

CLIP with Generative Latent Replay: a Strong Baseline for Incremental Learning

Emanuele Frascaroli*
emanuele.frascaroli@unimore.it

Aniello Panariello*
aniello.panariello@unimore.it

Pietro Buzzega
pietro.buzzega@unimore.it

Lorenzo Bonicelli
lorenzo.bonicelli@unimore.it

Angelo Porrello
angelo.porrello@unimore.it

Simone Calderara
simone.calderara@unimore.it

AlmageLab
University of Modena and Reggio
Emilia
Modena, Italy

Abstract

With the emergence of Transformers and Vision-Language Models (VLMs) such as CLIP, large pre-trained models have become a common strategy to enhance performance in Continual Learning scenarios. This led to the development of numerous prompting strategies to effectively fine-tune transformer-based models without succumbing to catastrophic forgetting. However, these methods struggle to specialize the model on domains significantly deviating from the pre-training and preserving its zero-shot capabilities. In this work, we propose **Continual Generative training for Incremental prompt-Learning**, a novel approach to mitigate forgetting while adapting a VLM, which exploits generative replay to align prompts to tasks. We also introduce a new metric to evaluate zero-shot capabilities within CL benchmarks. Through extensive experiments on different domains, we demonstrate the effectiveness of our framework in adapting to new tasks while improving zero-shot capabilities. Further analysis reveals that our approach can bridge the gap with joint prompt tuning. The codebase is available at <https://github.com/aimagelab/mammoth>.

1 Introduction

In real-world scenarios, data is not always readily available but arrives gradually and sequentially. Because of this, Continual Learning (CL) is becoming increasingly popular [4] as it mimics a production environment. The primary challenge in developing a neural network capable of continuously learning is the *catastrophic forgetting* [7] phenomenon, which describes the tendency of the model to replace previously acquired knowledge with that of

new data, becoming less proficient in tasks it has encountered in the past. With the advent of Transformers [8, 39] and Large Vision-Language Models (VLMs) [16, 33], the trend in Computer Vision is moving towards prompt learning [17, 29, 49, 50] and parameter efficient tuning [15], allowing the exploitation of available pre-trained models. Among this category of methods, the most widespread approach is that of *soft prompts*, in which a few learnable vectors are employed to adapt the model. For instance, CoOp [50] learns a context prompt, which is concatenated to the textual name of the class and fed to the CLIP [53] text encoder. Similarly, VPT [17] learns a visual prompt for each task, which is concatenated at each layer of a Vision Transformer (ViT) [8] pre-trained on ImageNet. Finally, CoCoOp [49] extends CoOp by conditioning the prompts on the input image.

One key capability of Vision-Language Models is to perform zero-shot classification on unseen classes. This becomes crucial in a Continual Learning scenario, where such ability is not affected by Catastrophic Forgetting. Thus, improving upon the zero-shot performance of VLMs while learning in a continuous fashion presents a non-trivial challenge. Current prompt-based state-of-the-art approaches devise a learnable *prompt pool* and a query function to map a given image embedding to its corresponding key within the pool [18, 57, 42, 43, 44]. However, when exploiting a VLM as the backbone, the query-key matching strategy privileges prompts for already seen classes, thus failing to generalize to unseen ones. Notably, only a few approaches [42, 46, 48] can perform zero-shot classification after the adaptation phase; however, they still exhibit limited improvement over CLIP. In this work, we propose a new parameter-efficient CL approach that addresses such a shortcoming. Specifically, we fine-tune CLIP on the available data while also transferring knowledge to future tasks, thus extending the performance of the frozen backbone.

Inspired by CoOp [50], we propose to learn class-specific prompts and keep the model’s visual and text encoders frozen. This technique allows the model to adapt to new domains while preserving CLIP’s zero-shot capabilities. Indeed, we employ a hybrid approach that leverages hand-crafted prompts (“a photo of a <CLS>”) for unseen classes and learned prompts for previously seen classes. However, as we show in Sec. 5, prompt learning alone is not enough to overcome the challenges of a CL scenario and cannot compete with current state-of-the-art CL methods. In this respect, we bridge the gap with joint training by exploiting generative replay. Unlike standard rehearsal methods [3, 9, 34], which store a subset of samples in a memory buffer, we eliminate the buffer and employ a generative model to learn the underlying distribution of input data. This approach offers two significant advantages: *i*) having access to the distributions of the past features, rather than just a subset of the input data and *ii*) ensuring data anonymity, thus meeting privacy constraints over data.

Moreover, instead of generating images like previous generative-replay methods [11, 36], we directly model the distribution of the latent representation of the data. As our framework only tunes prompts for CLIP’s text encoder, the visual embeddings of images do not change over the continual training phase. Thus, we can avoid generating images. Moreover, modelling the latent space instead of the input one dramatically improves the generative models’ computational requirements given the lower dimensionality of the latent representations *w.r.t.* raw images. Specifically, we adopt Variational Autoencoders (VAEs) [21] to model the visual embeddings. In Sec. 6, we compare various generative approaches, along with different prompt-learning techniques, to validate our choices.

We evaluate our framework, called **Continual Generative training for Incremental prompt-Learning** (CGIL), on various standard *class-incremental* CL benchmarks, showing state-of-the-art performance even on domains where zero-shot CLIP fails. Indeed, it overcomes the previous best performer by a wide margin (+11% on average). Inspired

by [46, 48], we devise an additional metric to current CL benchmarks that measures zero-shot performance on future tasks. We evaluate in this setting all competitors with zero-shot capabilities (*i.e.*, the ones adopting a VLM), showing the superiority of our prompting strategy. We remark on the following contributions:

- We propose CGIL, a CL framework for adapting a VLM to a CL scenario that leverages prompt-learning techniques and a generative replay strategy based on Variational Autoencoders.
- We devise a new metric for CL benchmarks that measures zero-shot performance on future tasks.
- Through extensive experiments, we demonstrate the validity of our approach, achieving state-of-the-art performance in the common class incremental scenario and over future evaluation.

2 Related Works

Continual Learning (CL) methods are designed to tackle the issue of *catastrophic forgetting* [27], which prevents the continuous transfer of previously acquired knowledge when data comes as a stream. Traditional CL approaches can be broadly classified in three main categories: i) regularization techniques [27, 24] prevent the most important parameters from drifting too far from the optimum; ii) architectural-based method allocate specific sets of parameters for each incremental task [26, 65]; iii) rehearsal-based methods adopt a small memory buffer to store past exemplars that are used for later replay [0, 6, 9, 30, 32, 34]. At the cost of bending the rules of continual learning, the latter models have been established as the state of the art when continuously training from scratch.

Recently, the advent of (large) pre-trained models based on the Vision Transformer (ViT) architecture [8, 33] has changed this paradigm, fostering the emergence of buffer-free alternatives [28, 37, 43, 44, 47] that can achieve minimal forgetting without compromising privacy. These approaches are designed to tackle *class-incremental learning* [38], whose goal is to continuously expand the set of recognizable classes with each incoming task.

While the class-incremental scenario is typically regarded as the most difficult for CL [0, 9, 38], it does not take into account the potential zero-shot capabilities of Visual Language Models (VLMs). In particular, recent works have shown that it is possible to tune VLMs while maintaining – or even improving – their zero-shot capabilities for a single task [6, 19, 45], or multiple tasks [46, 48]. Similarly to the latter, we aim to encourage the continuous adaptation of a VLM to the incoming data without neglecting its zero-shot capabilities.

3 Preliminaries

Contrastive Language-Image Pre-Training (CLIP). We employ CLIP [33] as our VLM, which is composed of a visual encoder $E_{vis}(\cdot)$ and a text encoder $E_{txt}(\cdot)$. These encoders are trained on image-text pairs by aligning their latent representations via a contrastive objective. To use CLIP for classification, the input image x is fed to the visual encoder, yielding the visual representation $z_{vis} = E_{vis}(x)$. In parallel, different text prompts are crafted by embedding every class label (*e.g.*, “cat”, “car”) into a template such as “a photo of a

<CLS>”. These prompts are tokenized and fed to the text encoder, obtaining a text representation z_{txt}^j for each class y^j . Finally, the posterior probabilities of each class are computed as the cosine similarity (noted as $\langle \cdot, \cdot \rangle$) between their respective text representations and the visual representation. Formally:

$$p(y^j|x) = \frac{\exp(\langle z_{txt}^j, z_{vis} \rangle / \tau)}{\sum_{j=1}^{|\mathcal{Y}|} \exp(\langle z_{txt}^j, z_{vis} \rangle / \tau)}, \quad (1)$$

where τ is the temperature parameter fixed by CLIP and \mathcal{Y} represents the set of classes.

Prompt-learning. With prompt-learning, it is possible to fine-tune large models (*e.g.*, VLMs) efficiently and with reduced computational cost. Among these methods, CoOp [50] stands out as particularly effective. In a nutshell, CoOp learns context vectors V that, when concatenated before the label token [CLS], construct the prompt for the CLIP text encoder. These context vectors are learned through gradient descent, with the standard cross-entropy objective applied to the similarity scores of the different classes.

4 Method

Problem setting. In Continual Learning (CL), a deep model $f(\cdot; \theta)$ parametrized by θ is presented with a sequence of tasks \mathcal{T}_i with $i := \{1, \dots, T\}$, where T denotes the number of tasks. Each t -th task provides N_t data entries, composing the task dataset $\mathcal{D}_t := \{x^{(n)}, y^{(n)}\}_{n=1}^{N_t}$ with label $y^{(n)} \in \mathcal{Y}_t$. Importantly, each task relies on a set of classes disjoint from others such that $\mathcal{Y}_i \cap \mathcal{Y}_j = \emptyset$ if $i \neq j$. The objective of CL is to minimize the empirical risk on all tasks:

$$\mathcal{L}_{\text{CL}} = \sum_{i=1}^T \mathbb{E}_{(x,y) \sim \mathcal{T}_i} [\mathcal{L}(f(x; \theta), y)], \quad (2)$$

where \mathcal{L} is the loss function (*e.g.*, the cross entropy for classification). Since the model observes one task at a time, tailored strategies are required to prevent catastrophic forgetting.

4.1 Continual Generative training for Incremental prompt-Learning

In Fig. 1, we present **Continual Generative training for Incremental prompt-Learning (CGIL)**, which comprises two main phases. First, we extract the image embeddings of all images within the current task with the CLIP Visual Encoder. With such features, we fit the distribution of the latent representation for each class with a VAE. In the second phase, inspired by [47] and [50], we learn the prompts for CLIP’s text encoder using feature vectors sampled by our VAEs. These two phases are repeated at each task to improve previously learned prompts with knowledge from subsequent tasks.

Learning the Distributions of Latent Representations. In the first phase of our approach, we aim to learn the underlying distribution of the latent representation of each class. We do so by extracting all the visual features for each image in the current task. Formally, given the samples \mathcal{D}_t , where t is the current task with the label set \mathcal{Y}_t , we obtain the visual features $z_{vis} = E_{vis}(x)$, $\forall x \in \mathcal{D}_t$. Subsequently, we fit $|\mathcal{Y}_t|$ VAEs (*i.e.*, the number of classes within the current task) to learn the distribution of each class $c \in \mathcal{Y}_t$. The objective function for the VAEs is the standard ELBO. Since we work solely in the latent space, we can discard

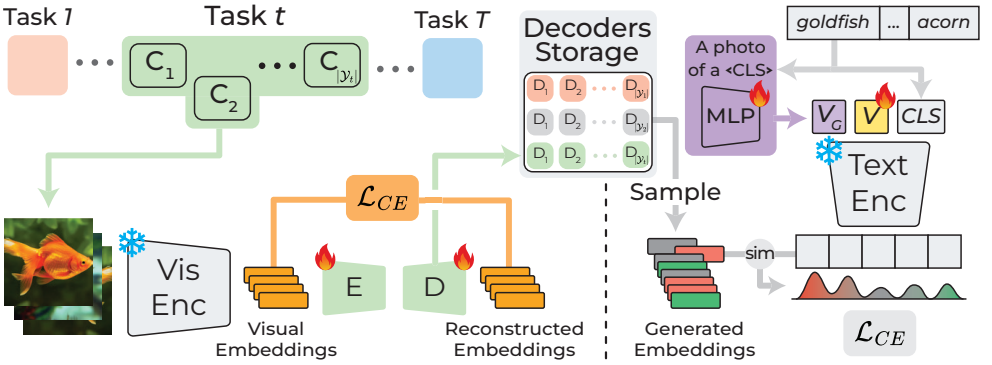


Figure 1: Training of the generative models (left) and prompt alignment (right). At each task t , we train a generative model (*i.e.*, VAE) for each class C_i in the task. While the decoders are stored, the encoders can be discarded after train. In the second phase, we perform prompt alignment between the features sampled from all the stored decoders up to task t and the text features of the learnable prompts V_G and V .

the visual encoder E_{vis} once we extract the visual feature for all samples. This approach significantly reduces our method’s computational costs. Once the VAEs are trained, we store their decoders and use them for each alignment phase.

Prompts Alignment. In the second phase, we build a synthetic dataset by generating feature vectors through all stored decoders. We then perform prompt tuning only on synthetic data by modeling the context words with continuous vectors. Specifically, we train *one* class-specific context V to capture fine-grained details regarding the classes and a shared Multi-Layer Perceptron (MLP) to recover cross-domain knowledge. The latter processes the embedding z_{txt} related to the hand-crafted prompt “a photo of a <CLS>”. The prompt \mathbf{t}_c for the class c that we feed the text encoder E_{txt} is:

$$\mathbf{t}_c = [V_G] [V] [CLS], \quad (3)$$

$$\text{where } V_G = \text{MLP}(E_{txt}(\text{“a photo of a <CLS>”})). \quad (4)$$

The posterior probability of the class c is obtained as in Eq. (1), using the text embeddings obtained with our prompts $z_{txt}^c = E_{txt}(\mathbf{t}_c)$. Since the visual embeddings are not obtained through the CLIP visual encoder but directly from the generated featured dataset, we avoid carrying out expensive forward steps through such an encoder. Both V and V_G are trained via gradient descent on synthetic data of all seen tasks. Thus, previously learned contexts are further fine-tuned, incorporating knowledge from subsequent tasks without incurring forgetting. During inference, we feed the image through the visual encoder and compute the posterior probability for each class.

Zero-shot Inference. When dealing with unseen classes (*i.e.*, zero-shot classification) can exploit the learned contexts for the classes the model has already seen. At the same time, we can still use handcrafted prompts (*e.g.*, “a photo of a <CLS>”) to classify unseen classes. This approach allows us to i) preserve the zero-shot capabilities of CLIP while adapting to the classes that arrive sequentially and ii) exploit past knowledge to exclude

Model	Imagenet-R	Cars-196	CUB-200	EuroSAT	ISIC	Avg.
Zero-shot CLIP [33]	81.95	64.99	50.52	53.32	26.59	55.47
LwF † [24]	19.09	23.24	16.73	25.13	33.06	23.45
GDumb † [32]	44.28	28.74	61.34	90.99	61.64	57.40
DER++ † [8]	56.66	53.66	74.62	93.08	65.68	68.74
L2P [44]	66.49	38.18	62.21	46.34	47.13	52.07
DualPrompt [43]	68.50	40.14	66.00	71.39	49.99	59.20
CODA-Prompt [37]	75.45	31.99	67.30	63.12	44.87	56.55
AttriCLIP [42]	87.39	75.63	58.28	72.33	28.26	64.38
SLCA † [47]	77.00	67.73	84.71	88.69	59.19	75.46
MoE Adapters [46]	90.67	77.76	64.98	80.56	34.52	69.70
CGIL	89.42	89.27	83.12	96.17	73.03	86.20

Table 1: The Final Avg. Accuracy on the tested benchmarks. † denotes methods that fine-tune the whole model, while other methods apply parameter-efficient techniques.

already seen classes from the classification problem, thus reducing interference and avoiding misclassification.

5 Experiments

Datasets. We evaluate our approach across a wide range of datasets with different levels of similarity *w.r.t.* the ImageNet pre-train [0, 31]. In particular, we test on:

- *Split Imagenet-R* [13], is a general-knowledge dataset frequently adopted in recent CL benchmarks [37, 43, 44, 47], with 200 classes split across 10 tasks.
- *Split Cars-196* [23] and *Split CUB-200* [40], are fine-grained classification datasets regarding car models and bird species, respectively. For both scenarios, the classes are split into 10 tasks.
- *Split EuroSAT* [11, 12], which features RGB satellite images and defines a land cover classification problem consisting of 5 binary tasks.
- *Split ISIC* [8], which includes images depicting 6 skin diseases equally split into 3 tasks.

Metrics. We chose the more challenging scenario of class incremental learning [38] (CIL), where there is no knowledge of the task to which data belong during evaluation. To quantitatively assess our model’s performance in this setting, we report the average accuracy of each task, computed after the end of the last training phase, that is, the **Final Average Accuracy** (FAA), also called *Last Accuracy*. Additionally, we adapt the **Transfer** metric from [46, 48], originally used across datasets, to a class incremental training regime to evaluate the zero-shot performance on future tasks. Specifically, let A_i^t be the CIL accuracy on the i -th task

CI-Transfer	Imagenet-R	Cars-196	CUB-200	EuroSAT	ISIC	Avg.
Zero-shot CLIP [53]	82.14	66.16	50.86	55.00	22.42	55.32
AttriCLIP [42]	85.75	73.98	54.07	59.69	24.14	59.53
MoE Adapters [46]	88.25	75.82	61.73	55.77	21.06	60.53
CGIL	86.71	78.80	66.34	71.52	48.18	70.31

Table 2: The Class Incremental Transfer on the tested benchmarks. Only methods with zero-shot capabilities (*i.e.*, with CLIP as a backbone) could be tested.

after being trained until task t , the **Class Incremental Transfer** is defined as:

$$\text{CI-Transfer} = \frac{1}{T-1} \sum_{t=1}^{T-1} \left(\frac{1}{T-t} \sum_{i=t+1}^T A_t^i \right). \quad (5)$$

Implementation Details. We employ Adam optimizer [24] with a learning rate of 0.0002 for 500 epochs to train our VAEs. Each VAE consists of an encoder and a decoder, comprising only 3 fully connected layers interleaved with LeakyReLU activations. The hidden and latent sizes are 512 and 256, respectively. During the prompt-learning phase, we use the Adam optimizer with a learning rate of 0.03, and the synthetic features generated for alignment are roughly 15K per class, mixed together and split into batches of size 128. We employ CLIP with the ViT-L/14 backbone for each model. Finally, all results reported are the average across 3 runs with different seeds, changing the composition of tasks. The standard deviations of these experiments are reported in the Supplementary Material, along with more technical details of the experiments.

Comparison Methods. We benchmark our model against several state-of-the-art prompt-tuning approaches, including L2P [44], DualPrompt [43], CODA-Prompt [57], and AttriCLIP [42]. Additionally, we assess models that fine-tune the entire architecture, namely LwF [24], GDumb [52], DER++ [3], and SLCA [47]. In addition to such methods, we integrate MoE Adapters [46] into our framework, a parameter-efficient approach originally designed to prevent zero-shot accuracy degradation across datasets. To ensure a fair comparison, we train all competing models, tuning their hyperparameters for optimal performance. We include the performance of zero-shot CLIP on our datasets as a baseline to emphasize the efficacy of prompt tuning methods over the frozen CLIP backbone. Additionally, AttriCLIP, Moe Adapters, and our CGIL are also evaluated on future tasks, measuring how their zero-shot capabilities are affected by incremental training.

5.1 Comparison with the State of the Art

In Tab. 1, we report the CIL final accuracies for all evaluated competitors across each benchmark. The last column shows the average of the performance of each method. Despite Zero-shot CLIP’s impressive performance on Split Imagenet-R and Split Cars-196, it fails in other domains, particularly in the medical field. Consequently, competitors that rely on CLIP as their backbone are heavily influenced by this limitation and exhibit a similar pattern. On the other hand, CGIL successfully addresses these CLIP-related issues, delivering top-tier performance in all scenarios. Considering average performance, our framework stands

	Imagenet-R	Cars-196	CUB-200	EuroSAT	ISIC	Avg.
CoOp (<i>Joint</i>)	89.24	89.89	82.52	96.25	73.57	86.29
CoOp (<i>Fine-tune</i>)	84.94	68.61	59.84	79.27	37.08	65.95
CGIL	89.42	89.27	83.12	96.17	73.03	86.20
Different generative approaches						
Multinomial Gaussian	83.89	82.59	80.06	85.70	51.89	76.83
Mixture of Gaussians	88.54	88.82	82.10	93.04	62.42	82.98
Diffusion Models	89.28	90.14	83.48	95.73	68.99	85.52
Variational Autoencoders (CGIL)	89.42	89.27	83.12	96.17	73.03	86.20
Different prompt-learning techniques						
Class-specific Context (w/o MLP)	89.09	88.96	83.06	95.59	72.21	85.78
Generated Context (only MLP)	89.41	88.91	82.82	95.69	70.79	85.52
Unified Context	89.02	88.08	81.50	95.47	70.18	84.85
Class-specific + Generated Contexts (CGIL)	89.42	89.27	83.12	96.17	73.03	86.20

Table 3: Ablative studies on CGIL regarding the generative models and the prompt-learning techniques. Performance is measured in terms of Final Average Accuracy.

out significantly, outperforming all other methods. Notably, it achieves a substantial lead (+11) over the best competitor, namely SLCA [47], while other prompt-based competitors fall behind.

Zero-shot performance. Tab. 2 displays the average CIL accuracies on unseen tasks, as measured by the CI-Transfer metric (Eq. (5)). This metric serves as an indicator of zero-shot capabilities. Thus, only methods with a VLM as their backbone could be evaluated. Nevertheless, similar patterns emerge, with CLIP and other competitors excelling in Split Imagenet-R and Split Cars-196 but struggling with other datasets. CGIL exhibits remarkable ability in leveraging both the zero-shot expertise of CLIP and its knowledge from previously learned tasks, achieving superior performance in nearly all benchmarks.

6 Analysis

To better validate the effectiveness of CGIL and its architectural design, we report additional experiments in Tab. 3. We evaluated vanilla CoOp (specifically, their class-specific version) under two distinct benchmarks: one trained jointly, *i.e.* without partitioning the dataset into tasks (*Joint*), and the other trained in the conventional CIL scenario (*Fine-tune*). For the latter, we implemented a minor adjustment: after each task, the learned prompts are frozen. This strategy, also employed by [67], helps prevent forgetting. Conversely, training all contexts in subsequent tasks can overwrite previous knowledge by altering learned prompts. The insights derived from the outcomes of these two baselines highlight the proficiency of CGIL in effectively bridging the gap between fine-tuning and joint training in prompt learning, as our method matches the performance of the *joint* approach. As other ablation studies indicate, this success is primarily attributed to the generative replay of latent features. On the other hand, the various prompting methods tested slightly affect the performance.

Different generative approaches. Along with Variational Autoencoders, we experimented with various generative models to identify the one that best complements our method. The most straightforward approach involves fitting a multivariate Gaussian distribution for each

class, which can be used for later sampling. As indicated in Tab. 3, this approach alone achieves state-of-the-art results (76.83 vs. 75.46 on Avg.), underscoring the effectiveness of our methodology regardless of the generator. However, exploiting a more powerful generative model considerably improves the effectiveness of the alignment procedure, demonstrating that the quality of the generated data is crucial. We first experimented with Mixture of Gaussians (MoG), fitted with expectation-maximization until convergence, which can better capture intra-feature relationships by combining multiple multivariate Gaussians. Then, we investigated Variational Autoencoders (VAEs) (CGIL) and Denoising Diffusion Probabilistic Models (DDPMs) [14], both saturating the required performance in generation. We train the DDPMs with the same hyper-parameters as VAEs, described in Sec. 5. Among the two, we stick to VAEs due to their faster training and reduced number of parameters *w.r.t.* DDPMs.

Different prompt-learning techniques. Our context comprises a class-specific context and a generated context; in the lower section of Tab. 3, we present the results with different choices. Specifically we evaluate: *i*) just one class-specific context, as in *CoOp Fine-tune*; *ii*) just one MLP generated context; *iii*) the unified-context variant of CoOp, *i.e.*, using multiple shared contexts. We also varied the number of contexts for each baseline. Interestingly, while the unified context benefited from this modification, the performance for the first two strategies did not improve with an increased number of contexts. Thus, we report only the results with a single context. The outcomes of these experiments fall shortly behind CGIL, showcasing that the main contributors to CGIL’s performance are the generative rehearsal and the alignment phase. This contributions are highlighted by the difference between the *CoOp fine-tune* and *Class-specific Context* lines in Tab. 3, as they share the same prompting strategy.

6.1 Discussion and Limitations

On the memory and computational costs. Despite our VAEs being relatively lightweight, with only half a million parameters after the encoders are discarded, CGIL may encounter memory constraints if the number of classes continues to grow, potentially restricting its applicability in scenarios with an extreme amount of classes. Compared with a standard rehearsal approach, the memory requirement of one decoder is comparable to a buffer of 14 RGB images of size 224×224 (*i.e.*, $\sim 2\text{MB}$). However, we note that these decoders are only needed for the training phase, while during inference, only the CLIP visual encoder and the embeddings z_{txt}^c of our prompts are required. Thus, the computational cost for inference is equivalent to a single pass through the visual encoder plus a matrix multiplication to obtain the similarity scores. This represents a significant advancement over other CL-prompting methods. Indeed, L2P, DualPrompt, and CODA-Prompt execute the forward pass on the image twice, while AttriCLIP computes both visual and textual embeddings during test time.

When considering the computational costs of training, our VAEs are notably lightweight, enabling the learning of latent embedding distribution to be remarkably swift, even without the need for a GPU. The alignment phase for prompt learning presents a higher level of complexity, as it performs gradient descent through the CLIP text encoder. Nonetheless, the duration of this phase can be controlled by adjusting the size of the synthetic dataset generated, thereby permitting flexible utilization and accommodation of time constraints.

Online CL Setting. We would like to highlight that CGIL may be classified under the category of Online CL methods [11, 15], as training images are fed only once to the visual encoder. While this requires temporarily storing the feature vectors of all samples from the

current task, the burden incurred is limited as the size of the features is significantly smaller than that of the images. These embeddings are employed to train our generative models and are subsequently discarded.

7 Conclusions

In this work, we introduce a novel framework, **Continual Generative training for Incremental prompt-Learning**, which allows the adaptation of Vision-Language Models to new tasks without incurring forgetting. Our model employs Variational Autoencoders to learn the latent distributions of input images, which enables the generation of synthetic latent embeddings. Such data is exploited to fine-tune our backbone efficiently through prompt learning. Specifically, CGIL learns textual contexts, keeping the backbone parameters frozen, thereby preserving the zero-shot capabilities of CLIP.

Our approach significantly outperforms previous state-of-the-art CL methods when tested across a broad spectrum of benchmarks in various domains. By introducing a new metric, the **Class Incremental Transfer**, we evaluate zero-shot performance on future tasks during training, demonstrating that our framework is the most effective at leveraging past knowledge to predict unseen classes. Further analysis validates our architectural choices and shows that CGIL bridges the gap with the performance of prompts learned jointly.

References

- [1] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient Based Sample Selection for Online Continual Learning. In *Advances in Neural Information Processing Systems*, 2019.
- [2] Matteo Boschini, Lorenzo Bonicelli, Pietro Buzzega, Angelo Porrello, and Simone Calderara. Class-incremental continual learning into the extended der-verse. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [3] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. *Advances in Neural Information Processing Systems*, 2020.
- [4] Lucas Caccia, Rahaf Aljundi, Nader Asadi, Tinne Tuytelaars, Joelle Pineau, and Eugene Belilovsky. New Insights on Reducing Abrupt Representation Change in Online Continual Learning. In *International Conference on Learning Representations*, 2022.
- [5] Sanjoy Chowdhury, Sayan Nag, and Dinesh Manocha. APoLLO : Unified adapter and prompt learning for vision language models. In *Empirical Methods in Natural Language Processing*, 2023.
- [6] Noel CF Codella, David Gutman, M Emre Celebi, Brian Helba, Michael A Marchetti, Stephen W Dusza, Aadi Kalloo, Konstantinos Liopyris, Nabin Mishra, Harald Kittler, et al. Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic). In *IEEE International Symposium on Biomedical Imaging*, 2018.

- [7] Yin Cui, Yang Song, Chen Sun, Andrew Howard, and Serge Belongie. Large scale fine-grained categorization and domain-specific transfer learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018.
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- [9] Sebastian Farquhar and Yarin Gal. Towards robust evaluations of continual learning. In *arXiv preprint arXiv: 1805.09733*, 2018.
- [10] Rui Gao and Weiwei Liu. Ddgr: Continual learning with deep diffusion-based generative replay. In *International Conference on Machine Learning*, 2023.
- [11] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Introducing eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. In *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*, 2018.
- [12] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2019.
- [13] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The many faces of robustness: A critical analysis of out-of-distribution generalization. *IEEE International Conference on Computer Vision*, 2021.
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 2020.
- [15] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- [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*, 2021.
- [17] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge J. Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *Proceedings of the European Conference on Computer Vision*, 2022.
- [18] Dahuin Jung, Dongyoon Han, Jihwan Bang, and Hwanjun Song. Generating instance-level prompts for rehearsal-free continual learning. In *IEEE International Conference on Computer Vision*, 2023.

- [19] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2023.
- [20] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- [21] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.
- [22] 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.
- [23] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *IEEE International Conference on Computer Vision and Pattern Recognition Workshops*, 2013.
- [24] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [25] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in Neural Information Processing Systems*, 2017.
- [26] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018.
- [27] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, 1989.
- [28] Mark D McDonnell, Dong Gong, Amin Parvaneh, Ehsan Abbasnejad, and Anton van den Hengel. Ranpac: Random projections and pre-trained models for continual learning. *Advances in Neural Information Processing Systems*, 2024.
- [29] Martin Menabue, Emanuele Frascaroli, Matteo Boschini, Enver Sangineto, Lorenzo Bonicelli, Angelo Porrello, and Simone Calderara. Semantic residual prompts for continual learning. In *Proceedings of the European Conference on Computer Vision*, 2024.
- [30] Matteo Mosconi, Andriy Sorokin, Aniello Panariello, Angelo Porrello, Jacopo Bonato, Marco Cotogni, Luigi Sabetta, Simone Calderara, and Rita Cucchiara. Mask and compress: Efficient skeleton-based action recognition in continual learning. In *International Conference on Pattern Recognition*, 2024.
- [31] Jaehoon Oh, Sungnyun Kim, Namgyu Ho, Jin-Hwa Kim, Hwanjun Song, and Se-Young Yun. Understanding cross-domain few-shot learning based on domain similarity and few-shot difficulty. *Advances in Neural Information Processing Systems*, 2022.
- [32] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *Proceedings of the European Conference on Computer Vision*, 2020.

- [33] 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*, 2021.
- [34] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. iCaRL: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2017.
- [35] 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.
- [36] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *Advances in Neural Information Processing Systems*, 2017.
- [37] James Seale Smith, Leonid Karlinsky, Vyshnavi Gutta, Paola Cascante-Bonilla, Donghyun Kim, Assaf Arbel, Rameswar Panda, Rogerio Feris, and Zolt Kira. Coda-prompt: Continual decomposed attention-based prompting for rehearsal-free continual learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2023.
- [38] Gido M van de Ven, Tinne Tuytelaars, and Andreas S Tolias. Three types of incremental learning. *Nature Machine Intelligence*, 2022.
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- [40] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds-200-2011 dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- [41] Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. A comprehensive survey of continual learning: Theory, method and application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [42] Runqi Wang, Xiaoyue Duan, Guoliang Kang, Jianzhuang Liu, Shaohui Lin, Songcen Xu, Jinhu Lü, and Baochang Zhang. Attriclip: A non-incremental learner for incremental knowledge learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2023.
- [43] 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. In *European Conference on Computer Vision*, 2022.
- [44] 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 conference on Computer Vision and Pattern Recognition*, 2022.

- [45] Hantao Yao, Rui Zhang, and Changsheng Xu. Visual-language prompt tuning with knowledge-guided context optimization. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2023.
- [46] Jiazuo Yu, Yunzhi Zhuge, Lu Zhang, Ping Hu, Dong Wang, Huchuan Lu, and You He. Boosting continual learning of vision-language models via mixture-of-experts adapters. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2024.
- [47] Gengwei Zhang, Liyuan Wang, Guoliang Kang, Ling Chen, and Yunchao Wei. SLCA: slow learner with classifier alignment for continual learning on a pre-trained model. In *IEEE International Conference on Computer Vision*, 2023.
- [48] Zangwei Zheng, Mingyu Ma, Kai Wang, Ziheng Qin, Xiangyu Yue, and Yang You. Preventing zero-shot transfer degradation in continual learning of vision-language models. *IEEE International Conference on Computer Vision*, 2023.
- [49] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2022.
- [50] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 2022.