

Just a Glimpse: Rethinking Temporal Information for Video Continual Learning

Lama Alssum, Juan León Alcázar, Meredy Ramazanova, Chen Zhao, Bernard Ghanem
King Abdullah University of Science and Technology (KAUST)

{lama.alssum.1, juancarlo.alcazar, meredy.ramazanova, chen.zhao, bernard.ghanem}@kaust.edu.sa

Abstract

Class-incremental learning is one of the most important settings for the study of Continual Learning, as it closely resembles real-world application scenarios. With constrained memory sizes, catastrophic forgetting arises as the number of classes/tasks increases. Studying continual learning in the video domain poses even more challenges, as video data contains a large number of frames, which places a higher burden on the replay memory. The current common practice is to sub-sample frames from the video stream and store them in the replay memory. In this paper, we propose SMILE a novel replay mechanism for effective video continual learning based on individual/single frames. Through extensive experimentation, we show that under extreme memory constraints, video diversity plays a more significant role than temporal information. Therefore, our method focuses on learning from a small number of frames that represent a large number of unique videos. On three representative video datasets, Kinetics, UCF101, and ActivityNet, the proposed method achieves state-of-the-art performance, outperforming the previous state-of-the-art by up to 21.49%.

1. Introduction

Recently, a large amount of data has become available on the Internet due to high-speed internet access and the rapid growth of social media platforms. In particular, sharing and uploading video data has become popular with the availability of video cameras in mobile and wearable devices. As the number of available videos grows, so does the required human labor to analyze them. This expensive manual labor motivates the development of intelligent systems for video understanding and analysis. Recently, a growing interest of the research community in human action recognition has led to the development of massive video datasets and large-scale models with a large number of action classes [5, 16, 38, 40].

Most of the previous work in video action recognition assumes that a large amount of labeled data from a prede-

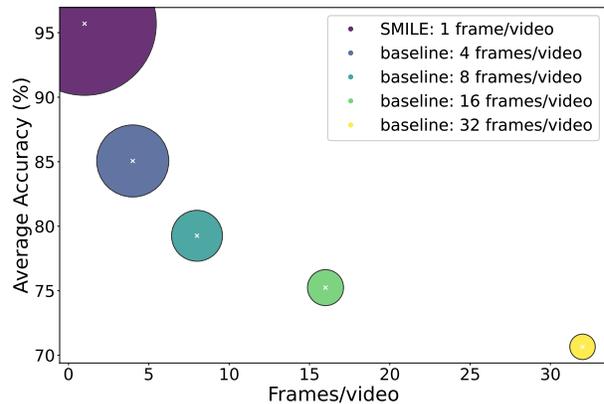


Figure 1. **SMILE vs. the baseline performance (iCaRL)**. Although our method (SMILE) is restricted to storing a single frame per video, this memory-efficient sampling outperforms the scenarios that focus on preserving the temporal progression of the video data in the UCF101 dataset, regardless of the temporal resolution of the videos stored in the memory buffer. In a restricted memory setting (e.g., rehearsal methods for continual learning), video diversity (represented by the disc size) is the most crucial factor in achieving high performance in video class incremental learning.

defined number of classes is available, and this class set remains fixed once the model is deployed. These assumptions do not always hold in real-world applications where novel classes are continually identified, and collecting enough labeled data is expensive and time-consuming. With a continuous stream of novel data and classes, the development of models that can learn sequentially from a set of tasks is needed. Such a setup requires the sequential fine-tuning of the model on a set of tasks. Under these circumstances, it has been shown that Neural Networks suffer from a phenomenon known as catastrophic forgetting [9, 10, 27], where fine-tuning the model on a new task reduces its performance drastically on previously learned tasks. Continual Learning (CL) is the field of study that addresses the challenging setup of training a model on a set of sequential tasks while mitigating catastrophic forgetting [8].

We focus on a special case of continual learning called Class Incremental Learning (CIL) [26], where the model

learns on a set of tasks that contain disjoint classes with access to a limited amount of previously seen data, named experience replay (ER) [33]. Currently, ER methods have shown superior performance in the CIL setup [22, 28] compared to their regularization-based counterparts [1, 6, 17, 43].

Previous work has already explored the CIL setup in the video domain [24, 29, 30, 39, 44]. However, video data requires a significantly larger memory size in comparison to image data; due to the additional temporal dimension. This drawback is a central motivation for the development of video CL strategies endowed with an efficient working memory management strategy. Despite the clear relevance of memory efficiency in video CL, it remains an understudied topic.

In this work, we set our attention towards one of the main limitations of ER-based methods in the video continual learning setup, namely the memory constraint. ER methods suffer from catastrophic forgetting, as the limited memory capacity allows only a fraction of previously seen data to be stored, and the in-memory training subset typically diverges from the original distribution of the training data [2, 4, 19, 21]. Moreover, as the task number increases, some elements must be evicted from the memory buffer to accommodate examples of novel classes. The model is then updated with increasingly skewed versions of the original distribution.

We depart from the golden standard in video continual learning, as we no longer attempt to retain any temporal information in the memory buffer. In fact, we propose an extreme sampling strategy that completely ignores temporal data in favor of greedily sampling individual frames from as many training videos as possible. We empirically show that our extreme sampling strategy, paired with an inexpensive training time regularization, can outperform the well-established ER strategies that focus on preserving the temporal progression of the video data, as shown in Figure 1, while also reducing the total amount of required memory. Moreover, since we store a single frame from every training sample, we create a direct equivalence between the memory buffers devised for image continual learning and those for video continual learning. This enables the direct application of image CIL methods for video CIL. This paper introduces SMILE (Single fraMe Incremental LEarner) for video CIL. Our work brings the following contributions:

- We introduce an extreme memory sampling strategy for video CIL, where a single representative frame is stored in memory. Paired with our regularization technique, this sampling can reach state-of-the-art performance across all datasets in the vClimb benchmark while using far less replay memory.
- We show that storing a single frame in memory allows the direct use of image-domain CL methods for video

data. This direct adaptation outperforms all other CIL methods in the challenging vClimb benchmark.

2. Related Work

Most of the literature approaches the CL problem in the image domain. For image classification, current methods can be divided into three sub-categories [7]: regularization, replay memory, and parameter isolation.

i) Regularization methods aim at building and retaining a parameter-set that can perform consistently across tasks [1, 17, 19, 43]. That is, for every individual task, a subset of the most relevant network weights is identified, and a penalty is applied for updating these parameters. It is expected that the critical parameter sub-sets for individual tasks are kept partially stable across the continual learning process.

ii) Rehearsal methods mitigate catastrophic forgetting by storing a subset of the original training samples [14, 23, 31, 32, 41] or by fitting a generative model to training data of previous tasks [12, 18, 36], thus enabling the generation and storage of surrogate training samples. This memory is replayed while learning a novel task, typically by using it as additional training data or by estimating boundary points that guide the optimization process on novel tasks.

iii) Parameter isolation methods tackle catastrophic forgetting by segmenting and specializing the parameter space in the Neural Network [25, 34, 35]. These models generate parameter sub-spaces specialized on an individual task, thus preventing performance loss across tasks. In its most simple way, they create alternative branches for data processing inside the network while freezing the existing parameter set from previous tasks [15, 42].

Video Continual Learning. Recently, some works have started approaching the Continual Learning problem in the video domain [24, 29, 39, 44]. Notably, the release of the vCLIMB benchmark [39] has provided the first standard test-bed for video CIL, along with the initial baseline results of classic CIL methods [32, 41] applied into video data.

Currently, the best-performing methods in vCLIMB are rehearsal approaches where a subset of the video frames is stored in the replay memory. Such an approach suffers from two drawbacks: First, methods must sub-sample the frames to be stored as pushing a full video in the replay memory is prohibitively expensive. Second, video data might contain some background frames which are not relevant to the classification task. However, memory methods still lack a frame selection strategy that could directly select the best frames to store in memory. Third, it is still unknown how to deal with the temporal dimension of videos; in other words, how to sample a memory that retains the relevant temporal cues contained in the original video.

To mitigate these drawbacks, we follow a rehearsal strategy but radically depart from the standard approach of stor-

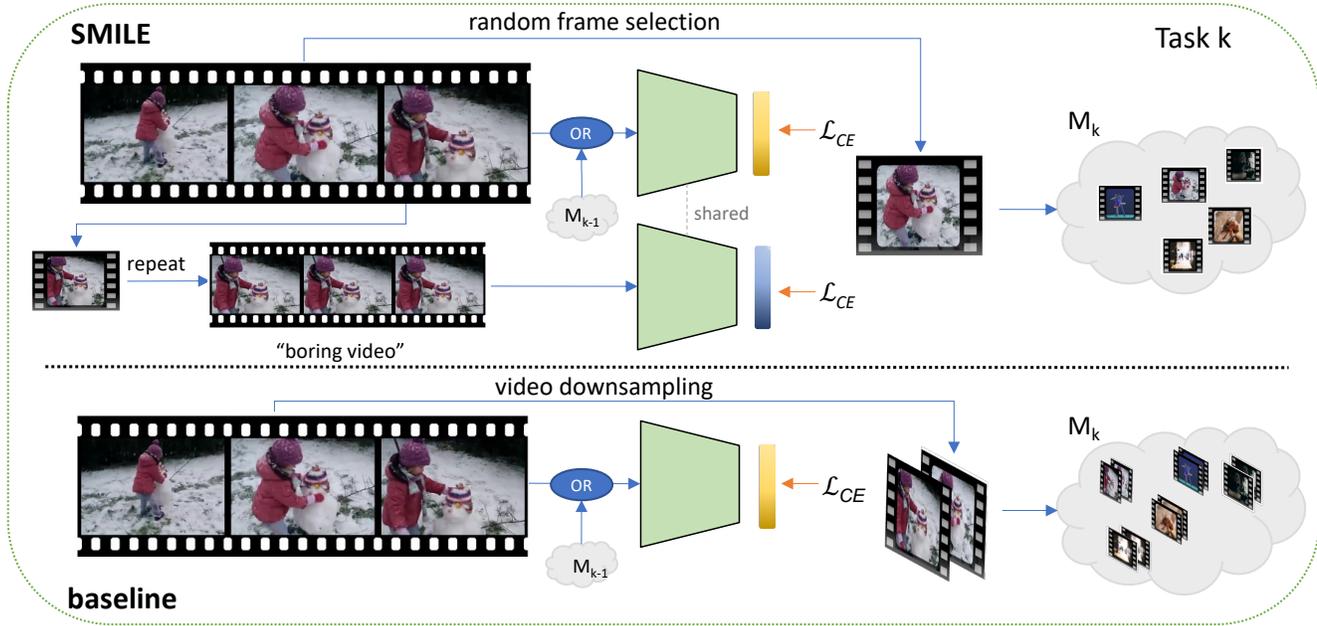


Figure 2. **SMILE vs. baseline methods.** SMILE relies on sampling a single frame per video into the replay memory. This sampling favors video diversity and improves the performance in the CIL setup. In addition, it drastically reduces memory usage. However, it removes the temporal information from memory. We develop a secondary forward pass that helps the video encoder to adapt to both modalities (video and static images). Under this sampling strategy, we can directly use image-based continual learning methods in the video domain.

ing a sub-set of frames in memory. Instead, we explore the extreme scenario where a single frame is stored in the replay memory. We show that without any bells and whistles, this extreme memory setup can outperform state-of-the-art methods in the vClimb benchmark.

Rehearsal Methods for Video Action Recognition. Following up with the previous line of research on the continual learning in videos [24, 29, 39, 44], we benchmark the performance on the task of action recognition. Previous works used frame-based approaches for the video representation methods [20, 40]. In this setup, several frames are sampled from temporal segments in a video and stored in the memory. In our work, we show how storing in the memory as few as one frame per video can result in a better trade-off, as more diverse video information could be preserved then.

3. SMILE: Single Frame Incremental Learner

In this section, we outline the details of our proposed approach. Since we rely on single-frame memories, we name our method SMILE (Single fraMe Incremental LEarner). Before delving into the details of SMILE, we first provide a formal definition of the video class incremental learning (CIL) problem and then motivate our extreme memory design with an initial experiment on the UCF101 dataset [37].

Notation and Problem Formulation. In class incremental learning, a neural network, $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ param-

eterized by θ , is trained on a set of m sequential tasks $T = \{t_1, t_2, \dots, t_k, \dots, t_m\}$. Each t_k is composed of n_k samples $\{(\mathbf{x}_1^k, y_1^k), (\mathbf{x}_2^k, y_2^k), \dots, (\mathbf{x}_{n_k}^k, y_{n_k}^k)\} \sim \mathcal{D}_k$ where \mathbf{x}_i^k represents a training sample and y_i^k is its corresponding ground truth label in the dataset \mathcal{D}_k . The classes of samples introduced at each task $\{y_i^k\}_{i=1}^{n_k} \in \mathcal{Y}_k$ are unique, *i.e.* $\bigcap_{i=1}^m \mathcal{Y}_i = \emptyset$.

3.1. Analyzing Replay Memories for Video CIL

We motivate the single frame sampling by summarizing the memory constraints defined in vCLIMB [39] and explaining the trade-off between dense temporal sampling and memory diversity. We define memory diversity as the number of unique video clip samples stored in the memory.

The vCLIMB Benchmark defines the replay memory in terms of frames. That is, a video must be stored as a full set of contiguous frames or as a temporally sub-sampled version of it. This design choice of using the number of frames to define the memory size enforces a common ground for comparison between methods, regardless of the length of the video data. Moreover, it creates a trade-off unique to video CIL, we must choose whether to store full videos (fewer unique videos, with full temporal data) or store sub-sampled videos (more unique videos, incomplete temporal data).

In Table 1, we provide the initial insight into our approach. We set a fixed memory budget of 16160 frames

Total Memory Size (Frames)	16160	16160	16160	16160
Unique Videos in Memory	4040	2020	1010	505
Stored Frames per Video	4	8	16	32
Average Accuracy (Acc)	85.06	79.26	75.24	70.65

Table 1. **Temporal Resolution vs. Memory Diversity.** We empirically assess that the CIL methods seem to favor video diversity over dense temporal sampling. For a fixed memory budget, the direct application of iCaRL to the UCF101 dataset shows that better results are obtained when pushing fewer frames per video (but more unique videos) into replay memory.

and analyze the effect of pushing 4, 8, 16, or 32 frames per video into the replay memory. We also report the number of unique videos stored in memory after the final task. We deploy a CIL setup on UCF101 [37] using iCaRL [32]. This simple analysis suggests that experience replay (ER) methods might benefit from a regime where the replay memory is filled with less temporal data (4 frames per video) while favoring video diversity (4040 unique videos in memory), instead of a regime with more temporal information (32 frames per video) and less video diversity (505 unique videos).

Based on these initial observations, we propose to improve the working memory usage by maximizing the diversity of samples contained in it. We focus on the extreme case of sub-sampling, where we maximize the diversity of in-memory elements by sampling a single frame into the working memory. Such an extreme sampling strategy represents a two-fold improvement as follows. First, it drastically reduces the total amount of required memory. Second, it enables the direct application of ER methods from the image domain. However, by sampling a single frame, we will lose the temporal dimension for the samples stored in the working memory. To mitigate the forgetting associated with the lack of temporal information in the memory bank, we develop a novel training-time regularization detailed in the following section.

3.2. Single Frame Memories for Video CIL

Following our memory sampling strategy, two types of data will be available during the CIL process. For every new task, video clips with full temporal resolution will be available. Meanwhile, the data from previous tasks will be represented exclusively by images (single frame per video) stored in the replay memory. To adapt to these two data sources, we begin by “inflating” the input tensor for image data. We follow [5] to expand the frame data stored in memory into a “boring video”, namely, a video clip composed of n copies of the original frame.

Boring videos enable us to learn from both types of data

Dataset	Tasks	Videos Per Task			Classes	Avg. Frames
		Train	Val	Test	Per Task	Per Video
UCF101	10	928	131	272	10	183
	20	464	65	136	5	183
ActivityNet	10	1541	765	–	20	3879
	20	770	383	–	10	3879
Kinetics	10	24628	1988	3977	40	250
	20	12314	994	1988	20	250

Table 2. **Tasks Statistics.** We outline the key statistics of the task sets proposed in vCLIMB [39]. The benchmark proposes two task schedules containing a total of 10-tasks and 20-tasks in every dataset.

with the same video encoder, however, these boring videos introduce a distribution shift between real video clips and in-memory images (represented as boring videos). To address this undesired shift, we propose the following loss:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{ce}(X_k, Y_k) + \lambda_2 \mathcal{L}_{ce}(\hat{X}_k, Y_k) + \lambda_3 \mathcal{L}_{ce}(X_m, Y_m),$$

where \mathcal{L}_{ce} is cross-entropy loss. X_k and \hat{X}_k are the video samples of the task t_k and the boring video version of them, respectively. Y_k is the ground truth labels for X_k . X_m and Y_m are the samples stored in memory and their ground truth labels. λ_1 , λ_2 , and λ_3 are weights to control the contribution of each loss term. The secondary loss ($\mathcal{L}_{ce}(\hat{X}_k, Y_k)$) directly addresses the domain shift, as it allows the video encoder to slowly adapt to the two data sources (standard and boring video) as the CIL tasks progress.

In every new task, we optimize our model using the original video samples (available as the task training data), the boring video version of them, and the boring videos from previous tasks (stored as images in replay memory). Since the boring and the original videos can be forwarded through the same encoder, we simply gather the samples after the forward pass and re-group them in the appropriate factor of the loss function. The pipeline of our approach is illustrated in Figure 2.

CIL Methods. Our single-frame working memory creates a direct link between the image domain and the video domain. We choose two representative methods from the image domain, BiC [41] and iCaRL [32]. We directly apply both methods to our single-frame memory approach without any changes.

4. Experimental Evaluation

In this section, we report the results across all the datasets included in the vCLIMB benchmark: UCF101

Model	Frames Per Video	UCF101		ActivityNet		Kinetics	
		Acc \uparrow	BWF \downarrow	Acc \uparrow	BWF \downarrow	Acc \uparrow	BWF \downarrow
iCaRL+SMILE (Ours)	1	95.70	2.59	50.26	15.87	46.58	7.34
BiC+SMILE (Ours)	1	92.49	2.15	54.83	7.69	52.24	6.25
vCLimb (iCaRL) [39]	N.A.*	80.97	18.11	48.53	19.72	32.04	38.74
vCLimb (BiC) [39]	N.A.*	78.16	18.49	51.96	24.27	27.90	51.96
vCLimb (iCaRL+TC) [39]	16	75.84	23.23	44.04	22.82	36.54	33.53
vCLimb (iCaRL+TC) [39]	8	74.25	25.27	45.73	18.90	36.24	33.83
vCLimb (iCaRL+TC) [39]	4	73.85	26.35	42.99	23.82	35.32	34.07

Table 3. **State-of-the-art comparison vCLIMB (10-tasks).** We report our results on the vCLIMB benchmark, along with the state-of-the-art outlined in [39]. Despite storing at least 87.7% less frames in memory, SMILE reports state-of-the-art performance for every dataset included in the vCLIMB benchmark. SMILE also outperforms the low-memory setting of [39] in every dataset and memory budget.*: N.A. indicates that the number of saved frames per video is not fixed.

[37], ActivityNet [3], and Kinetics [5], and provide ablation results to support our design decisions.

Datasets. UCF101 has 13.3K videos from 101 classes. ActivityNet consists of 20k videos from 200 classes and can be used for both trimmed and untrimmed action recognition. In this work, we use the trimmed version, where every frame of the video is part of an action instance. Kinetics is a large-scale dataset that has more than 300K short videos from 400 classes. For every dataset, we train SMILE sequentially on task-sets composed of 10 and 20 tasks as defined in [39]. Statistics for these tasks-sets are presented in Table 2.

Evaluation Metrics. To evaluate the performance of SMILE, we use the Average Accuracy metric (Acc) [23], which calculates the mean accuracy of all learned tasks after finishing the training on the sequence of tasks $T = \{t_1, t_2, \dots, t_m\}$:

$$Acc = \frac{1}{m} \sum_{i=1}^m acc_i,$$

where m is the number of available tasks and acc_i is the accuracy of $t_i \in T$. A higher Acc indicates better performance. We also use Backward Forgetting (BWF) [23] to measure the impact of learning a new task t_k on the performance of previously learned tasks, it is calculated as:

$$BWF = \frac{1}{m-1} \sum_{i=1}^{m-1} acc_{i,i} - acc_{m,i},$$

where m is the number of tasks learned, and $acc_{m,i}$ and $acc_{i,i}$ are the accuracy on task i after training on task m and the accuracy on task i after training on task i , respectively. A higher value of BWF indicates a larger degree of forgetting.

Implementation Details. Both iCaRL and BiC are trained sequentially on each task following the original implemen-

tation outlined in [41] and [32]. SMILE video encoder is based on TSN [40] with a ResNet-34 backbone [13], pre-trained with ImageNet weights. We operate on 8 segments per video (*i.e.* we sample $N = 8$ frames across evenly spaced temporal segments on a video). We use the SGD optimizer with a batch size of 80 and a learning rate of 1×10^{-3} for UCF101 and 6×10^{-4} for both ActivityNet and Kinetics. For the loss, we set $\lambda_1 = 1$, $\lambda_2 = 1$, and $\lambda_3 = 2$. For 10-tasks setup, we train iCaRL for [25, 25, 15] epochs and BiC for [50, 25, 50] epochs on [UCF101, Kinetics, ActivityNet], respectively. For 20-tasks setup, we train iCaRL for [25,15,15] epochs and BiC for [60, 20, 55] epochs on [UCF101, Kinetics, ActivityNet], respectively.

Efficient Calculation of the Secondary Loss. Since we work with TSN, the calculation of the secondary backward pass does not incur additional computational costs. The TSN framework already makes independent forward passes for a set of frames sampled from the original video, then averages their logits to get the video’s prediction. As a result, we can simply obtain the logits of any boring version of the video as a copy of the logits found in the original TSN forward pass. We simply perform an extra loss calculation on the copy of the logits and accumulate the losses in the backward pass.

4.1. Comparison Against the State-of-the-art

Video CIL on 10-Tasks. We first evaluate SMILE on the 10-task setup proposed in [39] using BiC and iCaRL. In Table 3, we report the results for SMILE along with the current state-of-the-art on UCF101, ActivityNet, and Kinetics. We also include the performance of the memory-efficient regime proposed in [39].

On every dataset, SMILE outperforms the current state-of-the-art. For UCF101, SMILE shows a significant improvement in acc over both baselines by 14.73% (iCaRL)

and 14.33% (BiC). We highlight that SMILE is only marginally below the standard fully supervised training setup [11] (training on the whole dataset in one stage) by 2.94%. Our experiments on ActivityNet show that SMILE outperforms baselines with an Acc of 54.83% for BiC and 50.26% for iCaRL, improving over the previous state-of-the-art vCLimb (BiC) by 2.87%. Finally, SMILE shows a significant improvement on the challenging Kinetics dataset, improving by 20.2% (BiC) over the previous state-of-the-art vCLimb (iCaRL).

We also note that backward forgetting is significantly reduced. In the case of the UCF101 dataset, SMILE reports BWF of 2.59% for iCaRL and 2.15% for BiC. A similar situation is observed in ActivityNet, where the BWF is reduced by 3.85% for iCaRL and 16.58% for BiC. Finally, the backward forgetting for Kinetics drops from 38.74% to 7.34% for iCaRL and from 51.96% to 6.25% for BiC.

Furthermore, [39] also proposed a memory efficient regime named 'Temporal Consistency' (TC). This regime explored pushing only 4, 8 and 16 frames and applying a feature regularization for sub-sampled video data. Although our proposal relies on single frame memory, SMILE outperforms any memory regime explored in [39] by at least 19.86% in UCF101, 9.10% in ActivityNet and 15.7% in Kinetics.

In Figure 3, we provide further analysis of the forgetting of SMILE using iCaRL on UCF101 dataset for the 10 task setup. We show that SMILE suffers far less from forgetting when compared to the baseline approaches that push 4, 8, 16 and 32 frames into the working memory. We also note that there is a significant difference in forgetting as the task set progresses. While the 4 and 8 frame baselines are close to SMILE during the first 4 tasks, they exhibit significant forgetting in the final 4 tasks (tasks 7 to 10). Meanwhile, SMILE suffers only marginally from forgetting in the same task-subset.

Video CIL on 20-Tasks. We conclude the state-of-the-art comparison by addressing the more challenging 20-tasks setup proposed in [39]. We report the performance of SMILE on UCF101, ActivityNet, and Kinetics along with the current state-of-the-art in Table 4. Despite the increased difficulty of the 20-tasks scenario, we observe that SMILE still outperforms the state-of-the-art on all datasets by up to 21.49% (Kinetics). Our approach outperforms the baselines in every scenario, except when we use iCaRL with ActivityNet where we get nearly on-par results, an Acc of 43.45% compared to the baseline's Acc of 43.33%.

4.2. Effect of Replay Memory Size

After performing the state-of-the-art comparison, we ablate our main design choices and explore their individual contributions. In Table 3, the memory budget for SMILE is defined by the scale of the training dataset (1 frame per

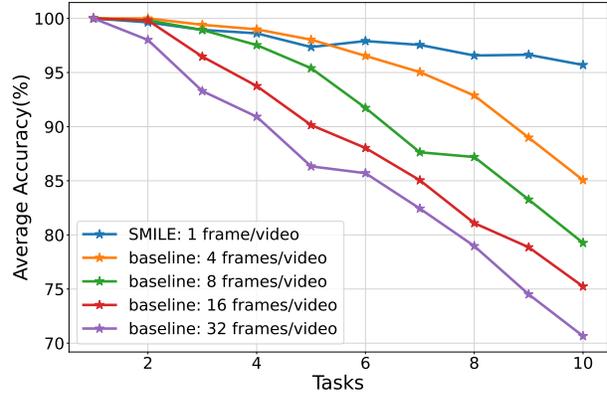


Figure 3. Average Accuracy for the 10-task setup with different numbers of frames stored in the replay memory compared to iCaRL+SMILE (UCF101). The trend clearly shows how storing fewer frames in the memory (thus allowing to preserve more videos), leads to performance boosts.

video). These memory sizes are 9280, 15410, and 246282 frames for UCF101, ActivityNet, and Kinetics, respectively. This already represents a reduction of 97.5%, 99.9%, and 87.7% of the memory budget proposed in vCLIMB. In this section, we investigate the performance of SMILE as we reduce the memory budget further.

Tables 5 and 6, summarize the results of running iCaRL on UCF101 and ActivityNet using the 10-tasks setup for smaller memory budgets. For UCF101, we can see that by reducing the memory size from 9280 to 4040 frames, which represents a relative reduction of 56.47% in memory, we get an Acc of 86.07% which is still higher than the results reported by [39] (80.97%). Moreover, by setting the memory budget to 2020 frames, which is about 21.77% of the original memory budget, we are on par with the results of [39]. For ActivityNet, Table 6 shows that by reducing the memory size from 15410 to 12000 frames which represents a relative reduction of 22.13% in memory, we get an Acc of 49.10% compared to [39] (48.53%). A further reduction of the memory to 8000 frames will report an Acc of 49.39%, which is slightly higher than the results of [39].

4.3. Effect of Frame Selection

It would be expected that, under such extreme sampling circumstances, the selection strategy for memory frames would play a key role in boosting the effectiveness of our approach. We empirically assess that a structured frame selection does not bring noticeable improvements to our approach. Using TSN as our backbone, we simultaneously obtain global video features (for classification) and frame-level features (for the estimation of the global feature) of any frame included in the forward pass. On this basis, we can select the frame whose feature embedding most closely

Model	UCF101		ActivityNet		Kinetics	
	Acc \uparrow	BWF \downarrow	Acc \uparrow	BWF \downarrow	Acc \uparrow	BWF \downarrow
iCaRL+SMILE (Ours)	95.09	1.72	43.45	23.02	45.77	4.57
BiC+SMILE (Ours)	90.90	1.62	51.14	1.33	48.22	0.31
vCLimb (iCaRL) [39]	76.59	21.83	43.33	21.57	26.73	42.25
vCLimb (BiC) [39]	70.69	24.90	46.53	15.95	23.06	58.97

Table 4. **SMILE vs. the state-of-the-art performance (20-tasks)**. We report our results along with vCLIMB’s using both average accuracy (Acc) and backward forgetting (BWF) on the challenging 20-tasks setup. We outperform the state-of-the-art for all datasets.

Memory Size (Frames)	2020	4040	8080	9280
Acc \uparrow	78.01	86.07	94.97	95.70
BWF \downarrow	15.68	10.36	2.32	2.59

Table 5. **Effect of memory budget (UCF101)**. We tested even smaller memory budgets for SMILE on UCF101 using iCaRL as CL method. We discover that even with just 2020 total frames (storing a single frame for about 1/4 of the dataset) SMILE is just below the current state-of-the-art. SMILE actually outperforms the state-of-the-art in every other scenario.

Memory Size (Frames)	4000	8000	12000	15410
Acc \uparrow	45.04	49.39	49.10	50.26
BWF \downarrow	20.82	17.56	18.39	15.87

Table 6. **Effect of memory budget (ActivityNet)**. We also test small memory regimes for the ActivityNet dataset using iCaRL. When compared against the iCaRL methods in [39], we still outperform their best performance when sampling a single frame for about every 4th video in the dataset.

Frame Selection	UCF101		ActivityNet	
	Acc \uparrow	BWF \downarrow	Acc \uparrow	BWF \downarrow
Random	95.70	2.59	50.26	15.87
Euclidean	95.34	2.37	50.25	16.37
Cosine Similarity	96.25	2.00	49.98	16.02

Table 7. **Analyzing Frame Selection Strategies**. We test different frame selection strategies to select a single frame that most closely resembles the full video features and then push it into the memory.

resembles the original video embedding. Intuitively, this will simplify the task of the secondary loss, favoring faster convergence, and will push into the working memory only those frames that contain relevant information for video classification.

We explore such an approach and use two metrics to test how similar a frame is to a video in the feature space. Our first metric is the Euclidean distance, where a smaller value indicates a higher similarity. The second metric is cosine similarity, where the larger the value, the higher the similarity. We perform the experiments using such a frame selection strategy over iCaRL on both UCF101 and ActivityNet datasets. Table 7 shows that there is no clear benefit in doing this frame selection. Regardless of the dataset or the selected metric, we could only get slight improvement for UCF101 using the cosine similarity. In fact, using this selection criterion for ActivityNet resulted in no improvement over random selection. Based on these observations, we adopted random sampling for frame selection in SMILE.

4.4. Effect of The Secondary Loss

As outlined in Section 3, we perform a secondary backward pass over the boring video data and supervise it with the original ground truth of the full-length video clip. To conclude the result section, we investigate the effect of this additional cross-entropy loss on the performance of our approach. We analyze the effect of this loss using iCaRL on both UCF101 and ActivityNet, and we assess the effectiveness of SMILE while enabling and disabling this additional loss. The results are summarized in Table 8. Overall, the

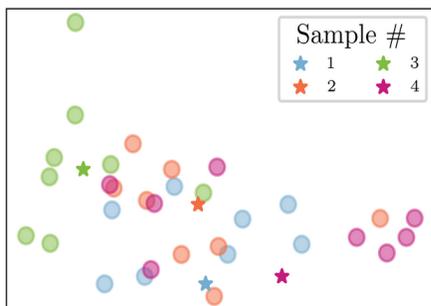
Secondary Loss	UCF101		ActivityNet	
	Acc \uparrow	BWF \downarrow	Acc \uparrow	BWF \downarrow
\times	84.40	0.78	46.40	9.36
\checkmark	95.70	2.59	50.26	15.87

Table 8. **The Effect of Secondary Loss**. We evaluate the effect of adding the secondary loss using iCaRL on both UCF101 and ActivityNet.

use of this additional loss leads to an improvement of 11.3% and 3.86% on UCF101 and ActivityNet, respectively.

We further explore the effect of including the secondary

Without Secondary CrossEntropy



With Secondary CrossEntropy

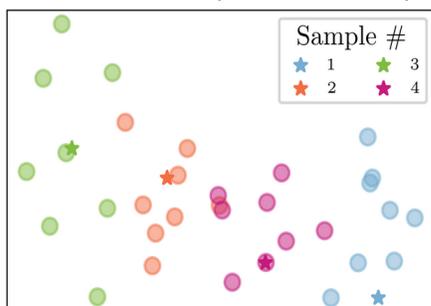


Figure 4. **Projection of Video and Boring Videos Features using PCA.** We show that training using iCaRL on UCF101 with an additional cross-entropy helps us learn a representation where the PCA projected features of a video and its corresponding boring videos are more similar. The projected features of the original video samples are marked with \star , whereas those of the boring videos are marked by a circle.

loss in SMILE by analyzing the PCA projected features of the original videos and the boring videos associated with them. In Figure 4, we show the arrangement of these PCA projections for the original video (shown as an \star in the plot) and the multiple boring videos (shown as circles) that can be built from it. Overall, we see that the plot at the bottom (with secondary cross-entropy) shows a spatial arrangement where the distance between the original video features and the multiple boring videos associated with the same video is smaller than the distance to boring videos of other original video clips. Without the secondary loss, the arrangement of such videos does not seem to follow any pattern.

5. Conclusion

In this work, we introduced SMILE, a simple yet effective replay mechanism for video continual learning. We showed that, by keeping a single video frame in memory, state-of-the-art results can be achieved in every dataset of the challenging vCLIMB benchmark. Despite reducing the working memory size by at least 87.7%, SMILE outperforms every other video CIL in the 10 and 20 task schedules proposed in vCLIMB. Remarkably, SMILE

achieves this performance by using standard image CIL methods in the video domain. This is possible as SMILE’s single frame memory enables the direct use of image CIL methods in the video domain.

Acknowledgments. This work is 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).

References

- [1] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 139–154, 2018. 2
- [2] Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3366–3375, 2017. 2
- [3] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–970, 2015. 5
- [4] Zhipeng Cai, Ozan Sener, and Vladlen Koltun. Online continual learning with natural distribution shifts: An empirical study with visual data. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8261–8270, 2021. 2
- [5] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 1, 4, 5
- [6] Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 532–547, 2018. 2
- [7] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385, 2021. 2
- [8] Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, Ziyang Wu, and Rama Chellappa. Learning without memorizing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5138–5146, 2019. 1
- [9] Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999. 1

- [10] Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013. [1](#)
- [11] Shreyank N Gowda, Marcus Rohrbach, and Laura Sevilla-Lara. Smart frame selection for action recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1451–1459, 2021. [6](#)
- [12] Tyler L Hayes, Kushal Kafle, Robik Shrestha, Manoj Acharya, and Christopher Kanan. Remind your neural network to prevent catastrophic forgetting. In *European Conference on Computer Vision*, pages 466–483. Springer, 2020. [2](#)
- [13] Kaiying He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [5](#)
- [14] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 831–839, 2019. [2](#)
- [15] Wenpeng Hu, Zhou Lin, Bing Liu, Chongyang Tao, Zhengwei Tao, Dongyan Zhao, Jinwen Ma, and Rui Yan. Overcoming catastrophic forgetting for continual learning via model adaptation. In *International conference on learning representations*, 2019. [2](#)
- [16] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014. [1](#)
- [17] 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. [2](#)
- [18] Timothée Lesort, Hugo Caselles-Dupré, Michael Garcia-Ortiz, Andrei Stoian, and David Filliat. Generative models from the perspective of continual learning. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2019. [2](#)
- [19] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. [2](#)
- [20] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7083–7093, 2019. [3](#)
- [21] Zhiqiu Lin, Jia Shi, Deepak Pathak, and Deva Ramanan. The clear benchmark: Continual learning on real-world imagery. *ArXiv*, abs/2201.06289, 2021. [2](#)
- [22] Vincenzo Lomonaco, Lorenzo Pellegrini, Pau Rodríguez López, Massimo Caccia, Qi She, Yu Chen, Quentin Jodelet, Ruiping Wang, Zheda Mai, David Vázquez, German Ignacio Parisi, Nikhil Churamani, Marc Pickett, Issam H. Laradji, and Davide Maltoni. Cvpr 2020 continual learning in computer vision competition: Approaches, results, current challenges and future directions. *Artif. Intell.*, 303:103635, 2022. [2](#)
- [23] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30, 2017. [2](#), [5](#)
- [24] Jiawei Ma, Xiaoyu Tao, Jianxing Ma, Xiaopeng Hong, and Yihong Gong. Class incremental learning for video action classification. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 504–508. IEEE, 2021. [2](#), [3](#)
- [25] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 67–82, 2018. [2](#)
- [26] Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost van de Weijer. Class-incremental learning: survey and performance evaluation on image classification. *arXiv preprint arXiv:2010.15277*, 2020. [1](#)
- [27] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989. [1](#)
- [28] Gabriele Merlin, Vincenzo Lomonaco, Andrea Cossu, Antonio Carta, and Davide Bacciu. Practical recommendations for replay-based continual learning methods. In *ICIAP Workshops*, 2022. [2](#)
- [29] Jaeyoo Park, Minsoo Kang, and Bohyung Han. Class-incremental learning for action recognition in videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13698–13707, 2021. [2](#), [3](#)
- [30] Yixuan Pei, Zhiwu Qing, Jun Cen, Xiang Wang, Shiwei Zhang, Yaxiong Wang, Mingqian Tang, Nong Sang, and Xueming Qian. Learning a condensed frame for memory-efficient video class-incremental learning. *arXiv preprint arXiv:2211.00833*, 2022. [2](#)
- [31] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *European conference on computer vision*, pages 524–540. Springer, 2020. [2](#)
- [32] 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*, pages 2001–2010, 2017. [2](#), [4](#), [5](#)
- [33] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. *Advances in Neural Information Processing Systems*, 32, 2019. [2](#)
- [34] 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](#)
- [35] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, pages 4548–4557. PMLR, 2018. [2](#)

- [36] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30, 2017. [2](#)
- [37] 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. [3](#), [4](#), [5](#)
- [38] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015. [1](#)
- [39] Andrés Villa, Kumail Alhamoud, Victor Escorcía, Fabian Caba, Juan León Alcázar, and Bernard Ghanem. vclimb: A novel video class incremental learning benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19035–19044, 2022. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [40] 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 transactions on pattern analysis and machine intelligence*, 41(11):2740–2755, 2018. [1](#), [3](#), [5](#)
- [41] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 374–382, 2019. [2](#), [4](#), [5](#)
- [42] Ju Xu and Zhanxing Zhu. Reinforced continual learning. *Advances in Neural Information Processing Systems*, 31, 2018. [2](#)
- [43] Junting Zhang, Jie Zhang, Shalini Ghosh, Dawei Li, Serafettin Tasci, Larry Heck, Heming Zhang, and C-C Jay Kuo. Class-incremental learning via deep model consolidation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1131–1140, 2020. [2](#)
- [44] Hanbin Zhao, Xin Qin, Shihao Su, Yongjian Fu, Zibo Lin, and Xi Li. When video classification meets incremental classes. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 880–889, 2021. [2](#), [3](#)