24.19	24.04	26.90	28.26	<b>31.84</b>	30.72
30%	100%	70%	19.85	32.14	20.01	34.87	10.68	8.03	22.96	23.96	24.75	25.86	<b>29.28</b>	32.79
	70%	100%	17.19	31.09	18.31	32.45	10.98	7.27	20.98	19.67	26.68	27.89	<b>34.56</b>	32.38
	85%	85%	14.44	33.29	14.20	35.57	11.58	7.90	22.18	20.72	25.80	26.95	<b>30.62</b>	29.23
50%	100%	50%	19.19	36.18	15.93	35.29	9.54	5.54	20.04	21.48	23.89	25.27	<b>28.88</b>	33.18
	50%	100%	18.83	32.02	16.78	34.14	9.77	4.88	19.97	20.24	26.45	27.53	<b>32.83</b>	31.09
	75%	75%	17.36	29.67	16.24	31.81	10.08	5.23	20.32	19.09	25.31	25.88	<b>30.15</b>	33.72
70%	100%	30%	17.45	27.14	18.01	27.32	7.26	5.37	16.80	13.82	22.41	24.39	<b>27.21</b>	33.39
	30%	100%	16.09	34.59	14.30	35.72	7.49	5.20	18.42	15.32	26.27	27.56	<b>31.03</b>	32.76
	65%	65%	17.20	29.39	18.27	29.34	6.93	6.23	16.19	15.76	23.77	25.26	<b>28.97</b>	31.84
90%	100%	10%	18.85	25.03	18.01	29.14	5.98	5.82	15.50	11.27	21.37	23.66	<b>26.11</b>	35.11
	10%	100%	17.79	31.65	17.88	32.19	6.29	5.09	21.22	18.38	26.22	27.50	<b>30.23</b>	33.85
	55%	55%	16.28	27.89	16.20	31.47	6.84	6.39	16.82	8.44	25.42	25.82	<b>28.59</b>	32.79

Table 3: **Performance comparison on the MM-IMDb-CMML dataset.** We report AP and FG based on F1-Macro. Missing Modality can occur in both training and testing. The best performance is in **bold**.

AP is the average performance across all  $T$  tasks after the model has completed training on the final task, calculated as  $AP = \frac{1}{T} \sum_{t=1}^T a_{t,T}$ , where  $a_{i,j}$  denotes the performance on task  $\mathcal{T}_i$  after training up to task  $\mathcal{T}_j$ . FG measures the performance degradation on past tasks and is calculated as:

$$FG = \frac{1}{T-1} \sum_{t=1}^{T-1} \max_{z \in \{t, \dots, T-1\}} (a_{t,z} - a_{t,T}). \quad (11)$$

**Comparison Methods.** To comprehensively evaluate our DeLo framework, we benchmark it against a strong set of baselines from two relevant research domains: missing modality and continual learning. For the missing modality problem, we include MAP (Lee et al. 2023) and MSP (Jang, Wang, and Kim 2024), which represent modern methods that adapt pre-trained LMMs for modality-incomplete data. From the continual learning perspective, we compare against leading prompt-based frameworks to contrast with our LoRA-based approach, including L2P (Wang et al. 2022b) and DualPrompt (Wang et al. 2022a). Importantly, we also conduct a direct comparison with RebQ (Zhao et al. 2024), the state-of-the-art method that first introduced the CMML task and its corresponding benchmark.

**Implementation Details.** Following (Lee et al. 2023), we utilize ViLT-B/32<sup>1</sup> (Kim, Son, and Kim 2021) as the pre-trained LMM backbone. Following the parameter-efficient paradigm, the parameters of the ViLT backbone remain frozen throughout all continual learning stages. A new, learnable fully-connected layer is instantiated as a dedicated classification head for each task. An empty string denotes the dummy text input (i.e.,  $\hat{\mathbf{t}}$ ), and an image with all pixel values equal one as the dummy image input (i.e.,  $\hat{\mathbf{v}}$ ). We set the expert pool size to  $E = 16$  and the selected rank to  $r = 4$ . The loss hyperparameters are  $\lambda_1 = 0.1$  and  $\lambda_2 = 0.1$ , and the EMA momentum is  $\beta = 0.99$ . We train all model

using the AdamW (Loshchilov and Hutter 2017) optimizer with a batch size of 4. The learning rate is initialized to  $1e-4$  and managed by a cosine annealing scheduler with a linear warmup phase over the first 10% of training steps. All experiments are conducted on one NVIDIA A100 GPU.

## Main Results

Table 2 and Table 3 summarizes the experimental results on both benchmarks, showing that our proposed DeLo framework consistently achieves the highest AP and maintains comparable or lower FG across all configurations. Unlike methods such as MAP and MSP which are not designed for continual learning, DeLo avoids the catastrophic forgetting that arises from overfitting adaptable parameters to each new task. Furthermore, DeLo improves upon advanced prompt-based continual learning methods. While approaches like L2P and DualPrompt utilize prompt pools, their performance can be hindered when missing modalities prevent effective pool access, our dynamic low-rank experts remain robust. Even RebQ, which extends this paradigm to CMML, still exhibits measurable forgetting in certain settings. The slightly higher FG of DeLo on the MM-IMDb-CMML can be attributed to its multi-label nature and richer semantic diversity, which make continual adaptation more challenging. We also observe that for all methods, performance correlates with the availability of text, likely due to the richer semantic information contained in textual descriptions.

## Ablation Study

In this section, we explore various DeLo variants to assess the importance of each component. Experiments are all based on the UPMC-Food101-CMML with a missing ratio  $\eta = 70\%$ , where both image and text modalities are missing.

**Effectiveness of DeLo Components** As presented in Table 4, we compare our full DeLo model against several variants, each with a crucial component removed or simplified.

<sup>1</sup><https://huggingface.co/dandelin/vilt-b32-mlm>

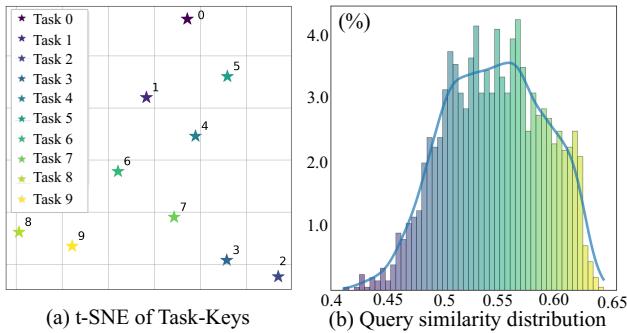


Figure 4: (a) t-SNE visualization of the final representations learned by the Task-Key Memory for each of the 10 continual tasks in UPMC-Food101-CMML. (b) The distribution of cosine similarity between visual and textual query signals, evaluated on modality-complete data.

The importance of our modality-decomposed architecture is evident when replacing the separate vision and text pools with a single unified pool (w/o Modality-Specific Factor Pool), which results in a significant performance drop due to modality interference. Performance collapses even more severely when removing the Cross-Modal Guided Routing mechanism (w/o Cross-Modal Guide), demonstrating that our strategy to preserve multimodal collaboration with incomplete data is critical. As illustrated in Figure 4, the Task-Key Memory forms well-separated clusters for different continual tasks, while the visual and textual query signals exhibit strong semantic alignment. Additionally, replacing our dynamic decomposed LoRA experts with conventional static LoRA (DeLo-LoRA) also degrades performance. These results collectively confirm that each component of DeLo provides an essential contribution to its overall effectiveness.

Method	AP ( $\uparrow$ )	FG ( $\downarrow$ )
DeLo ( $E = 16, r = 4$ )	<b>67.42</b>	<b>9.41</b>
DeLo-LoRA ( $r = 4$ )	60.39	12.08
w/o Cross-Modal Guide	48.84	20.13
w/o Modality-Specific Factor Pool	56.52	13.76

Table 4: Component-wise analysis of the DeLo framework.

**Impact of the Number of Rank and Factors** We analyze the impact of the total number of available expert factors ( $E$ ) and the selected rank ( $r$ ) to determine an optimal balance between performance and parameter efficiency. As shown in Table 5, our results reveal a clear trend: increasing the rank from 2 to 4 yields a substantial improvement in AP, but this gain diminishes or slightly reverses when the rank is further increased to 8. This suggests that  $r = 4$  provides a sweet spot for expressive capacity without the potential optimization challenges of a larger rank. Meanwhile, a larger factor pool ( $E = 32$ ) consistently offers a modest performance advantage over a smaller one ( $E = 16$ ) by providing the router with a more diverse set of experts. Given that a core goal of

Factors ( $E$ )	Rank ( $r$ )	AP ( $\uparrow$ )	FG ( $\downarrow$ )	Trainable Parameters	Total Parameters
16	2	63.78	11.20	7.1M	182.8M
16	4	67.42	9.41	12.4M	235.9M
16	8	67.11	10.89	23.0M	342.1M
32	2	64.53	13.37	14.2M	253.6M
32	4	68.03	8.92	24.8M	359.8M
32	8	67.92	9.24	46.0M	572.1M

Table 5: Impact of factor pool size and selected rank on DeLo’s performance and parameter efficiency.

our work is to maintain high parameter efficiency, we identify the setting of  $E = 16$  and  $r = 4$  as the optimal trade-off between predictive accuracy and computational cost.

**Effectiveness of three loss components.** Table 6 demonstrates that our regularization terms are crucial for the effectiveness of the Cross-Modal Guided Routing mechanism. When training with only the standard classification loss ( $\mathcal{L}_c$ ), the model achieves a baseline performance. Without explicit constraints, the semantic alignment of the proxy queries used for missing modalities is not guaranteed, which limits the guidance mechanism’s effectiveness. Introducing either the Alignment Loss ( $\mathcal{L}_{\text{align}}$ ) or the Consistency Loss ( $\mathcal{L}_{\text{con}}$ ) individually leads to a clear improvement in performance. The former enhances the proxy signal’s quality by explicitly enforcing semantic similarity between queries, while the latter reinforces a consistent reasoning process by aligning the model’s output distributions. Notably, the full model achieves the best results, suggesting that these two losses are complementary and work synergistically to enhance the reliability of our guided routing.

$\mathcal{L}_c$	$\mathcal{L}_{\text{align}}$	$\mathcal{L}_{\text{con}}$	AP ( $\uparrow$ )	FG ( $\downarrow$ )
✓	-	-	61.95	12.37
✓	-	✓	65.15	10.85
✓	✓	-	65.88	10.74
✓	✓	✓	<b>67.42</b>	<b>9.41</b>

Table 6: Impact of each loss component on performance.

## Conclusion

In this paper, we propose DeLo, a novel framework that introduces a dual-decomposed low-rank expert architecture to the challenging problem of CMML. DeLo dynamically constructs low-rank adaptations for each input by selecting factors from expert pools that are partitioned by both task and modality. This architecture, combined with a Cross-Modal Guided Routing, enables knowledge preservation while robustly handling modality-incomplete data. Extensive experiments on two established CMML benchmarks demonstrate that DeLo significantly outperforms existing state-of-the-art methods, achieving superior average performance while showing robust knowledge retention across previous tasks.

## References

- Aljundi, R.; Babiloni, F.; Elhoseiny, M.; Rohrbach, M.; and Tuytelaars, T. 2018. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European conference on computer vision (ECCV)*, 139–154.
- Aljundi, R.; Belilovsky, E.; Tuytelaars, T.; Charlin, L.; Caccia, M.; Lin, M.; and Page-Caccia, L. 2019. Online continual learning with maximal interfered retrieval. *Advances in neural information processing systems*, 32.
- Arevalo, J.; Solorio, T.; Montes-y Gómez, M.; and González, F. A. 2017. Gated multimodal units for information fusion. *arXiv preprint arXiv:1702.01992*.
- Boschini, M.; Bonicelli, L.; Porrello, A.; Bellitto, G.; Pennisi, M.; Palazzo, S.; Spampinato, C.; and Calderara, S. 2022. Transfer without forgetting. In *European conference on computer vision*, 692–709. Springer.
- Buzzega, P.; Boschini, M.; Porrello, A.; Abati, D.; and Calderara, S. 2020. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33: 15920–15930.
- Cai, L.; Wang, Z.; Gao, H.; Shen, D.; and Ji, S. 2018. Deep adversarial learning for multi-modality missing data completion. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 1158–1166.
- Cha, H.; Lee, J.; and Shin, J. 2021. Co2l: Contrastive continual learning. In *Proceedings of the IEEE/CVF International conference on computer vision*, 9516–9525.
- Chaudhry, A.; Dokania, P. K.; Ajanthan, T.; and Torr, P. H. 2018a. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European conference on computer vision (ECCV)*, 532–547.
- Chaudhry, A.; Ranzato, M.; Rohrbach, M.; and Elhoseiny, M. 2018b. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*.
- Chen, H.; Li, J.; Gazagnadou, N.; Zhuang, W.; Chen, C.; and Lyu, L. 2024. Dual low-rank adaptation for continual learning with pre-trained models. *arXiv preprint arXiv:2411.00623*.
- Chrysakakis, A.; and Moens, M.-F. 2020. Online continual learning from imbalanced data. In *International Conference on Machine Learning*, 1952–1961. PMLR.
- Dong, S.; Luo, H.; He, Y.; Wei, X.; Cheng, J.; and Gong, Y. 2023. Knowledge restore and transfer for multi-label class-incremental learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 18711–18720.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Ebrahimi, S.; Meier, F.; Calandra, R.; Darrell, T.; and Rohrbach, M. 2020. Adversarial continual learning. In *European conference on computer vision*, 386–402. Springer.
- Guo, Z.; Wang, S.; Lin, W.; Yan, W.; Wu, Y.; and Jin, T. 2025. Efficient prompting for continual adaptation to missing modalities. *arXiv preprint arXiv:2503.00528*.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16000–16009.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Hung, C.-Y.; Tu, C.-H.; Wu, C.-E.; Chen, C.-H.; Chan, Y.-M.; and Chen, C.-S. 2019. Compacting, picking and growing for unforgetting continual learning. *Advances in neural information processing systems*, 32.
- Jang, J.; Wang, Y.; and Kim, C. 2024. Towards robust multimodal prompting with missing modalities. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 8070–8074. IEEE.
- Jin, T.; Cheng, X.; Li, L.; Lin, W.; Wang, Y.; and Zhao, Z. 2023. Rethinking missing modality learning from a decoding perspective. In *Proceedings of the 31st ACM international conference on multimedia*, 4431–4439.
- Jung, S.; Ahn, H.; Cha, S.; and Moon, T. 2020. Continual learning with node-importance based adaptive group sparse regularization. *Advances in neural information processing systems*, 33: 3647–3658.
- Ke, G.; He, S.; Wang, X.; Wang, B.; Chao, G.; Zhang, Y.; Xie, Y.; and Su, H. 2025. Knowledge Bridger: Towards Training-Free Missing Modality Completion. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 25864–25873.
- Khan, M. G. Z. A.; Naeem, M. F.; Van Gool, L.; Stricker, D.; Tombari, F.; and Afzal, M. Z. 2023. Introducing language guidance in prompt-based continual learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, 11463–11473.
- Kim, W.; Son, B.; and Kim, I. 2021. Vilt: Vision-and-language transformer without convolution or region supervision. In *International conference on machine learning*, 5583–5594. PMLR.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13): 3521–3526.
- Lee, Y.-L.; Tsai, Y.-H.; Chiu, W.-C.; and Lee, C.-Y. 2023. Multimodal prompting with missing modalities for visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14943–14952.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. BliP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, 19730–19742. PMLR.
- Li, X.; Zhou, Y.; Wu, T.; Socher, R.; and Xiong, C. 2019. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *International conference on machine learning*, 3925–3934. PMLR.

- Liang, Y.-S.; and Li, W.-J. 2023a. Adaptive plasticity improvement for continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7816–7825.
- Liang, Y.-S.; and Li, W.-J. 2023b. Loss decoupling for task-agnostic continual learning. *Advances in Neural Information Processing Systems*, 36: 11151–11167.
- Liang, Y.-S.; and Li, W.-J. 2024. Inflora: Interference-free low-rank adaptation for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23638–23647.
- Lopez-Paz, D.; and Ranzato, M. 2017. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Ma, M.; Ren, J.; Zhao, L.; Tulyakov, S.; Wu, C.; and Peng, X. 2021. Smil: Multimodal learning with severely missing modality. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 2302–2310.
- McCloskey, M.; and Cohen, N. J. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, 109–165. Elsevier.
- Pham, H.; Liang, P. P.; Manzini, T.; Morency, L.-P.; and Póczos, B. 2019. Found in translation: Learning robust joint representations by cyclic translations between modalities. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 6892–6899.
- Qian, Z.; Wang, X.; Duan, X.; Qin, P.; Li, Y.; and Zhu, W. 2023. Decouple before interact: Multi-modal prompt learning for continual visual question answering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2953–2962.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Rolnick, D.; Ahuja, A.; Schwarz, J.; Lillicrap, T.; and Wayne, G. 2019. Experience replay for continual learning. *Advances in neural information processing systems*, 32.
- Serra, J.; Suris, D.; Miron, M.; and Karatzoglou, A. 2018. Overcoming catastrophic forgetting with hard attention to the task. In *International conference on machine learning*, 4548–4557. PMLR.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Shen, Y.; Xu, Z.; Wang, Q.; Cheng, Y.; Yin, W.; and Huang, L. 2024. Multimodal instruction tuning with conditional mixture of lora. *arXiv preprint arXiv:2402.15896*.
- Smith, J. S.; Karlinsky, L.; Gutta, V.; Cascante-Bonilla, P.; Kim, D.; Arbelle, A.; Panda, R.; Feris, R.; and Kira, Z. 2023. Coda-prompt: Continual decomposed attention-based prompting for rehearsal-free continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11909–11919.
- Wang, X.; Kumar, D.; Thome, N.; Cord, M.; and Precioso, F. 2015. Recipe recognition with large multimodal food dataset. In *2015 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, 1–6. IEEE.
- Wang, Y.; Huang, Z.; and Hong, X. 2022. S-prompts learning with pre-trained transformers: An occam’s razor for domain incremental learning. *Advances in Neural Information Processing Systems*, 35: 5682–5695.
- Wang, Z.; Zhang, Z.; Ebrahimi, S.; Sun, R.; Zhang, H.; Lee, C.-Y.; Ren, X.; Su, G.; Perot, V.; Dy, J.; et al. 2022a. Dual-prompt: Complementary prompting for rehearsal-free continual learning. In *European conference on computer vision*, 631–648. Springer.
- Wang, Z.; Zhang, Z.; Lee, C.-Y.; Zhang, H.; Sun, R.; Ren, X.; Su, G.; Perot, V.; Dy, J.; and Pfister, T. 2022b. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 139–149.
- Wistuba, M.; Sivaprasad, P. T.; Balles, L.; and Zappella, G. 2023. Continual learning with low rank adaptation. *arXiv preprint arXiv:2311.17601*.
- Wistuba, M.; Sivaprasad, P. T.; Balles, L.; and Zappella, G. 2024. Choice of peft technique in continual learning: Prompt tuning is not all you need. *arXiv preprint arXiv:2406.03216*.
- Woo, S.; Lee, S.; Park, Y.; Nugroho, M. A.; and Kim, C. 2023. Towards good practices for missing modality robust action recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 2776–2784.
- Yue, X.; Chen, Y.; Zhang, X.; Gao, X.; Feng, M.; Lao, M.; Zhuang, H.; and Li, H. 2025. PAL: Prompting Analytic Learning with Missing Modality for Multi-Modal Class-Incremental Learning. *arXiv preprint arXiv:2501.09352*.
- Zenke, F.; Poole, B.; and Ganguli, S. 2017. Continual learning through synaptic intelligence. In *International conference on machine learning*, 3987–3995. PMLR.
- Zhang, Y.; Peng, C.; Wang, Q.; Song, D.; Li, K.; and Zhou, S. K. 2024. Unified multi-modal image synthesis for missing modality imputation. *IEEE Transactions on Medical Imaging*, 44(1): 4–18.
- Zhao, J.; Li, R.; and Jin, Q. 2021. Missing modality imagination network for emotion recognition with uncertain missing modalities. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2608–2618.
- Zhao, S.; Zou, X.; Yu, T.; and Xu, H. 2024. Reconstruct before query: Continual missing modality learning with decomposed prompt collaboration. *arXiv preprint arXiv:2403.11373*.
- Zheng, Z.; Ma, M.; Wang, K.; Qin, Z.; Yue, X.; and You, Y. 2023. Preventing zero-shot transfer degradation in continual learning of vision-language models. In *Proceedings of the IEEE/CVF international conference on computer vision*, 19125–19136.