

# Continual 3D Convolutional Neural Networks for Real-time Processing of Videos

Lukas Hedegaard<sup>✉</sup> and Alexandros Iosifidis<sup>✉</sup>

Department of Electrical and Computer Engineering, Aarhus University, Denmark  
{lh, ai}@ece.au.dk

**Abstract.** We introduce *Continual* 3D Convolutional Neural Networks (*Co3D* CNNs), a new computational formulation of spatio-temporal 3D CNNs, in which videos are processed frame-by-frame rather than by clip. In online tasks demanding frame-wise predictions, *Co3D* CNNs dispense with the computational redundancies of regular 3D CNNs, namely the repeated convolutions over frames, which appear in overlapping clips. We show that *Continual* 3D CNNs can reuse preexisting 3D-CNN weights to reduce the per-prediction floating point operations (FLOPs) in proportion to the temporal receptive field while retaining similar memory requirements and accuracy. This is validated with multiple models on Kinetics-400 and Charades with remarkable results: *CoX3D* models attain state-of-the-art complexity/accuracy trade-offs on Kinetics-400 with 12.1–15.3× reductions of FLOPs and 2.3–3.8% improvements in accuracy compared to regular X3D models while reducing peak memory consumption by up to 48%. Moreover, we investigate the transient response of *Co3D* CNNs at start-up and perform extensive benchmarks of on-hardware processing characteristics for publicly available 3D CNNs.

**Keywords:** 3D CNN, Human Activity Recognition, Efficient, Stream Processing, Online Inference, Continual Inference Network.

## 1 Introduction

Through the availability of large-scale open-source datasets such as ImageNet [37] and Kinetics [25], [4], deep, over-parameterized Convolutional Neural Networks (CNNs) have achieved impressive results in the field of computer vision. In video recognition specifically, 3D CNNs have led to multiple breakthroughs in the state-of-the-art [3], [43], [11], [10]. Despite their success in competitions and benchmarks where only prediction quality is evaluated, computational cost and processing time remains a challenge to the deployment in many real-life use-cases with energy constraints and/or real-time needs. To combat this general issue, multiple approaches have been explored. These include computationally efficient architectures for image [17], [48], [42] and video recognition [28], [10], [49], pruning of network weights [6], [13], [14], knowledge distillation [16], [47], [36], and network quantisation [19], [2], [12].

The contribution in this paper is complementary to all of the above. It exploits the computational redundancies in the application of regular spatio-temporal 3D CNNs to a continual video stream in a sliding window fashion (Fig. 2). This redundancy was also explored recently [26], [39] using specialised architectures. However, these are not weight-compatible with regular 3D CNNs. We present a weight-compatible reformulation of the 3D CNN and its components as a *Continual* 3D Convolutional Neural Network (*Co3D* CNN). *Co3D* CNNs process input videos frame-by-frame rather than clip-wise and can reuse the weights of regular 3D CNNs, producing identical outputs for networks without temporal zero-padding. Contrary to most deep learning papers, the work presented here needed no training; our goal was to validate the efficacy of converting regular 3D CNNs to Continual CNNs directly, and to explore their characteristics in the online recognition domain. Accordingly, we perform conversions from five 3D CNNs, each at different points on the accuracy/speed pareto-frontier, and evaluate their frame-wise performance. While there is a slight reduction in accuracy after conversion due to zero-padding in the regular 3D CNNs, a simple network modification of extending the temporal receptive field recovers and improves the accuracy significantly *without* any fine-tuning at a negligible increase in computational cost. Furthermore, we measure the transient network response at start-up, and perform extensive benchmarking on common hardware and embedded devices to gauge the expected inference speeds for real-life scenarios. Full source code is available at <https://github.com/lukashedegaard/co3d>.

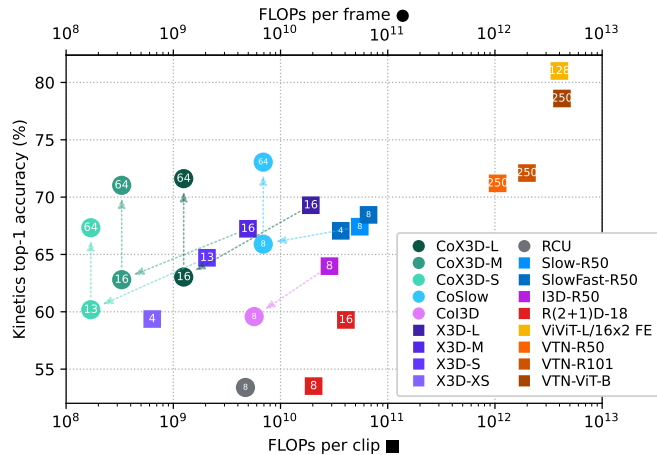


Fig. 1: **Accuracy/complexity trade-off** for *Continual* 3D CNNs and recent state-of-the-art methods on Kinetics-400 using 1-clip/frame testing. ■ FLOPs per *clip* are noted for regular networks, while ● FLOPs per *frame* are shown for the Continual 3D CNNs. Frames per clip / global average pool size is noted in the representative points. Diagonal and vertical arrows indicate a direct weight transfer from regular to Continual 3D CNN and an extension of receptive field.

## 2 Related Works

### 2.1 3D CNNs for video recognition

Convolutional Neural Networks with spatio-temporal 3D kernels may be considered the natural extension of 2D CNNs for image recognition to CNNs for video recognition. Although they did not surpass their 2D CNN + RNN competitors [7], [21] initially [20], [23], [44], arguably due to a high parameter count and insufficient dataset size, 3D CNNs have achieved state-of-the-art results on Human Action Recognition tasks [3], [43], [11] since the Kinetics dataset [25] was introduced. While recent large-scale Transformer-based methods [1], [32] have become leaders in terms of accuracy, 3D CNNs still achieve state-of-the-art accuracy/complexity trade-offs. Nevertheless, competitive accuracy comes with high computational cost, which is prohibitive to many real-life use cases.

In image recognition, efficient architectures such as MobileNet [17], ShuffleNet [48], and EfficientNet [42] attained improved accuracy-complexity trade-offs. These architectures were extended to the 3D-convolutional versions 3D-MobileNet [28], 3D-ShuffleNet [28] and X3D [10] ( $\approx$ 3D-EfficientNet) with similarly improved pareto-frontier in video-recognition tasks. While these efficient 3D CNNs work well for offline processing of videos, they are limited in the context of online processing, where we wish to make updates predictions for each frame; real-time processing rates can only be achieved with the smallest models at severely reduced accuracy. 3D CNNs suffer from the restriction that they must process a whole “clip” (spatio-temporal volume) at a time. When predictions are needed for each frame, this imposes a significant overhead due to repeated computations. In our work, we overcome this challenge by introducing an alternative computational scheme for spatio-temporal convolutions, -pooling, and -residuals, which lets us compute 3D CNN outputs frame-wise (continually) and dispose of the redundancies produced by regular 3D CNNs.

### 2.2 Architectures for online video recognition

A well-explored approach to video-recognition [7], [21], [22], [40] is to let each frame pass through a 2D CNN trained on ImageNet in one stream alongside a second stream of Optical Flow [9] and integrate these using a recurrent network. Such architectures requires no network modification for deployment in online-processing scenarios, lends themselves to caching [46], and are free of the computational redundancies experienced in 3D CNNs. However, the overhead of Optical Flow and costly feature-extractors pose a substantial disadvantage.

Another approach is to utilise 3D CNNs for feature extraction. In [31], spatio-temporal features from non-overlapping clips are used to train a recurrent network for hand gesture recognition. In [27], a 3D CNN processes a sliding window of the input to perform spatio-temporal action detection. These 3D CNN-based methods have the disadvantage of either not producing predictions for each input frame [31] or suffering from redundant computations from overlapping clips [27].

Massively Parallel Video Networks [5] split a DNN into depth-parallel sub-networks across multiple computational devices to improve online multi-device parallel processing performance. While their approach treats network layers as atomic operations and doesn’t tackle the fundamental redundancy of temporal convolutions, *Continual* 3D CNNs reformulate the network layers, remove redundancy, and accelerate inference on single devices as well.

Exploring modifications of the spatio-temporal 3D convolution, the Recurrent Convolutional Unit (RCU) [39] replaces the 3D convolution by aggregating a spatial 2D convolution over the current input with a 1D convolution over the prior output. Dissected 3D CNNs [26] (D3D) cache the  $1 \times n_H \times n_W$  frame-level features in network residual connections and aggregate them with the current frame features via  $2 \times 3 \times 3$  convolutions. Like our proposed *Continual* 3D CNNs, both RCU and D3D are causal and operate frame-by-frame. However, they are speciality architectures, which are incompatible with pre-trained 3D CNNs, and must be trained from scratch. We reformulate spatio-temporal convolutions in a one-to-one compatible manner, allowing us to reuse existing model weights.

### 3 Continual Convolutional Neural Networks

#### 3.1 Regular 3D-convolutions lead to redundancy

Currently, the best performing architectures (e.g., X3D [10] and SlowFast [11]) employ variations on 3D convolutions as their main building block and perform predictions for a spatio-temporal input volume (video-clip). These architectures achieve high accuracy with reasonable computational cost for predictions on clips in the offline setting. They are, however, ill-suited for online video classification, where the input is a continual stream of video frames and a class prediction is needed for each frame. For regular 3D CNNs processing clips of  $m_T$  frames to be used in this context, prior  $m_T - 1$  input frames need to be stored between temporal time-steps and assembled to form a new video-clip when the next frame is sampled. This is illustrated in Fig. 2.

Recall the computational complexity for a 3D convolution:

$$\Theta([k_H \cdot k_W \cdot k_T + b] \cdot c_I \cdot c_O \cdot n_H \cdot n_W \cdot n_T), \quad (1)$$

where  $k$  denotes the kernel size,  $T$ ,  $H$ , and  $W$  are time, height, and width dimension subscripts,  $b \in \{0, 1\}$  indicates whether bias is used, and  $c_I$  and  $c_O$  are the number of input and output channels. The size of the output feature map is  $n = (m + 2p - d \cdot (k - 1) - 1)/s + 1$  for an input of size  $m$  and a convolution with padding  $p$ , dilation  $d$ , and stride  $s$ . During online processing, every frame in the continual video-stream will be processed  $n_T$  times (once for each position in the clip), leading to a redundancy proportional with  $n_T - 1$ . Moreover, the memory-overhead of storing prior input frames is

$$\Theta(c_I \cdot m_H \cdot m_W \cdot [m_T - 1]), \quad (2)$$

and during inference the network has to transiently store feature-maps of size

$$\Theta(c_O \cdot n_H \cdot n_W \cdot n_T). \quad (3)$$

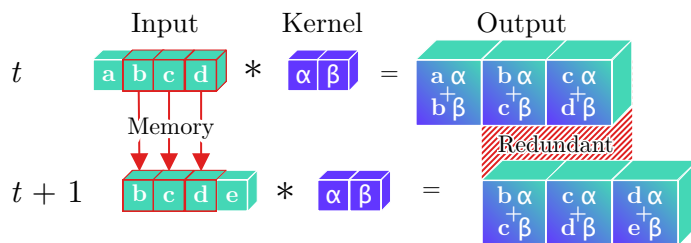


Fig. 2: **Redundant computations** for a temporal convolution during online processing, as illustrated by the repeated convolution of inputs (green  $\mathbf{b}, \mathbf{c}, \mathbf{d}$ ) with a kernel (blue  $\alpha, \beta$ ) in the temporal dimension. Moreover, prior inputs ( $\mathbf{b}, \mathbf{c}, \mathbf{d}$ ) must be stored between time-steps for online processing tasks.

### 3.2 Continual Convolutions

We can remedy the issue described in Sec. 3.1 by employing an alternative sequence of computational steps. In essence, we reformulate the repeated convolution of a (3D) kernel with a (3D) input-clip that continually shifts along the temporal dimension as a *Continual Convolution (CoConv)*, where all convolution computations (bar the final sum) for the (3D) kernel with each (2D) input-frame are performed in one time-step. Intermediary results are stored as states to be used in subsequent steps, while previous and current results are summed up to produce the output. The process for a 1D input and kernel, which corresponds to the regular convolution in Fig. 2, is illustrated in Fig. 3. In general, this scheme can be applied for online-processing of any  $ND$  input, where one dimension is a temporal continual stream. Continual Convolutions are causal [34] with no information leaking from future to past and can be efficiently implemented by zero-padding the input frame along the temporal dimension with  $p = \text{floor}(k/2)$ . Python-style pseudo-code of the implementation is shown in Listing 1.1.

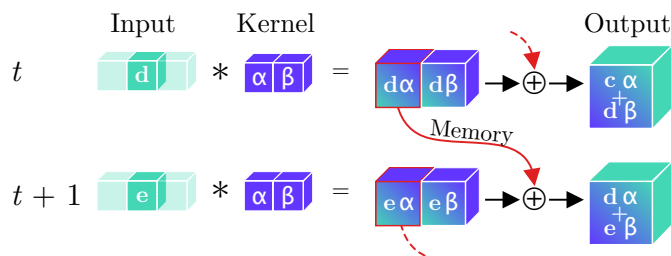


Fig. 3: **Continual Convolution**. An input (green  $\mathbf{d}$  or  $\mathbf{e}$ ) is convolved with a kernel (blue  $\alpha, \beta$ ). The intermediary feature-maps corresponding to all but the last temporal position are stored, while the last feature map and prior memory are summed to produce the resulting output. For a continual stream of inputs, Continual Convolutions produce identical outputs to regular convolutions.

---

```

def coconv3d(frame, prev_state = (mem, i)):
    frame = spatial_padding(frame)
    frame = temporal_padding(frame)
    feat = conv3d(frame, weights)
    output, rest_feat = feat[0], feat[1:]
    mem, i = prev_state or init_state(output)
    M = len(mem)
    for m in range(M):
        output += mem[(i + m) % M, M - m - 1]
    output += bias
    mem[i] = rest_feat
    i = (i + 1) % M
    return output, (mem, i)

```

---

Listing 1.1: **Pseudo-code** for Continual Convolution. Ready-to-use modules are available in the Continual Inference library [15].

In terms of computational cost, we can now perform frame-by-frame computations much more efficiently than a regular 3D convolution. The complexity of processing a frame becomes:

$$\Theta([k_H \cdot k_W \cdot k_T + b] \cdot c_I \cdot c_O \cdot n_H \cdot n_W). \quad (4)$$

This reduction in computational complexity comes at the cost of a memory-overhead in each layer due to the state that is kept between time-steps. The overhead of storing the partially computed feature-maps for a frame is:

$$\Theta(d_T \cdot [k_T - 1] \cdot c_O \cdot n_H \cdot n_W). \quad (5)$$

However, in the context of inference in a deep neural network, the transient memory usage within each time-step is reduced by a factor of  $n_T$  to

$$\Theta(c_O \cdot n_H \cdot n_W). \quad (6)$$

The benefits of Continual Convolutions include the independence of clip length on the computational complexity, state overhead, and transient memory consumption. The change from (non-causal) regular convolutions to (causal) Continual Convolutions has the side-effect of introducing a delay to the output. This is because some intermediary results of convolving a frame with the kernel are only added up at a later point in time (see Fig. 3). The delay amounts to

$$\Theta(d_T \cdot [k_T - p_T - 1]). \quad (7)$$

### 3.3 Continual Residuals

The delay from Continual Convolutions has an adverse side-effect on residual connections. Despite their simplicity in regular CNNs, we cannot simply add the input to a Continual Convolution with its output because the *CoConv* may delay the output. Residual connections to a *CoConv* must therefore be delayed by an equivalent amount (see Eq. (7)). This produces a memory overhead of

$$\Theta(d_T \cdot [k_T - 1] \cdot c_O \cdot m_H \cdot m_W). \quad (8)$$

### 3.4 Continual Pooling

The associative property of pooling operations allows for pooling to be decomposed across dimensions, i.e.  $\text{pool}_{T,H,W}(\mathbf{X}) = \text{pool}_T(\text{pool}_{H,W}(\mathbf{X}))$ . For continual spatio-temporal pooling, the pooling over spatial dimensions is equivalent to a regular pooling, while the intermediary pooling results must be stored for prior temporal frames. For a pooling operation with temporal kernel size  $k_T$  and spatial output size  $n_H \cdot n_W$ , the memory consumption and delays are

$$\Theta([k_T - 1] \cdot n_H \cdot n_W), \quad (9)$$

$$\Theta(k_T - p_T - 1). \quad (10)$$

Both memory consumption and delay scale linearly with the temporal kernel size. Fortunately, the memory consumed by temporal pooling layers is relatively modest for most CNN architectures (1.5% for *CoX3D-M*, see Appendix A). Hence, the delay rather than memory consumption may be of primary concern for real-life applications. For some network modules it may even make sense to skip the pooling in the conversion to a Continual CNN. One such example is the 3D Squeeze-and-Excitation (SE) block [18] in X3D, where global spatio-temporal average-pooling is used in the computation of channel-wise self-attention. Discarding the temporal pooling component (making it a 2D SE block) shifts the attention slightly (assuming the frame contents change slowly relative to the sampling rate) but avoids a considerable temporal delay.

### 3.5 The issue with temporal padding

Zero-padding of convolutional layers is a popular strategy for retaining the spatio-temporal dimension of feature-maps in consecutive CNN layers. For Continual CNNs, however, temporal zero-padding poses a problem, as illustrated in Fig. 4. Consider a 2-layer 1D CNN where each layer has a kernel size of 3 and zero padding of 1. For each new frame in a continual stream of inputs, the first layer  $l$  should produce two output feature-maps: One by the convolution of the two prior frames and the new frame, and another by convolving with one prior frame, the new frame, and a zero-pad. The next layer  $l + 1$  thus receives two inputs and produces three outputs which are dependent on the new input frame of the first layer (one for each input and another from zero-padding). In effect, each zero padding in a convolution forces the next layer to retrospectively update its output for a previous time-step in a non-causal manner. Thus, there is a considerable downside to the use of padding. This questions the necessity of zero padding along the temporal dimension. In regular CNNs, zero padding has two benefits: It helps to avoid spatio-temporal shrinkage of feature-maps when propagated through a deep CNN, and it prevents information at the borders from “washing away” [24]. The use of zero-padding, however, has the downside that it alters the input-distribution along the borders significantly [29], [33]. For input data which is a continual stream of frames, a shrinkage of the feature-size in the temporal dimension is not a concern, and an input frame (which may be

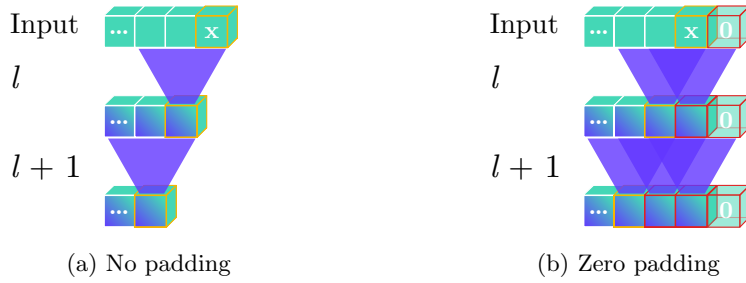


Fig. 4: **Issue with temporal padding:** The latest frame  $\mathbf{x}$  is propagated through a CNN with (purple) temporal kernels of size 3 (a) without or (b) with zero padding. Highlighted cubes can be produced only in the latest frame, with yellow boarder indicating independence of padded zero and red boarders dependence. In the zero-padded case (b), the number of frame features dependent on  $\mathbf{x}$  following a layer  $l$  increases with the number of padded zeros.

considered a border frame in a regular 3D CNN) has no risk of “washing away” because it is a middle frame in subsequent time steps. Temporal padding is thus omitted in Continual CNNs. As can be seen in the experimental evaluations presented in the following, this constitutes a “model shift” in the conversion from regular to Continual 3D CNN if the former was trained with temporal padding.

### 3.6 Initialisation

Before a Continual CNN reaches a steady state of operation, it must have processed  $r_T - p_T - 1$  frames where  $r_T$  and  $p_T$  are the aggregated temporal receptive field and padding of the network. For example, Continual X3D- $\{S, M, L\}$  models have receptive fields of size  $\{69, 72, 130\}$ , aggregated padding  $\{28, 28, 57\}$ , and hence need to process  $\{40, 43, 72\}$  frames prior to normal operation. The initial response depends on how internal state variables are initialised. In Sec. 4.2, we explore this further with two initialisation variants: 1) Initialisation with *zeros* and 2) by repeating a *replicate* of the features corresponding to the first input-frame. The latter corresponds to operating in a steady state for a “boring video” [3] which has one frame repeated in the entire clip.

### 3.7 Design considerations

Disregarding the storage requirement of model weights (which is identical between for regular and continual 3D CNNs), X3D-M has a worst-case total memory-consumption of 7,074,816 floats when prior frames and the transient feature-maps are taken into account. Its continual counterpart, *CoX3D-M*, has a worst case memory only 5,072,688 floats. How can this be? Since Continual 3D CNNs do not store prior input frames and have smaller transient feature maps, memory savings outweigh the cost of caching features in each continual layer. Had the



clip size been four instead of sixteen, X3D-M<sub>4</sub> would have had a worst-case memory consumption of 1,655,808 floats and *CoX3D-M*<sub>4</sub> of 5,067,504 floats. For clip size 64, X3D-M<sub>64</sub> consumes 28,449,792 floats and *CoX3D-M*<sub>64</sub> uses 5,093,424 floats. The memory load of regular 3D CNNs is thus highly dependent on clip size, while that of *Co3D* CNNs is not. Continual CNNs utilise longer receptive fields much more efficiently than regular CNNs in online processing scenarios. In networks intended for embedded systems or online processing, we may increase the clip size to achieve higher accuracy with minimal penalty in computational complexity and worst-case memory.

Another consideration, which influences memory consumption is the temporal kernel size and dilation of *CoConv* layers. Fortunately, the trend to employ small kernel sizes leaves the memory consumption reasonable for recent 3D CNNs [3], [43], [11], [10]. A larger temporal kernel size would not only affect the memory growth through the *CoConv* filter, but also for co-occurring residual connections. These consume a significant fraction of the total state-memory for real-life networks: in a Continual X3D-M model (*CoX3D-M*) the memory of residuals constitutes 20.5% of the total model state memory (see Appendix A).

### 3.8 Training

*Co3D* CNNs are trained with back-propagation like other neural networks. However, special care must be taken in the estimation of data statistics in normalisation layers: 1) Momentum should be adjusted to  $\text{mom}_{\text{step}} = 2/(1 + \text{timesteps} \cdot (2/\text{mom}_{\text{clip}} - 1))$  to match the exponential moving average dynamics of clip-based training, where T is the clip size; 2) statistics should not be tracked for the transient response. Alternatively, they can be trained offline in their “unrolled” regular 3D-CNN form with no temporal padding. This is similar to utilising pre-trained weights from a regular 3D CNN, as we do in our experiments.

## 4 Experiments

The experiments in this section aim to show the characteristics and advantages of Continual 3D CNNs as compared with regular 3D CNNs. One of the main benefits of *Co3D* CNNs is their ability to reuse the network weights of regular 3D CNNs. As such, all *Co3D* CNNs in these experiments use publicly available pre-trained network weights of regular 3D CNNs [11], [10], [8] without further fine-tuning. Data pre-processing follows the respective procedures associated with the originating weights unless stated otherwise. The section is laid out as follows: First, we showcase the network performance following weight transfer from regular to Continual 3D on multiple datasets for Human Activity Recognition. This is followed by a study on the transient response of *Co3D* CNNs at startup. Subsequently, we show how the computational advantages of *Co3D* CNNs can be exploited to improve accuracy by extending the temporal receptive field. Finally, we perform an extensive on-hardware benchmark of prior methods and Continual 3D CNNs, measuring the 1-clip/frame accuracy of publicly available models, as well as their inference throughput on various computational devices.

Model	Acc. (%)	Par. (M)	Mem. (MB)	FLOPs (G)	Throughput (preds/s)			
					CPU	TX2	Xavier	2080Ti
I3D-R50	63.98	28.04	191.59	28.61	0.93	2.54	9.20	77.15
R(2+1)D-18 <sub>8</sub>	53.52	31.51	168.87	20.35	1.75	3.19	6.82	130.88
R(2+1)D-18 <sub>16</sub>	59.29	31.51	215.44	40.71	0.83	1.82	3.77	75.81
Slow-8×8-R50	67.42	32.45	266.04	54.87	0.38	1.34	4.31	61.92
SlowFast-8×8-R50	68.45	66.25	344.84	66.25	0.34	0.87	2.72	30.72
SlowFast-4×16-R50	67.06	34.48	260.51	36.46	0.55	1.33	3.43	41.28
X3D-L	69.29	6.15	240.66	19.17	0.25	0.19	4.78	36.37
X3D-M	67.24	3.79	126.29	4.97	0.83	1.47	17.47	116.07
X3D-S	64.71	3.79	61.29	2.06	2.23	2.68	42.02	276.45
X3D-XS	59.37	3.79	28.79	0.64	8.26	8.20	135.39	819.87
RCU <sub>8</sub> [39] <sup>†</sup>	53.40	12.80	-	4.71	-	-	-	-
CoI3D <sub>8</sub>	59.58	28.04	235.87	5.68	3.00	2.41	14.88	125.59
CoI3D <sub>64</sub>	56.86	28.04	236.08	5.68	3.15	2.41	14.89	126.32
CoSlow <sub>8</sub>	65.90	32.45	175.98	6.90	2.80	1.60	6.18	113.77
<b>CoSlow<sub>64</sub></b>	<b>73.05</b>	<b>32.45</b>	<b>176.41</b>	<b>6.90</b>	<b>2.92</b>	<b>1.60</b>	<b>6.19</b>	<b>102.00</b>
CoX3D-L <sub>16</sub>	63.03	6.15	184.29	1.25	2.30	0.99	25.17	206.65
<b>CoX3D-L<sub>64</sub></b>	<b>71.61</b>	<b>6.15</b>	<b>184.37</b>	<b>1.25</b>	<b>2.30</b>	<b>0.99</b>	<b>27.56</b>	<b>217.53</b>
CoX3D-M <sub>16</sub>	62.80	3.79	68.88	0.33	7.57	7.26	88.79	844.73
<b>CoX3D-M<sub>64</sub></b>	<b>71.03</b>	<b>3.79</b>	<b>68.96</b>	<b>0.33</b>	<b>7.51</b>	<b>7.04</b>	<b>86.42</b>	<b>796.32</b>
CoX3D-S <sub>13</sub>	60.18	3.79	41.91	0.17	13.16	11.06	219.64	939.72
<b>CoX3D-S<sub>64</sub></b>	<b>67.33</b>	<b>3.79</b>	<b>41.99</b>	<b>0.17</b>	<b>13.19</b>	<b>11.13</b>	<b>213.65</b>	<b>942.97</b>

Table 1: **Kinetics-400 benchmark**. The noted accuracy is the single clip or frame top-1 score using RGB as the only input-modality. The performance was evaluated using publicly available pre-trained models without any further fine-tuning. For speed comparison, predictions per second denote frames per second for the *CoX3D* models and clips per second for the remaining models. Throughput results are the mean of 100 measurements. Pareto-optimal models are marked with bold. Mem. is the maximum allocated memory during inference noted in megabytes. <sup>†</sup>Approximate FLOPs derived from paper (see Appendix C).

#### 4.1 Transfer from regular to Continual CNNs

To gauge direct transferability of 3D CNN weights, we implement continual versions of various 3D CNNs and initialise them with their publicly available weights for Kinetics-400 [25] and Charades [38]. While it is common to use an ensemble prediction from multiple clips to boost video-level accuracy on these benchmarks, we abstain from this, as it doesn't apply to online-scenarios. Instead, we report the single-clip/frame model performance.

*Kinetics-400*. We evaluate the X3D network variants XS, S, M, and L on the test set using one temporally centred clip from each video. The XS network is omitted in the transfer to *CoX3D*, given that it is architecturally equivalent to S, but with fewer frames per clip. In evaluation on Kinetics-400, we faced the challenge that videos were limited to 10 seconds. Due to the longer transient response of Continual CNNs (see Sec. 4.2) and low frame-rate used for training X3D models

	Model	FLOPs (G) $\times$ views	mAP (%)
Clip	Slow-8 $\times$ 8 [11]	$54.9 \times 30$	39.0
	Slow-8 $\times$ 8 [11] <sup>†</sup>	$54.9 \times 1$	21.4
	Slow-8 $\times$ 8 (ours)	$54.9 \times 1$	24.1
Fr.	CoSlow <sub>8</sub>	$6.9 \times 1$	21.5
	CoSlow <sub>64</sub>	$6.9 \times 1$	25.2

Table 2: **Charades benchmark.** Noted are the FLOPs  $\times$  views and video-level mean average precision (mAP) on the validation set using pre-trained model weights. <sup>†</sup>Results achieved using the publicly available SlowFast code [11].

(5.0, 6.0, 6.0 FPS for S, M, and L), the video-length was insufficient to reach steady-state for some models. As a practical measure to evaluate near steady-state, we repeated the last video-frame for a padded video length of  $\approx 80\%$  of the network receptive field as a heuristic choice. The Continual CNNs were thus tested on the last frame of the padded video and initialised with the prior frames. The results of the X3D transfer are shown in Tab. 1 and Fig. 1.

For all networks, the transfer from regular to Continual 3D CNN results in significant computational savings. For the S, M, and L networks the reduction in FLOPs is  $12.1\times$ ,  $15.1\times$ , and  $15.3\times$  respectively. The savings do not quite reach the clip sizes since the final pooling and prediction layers are active for each frame. As a side-effect of the transfer from zero-padded regular CNN to Continual CNN without zero-padding, we see a notable reduction in accuracy. This is easily improved by using an extended pooling size for the network (discussed in Sec. 3.7 and in Sec. 4.2). Using a global average pooling with temporal kernel size 64, we improve the accuracy of X3D by 2.6%, 3.8%, and 2.3% in the Continual S, M, and L network variants. As noted, Kinetics dataset did not have sufficient frames to fill the temporal receptive field of all models in these tests. We explore this further in Sections 4.2 and 4.2.

*Charades.* To showcase the generality of the approach, we repeat the above described procedure with another 3D CNN, the CoSlow network [11]. We report the video-level mean average precision (mAP) of the validation split alongside the FLOPs per prediction in Tab. 2. Note the accuracy discrepancy between 30 view (10 temporal positions with 3 spatial positions each) and 1 view (spatially and temporally centred) evaluation. As observed on Kinetics, the CoSlow network reduces the FLOPs per prediction proportionally with the original clip size (8 frames), and can recover accuracy by extending the global average pool size.

## 4.2 Ablation Experiments

As described in Sec. 3.6, Continual CNNs exhibit a transient response during their up-start. In order to gauge this response, we perform ablations on the

Kinetics-400 validation set, this time sampled at 15 FPS to have a sufficient number of frames available. This corresponds to a data domain shift [45] relative to the pre-trained weights, where time advances slower.

**Transient response of Continual CNNs.** Our expected upper bound is given by the baseline X3D network 1-clip accuracy at 15 FPS. The transient response is measured by varying the number of prior frames used for initialisation before evaluating a frame using the *CoX3D* model. Note that temporal center-crops of size  $T_{\text{init}} + 1$ , where  $T_{\text{init}}$  is the number of initialisation frames, are used in each evaluation to ensure that the frames seen by the network come from the centre. This precaution counters a data-bias, we noticed in Kinetics-400, namely that the start and end of a video are less informative and contribute to worse predictions than the central part. We found results to vary up to 8% for a X3D-S network evaluated at different video positions. The experiment is repeated for two initialisation schemes, “zeros” (used in other experiments) and “replicate”, and two model sizes, S and M. The transient responses are shown in Fig. 5.

For all responses, the first  $\approx 25$  frames produce near-random predictions, before rapidly increasing at 25–30 frames until a steady-state is reached at 49.2% and 56.2% accuracy for S and M. Relative to the regular X3D, this constitutes a steady-state error of  $-1.7\%$  and  $-5.8\%$ . Comparing initialisation schemes, we see that the “replicate” scheme results in a slightly earlier rise. The rise sets in later for the “zeros” scheme, but exhibits a sharper slope, topping with peaks of 51.6% and 57.6% at 41 and 44 frames seen as discussed in Sec. 3.6. This

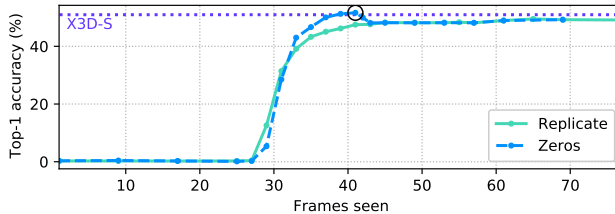
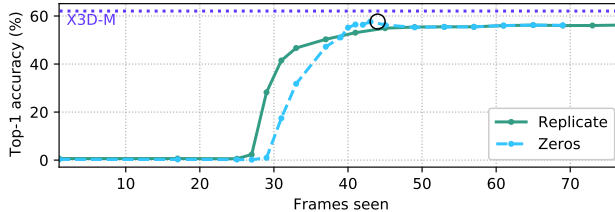
(a) *CoX3D-S*(b) *CoX3D-M*

Fig. 5: **Transient response** for Continual X3D- $\{S,M\}$  on the Kinetics-400 val at 15 FPS. Dotted horizontal lines denote X3D validation accuracy for 1-clip predictions. Black circles highlight the theoretically required initialisation frames.

makes sense considering that the original network weights were trained with this exact amount of zero-padding. Adding more frames effectively replaces the padded zeros and causes a slight drop of accuracy in the steady state, where the accuracy settles at the same values as for the “replication” scheme.

**Extended receptive field.** Continual CNNs experience a negligible increase in computational cost when larger temporal receptive field are used (see Sec. 3.7). For *CoX3D* networks, this extension can be trivially implemented by increasing the temporal kernel size of the last pooling layer. In this set of experiments, we extend *CoX3D*-{S,M,L} to have temporal pooling sizes 32, 64, and 96, and evaluate them on the Kinetics-400 validation set sampled at 15 FPS. The Continual CNNs are evaluated at frames corresponding to the steady state.

Tab. 3 shows the measured accuracy and floating point operations per frame (*CoX3D*) / clip (*X3D*) as well as the pool size for the penultimate network layer (global average pooling) and the total receptive field of the network in the temporal dimension. As found in Sec. 4.1, each transfer results in significant computational savings alongside a drop in accuracy. Extending the kernel size of the global average pooling layer increases the accuracy of the Continual CNNs by 11.0–13.3% for 96 frames relative the original 13–16 frames, surpassing that of the regular CNNs. Lying at 0.017–0.009%, the corresponding computational increases can be considered negligible.

Model	Size	Pool	Acc.	FLOPs (K)	Rec. Field
<i>X3D</i>	S	13	51.0	2,061,366	13
	M	16	62.1	4,970,008	16
	L	16	64.1	19,166,052	16
<i>CoX3D</i>	S	13	49.2	166,565	69
		16	50.1	166,567	72
		32	54.7	166,574	88
		64	59.8	166,587	120
		96	<i>61.8</i>	166,601	152
		16	56.3	325,456	72
	M	32	60.7	325,463	88
		64	64.9	325,477	120
		96	<i>67.3</i>	325,491	152
		16	53.0	1,245,549	130
	L	32	58.5	1,245,556	146
		64	<i>64.3</i>	1,245,570	178
96		<i>66.3</i>	1,245,584	210	

Table 3: **Effect of extending pool size.** Note that the model weights were trained at different sampling rates than evaluated at (15 FPS), resulting in a lower top-1 val. accuracy. *Italic numbers* denote measurement taken within the transient response due to a lack of frames in the video-clip.

### 4.3 Inference benchmarks

Despite their high status in activity recognition leader-boards [35], it is unclear how recent 3D CNNs methods perform in the online setting, where speed and accuracy constitute a necessary trade-off. To the best of our knowledge, there has not yet been a systematic evaluation of throughput for these video-recognition models on real-life hardware. In this set of experiments, we benchmark the FLOPs, parameter count, maximum allocated memory and 1-clip/frame accuracy of I3D [3], R(2+1)D [43], SlowFast[41], X3D [10], *CoI3D*, *CoSlow*, and *CoX3D*. To gauge achievable throughputs at different computational budgets, networks were tested on four hardware platforms as described in Appendix B.

As seen in the benchmark results found in Tab. 1, the limitation to one clip markedly lowers accuracy compared with the multi-clip evaluation published in the respective works [3], [43], [11], [10]. Nonetheless, the Continual models with extended receptive fields attain the best accuracy/speed trade-off by a large margin. For example, *CoX3D-L<sub>64</sub>* on the Nvidia Jetson Xavier achieves an accuracy of 71.3% at 27.6 predictions per second compared to 67.2% accuracy at 17.5 predictions per second for X3D-M while reducing maximum allocated memory by 48%! Confirming the observation in [30], we find that the relation between model FLOPs and throughput varies between models, with better ratios attained for simpler models (e.g., I3D) than for complicated ones (e.g., X3D). This relates to different memory access needs and their cost. Tailor-made hardware could plausibly reduce these differences. Supplementary visualisation of the results in Tab. 1 are found in Appendix C.

## 5 Conclusion

We have introduced Continual 3D Convolutional Neural Networks (*Co3D CNNs*), a new computational model for spatio-temporal 3D CNNs, which performs computations frame-wise rather than clip-wise while being weight-compatible with regular 3D CNNs. In doing so, we are able dispose of the computational redundancies faced by 3D CNNs in continual online processing, giving up to a  $15.1\times$  reduction of floating point operations, a  $9.2\times$  real-life inference speed-up on CPU, 48% peak memory reduction, and an accuracy improvement of 5.6% on Kinetics-400 through an extension in the global average pooling kernel size.

While this constitutes a substantial leap in the processing efficiency of energy-constrained and real-time video recognition systems, there are still unanswered questions pertaining to the dynamics of *Co3D CNNs*. Specifically, the impact of extended receptive fields on the networks ability to change predictions in response to changing contents in the input video is untested. We leave these as important directions for future work.

## Acknowledgement

This work has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 871449 (OpenDR).

## References

1. Arnab, A., Deghani, M., Heigold, G., Sun, C., Lučić, M., Schmid, C.: Vivit: A video vision transformer. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 6836–6846 (2021)
2. Cai, Z., He, X., Sun, J., Vasconcelos, N.: Deep learning with low precision by half-wave gaussian quantization. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5406–5414 (2017)
3. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4724–4733 (2017)
4. Carreira, J., Noland, E., Banki-Horvath, A., Hillier, C., Zisserman, A.: A short note about kinetics-600. preprint, arXiv:1808.01340 (2018)
5. Carreira, J., Pătrăucean, V., Mazare, L., Zisserman, A., Osindero, S.: Massively parallel video networks. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) Proceedings of the European Conference on Computer Vision (ECCV). pp. 680–697 (2018)
6. Chen, W., Wilson, J.T., Tyree, S., Weinberger, K.Q., Chen, Y.: Compressing neural networks with the hashing trick. In: International Conference on International Conference on Machine Learning (ICML). p. 2285–2294 (2015)
7. Donahue, J., Hendricks, L.A., Guadarrama, S., Rohrbach, M., Venugopalan, S., Darrell, T., Saenko, K.: Long-term recurrent convolutional networks for visual recognition and description. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2625–2634 (2015)
8. Fan, H., Murrell, T., Wang, H., Alwala, K.V., Li, Y., Li, Y., Xiong, B., Ravi, N., Li, M., Yang, H., Malik, J., Girshick, R., Feiszli, M., Adcock, A., Lo, W.Y., Feichtenhofer, C.: PyTorchVideo: A deep learning library for video understanding. In: ACM International Conference on Multimedia (2021)
9. Farnebäck, G.: Two-frame motion estimation based on polynomial expansion. In: Image Analysis. pp. 363–370. Springer Berlin Heidelberg (2003)
10. Feichtenhofer, C.: X3D: Expanding architectures for efficient video recognition. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
11. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: IEEE/CVF International Conference on Computer Vision (ICCV) (October 2019)
12. Floropoulos, N., Tefas, A.: Complete vector quantization of feedforward neural networks. *Neurocomputing* **367**, 55–63 (2019)
13. Han, S., Mao, H., Dally, W.J.: Deep compression: Compressing deep neural network with pruning, trained quantization and Huffman coding. In: International Conference on Learning Representations (ICLR) (2016)
14. He, Y., Zhang, X., Sun, J.: Channel pruning for accelerating very deep neural networks. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 1398–1406 (2017)
15. Hedegaard, L., Iosifidis, A.: Continual inference: A library for efficient online inference with deep neural networks in pytorch. preprint, arXiv:2204.03418 (2022)
16. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. In: NIPS Deep Learning and Representation Learning Workshop (2015)
17. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. preprint, arXiv:1704.04861 **abs/1704.04861** (2017)

18. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7132–7141 (2018)
19. Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., Bengio, Y.: Binarized neural networks. In: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 29. Curran Associates, Inc. (2016)
20. Ji, S., Xu, W., Yang, M., Yu, K.: 3d convolutional neural networks for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **35**(1), 221–231 (2013)
21. Joe Yue-Hei Ng, Hausknecht, M., Vijayanarasimhan, S., Vinyals, O., Monga, R., Toderici, G.: Beyond short snippets: Deep networks for video classification. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4694–4702 (2015)
22. Kalogeiton, V., Weinzaepfel, P., Ferrari, V., Schmid, C.: Action tubelet detector for spatio-temporal action localization. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 4415–4423 (2017)
23. Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L.: Large-scale video classification with convolutional neural networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1725–1732 (2014)
24. Karpathy, A.: CS231n convolutional neural networks for visual recognition, <https://cs231n.github.io/convolutional-networks/>. Last visited on 2021/01/26
25. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., Suleyman, M., Zisserman, A.: The kinetics human action video dataset. preprint, arXiv:1705.06950 (2017)
26. Köpüklü, O., Hörmann, S., Herzog, F., Cevikalp, H., Rigoll, G.: Dissected 3d cnns: Temporal skip connections for efficient online video processing. preprint, arXiv:2009.14639 (2020)
27. Köpüklü, O., Wei, X., Rigoll, G.: You only watch once: A unified cnn architecture for real-time spatiotemporal action localization. preprint, arXiv:1911.06644 (2019)
28. Köpüklü, O., Kose, N., Gunduz, A., Rigoll, G.: Resource efficient 3d convolutional neural networks. In: IEEE/CVF International Conference on Computer Vision Workshop (ICCVW). pp. 1910–1919 (2019)
29. Liu, G., Shih, K.J., Wang, T.C., Reda, F.A., Sapra, K., Yu, Z., Tao, A., Catanzaro, B.: Partial convolution based padding. preprint, arXiv:1811.11718 pp. 1–11 (2018)
30. Ma, N., Zhang, X., Zheng, H.T., Sun, J.: Shufflenet v2: Practical guidelines for efficient cnn architecture design. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (September 2018)
31. Molchanov, P., Yang, X., Gupta, S., Kim, K., Tyree, S., Kautz, J.: Online detection and classification of dynamic hand gestures with recurrent 3d convolutional neural networks. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4207–4215 (2016)
32. Neimark, D., Bar, O., Zohar, M., Asselmann, D.: Video transformer network. In: 2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW). pp. 3156–3165 (2021)
33. Nguyen, A., Choi, S., Kim, W., Ahn, S., Kim, J., Lee, S.: Distribution padding in convolutional neural networks. In: *International Conference on Image Processing (ICIP)*. pp. 4275–4279 (2019)
34. van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., Kavukcuoglu, K.: Wavenet: A generative model for raw audio. preprint, arXiv:1609.03499 (2016)



35. Papers with Code: Kinetics-400 leaderboard, <https://paperswithcode.com/sota/action-classification-on-kinetics-400>. Last visited on 2021/02/03.
36. Passalis, N., Tefas, A.: Learning deep representations with probabilistic knowledge transfer. In: Proceedings of the European Conference on Computer Vision (ECCV) (2018)
37. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (ICCV)* **115**(3), 211–252 (2015)
38. Sigurdsson, G.A., Varol, G., Wang, X., Farhadi, A., Laptev, I., Gupta, A.: Hollywood in homes: Crowdsourcing data collection for activity understanding. In: Proceedings of the European Conference on Computer Vision (ECCV) (2016)
39. Singh, G., Cuzzolin, F.: Recurrent convolutions for causal 3d cnns. In: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW). pp. 1456–1465 (2019)
40. Singh, G., Saha, S., Sapienza, M., Torr, P., Cuzzolin, F.: Online real-time multiple spatiotemporal action localisation and prediction. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 3657–3666 (2017)
41. Sovrasov, V.: Ptflops, <https://github.com/sovrasov/flops-counter.pytorch>. MIT License. Last visited on 2021/03/02
42. Tan, M., Le, Q.: EfficientNet: Rethinking model scaling for convolutional neural networks. In: Proceedings of Machine Learning Research. vol. 97, pp. 6105–6114 (2019)
43. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A closer look at spatiotemporal convolutions for action recognition. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6450–6459 (2018)
44. Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3d convolutional networks. In: IEEE International Conference on Computer Vision (ICCV). pp. 4489–4497 (2015)
45. Wang, M., Deng, W.: Deep visual domain adaptation: A survey. *Neurocomputing* **312**, 135–153 (2018)
46. Xu, M., Zhu, M., Liu, Y., Lin, F., Liu, X.: Deepcache: Principled cache for mobile deep vision. *International Conference on Mobile Computing and Networking* (2018)
47. Yim, J., Joo, D., Bae, J., Kim, J.: A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7130–7138 (2017)
48. Zhang, X., Zhou, X., Lin, M., Sun, J.: Shufflenet: An extremely efficient convolutional neural network for mobile devices. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6848–6856 (2018)
49. Zhu, L., Sevilla-Lara, L., Yang, Y., Feiszli, M., Wang, H.: Faster recurrent networks for efficient video classification. *Proceedings of the AAAI Conference on Artificial Intelligence* **34**, 13098–13105 (2020)