

ConStruct-VL: Data-Free Continual Structured VL Concepts Learning*

James Seale Smith^{†1,2} Paola Cascante-Bonilla^{1,3} Assaf Arbelle⁴
 Donghyun Kim^{1,4} Rameswar Panda^{1,4} David Cox^{1,4} Diyi Yang⁵
 Zsolt Kira² Rogerio Feris^{1,4} Leonid Karlinsky^{†1,4}
¹MIT-IBM Watson AI Lab ²Georgia Institute of Technology ³Rice University
⁴IBM Research ⁵Stanford University

Abstract

Recently, large-scale pre-trained Vision-and-Language (VL) foundation models have demonstrated remarkable capabilities in many zero-shot downstream tasks, achieving competitive results for recognizing objects defined by as little as short text prompts. However, it has also been shown that VL models are still brittle in Structured VL Concept (SVLC) reasoning, such as the ability to recognize object attributes, states, and inter-object relations. This leads to reasoning mistakes, which need to be corrected as they occur by teaching VL models the missing SVLC skills; often this must be done using private data where the issue was found, which naturally leads to a data-free continual (no task-id) VL learning setting. In this work, we introduce the first Continual Data-Free Structured VL Concepts Learning (ConStruct-VL) benchmark¹ and show it is challenging for many existing data-free CL strategies. We, therefore, propose a data-free method comprised of a new approach of Adversarial Pseudo-Replay (APR) which generates adversarial reminders of past tasks from past task models. To use this method efficiently, we also propose a continual parameter-efficient Layered-LoRA (LaLo) neural architecture allowing no-memory-cost access to all past models at train time. We show this approach outperforms all data-free methods by as much as $\sim 7\%$ while even matching some levels of experience-replay (prohibitive for applications where data-privacy must be preserved).

1. Introduction

Recently, large Vision-and-Language (VL) models achieved great advances in zero-shot learning [4, 16, 16, 19, 23, 25, 40, 51, 58]. Pre-trained on hundreds of millions [40] or billions [45] of image-and-text pairs collected from the

*This work is supported by the Defense Advanced Research Projects Agency (DARPA) Contract No. FA8750-19-C-1001. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA.

[†]Equal contribution

¹Our code is publicly available at <https://github.com/jamessealesmith/ConStruct-VL>

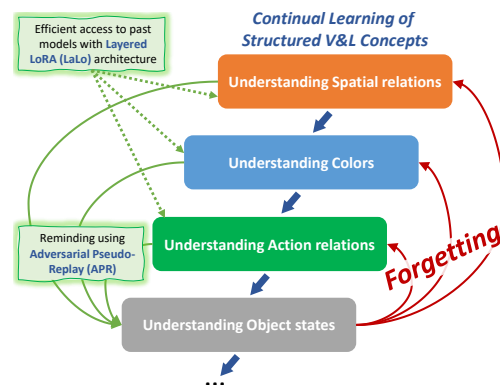


Figure 1. Illustration for the Continual Data-Free Structured VL Concepts Learning (ConStruct-VL). Structured VL Concept (SVLC) understanding skills are added / refined over time, with Adversarial Pseudo-Replay (APR) effectively countering catastrophic forgetting using our Layered-LoRA (LaLo) architecture’s ability of efficient no-memory-cost access to past task models.

web, these VL models have demonstrated remarkable capabilities in understanding (recognizing [16, 40], detecting [63], segmenting [57], etc.) objects appearing in images or videos [39], thus moving beyond the previous paradigm of fixed object classes to open-vocabulary models.

Despite these great advances, several recent works [52, 67] have found VL models to be brittle with respect to understanding Structured VL Concepts (SVLCs) - non-object textual content such as object attributes, states, inter-object relations (e.g. spatial, interaction, etc.), and more. Naturally, it is important to alleviate these shortcomings, as this lack of understanding of SVLCs can lead to embarrassing errors on the VL model’s part for many applications, such as analysis of multi-modal social networks data, multi-modal chat, multi-modal document analysis, and more. Importantly, in most of these applications: (i) different errors made by the model surface sequentially in the course of time, each time identifying another SVLC ‘skill’ missing in the model; (ii) typically, to fix the errors, one would attempt additional fine-tuning of the model on a data collected from the source on which the errors were observed (hoping to avoid catastrophic forgetting of the previous model im-

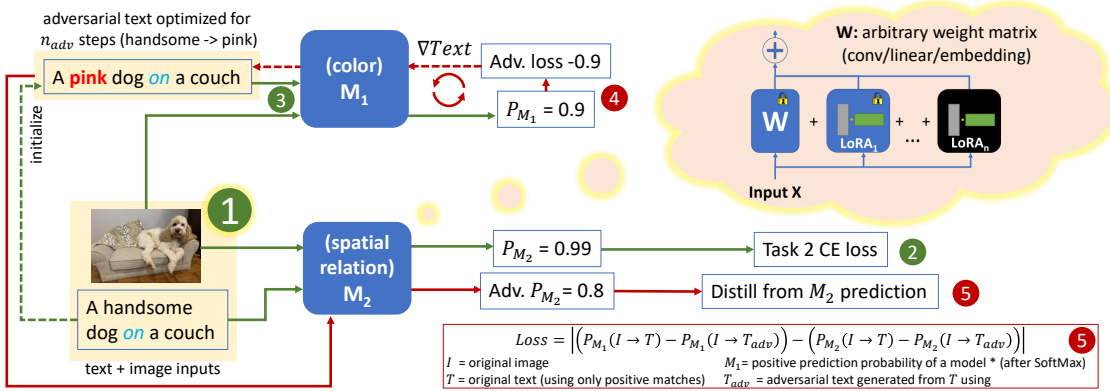


Figure 2. Layered-LoRA (LaLo) architecture and Adversarial Pseudo-Replay (APR), illustrated for a two-task ConStruct-VL sequence teaching ‘color’ \rightarrow ‘spatial relation’ understanding skills. (1) A text+image pair is entering the model M_2 currently training to understand spatial relations; M_2 is an extension of the model M_1 (previously trained to understand color) via adding a layer of low-rank (LoRA) adapters to every parametric function of M_1 while keeping M_1 frozen. (2) M_2 produces positive prediction probability P_{M_2} (after softmax) and applies CE loss. (3) Text+image pair are passed through M_1 (done without reloading in LaLo). (4) M_1 produces a positive prediction probability P_{M_1} ; using $-P_{M_1}$ as adversarial loss, we perform a sign-gradient attack on the text input (after tokenizer+first embedding layer). The produced text is expected to be adversarial in the direction of the SVLC used to train M_1 , color in this case. We, therefore, expect some embedded text tokens to become ‘colored’ inconsistently with the image. (5) We use the produced adversarial sample to distill from M_1 into M_2 via the proposed pseudo-replay positive prediction probability drop preserving loss.

provement rounds); (iii) data sources where errors are observed are typically private and cannot be passed between these fine-tuning tasks. We, therefore, ask the question: Can we sequentially add SVLC skills to multi-modal VL models, in a privacy-preserving manner? This leads to the *data-free continual learning problem* (e.g., [55, 56]) cast in the multi-modal VL domain.

To this end, we introduce the first *Continual Data-Free Structured VL Concepts Learning (ConStruct-VL) multi-modal benchmark*, built on top of the popular Visual Genome [21] and Visual Attributes in the Wild [37] datasets using the protocol proposed in VL-Checklist [67], and show it is challenging for many existing data-free CL strategies, including recent SOTA prompting-based CL methods [55, 56]. We then offer a novel data-free CL method, leveraging the multi-modal nature of the problem, to effectively avoid forgetting. We propose the concept of Adversarial Pseudo-Replay (APR), that (as opposed to the previous pseudo-replay works [5, 46, 60]) generates *negative* examples to past task models conditioned on the current batch data. For continual VL training we generate negatives in one of the modalities by making it inconsistent with the other modality via a series of adversarial attack steps utilizing past models (see Fig. 2 for an example). Intuitively, generating (inconsistent) negatives (APR) is easier than generating (consistent) positives (pseudo-replay). Also, past task model attacks are likely to generate inconsistencies corresponding to their tasks, thus leading to reduced forgetting when we use the generated negative samples to distill from the past task models the drop in prediction probabilities after adversarial examples are applied (Fig. 1). To use the proposed APR technique we need the ability to efficiently invoke past models at training time. We, therefore, pro-

pose a Layered-LoRA (LaLo) continual learning neural architecture utilizing layered parameter-efficient (low-rank) residual adapters, supporting invocation of any of the past task models on any given training batch at no additional memory cost (without the need to reload these past models into memory). Moreover, our proposed architecture can be collapsed to the original model size by collapsing all the adapters into their adapted parametric functions, thus supporting inference on the final model at no additional cost.

Contributions: (i) we propose the challenging ConStruct-VL benchmark, and show that existing data-free CL methods struggle in this setting; (ii) we propose a new concept of Adversarial Pseudo-Replay (APR) specifically designed for multi-modal continual learning, alongside a Layered-LoRA (LaLo) architecture allowing invoking any of the past task models efficiently without reloading and having no additional inference time or parameters cost; and (iii) we demonstrate significant improvements (over 6.8% increase in final accuracy and $\times 5$ smaller average forgetting) of the proposed approach compared to all the popular data-free CL baselines, as well as some amounts of experience replay.

2. Related Work

Vision-and-Language (VL) Models. VL models pre-trained on large-scale noisy image-text data with contrastive loss (e.g., CLIP [40] and ALIGN [16]) show remarkable performance in many zero-shot downstream tasks. Some of the methods improve image-text alignment by focusing on region features from the off-shelf object detectors [4, 25, 51] or use end-to-end cross-attention layers with additional objective functions such as image-text matching and masked language modeling [19, 23, 24, 58], or filtering noisy captions (e.g., BLIP [23]). Some VL models [8–10, 28, 59]

utilize additional properties of language structure. DeCLIP [28] uses textual nearest-neighbor to find additional positives. CyClip [10] learns geometrically consistent representations in the image and text embedding space. However, recent studies [52, 67] have demonstrated the limitation of all of these models in learning SVLC reasoning skills. Even the most recent and sophisticated VL methods like CyClip [10] still struggles with SVLC understanding (Tab. 1). In our work, we study ways for continual improvement of VL models in SVLC reasoning tasks and propose the first Continual Data-Free Structured VL Concepts Learning benchmark to facilitate this study in future works.

Continual Learning. Continual learning aims to continuously adapt a model to non-stationary data distributions. Prior works can be categorized by their method for avoiding catastrophic forgetting [33]. Regularization-based methods (e.g., [17, 20, 29, 64]) introduce extra regularization terms in the objective functions when learning a new task. LwF [29] uses soft labels for the previous task to avoid forgetting. EWC [20] estimates the importance of parameters and uses this for per-parameter weight decay. Rehearsal-based methods (e.g., [2, 3, 12, 18, 34, 38, 41, 43, 53]) use a data buffer to store or generate some samples of the previous tasks. These samples are replayed with new task data. However, in many cases storing samples may not be allowed due to privacy or copyright concerns. Architecture-based methods (e.g., [1, 26, 44, 55, 56, 61]) isolate model parameters for each task. Recently, prompt-based continual learning methods (e.g., L2P [56], DualPrompt [55], S-Prompt² [54]) outperformed rehearsal-based methods without using a replay buffer. Additional detailed related work can be found in [6]. All of the above methods target uni-modal (e.g. vision only) CL setting, while we focus on multi-modal VL setting in this work. REMIND [11] proposed Continual VQA tasks with latent replay, but unlike ours requires storing compressed training data. CLiMB [48] proposed CL adaptation to coarsely different VL tasks, such as VQA, NLVR, SNLIVE, VCR, and assumed knowledge of the evaluated task-id at inference time. In contrast, we target continual adaptation to a fine-grained sequence of SVLC reasoning skills, with no task-id knowledge during inference.

Parameter-Efficient Adaptation. Several approaches [7, 32, 49] have been proposed for efficient model fine-tuning using fewer parameters. Either fine-tuning a subset of the model [50, 62], adding adapters [13, 36, 42], low-rank adapters (LoRA [15]), or learning prompts [22, 27, 68]. UNIPELT [32] proposes a unified framework for NLP that subsumes the method of adapters, prompt tuning, and LoRA with gating mechanism. VL-ADAPTER [49] presents a unified multi-task learning framework for VL tasks evaluating the use of adapters, Hyperformer [31], and Compactor [65] for VL tasks. In this work, we propose an ef-

ficient CL neural architecture built on the concept of layering low-rank adapters on a VL model effectively allowing replaying past task models at no additional memory cost.

3. Method

3.1. Tasks

In this work we target adapting VL models to an arbitrary sequence of SVLC reasoning tasks. Recently, several datasets (e.g., [52] and [67]) have been proposed for these tasks, given their importance. Each task \mathcal{T}^i is defined by a set of image I and text T pairs, and a binary ground-truth function $\psi^i(T, I) \in \{0, 1\}$ indicating whether the text of the pair corresponds to the image content or not (Fig. 3a). The difference between the tasks is the type of SVLC that they are ‘teaching’. For example, there could be tasks teaching the understanding of color, material, size, state, or other object attributes by explicitly having words related to these attributes contained inside the task’s texts T (Fig. 3a). Other tasks could teach skills for understanding positional or transitive action relations between multiple objects being present in the image and interacting with each other. In practical applications, tasks could teach understanding the presence of unwanted attributes or relations, such as nudity, profanity, violence, etc. Alternatively, they could teach understanding positive, but personal, aspects including love, affection, or any type of personal preference from private captioned photo collections. As their nature, these types of task data (both images and text) could be very sensitive and private information, prohibiting maintaining any of these task’s data between tasks training (e.g. tasks data could be hosted by different institutions, or privately belong to different users of a social network). This naturally leads to the *continual data-free setting with no task-id*, which is exactly the setting on which we focus. Note that methods that require detecting to which task a given input belongs at inference time are not reasonable here, as it is difficult to estimate what SVLC skill is needed from free-form text and a mixture of skills may be needed.

3.2. Model

Our model architecture is based on BLIP [23] (an openly available VLP with strongest support for SVLC understanding as we show in Tab. 1). BLIP is comprised of separate text and image encoders followed by a cross-attention decoder admitting both streams and followed by a classifier (for SVLC reasoning). Thanks to its modular architecture, the model supports both zero-shot retrieval and SVLC reasoning tasks. Our proposed method of model adaptation by sequentially improving the model’s SVLC understanding skills maintains this support leaving the zero-shot retrieval pipeline intact and unchanged.

3.2.1 Layered-LoRA (LaLo) Architecture

Our model admits a text-and-image pair (T, I) and is comprised of four parts: (i) image encoder $e_I = \mathcal{E}_I(I)$; (ii)

²S-Prompt is designed for a different setting than ours.

text encoder $e_T = \mathcal{E}_T(T)$; (iii) cross-attention decoder $d_{I,T} = \mathcal{D}(e_T, e_I)$; and (iv) binary classification head $\mathcal{C}(d_{I,T}) \in \{0, 1\}$. The final model is, therefore:

$$\mathcal{M}(T, I) = \mathcal{C}(\mathcal{D}(\mathcal{E}_T(T), \mathcal{E}_I(I))) \quad (1)$$

The decoder \mathcal{D} is an extension of the text encoder \mathcal{E}_T , with cross-attention blocks inserted between each two consecutive self-attention blocks as well as self-attention block parameters being shared between \mathcal{D} and \mathcal{E}_T . Each of the cross-attention blocks of \mathcal{D} also admits the embedded image tokens e_I as additional hidden states and employs them in the standard cross-attention computation of the prefix-tuning type (meaning output size is the same as input, and only hidden states of the text tokens are passed between the blocks). Each of the \mathcal{E}_T , \mathcal{E}_I , and \mathcal{D} networks is comprised of a mix of non-parametric functions (here referring to any kinds of data norms, e.g. *LayerNorm*, as non-parametric since we have them frozen) and two types of parametric functions: linear (e.g. *torch.nn.Linear* in PyTorch [35]) and embedding (e.g. *torch.nn.Embedding*). Roughly, each of those functions $\mathcal{F}_k^*(x)$ (where $*$ is either *lin* or *emb*) is parameterized by a weight matrix \mathcal{W}_k so that:

$$\mathcal{F}_k^{lin}(x) = \mathcal{W}_k \cdot x \quad (2)$$

$$\mathcal{F}_k^{emb}(x) = EMB(x; \mathcal{W}_k) \quad (3)$$

where EMB is the embedding operator picking the respective columns of \mathcal{W}_k assuming x is a stream of integers. Also, we disregard possible bias terms of the linear functions, as if they are present we keep them frozen. Assuming a sequence of n tasks $[\mathcal{T}^i]_{i=1}^n$, we train a sequence of models $[\mathcal{M}^i]_{i=1}^n$ where the weights \mathcal{W}_k^i of the model \mathcal{M}^i employed in the task \mathcal{T}^i training are parameterized as:

$$\mathcal{W}_k^i = \mathcal{W}_k^{i-1} + \mathcal{A}_k^i \cdot \mathcal{B}_k^i \quad (4)$$

where \mathcal{W}_k^0 are the weights of the base model before training any task, and, following the idea proposed in [15] for efficient LLM fine-tuning using Low-rank Residual Adapters (LoRA), \mathcal{A}_k^i and \mathcal{B}_k^i are (low) rank r matrices of sizes $m \times r$ and $r \times l$ respectively (with $m \times l$ being the dimensions of \mathcal{W}_k^i). These LoRA adapters can be applied efficiently as:

$$\mathcal{F}_k^{i,lin}(x) = \mathcal{F}_k^{i-1,lin}(x) + \mathcal{A}_k^i \cdot (\mathcal{B}_k^i \cdot x) \quad (5)$$

$$\mathcal{F}_k^{i,emb}(x) = \mathcal{F}_k^{i-1,emb}(x) + \mathcal{A}_k^i \cdot EMB(x; \mathcal{B}_k^i) \quad (6)$$

During the training of task \mathcal{T}^i , we freeze all previous tasks' (and base's) model parameters: $\forall k, \{\mathcal{W}_k^j\}_{j=0}^{i-1}$, and only learn current task LoRA adapters: $\forall k, \{(\mathcal{A}_k^i, \mathcal{B}_k^i)\}$.

There are several interesting things to note about the proposed LaLo architecture: (i) as opposed to [49], who evaluated the use of LoRA for (single task) VL models finding it ineffective, we also add our LoRA adapters to all the

layers of the image encoder, and not only to the text encoder/decoder as done in [49]; (ii) as opposed to [66] who attach a small side network only to the *output* of the adapted model (also requiring task-id in continual setting), our LoRA adapters are added to all parametric functions inside the model and affect all the intermediate computations; (iii) at the end of training, we expect our last model \mathcal{M}^n to have acquired all the SVLC understanding skills taught by the task sequence $[\mathcal{T}^i]_{i=1}^n$ with minimal forgetting (Sec. 3.2.2), and hence, at inference, all the LoRA adapters can be folded back into the original weight matrices by simple summation leading to *zero* additional inference cost; (iv) with rank r kept low, the number of extra parameters added by all the LoRA adapters is very low (e.g. 2.8% in our experiments); finally, (v) we keep our classifier $\mathcal{C}^i = \mathcal{C}$ shared across all task models - in practice, we only train it for the first task \mathcal{T}^1 and keep it frozen for all later tasks - this reduces some model capacity and could be improved in future work.

3.2.2 Adversarial Pseudo-Replay (APR)

One of the benefits of our proposed LaLo architecture is the ability to efficiently invoke any past tasks model \mathcal{M}^i by stopping the iterative (on j) computation of $\mathcal{F}_k^{j,*}(x)$ at $j = i$. This gives us an effective tool for battling forgetting by allowing us to revisit any past model as part of a single training batch computation without the need to reload any weights (infeasible for modern model sizes). We use this neat potential for our data-free APR multi-modal continual learning strategy to make our current model \mathcal{M}^i training on task \mathcal{T}^i not forget the 'lessons' previously 'taught' by tasks $\{\mathcal{T}^j\}_{j=1}^{i-1}$. In APR, we use task \mathcal{T}^i positive (T, I) pairs ($\psi^i(T, I) = 1$) for adversarially simulating negatives with respect to past task models, and enforcing a loss requiring the current model's positive prediction probability to drop in the presence of the simulated negatives by a similar amount as it would drop for the past task model.

Next, we describe how to simulate a negative sample for a past task model $\mathcal{M}^j = (\mathcal{E}_T^j, \mathcal{E}_I^j, \mathcal{D}^j, \mathcal{C})$, with $j < i$, given a text and image pair $(T, I) \in \mathcal{T}^i$, s.t. $\psi^i(T, I) = 1$. Let the adversarial gradient be defined as:

$$\nabla_*^j(T, I) = -\frac{\partial}{\partial_*} POS(\mathcal{C}(\mathcal{D}^j(\mathcal{E}_T^j(T), \mathcal{E}_I^j(I)))) \quad (7)$$

where POS is the positive prediction probability of the classifier \mathcal{C} following *softmax* (taking it before softmax would not regularize the negative prediction probability), and $*$ can be either image I or text T (since T is a discrete output of a tokenizer, when computing adversarial samples using text we freeze the text embedding layer for all models and compute adversarial gradients with respect to its continuous output). Without loss of generality, we describe how to obtain an adversarial text T_{adv}^j with respect to a past model \mathcal{M}^j from a text T corresponding image I in task \mathcal{T}^i . Similarly, we could compute an adversarial image I_{adv}^j . In prac-

tice, we use adversarial text in all of our experiments, which leads to less forgetting and better performance (Sec. 4.2). We initialize $T_{adv}^j = T$ and update T_{adv}^j in each adversarial step via the sign-gradient adversarial attack:

$$T_{adv}^j \leftarrow T_{adv}^j + \lambda_{adv} \cdot \text{sign}(\nabla_T^j(T_{adv}^j, I)) \quad (8)$$

where λ_{adv} is the step size, and we use n_{adv} update steps. Note that this process is multi-modal as the adversarial text T_{adv}^j is computed with respect to the image I , originally a positive pair of the starting text T . Some qualitative examples of generated T_{adv}^j are in Fig. 3b. We compute all the T_{adv}^j for all positive pairs in the training batch simultaneously. Having obtained the adversarial texts, we define the following APR loss \mathcal{L}_{adv}^j w.r.t. any past model $j < i$ as:

$$\delta^k(T, T_{adv}^j, I) = \text{POS}(\mathcal{M}^k(T, I)) - \text{POS}(\mathcal{M}^k(T_{adv}^j, I)) \quad (9)$$

$$\mathcal{L}_{adv}^j = \text{AVG}_{(T,I)} |\delta^j(T, T_{adv}^j, I) - \delta^i(T, T_{adv}^j, I)| \quad (10)$$

where POS is the positive prediction probability of the classifier head \mathcal{C} (after softmax), and AVG_* computes average over $*$. The intuition behind \mathcal{L}_{adv}^j is that T_{adv}^j will move the text T to contain words corresponding to SVLC taught by the past task \mathcal{T}^j , but not present on the image I . For example, a text describing a ‘handsome dog on a couch’ as part of *spatial relations* training task, might be modified to include the color ‘pink’ as adversarial modification by the past task model that has been taught the *color* SVLC (Fig. 2). Maintaining similar response to this modification in the text between the normalized prediction probabilities of the current model \mathcal{M}^i and the normalized prediction probabilities of the past model \mathcal{M}^j will likely reduce forgetting, which is indeed verified by the strong effects of this loss in our experiments and ablations. An additional intuition here is that APR is easier to do than positive pseudo-replay proposed in past works for the uni-modal continual setting [5, 46, 60]. Indeed, producing adversarial effects, that is modifying the text to make it reference things not on the image, is easier than modifying the text to contain words corresponding to past task, but actually appearing on the image. This is verified in our ablations (Sec. 4.2).

Finally, combined with the standard cross-entropy loss for the task \mathcal{T}^i with respect to its ground truth annotations $\psi^i(T, I)$ we arrive at the final loss \mathcal{L}^i used for training \mathcal{M}^i :

$$\mathcal{L}^i = \text{AVG}_{(T,I)} (\text{CE}(\mathcal{M}^i(T, I), \psi^i(T, I))) + \rho \cdot \sum_{j=0}^{i-1} \mathcal{L}_{adv}^j \quad (11)$$

where ρ is a parameter. In our experiments we also evaluate a simplified variant of this loss where we use only the adversarial loss w.r.p.t. the previous model \mathcal{L}_{adv}^{i-1} instead of the last summation term in Eq. (11).

3.3. ConStruct-VL Benchmark

To the best of our knowledge, no prior works have proposed a benchmark to evaluate the continual adaptation of VL models, teaching them reasoning over SVLCs (such as

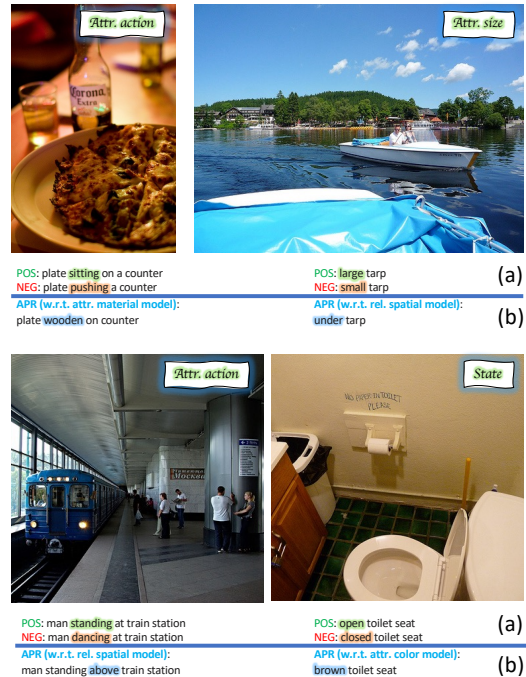


Figure 3. Some examples of the ConStruct-VL tasks data and negatives generated by our proposed APR technique w.r.t. past task models. (a) Showing image + POS text + NEG text triplets for (*in-transitive*) *action*, *size*, and *state* tasks; (b) Showing the adversarial text T_{adv} generated w.r.t. previously trained *material*, *spatial relations*, or *color* models, modifications highlighted

understanding of attributes, relations, and states) in a continual manner without task-id knowledge at test time (as we want the final model to understand all the concepts taught without informing it on which concept it is evaluated). Therefore, we propose the first such Continual Data-Free Structured VL Concepts Learning (ConStruct-VL) benchmark. We construct ConStruct-VL from the two popular publicly available VL datasets: Visual Genome (VG) [21] and Visual Attributes in the Wild (VAW) [37], building on protocols proposed in the VL-Checklist [67].

Specifically, we use 7 SVLC subsets defined by [67] for measuring VL models’ performance on understanding relations (spatial, inter-object transitive actions), attributes (size, color, material, intransitive single object actions), and object state. Each subset is constructed from triplets including a single image and two texts, one of them corresponding to the image content (positive) and another non-correspondent (negative) - manipulated by changing the value of a single SVLC (a specific concept for each subset). Some examples of the triplets for different subsets are provided in Fig. 3a. The subsets are separately constructed from each of the VG and VAW datasets image-text pairs, with all 7 SVLC subsets available for VG and only the 5 attribute and object state subsets defined for VAW. The sizes of the subsets range between 1K to 31K in VG and 3K to 64K in VAW. In addition to VG and VAW, we also define the combined VG+VAW dataset uniting the corresponding sub-

sets of the same nature. The resulting VG+VAW dataset has the aforementioned 7 subsets with sizes ranging between 5K to 75K. We split each subset to train / val / test according to the 80 / 10 / 10 percent rule guaranteeing no overlap between the splits. In ConStruct-VL for each of the 3 datasets (VG, VAW, VG+VAW) each subset is treated as a separate continual learning task, thus forming 7 total available tasks used to construct the task sequences in our experiments.

3.4. Baselines

For a fair comparison, we compare to the state-of-the-art *data-free* continual learning baselines *which do not rely on task id*. These methods are: Learning-without-Forgetting (LwF) [29], a method which distills soft model predictions from a prior, frozen model copy to the current model; Elastic-Weight-Consolidation (EWC) [20], a method which penalizes changes to model parameters weighted by an importance matrix³; and L2, which is EWC without the importance-weighting matrix (which can sometimes outperform EWC in the presence of pre-training [47]).

We also compare to state-of-the-art *prompting for continual learning* methods L2P [56] and DualPrompt [55]. These approaches continuously learn a set of prompts which are matched with input data in an instance-wise fashion. To our surprise, we found that these methods struggled in our benchmark, despite a heavy search for prompt depth and length. Our intuition is that these methods rely on encoder knowledge that is strongly related to the continual learning task sequence and, unlike image classification, our proposed task sequences requires substantial semantic modifications to the transformer encoders.

Finally, we also compare to continual fine-tuning of the full pre-trained model over the task sequence using a small learning rate (referred to as CFT). CFT was found to be more effective than the prompting methods L2P/DualPrompt. We also compare to two variants of CFT: (i) CFT Frozen Head (CFT-FH), where the classifier head is frozen after task 1 (as done in our approach); and (ii) CFT Linear Probing (CFT-LP), where the encoders are frozen and only the classifier head is learned.

4. Experiments

Implementation Details. Our code is built on top the official BLIP [23] repository. We start all compared methods (including ours) from BLIP pre-trained weights. Our \mathcal{E}_I image encoder is ViTB/16 and our text encoder \mathcal{E}_T is BERT with 12-layer encoder and 768 hidden size. As explained in Sec. 3.2.1, our decoder \mathcal{D} is an extension of the text encoder \mathcal{E}_T , with cross-attention layers added between every two self-attention layers, each receiving \mathcal{E}_I -encoded image tokens as additional input. Our binary classifier \mathcal{C} is a 2-

³The importance matrix is of the same dimension of the model parameter weights and is calculated using the diagonal terms of the Fisher-Information matrix w.r.t. past task loss functions and data.

Table 1. **Results (%) on 7 Task VG+VAW.** Fin. Acc. is the no-task-id performance of the final model on all tasks jointly; A_N is the accuracy averaged over tasks; F_N is the average forgetting; and N_{prm} is the % of trainable parameters. For L2P [56] and DP [55] we swept over multiple % of N_{prm} and report the highest result. [†] *measured using original publicly available codes + weights*

	Method (N_{prm})	Fin. Acc. (↑)	A_N (↑)	F_N (↓)
VLPs	BLIP [†] [23]	77.8	72.33	—
	CLIP [†] [40]	62.99	59.93	—
	CyCLIP [†] [10]	57.76	56.36	—
	ALBEF [†] [24]	71.2	67.9	—
	OSCAR [†] [25]	71.04	67.69	—
	ViLT [†] [19]	70.35	68.41	—
	UNITER [†] [4]	67.23	65.29	—
	LXMERT [†] [51]	68	64.39	—
	TCL [†] [58]	70.29	65.42	—
Baselines	CFT (100%)	73.91	80.76	6.94
	CFT-FH (99.7%)	77.43	83.73	6.14
	CFT-LP (0.3%)	64.85	71.32	12.26
	L2 (100%)	76.36	85.59	4.57
	LwF (100%)	79.39	86.11	4.28
	EWC (100%)	73.29	81.91	6.13
	L2P [56] ($\leq 5\%$)	66.96	76.46	6.7
	DP [55] ($\leq 5\%$)	58.59	68.65	9.52
Ours	Ours ^{rand} (2.8%)	78.6	86.19	1.5
	Ours ⁱ⁻¹ (2.8%)	86.2	90.37	1.4
	Ours (2.8%)	85.4	90.88	0.75

layer MLP with 768 hidden size. The entire model is comprised of 224M trainable parameters. All baselines including our method were tuned once on the VG *color* \rightarrow *material* \rightarrow *size* attributes task sequence (i.e., they were tuned on a *single* sequence in this particular order) using validation data, and afterward the same parameters were used throughout. We used WD 0.05 and an initial LR of 1e-5 and cosine scheduler in case of full model FT, and initial LR 1.25e-3 in all experiments involving the low-rank adapters (so in all instances of our proposed method). For our method, we used low-rank adapters rank $r = 16$, $n_{adv} = 10$, $\lambda_{adv} = 0.01$, and $\rho = 0.2$. Max of 12 epochs were used to train each task, in practice all compared methods converged earlier. For all the compared methods, for each task in the sequence, the val set of the current training task was used to automatically choose the best epoch.

Metrics. We employ 4 metrics to evaluate all models in all the experiments. The Final Accuracy is the accuracy of the final model (after training on all tasks in the sequence) evaluated on the test data of all the tasks without knowledge of task id. It is arguably the most interesting metric for practical applications. The A_N averages the accuracy also over models corresponding to intermediate task sequences arising in the middle of the full sequence (it is commonly higher than final accuracy as it also averages over shorter and easier task sequences), F_N provides the average forgetting [2, 14, 30] (average over performance drops of consec-

Table 2. **Results (%) on 7 Task VG and 4 Task VAW**. Fin. Acc. is the no-task-id performance of the final model on all tasks jointly; A_N is the accuracy averaged over tasks; F_N is the average forgetting; and N_{prm} is the % of trainable parameters. For L2P [56] and DP [55] we swept over multiple % of N_{prm} and report the highest result.

(a) 7 Task VG					(b) 4 tasks VAW				
	Method (N_{prm})	Fin. Acc. (\uparrow)	A_N (\uparrow)	F_N (\downarrow)		Method (N_{prm})	Final Acc. (\uparrow)	A_N (\uparrow)	F_N (\downarrow)
VLPs	BLIP [23]	79.58	74.24	–	VLPs	BLIP [23]	74.73	78.3	–
	CFT (100%)	77.99	83.5	6.94		Baselines	CFT (100%)	80.15	83.88
Baselines	CFT-FH (99.7%)	81.30	85.51	5.88	CFT-FH (99.7%)		80.87	84.08	6.60
	CFT-LP (0.3%)	65.13	72.70	11.73	CFT-LP (0.3%)	65.46	74.98	15.28	
	L2 (100%)	77.40	87.12	5.21	L2 (100%)	83.72	86.01	2.95	
	LwF (100%)	77.01	85.24	5.80	LwF (100%)	81.20	85.38	6.26	
	EWC (100%)	72.78	82.61	7.23	EWC (100%)	86.00	87.26	1.05	
	L2P ($\leq 5\%$)	60.88	69.23	17.07	L2P ($\leq 5\%$)	62.27	68.68	16.02	
	DP ($\leq 5\%$)	65.55	66.94	7.05	DP ($\leq 5\%$)	60.07	71.98	8.07	
	Ours	Ours ^{rand} (2.8%)	77.2	88.62	3.7	Ours	Ours ^{rand} (2.8%)	85.43	88.78
Ours ⁱ⁻¹ (2.8%)		85.87	92.29	1.83	Ours ⁱ⁻¹ (2.8%)		86.00	89.43	1.9
Ours (2.8%)		86.88	93	1.12	Ours (2.8%)		87.50	89.77	1.19

utive task models in the sequence), and finally N_{prm} gives the % of trainable parameters for each model.

4.1. Comparison to Baselines

For our main evaluations, we use the task sequences formed from the SVLC subsets defined for the VG, VAW, and VG+VAW datasets as part of our ConStruct-VL benchmark as explained in Sec. 3.3. Specifically, we use the *rel. spatial* \rightarrow *attr. size* \rightarrow *attr. material* \rightarrow *rel. action* \rightarrow *attr. color* \rightarrow *object state* \rightarrow *attr. action* 7-task sequence on VG and VG+VAW, *attr. action* \rightarrow *object state* \rightarrow *attr. color* \rightarrow *attr. size* 4-task sequence on VAW, and use *attr. color* \rightarrow *attr. material* \rightarrow *attr. size* \rightarrow *rel. action* \rightarrow *object state* \rightarrow *attr. action* for our ablation study. The task order was set arbitrarily, and we evaluate the effect of tasks order in Sec. 4.2. As VAW is a ‘clean’ dataset, being comprised of pairs of explicit object and attribute annotations (as opposed to VG where annotations were collected in natural language form as unconstrained sentences), we make our VAW task sequence harder by excluding the *attr. material* task, which has the highest performance in the original pre-trained BLIP that all compared methods use, as the starting point. The results of our method and comparisons to baselines for all the tested datasets and task sequences are provided in Tab. 1 (VG+VAW), Tab. 2a (VG) and Tab. 2b (VAW).

VL Pre-trained (VLP) models. We evaluate multiple openly available VLPs, including the very popular CLIP [40] and BLIP [23] on our longest and largest VG+VAW task sequence in Tab. 1. We use their original implementations and pre-trained weights. All of these models are pre-trained on with millions of images, and yet, as noticed in [52] and [67], without additional training are brittle with respect to understanding SVLCs which are the focus of our ConStruct-VL benchmark. From Tab. 1 we clearly see that BLIP has the highest out-of-the-box performance on our tasks among all the VLPs (and in general, being a more recent released model), hence we choose it as a starting point for all our training CL baselines as well as for our method.

For the same reason, we use only BLIP to represent VLPs in smaller VG only and VAW only evaluations. Interestingly, naively fine-tuning BLIP in a continual setting (CFT baseline in Tab. 1) reduces its final accuracy below the original non-fine-tuned BLIP by 3.88%. Also interestingly, our proposed approach improves the final accuracy of the CFT and the original BLIP by 12.28% and by 8.4% respectively. Similar strong gains in the final accuracy are observed for our method w.r.t. original BLIP and CFT in VAW and VG experiments in Tab. 2b and in Tab. 2a respectively.

Our model variants. In all of Tab. 1, Tab. 2a, and Tab. 2b we compare to 3 variants of our method. *Ours* refers to our full approach detailed in Sec. 3. *Oursⁱ⁻¹* refers to a simplified version of our approach that performs the APR w.r.t. only the previous task model, which in turn incorporates the knowledge of all previous tasks in a good way thanks to our proposed technique of reducing forgetting (apparent from the comparable performance of this variant compared to full approach). *Ours^{rand}* refers to a variant of our approach that employs the same LaLo architecture and adversarial loss \mathcal{L}_{adv}^j (defined in Eq. (10)), but generates the adversarial texts by random sampling of other texts in the batch. Although this approach obtains lower gains than our full method or *Oursⁱ⁻¹*, it still competes strongly with most of the baselines, highlighting the benefits of our LaLo architecture and the proposed adversarial loss \mathcal{L}_{adv}^j .

CL baselines comparison. We compare to a wide variety of baselines (detailed in Sec. 3.4), also including the very recent data-free visual CL SOTA prompting-based methods L2P [56] and DP [55] (using their respective methods PyTorch adaptation to VL, which, as detailed in our supplementary material, was verified to reproduce their vision-only results on their respective original papers benchmarks and was thoroughly tuned on the 3 tasks same as all the compared methods). As can be seen from Tab. 1, Tab. 2a, and Tab. 2b, our method has significant advantages over all of the baselines, especially apparent in $\times 5$ reduction in the




Figure 4. Analysis of our approach versus replay of stored training data with A_N on the y-axis and data stored per task on the x-axis.

Table 3. Ablations - 6 tasks VG+VAW. Fin. Acc. is the performance of the final model on all tasks without task-id knowledge; A_N is the accuracy averaged over tasks; F_N is the average forgetting; all methods have $N_{param} = 2.8\%$ trainable parameters.

	Method	Fin. Acc. (\uparrow)	A_N (\uparrow)	F_N (\downarrow)
Ablations	LaLo	73.68	78.7	6.21
	LWF LaLo	70.83	74.86	5.72
	Pos. Adv.	76.3	82.15	5.8
	Pos. Adv. ^{<i>i-1</i>}	80.98	84.84	5.19
	Adv. Image	81.2	86.12	4.22
Ours	Ours ^{<i>rand</i>}	82.71	86.97	2.25
	Ours ^{<i>i-1</i>}	84.69	87.73	2.29
	Ours	87.11	89.69	1.55

Table 4. Task order ablation - mean over all 3-tasks orders with confidence intervals. Final Accuracy measures the performance of the final model on all tasks without task-id knowledge; A_N is the accuracy averaged over tasks; F_N is the average forgetting.

Method	Fin. Acc. (\uparrow)	A_N (\uparrow)	F_N (\downarrow)
CFT	84.64 \pm 1.3	88.09 \pm 1.75	6.32 \pm 1.16
Ours	89.6 \pm 0.29	90.56 \pm 1.32	1.23 \pm 0.27

forgetting metric F_N and 6.8% to 5.5% gains in final accuracy (VG+VAW and VG) with respect to the strongest baseline. Smaller 1.5% and 2.5% gains in final accuracy and A_N (w.r.t. strongest baseline) on VAW can likely be attributed to VAW being very clean data constructed from combining explicit object + attribute annotations, as opposed to VG whose annotations were collected as free-form natural language texts. Another interesting aspect is that our model uses significantly less parameters for adapting the base model, making it parameter efficient, and yet, as we show in our ablations Sec. 4.2, the reduced number of trained parameters is clearly not the source of our methods greatly reduced forgetting. As revealed there, the source of reduced forgetting is certainly the proposed APR technique building on our efficient LaLo architecture.

4.2. Ablations

To explore the different aspects of what makes the proposed approach successful, we use a reduced size VG+VAW sequence of 6 tasks, removing the large *rel. spatial* task for quicker experiments. We present the results of the ablation study in Tab. 3, exploring the following variations of our method: (i) **LaLo** - evaluates our proposed Layered-LoRA architecture on ConStruct-VL without employing the proposed APR, as we can see the reduced number of trained parameters in LaLo alone is not sufficient to alleviate for-

getting in ConStruct-VL; (ii) **LWF LaLo** - explores if the proposed APR technique could be replaced by simple distillation from past models by applying LWF, as can be seen distillation does not help and even reduces the result due to reduced model plasticity; (iii) **Pos. Adv.** - simulates *positive* pseudo-replay by reversing the adversarial gradient computed for APR in our method and thus trying to generate samples more positive for past tasks, as we can see this significantly under-performs our proposed APR, supporting the intuition that generating negatives for past tasks is much easier; (iv) **Pos. Adv.^{*i-1*}** - *positive* pseudo-replay could be significantly improved if only using single past task model, still strongly under-performing our APR; Finally, (v) **Adv. Image** - verifies that generating adversarial texts in APR is more beneficial than generating adversarial images, this is intuitive as an image is a more complex and less structured signal and hence harder to effectively manipulate.

We examine the effects of tasks order by evaluating our approach on *all* possible task orders of *color*, *material*, and *size* attributes task of VG+VAW, the results are presented in Tab. 4. Clearly, our method is very robust to task order with only minimal changes in the performance metrics.

Finally, although our proposed method and the ConStruct-VL setting are data-free, it is still interesting to compare our results to increasing amounts of experience replay (prohibitive when data privacy and/or copyright need to be preserved) applied to the same sequences of tasks. We present this comparison in Fig. 4. We observe that our method is comparable to the highest amounts of experience replay, also improving over some replay amounts.

5. Conclusions

We have presented a benchmark and a method for exploring a very practical and important task of continual improvement of VL models in understanding Structured VL Concept. Our proposed approach builds upon the proposed efficient LaLo architecture combined with the APR technique to significantly reduce forgetting especially designed for LaLo and the challenging multi-modal continual VL setting. We have shown the advantages of the proposed approach w.r.t. strong CL baselines applied to the same setting demonstrating its effectiveness, as well as exploring different aspects of its design in our ablation studies and other experiments. Potential future work includes exploring applications of the proposed LaLo and APR techniques to uni-modal settings, as well as exploring potential additional improvements discussed in the text.

References

- [1] Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3366–3375, 2017. [3](#)
- [2] Arslan Chaudhry, Marc’ Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-GEM. In *International Conference on Learning Representations*, 2019. [3](#), [6](#)
- [3] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’ Aurelio Ranzato. Continual learning with tiny episodic memories. *arXiv preprint arXiv:1902.10486*, 2019. [3](#)
- [4] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer, 2020. [1](#), [2](#), [6](#)
- [5] Yoojin Choi, Mostafa El-Khamy, and Jungwon Lee. Dual-teacher class-incremental learning with data-free generative replay. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3543–3552, 2021. [2](#), [5](#)
- [6] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385, 2021. [3](#)
- [7] Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin Chen, Chi-Min Chan, Weize Chen, et al. Delta tuning: A comprehensive study of parameter efficient methods for pre-trained language models. *arXiv preprint arXiv:2203.06904*, 2022. [3](#)
- [8] Andreas Furst, Elisabeth Rumetshofer, Viet Tran, Hubert Ramsauer, Fei Tang, Johannes Lehner, David Kreil, Michael Kopp, Günter Klambauer, Angela Bitto-Nemling, et al. Cloob: Modern hopfield networks with infoloob outperform clip. *arXiv preprint arXiv:2110.11316*, 2021. [2](#)
- [9] Yuting Gao, Jinfeng Liu, Zihan Xu, Jun Zhang, Ke Li, and Chunhua Shen. Pyramidclip: Hierarchical feature alignment for vision-language model pretraining. *arXiv preprint arXiv:2204.14095*, 2022. [2](#)
- [10] Shashank Goel, Hritik Bansal, Sumit Bhatia, Ryan A Rossi, Vishwa Vinay, and Aditya Grover. Cyclic: Cyclic contrastive language-image pretraining. *arXiv preprint arXiv:2205.14459*, 2022. [2](#), [3](#), [6](#)
- [11] Tyler L Hayes, Kushal Kafle, Robik Shrestha, Manoj Acharya, and Christopher Kanan. Remind your neural network to prevent catastrophic forgetting. In *European Conference on Computer Vision*, pages 466–483. Springer, 2020. [3](#)
- [12] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 831–839, 2019. [3](#)
- [13] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR, 2019. [3](#)
- [14] Yen-Chang Hsu, Yen-Cheng Liu, Anita Ramasamy, and Zsolt Kira. Re-evaluating continual learning scenarios: A categorization and case for strong baselines. *arXiv preprint arXiv:1810.12488*, 2018. [6](#)
- [15] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. [3](#), [4](#)
- [16] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021. [1](#), [2](#)
- [17] Heechul Jung, Jeongwoo Ju, Minju Jung, and Junmo Kim. Less-forgetting learning in deep neural networks. *arXiv preprint arXiv:1607.00122*, 2016. [3](#)
- [18] Nitin Kamra, Umang Gupta, and Yan Liu. Deep generative dual memory network for continual learning. *arXiv preprint arXiv:1710.10368*, 2017. [3](#)
- [19] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning*, pages 5583–5594. PMLR, 2021. [1](#), [2](#), [6](#)
- [20] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 2017. [3](#), [6](#)
- [21] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. [2](#), [5](#)
- [22] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021. [3](#)
- [23] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. *arXiv preprint arXiv:2201.12086*, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)
- [24] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021. [2](#), [6](#)
- [25] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu

- Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *European Conference on Computer Vision*, pages 121–137. Springer, 2020. 1, 2, 6
- [26] Xilai Li, Yingbo Zhou, Tianfu Wu, Richard Socher, and Caiming Xiong. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *International Conference on Machine Learning*, pages 3925–3934. PMLR, 2019. 3
- [27] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021. 3
- [28] Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. *arXiv preprint arXiv:2110.05208*, 2021. 2, 3
- [29] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. 3, 6
- [30] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS’17, pages 6470–6479, USA, 2017. Curran Associates Inc. 6
- [31] Rabeeh Karimi Mahabadi, Sebastian Ruder, Mostafa Dehghani, and James Henderson. Parameter-efficient multi-task fine-tuning for transformers via shared hypernetworks. *arXiv preprint arXiv:2106.04489*, 2021. 3
- [32] Yuning Mao, Lambert Mathias, Rui Hou, Amjad Almahairi, Hao Ma, Jiawei Han, Wen-tau Yih, and Madian Khabsa. Unipelt: A unified framework for parameter-efficient language model tuning. *arXiv preprint arXiv:2110.07577*, 2021. 3
- [33] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989. 3
- [34] Oleksiy Ostapenko, Mihai Puscas, Tassilo Klein, Patrick Jah-nichen, and Moin Nabi. Learning to remember: A synaptic plasticity driven framework for continual learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11321–11329, 2019. 3
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. 4
- [36] Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. Adapterfusion: Non-destructive task composition for transfer learning. *arXiv preprint arXiv:2005.00247*, 2020. 3
- [37] Khoi Pham, Kushal Kafle, Zhe Lin, Zhihong Ding, Scott Cohen, Quan Tran, and Abhinav Shrivastava. Learning to predict visual attributes in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13018–13028, 2021. 2, 5
- [38] Quang Pham, Chenghao Liu, and Steven Hoi. Dualnet: Continual learning, fast and slow. *Advances in Neural Information Processing Systems*, 34:16131–16144, 2021. 3
- [39] Rui Qian, Yeqing Li, Zheng Xu, Ming-Hsuan Yang, Serge Belongie, and Yin Cui. Multimodal open-vocabulary video classification via pre-trained vision and language models. *arXiv preprint arXiv:2207.07646*, 2022. 1
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 1, 2, 6, 7
- [41] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. icarl: Incremental classifier and representation learning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition*, CVPR’17, pages 5533–5542, 2017. 3
- [42] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. *Advances in neural information processing systems*, 30, 2017. 3
- [43] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. In *Advances in Neural Information Processing Systems*, pages 348–358, 2019. 3
- [44] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016. 3
- [45] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. 1
- [46] James Smith, Yen-Chang Hsu, Jonathan Balloch, Yilin Shen, Hongxia Jin, and Zsolt Kira. Always be dreaming: A new approach for data-free class-incremental learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9374–9384, October 2021. 2, 5
- [47] James Seale Smith, Junjiao Tian, Yen-Chang Hsu, and Zsolt Kira. A closer look at rehearsal-free continual learning. *arXiv preprint arXiv:2203.17269*, 2022. 6
- [48] Tejas Srinivasan, Ting-Yun Chang, Leticia Leonor Pinto Alva, Georgios Chochlakis, Mohammad Rostami, and Jesse Thomason. Climb: A continual learning benchmark for vision-and-language tasks, 2022. 3
- [49] Yi-Lin Sung, Jaemin Cho, and Mohit Bansal. VI-adapter: Parameter-efficient transfer learning for vision-and-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5227–5237, 2022. 3, 4

- [50] Yi-Lin Sung, Varun Nair, and Colin A Raffel. Training neural networks with fixed sparse masks. *Advances in Neural Information Processing Systems*, 34:24193–24205, 2021. [3](#)
- [51] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*, 2019. [1](#), [2](#), [6](#)
- [52] Tristan Thrush, Ryan Jiang, Max Bartolo, Amanpreet Singh, Adina Williams, Douwe Kiela, and Candace Ross. Winoground: Probing vision and language models for visiolinguistic compositionality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5238–5248, 2022. [1](#), [3](#), [7](#)
- [53] Gido M van de Ven, Hava T Siegelmann, and Andreas S Tolia. Brain-inspired replay for continual learning with artificial neural networks. *Nature communications*, 11(1):1–14, 2020. [3](#)
- [54] Yabin Wang, Zhiwu Huang, and Xiaopeng Hong. S-prompts learning with pre-trained transformers: An occam’s razor for domain incremental learning. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2022. [3](#)
- [55] Zifeng Wang, Zizhao Zhang, Sayna Ebrahimi, Ruoxi Sun, Han Zhang, Chen-Yu Lee, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, et al. Dualprompt: Complementary prompting for rehearsal-free continual learning. *arXiv preprint arXiv:2204.04799*, 2022. [2](#), [3](#), [6](#), [7](#)
- [56] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 139–149, 2022. [2](#), [3](#), [6](#), [7](#)
- [57] Mengde Xu, Zheng Zhang, Fangyun Wei, Yutong Lin, Yue Cao, Han Hu, and Xiang Bai. A simple baseline for zero-shot semantic segmentation with pre-trained vision-language model. *arXiv preprint arXiv:2112.14757*, 2021. [1](#)
- [58] Jinyu Yang, Jiali Duan, Son Tran, Yi Xu, Sampath Chanda, Liqun Chen, Belinda Zeng, Trishul Chilimbi, and Junzhou Huang. Vision-language pre-training with triple contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15671–15680, 2022. [1](#), [2](#), [6](#)
- [59] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. *arXiv preprint arXiv:2111.07783*, 2021. [2](#)
- [60] Hongxu Yin, Pavlo Molchanov, Jose M Alvarez, Zhizhong Li, Arun Mallya, Derek Hoiem, Niraj K Jha, and Jan Kautz. Dreaming to distill: Data-free knowledge transfer via deep-inversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8715–8724, 2020. [2](#), [5](#)
- [61] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*, 2017. [3](#)
- [62] Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. *arXiv preprint arXiv:2106.10199*, 2021. [3](#)
- [63] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14393–14402, 2021. [1](#)
- [64] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, 2017. [3](#)
- [65] Aston Zhang, Yi Tay, Shuai Zhang, Alvin Chan, Anh Tuan Luu, Siu Cheung Hui, and Jie Fu. Beyond fully-connected layers with quaternions: Parameterization of hypercomplex multiplications with $1/n$ parameters. *arXiv preprint arXiv:2102.08597*, 2021. [3](#)
- [66] Jeffrey O Zhang, Alexander Sax, Amir Zamir, Leonidas Guibas, and Jitendra Malik. Side-tuning: a baseline for network adaptation via additive side networks. In *European Conference on Computer Vision*, pages 698–714. Springer, 2020. [4](#)
- [67] Tiancheng Zhao, Tianqi Zhang, Mingwei Zhu, Haozhan Shen, Kyusong Lee, Xiaopeng Lu, and Jianwei Yin. Vl-checklist: Evaluating pre-trained vision-language models with objects, attributes and relations. *arXiv preprint arXiv:2207.00221*, 2022. [1](#), [2](#), [3](#), [5](#), [7](#)
- [68] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. [3](#)