

PIVOT: Prompting for Video Continual Learning

Andrés Villa^{1,2}, Juan León Alcázar², Motasem Alfarra², Kumail Alhamoud², Julio Hurtado³,
 Fabian Caba Heilbron⁴, Alvaro Soto¹, Bernard Ghanem²

¹Pontificia Universidad Católica de Chile, ²King Abdullah University of Science and Technology (KAUST),

³University of Pisa, ⁴Adobe Research

afvilla@uc.cl, {juancarlo.alcazar,motasem.alfarra,kumail.hamoud}@kaust.edu.sa
 julio.hurtado@di.unipi.it, caba@adobe.com, asoto@ing.puc.cl, bernard.ghanem@kaust.edu.sa

Abstract

Modern machine learning pipelines are limited due to data availability, storage quotas, privacy regulations, and expensive annotation processes. These constraints make it difficult or impossible to train and update large-scale models on such dynamic annotated sets. Continual learning directly approaches this problem, with the ultimate goal of devising methods where a deep neural network effectively learns relevant patterns for new (unseen) classes, without significantly altering its performance on previously learned ones. In this paper, we address the problem of continual learning for video data. We introduce PIVOT, a novel method that leverages extensive knowledge in pre-trained models from the image domain, thereby reducing the number of trainable parameters and the associated forgetting. Unlike previous methods, ours is the first approach that effectively uses prompting mechanisms for continual learning without any in-domain pre-training. Our experiments show that PIVOT improves state-of-the-art methods by a significant 27% on the 20-task ActivityNet setup.

1. Introduction

Modern Deep Neural Networks (DNNs) are at the core of many state-of-the-art methods in machine vision [10, 20, 21, 38, 44, 46] tasks. To achieve their remarkable performance, most DNNs rely on large-scale pre-training [11, 26, 36, 46], thereby enabling feature reuse in related downstream tasks [4]. However, adapting and fine-tuning a pre-trained DNN on a novel dataset commonly leads to *catastrophic forgetting* [16]. This phenomenon explains how the effectiveness of the fine-tuned DNNs drastically reduces in the original training distribution, in favor of increased performance on the downstream task.

This undesirable property has driven multiple research efforts in the area of Continual Learning (CL) [7, 8, 18, 25,




Figure 1. **Performance improvement by each PIVOT component.** We report the average accuracy on all tasks under the 10-task CIL on UCF101 and ActivityNet. We report the performance of basic CLIP, and then gradually equip it with other components: **S**patial **P**rompting, **M**emory **B**uffer, **M**ulti-modal **C**ontrastive Learning, **T**emporal **E**ncoder, to finally reach our proposed PIVOT method. The addition of each proposed component generally boosts the performance on both datasets. Stars represent the upper bound performance on each benchmark.

29], yielding techniques that enable fine-tuning a DNN on a sequence of tasks while mitigating the performance drop along the intermediate steps. One of the most challenging scenarios for the study of CL is Class Incremental Learning (CIL), where the labels and data are mutually exclusive between tasks, training data is available only for the current task, and there are no task identifiers on the validation step (*i.e.* task boundaries are not available at test time). Such a setup requires learning one model that, despite the continuous adaptation to novel tasks, performs well on all the seen classes. Recent advances in mitigating catastrophic forgetting rely on deploying episodic memory or regularization techniques [2, 8, 25, 37]. Nevertheless, most of this progress has been directed toward analyzing the catastrophic forgetting of DNNs in the image domain. The works of [34, 35, 47] introduced the CIL setup to the video domain, in particular the action recognition task. Unlike its image counterpart,

video CIL requires careful modeling of the temporal information, making it an even more challenging setup. Despite the success in small-scale datasets (like UCF101), state-of-the-art methods have shown limited effectiveness on more challenging video test-beds built upon larger action taxonomies, such as Kinetics and ActivityNet [47].

Currently, state-of-the-art methods for video CIL rely on temporal masking and feature distillation to mitigate catastrophic forgetting [34,47]. In this paper, we take inspiration from recent advances in large-scale DNNs for zero-shot image classification [36] and learnable prompts for continual learning [50], and we propose a novel strategy for CIL in the video domain. We show that a zero-shot baseline pre-trained in the image domain already outperforms the best CIL methods in the action recognition task*. Moreover, we show that this baseline can be significantly improved by enabling temporal reasoning and augmenting the modality encoders with a novel prompting strategy. As a consequence, our proposed method, PIVOT, outperforms every other baseline, setting a new state-of-the-art in all the 3 datasets included in the challenging vCLIMB benchmark for video CIL [47]. Figure 1 summarizes the performance improvements of our approach.

Notably and following the core ideas of the vCLIMB [47] benchmark, PIVOT does not rely on any in-distribution pre-training (a common feature of prompting methods for CL [45,50]). Rather, it leverages the vast and general visual knowledge contained in the CLIP visual encoder (trained on massive amounts of paired static images and text) and maps that knowledge into a feature space suitable for video understanding in a continual learning setup.

Contributions. This paper proposes PIVOT (Prompting for Video cOntinual learning), a novel strategy for continual learning in the video domains that leverages large-scale pre-trained networks in the image domain. Our work brings the following contributions: **(i)** We show that a multi-modal classifier (Video-Text) mitigates catastrophic forgetting while greatly increasing the final average CIL accuracy. **(ii)** We design the first prompt-based strategy for video CIL. Our approach leverages image pre-training to significantly mitigate forgetting when learning a sequence of video action recognition tasks. **(iii)** We conduct extensive experimental analysis to demonstrate the effectiveness of PIVOT. Our results show that PIVOT outperforms state-of-the-art methods in the challenging vCLIMB benchmark by 31%, 27%, and 17.2% in the 20-task setups of Kinetics, ActivityNet, and UCF101, respectively.

2. Related Work

Image Continual Learning. Different approaches to tackle Continual Learning have been proposed over the last

*A similar result for image domain CIL is outlined in the concurrent work Thengane *et al.* [45]

few years. These methods can broadly be classified into three categories: regularization, memory, and parameter isolation methods [9]. Regularization-based approaches penalize changes to weights associated with previous tasks. Among these methods, there is a first sub-group that attempted to minimize the differences in weights that are relevant to previous tasks. There are several techniques to determine the relevance of each weight, including the Fisher matrix [25], the gradient [2,55], the uncertainty [12], among others [3,40]. An alternative to previous approaches is to use distillation over previous task representations while training on the new one [17,28,42]. While regularization methods can mitigate forgetting, they usually underperform in the video domain, as their extension to temporal data remains an open research topic [47].

Memory-based methods store a subset of training samples from past tasks to be used in the future as replay [8,14] or regularization [7,29]. Alternatively, some methods train generative models to approximate the image distribution of past tasks. Synthetic images drawn from the generative model are then used as replay [1,19,27,41]. While showing high effectiveness in the image domain, these methods are not directly applicable to the video domain for two main reasons. First, due to the higher memory usage when storing video samples and high complexity of generating videos. Second, due to the associated challenge of selecting the key video segments to store in memory [47].

On the other hand, parameter-isolation methods aim to overcome the distribution shift by learning isolated parameters for each task [22]. While some of these methods clone the model when a new task arrives [15,39], others prefer to use masks [30,52] or conditioning a trainable knowledge base [13,23] to select relevant information for the current input. A critical problem with these approaches is identifying the corresponding weights for each input. Given that only a sub-group of weights is specialized for a task, the task-id is needed during training and validation to identify the corresponding group.

Adapting General Visual Knowledge. Instead of continually learning reusable representations [23,32], some approaches use fixed large-scale pre-trained models as knowledge bases [31,45,50]. The large amount of data used to train these models helps them generalize to multiple concepts, making it possible to apply zero-shot classification for continual learning [45] without fine-tuning. Other works have augmented these pre-trained models with learnable prompts [49,50], extracting better representations, significantly improving performance, and reducing forgetting in image CIL. However, these methods do not consider the temporal nature of videos, so they are not directly applicable. Although some works [33,54] have explored transferring image domain knowledge of [45] to video domain, [6] has demonstrated that CLIP pre-training on image-text pairs

is more transferable than a pre-training on noisy video-text pairs. In this work, we take a first step in leveraging large-scale pre-trained networks, *i.e.* CLIP, and equip them with temporal reasoning and learnable prompts to mitigate forgetting in video CIL.

Video Continual Learning. Only a few works have addressed the unique challenges of video continual learning by proposing video-specific strategies to mitigate forgetting. In the case of [34], the authors proposed to encode the temporal information through temporal-based knowledge distillation. A second work [47] resolves the complexity of storing videos in memory by employing a temporal consistency loss, which reduces the number of frames stored. Despite the progress, such methods do not scale well on real-world action recognition benchmarks such as Kinetics [24] and ActivityNet [5]. In this paper, we extend the ideas presented in previous works and propose PIVOT, which exploits an image-language pre-trained model (CLIP) and extends it temporally. PIVOT enables significant performance gains on both small and large-scale video CIL benchmarks.

3. Prompting for Video Continual Learning

Our main objective is to leverage the knowledge available in large-scale image-language models for video CIL. That is, we want to learn how to adapt large visual-language DNNs available in the image domain into a single end-to-end network that is effective in video CIL. Furthermore, we want to make sure that such a transfer of knowledge involves no fine-tuning of the image encoder, and requires minimal extensions over the original network. Such constraints will ultimately favor less forgetting in the CIL setup, as the key parameters will suffer less drift during fine-tuning on different tasks.

Our approach, PIVOT, is divided into three major components. **(i)** We adapt a large-scale image-language pre-trained model, CLIP [36], to extract visual information from video data. **(ii)** We aggregate the temporal information extracted from several frames within a video via a transformer encoder. **(iii)** We reinforce the learned knowledge from each task by deploying learnable tokens, *i.e.* prompts, to prevent the forgetting of previously learned tasks. Despite its multiple components, we train PIVOT in an end-to-end fashion. We highlight that this training process is lightweight since the image and text encoders are frozen. Before we delve into the details of each of our model components, we give a brief preliminary on the Video Class Incremental Learning (VCIL) problem and the multi-modal contrastive classifier proposed in CLIP, as it plays a key role in our approach.

3.1. Problem Definition

We follow the recent vCLIMB benchmark [47] training a single model $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$, *i.e.* a DNN parameterized

with θ , that learns to predict an action label $y \in \mathcal{Y}$ given an input video $x \in \mathcal{X}$. The model learns from a sequence of tasks $\{(X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n)\}$ where $X_i \subseteq \mathcal{X}$ and $Y_i \subseteq \mathcal{Y} \forall i$. In the class-incremental setup, we assume that different tasks present the model with a different set of labels, *i.e.* $Y_i \cap Y_j = \emptyset \forall i \neq j$. We define Acc_i as the average classification accuracy of all the tasks up to and including task i . Our objective is to train a single model f_θ that maximizes the average accuracy Acc , which is the average over all observed tasks.

3.2. Multi-modal Contrastive Classifier for Video Continual Learning

Our method builds on top of the image-language pre-trained CLIP encoder. CLIP [36] demonstrated that textual information promotes the learning of transferable visual features. PIVOT re-purposes the CLIP encoder for video CIL by using its multi-modal contrastive classifier as a video feature generator. We use CLIP’s language encoder to extract text embeddings from the labels of all the available classes, and its image encoder to extract visual embeddings for arbitrary frame data. These paired embeddings allow us to perform a pseudo-classification by looking for tuples that minimize the distance between visual and textual embeddings. By keeping the CLIP’s encoder frozen and forwarding video related data (in the visual and language streams), we allow the transfer of its extensive knowledge into the video domain in a zero-shot fashion.

We observe three advantages in Multi-modal Contrastive Learning (MCL) for the continual learning setup. First, this approach is parameter-free. Hence, our classifier would not suffer from *any* forgetting when learning different tasks. This mitigates the need to deploy a large replay buffer to maintain the performance of the classifier on previously learned tasks. Second, this approach directly leverages the frozen CLIP encoder in the video CIL setup. We empirically find that, despite the simplicity of this approach and the domain gap between image and video data, a zero-shot CLIP significantly outperforms all continual learning methods in the literature on both Kinetics and ActivityNet VCIL setups. Third, we find the core idea of the zero-shot approach (multi-modal contrastive loss) could be used to augment the zero-shot strategy, thus enabling prompting mechanisms, temporal modeling, and replay memories that still rely on a frozen CLIP encoder.

3.3. PIVOT Model

Prompting methods [49, 50] have emerged as an alternative to traditional rehearsal methods. While prompting mechanisms have been shown effective in drastically reducing memory sizes, they rely on in-distribution pre-training to make suitable prompt selections [50]. We approach a more general case, where we learn to prompt for video

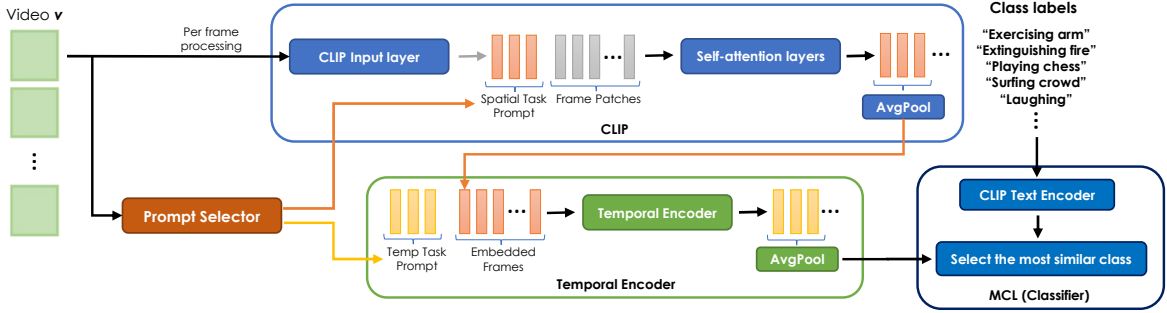


Figure 2. **Outline of the PIVOT Model at inference time.** Our model follows 3 main stages to make a prediction in video CIL setup: (1) PIVOT selects the specific task prompts for a video instance based on its similarity with the keys K of each task prompt (Blue box). (2) Then, PIVOT adds the selected spatial and temporal task prompts to the corresponding encoder spatial (CLIP) or Temporal Encoder (green box). (3) Finally, PIVOT classifies the video following an MCL, considering the similarity with text embeddings from the labels of all the available classes computed with the CLIP text encoder (black box).

data while relying exclusively on image-language knowledge bases. To this end, we propose PIVOT, a Prompting-based Model for Video Continual Learning. PIVOT leverages CLIP (f_{sp}) to encode spatial information and a transformer as temporal encoder (f_{tp}) to model the temporal relations essential to understanding actions. Additionally, we define a novel spatial and temporal prompting scheme that captures specific properties of each task to increase the model performance while mitigating the forgetting of previously learned tasks.

We outline a 3-stage process to effectively train and select prompts for the spatial and temporal dimensions of videos by leveraging extensive image domain pre-training. Despite having no video domain pre-training, our prompting strategy is effective on the action recognition task. To the best of our knowledge, this is the first prompting method that can work in the presence of such a domain shift.

Out-of-distribution Adaptation. The first training stage adapts an image pre-trained model to extract video features. Given an input video $v \in \mathbb{R}^{T \times H \times W \times C}$, where T is the number of video frames, we utilize CLIP to encode the video frames into meaningful spatial features $v_{sp} = f_{sp}(v) \in \mathbb{R}^{T \times D_m}$. These spatial features are then passed through the temporal encoder where we append a $[class]$ token that aggregates the temporal information of the feature sequence, resulting in the temporal feature of the video $v_{tp} = f_{tp}(v_{sp})[class] \in \mathbb{R}^{D_m}$. During this process, the CLIP encoder is completely frozen to avoid any distribution shift. In contrast, the temporal encoder is trained with the current task data (X_i, Y_i) . Specifically, we train the temporal encoder following the MCL approach outlined in Section 3.2. By using an approach that is informed by the CLIP frozen features, we ensure that the learned features of the temporal model stay consistent with the knowledge of the original image-language model.

Task-specific Prompt Generation. After tuning the temporal component of our model, the second training stage aims at learning task-specific prompts. Current prompting methods [50] would rely on continuously updating a prompt pool along the full task set. Since we can not rely on in-distribution pre-training to update and select prompts, we adopt task-specific prompts that preserve individual task knowledge along the CL process. Essentially, we train a new set of prompts as the data for a new task becomes available, then we append these new prompts into the pre-existent pool.

In addition to maintaining task-specific prompts, we propose to train the prompts and the prompt selector separately. At training time, we optimize the task-specific prompts P_n for task n following the MCL approach, while the prompts associated to previous tasks ($P_{n-1} \cup P_{n-2} \cup \dots \cup P_0$) remain fixed. To address the temporal nature of video data, each task will estimate independent spatial and temporal prompts $P_n = (P_n^{sp}, P_n^{tp})$, where $P_n^{sp} \in \mathbb{R}^{N_p \times L_{sp} \times D_{in}}$ and $P_n^{tp} \in \mathbb{R}^{N_p \times L_{tp} \times D_m}$, N_p is the number of prompts per task, L_{sp} and L_{tp} are the lengths of spatial and temporal prompts, respectively, and D_{in} is the output dimension of the CLIP input layer f_{sp}^e .

Prompt Selection Unlike L2P [49], we can not use extensive in-domain knowledge to select the best prompts for each input. To alleviate CLIP’s lack of domain comprehension, and thus improve its selection performance, we train the temporal encoder to select the prompts. We first associate each task-specific prompt with the feature set estimated from the frozen CLIP representations. To this end, we use the feature estimated on the language stream of CLIP, using the labels as textual input to the stream, this feature set will define the key K_n to index a prompt: $\{(K_1, P_1), (K_2, P_2), \dots, (K_n, P_n)\}$, where $K_n \in \mathbb{R}^{M_n \times D_m}$, M_n is the number classes on task n , and D_m is the output dimension of CLIP Text encoder.

We train the temporal encoder to select the corresponding task prompts for each input following a pseudo-classification based on the distance estimated by the MCL. We select the task prompts with the most similar key to the final video representation v_{tp} . Likewise, once the model selects the prompts, we train the temporal encoder to support them with an MCL. Figure 2 illustrates the PIVOT pipeline.

3.4. PIVOT Forward Pass

In Section 3.3, we outlined the construction and selection of both spatial and temporal prompts. Next, we discuss how we use these prompts during inference.

First, we note that CLIP f_{sp} is based on the ViT architecture, so, it is composed of input layer f_{sp}^e and self-attention layers f_{sp}^s . Given an input video v , we pass it through f_{sp}^e , which divides each frame into L patches and encodes them to $v_{sp} = f_{sp}^e(v) \in \mathbb{R}^{T \times L \times D_{in}}$. We leverage the spatial prompts at this stage, and concatenate P_n^{sp} with v_{sp} resulting in $v_{sp}^e = [P_n^{sp}; v_{sp}] \in \mathbb{R}^{T \times (N_p L_{sp} + L) \times D_{in}}$ and pass the extended sequence to the self-attention layers $v_{sp}^s = f_{sp}^s(v_{sp}^e) \in \mathbb{R}^{T \times (N_p L_{sp} + L) \times D_m}$. Following [50], for the final spatial representation we take an average pooling (AvgPool) through the first $N_p L_{sp}$ tokens of v_{sp}^s resulting in $v_{sp}^{avg} \in \mathbb{R}^{T \times D_m}$. It is important to note that, we use the same spatial prompt for all frames in a given video.

The previous process results in an embedding vector for each frame that takes into account both the extracted knowledge from CLIP and the learnt spatial prompts. We then aim at aggregating the information across all frames by leveraging a temporal encoder f_{tp} . Before we do so, we employ the learnt temporal prompts in a similar fashion to the spatial ones. In particular, we first concatenate $x = [P_n^{tp}; v_{sp}^{avg}]$ and then pass it through the temporal encoder followed by an average pooling layer. Formally $v_{tp} = AvgPool(f_{tp}(x); [N_p L_{tp}]) \in \mathbb{R}^{D_m}$. At last, we use v_{tp} as a global feature representation for the input video to classify it with the MCL classifier outlined in Section 3.2. See Section A in supplementary material for more details.

3.5. Training PIVOT

Our training pipeline is divided into three stages for each task in the VCIL setup. We first train our temporal encoder to correctly classify input data within the current task. Then, we leverage the knowledge of both CLIP and the trained temporal encoder to learn per task learnable tokens, *i.e.* prompts. Finally, we fine tune the temporal encoder to select and leverage the learned prompts to enhance the performance and reduce forgetting. It is worth mentioning that we freeze CLIP throughout the entire training process.

First Stage: Out-of-distribution Adaptation. In this stage, we want to align the temporal encoder with the extracted representation of a given video from CLIP. To do so, we train our temporal encoder to classify a video based on

the features extracted from CLIP with an MCL loss. Let θ_{tp} be the learnable parameters for f_{tp} , we seek to minimize the loss for the n^{th} task:

$$\min_{\theta_{tp}} \mathbb{E}_{(v,y) \sim (X_n, Y_n)} [\mathcal{L}(f_{tp}(f_{sp}(v)), y)], \quad (1)$$

where y is the representation of the class label of v obtained with the CLIP text encoder, \mathcal{L} is the MCL loss, and (X_n, Y_n) is the data distribution for the n^{th} task. We note here that relying exclusively on Eq. (1) for training might result in forgetting previously learnt tasks.

Second Stage: Task-specific Prompt Generation. In this stage, we leverage the knowledge encoded on the temporal encoder to learn meaningful prompts which are specific for task n . To that end, we freeze the temporal encoder during this stage. Then we randomly initialize a set of prompts for task n denoted as $P_n = (P_n^{sp}, P_n^{tp})$. We include P_n in the forward propagation of our pipeline as discussed in Section 3.4. We then learn the prompts P_n by minimizing the following loss:

$$\min_{P_{sp}, P_{tp}} \mathbb{E}_{(v,y) \sim (X_n, Y_n)} [\mathcal{L}(f_{tp}(f_{sp}(v, P_n^{sp}), P_n^{tp}), y)]. \quad (2)$$

Note that Eq. (2) aims at aligning the learnable prompts with the knowledge of both CLIP and the learnt temporal encoder for the same task. This is accomplished by the prompts P_n . We note here that we freeze the prompts for a given task while learning prompts for other tasks.

Third Stage: Prompt Selection. In the last stage, we fine-tune the temporal encoder over the replay memory data M , M includes data samples for the current task. This is accomplished through minimizing the combined MCL loss from Equations 1 and 2:

$$\min_{\theta_{tp}} \mathbb{E}_{(v,y) \sim M} [\mathcal{L}(f_{tp}(f_{sp}(v, P_{sp}), P_{tp}), y) + \mathcal{L}(f_{tp}(f_{sp}(v)), y)]. \quad (3)$$

It is worth mentioning that the last training step in Eq. (3) aligns the parameters of the temporal encoder to solidify the knowledge of a given task. This is achieved by training over the selected prompts and the network’s input. This will reduce forgetting previously learnt tasks as their corresponding prompts are frozen.

4. Experiments

We perform extensive experiments to benchmark existing methods against PIVOT in the VCIL action recognition task. In addition to comparing PIVOT directly to the existing methods reported in vCLIMB, we benchmark the performance of our frozen CLIP baseline. Then, we conduct thorough ablation studies to understand the role that each component of PIVOT contributes to its final performance. See our supplementary material for additional experiments.

Table 1. **Results on the two most challenging datasets of vCLIMB.** We follow the vCLIMB Benchmark [47] and report the average accuracy (Acc) and the backward forgetting (BWF) at 10 and 20 tasks. We highlight that CLIP zero-shot [36] outperforms regularized-based methods such as EWC [25] and MAS [2] and Memory-based methods such as iCaRL [37] and BiC [53]. Our approach PIVOT, which builds on top of CLIP, outperforms all the previous methods, by up to 31% in the Kinetics and ActivityNet 20-task sequence.

Model	Mem. Video Instances	Kinetics				ActivityNet-Trim				
		10 Tasks		20 Tasks		10 Tasks		20 Tasks		
		Acc ↑	BWF ↓	Acc ↑	BWF ↓	Acc ↑	BWF ↓	Acc ↑	BWF ↓	
EWC	None	5.81%	16.05%	2.95%	32.70%	None	4.02%	5.32%	1.28%	3.77%
MAS	None	7.81%	10.12%	4.25%	5.54%	None	8.11%	0.18%	4.61%	0.1%
BiC	8000	27.90%	51.96%	23.06%	58.97%	4000	51.96%	24.27%	46.53%	15.95%
iCaRL	8000	32.04%	38.74%	26.73%	42.25%	4000	48.53%	19.72%	43.33%	21.57%
CLIP	None	46.50%	9.80%	46.52%	11.32%	None	64.1%	10.47%	64.13%	11.71%
PIVOT	4000	55.13%	26.50%	55.04%	26.41%	2000	73.8%	11.41%	73.84%	11.94%
PIVOT w/o prompts	4000	58.61%	20.78%	57.51%	21.98%	2000	72.22%	13.71%	73.43%	13.45%
Upper bound	–	73.9%	–	73.9%	–	2000	88%	–	88%	–

4.1. vCLIMB Benchmark

We follow the vCLIMB benchmark [47] and conduct our experiments on Kinetics [24], ActivityNet [5], and UCF101 [43] datasets where classes of each dataset are presented sequentially over 10 or 20 tasks. For example, in the 10-task schedule of ActivityNet, the 200 classes are broken into 10 sets, each containing 20 non-overlapping categories. We report the average accuracy (Acc) of f_θ on all observed tasks (n), where higher values indicate better classification performance. Additionally, we follow the standard practice in VCIL [47], and report the backward forgetting (BWF), where smaller values indicate less forgetting and hence are better. Equation 4 shows the formulas for the metrics, where $A_{i,j}$ is the accuracy of task j after training task i .

$$\begin{aligned}
 Acc &= \frac{1}{n} \sum_{i=1}^n A_{n,i} \\
 BWF &= \frac{1}{n-1} \sum_{i=1}^{n-1} A_{i,i} - A_{n,i}
 \end{aligned}
 \tag{4}$$

Implementation details. PIVOT is augmented with task-specific prompts (see Algorithm 2 in the supplementary material). We add $N_p = 1$ set of spatial and temporal prompts per task with lengths $L_{sp} = 3$ and $L_{tp} = 3$, respectively. Therefore, at the end of the CL training process, PIVOT has 10 and 20 sets of spatial and temporal prompts in the 10-tasks and 20-tasks scenarios, respectively. Hence, PIVOT incorporates all three stages outlined in 3.3, each stage taking 40 epochs using SGD optimizer with a constant learning rate of 0.01 and a batch size of 50. The PIVOT temporal encoder is a transformer encoder with 2 heads and 3 layers. We use the CLIP model for all experiments, employing a ViT-B/32 architecture. $T = 8$ frames per video are sampled based on the strategy proposed in [48]. Moreover, we

define "PIVOT without prompts" (see Algorithm 1 in the supplementary material). It uses the very same architecture as PIVOT but lacks any prompting mechanisms. It is trained for 40 epochs using the same optimizer, learning rate, and batch size as PIVOT, following the MCL approach on every CIL task. It's worth noting that all methods with a temporal encoder in our ablation study (see Section 4.3) follow the same optimization setup as PIVOT w/o prompts.

On the contrary, the augmented CLIP model with spatial prompts follows the prompting approach of [50]. Therefore, it follows a similar setup as [50]: an Adam optimizer with a learning rate of 0.03, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and a batch size of 50. For this spatial prompting, we consider 10 prompts in total, each with a length $L_{sp} = 5$, and select 5 prompts per video instance using the same selection approach as [50].

4.2. Comparison with Baselines.

We compare the performance of our proposed methods PIVOT and PIVOT w/o prompts, and our CLIP baseline, against the approaches reported in the recent video class incremental learning benchmark vCLIMB [47], namely EWC, MAS, BiC, and iCaRL. PIVOT differs from PIVOT w/o prompts by employing task-specific prompts while sequentially learning the video tasks. We further report the upper bound performance that is acquired by training the TSN backbone [48] offline, i.e., on all the dataset classes simultaneously. The results for Kinetics and ActivityNet are shown in Table 1, while the results for UCF101 are reported in Table 2. Note that some methods store past-task exemplars to mitigate forgetting. We set the memory budget for these replay-based methods proportionally to the amount of classes of the dataset. The number of video instances used in each setup is stated in the respective table.

SOTA Comparisons to PIVOT. Our results show that PIVOT and PIVOT w/o prompts outperform all the meth-

Table 2. **Results on UCF101.** We report the average accuracy (Acc) and the backward forgetting (BWF) at 10 and 20 tasks. It is important to note that in this benchmark, the CLIP zero-shot does not outperform the baselines proposed in [47]. However, our PIVOT model outperforms the previous baselines with great consistency through the number of tasks.

Model	Mem. Video Instances	UCF101			
		10 Tasks		20 Tasks	
		Acc ↑	BWF ↓	Acc ↑	BWF ↓
EWC	None	9.51%	98.94%	4.71%	92.12%
MAS	None	10.89%	11.11%	5.90%	5.31%
BiC	2020	78.16%	18.49%	70.69%	24.90%
iCaRL	2020	80.97%	18.11%	76.59%	21.83%
CLIP	None	65.73%	11.14%	65.38%	12.35%
PIVOT	1010	93.36%	4.47%	93.07%	3.90%
PIVOT w/o prompts	1010	94.80%	3.89%	93.70%	4.77%
Upper bound	–	94.9%	–	94.9%	–

ods in vCLIMB by large margins. We first highlight that the highest performing method in vCLIMB was iCaRL, which had an average accuracy of 26.73% and 43.33% on 20-task Kinetics and 20-task ActivityNet respectively. The average accuracy of 57.51% obtained by PIVOT w/o prompts is 31% higher than iCaRL in the Kinetics 20 task set. Moreover, using PIVOT results in a remarkable improvement of 30% over the best result reported in vCLIMB for the 20 task set on the ActivityNet dataset. Regardless of the dataset used, our approach leverages image-pretraining, and a multi-modal contrastive learning approach to consistently outperform the prior VCIL art.

Image CLIP is a strong VCIL baseline. Following [45], we also propose to use a frozen CLIP [36] as a baseline for VCIL. Note that this baseline is trained on images and evaluated in a zero-shot manner, so it requires no fine-tuning on the sequence of video tasks. We denote this as CLIP in Table 1. Despite not being trained on video data, the CLIP baseline significantly outperforms the memory-based methods iCaRL and BiC reported in the vCLIMB benchmark. The performance gap is especially large in the 20-task setups, where the accuracy of CLIP outperforms iCaRL by 20% on both Kinetics and ActivityNet. This confirms the importance of having models trained on a large volume of data, even when the domains are different, and the relevance of the MCL that allows using it in a zero-shot manner.

The Effect of Long CIL Sequences. Our results in Table 1 show that a frozen image-pretrained model does not experience extra forgetting when the length of the CIL sequence increases. Moving from 10 to 20 tasks in Kinetics does not affect CLIP’s performance as it only changes from 46.5% to 46.52%. Similarly, for ActivityNet, the accuracy values of CLIP barely differ, going from 64.1% to 64.13%. This makes sense as the full clip encoder is frozen, so presenting the action categories in a shorter or longer sequence does not change the model’s behavior along the task set.

Likewise, the last two rows of Table 1, show that PIVOT and PIVOT w/o prompts significantly improve the CLIP baseline performance without resulting in a performance gap between the 10-task and 20-task sequences. In both Kinetics and ActivityNet, our models PIVOT and PIVOT w/o prompts are consistent regardless of the task sequence length. They present at most 1.21% between the accuracy in the two scenarios, which is much less than the 5.43% presented on [47].

Results on UCF101. We verify the effectiveness of PIVOT by comparing its performance to SOTA methods on the UCF101 dataset. While iCaRL and BiC perform reasonably well in the 10-task CIL scenario on this less challenging dataset, the final accuracy of frozen CLIP on UCF is around 15, 24% behind them, suggesting that this baseline may not always beat traditional CIL approaches. Yet, by effectively leveraging the extensive image pre-training through temporally aware MCL and prompting, PIVOT results in a significant improvement of 28% in accuracy over the CLIP baseline. Consistent with our results on Kinetics and ActivityNet, both PIVOT and PIVOT w/o prompts outperform the methods in vCLIMB with up to 14% improvement over iCaRL, and exhibit a minimal drop in performance between the 10 task and 20 task scenario.

Comparison with the Fully-supervised Upper Bound. We would like to highlight from Table 1 and Table 2 that the significant improvements, acquired by using PIVOT or PIVOT w/o prompts, result in a final CIL performance that is comparable to the upper bound of training on the whole data offline. While there is a significant gap to the upper bound performance on Kinetics, the gap to the upper bound in the ActivityNet 20-task setting is decreased from around 45% to 15%. Moreover, the accuracy of 94.8% obtained by PIVOT w/o prompts in the 10-task UCF101 benchmark already matches the upper bound performance of 94.9%.

4.3. Relevance of each component of PIVOT

To demonstrate the effectiveness and contribution of each component of PIVOT, we provide a series of experiments on UCF101 and ActivityNet. We start from the frozen CLIP baseline and we iteratively add one element of PIVOT to test its effect. The results are summarized in Table 3.

Effect of spatial prompts. We first add spatial learnable prompts to modulate the frozen CLIP, and we replace the MCL strategy with a linear classifier. The goal of these prompts is to guide the prediction of the classifier based on task-specific spatial information. We see from Table 3 these spatial prompts are useful in UCF101, resulting in a 12% increase in accuracy in the 10-task scenario. However, spatial prompting alone is not useful in ActivityNet. This can be explained by the complicated temporal dependencies that are found in ActivityNet but not in UCF101. We show later that augmenting these spatial prompts with temporal

Table 3. **Ablation study results.** We analyze the impact of each component of PIVOT: the memory buffer, the temporal encoder, the Multimodal Contrastive Loss (MCL), and the prompting approach. For the prompting approach, we compare PIVOT with an extension of L2P for the video domain using CLIP as the backbone (Spatial Prompting). As can be observed, every component of our model contributes significantly to its final performance, which outperforms all the state-of-the-art approaches.

Method	Uses Memory	UCF101				ActivityNet-Trim			
		10 Task		20 Task		10 Task		20 Task	
		Acc ↑	BWF ↓	Acc ↑	BWF ↓	Acc ↑	BWF ↓	Acc ↑	BWF ↓
CLIP Baseline	✗	65.73%	11.14%	65.38%	12.35%	64.1%	10.47%	64.13%	11.71%
+ Spatial Prompting	✗	77.66%	9.79%	68.66%	15.19%	64.21%	12.57%	57.89%	15.14%
+ Memory Buffer	✓	88.42%	11.91%	87.52%	11.94%	63.33%	26.17%	62.10%	28.29%
+ Classification with MCL	✓	86.17%	10.65%	87.59%	9.05%	69.53%	12.44%	69.32%	13.51%
+ Temporal encoder	✓	91.09%	5.95%	90.04%	5.19%	70.37%	10.83%	72.15%	10.02%
+ improved prompting (PIVOT)	✓	93.36%	4.47%	93.07%	3.90%	73.8%	11.41%	73.84%	11.94%
PIVOT w/o prompts	✓	94.80%	3.89%	93.70%	4.77%	72.22%	13.71%	73.43%	13.45%




Figure 3. **MCL robustness to the memory size.** All the experiments were done in the 10-task scenario of UCF101 considering different memory sizes (1010, 505, and 202) and two types of classifiers (Linear and MCL) for our PIVOT w/o prompts model.

prompts can increase the performance in ActivityNet.

Effect of memory replay. We extend the previous method by allowing the model to store a few video samples for future replay. As an incremental analysis, we keep using spatial prompts from the previous. This addition results in a further performance increase of 11% in the 10-task setting of UCF101 and 19% in the 20-task setting. This model is already ahead of the CLIP baseline by around 23%.

Effect of the Multimodal Contrastive Classifier. To address the temporal variation, we employ the MCL strategy explained in 3.2. Table 3 shows that this addition results in a 5% increase in final accuracy over the CLIP baseline in the challenging ActivityNet dataset. Moreover, as observed in Figure 3, the MCL is more robust to the number of instances in the memory than the traditional linear classifier.

Effect of the Temporal Modeling. As one of the limitations for using CLIP in videos context is the lack of temporal modeling in its architecture. To improve this point in VCIL, we utilize a transformer encoder to model the tem-

poral information of the input videos. This encoder further mitigates forgetting, decreasing the BWF from 10.65% to 5.95% in the 10-task UCF101 benchmark.

Effect of the improved Prompts. Given the transformer encoder, we can learn the temporal prompts. The resulting model is PIVOT, which further improves the accuracy on ActivityNet by around 3.5%. Finally, we show that PIVOT w/o prompts, which removes the prompting strategy from this model but keeps all the other components, results in the best performance on UCF101.

5. Conclusion and Limitations

We proposed PIVOT, a novel method that leverages large-scale pretraining from the image domain to mitigate catastrophic forgetting in video class incremental learning. PIVOT is the first prompting-based approach for video continual learning. While requiring much less memory than existing baselines, PIVOT outperforms these baselines by 31%, 27%, and 17.2% in the 20-task setups of the vCLIMB benchmark. Despite PIVOT’s advantages, we found it difficult to fine-tune the CLIP encoder for VCIL. PIVOT overcomes this problem by leveraging learnable prompts and a parameter-free multimodal contrastive classifier. However, we believe that recent advances in robustly fine-tuning large pre-trained models under distribution shifts [51] could lead to even greater improvements for PIVOT.

Acknowledgments. This work was supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research (OSR) under Award No. OSR-CRG2021-4648, as well as the SDAIA-KAUST Center of Excellence in Data Science and Artificial Intelligence (SDAIA-KAUST AI). Likewise, it was also partially funded by FONDECYT grant 1221425 and the National Center for Artificial Intelligence CENIA FB210017, Basal ANID.

References

- [1] Motasem Alfarra, Zhipeng Cai, Adel Bibi, Bernard Ghanem, and Matthias Müller. Simcs: Simulation for online domain-incremental continual segmentation. *arXiv preprint arXiv:2211.16234*, 2022. [2](#)
- [2] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Eur. Conf. Comput. Vis.*, pages 139–154, 2018. [1](#), [2](#), [6](#)
- [3] Rahaf Aljundi, Marcus Rohrbach, and Tinne Tuytelaars. Selfless sequential learning. In *Int. Conf. Learn. Represent.*, 2019. [2](#)
- [4] Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Multi-task feature learning. *Adv. Neural Inform. Process. Syst.*, 19, 2006. [1](#)
- [5] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 961–970, 2015. [3](#), [6](#)
- [6] Santiago Castro and Fabian Caba. Fitclip: Refining large-scale pretrained image-text models for zero-shot video understanding tasks. In *Brit. Mach. Vis. Conf.*, 2022. [2](#)
- [7] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*, 2018. [1](#), [2](#)
- [8] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*, 2019. [1](#), [2](#)
- [9] M. Delange, R. Aljundi, M. Masana, S. Parisot, X. Jia, A. Leonardis, G. Slabaugh, and T. Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE Trans. Pattern Anal. Mach. Intell.*, pages 1–1, 2021. [2](#)
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. [1](#)
- [11] 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. *arXiv preprint arXiv:2010.11929*, 2020. [1](#)
- [12] Sayna Ebrahimi, Mohamed Elhoseiny, Trevor Darrell, and Marcus Rohrbach. Uncertainty-guided continual learning in bayesian neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, 2019. [2](#)
- [13] Sayna Ebrahimi, Franziska Meier, Roberto Calandra, Trevor Darrell, and Marcus Rohrbach. Adversarial continual learning. *Eur. Conf. Comput. Vis.*, 2020. [2](#)
- [14] Sayna Ebrahimi, Suzanne Petryk, Akash Gokul, William Gan, Joseph E. Gonzalez, Marcus Rohrbach, and trevor darrell. Remembering for the right reasons: Explanations reduce catastrophic forgetting. In *International Conference on Learning Representations*, 2021. [2](#)
- [15] Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A Rusu, Alexander Pritzel, and Daan Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734*, 2017. [2](#)
- [16] Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999. [1](#)
- [17] Qiankun Gao, Chen Zhao, Bernard Ghanem, and Jian Zhang. R-dfcil: Relation-guided representation learning for data-free class incremental learning. In *Eur. Conf. Comput. Vis.*, 2022. [2](#)
- [18] Yasir Ghunaim, Adel Bibi, Kumail Alhamoud, Motasem Alfarra, Hasan Abed Al Kader Hammoud, Ameya Prabhu, Philip HS Torr, and Bernard Ghanem. Real-time evaluation in online continual learning: A new paradigm. *arXiv preprint arXiv:2302.01047*, 2023. [1](#)
- [19] Tyler L Hayes, Kushal Kafle, Robik Shrestha, Manoj Acharya, and Christopher Kanan. Remind your neural network to prevent catastrophic forgetting. In *Eur. Conf. Comput. Vis. Springer*, 2020. [2](#)
- [20] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Int. Conf. Comput. Vis.*, pages 2961–2969, 2017. [1](#)
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–778, 2016. [1](#)
- [22] Wenpeng Hu, Zhou Lin, Bing Liu, Chongyang Tao, Zhengwei Tao Tao, Dongyan Zhao, Jinwen Ma, and Rui Yan. Overcoming catastrophic forgetting for continual learning via model adaptation. In *Int. Conf. Learn. Represent.*, 2019. [2](#)
- [23] Julio Hurtado, Alain Raymond, and Alvaro Soto. Optimizing reusable knowledge for continual learning via metalearning. *Adv. Neural Inform. Process. Syst.*, 34:14150–14162, 2021. [2](#)
- [24] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. [3](#), [6](#)
- [25] 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*, 114(13):3521–3526, 2017. [1](#), [2](#), [6](#)
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. [1](#)
- [27] Timothée Lesort, Hugo Caselles-Dupré, Michael Garcia-Ortiz, Andrei Stoian, and David Filliat. Generative models from the perspective of continual learning. In *Int. Joint Conf. on Neural Networks*. IEEE, 2019. [2](#)
- [28] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(12):2935–2947, 2017. [2](#)

- [29] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *Adv. Neural Inform. Process. Syst.*, 30, 2017. 1, 2
- [30] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. *Eur. Conf. Comput. Vis.*, 2018. 2
- [31] Sanket Vaibhav Mehta, Darshan Patil, Sarath Chandar, and Emma Strubell. An empirical investigation of the role of pre-training in lifelong learning. *arXiv preprint arXiv:2112.09153*, 2021. 2
- [32] Jorge A Mendez and Eric Eaton. How to reuse and compose knowledge for a lifetime of tasks: A survey on continual learning and functional composition. *arXiv preprint arXiv:2207.07730*, 2022. 2
- [33] Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. St-adapter: Parameter-efficient image-to-video transfer learning. In *Adv. Neural Inform. Process. Syst.*, 2022. 2
- [34] Jaeyoo Park, Minsoo Kang, and Bohyung Han. Class-incremental learning for action recognition in videos. In *Int. Conf. Comput. Vis.*, pages 13698–13707, October 2021. 1, 2, 3
- [35] Lorenzo Pellegrini, Gabriele Graffieti, Vincenzo Lomonaco, and Davide Maltoni. Latent replay for real-time continual learning. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10203–10209. IEEE, 2020. 1
- [36] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In Marina Meila and Tong Zhang, editors, *Int. Conf. on Mach. Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR, 18–24 Jul 2021. 1, 2, 3, 6, 7
- [37] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2001–2010, 2017. 1, 6
- [38] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Adv. Neural Inform. Process. Syst.*, 28, 2015. 1
- [39] 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. 2
- [40] Gobinda Saha and Kaushik Roy. Gradient projection memory for continual learning. In *Int. Conf. Learn. Represent.*, 2021. 2
- [41] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *Adv. Neural Inform. Process. Syst.*, 30, 2017. 2
- [42] 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. *Int. Conf. Comput. Vis.*, pages 9374–9384, October 2021. 2
- [43] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 6
- [44] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *Int. Conf. on Mach. Learning*, pages 6105–6114. PMLR, 2019. 1
- [45] Vishal Thengane, Salman Khan, Munawar Hayat, and Fahad Khan. Clip model is an efficient continual learner. *arXiv preprint arXiv:2210.03114*, 2022. 2, 7
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Adv. Neural Inform. Process. Syst.*, 30, 2017. 1
- [47] Andrés Villa, Kumail Alhamoud, Victor Escorcia, Fabian Caba, Juan León Alcázar, and Bernard Ghanem. vclimb: A novel video class incremental learning benchmark. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 19035–19044, June 2022. 1, 2, 3, 6, 7
- [48] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks for action recognition in videos. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(11):2740–2755, 2019. 6
- [49] 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. *Eur. Conf. Comput. Vis.*, 2022. 2, 3, 4
- [50] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 139–149, June 2022. 2, 3, 4, 5, 6
- [51] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 8
- [52] Mitchell Wortsman, Vivek Ramanujan, Rosanne Liu, Aniruddha Kembhavi, Mohammad Rastegari, Jason Yosinski, and Ali Farhadi. Supermasks in superposition. *Adv. Neural Inform. Process. Syst.*, 33, 2020. 2
- [53] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2019. 6
- [54] Hu Xu, Gargi Ghosh, Po-Yao Huang, Dmytro Okhonko, Armen Aghajanyan, Florian Metze, Luke Zettlemoyer, and Christoph Feichtenhofer. Videoclip: Contrastive pre-training for zero-shot video-text understanding. *Conf. on Empirical Methods in Natural Language Processing*, abs/2109.14084, 2021. 2
- [55] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *Int. Conf. on Mach. Learning*, 2017. 2