

MacVQA: Adaptive Memory Allocation and Global Noise Filtering for Continual Visual Question Answering

Zhifei Li^{1,4,5}, Yiran Wang¹, Chenyi Xiong¹, Yujing Xia¹, Xiaoju Hou^{2,*},
Yue Zhao³, Miao Zhang^{1,*}, Kui Xiao^{1,*}, Bing Yang¹

¹School of Computer Science, Hubei University, Wuhan 430062, China

²Institute of Vocational Education, Guangdong Industry Polytechnic University, Guangzhou 510300, China

³Shandong Police College, Ji'nan 250200, China

⁴Hubei Key Laboratory of Big Data Intelligent Analysis and Application (Hubei University), Wuhan 430062, China

⁵Key Laboratory of Intelligent Sensing System and Security (Hubei University), Ministry of Education, Wuhan 430062, China

{zhifei1993, zhangmiao, xiaokui}@hubu.edu.cn, 2023030010@gdip.edu.cn,

zhaoy@sdpc.edu.cn, {yiranwang, xiongchenyi, xiayujing}@stu.hubu.edu.cn, yangbing@126.com

Abstract

Visual Question Answering (VQA) requires models to reason over multimodal information, combining visual and textual data. With the development of continual learning, significant progress has been made in retaining knowledge and adapting to new information in the VQA domain. However, current methods often struggle with balancing knowledge retention, adaptation, and robust feature representation. To address these challenges, we propose a novel framework with adaptive memory allocation and global noise filtering called MacVQA for visual question answering. MacVQA fuses visual and question information while filtering noise to ensure robust representations, and employs prototype-based memory allocation to optimize feature quality and memory usage. These designs enable MacVQA to balance knowledge acquisition, retention, and compositional generalization in continual VQA learning. Experiments on ten continual VQA tasks show that MacVQA outperforms existing baselines, achieving 43.38% average accuracy and 2.32% average forgetting on standard tasks, and 42.53% average accuracy and 3.60% average forgetting on novel composition tasks.

Code — <https://github.com/HubuKG/MacVQA>

Introduction

Visual Question Answering (VQA) integrates visual and textual information to answer questions about images. As a representative task in multimodal learning, VQA has garnered significant attention due to its wide range of real-world applications, including autonomous driving (Lv et al. 2023), medical diagnosis (Bai, Islam, and Ren 2023), accessibility tools for visually impaired individuals (Huh et al. 2024), and intelligent human-computer interaction systems (Wu et al. 2022). For instance, a VQA system can assist medical professionals in interpreting complex diagnostic imaging by answering questions about abnormalities or help autonomous vehicles analyze scenes through visual queries. These applications highlight the importance of VQA as a bridge between visual perception and natural question understanding.

*Corresponding Authors.

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

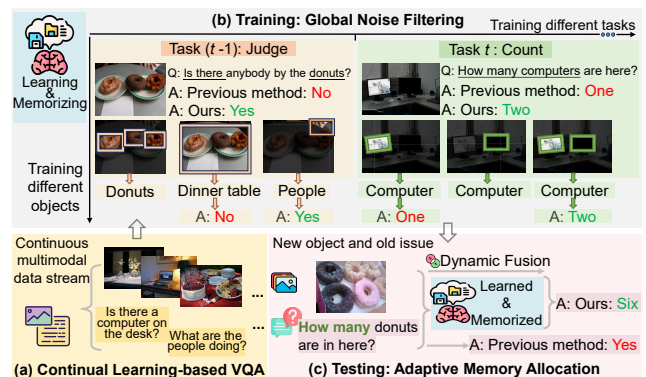


Figure 1: Illustration with continual learning methods for VQA. (a) Continual Learning-based VQA: The system learns new tasks while retaining prior knowledge. (b) Global Noise Filtering: Refines visual features by removing noise for better task performance. (c) Adaptive Memory Allocation: Combines learned and memorized knowledge for question answering.

Continual learning (CL) has emerged as a promising paradigm in the VQA domain, offering solutions to the challenges posed by evolving environments and dynamic tasks (Sun et al. 2023). By enabling models to learn and adapt to new tasks without forgetting previously acquired knowledge (Greco et al. 2019; Xu et al. 2020), CL addresses the inherent requirements of VQA systems for long-term adaptability and scalability (Kim et al. 2024; Tiwari et al. 2022). Its value is particularly evident in scenarios where VQA systems need to handle diverse datasets, adapt to new visual concepts, and retain previously learned reasoning capabilities (Hayes et al. 2020; Lei et al. 2023). The integration of CL into VQA not only enhances the robustness and flexibility of these systems but also paves the way for more practical and efficient applications in real-world settings.

Existing methods primarily adopt regularization-based or rehearsal-based methods (Yang et al. 2025; Singh et al. 2019; Radford et al. 2021). Regularization techniques pre-

serve prior knowledge by protecting critical parameters, with Memory Aware Synapses (MAS) (Aljundi et al. 2018) maintaining important weights and Neuronal Plasticity Control (NPC) (Paik et al. 2020) regulating synaptic adaptability. Complementary rehearsal strategies reinforce learning through stored task memories, exemplified by Experience Replay (ER) (Chaudhry et al. 2019) and its enhanced variant Dark Experience Replay (DER) (Chaudhry et al. 2018). Recent VQA-specific innovations include: VQACL (Zhang, Zhang, and Xu 2023) evaluating unseen skill-concept combinations; PROOF (Zhou et al. 2025) employing expandable projections; QUAD (Marouf et al. 2025) eliminating visual storage via attention consistency; and ProtoGroup (Zhang et al. 2025) stabilizing representations through multi-prototype grouping with dual-score memory selection.

Despite these advancements, existing methods face significant limitations in multimodal continual learning. Regularization struggles with fine-grained interactions, while rehearsal depends heavily on memory replay, limiting scalability (Naik and Nushi 2023; Li et al. 2019; Krishna, Bernstein, and Fei-Fei 2019; Smith et al. 2021). VQACL and PROOF adapt poorly to dynamic tasks; QUAD is less robust to noisy features; ProtoGroup incurs computational overhead and is sensitive to clustering parameters. Such constraints hinder balancing retention, adaptation, and robust reasoning.

To address these challenges, we propose MacVQA, an innovative framework designed to enhance multimodal feature representation and optimize memory allocation for continual visual question answering, as illustrated in Figure 1. MacVQA integrates two key components: Global Noise Filtering (GonF) and Adaptive Memory Allocation (AMA). GonF improves multimodal feature quality by removing irrelevant noise and fusing global visual and linguistic information, ensuring robust feature representations for complex reasoning tasks. AMA employs adaptive memory strategies, including prototype retrieval and memory decay, to efficiently update knowledge, mitigate forgetting, and maintain scalability. Together, these components enable MacVQA to balance multimodal information adaptively while effectively integrating new knowledge and preserving prior learning.

In this paper, we have made the following contributions:

- We propose Global Noise Filtering, which combines global feature fusion and noise removal to enhance multimodal feature robustness and reasoning quality significantly.
- We design Adaptive Memory Allocation, which uses prototype retrieval, memory decay, and adaptive updates to optimize knowledge allocation in dynamic continual learning environments.
- We conduct experiments on ten tasks in VQA v2 dataset (Goyal et al. 2017). MacVQA achieves accuracies of 43.38% and 42.53% on standard and novel task combinations, while reducing forgetting to 2.32% and 3.60%.

Related Work

Visual Question Answering

Visual Question Answering (VQA) is a multimodal task that combines visual and textual inputs to answer questions about images. Early methods established visual-textual alignment through attention mechanisms (Theeuwes 2010; Kim, Jun, and Zhang 2018), while transformer architectures (Tan and Bansal 2019; Chen et al. 2020; Gómez-Pérez and Ortega 2020; Liang et al. 2020) advanced cross-modal fusion via self-attention. Recent generative paradigms unify diverse tasks under a text generation framework, eliminating specialized modules while enabling flexible, open-ended responses. This evolution reflects progressive architectural simplification from component-specific designs to unified processing.

However, as VQA applications expand to more complex and dynamic scenarios, traditional methods face challenges in handling diverse and evolving tasks (Hudson and Manning 2019). Multimodal data in real-world environments often exhibit high variability, requiring models to establish accurate connections between rich visual content and diverse linguistic expressions (Liang et al. 2020; Agrawal et al. 2018). For instance, many datasets highlight the difficulties posed by incomplete or noisy images (Zhu et al. 2016), while open-ended tasks reveal limitations in predefined structures and fixed setups. Addressing these challenges requires models capable of adapting to dynamic data distributions and generalizing to unseen tasks (Whitehead et al. 2021; Zhang et al. 2022), marking a crucial direction for future research in VQA.

Continual Learning in VQA

Continual learning enables sequential knowledge acquisition while preserving prior representations. Current VQA continual learning methods primarily employ two paradigms. Regularization-based methods constrain parameter updates to protect critical knowledge, exemplified by MAS (Aljundi et al. 2018) computing parameter importance via Fisher-based metrics, and NPC (Paik et al. 2020) modulating layer-specific learning rates through activation stability. However, these methods struggle with modality conflicts in multimodal architectures. Rehearsal-based methods mitigate forgetting through stored exemplars, ER (Chaudhry et al. 2019) replays random past data (high storage), while DER (Chaudhry et al. 2018) augments ER via logit distillation (privacy-constrained).

Recent frameworks tailored for continual learning in VQA address these challenges. VQACL (Zhang, Zhang, and Xu 2023) introduces a dual-level task sequence and compositional testing protocol; QUAD (Marouf et al. 2025) employs question-only replay and attention consistency distillation; while ProtoGroup (Zhang et al. 2025) leverages multi-prototype grouping for stable representations and dual-score memory selection. Nonetheless, challenges such as limited feature robustness, catastrophic forgetting, and computational overhead persist in dynamic environments.

To fundamentally address these issues, our method enhances representation robustness and strengthens knowledge

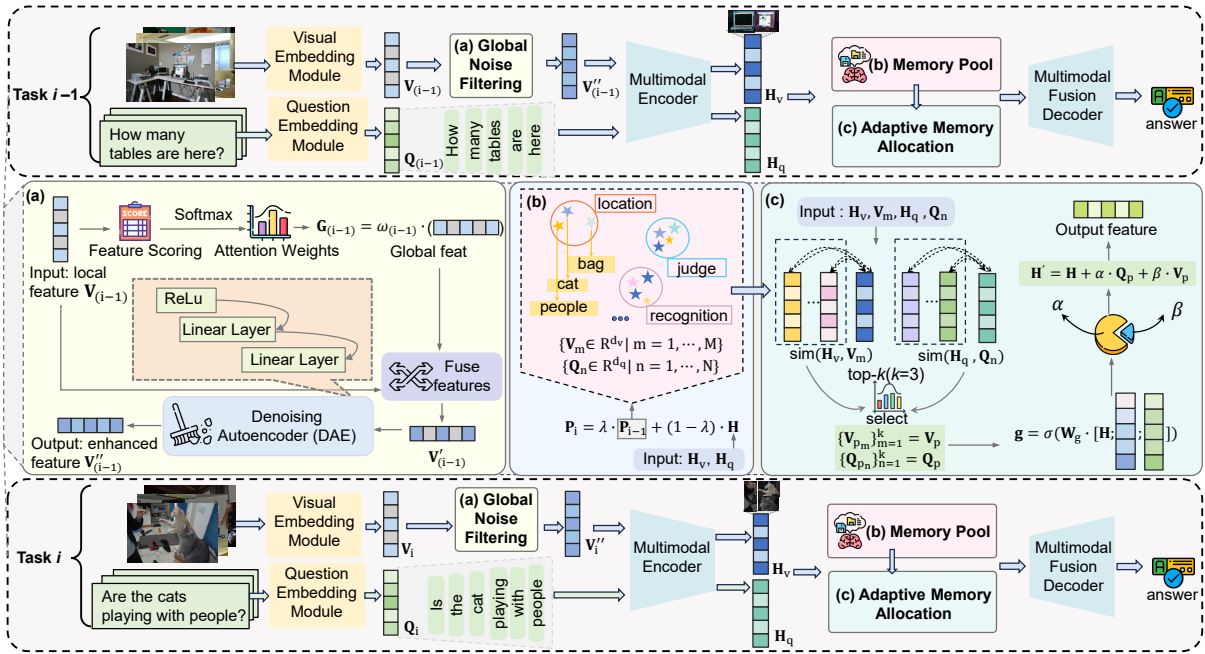


Figure 2: The proposed framework consists of three key modules: (a) Global Noise Filtering, which refines visual features by removing noise through scoring, fusion, and a Denoising Autoencoder (DAE); (b) Memory Pool, storing visual and textual prototypes as dynamic references; (c) Adaptive Memory Allocation, which dynamically allocates prototype features from the memory pool to enhance adaptability and generalization.

retention across the learning process. It enables stable task adaptation, preserves long-term knowledge, and supports generalization, offering a comprehensive solution for continual learning in dynamic visual question scenarios.

Methodology

We propose a novel multimodal continual learning framework (Figure 2) for visual question answering, addressing feature representation, memory, and knowledge update challenges. By incorporating Global Noise Filtering and Adaptive Memory Allocation modules, our framework enhances feature robustness, memory efficiency, and task adaptation. Module details are described below.

Global Noise Filtering

The Global Noise Filtering (GonF) module addresses the problem of noisy and irrelevant region features in visual representation learning. By globally highlighting important information and filtering out noise, GonF improves the robustness and quality of visual features for visual-question tasks.

For each outer task (e.g., Count), the same process is applied to all its image-question pairs (inner tasks). Here, we take the $(i-1)$ -th outer task as an example to illustrate the feature extraction and aggregation process.

Visual feature extraction employs a pre-trained Faster R-CNN (Ren et al. 2015) model to extract n region features from each image, represented as $\mathbf{V}_{(i-1),j} \in \mathbb{R}^{n \times d}$, where n denotes the number of detected image regions, d represents the feature embedding dimension (typically 2048),

$(i-1)$ indicates the index of the previous (outer) task, and j corresponds to the attention head index in multi-head architectures. Textual features are obtained by encoding input queries into embeddings $\mathbf{Q}_{(i-1),j} \in \mathbb{R}^{L \times d}$, with L signifying the number of words in the query.

The noise filtering process begins with feature scoring, which assigns attention weights $\omega_{(i-1),j,m}$ to each image region through a linear transformation and softmax normalization:

$$\omega_{(i-1),j,m} = \frac{\exp(\text{score}(\mathbf{V}_{(i-1),j,m}))}{\sum_{l=1}^n \exp(\text{score}(\mathbf{V}_{(i-1),j,l}))}, \quad (1)$$

where $\omega_{(i-1),j,m}$ quantifies the attention weight for the m -th region in the $(i-1)$ -th task and head j . These weights are then used to generate a global feature vector $\mathbf{G}_{(i-1),j} \in \mathbb{R}^d$ via weighted summation:

$$\mathbf{G}_{(i-1),j} = \sum_{m=1}^n \omega_{(i-1),j,m} \cdot \mathbf{V}_{(i-1),j,m}, \quad (2)$$

where $\mathbf{G}_{(i-1),j}$ captures contextual information across all image regions for the $(i-1)$ -th task.

Subsequent denoising utilizes a Denoising Autoencoder (DAE) to refine $\mathbf{V}_{(i-1),j}$ into denoised features $\mathbf{V}'_{(i-1),j} \in \mathbb{R}^{n \times d}$. To further enhance the representation, the global feature $\mathbf{G}_{(i-1),j}$ and the denoised region features $\mathbf{V}'_{(i-1),j}$ are combined into the final enhanced features $\mathbf{V}''_{(i-1),j}$.

For optimization, GonF employs a unified loss function that simultaneously encourages accurate feature reconstruction and smooth attention distribution. Specifically, the loss

is defined as the squared reconstruction error penalized by an entropy-based regularization term on the attention weights:

$$\mathcal{L}_{\text{GonF}} = \frac{1}{n} \sum_{m=1}^n \left(\|\mathbf{V}_{(i-1),j,m} - \mathbf{V}'_{(i-1),j,m}\|_2^2 - \theta_1 \omega_{(i-1),j,m} \log \omega_{(i-1),j,m} \right), \quad (3)$$

where θ_1 is a hyperparameter (typically within $[0.1, 0.5]$) that balances the trade-off between denoising precision and attention distribution smoothness.

The enhanced visual features $\mathbf{V}''_{(i-1),j}$, together with textual features $\mathbf{Q}_{(i-1),j}$, are subsequently processed by the Multimodal Encoder, which leverages a Transformer backbone to fuse modalities into the final multimodal representation $\mathbf{F}_{(i-1),j} \in \mathbb{R}^{M \times d}$, where M is the total number of fused tokens.

Adaptive Memory Allocation

Conventional memory or prototype-based mechanisms often struggle to cope with the dynamic and heterogeneous nature of continual multimodal tasks. To overcome this limitation, we propose the Adaptive Memory Allocation (AMA) module, which enables the model to selectively leverage and update relevant prior knowledge. This design enhances generalization and robustness across a wide range of visual-question scenarios.

The input to AMA is a hidden state feature $\mathbf{H} \in \mathbb{R}^d$, which encodes the global context of a given sample. This feature is obtained from the multimodal encoder by fusing visual and textual information. To facilitate fine-grained prototype matching, \mathbf{H} is projected into two modality-specific subspaces:

$$\mathbf{H}_v = \mathbf{W}_v \cdot \mathbf{H}, \quad \mathbf{H}_q = \mathbf{W}_q \cdot \mathbf{H}, \quad (4)$$

where $\mathbf{W}_v \in \mathbb{R}^{d_v \times d}$ and $\mathbf{W}_q \in \mathbb{R}^{d_q \times d}$ are learnable transformation matrices. Here, $\mathbf{H}_v \in \mathbb{R}^{d_v}$ and $\mathbf{H}_q \in \mathbb{R}^{d_q}$ represent the visual- and textual-specific features, respectively.

AMA maintains a memory pool consisting of two distinct prototype sets: visual prototypes $\{\mathbf{V}_m \in \mathbb{R}^{d_v} \mid m = 1, \dots, M\}$ and textual prototypes $\{\mathbf{Q}_n \in \mathbb{R}^{d_q} \mid n = 1, \dots, N\}$. Each incoming sample compares its modality-specific features to stored prototypes using cosine similarity:

$$\begin{cases} \text{sim}(\mathbf{H}_v, \mathbf{V}_m) = \frac{\mathbf{H}_v \cdot \mathbf{V}_m}{\|\mathbf{H}_v\| \cdot \|\mathbf{V}_m\|}, \\ \text{sim}(\mathbf{H}_q, \mathbf{Q}_n) = \frac{\mathbf{H}_q \cdot \mathbf{Q}_n}{\|\mathbf{H}_q\| \cdot \|\mathbf{Q}_n\|}, \end{cases} \quad (5)$$

where \cdot denotes dot product and $\|\cdot\|$ is the L_2 norm. The top- k most similar visual and textual prototypes, denoted as $\{\mathbf{V}_{p_m}\}_{m=1}^k$ and $\{\mathbf{Q}_{p_n}\}_{n=1}^k$, are retrieved for further fusion.

To adaptively integrate retrieved prototypes with the current input, AMA employs a gating-based fusion mechanism. A dynamic weight vector $\mathbf{g} \in \mathbb{R}^{d_g}$ is computed as:

$$\mathbf{g} = \sigma(\mathbf{W}_g \cdot [\mathbf{H}; \mathbf{Q}_p; \mathbf{V}_p]), \quad (6)$$

where $\mathbf{W}_g \in \mathbb{R}^{d_g \times (d+d_q+d_v)}$ is a learnable matrix, σ is a non-linear activation (e.g., sigmoid), and $[\cdot; \cdot; \cdot]$ denotes vector concatenation. This gating mechanism allows the model

to selectively emphasize useful modalities and ignore less relevant signals.

To ensure that knowledge is accumulated effectively across tasks, we further update the memory pool using a temporal interpolation strategy:

$$\mathbf{P}_i = \lambda \cdot \mathbf{P}_{i-1} + (1 - \lambda) \cdot \mathbf{H}, \quad (7)$$

where \mathbf{P}_i denotes the updated memory state after processing task i , and $\lambda \in [0, 1]$ controls the balance between past memory \mathbf{P}_{i-1} and current context \mathbf{H} . This incremental update rule ensures both long-term memory retention and adaptability to new information, making the memory pool more resilient to catastrophic forgetting.

The final fused feature $\mathbf{H}' \in \mathbb{R}^d$ is computed as:

$$\mathbf{H}' = \mathbf{H} + \alpha \cdot \mathbf{Q}_p + \beta \cdot \mathbf{V}_p, \quad (8)$$

where $\alpha \in \mathbb{R}^{d_q}$ and $\beta \in \mathbb{R}^{d_v}$ are dynamically adjusted weights for the textual and visual prototypes. These weights are defined as:

$$\begin{cases} \alpha = g_v \cdot \text{softmax}(\mathbf{W}_\alpha \cdot \mathbf{V}_p) \\ \beta = g_q \cdot \text{softmax}(\mathbf{W}_\beta \cdot \mathbf{Q}_p) \end{cases}, \quad (9)$$

where g_v and g_q are portions of \mathbf{g} allocated to the visual and textual modalities, respectively. $\mathbf{W}_\alpha \in \mathbb{R}^{d_q \times d_v}$ and $\mathbf{W}_\beta \in \mathbb{R}^{d_v \times d_q}$ are learnable weight matrices, and softmax ensures that the contributions of all prototypes are normalized to sum to 1.

The overall loss function for the AMA module is defined as:

$$\begin{aligned} \mathcal{L}_{\text{AMA}} = & - \sum_{m=1}^k \text{sim}(\mathbf{H}_v, \mathbf{V}_{p_m}) - \sum_{n=1}^k \text{sim}(\mathbf{H}_q, \mathbf{Q}_{p_n}) \\ & + \theta_2 (g_q + g_v - 1)^2 + \theta_3 \|\mathbf{H}' - \mathbf{H}\|_2^2, \end{aligned} \quad (10)$$

where θ_2 and θ_3 are hyperparameters used to balance the contributions of the regularization and update terms. In practice, θ_2 and θ_3 are typically selected from the range $[0.01, 1.0]$ based on validation performance.

Multimodal Decoding for Answers

The multimodal decoding module generates task predictions based on fused visual and textual features. Given an input representation $\mathbf{H}' \in \mathbb{R}^d$, the model first applies two parallel linear projections:

$$\mathbf{E}_v = \mathbf{W}_v \mathbf{H}', \quad \mathbf{E}_q = \mathbf{W}_q \mathbf{H}', \quad (11)$$

where $\mathbf{W}_v, \mathbf{W}_q \in \mathbb{R}^{d_e \times d}$ are learnable matrices, and $\mathbf{E}_v, \mathbf{E}_q \in \mathbb{R}^{d_e}$ are the corresponding visual and question embeddings.

These embeddings are concatenated to form a unified decoder input:

$$\mathbf{E} = [\mathbf{E}_v; \mathbf{E}_q], \quad (12)$$

where $\mathbf{E} \in \mathbb{R}^{2d_e}$ combines modality-specific cues into a single feature for downstream VQA prediction.

The decoder receives the fused multimodal embeddings \mathbf{E} and generates task-specific outputs, such as answers to

| Methods | Various tasks in VQA v2 | | | | | | | | | | AP (↑) | AF (↓) |
|-------------------------|-------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|-------------|
| | Rec. | Loc. | Jud. | Com. | Cou. | Act. | Col. | Typ. | Sub. | Cau. | | |
| <i>Standard Test</i> | | | | | | | | | | | | |
| Vanilla | 7.39 | 4.94 | 22.29 | 32.30 | 0.71 | 12.14 | 12.10 | 10.69 | 27.29 | 15.10 | 14.49 | 30.15 |
| EWC (PNAS'17) | 6.73 | 8.43 | 27.22 | 47.10 | 0.14 | 12.40 | 1.76 | 10.98 | 31.05 | 11.85 | 15.77 | 28.38 |
| MAS (ECCV'18) | 30.81 | 8.07 | 25.50 | 4.00 | 31.90 | 32.39 | 26.24 | 24.75 | 19.85 | 2.75 | 20.56 | 21.97 |
| ER (MTLR'19) | 18.64 | 21.36 | 61.27 | 64.17 | 30.29 | 52.84 | 43.39 | 23.31 | 42.75 | 11.85 | 36.99 | 4.80 |
| DER (NeurIPS'20) | 14.55 | 13.83 | 62.88 | 65.16 | 30.96 | 51.19 | 40.51 | 19.04 | 42.87 | 12.55 | 35.35 | 6.58 |
| VQACL (CVPR'23) | 20.47 | 28.02 | 62.55 | 68.61 | 29.35 | 50.66 | 44.45 | 26.36 | 44.65 | 12.60 | 38.77 | 2.90 |
| QUAD (arXiv'25) | 20.51 | 28.04 | 62.62 | 68.65 | 29.39 | 50.67 | 44.48 | 26.38 | 44.72 | 12.64 | 39.25 | 3.91 |
| ProtoGroup (ICASSP'25) | 20.66 | 28.77 | 62.55 | 69.21 | 29.44 | 50.79 | 44.66 | 26.44 | 44.86 | 12.70 | 39.81 | 2.87 |
| MacVQA (Ours) | 26.19 | 28.96 | 66.64 | 66.22 | 32.30 | 58.13 | 56.95 | 35.90 | 50.78 | 15.95 | 43.38 | 2.32 |
| <i>Novel Comp. Test</i> | | | | | | | | | | | | |
| Vanilla | 7.23 | 4.55 | 21.56 | 31.98 | 0.66 | 12.08 | 11.88 | 8.99 | 27.12 | 14.70 | 11.79 | 27.16 |
| EWC (PNAS'17) | 6.44 | 8.14 | 26.93 | 46.81 | 0.05 | 12.02 | 1.67 | 10.60 | 30.76 | 11.56 | 12.83 | 28.16 |
| MAS (ECCV'18) | 31.06 | 8.32 | 25.75 | 4.25 | 32.15 | 32.64 | 26.49 | 25.00 | 20.10 | 3.00 | 23.90 | 6.24 |
| ER (MTLR'19) | 18.32 | 21.04 | 60.79 | 64.01 | 29.97 | 52.52 | 43.07 | 22.99 | 42.43 | 11.53 | 33.78 | 5.76 |
| DER (NeurIPS'20) | 14.17 | 13.45 | 62.50 | 64.78 | 30.58 | 50.81 | 39.94 | 18.85 | 42.49 | 12.17 | 31.52 | 8.59 |
| VQACL (CVPR'23) | 20.13 | 27.68 | 62.38 | 68.10 | 29.01 | 50.32 | 44.11 | 26.02 | 44.31 | 12.26 | 35.40 | 4.90 |
| QUAD (arXiv'25) | 20.59 | 28.08 | 63.01 | 68.73 | 29.47 | 50.75 | 44.56 | 26.46 | 44.80 | 12.72 | 40.00 | 3.81 |
| ProtoGroup (ICASSP'25) | 20.33 | 27.89 | 62.22 | 68.64 | 29.23 | 50.45 | 44.27 | 26.14 | 44.42 | 12.34 | 36.81 | 4.09 |
| MacVQA (Ours) | 24.35 | 28.09 | 67.05 | 69.65 | 30.06 | 48.41 | 54.05 | 29.43 | 47.21 | 17.00 | 42.53 | 3.60 |

Table 1: Comparison of various tasks in VQA v2 across Standard Test and Novel Compositional Test.

| Method | | Various task in VQA v2 | | | | | | | | | | AP (↑) | AF (↓) |
|--------|-----|------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|---------------------------------|--------------------------------|
| GonF | AMA | Rec. | Loc. | Jud. | Com. | Cou. | Act. | Col. | Typ. | Sub. | Cau. | | |
| × | × | 20.47 | 28.02 | 62.55 | 68.61 | 29.35 | 50.66 | 44.45 | 26.36 | 44.65 | 12.60 | 38.77 | 2.90 |
| × | × | 23.65 | 36.08 | 66.09 | 65.90 | 32.24 | 47.68 | 55.37 | 32.09 | 44.76 | 14.30 | 41.75 | 2.14 |
| × | ✓ | 22.54 | 35.05 | 65.79 | 65.90 | 32.15 | 45.30 | 55.97 | 31.96 | 44.05 | 13.78 | 40.97 | 2.34 |
| ✓ | ✓ | 26.19 | 28.96 | 66.64 | 66.22 | 32.30 | 58.13 | 56.95 | 35.90 | 50.78 | 15.95 | 43.38 _(+3.57) | 2.32 _(-0.58) |

Table 2: Ablation study of GonF and AMA modules on VQA performance.

visual questions. It integrates visual and textual information through the cross-attention mechanism, expressed as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}}\right)\mathbf{V}, \quad (13)$$

where \mathbf{Q} , \mathbf{K} , and \mathbf{V} represent the query, key, and value matrices derived from the multimodal embeddings \mathbf{E} , and d_k is the dimensionality of the key vectors.

During training, the decoder minimizes the negative log-likelihood of the ground truth answer A_{true} :

$$\mathcal{L}_{\text{decoder}} = -\sum_{t=1}^T \log P(A_{\text{true},t} | \mathbf{Q}, \mathbf{I}, A_{\text{true},<t}), \quad (14)$$

where T is the length of the answer sequence, and $A_{\text{true},<t}$ represents the partial sequence of the ground truth answer up to time step $t - 1$.

The total loss of the framework is defined as:

$$\mathcal{L}_{\text{total}} = \phi_1 \mathcal{L}_{\text{GonF}} + \phi_2 \mathcal{L}_{\text{AMA}} + \phi_3 \mathcal{L}_{\text{decoder}}, \quad (15)$$

where ϕ_1 , ϕ_2 , and ϕ_3 are hyperparameters balancing the contributions of the three modules, and satisfy $\phi_1, \phi_2, \phi_3 \geq 0$ and $\phi_1 + \phi_2 + \phi_3 = 1$.

In summary, the Multimodal Decoding module integrates visual and textual information through cross-attention to

generate accurate answers. Optimizing and balancing the total loss ensures robust multimodal understanding and enhances overall framework performance.

Experiments

To comprehensively evaluate our proposed MacVQA framework, we conduct extensive experiments addressing six core research questions:

- **RQ1:** How does MacVQA’s performance compare with baselines under both standard and novel composition training paradigms?
- **RQ2:** What contribution does each core component make to MacVQA’s overall effectiveness?
- **RQ3:** How do different prototype selection strategies impact model accuracy and stability?
- **RQ4:** How does memory buffer capacity influence knowledge retention and forgetting in MacVQA?
- **RQ5:** How sensitive is MacVQA to key hyperparameter configurations?
- **RQ6:** How robustly does MacVQA generalize to real-world datasets with dynamically evolving tasks?

| Method | Mechanism | standard | | novel | |
|------------|----------------|--------------|-------------|--------------|-------------|
| | | AP | AF | AP | AF |
| ProtoGroup | Random | 37.22 | 3.89 | 35.78 | 4.54 |
| | Max-Similarity | 39.81 | 2.87 | 36.81 | 4.09 |
| MacVQA | Random | 40.29 | 3.46 | 41.25 | 3.76 |
| | Max-Similarity | 43.38 | 2.32 | 42.53 | 3.60 |

Table 3: Performance comparison of prototype selection strategies (Random vs. Max-Similarity) under standard and novel composition scenarios.




Figure 3: Similarity score radar plots for two examples (the k -th data from task $i-1$ and task i). Left: question memory; right: visual memory. Top-3 prototypes are selected by similarity.

Experimental Settings

Datasets We evaluate on the VQA v2 benchmark, which includes over 200k COCO images (Lin et al. 2014) and 1.1M annotated QA pairs. Following VQACL, the dataset is partitioned into 10 tasks by question type to simulate real-world sequential learning (recognition, location, judge, commonsense, count, action, color, type, subcategory, and causal.).

Evaluation Metrics We report Final Average Performance (AP) and Average Forgetting (AF) (Devlin et al. 2019), which respectively reflect the overall accuracy after continual learning and the degradation on previously learned tasks.

Baselines Methodologies are categorized by core mechanisms: (1) Vanilla (sequential fine-tuning without continual learning mechanisms, serving as lower-bound benchmark); (2) Regularization-based methods: EWC (Kirkpatrick et al. 2017), MAS (Aljundi et al. 2018); (3) Replay-based methods: ER (Chaudhry et al. 2019), DER (Chaudhry et al. 2018), VQACL (Zhang, Zhang, and Xu 2023), QUAD (Marouf et al. 2025), ProtoGroup (Zhang et al. 2025).

Implementation Details All methods use the same pre-trained VL-T5 backbone (Cho et al. 2021). Evaluation follows two paradigms: standard testing on seen skill-concept combinations, and novel composition testing on held-out combinations. A dynamically optimized memory buffer is employed. We use Adam (Adam et al. 2014) with a $3e-5$ learning rate, gradient clipping of 5, and a 0.1 warmup ratio.

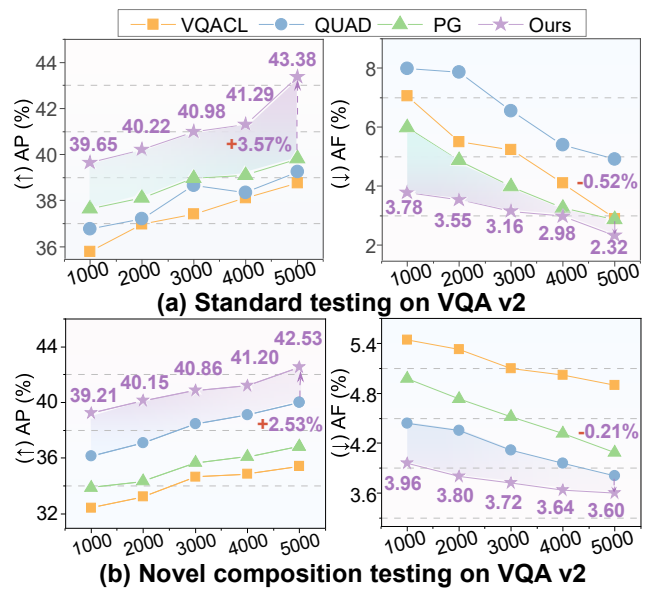


Figure 4: Memory size sensitivity analysis for VQA v2 under standard and novel composition testing paradigms

Main Results (RQ1)

As shown in Table 1, MacVQA achieves an AP of 43.38% and an AF of 2.32% on the Standard Test, outperforming all baseline methods. On the Novel Composition Test, it maintains strong results with 42.53% AP and 3.60% AF. In standard settings, MacVQA achieves the highest accuracy on tasks such as *judgment* (66.64%), *action* (58.13%), and *color* (56.95%), reflecting its strength in both reasoning and perception. Under the more challenging novel setting, it delivers top performance in *commonsense* reasoning with 67.05% accuracy, and shows clear advantages in *subcategory* classification and *causal* reasoning with 47.21% and 17.00%, respectively. These results confirm MacVQA’s ability to retain prior knowledge, adapt to new tasks, and generalize effectively across all ten continual VQA tasks. This highlights the effective coordination between visual filtering and memory adaptation components.

Effects of Modality (RQ2)

Table 2 reports the ablation study on the Global Noise Filtering (GonF) and Adaptive Memory Allocation (AMA) modules. Introducing either module individually improves performance over the baseline, with GonF reducing AF from 2.90% to 2.14% and AMA increasing AP from 38.77% to 40.97%. When combined, the full MacVQA model achieves the best results with 43.38% AP and 2.32% AF. GonF enhances perception tasks such as recognition and location, while AMA supports memory-sensitive reasoning like color and typicality. GonF’s early filtering curbs error propagation, and AMA promotes long-term retention by prioritizing salient prototypes. Together, they form a complementary mechanism that balances feature robustness and continual learning.




Figure 5: Task-Specific Hyperparameter Sensitivity of MacVQA.

Prototype Selection Strategies (RQ3)

As shown in Table 3, prototype selection plays a crucial role in continual VQA. Compared to random sampling, the Max-Similarity strategy consistently yields better performance across both standard and novel settings. On the Standard Test, it achieves 43.38% AP and 2.32% AF, outperforming the random variant by notable margins. Figure 3 shows that selecting top-3 prototypes by similarity enables retrieval of semantically aligned exemplars, improving rehearsal relevance and model stability. These results confirm that informed prototype selection is key to effective knowledge retention and generalization.

Memory Capacity Sensitivity Analysis (RQ4)

Figure 4 illustrates the impact of memory size on model performance under both standard and novel composition settings. While all methods benefit from increased buffer capacity, the degree and stability of improvement vary. VQACL shows rapid AP gains at smaller sizes but suffers from fluctuating AF, while QUAD consistently exhibits high AF, indicating inefficient memory use. ProtoGroup yields modest gains. In contrast, MacVQA maintains the highest AP and lowest AF across all memory scales. At the largest buffer size (5000), it achieves 43.38% AP and 2.32% AF on the standard test, and 42.53% AP with 3.60% AF on the novel test, outperforming the best baseline by +3.57% AP and -0.52% AF. These findings underscore MacVQA’s efficient memory use and its robustness across varying capacities.

Hyperparameter Sensitivity Analysis (RQ5)

Figure 5 presents the sensitivity analysis of MacVQA under different combinations of visual contribution α and question contribution β , showing clear task-specific preferences. For visual-dominant tasks such as recognition or location (Figure 5(a)), performance peaks at $\alpha = 1.0$, $\beta = 0.4$ with an AP of 43.53%. In contrast, question-dominant tasks such as commonsense or judgment (Figure 5(b)) favor higher textual weights, with a best performance of 48.5% at $\alpha = 0.4$, $\beta = 1.0$. This highlights MacVQA’s flexibility in adapting to diverse reasoning needs through multimodal fusion. These findings support the use of modality-aware weighting and validate the design of MacVQA’s dynamic prototype integration.

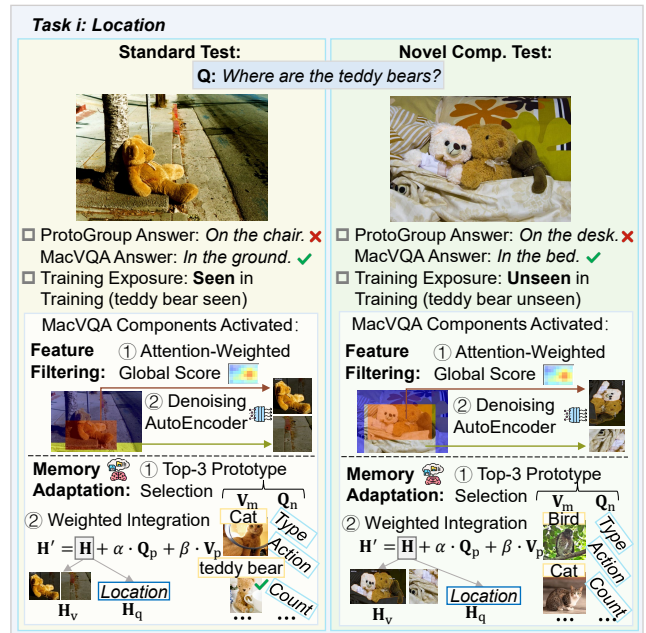


Figure 6: MacVQA exemplars in the Location task under standard and novel test settings.

Real-world Robustness Analysis (RQ6)

Figure 6 presents qualitative exemplars that illustrate MacVQA’s robustness across both standard and novel testing scenarios. In the location task, the model correctly answers “in the chair” for a standard test image where the object type (teddy bear) was seen during training. It also succeeds on a novel case involving an unseen combination—answering “in the bed” when the teddy bear type had not appeared in training for this reasoning context. In contrast, baseline methods fail in both settings. These examples highlight MacVQA’s ability to generalize to unseen object-reasoning pairs while maintaining accuracy on familiar inputs. This stems from its feature filtering mechanism, which reduces visual noise, and its adaptive memory strategy, which retrieves semantically aligned prototypes. Together, these components enable robust performance under distribution shifts in continual learning.

Conclusion

We propose MacVQA to address the challenges of balancing knowledge retention, adaptation efficiency, and robust reasoning in multimodal continual learning. By introducing Global Noise Filtering for robust feature representation and Adaptive Memory Allocation for efficient and effective memory management, MacVQA achieves a balanced fusion of visual and textual information while mitigating the impact of noise. Experiments on ten continual VQA tasks from the VQA v2 dataset demonstrate significant improvements in both accuracy and resistance to forgetting, validating the effectiveness of MacVQA. Future work will extend this framework to broader modalities and explore dynamic memory strategies for large-scale scenarios.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (No. 62207011, 62407013, 62377009, 62101179), the Natural Science Foundation of Hubei Province of China (No. 2025AFB653), the Natural Science Foundation of Shandong Province of China (No. ZR2024QF257), the Science and Technology Support Plan for Youth Innovation of Colleges and Universities of Shandong Province of China (No. 2023KJ370), the Open Fund of Hubei Key Laboratory of Big Data Intelligent Analysis and Application, Hubei University (No. 2024BDIAA05), and the Open Fund of Key Laboratory of Intelligent Sensing System and Security of Hubei University, Ministry of Education (No. KLISSS202410).

References

- Adam, K. D. B. J.; et al. 2014. A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 1412(6).
- Agrawal, A.; Batra, D.; Parikh, D.; and Kembhavi, A. 2018. Don't just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4971–4980.
- 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*, 139–154.
- Bai, L.; Islam, M.; and Ren, H. 2023. Revisiting distillation for continual learning on visual question localized-answering in robotic surgery. In *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*, 68–78.
- Chaudhry, A.; Dokania, P. K.; Ajanthan, T.; and Torr, P. H. 2018. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European conference on computer vision*, 532–547.
- Chaudhry, A.; Rohrbach, M.; Elhoseiny, M.; Ajanthan, T.; Dokania, P.; Torr, P.; and Ranzato, M. 2019. Continual learning with tiny episodic memories. In *Workshop on Multi-Task and Lifelong Reinforcement Learning*.
- Chen, Y.-C.; Li, L.; Yu, L.; El Kholi, A.; Ahmed, F.; Gan, Z.; Cheng, Y.; and Liu, J. 2020. Uniter: Universal image-text representation learning. In *Proceedings of the European conference on computer vision*, 104–120.
- Cho, J.; Lei, J.; Tan, H.; and Bansal, M. 2021. Unifying vision-and-language tasks via text generation. In *Proceedings of the International Conference on Machine Learning*, 1931–1942.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies*, 4171–4186.
- Gómez-Pérez, J. M.; and Ortega, R. 2020. ISAAQ - Mastering Textbook Questions with Pre-trained Transformers and Bottom-Up and Top-Down Attention. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, 5469–5479.
- Goyal, Y.; Khot, T.; Summers-Stay, D.; Batra, D.; and Parikh, D. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6904–6913.
- Greco, C.; Plank, B.; Fernández, R.; and Bernardi, R. 2019. Psycholinguistics Meets Continual Learning: Measuring Catastrophic Forgetting in Visual Question Answering. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, 3601–3605.
- Hayes, T. L.; Kafle, K.; Shrestha, R.; Acharya, M.; and Kanan, C. 2020. Remind your neural network to prevent catastrophic forgetting. In *Proceedings of the European conference on computer vision*, 466–483.
- Hudson, D. A.; and Manning, C. D. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6700–6709.
- Huh, M.; Xu, F.; Peng, Y.-H.; Chen, C.; Murugu, H.; Gurari, D.; Choi, E.; and Pavel, A. 2024. Long-Form Answers to Visual Questions from Blind and Low Vision People. In *Proceedings of the First Conference on Language Modeling*.
- Kim, J.; Ku, Y.; Kim, J.; Cha, J.; and Baek, S. 2024. Vlm-pl: Advanced pseudo labeling approach for class incremental object detection via vision-language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4170–4181.
- Kim, J.-H.; Jun, J.; and Zhang, B.-T. 2018. Bilinear attention networks. *Advances in neural information processing systems*, 31.
- 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.
- Krishna, R.; Bernstein, M.; and Fei-Fei, L. 2019. Information maximizing visual question generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2008–2018.
- Lei, S. W.; Gao, D.; Wu, J. Z.; Wang, Y.; Liu, W.; Zhang, M.; and Shou, M. Z. 2023. Symbolic replay: Scene graph as prompt for continual learning on vqa task. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1250–1259.
- 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 *Proceedings of the International conference on machine learning*, 3925–3934.
- Liang, Z.; Jiang, W.; Hu, H.; and Zhu, J. 2020. Learning to contrast the counterfactual samples for robust visual question answering. In *Proceedings of the 2020 conference on empirical methods in natural language processing*, 3285–3292.

- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Proceedings of the European conference on computer vision*, 740–755.
- Lv, G.; Jiang, X.; Sun, Y.; Ni, W.; and Nian, F. 2023. VQA-CLPR: Turning a visual question answering model into a Chinese license plate recognizer. In *Proceedings of the International Conference on Image and Graphics*, 348–359.
- Marouf, I. E.; Tartaglione, E.; Lathuiliere, S.; and van de Weijjer, J. 2025. No Images, No Problem: Retaining Knowledge in Continual VQA with Questions-Only Memory. *arXiv preprint arXiv:2502.04469*.
- Naik, R.; and Nushi, B. 2023. Social biases through the text-to-image generation lens. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 786–808.
- Paik, I.; Oh, S.; Kwak, T.; and Kim, I. 2020. Overcoming catastrophic forgetting by neuron-level plasticity control. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 5339–5346.
- 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 *Proceedings of the International conference on machine learning*, 8748–8763.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28.
- Singh, A.; Natarajan, V.; Shah, M.; Jiang, Y.; Chen, X.; Batra, D.; Parikh, D.; and Rohrbach, M. 2019. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8317–8326.
- Smith, J.; Hsu, Y.-C.; Balloch, J.; Shen, Y.; Jin, H.; and Kira, Z. 2021. Always be dreaming: A new approach for data-free class-incremental learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9374–9384.
- Sun, L.; Zhang, H.; Xu, W.; and Tomizuka, M. 2023. Efficient multi-task and transfer reinforcement learning with parameter-compositional framework. *IEEE Robotics and Automation Letters*, 8(8): 4569–4576.
- Tan, H.; and Bansal, M. 2019. LXMERT: Learning Cross-Modality Encoder Representations from Transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, 5099–5110.
- Theeuwes, J. 2010. Top-down and bottom-up control of visual selection. *Acta psychologica*, (2): 77–99.
- Tiwari, R.; Killamsetty, K.; Iyer, R.; and Shenoy, P. 2022. Gcr: Gradient coreset based replay buffer selection for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 99–108.
- Whitehead, S.; Wu, H.; Ji, H.; Feris, R.; and Saenko, K. 2021. Separating skills and concepts for novel visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5632–5641.
- Wu, D.; Dong, X.; Shao, L.; and Shen, J. 2022. Multi-level representation learning with semantic alignment for referring video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4996–5005.
- Xu, Y.; Zhong, X.; Yepes, A. J. J.; and Lau, J. H. 2020. Forget me not: Reducing catastrophic forgetting for domain adaptation in reading comprehension. In *Proceedings of the 2020 International Joint Conference on Neural Networks*, 1–8.
- Yang, Y.; Zhou, J.; Ding, X.; Huai, T.; Liu, S.; Chen, Q.; Xie, Y.; and He, L. 2025. Recent advances of foundation language models-based continual learning: A survey. *ACM Computing Surveys*, 57(5): 1–38.
- Zhang, F.; Yan, M.; Zhang, J.; and Xu, C. 2022. Comprehensive relationship reasoning for composed query based image retrieval. In *Proceedings of the 30th ACM international conference on multimedia*, 4655–4664.
- Zhang, L.; Mao, Z.; Peng, Y.; Fu, Z.; and Zhang, Y. 2025. Multi-Prototype Grouping for Continual Learning in Visual Question Answering. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, 1–5.
- Zhang, X.; Zhang, F.; and Xu, C. 2023. Vqacl: A novel visual question answering continual learning setting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19102–19112.
- Zhou, D.-W.; Zhang, Y.; Wang, Y.; Ning, J.; Ye, H.-J.; Zhan, D.-C.; and Liu, Z. 2025. Learning without forgetting for vision-language models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zhu, Y.; Groth, O.; Bernstein, M.; and Fei-Fei, L. 2016. Visual7w: Grounded question answering in images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4995–5004.