

Enhancing Continual Learning in Visual Question Answering with Modality-Aware Feature Distillation

Malvina Nikandrou¹ Georgios Pantazopoulos^{1,2} Ioannis Konstas^{1,2} Alessandro Suglia^{1,2}

¹Heriot-Watt University; ²Alana AI

{mn2002, gmp2000, i.konstas, a.suglia}@hw.ac.uk

Abstract

Continual learning focuses on incrementally training a model on a sequence of tasks with the aim of learning new tasks while minimizing performance drop on previous tasks. Existing approaches at the intersection of Continual Learning and Visual Question Answering (VQA) do not study how the multimodal nature of the input affects the learning dynamics of a model. In this paper, we demonstrate that each modality evolves at different rates across a continuum of tasks and that this behavior occurs in established encoder-only models as well as modern recipes for developing Vision & Language (VL) models. Motivated by this observation, we propose a modality-aware feature distillation (MAFED) approach which outperforms existing baselines across models of varying scale in three multimodal continual learning settings. Furthermore, we provide ablations showcasing that modality-aware distillation complements experience replay. Overall, our results emphasize the importance of addressing modality-specific dynamics to prevent forgetting in multimodal continual learning.

1 Introduction

Large Language Models (LLMs) (Touvron et al., 2023; Jiang et al., 2023) and Visual Language Models (VLMs) (Bai et al., 2023b; Liu et al., 2024), have achieved unprecedented performance and have become the go-to option for most NLP and Vision & Language (VL) tasks. However, once they have been trained, it is not straightforward how to update them to accommodate for novel concepts or concepts with reworked meanings. As a result, over time, the knowledge of these models may be obsolete or needs to be refined periodically to maintain their relevance. Consider two examples: 1) As of July 2023, Twitter has been re-branded to X with a new logo. 2) According to the Oxford English Dictionary (OED), as of March 2024, more than 1000 English words or phrases have been either




Figure 1: Overview of MAFED. Along with training on the data from the current task and a memory buffer, we apply feature distillation using the previous checkpoint as the teacher. The distillation losses applied to the representations from question and visual tokens are weighted separately to compensate for modality-specific training dynamics.

been revised or included as novel entries¹. In these cases, models trained with data from a preceding period will inevitably show performance deterioration (Lazaridou et al., 2021). Commercialized LLMs circumvent this limitation (OpenAI, 2022; Gemini Team et al., 2023; Anthropic, 2024) with statements regarding the knowledge cutoff of these models. On the other hand, humans continuously update their knowledge and acquire new skills over time. Continual learning is a paradigm that aims to simulate this behavior, focusing on models that learn incrementally from a sequence of tasks with minimal catastrophic forgetting (McCloskey and Cohen, 1989; Ratcliff, 1990).

Continual learning has started being explored more widely in VL settings (Greco et al., 2019; Srinivasan et al., 2022; Nikandrou et al., 2022;

¹OED March 2024 update lists entries appearing for the first time, or entries with updated meanings.

Zhang et al., 2023; Lei et al., 2023; Cui et al., 2024). However, existing approaches do not explicitly account for the dissimilarities in the representation space of multimodal inputs and their effect on the learning dynamics, which we argue is necessary for effective continual learning of VL tasks. Previous work on the optimization dynamics of multimodal learning has demonstrated that different modalities are learned at different speeds (Wang et al., 2020; Wu et al., 2022). Using encoder and decoder-only VLMs, we empirically showcase a similar forgetting discrepancy (Section 4.1), which indicates that representations from each modality evolve at different rates across a sequence of tasks.

Motivated by this observation, we propose *MAFED*, a Modality-Aware FEature Distillation approach summarized in Figure 1. We explore different strategies for weighting the distillation losses derived from the tokens of each modality, using either fixed balanced weights or adaptive weights derived from the loss gradients computed with respect to the inputs. We combine both variants with experience replay and show promising results across all VLM families compared to established continual learning methods. Additionally, we conduct experiments with decoder-only VLMs ranging from 100M to 1B parameters, showing that although scale alleviates forgetting, certain settings remain challenging. Our ablations comparing experience replay and feature distillation approach showcase that these methods are complementary and yield greater performance when combined. Overall, our results emphasize the need to address modality-specific dynamics to effectively mitigate forgetting in multimodal continual learning.

2 Related Work

2.1 Continual Learning

Continual Learning Approaches Continual learning approaches can be categorized as regularization, replay, and architecture-based (Delange et al., 2021). Regularization-based approaches introduce auxiliary losses that aim to constrain the model weights (Kirkpatrick et al., 2017; Zenke et al., 2017; Aljundi et al., 2018), outputs (Li and Hoiem, 2018; Rebuffi et al., 2017), or internal representations (Hou et al., 2019). Replay-based approaches rely on storing (Chaudhry et al., 2019; Buzzega et al., 2020; Bagus and Gepperth, 2021) or generating samples (Van de Ven and Tolia, 2018; Sun et al., 2019) from past tasks so that they can be

sampled along with new samples during training. Finally, architecture-based approaches introduce task-specific parameters, either by masking the model parameters (Yoon et al., 2020; Serrà et al., 2018) or adding new ones for each task (Fernando et al., 2017; Madotto et al., 2021).

Most recent work tends to combine techniques from multiple categories in order to maximize performance. Similarly, our work utilizes replay and regularization through feature distillation. Distillation has been used in various continual learning approaches. Some focus on knowledge distillation on the output level, using the logits (Li and Hoiem, 2018) or pseudo-labels from a past checkpoint (Wang et al., 2022; Karim et al., 2022). Other work applies distillation on the internal model representations (Dhar et al., 2019; Douillard et al., 2020; Kang et al., 2022). *MAFED* expands this line of work by introducing different weighting schemes to balance the distillation loss from visual and textual representations.

VL Continual Learning Previous work has studied varying instantiations of VL continual learning problems, including image captioning (Del Chiaro et al., 2020; Nguyen et al., 2019), compositional phrase generalization (Jin et al., 2020), and task-incremental learning (Srinivasan et al., 2022). Within VQA, prior work has investigated continual learning based on question types (Greco et al., 2019), across varying domains (Zhang et al., 2022; Lao et al., 2023), and from a compositionality perspective (Zhang et al., 2023). Lei et al. (2023); Nikandrou et al. (2022) further study how VQA models evolve in different settings, including novel visual scenes or different question types. However, these works have not investigated the effect of modality-aware methods with the exception of Qian et al. (2023) that focus on multimodal prompt learning for vision, text and fusion modules. In contrast, feature distillation does not assume separate modality-specific and multimodal parameters and can be applied to more varied VLM architectures.

2.2 VLMs

Progress in representation learning has led to models that achieve impressive performance across multimodal benchmarks (Goyal et al., 2017; Hudson and Manning, 2019; Li et al., 2023). Early approaches relied on complicated architectures (Tan and Bansal, 2019; Lu et al., 2019), and multiple learning objectives (Chen et al., 2020; Li et al.,




Figure 2: Illustration of tasks in each of the three continual learning settings for VQA. Each of these settings consists of five tasks. The first two settings are defined based on the visual categories. In Diverse Content, the objects present in each task are grouped randomly, while in Taxonomy Content, the objects are grouped based on their supercategory. Finally, in Question Types, the tasks are defined according to the type of the questions.

2021; Jia et al., 2021). More recently, given the rapid development of increasingly capable LLMs (Touvron et al., 2023; Jiang et al., 2023; Bai et al., 2023a), these approaches have been superseded by a new paradigm where representations from visual experts (Radford et al., 2021; Oquab et al., 2024) are treated as input tokens for the LLM. This shift has led the development of modern VLMs (Liu et al., 2024; Dai et al., 2024; Laurençon et al., 2024a) that are based on the same underlying principles with deviations regarding the choice of the experts, or how the patch tokens are integrated into the language model. Our experiments demonstrate the effectiveness of our approach in both encoder-only (Chen et al., 2020; Kim et al., 2021) as well as decoder-only models.

3 Preliminaries

3.1 Data

We leverage an existing evaluation suite for Continual Learning in VQA (Nikandrou et al., 2022) comprised of three settings based on the visual and the language input. In particular, each of these settings consists of five tasks and is designed to test the model’s performance on learning varying concepts or question types. Figure 2 illustrates exemplary images and questions from different tasks in each of the settings. Below, we provide a brief summary for each of these settings.

Diverse Content corresponds to a real-world use case as well as a common standard procedure within continual learning (Lomonaco and Maltoni, 2017; Rebuffi et al., 2017; Zenke et al., 2017; Lin et al., 2021), where a model is trained on new sets

of concepts progressively that do not necessarily comply to a taxonomy. Each task in this setting covers 10 distinct object categories from the COCO dataset (Lin et al., 2014).

Taxonomy Content In this setting, each task consists of questions regarding objects based on the same super-category. This setting contains questions from the following categories: Animals, Food, Interior, Sports, and Transport, and simulates a more progressive approach, where a model learns about a fixed (and similar) pool of concepts before being applied to a different domain. Importantly, we note that in both Diverse Content and Taxonomy Content, images containing objects shared between tasks are discarded to create clean task splits, preventing contamination between them. In total, there are 181K train, 45K validation, and 110K test samples for both settings.

Question Types The final setting resembles a scenario where the model learns to answer different categories of questions. In this setting, the model is tasked with learning from a sequence of five tasks: Count, Color, Scene-level, Subcategory, and Action recognition. Question Types have a total of 140K train, 35K validation, and 84K test samples.

3.2 Models

Throughout our experiments, we use two families of models, including encoder- and decoder-only pretrained models. Specifically, we use the encoder-only models, UNITER-base (Chen et al., 2020) and ViLT-base (Kim et al., 2021), that differ in terms of how the visual input is encoded. UNITER uses region features extracted from the Faster R-CNN

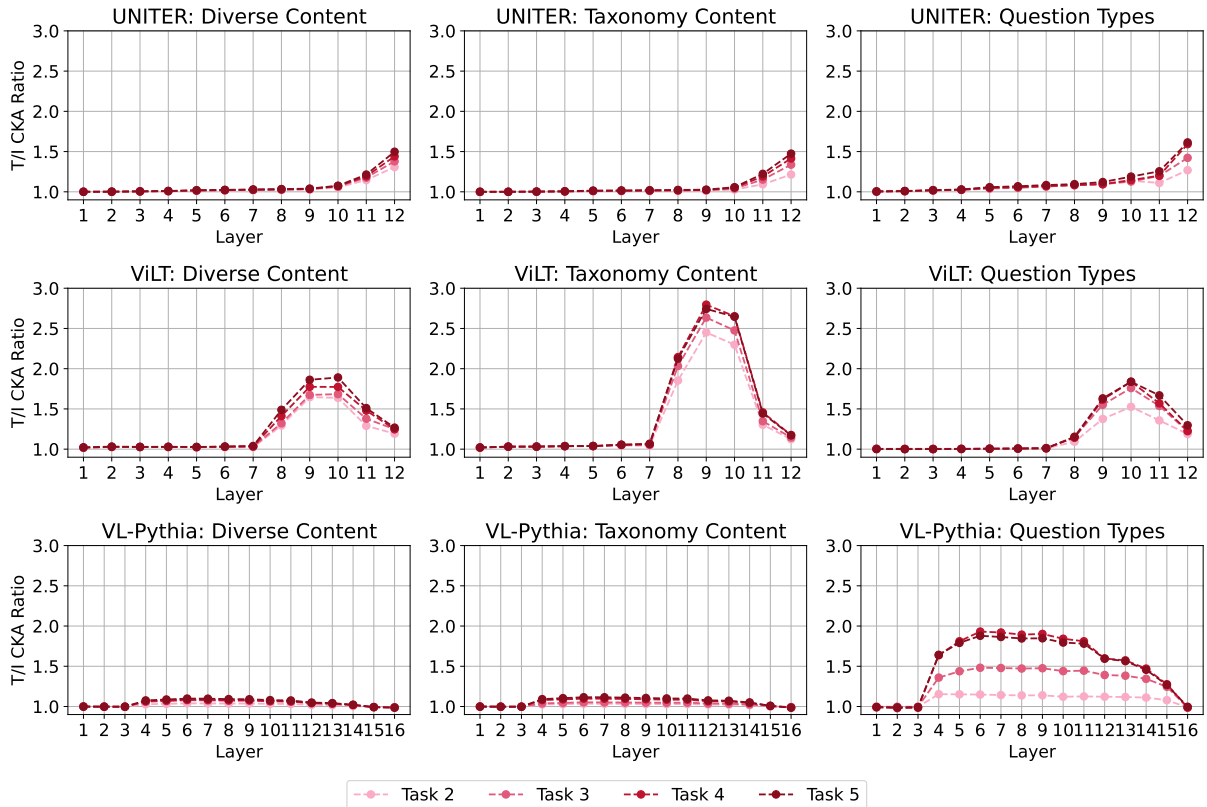


Figure 3: Ratio of text-to-image representation similarity across layers and tasks, for UNITER (first row), ViLT (second-row), and VL-Pythia (third-row). We consistently observe that in the earlier layers, the ratio is close to one, indicating that representations from both modalities change at a similar rate. However, in intermediate or deeper layers, text representations seem to retain larger similarities.

object detector (Anderson et al., 2018). On the other hand, ViLT is a patch-based model that does not use an additional vision encoder.

However, recent trends in the development of VLMs have moved towards decoder-only architectures that combine an expert vision encoder with an LLM using a connector module that learns a mapping between them (Liu et al., 2024). We employ a similar recipe to combine EVA02 (Fang et al., 2023) as the visual encoder and Pythia (Biderman et al., 2023) as our LLM. Regarding the connection module, we follow the LLaVA-1.5 (Liu et al., 2023) approach with a two-layer MLP that matches the dimensionality of the visual and the language embeddings. We refer to this model as VL-Pythia.

We refrain from using other existing VLMs for two reasons. First, we aim to match the parameters and pretraining data between the encoder-only and the generative models. Pythia models provide a collection of checkpoints covering a wide range of sizes that have been pretrained following a state-of-art transformer recipe (Su et al., 2024; Dao et al., 2022). For a controlled comparison with encoder-only models, we train VL-Pythia us-

ing the same data used to train UNITER and ViLT (see Appendix A.1 for additional details regarding the pretraining of the model). In particular, we use the checkpoints with 160M, 410M, and 1B parameters. The smaller model is on par with the encoder-only ones, while the larger models allow us to explore the role of model capacity. Secondly, existing VLMs (Liu et al., 2024; Dai et al., 2024; Bai et al., 2023a; Laurençon et al., 2024a,b) are typically instruction-tuned on datasets that include VQA-v2 (Goyal et al., 2017) on which the continual learning settings are based. This overlap is undesirable since we want to prevent any form of data contamination between tasks.

During continual learning, we keep the vision encoders of UNITER and VL-Pythia frozen. In encoder-only models, VQA is treated as a classification task. The classification (CLS) token is passed to a classification head that gets expanded with the new answers from each task. On the contrary, VL-Pythia is fine-tuned to generate answers autoregressively. In our experiments, we follow a greedy decoding strategy during inference.

4 Method

4.1 Motivation

In this work, we argue that modality-specific learning dynamics, and more specifically, the different speeds at which each modality is learned (Wu et al., 2022) or forgotten, should be accounted for in multimodal continual learning settings. We demonstrate this behavior by measuring the similarities of the question Q and image V representations from sequential model checkpoints using Centered Kernel Alignment (CKA) (Kornblith et al., 2019). In particular, we extract text Q_t and image V_t representations for data of the first task after training on the first $t = 1$ and each subsequent task $t = 2 \dots T$, and we compute the CKA similarity across time. Finally, we visualize the ratio R_t of the text over image similarities:

$$R_t = \frac{CKA(Q_1, Q_t)}{CKA(V_1, V_t)} \quad \forall t = 2 \dots T \quad (1)$$

Figure 3 shows the Text-to-Image CKA ratio per layer. We observe that there are differences in how the modalities evolve within the models and across settings. First, we note that the ratio is greater or equal to one in all cases, meaning that visual tokens exhibit decreasing similarity throughout the learning process. In UNITER, the ratio remains close to one for earlier layers and increases only for the last three layers. ViLT and Pythia exhibit a different trend, where the ratio peaks for intermediate layers. As a result, we hypothesize that incorporating the variability of each modality in a regularization technique can benefit continual learning strategies. Due to parameter sharing between the two modalities, we materialize this in a modality-aware feature distillation strategy, which we elaborate on below.

4.2 Modality-Aware Feature Distillation

Feature Distillation (FD) is an established continual learning technique (Hou et al., 2019; Douillard et al., 2020; Kang et al., 2022) that adds a regularization loss term to prevent the drift of model representations. Given two model checkpoints f_{t-1} and f_t from consecutive tasks, we extract representations H from an intermediate layer and compute the feature distillation loss l_i using the representations of each token h_i :

$$L_{FD} = \frac{1}{N} \sum_{i=1}^N l_i = \frac{1}{N} \sum_{i=1}^N \|h_{i,t} - h_{i,t-1}\|_2^2 \quad (2)$$

Assuming an example has Q text tokens and V visual tokens, we can rewrite Equation (2) in terms of the average loss contributed by the language and the vision tokens, $L_{FD,Q}$ and $L_{FD,V}$ respectively:

$$L_{FD,weighted} = \alpha \cdot L_{FD,Q} + (1 - \alpha) \cdot L_{FD,V} \quad (3)$$

In the simplest case, α is proportional to the number of tokens available from each input modality. In practice, this might be suboptimal since it depends on the input tokenization strategy. In fact, across the examined settings, the average tokenized inputs have approximately 9 question tokens and 33 or 199 visual tokens for region and patch-based image features, respectively. Consequently, visual tokens will dominate the distillation loss of Equation (3) potentially leading to inferior performance.

Therefore, we experiment with two modifications: i) MAFED-B which balances the losses from each modality by fixing α to 0.5, and ii) MAFED-A which uses an adaptive weighting approach based on modality importance inspired by Kang et al. (2022). Modality importances I_Q and I_V are estimated using the gradient of the VQA classification loss², with respect to the intermediate model representations H_m from each modality m . I_Q and I_V are updated at the beginning of each training task using the available memory data M_t as follows:

$$I_m = \mathbb{E}_{(x,y) \sim M_t} [\|\nabla_{H_m} L_{cls}(f_t(x), y)\|_F^2] \quad (4)$$

where $\|\cdot\|$ corresponds to the Frobenius norm. Finally, the weight α is computed by normalizing the importance of the question tokens:

$$\alpha = \frac{I_Q}{I_Q + I_V} \quad (5)$$

We apply feature distillation to all layers except the last since only the representation of CLS token in encoder-only and the final question token in decoder-only models is propagated to the model output’s head. Furthermore, in Section 4.1, we showcased that the representations of deeper layers are affected more during continual learning. Therefore, we introduce a discount factor w_d that is used to weigh the contribution of the loss from each layer proportionally to its distance d from the model’s head:

$$w_d = \frac{\gamma^d}{\sum_{d=0}^D \gamma^d} \quad (6)$$

²Or the language model head in the case of Pythia.

Model	Method	Diverse Content		Taxonomy Content		Question Types	
		Accuracy	SBWT	Accuracy	SBWT	Accuracy	SBWT
UNITER	FT*	64.59 ± 0.56	-1.93 ± 0.39	63.65 ± 0.63	-3.89 ± 0.53	48.81 ± 5.56	-22.43 ± 7.02
	EWC*	66.26 ± 0.55	-0.67 ± 0.29	67.70 ± 0.29	-0.62 ± 0.19	66.77 ± 3.54	-2.62 ± 2.28
	ER*	66.47 ± 0.51	-0.29 ± 0.18	66.76 ± 0.16	-1.22 ± 0.10	69.01 ± 0.76	-1.42 ± 0.31
	FD	66.67 ± 0.38	-0.17 ± 0.19	66.94 ± 0.23	-0.68 ± 0.17	69.53 ± 0.12	-1.22 ± 0.51
	MAFED-B	66.77 ± 0.24	-0.12 ± 0.17	67.05 ± 0.23	-0.57 ± 0.08	69.58 ± 0.55	-1.17 ± 0.31
	MAFED-A	66.52 ± 0.26	-0.23 ± 0.12	66.84 ± 0.25	-0.70 ± 0.45	69.34 ± 0.43	-0.94 ± 0.48
	Multitask*	69.76 ± 0.18	-	70.08 ± 0.18	-	72.54 ± 0.15	-
ViLT	FT*	61.07 ± 0.41	-2.80 ± 0.41	61.25 ± 0.50	-4.09 ± 0.50	36.95 ± 11.09	-32.86 ± 11.09
	EWC*	61.80 ± 0.96	-1.14 ± 0.96	63.69 ± 0.46	-0.92 ± 0.46	60.25 ± 2.86	-8.19 ± 2.86
	ER*	64.22 ± 0.10	-0.25 ± 0.10	63.52 ± 0.20	-1.46 ± 0.20	65.61 ± 0.76	-2.86 ± 0.76
	FD	64.57 ± 0.57	-0.51 ± 0.12	64.24 ± 0.73	-1.07 ± 0.46	67.70 ± 0.54	-1.98 ± 0.70
	MAFED-B	64.78 ± 0.55	-0.34 ± 0.27	64.51 ± 0.36	-1.02 ± 0.17	67.76 ± 0.27	-1.85 ± 0.61
	MAFED-A	65.00 ± 0.41	-0.28 ± 0.19	64.63 ± 0.37	-0.89 ± 0.21	67.67 ± 0.46	-2.01 ± 0.84
	Multitask*	67.51 ± 1.94	-	67.84 ± 3.92	-	72.41 ± 3.75	-

Table 1: UNITER and ViLT average accuracy and semantic backward transfer over five task orders. * results reported in (Nikandrou et al., 2022).

where $\gamma \in (0, 1]$ is a hyperparameter such that lower γ values assign more weight to deeper layers, and $\gamma = 1$ weighs the losses from all layers equally. Unless stated otherwise, we combine feature distillation with replay since it requires no computational overhead and helps mitigate the miscalibration of the output layer, which can negatively impact performance on past tasks (Wu et al., 2019).

5 Experiments

5.1 Baselines

We compare our methods against naive *Fine Tuning* (FT), where a model is trained sequentially on each task. Our continual learning baselines include *Elastic Weight Consolidation* (EWC) (Kirkpatrick et al., 2017) and *Experience Replay* (ER) (Chaudhry et al., 2019), which have been shown to perform competitively. Finally, we report the upper bound performance of *Multitask* learning, where the model is trained on all tasks simultaneously. Note that this is the de facto standard for instruction-tuning in VLMs (Dai et al., 2024; Liu et al., 2023).

5.2 Evaluation Metrics

We measure the performance of a model using three metrics. First, we report the macro-average accuracy at the end of a training sequence, $A = \frac{1}{T} \sum_{i=1}^T A_{T,i}$, where $A_{T,i}$ depicts the performance of the model on the data from task i after training on final task T . Additionally, we report the Semantic Backward Transfer (SBWT) (Nikandrou et al., 2022), which captures the impact of catastrophic forgetting, weighted by the semantic similarity of the prediction and the target:

$$\text{SBWT} = \frac{1}{T-1} \sum_{i=1}^{T-1} S_{T,i} \quad (7)$$

where $S_{T,i}$ is the average weighted accuracy difference for task i .

5.3 Implementation Details

We train all models using the Adam optimizer (Kingma and Ba, 2014) and a learning rate schedule that follows linear decay after a warmup for 10% of the training steps. The maximum learning rate is optimized through grid search separately for each setting based on the performance of the fine-tuning method. For UNITER and ViLT, we set the number of epochs per task to 60 with a patience of 5. We found that VL-Pythia models reach their peak accuracy for fewer updates, possibly because they do not use a randomly initialized classification head. As a result, we set the maximum number of epochs to 15. For all replay and feature distillation runs, we keep a memory of 1000 randomly selected samples per task, ensuring that the same samples are stored across methods for the same task order. Further details about the selected hyperparameters are listed in Appendix A.2.

5.4 Results

5.4.1 Encoder-only Models

Table 1 reports the results across the three settings using the encoder-only models UNITER and ViLT. Adding the feature distillation loss improves upon the ER baseline in all settings. Although the benefits with UNITER are moderate, as ER already achieves low forgetting, when using ViLT,

Setting	Method	VL-Pythia 160M		VL-Pythia 410M		VL-Pythia 1B	
		Accuracy	SBWT	Accuracy	SBWT	Accuracy	SBWT
Question Types	FT	25.98 \pm 8.23	-31.35 \pm 7.27	63.20 \pm 2.10	-7.22 \pm 1.98	65.52 \pm 5.60	-6.16 \pm 5.42
	EWC	41.55 \pm 8.31	-8.58 \pm 6.96	66.83 \pm 2.45	-3.85 \pm 2.48	66.78 \pm 2.98	-5.30 \pm 2.93
	ER	53.56 \pm 0.72	-5.20 \pm 1.06	70.25 \pm 1.00	-0.52 \pm 0.67	69.66 \pm 3.34	-2.27 \pm 1.88
	FD	56.19 \pm 1.59	-3.04 \pm 1.21	70.76 \pm 0.50	-0.22 \pm 0.25	71.85 \pm 1.07	-0.62 \pm 0.54
	MAFED-B	57.53 \pm 0.76	-2.83 \pm 0.42	70.82 \pm 0.38	-0.17 \pm 0.10	72.19 \pm 1.40	-0.55 \pm 0.80
	MAFED-A	57.65 \pm 0.24	-2.46 \pm 0.29	71.06 \pm 0.30	-0.19 \pm 0.28	72.69 \pm 0.12	-0.10 \pm 0.07
	Multitask	65.65 \pm 0.14	-	71.96 \pm 0.15	-	73.44 \pm 0.18	-

Table 2: Performance of different VL-Pythia model sizes across three task orders.

FD offers substantial accuracy gains of up to 2.6 in Question Types. Comparing the feature distillation variants, modality-aware weighting (MAFED-A or MAFED-B) consistently boosts performance. For UNITER and ViLT in Question Types, equally balancing the modality losses with MAFED-B shows the best performance. In the image-based settings with ViLT, adaptive weighting performs better. Given the similarity ratios shown in Figure 3, these results suggest that MAFED-B is more effective when the relative change of text and vision representation is small, while MAFED-A is more appropriate in cases of larger discrepancies.

5.4.2 Scaling to larger decoder-only VLMs

Setting	Accuracy	BWT
Diverse Content	70.11	1.58
Taxonomy Content	69.03	0.44
Question Types	66.01	-9.70

Table 3: Average accuracy and backward transfer for finetuning VL-Pythia 1B across settings. We report the accuracy of three task orders on the validation set.

As decoder-only architectures have become more widely used, we also experiment with three model sizes of VL-Pythia (160M, 410M, 1B parameters). In our initial results shown in Table 3, we find that larger models exhibit no forgetting in the image-based settings of Diverse and Taxonomy Content. As a result, we focus on more challenging Question Types setting.

Table 2 provides the results for different continual learning strategies using the VL-Pythia variants. We observe that scaling leads to higher final accuracy and less catastrophic forgetting similar to previous work (Mirzadeh et al., 2022; Ramasesh et al., 2022). Nevertheless, even the largest explored model has a gap of almost 8% between naive finetuning and multitask learning. Compared to experience replay, we find that the benefit of EWC

diminishes for larger models. As in encoder-only models, the inclusion of feature distillation further improves performance, and modality-aware weighting of the distillation losses is beneficial in all cases. For VL-Pythia, MAFED-A leads to the best performance, improving the accuracy on average by +2.6 compared to ER and +0.87 compared to FD. As mentioned in Section 5.4.1, we hypothesize that adaptive modality weighting can be particularly effective where the text and visual representations change more unequally. Overall, our results indicate that stabilizing the representations from vision and language tokens separately can improve multi-modal continual learning.

6 Analysis

6.1 Distillation without Replay

Method	Accuracy	SBWT
FT	62.77	-5.27
ER	66.18	-3.42
FD	72.05	-1.00
w/o Replay	63.53	-4.87
MAFED-B	72.91	-0.16
w/o Replay	67.66	-3.67
MAFED-A	72.56	-0.24
w/o Replay	64.14	-4.09

Table 4: Ablation of feature distillation methods without replay using VL-Pythia (1B) on Question Types.

In the previous section, we applied feature distillation in conjunction with experience replay. Table 4 shows the performance when applying feature distillation with and without experience replay on VL-Pythia (1B) for one task order on the Question Types setting. Feature distillation improves the model performance, while standalone modality-aware methods are competitive with experience replay. The combination of the two approaches yields the greatest performance with no additional

cost over standalone feature distillation. Both methods require maintaining a memory of past samples that are passed through the current model. Therefore, applying replay puts no additional overhead on training time and memory.

6.2 Distillation Layer Ablation




Figure 4: Ablation of feature distillation from a single or cumulative (+) model layers.

Next, we explore the effectiveness of applying modality-aware feature distillation from a single layer, as well as a subset of layers. Figure 4 illustrates the performance of VL-Pythia (1B) after applying MAFED-B and MAFED-A every three layers. As expected, applying both methods on a deeper layer of the model but also distilling from all previous layers yields greater performance. However, the performance of all four variants does not increase monotonically with the layer depth. More specifically, we observe that distilling from layer 6 leads to performance degradation. This behavior correlates with the results in Figure 3, where the per-modality similarities diverge the most at layer 6 and gradually align throughout the deeper layers of the model.

6.3 Modality Weights in MAFED-A

Figure 5 shows the weight α placed on the distillation loss of the language tokens using the MAFED-A method. In all models, MAFED-A assigns more weight to language tokens. Interestingly, we observe that in encoder-only models, the language weight progressively increases up to layer 8 and then drops. On the other hand, in VL-Pythia (1B), more than 90% of the weight is assigned to language tokens for all layers. We hypothesize that this is because the model is causal, and the last token before the answer is a text token.




Figure 5: Language weight during MAFED-A. Note that language tokens in the encoder family (UNITER and ViLT) are weighted similarly across the layers of the models. For the causal VL-Pythia model, the language tokens have higher weights.

7 Conclusion

In this paper, we argued that applying approaches that were developed with unimodal models in mind is suboptimal for continual learning in VQA since this ignores modality-specific learning dynamics. We empirically showcased that the visual and the textual representations evolve at different rates – a phenomenon that occurs in both encoder-only and decoder-only VLMs. Given this observation, we proposed two modality-aware feature distillation approaches that equally weigh the distillation loss from each modality or adaptively estimate the importance of a modality based on the gradients with respect to the inputs. We believe this is a promising direction towards closing the gap with multitask training in multimodal continual learning.

7.1 Limitations & Future Work

Despite the promising results, our method has certain limitations. First, distillation is more computationally expensive than replay, as it requires accessing the representations from the previous model. However, compared to established distillation methods, MAFED-B improves performance with no overhead, while MAFED-A requires computing importance weights, which are only updated between tasks. Furthermore, our work does not investigate the potential effectiveness of architecture-based approaches, which could offer greater control over the learning of each modality through novel parameter-isolation approaches. Finally, we show that larger models exhibit less or even no forgetting depending on the setting. Future work should explore whether increasing the model size or the VL pretraining data (Ostapenko et al., 2022) can further decrease forgetting in VQA settings.

Acknowledgements

This work was supported by the Edinburgh International Data Facility (EIDF) and the Data-Driven Innovation Programme at the University of Edinburgh.

References

- Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. 2018. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 139–154.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. [Bottom-up and top-down attention for image captioning and visual question answering](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Anthropic. 2024. [Introducing the next generation of claude](#).
- Benedikt Bagus and Alexander Gepperth. 2021. An investigation of replay-based approaches for continual learning. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–9. IEEE.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023a. [Qwen technical report](#). *arXiv preprint arXiv:2309.16609*.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023b. [Qwen-vl: A frontier large vision-language model with versatile abilities](#). *arXiv preprint arXiv:2308.12966*.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. 2023. [Pythia: A suite for analyzing large language models across training and scaling](#). In *International Conference on Machine Learning*, pages 2397–2430. PMLR.
- Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. 2020. [Dark experience for general continual learning: a strong, simple baseline](#). In *NeurIPS*.
- Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaisyasingam Ajanthan, Puneet Kumar Dokania, Philip H. S. Torr, and Marc’Aurelio Ran-zato. 2019. [Continual learning with tiny episodic memories](#). *CoRR*, abs/1902.10486.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholly, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. [Uniter: Universal image-text representation learning](#). In *Computer Vision – ECCV 2020*, pages 104–120, Cham. Springer International Publishing.
- Zhenyu Cui, Yuxin Peng, Xun Wang, Manyu Zhu, and Jiahuan Zhou. 2024. [Continual vision-language retrieval via dynamic knowledge rectification](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 11704–11712.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale N Fung, and Steven Hoi. 2024. [Instructblip: Towards general-purpose vision-language models with instruction tuning](#). *Advances in Neural Information Processing Systems*, 36.
- Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. [Flashattention: Fast and memory-efficient exact attention with io-awareness](#). *Advances in Neural Information Processing Systems*, 35:16344–16359.
- Riccardo Del Chiaro, Bartłomiej Twardowski, Andrew Bagdanov, and Joost van de Weijer. 2020. [Ratt: Recurrent attention to transient tasks for continual image captioning](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 16736–16748. Curran Associates, Inc.
- Matthias Delange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Greg Slabaugh, and Tinne Tuytelaars. 2021. [A continual learning survey: Defying forgetting in classification tasks](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1.
- Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, Ziyang Wu, and Rama Chellappa. 2019. [Learning without memorizing](#). In *Conference on Computer Vision and Pattern Recognition*.
- Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. 2020. [Podnet: Pooled outputs distillation for small-tasks incremental learning](#). In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pages 86–102. Springer.
- Yuxin Fang, Quan Sun, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. 2023. [Eva-02: A visual representation for neon genesis](#). *arXiv preprint arXiv:2303.11331*.
- Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A. Rusu, Alexander Pritzel, and Daan Wierstra. 2017. [Pathnet: Evolution channels gradient descent in super neural networks](#). *CoRR*, abs/1701.08734.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Slav Petrov,

- Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy Lillicrap, Angeliki Lazaridou, Orhan Firat, James Molloy, Michael Isard, Paul R. Barham, Tom Hennigan, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, Ryan Doherty, Eli Collins, Clemens Meyer, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha Goel, et al. 2023. [Gemini: A Family of Highly Capable Multimodal Models](#). Preprint, arxiv:2312.11805.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 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 (CVPR)*.
- Claudio Greco, Barbara Plank, Raquel Fernández, and Raffaella Bernardi. 2019. [Psycholinguistics meets continual learning: Measuring catastrophic forgetting in visual question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3601–3605, Florence, Italy. Association for Computational Linguistics.
- Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. 2019. [Learning a unified classifier incrementally via rebalancing](#). In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Drew A Hudson and Christopher D Manning. 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*, pages 6700–6709.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. [Scaling up visual and vision-language representation learning with noisy text supervision](#). In *International conference on machine learning*, pages 4904–4916. PMLR.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7b](#). *arXiv preprint arXiv:2310.06825*.
- Xisen Jin, Junyi Du, Arka Sadhu, Ram Nevatia, and Xiang Ren. 2020. [Visually grounded continual learning of compositional phrases](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2018–2029, Online. Association for Computational Linguistics.
- Minsoo Kang, Jaeyoo Park, and Bohyung Han. 2022. [Class-incremental learning by knowledge distillation with adaptive feature consolidation](#). In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16050–16059. IEEE.
- Siddharth Karamcheti, Suraj Nair, Ashwin Balakrishna, Percy Liang, Thomas Kollar, and Dorsa Sadigh. 2024. [Prismatic vlms: Investigating the design space of visually-conditioned language models](#). *arXiv preprint arXiv:2402.07865*.
- Nazmul Karim, Umar Khalid, Ashkan Esmaeili, and Nazanin Rahnavard. 2022. [Cnll: A semi-supervised approach for continual noisy label learning](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3878–3888.
- Wonjae Kim, Bokyung Son, and Ildoo Kim. 2021. [Vilt: Vision-and-language transformer without convolution or region supervision](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 5583–5594. PMLR.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. [Overcoming catastrophic forgetting in neural networks](#). *Proceedings of the national academy of sciences*, 114(13):3521–3526.
- Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. 2019. [Similarity of neural network representations revisited](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 3519–3529. PMLR.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. 2017. [Visual genome: Connecting language and vision using crowdsourced dense image annotations](#). *International journal of computer vision*, 123:32–73.
- Mingrui Lao, Nan Pu, Yu Liu, Zhun Zhong, Erwin M. Bakker, Nicu Sebe, and Michael S. Lew. 2023. [Multi-domain lifelong visual question answering via self-critical distillation](#). In *Proceedings of the 31st ACM International Conference on Multimedia*, MM '23, page 4747–4758, New York, NY, USA. Association for Computing Machinery.
- Hugo Laurençon, Lucile Saulnier, Léo Tronchon, Stas Bekman, Amanpreet Singh, Anton Lozhkov, Thomas Wang, Siddharth Karamcheti, Alexander Rush, Douwe Kiela, et al. 2024a. [Obelics: An open web-scale filtered dataset of interleaved image-text documents](#). *Advances in Neural Information Processing Systems*, 36.
- Hugo Laurençon, Léo Tronchon, Matthieu Cord, and Victor Sanh. 2024b. [What matters when building vision-language models?](#) *arXiv preprint arXiv:2405.02246*.

- Angeliki Lazaridou, Adhiguna Kuncoro, Elena Gribovskaya, Devang Agrawal, Adam Liska, Tayfun Terzi, Mai Gimenez, Cyprien de Masson d’Autume, Tomáš Kočiský, Sebastian Ruder, Dani Yogatama, Kris Cao, Susannah Young, and Phil Blunsom. 2021. [Mind the gap: Assessing temporal generalization in neural language models](#). In *Advances in Neural Information Processing Systems*.
- Stan Weixian Lei, Difei Gao, Jay Zhangjie Wu, Yuxuan Wang, Wei Liu, Mengmi Zhang, and Mike Zheng Shou. 2023. [Symbolic replay: Scene graph as prompt for continual learning on vqa task](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 1250–1259.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021. [Align before fuse: Vision and language representation learning with momentum distillation](#). *Advances in neural information processing systems*, 34:9694–9705.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Xin Zhao, and Ji-Rong Wen. 2023. [Evaluating object hallucination in large vision-language models](#). In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Zhizhong Li and Derek Hoiem. 2018. [Learning without forgetting](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. [Microsoft coco: Common objects in context](#). In *Computer Vision – ECCV 2014*, pages 740–755, Cham. Springer International Publishing.
- Zhiqiu Lin, Jia Shi, Deepak Pathak, and Deva Ramanan. 2021. [The CLEAR benchmark: Continual LEARNING on real-world imagery](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023. [Improved baselines with visual instruction tuning](#). *arXiv preprint arXiv:2310.03744*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024. [Visual instruction tuning](#). *Advances in neural information processing systems*, 36.
- Vincenzo Lomonaco and Davide Maltoni. 2017. [Core50: a new dataset and benchmark for continuous object recognition](#). In *Proceedings of the 1st Annual Conference on Robot Learning*, volume 78 of *Proceedings of Machine Learning Research*, pages 17–26. PMLR.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. [Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks](#). *arXiv preprint arXiv:1908.02265*.
- Andrea Madotto, Zhaojiang Lin, Zhenpeng Zhou, Seungwhan Moon, Paul Crook, Bing Liu, Zhou Yu, Eunjoon Cho, Pascale Fung, and Zhiguang Wang. 2021. [Continual learning in task-oriented dialogue systems](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7452–7467, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Michael McCloskey and Neal J Cohen. 1989. [Catastrophic interference in connectionist networks: The sequential learning problem](#). In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier.
- Seyed Iman Mirzadeh, Arslan Chaudhry, Dong Yin, Huiyi Hu, Razvan Pascanu, Dilan Gorur, and Mehrdad Farajtabar. 2022. [Wide neural networks forget less catastrophically](#). In *International conference on machine learning*, pages 15699–15717. PMLR.
- Giang Nguyen, Tae Joon Jun, Trung Tran, Tolcha Yalew, and Daeyoung Kim. 2019. [Contcap: A scalable framework for continual image captioning](#). *arXiv preprint arXiv:1909.08745*.
- Mavina Nikandrou, Lu Yu, Alessandro Suglia, Ioannis Konstas, and Verena Rieser. 2022. [Task formulation matters when learning continually: A case study in visual question answering](#). *arXiv preprint arXiv:2210.00044*.
- OpenAI. 2022. [Introducing ChatGPT](#).
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. 2024. [DINOv2: Learning robust visual features without supervision](#). *Transactions on Machine Learning Research*.
- Vicente Ordonez, Girish Kulkarni, and Tamara Berg. 2011. [Im2text: Describing images using 1 million captioned photographs](#). In *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc.
- Oleksiy Ostapenko, Timothee Lesort, Pau Rodríguez, Md Rifat Arefin, Arthur Douillard, Irina Rish, and Laurent Charlin. 2022. [Continual learning with foundation models: An empirical study of latent replay](#). In *Conference on lifelong learning agents*, pages 60–91. PMLR.
- Zi Qian, Xin Wang, Xuguang Duan, Pengda Qin, Yuhong Li, and Wenwu Zhu. 2023. [Decouple before interact: Multi-modal prompt learning for continual visual question answering](#). In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2953–2962.

- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. [Learning transferable visual models from natural language supervision](#). In *International conference on machine learning*, pages 8748–8763. PMLR.
- Vinay Venkatesh Ramasesh, Aitor Lewkowycz, and Ethan Dyer. 2022. [Effect of scale on catastrophic forgetting in neural networks](#). In *International Conference on Learning Representations*.
- Roger Ratcliff. 1990. [Connectionist models of recognition memory: constraints imposed by learning and forgetting functions](#). *Psychological review*, 97(2):285.
- Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. 2017. [icarl: Incremental classifier and representation learning](#). In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010.
- Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. 2018. [Overcoming catastrophic forgetting with hard attention to the task](#). In *ICML*, pages 4555–4564.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. [Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565.
- Tejas Srinivasan, Ting-Yun Chang, Leticia Pinto Alva, Georgios Chochlakakis, Mohammad Rostami, and Jesse Thomason. 2022. [Climb: A continual learning benchmark for vision-and-language tasks](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 29440–29453. Curran Associates, Inc.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. [Roformer: Enhanced transformer with rotary position embedding](#). *Neurocomputing*, 568:127063.
- Fan-Keng Sun, Cheng-Hao Ho, and Hung yi Lee. 2019. [Lamol: Language modeling for lifelong language learning](#). In *International Conference on Learning Representations*.
- Hao Tan and Mohit Bansal. 2019. [Lxmert: Learning cross-modality encoder representations from transformers](#). *arXiv preprint arXiv:1908.07490*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Gido M Van de Ven and Andreas S Tolias. 2018. [Generative replay with feedback connections as a general strategy for continual learning](#). *arXiv preprint arXiv:1809.10635*.
- Weiyao Wang, Du Tran, and Matt Feiszli. 2020. [What makes training multi-modal classification networks hard?](#) In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12695–12705.
- Zhen Wang, Liu Liu, Yiqun Duan, and Dacheng Tao. 2022. [Continual learning through retrieval and imagination](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8594–8602.
- Nan Wu, Stanislaw Jastrzebski, Kyunghyun Cho, and Krzysztof J Geras. 2022. [Characterizing and overcoming the greedy nature of learning in multi-modal deep neural networks](#). In *International Conference on Machine Learning*, pages 24043–24055. PMLR.
- Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. 2019. [Large scale incremental learning](#). In *Conference on Computer Vision and Pattern Recognition*.
- Jaehong Yoon, Saehoon Kim, Eunho Yang, and Sung Ju Hwang. 2020. [Scalable and order-robust continual learning with additive parameter decomposition](#). In *International Conference on Learning Representations*.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. 2017. [Continual learning through synaptic intelligence](#). In *International Conference on Machine Learning*, pages 3987–3995. PMLR.
- Xi Zhang, Feifei Zhang, and Changsheng Xu. 2023. [Vqacl: A novel visual question answering continual learning setting](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19102–19112.
- Yao Zhang, Haokun Chen, Ahmed Frikha, Yezi Yang, Denis Krompass, Gengyuan Zhang, Jindong Gu, and Volker Tresp. 2022. [Cl-crossvqa: A continual learning benchmark for cross-domain visual question answering](#). *arXiv preprint arXiv:2211.10567*.

A Experiments

A.1 Pretraining VL-Pythia

We pretrain all VL-Pythia models following the LLaVA-1.5 recipe (Liu et al., 2023). The only difference is that we skip the first stage for multimodal alignment as recent work (Karamcheti et al., 2024) has shown that the two-stage training can be redundant and the same performance can be achieved when omitting the first stage of training. Throughout VL pretraining, the vision encoder remains frozen, while the LLM and connector parameters are trained using the Adam optimizer (Kingma and Ba, 2014) with a batch size of 256 and a learning rate of $1e-3$. For all models, we used the same data to train UNITER and ViLT - COCO (Lin et al., 2014), SBU captions (Ordonez et al., 2011), Visual Genome captions (Krishna et al., 2017) and Conceptual Captions 3M (Sharma et al., 2018). We perform one epoch of pretraining and keep the final checkpoint.

A.2 Hyperparameters

Model	Setting	Batch Size	LR	EWC λ	FD γ
UNITER	Diverse Content	1024	$8e-5$	500	0.8
	Taxonomy Content	1024	$5e-5$	500	0.8
	Question Types	512	$5e-5$	20K	0.6
ViLT	Diverse Content	1024	$1e-5$	500	1
	Taxonomy Content	1024	$1e-5$	700	1
	Question Types	512	$8e-5$	10K	0.5
VL-Pythia	Diverse Content	128	$5e-5$	-	0.5
	Taxonomy Content	128	$5e-5$	-	0.5
	Question Types	128	$5e-5$	10K	0.5

Table 5: Selected hyperparameters.

We tune the hyperparameters using grid search based on the validation accuracy of a single task order. For VL-Pythia variants, we use the same hyperparameters for all model sizes, as we find them to perform reasonably well. For UNITER and ViLT, we keep the batch size, learning rate (LR), and EWC loss weight λ reported in prior work (Nikandrou et al., 2022). For the remaining values, we perform the following grid search: $lr \in \{1e-5, 5e-5, 8e-5, 1e-4\}$, EWC $\lambda \in \{500, 1K, 5K, 10K\}$, FD discount factor $\gamma \in [0.3, 1.0]$ with a step of 0.1.