Federated Self-supervised Learning for Video Understanding

Yasar Abbas Ur Rehman^{1,*}, Yan Gao^{2,*}, Jiajun Shen¹, Pedro Porto Buarque de Gusmão², and Nicholas Lane^{2,3}

¹ TCL AI Lab, Hong Kong,
 ² University of Cambridge, United Kingdom,
 ³ Samsung AI Cambridge, United Kingdom

Abstract. The ubiquity of camera-enabled mobile devices has lead to large amounts of unlabelled video data being produced at the edge. Although various self-supervised learning (SSL) methods have been proposed to harvest their latent spatio-temporal representations for taskspecific training, practical challenges including privacy concerns and communication costs prevent SSL from being deployed at large scales. To mitigate these issues, we propose the use of Federated Learning (FL) to the task of video SSL. In this work, we evaluate the performance of current state-of-the-art (SOTA) video-SSL techniques and identify their shortcomings when integrated into the large-scale FL setting simulated with kinetics-400 dataset. We follow by proposing a novel federated SSL framework for video, dubbed FedVSSL, that integrates different aggregation strategies and partial weight updating. Extensive experiments demonstrate the effectiveness and significance of FedVSSL as it outperforms the centralized SOTA for the downstream retrieval task by 6.66% on UCF-101 and 5.13% on HMDB-51.

Keywords: Self-Supervised Learning, Federated Learning, Video Understanding, Model Aggregation

1 Introduction

A plethora of video content, often unlabelled and of private nature, is generated everyday from cameras in cellphones, tablets, and other mobile devices [10,21,31,1,38]. Being able to repurpose this data to solve various tasks in computer vision has been of great interest to researchers since the last decade [4,13,23,7,32,41,40,10]. Self-Supervised Learning (SSL) allows us to harvest these data contents by learning *intermediate* visual representations from unlabelled data, which can then be used as a starting point to solve specific downstream tasks (e.g., human action recognition [36], temporal action detection [45]). In practice, however, deploying SSL in its naïve form would require massive amounts of data to be sent to a centralized server for processing, posing significant concerns around privacy, [17,15], communication, and storage costs, ultimately limiting the technology to small datasets.

^{*} Equal contribution, authors ordered alphabetically.

A natural way to mitigate those issues is to combine SSL with the novel decentralized machine learning technique known as Federated Learning (FL) [29]. In FL, distributed population of edge devices collaboratively train a shared model while keeping their personal data private. Essentially, FL dilutes the burden of training across devices and avoids privacy and storage issues by not collecting users' data samples. The potential integration of video-SSL and FL into one coherent system offers lots of benefits in addition to data privacy. It enables large-scale decentralized feature learning from real-world data without requiring any costly and laborious data annotations. This can practically improve the performance and enables a vast majority of vision models for video applications [10], which have been under the shadow otherwise. Surprisingly, bearing such potential benefits there is no prior work that has so far studied video-SSL in FL.

In this paper, we conduct the first comprehensive study on video-SSL training in *cross-device* FL environment [20]. Our key findings from this study shows that: (1) The vanilla FL pretraining of video-SSL approaches are surprisingly not affected by the distribution of the data either being IID or non-IID. (2) Video-SSL with FL performs significantly better than the corresponding centralized SSL pretraining on video retrieval tasks and comparatively worse when the model is fined-tuned for action recognition. (3) We also show that the video SSL methods in FL settings are computationally efficient, regularized, and resilient to small-scale perturbation, compared to their centralized counterparts.

Based on the findings of our study of video-SSL with FL, we propose a novel federated learning framework, *FedVSSL*, designed specifically for video SSL. This framework allows to transceive only the backbone parameters of the video-SSL model during each communication round in FL. It then leverages a novel aggregation strategy, inspired by stochastic weighted averaging (SWA) [16], to aggregate and update the weights of the clients (performing video-SSL) at the server. FedVSSL obtains the state-of-the-art (SOTA) performance in video clip retrieval and competitive performance on action recognition against FedAvg and centralized video-SSL. The main contributions of this work are as follows:

- 1. We conduct the first systematic study of training video-SSL methods in FL *cross-device* settings with a large number of distributed clients. This establishes a baseline for naïvely implementing various video-SSL techniques using FL; shedding light on the basic problems of integrating video-SSL with FL into one coherent system.
- Based on the above, we propose a general FL framework, FedVSSL, based on SWA [16] for pretraining video-SSL methods in FL. Our method obtains SOTA performance in video clip retrieval and competitive performance on action recognition against FedAvg and centralized video-SSL.
- 3. We release our code and models on GitHub ⁴ to allow for reproducibility and stimulate further research in the field.

⁴ https://github.com/yasar-rehman/FEDVSSL

2 Background and Related Work

2.1 Video Self-supervised Representation Learning

Video-SSL approaches often rely on solving a pretext task [8] in an unsupervised fashion to learn representations that can be reused in solving other downstream tasks. The pretext tasks in video-SSL are either based on contrastive methods, non-contrastive methods, or a combination of both. Contrastive methods exploit the similarity between the two augmented views of the input samples using contrastive learning to learn the spatio-temporal representation from the unlabeled data [12,11,28,8,34]. On the other hand, non-contrastive approaches utilize specialized pretext tasks to generate pseudo signals to learn the spatio-temporal representation in a supervised manner, usually requiring just a single input [19,30,42,18,2,43,40,39,26].

Regardless of the specific SSL method being deployed, a recent study has shown that SSL models become more robust on a wide range of vision-based tasks when pretrained on real-world and uncurated data [10]. Bearing such tremendous potential, video-SSL models can extend the horizons of the many vision-based applications. However, the utility of these video-SSL models is significantly limited by the scale of the datasets available on the training server due to issues such as data privacy, communication cost and large data storage requirements. In this paper, we extend the utility of video-SSL methods beyond the centralized servers. This would allow harvesting of information from an unprecedented amount of user data, offering new opportunities to advance the quality and robustness of video-SSL models.

2.2 Federated Visual Representation Learning

Federated Learning (FL) [29] has received a lot of attention in recent years. In this new paradigm, the server now needs to *aggregate* incoming weights from clients to progressively produce better networks. In its original form [29], aggregation was performed using FedAvg, an *aggregation strategy* that generates a model via a weighted sum of clients' parameters. However, this simple aggregation method can perform particularly poorly in realistic scenarios where clients have very different data distributions, i. e. the available datasets are essentially not Independent and Identically Distributed (IID).

In the past few years, a number of aggregation strategies have been proposed to improve upon the original FedAvg. Authors in [9] introduce local-model training loss as a weighting coefficient for aggregation. Adaptive federated optimization approaches, proposed in [33], incorporate knowledge of past iterations by applying a separate gradient-based optimization on the server-side. Specific to SSL, authors in [44] suggest using contrastive-learning on the server-side based on the logits sent by the clients, thus imposing some privacy issues.

Besides choosing a specific aggregation method when training using SSL in FL, we must also decide which parts of the networks need to be aggregated as some of the weights are associated with latent representations (backbone)

while others are simply used for solving the pretext tasks (head). To this end, authors in [46] proposed FedU, that determines the update of the clients' model classification head based on the degree of backbone model divergence between the server and the clients.

It is important to note that the above-mentioned works were applied to smallscale *image-based* datasets [24,35] and that their proposed aggregation methods *rely on the usual class-label based definition of non-IID*. However, SSL should not depend on class-based labels, which prompts the question of whether such a definition of non-IID datasets bears any impact in video-SSL in FL. To the best of our knowledge, ours is the first FL SSL training framework tailored to *video* that is capable of producing SOTA models better than their centralized counterparts, regardless of the pretext class-label non-IID partitioning. We do this by correctly selecting only the backbone weights for aggregation.

3 Methodology

In this section, we provide the details of our systematic study on video-SSL using Federated Learning. Based on our findings, we follow by describing our proposed FedVSSL approach.

3.1 Federated Video-SSL System Design

We begin with a vanilla FL system that can be integrated with the video-SSL learning. The resulting FL video-SSL system is depicted in Fig. 1-Stage 1. We consider having n partitions $\{d_i\}_{i=1}^n$ of dataset D distributed among $\{c_i|c_i \in C\}_{i=1}^n$ decentralized clients in a Non-IID fashion. Each decentralized client learns intermediate features collaboratively by training the video-SSL approach on their respective local data partition for a few epochs before performing synchronization through the server. The synchronization includes receiving the clients' model parameters, performing model aggregation, and sending back the global model again to clients. More specifically, our training pipeline can be described as follows:

1. Each client c_i holds a set of local parameters $\{\theta_i^b, \theta_i^{p_t}\}$ for a backbone network $f(\theta^b)$ parameterized by the weights θ^b followed by a predictor head network $f(\theta^{p_t})$ parameterized by the weights θ^{p_t} . Note that all clients use the same model architecture but the parameter weights $\{\theta_i^b, \theta_i^{p_t}\}_{1 \le i \le n}$ can be non-consensus. During each FL round r, local video-SSL pretraining and synchronization will be conducted intersectively.

2. During the communication/synchronization step, each client would receive global values for the weights θ_g^b and $\theta_g^{p_t}$ aggregated on the server based on certain FL aggregation strategies.

3. During each round r, a random subset of clients, M, is selected to perform video-SSL training. Each participant $m \in M$ optimizes its local model's parameters $\{\theta_m^b, \theta_m^{p_t}\}$ for E number of local epochs.



Fig. 1: System overview: Stage-1 (left) represents the FL video-SSL pretraining. For the downstream tasks; video clip retrieval is seen in Stage-2(a) and action recognition is depicted in Stage-2(b)

The local training steps in each client completely follows the setup of the video-SSL algorithms, including the pre-training task together with the task-specific loss function design. The optimization is performed with SGD by minimizing the objective $J_m^{p_t}(\theta^b, \theta^{pt})$:

$$J_m^{p_t}(\theta^b, \theta^{pt}) = \mathbb{E}_{\mathbf{x} \sim d_m} \mathcal{L}_m(\mathbf{x}, y_{pt}, \theta^b, \theta^{pt}), \tag{1}$$

where (\mathbf{x}, y_{pt}) is the data with its corresponding pseudo labels for the specific pretext task p_t . At the end of the R rounds, the final global model parameters θ_g^b for the backbone network can be used to extract spatio-temporal representations from videos for downstream tasks. We then follow the common practice of evaluating the performance of representation learning as done in traditional video-SSL work [42,41]. The representations learned by the global model parameters θ_g^b are evaluated by solving the downstream tasks tsk as is shown in Fig. 1 - Stage 2. In this work, we consider two downstream tasks:

Action Recognition: A classification head (fully connected layer) represented by $f(\theta^{tsk})$ is attached following the final layer of $f(\theta_g^b)$, resulting in the model F_{tsk} parameterized by $\{\theta_g^b, \theta^{tsk}\}$. The F_{tsk} is then fine-tuned in two ways: (1) fine-tuning the whole network $\{\theta_g^b, \theta^{tsk}\}$, and (2) fine-tuning only the linear classification layer $f(\theta^{tsk})$, i.e., linear probe.

Video Clip Retrieval: For the video clip retrieval, we determine the video label by majority voting of K-nearest neighbor (KNN) by directly using the parameters learned by $f(\theta_g^b)$ after the FL video-SSL pretraining and followed the protocol as described in [42].

The above generic video federated-SSL system allows us to perform a systematic analysis of the video-SSL approaches and its implications in FL environment against a number of key factors:

Video-SSL in Vanilla FL Settings. In this scenario, our main objective is to evaluate the performance gap between pretraining video-SSL algorithms in

5

FL and centralized environment. This evaluation will then quantify how different video-SSL methods behave in FL settings with IID and Non-IID data.

Aggregation Strategies. The aggregation method plays an important role towards superior performance of the final feature representations learned in FL video-SSL pretraining. We intend to study the effects of different aggregation strategies on the feature representations learned during FL pretraining of video-SSL approaches and hence the downstream tasks.

Loss-Surface and Model Stability of Video-SSL in FL. These analyses enable us to perform a side-by-side comparison of the key properties of pretraining video-SSL approaches in FL settings against their centralized counterpart. For example, recent work indicates that averaging the weights of the neural networks leads to wider minima [16,27] which subsequently provides better generalization. A natural question arises as to whether the aggregation strategies, based on weight averaging, by performing FL video-SSL pretraining could provide wider minima? Moreover, do such wider minima potentially lead to model stability against small-scale perturbations that naturally occur during the communication between edge devices and the server in FL?

Training Efficiency and Communication Cost. The total time for pretraining video-SSL in FL settings depends on the number of epochs during local training on each client and the number of communication rounds between the clients and the server. It is important to compute the number of communication rounds necessary to achieve a target performance with FL video-SSL training in order to find an optimal balance between the performance of video-SSL and FL system resources.

3.2 Proposed Method

Based on the observations of the systematic study of the video-SSL approaches in cross-device FL settings, we proposed a novel FL framework, FedVSSL, designed specifically for video-SSL in FL environment. The key characteristics of this framework are as follows: (1) Transceive only the backbone model parameters θ_g^b between the server and the clients. (2) FedVSSL aggregation strategy for pretraining video-SSL approaches in FL cross-device settings.

Transceive only Backbone Weights. Since each c_i client model contains two modules $(f(\theta_{c_i}^b)$ and $f(\theta_{c_i}^{p_t}))$, we analyse the characteristics of both modules and hypothesize that it is beneficial to only upload and aggregate the parameters θ^b of the backbone model at the server (see Table 4). The backbone module learns hierarchical deep features from the local data, representing the role of encoder. The classification head is more representative of local data distribution, which learns data-oriented features. We argue that only aggregating the backbone parts could increase generalization of the model, while retaining the classification head and updating it locally could capture more characteristics of non-IID data from clients. More evidence from Fig. 5 validates our hypothesis.

FedVSSL Aggregation. The motivation behind proposing FedVSSL aggregation is two folds: First, to integrate the FL aggregation strategies based on weighted averaging under a common framework. Second, to induce the knowledge from the past global models while performing model parameters aggregation.

After local pretraining at round r, the server collects the participating clients' locally updated weights θ_m^b , and the local gradients $g^{(m)}(r)$ from client m at round r is computed as: $g^{(m)}(r) = \theta_m^b(r) - \theta_g^b(r-1)$. The update rule for θ_g^b of the global model F_g for round r can then be stated as follows:

$$\theta_g^b(r) = \frac{\left(\sum_{i=1}^{\beta} \theta_g^b(r-i)\right) + \tilde{\theta}_g^b(r)}{\beta + 1},\tag{2}$$

where $\tilde{\theta}_{g}^{b}(r) = \theta_{g}^{b}(r-1) - \eta_{s}\Delta_{r}$ and η_{s} is the server learning rate. The averaging is performed over $\beta + 1$ global models. Eq. 2 simply represents the stochastic weight averaging (SWA) [16] of the global models. As [16] reported, simple averaging of the multiple checkpoints of training models obtains better generalisation than conventional training. The Δ_{r} represents overall local gradients over m clients computed from a weighted combination of different aggregation strategies. Here we choose two aggregation strategies: FedAvg [29] and Loss [9], which aggregate the client models based on the number of samples or local training loss respectively. Then, Δ_{r} can be written as follows:

$$\Delta_r = \alpha \Delta_r^{Loss} + (1 - \alpha) \Delta_r^{FedAvg}.$$
(3)

The α in Eq. 3 controls the amount of the contribution of each aggregation strategy. Eq. 2 is the generalized representation of our FedVSSL aggregation rule. For example, setting $\beta = 0$ in Eq. 2 reduces the update rule to a weighted combination of FedAvg and Loss. On the other hand, setting $\alpha = 0$ or 1 reduces the update rule to running mean of FedAvg or Loss, respectively. The overall algorithm is summarised in Algo. 1.

4 Experiments and Results

In this section, we first describe our experiment setup in Sec. 4.1 & Sec. 4.2). We conduct a systematic analysis of the behavior of video-SSL models in vanilla FL settings in Sec. 4.3, followed by a discussion for the results of our proposed FedVSSL in Sec. 4.4.

4.1 Datasets

For the pretraining stages of all video-SSL approaches, we utilize kinetics-400 (K400) dataset [21] with 219k training samples distributed among 400 action classes. For downstream task, we utilize the UCF-101 [36] (UCF) dataset containing 13, 320 video samples for 101 action classes, and HMDB-51 [25] (HMDB) dataset with 7k video samples distributed among 51 action classes.

Kinectics-400 Non-IID. Each video sample in K400 comes from a different source, which conforms to the definition of non-IID based on video source-level.

Algorithm 1 Federated Video Self-supervised Learning (FedVSSL)

Input: $R, M, N, n_m, \eta_s, \eta_l, \alpha, \beta$ **Output**: θ^b Central server does: 1: for r = 1, ..., R do 2: Server randomly sample M clients. for each m in M do 3: $\theta_m^b(r), n_m, \mathcal{L}_m^{pt} = \mathbf{TrainLocally}(m, \theta_q^b(r))$ 4: $g^{(m)}(r) = \theta_m^b(r) - \theta_g^b(r-1)$ $\Delta_r^{FedAvg} = \sum_{m=1}^{M} \frac{n_m}{\sum_{m=1}^{m} n_m} g^{(m)}(r)$ $\Delta_r^{Loss} = \sum_{m=1}^{M} \frac{\exp\left(-\mathcal{L}_{(m)}^{pt}\right)}{\sum_{m=1}^{M} \exp\left(-\mathcal{L}_{(m)}^{pt}\right)} g^{(m)}(r)$ $\Delta_r = \alpha \Delta_r^{Loss} + (1-\alpha) \Delta_r^{FedAvg}$ Update global model weights $\tilde{\theta}_g^b(r) \leftarrow \theta_g^b(r-1) - \eta_s \Delta_r$. 5:6: 7: 8: 9: Compute $\theta_g^b(r) = \frac{(\sum_{i=1}^{\beta} \tilde{\theta}_g^b(r-i)) + \theta_g^b(r)}{\beta+1}$. 10: **TrainLocally** $(m, \theta_q^b(r))$: 1: for k = 1, ..., E do $\{\theta_m^b, \theta_m^{p_t}\}(k+1) \leftarrow \text{SSL}(\theta_a^b(k), \theta_m^{p_t}(k), \eta_l)$ based on Eq. 1. 2: 3: Upload $\theta_m^b, n_m, \mathcal{L}_m^{pt}$ to the server.

To make it more realistic, we generate the non-IID version of K400 based on actual class-labels [29]. We randomly partition the dataset into 100 shards to mimic the setting of having 100 disjoint clients participating in FL. Each client contains 8 classes resulting in each client having 2285 samples on average. Note that there is no overlap of samples between different clients.

4.2 Architecture and Implementation

Video-SSL Approaches. In this work, we consider three SOTA video-SSL algorithms, all of which propose to solve different pretext tasks for video representation learning. More specifically, VCOP [42] learns to determine the permutation order of shuffled clips, Speed [2,43,5] learns to predict the playback speed of videos, and CtP [39] predicts the positions and sizes of a synthetic image patch in a sequence of video frames. For all video-SSL approaches, we use R3D-18 [37] architecture as the backbone $f(\theta^b)$. The architecture choices of prediction heads for different pretext tasks and downstream tasks follow the settings in the original papers. It should be noted that our federated framework is agnostic to different architectures and video-SSL approaches. We develop the FL version of the video-SSL approaches considered in this work on top of Flower [3] federated learning platform by incorporating various video-SSL algorithms developed in MMCV framework [6,39]. Unless otherwise specified, we keep the settings of the video-SSL approaches as provided by [39] during the pretraining tasks and downstream tasks.

Table 1: Action recognition accuracy (Top-1%) and video clip retrieval accuracy (Top-1%, Top5%) on UCF and HMDB for three video-SSL methods. F-T represents fine-tune and L-P stands for linear probe. Δ represents the difference between centralized and corresponding FL performance."+" and "-" show % improvement and degradation respectively. * means the results are reproduced using the implementation in [39]

	\mathbf{A}	ction R	ecognit	ion	Video Clip Retrieval					
	UCF		HMDB		UCF		HMDB			
SSL Method	F-T	L-P	F-T	L-P	R@1	R@5	R@1	R@5		
VCOP*	71.29	24.93	38.56	13.53	15.52	28.26	8.11	22.22		
VCOP(Fed)	69.26	20.00	33.27	12.22	13.72	24.85	6.41	19.94		
Δ	(-2.03)	(-4.93)	(-5.29)	(-1.31)	(-1.8)	(-3.41)	(-1.7)	(-2.28)		
Speed*	81.15	29.32	47.58	14.90	16.84	36.58	6.93	21.05		
Speed(Fed)	73.16	35.63	38.43	21.57	21.97	41.61	10.98	28.30		
Δ	(-7.99)	(+6.31)	(-9.05)	(+6.67)	(+4.05)	(+3.94)	(+3.92)	(+7.25)		
CtP^*	86.20	48.14	57.00	30.65	29.0	47.30	11.80	30.10		
CtP(Fed)	81.95	46.13	49.15	28.63	29.29	48.90	13.66	32.42		
Δ	(-4.25)	(-2.01)	(-7.85)	(-2.02)	(+0.29)	(+1.6)	(+1.86)	(+2.32)		

FL Pretraining. We perform the FL pretraining of the video-SSL pretext-task p_t using Algo. 1. The local pretraining on each client lasts for E epochs per FL round R, where we set E = 1 in our experiments. We set the total number of rounds R to 540 to ensure that each client acquires sufficient participation during FL pretraining. The selection of the number of E and R is based on our empirical observations. Each round, we randomly select M = 5 clients from the pool of 100 clients to participate in training and each client trains its local model using SGD optimizer without momentum. We set a constant learning rate of 0.01 for CtP and Speed and 0.001 for VCOP. Weight decay is set to 10^{-4} and training batch-size is set to 4. On the server side, in addition to our proposed method FedVSSL, we consider three existing aggregation strategies including FedAvg, Loss and FedU.

Downstream Tasks. For the fine-tuning stage of action recognition, we follow the configuration in CtP framework [39]. The $F_{d_{tsk}}$ is fine-tuned using the SGD optimizer with an initial learning rate of 0.01, momentum of 0.9, and weight decay of 5×10^{-4} . The learning rate is decayed by a factor of 0.1 after 60 and 120 epochs, respectively. The batch-size is set 32, and the fine-tuning stage lasts for 150 epochs. For linear probe, we keep the same settings as for the fine-tuning of the whole network, except that we train only $f(\theta^{d_{tsk}})$ layer of $F_{d_{tsk}}$ for 100 epochs. The learning rate is decayed by the factor of 0.1 after 60 and 80 epochs. For the video clip retrieval task, we follow the approach described in Sec. 3.1.

4.3 Video-SSL in Vanilla FL Settings

Here we investigate the performance of video-SSL in vanilla FL settings against the key factors listed in Sec.3.1.

			Fine-	tuning				
			UCF	HMDB	U	CF	HMDE	
Method	C/Cpc	Data	Top-1	Top-1	R@1	R@5	R@1	R@5
	100/-	IID	81.92	48.49	29.42	47.90	13.80	34.56
$\operatorname{CtP}(\operatorname{Fed.})$	100/8	$\operatorname{Non-IID}$	81.95	49.15	29.29	48.90	13.66	32.42
	100/4	$\operatorname{Non-IID}$	81.15	47.78	29.18	48.37	14.70	32.94
CtP(Cent.)	-	IID	86.2	57.00	29.00	47.30	11.80	30.10

Table 2: Action recognition and video clip retrieval accuracies (%) on UCF and HMDB for the federated CtP models pretrained with one local epoch per round. C represents the number of clients, and Cpc stands for the number of classes per client

Centralized vs. Federated Video-SSL. In this experiment, we draw a first investigation for the performance of video-SSL approaches using FedAvg in centralized and vanilla FL settings with Non-IID video data. We report this comparison in terms of three downstream tasks, i.e., fine-tuning, linear probe, video clip retrieval, on UCF and HMDB datasets (Table 1). First, CtP obtains the best performance for all tasks in both centralized and FL settings. Second, when the trained network is fine-tuned for the action recognition task, the centralized video-SSL approaches perform better compared to their corresponding FL counterparts. However, the degradation in the performance is not drastic. We conjecture that it is caused by the smoother and flatter manifold in the models in FL settings, which would be more challenging to fine-tune. Third, the linear-probe results for action recognition show mixed results due to the fact that only the classification head participates in the training. Finally, the video clip retrieval results are more competitive with the FL version of Speed and CtP, which achieve better performance than their centralized counterparts.

Given the fact that the video-SSL benefits from the large-scale datasets and the performance degradation of the FL version of the video-SSL is acceptable (even in the Non-IID case with FedAvg), it makes the FL a natural choice for video-SSL approaches.

Performance with IID vs. Non-IID Data. Conventional FL methods are designed to solve a supervised/semi-supervised learning task within an IID/Non-IID data distribution based on the actual class-labels. To understand the impact of data distribution on federated video-SSL training, we simulate IID and Non-IID settings based on class-labels and compare the performance of downstream task in Table 2. The Non-IID versions of K-400 are generated with two variations of the distribution of the samples among 100 clients, with 4 and 8 classes per client respectively. We report the fine-tuning and video clip retrieval accuracy on UCF and HMDB, by pretraining CtP video-SSL approach on all settings.

One can see from Table 2 that pretraining the CtP video-SSL approach using standard FedAvg with IID and different degrees of Non-IID levels achieve comparatively similar performance on the fine-tuning and video clip retrieval task. This could be explained by the generation process of IID/Non-IID data based the actual class-labels. The video-SSL methods learn representations by gener-

Table 3: Video clip retrieval accuracies (%) and fine-tuning accuracies (%) on UCF and HMDB for CtP video-SSL approach using various aggregation strategies. The SSL pretraining is performed on K400 (Non-IID). Cent[†] represents the centralized training for 27 epochs which equals to 540 rounds in our FL setting

		Retr	Fine-tuning			
	UCF		HMDB		UCF	HMDB
Method	R@1	R@5	R@1	R@5	Top-1	Top-1
$\begin{array}{c} \text{Centralized} \\ \text{Centralized}^{\dagger} \end{array}$	$29.00 \\ 27.65$	$47.3 \\ 47.67$	$\begin{array}{c} 11.80\\ 12.81 \end{array}$	$30.1 \\ 31.05$	$86.20 \\ 83.64$	$57.00 \\ 53.73$
FedAvg (Baseline) Loss-based FedU	32.62 32.54 34.07	50.41 50.01 52.29	16.54 14.44 14.90	35.29 34.97 36.67	79.91 79.43 80.17	52.88 50.63 53.73

ating pseudo labels based on the pretext task it solves, which is independent of actual class-labels. Hence, the IID/Non-IID data distribution has a slight impact to the SSL model training. In addition, we find that there exists a degradation in the fine-tuning performance of federated CtP video-SSL compared to its centralized counterpart. Interestingly, the FL version of CtP video-SSL approach gives better video clip retrieval performance when compared to its centralized counterpart with both IID and different degrees of Non-IID levels.

Performance of Aggregation Strategies. The aggregation method often plays an important role towards superior performance of models trained in FL environments. In this experiment, we investigate the impact of FL aggregation strategies on the pretraining by evaluating the final performance of video-SSL approaches in Non-IID FL settings. In Table 3, we show the performance of CtP video-SSL approach on UCF and HMDB against a range of aggregation strategies that include FedAvg [29], Loss [9], and FedU [46]. It can be observed that FedAvg performs better on HMDB and obtains similar performance with Loss method on UCF. Additionally, except for retrieval accuracy on HMDB, FedU outperforms others on both video clip retrieval and fine-tuning mainly due to its dynamic aggregation mechanism.

Loss Surface and Model Stability of Video-SSL in FL. To understand why the federated video-SSL models gain higher retrieval accuracy than centralized models, we further analyze the loss landscape around the pretrained model both in centralized and FL settings. To compute this, we utilize the filter normalization method as proposed in [27]. The results are shown in Fig. 2 for CtP video-SSL approach. We find that the loss landscape of model pretrained with FedAvg is flatter than the model pretrianed with the centralized video-SSL. The width of the optima is critically related to generalization, which enables the model to converge in a point centered in this region [16,22,14]. This often leads to slightly worse training loss but substantially better test accuracy.

We then explore whether such wider optima could increase model stability against small perturbations. To achieve this, we perturb the weights of



Fig. 2: Loss landscape of the final $f(\theta^b)$ pretrained in a centralized(left) and FL with FedAvg(right) settings



Fig. 3: Top1% action retrieval accuracy on UCF and HMDB by adding perturbation to all the weights of the pretrained $f(\theta^b)$ network. The perturbation is sampled from a normal distribution N(0, 1) and multiplied by a factor k

the backbone pretrained with centralized and federated video-SSL approach. The perturbations are sampled from a uniform normal distribution with zero mean and unit variance, i.e., $\mathcal{N}(0, 1)$. We start with the perturbation level of 0 and incrementally increase the level of perturbation by a factor of 0.1, i.e., $(\{\mathcal{N}(0,1) \times x | 0 \leq x \leq 0.5\})$. The results are shown in Fig. 3 for the top-1% video clip retrieval accuracy for the centralized video-SSL approach and its corresponding FL counterparts, on UCF and HMDB datasets. One can see from that as the level of perturbation is increased from 0 to 0.1 the top-1% accuracy drops significantly for both centralized CtP video-SSL and its FL counterparts. However, as the level for perturbation is further increased we find that the FL versions of the video-SSL approach show good stability compared to its centralized counterpart on both datasets. Overall, the federated training boosts the generalization and stability of video-SSL models.

Training Efficiency and Communication Cost. In this experiment, we analyze the computational efficiency of pretraining CtP approach in both FL and centralized settings. We find that our FL pretraining of CtP for 540 rounds, with 100 clients and 5% client sampling rate (in ideal scenarios), is equivalent to 27 epochs (30%) of centralized pretraining. This results in 62.5% of GPU time saving compared to the centralized pretraining of the CtP that lasts for 90 epochs. Additionally, compared with the performance of centralized pretraining of CtP for 27 epochs, our federated CtP model achieves significant boost in video clip retrieval performance while providing competitive performance on fine-tuning as shown in Table 3. We further show the number of communication rounds



Fig. 4: Top 1% retrieval accuracy (%) w.r.t communication rounds for our proposed FL methods on UCF(left) and HMDB (right). Our proposed methods require less than 200 rounds to reach the centralized target accuracy on both datasets

required by federated video-SSL models to achieve the centralized target video clip retrieval accuracy. One can see from Fig. 4 that the FL model trained with FedAvg requires less than 100 rounds to reach the centralized target accuracy on UCF and HMDB datasets. Indeed, our proposed FedVSSL shows even better better performance and convergence behaviour.

4.4 Results on FedVSSL

In this section, we investigate the performance of our proposed FVSSL method, conduct ablation studies by varying α and β in Eq. 2 and Eq. 3, and report the results on UCF and HMDB in Table 4. One can see that the performance of all proposed methods are competitive, which demonstrates that the learned representations are qualified for the downstream applications.



Fig. 5: Standard deviations of the L_2 difference between the global model weights and the locally trained model weights at each round of FL video-SSL pretraining with FedAvg. Both backbone θ^b and prediction head θ^{p_t} are aggregated on the server

First, all models trained with FedVSSL provide superior retrieval performance. Concretely, (FedVSSL with $\alpha = 0.9, \beta = 0$) we improve the video clip retrieval performance over FedAvg baseline by 2.88% (top1%) and 3.86% (top5%)

	Retrieval			Fine-tuning		Linear-probe		
	UCF		HMDB		UCF	HMDB	UCF	HMDB
Method	R@1	R@5	R@1	R@5	Top-1	Top-1	Top-1	Top-1
Centralized	29.00	47.3	11.80	30.1	86.20	57.00	48.14	30.65
FedAvg (Baseline)	32.62	50.41	16.54	35.29	79.91	52.88	45.31	31.44
$\operatorname{FedVSSL}(\alpha = 0, \beta = 0)$	34.34	51.71	15.82	36.01	79.91	52.94	47.95	31.12
$\operatorname{FedVSSL}(\alpha = 1, \beta = 0)$	34.23	52.21	16.73	38.30	79.14	51.11	47.90	29.48
$\operatorname{FedVSSL}(\alpha = 0, \beta = 1)$	35.61	52.18	16.93	37.78	79.43	51.90	47.66	30.00
$\operatorname{FedVSSL}(\alpha = 1, \beta = 1)$	35.66	52.34	16.41	36.93	78.99	51.18	48.93	31.44
$\text{FedVSSL}(\alpha = 0.9, \beta = 0)$	35.50	54.27	16.27	37.25	80.62	53.14	50.36	32.68
$\text{FedVSSL}(\alpha = 0.9, \beta = 1)$	35.34	52.34	16.93	37.39	79.41	51.50	50.30	32.42

Table 4: Video clip retrieval accuracies (%) and fine-tuning accuracies (%) on UCF101 and HMDB51 for CtP video-SSL approach using our proposed FedVSSL methods. The SSL pretraining is performed on K400 (Non-IID)

on UCF, and 1.96% (top5%) on HMDB. We obtained 0.71% and 0.26% improvement (top1%) in the category of fine-tuning on UCF and HMDB, respectively compared to FedAvg baseline. In the category of Linear-probe, we obtained 5.05% and 1.24% improvement (top1%) on UCF and HMDB, respectively compared to the FedAvg baseline. Second, the component of SWA ($\beta = 1$)in Fed-VSSL has a distinct benefit on the improvement of retrieval performance. Third, the top-1% retrieval accuracy of FedVSSL($\alpha = 0, \beta = 0$) outperforms the FedAvg baseline by 1.72% on UCF dataset, which highlights the benefits of only updating backbone.

One can see from Fig. 5 that the standard deviation of the L_2 distance for the backbone model parameter is more consistent throughout the FL pretraining compared to that of prediction model parameters. This indicates that the backbone weights cause less divergence in FL pretraining of video-SSL which may provide significant efficiency and performance boost in the FL scenarios with stringent communication budget.

5 Conclusions

In this paper, we presented the first systematic study on video-SSL for FL. Our key findings include (1) the importance of aggregating just the backbone network and that (2) non-IID definition based on class-labels bears no impact on pretext training as they are not used by that task. Based on these findings, we proposed FedVSSL, an aggregation strategy tailored to video. FedVSSL was able to outperform the centralized SOTA for the downstream retrieval task by 6.66% on UCF-101 dataset and by 5.13% on HMDB-51 using non-contrastive methods. We hope this work will enable future research towards further combining FL and SSL for video representation learning. These are complementary technologies that can together harvest rich visual information from edge devices while still preserving user privacy.

15

References

- 1. Aytar, Y., Vondrick, C., Torralba, A.: Soundnet: Learning sound representations from unlabeled video. Advances in neural information processing systems **29** (2016)
- Benaim, S., Ephrat, A., Lang, O., Mosseri, I., Freeman, W.T., Rubinstein, M., Irani, M., Dekel, T.: Speednet: Learning the speediness in videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9922–9931 (2020)
- Beutel, D.J., Topal, T., Mathur, A., Qiu, X., Parcollet, T., Lane, N.D.: Flower: A friendly federated learning research framework. arXiv preprint arXiv:2007.14390 (2020)
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
- 5. Cho, H., Kim, T., Chang, H.J., Hwang, W.: Self-supervised visual learning by variable playback speeds prediction of a video. IEEE Access (2021)
- Contributors, M.: MMCV: OpenMMLab computer vision foundation. https://github.com/open-mmlab/mmcv (2018)
- Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: Proceedings of the IEEE international conference on computer vision. pp. 1422–1430 (2015)
- Feichtenhofer, C., Fan, H., Xiong, B., Girshick, R., He, K.: A large-scale study on unsupervised spatiotemporal representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3299– 3309 (2021)
- Gao, Y., Parcollet, T., Zaiem, S., Fernandez-Marques, J., de Gusmao, P.P., Beutel, D.J., Lane, N.D.: End-to-end speech recognition from federated acoustic models. arXiv preprint arXiv:2104.14297 (2021)
- Goyal, P., Duval, Q., Seessel, I., Caron, M., Singh, M., Misra, I., Sagun, L., Joulin, A., Bojanowski, P.: Vision models are more robust and fair when pretrained on uncurated images without supervision. arXiv preprint arXiv:2202.08360 (2022)
- Han, T., Xie, W., Zisserman, A.: Video representation learning by dense predictive coding. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 0–0 (2019)
- Han, T., Xie, W., Zisserman, A.: Self-supervised co-training for video representation learning. Advances in Neural Information Processing Systems 33, 5679–5690 (2020)
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9729–9738 (2020)
- Hochreiter, S., Schmidhuber, J.: Flat minima. Neural computation 9(1), 1–42 (1997)
- Hu, Z., Xie, H., Yu, L., Gao, X., Shang, Z., Zhang, Y.: Dynamic-aware federated learning for face forgery video detection. ACM Transactions on Intelligent Systems and Technology (TIST) (2022)
- Izmailov, P., Podoprikhin, D., Garipov, T., Vetrov, D., Wilson, A.G.: Averaging weights leads to wider optima and better generalization. In: 34th Conference on Uncertainty in Artificial Intelligence 2018, UAI 2018. pp. 876–885. Association For Uncertainty in Artificial Intelligence (AUAI) (2018)

- 16 Y. Rehman and Y. Gao et al.
- 17. Jain, A.K., Deb, D., Engelsma, J.J.: Biometrics: Trust, but verify. arXiv preprint arXiv:2105.06625 (2021)
- Jenni, S., Meishvili, G., Favaro, P.: Video representation learning by recognizing temporal transformations. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVIII 16. pp. 425–442. Springer (2020)
- Jing, L., Yang, X., Liu, J., Tian, Y.: Self-supervised spatiotemporal feature learning via video rotation prediction. arXiv preprint arXiv:1811.11387 (2018)
- Kairouz, P., McMahan, H.B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A.N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al.: Advances and open problems in federated learning. arXiv preprint arXiv:1912.04977 (2019)
- Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017)
- Keskar, N.S., Mudigere, D., Nocedal, J., Smelyanskiy, M., Tang, P.T.P.: On largebatch training for deep learning: Generalization gap and sharp minima. arXiv preprint arXiv:1609.04836 (2016)
- Kolesnikov, A., Zhai, X., Beyer, L.: Revisiting self-supervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1920–1929 (2019)
- 24. Krizhevsky, A.: Learning multiple layers of features from tiny images (2009)
- Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: HMDB: a large video database for human motion recognition. In: Proceedings of the International Conference on Computer Vision (ICCV) (2011)
- Lee, H.Y., Huang, J.B., Singh, M., Yang, M.H.: Unsupervised representation learning by sorting sequences. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 667–676 (2017)
- Li, H., Xu, Z., Taylor, G., Studer, C., Goldstein, T.: Visualizing the loss landscape of neural nets. In: Proceedings of the 32nd International Conference on Neural Information Processing Systems. pp. 6391–6401 (2018)
- Li, T., Wang, L.: Learning spatiotemporal features via video and text pair discrimination. arXiv preprint arXiv:2001.05691 (2020)
- McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: Artificial intelligence and statistics. pp. 1273–1282. PMLR (2017)
- Misra, I., Zitnick, C.L., Hebert, M.: Shuffle and learn: unsupervised learning using temporal order verification. In: European Conference on Computer Vision. pp. 527–544. Springer (2016)
- Park, H., Sjosund, L., Yoo, Y., Monet, N., Bang, J., Kwak, N.: Sinet: Extreme lightweight portrait segmentation networks with spatial squeeze module and information blocking decoder. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2066–2074 (2020)
- Piergiovanni, A., Angelova, A., Ryoo, M.S.: Evolving losses for unsupervised video representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 133–142 (2020)
- Reddi, S.J., Charles, Z., Zaheer, M., Garrett, Z., Rush, K., Konečný, J., Kumar, S., McMahan, H.B.: Adaptive federated optimization. In: International Conference on Learning Representations (2020)
- Romijnders, R., Mahendran, A., Tschannen, M., Djolonga, J., Ritter, M., Houlsby, N., Lucic, M.: Representation learning from videos in-the-wild: An object-centric

approach. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 177–187 (2021)

- 35. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. International journal of computer vision 115(3), 211–252 (2015)
- Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild (2012)
- 37. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A closer look at spatiotemporal convolutions for action recognition. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 6450–6459 (2018)
- Vondrick, C., Pirsiavash, H., Torralba, A.: Anticipating visual representations from unlabeled video. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 98–106 (2016)
- Wang, G., Zhou, Y., Luo, C., Xie, W., Zeng, W., Xiong, Z.: Unsupervised visual representation learning by tracking patches in video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2563– 2572 (2021)
- Wang, J., Jiao, J., Liu, Y.H.: Self-supervised video representation learning by pace prediction. In: European conference on computer vision. pp. 504–521. Springer (2020)
- Wang, X., Gupta, A.: Unsupervised learning of visual representations using videos. In: Proceedings of the IEEE international conference on computer vision. pp. 2794–2802 (2015)
- 42. Xu, D., Xiao, J., Zhao, Z., Shao, J., Xie, D., Zhuang, Y.: Self-supervised spatiotemporal learning via video clip order prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10334–10343 (2019)
- Yao, Y., Liu, C., Luo, D., Zhou, Y., Ye, Q.: Video playback rate perception for self-supervised spatio-temporal representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6548– 6557 (2020)
- Zhang, F., Kuang, K., You, Z., Shen, T., Xiao, J., Zhang, Y., Wu, C., Zhuang, Y., Li, X.: Federated unsupervised representation learning. arXiv preprint arXiv:2010.08982 (2020)
- Zhao, Y., Xiong, Y., Wang, L., Wu, Z., Tang, X., Lin, D.: Temporal action detection with structured segment networks. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2914–2923 (2017)
- Zhuang, W., Gan, X., Wen, Y., Zhang, S., Yi, S.: Collaborative unsupervised visual representation learning from decentralized data. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4912–4921 (2021)