

Self-Supervised Video Representation Learning with Space-Time Cubic Puzzles

Dahun Kim, Donghyeon Cho, In So Kweon

Dept. of Electrical Engineering, KAIST, Daejeon, Korea
mcahny@kaist.ac.kr, cdh12242@gmail.com, iskweon77@kaist.ac.kr

Abstract

Self-supervised tasks such as colorization, inpainting and zig-saw puzzle have been utilized for visual representation learning for still images, when the number of labeled images is limited or absent at all. Recently, this worthwhile stream of study extends to video domain where the cost of human labeling is even more expensive. However, the most of existing methods are still based on 2D CNN architectures that can not directly capture spatio-temporal information for video applications. In this paper, we introduce a new self-supervised task called as *Space-Time Cubic Puzzles* to train 3D CNNs using large scale video dataset. This task requires a network to arrange permuted 3D spatio-temporal crops. By completing *Space-Time Cubic Puzzles*, the network learns both spatial appearance and temporal relation of video frames, which is our final goal. In experiments, we demonstrate that our learned 3D representation is well transferred to action recognition tasks, and outperforms state-of-the-art 2D CNN-based competitors on UCF101 and HMDB51 datasets.

Introduction

Recent progress in computer vision stems from a huge amount of labeled images as well as deep convolutional neural networks. Typically, a network trained with ImageNet (Russakovsky et al. 2015) consisting of one million images and label pairs learns the general features of the image and has been used to initialize the network for various kinds of downstream tasks. In fact, there are much more than one million images in the web, however, building large-scale annotated datasets is extremely expensive and impractical. Therefore, many researches have been attempted to minimize human supervision in computer vision. For example, (Oquab et al. 2015) and (Kim et al. 2017) proposed to use *weak* image tag information for object localization without using bounding boxes or pixel-level masks. In the same vein, unsupervised representation learning, which learns general-purpose semantic features without human annotation, has been regarded as a fundamental problem for years (Bengio, Courville, and Vincent 2013). Among them, a prominent paradigm is the so-called *self-supervised representation learning* that defines an annotation-free pretext task from the raw data in order to provide a free supervision signal

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

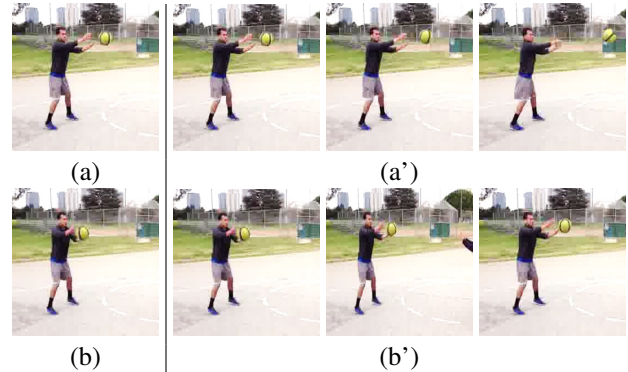


Figure 1: **Ambiguity in time direction.** Given a pair of images on the left, it is ambiguous to determine the time direction in between. One can finally identify the direction ($b' \rightarrow a'$) and action (*throwing*) when given video sequences.

for feature learning. For instance, a deep CNN is taught to complete zigsaw puzzles (Noroozi and Favaro 2016) and fill in missing pixels (Pathak et al. 2016). The rationale behind such self-supervised tasks is that solving them will force the CNN to learn semantic image features that can be useful for other vision tasks. In image domain, self-supervised learning is performed using only images from ImageNet (Russakovsky et al. 2015) without labels and is transferred to downstream tasks such as Pascal (Everingham et al. 2007; 2012). Recent methods have shown promising results, and significantly narrowed the gap with the fully supervised learning using ImageNet labels.

More recently, this worthwhile stream of research has extended to video domain, where the burden of human annotation is even more severe. Compared to images, videos provide additional temporal information. To illustrate, we introduce a problem of guessing the direction of time in Fig. 1. Given a pair of image on the left, one runs into a problem of determining whether the action is *catching* ($a \rightarrow b$) or *throwing* ($b \rightarrow a$). One can finally clarify when the neighboring frames are given together that it is *throwing* ($b' \rightarrow a$). In the past few years, various self-supervision signals using video frames have shown promising results that are better than random initialization in action recognition tasks. How-

ever, the scope of these approaches is still limited to using 2D CNN architectures which are appearance-based, leaving the ambiguity in the temporal dimension unsolved.

In this paper, we focus on 3D CNNs which can directly extract spatio-temporal features from raw videos. With the advent of large scale video datasets such as Kinetics (Kay et al. 2017), these 3D CNNs have recently begun to outperform 2D CNNs in action recognition as (Hara, Kataoka, and Satoh 2018). In the context of self-supervised learning, we propose a pretext task for 3D CNNs to close the gap with fully supervised Kinetics-pretraining in video domain. Given a randomly permuted 3D spatio-temporal crops extracted from each video clips, we train a network to predict their original spatio-temporal arrangement. We call this task as *Space-Time Cubic Puzzles*, and examples are shown in Fig. 4. By solving *Space-Time Cubic Puzzles*, the 3D CNN is forced to have an understanding of both spatial appearance and temporal relation in the video, which is our final goal.

We conduct extensive experimental validation to demonstrate the effectiveness of our self-supervised video feature learning. First, we compare the proposed method with baseline methods including the random initialization and fully supervised pretraining as well as alternative pretraining strategies. Also, we perform various ablation studies to provide deeper analysis on 3D spatio-temporal representation. Finally, we demonstrate that our learned 3D representation with comparable or fewer number of parameters outperforms state-of-the-art 2D CNN competitors on action recognition tasks of using UCF101 (Soomro, Zamir, and Shah 2012) and HMDB51 (Kuehne et al. 2011) benchmark datasets.

Our contributions can be summarized as follows:

- We propose a novel pretext task of solving 3D video cubic puzzles for self-supervised video representation learning from unlabeled videos. To our best knowledge, this is the first work to focus on the spatio-temporal 3D CNNs in self-supervised representation learning in videos.
- We provide various ablation studies and analysis for deeper understanding of 3D spatio-temporal representation.
- Our learned 3D CNN representation outperforms other self-supervised approaches on two publicly available action recognition datasets (UCF101, HMDB51), while having fewer or comparable number of parameters.
- We significantly close the gap between unsupervised representation learning and Kinetics-pretraining for 3D CNNs. When transferred onto UCF101, our self-supervised learning improves +23.4% over training from scratch, and shows comparable performances to the strong supervision that uses *one eighth* of the Kinetics labels.

Related Works

In this section, we review two categories of prior works: video recognition and self-supervised representation learning, which are the most relevant to our work.

Video Recognition and Kinetics Dataset

Recent progress in video recognition is rooted in the use of large-scale datasets that enable the *pretraining* of CNNs for a wide variety of downstream tasks. To date, ImageNet (Russakovsky et al. 2015) has contributed substantially to the pretraining of a generic feature representation in many video recognition algorithms. First of all, (Karpathy et al. 2014) introduced multiresolution CNN architecture for large-scale video classification. They also provided several schemes for time information fusion. (Simonyan and Zisserman 2014) proposed a two-stream architecture to capture spatial and motion information with a RGB stream and an optical flow stream respectively. (Wang et al. 2016) further improved the results by using temporal segments. These approaches are based on 2D CNNs that are pretrained on ImageNet.

Recently, CNNs with spatio-temporal 3D convolutional kernels (3D CNNs) have been actively touched for video applications. The first 3D CNN was proposed several years ago by (Ji et al. 2013). However, even the usage of well-organized models such as (Tran et al. 2015) has failed to outperform the advantages of 2D CNNs that combined both RGB and stacked flow (Simonyan and Zisserman 2014). The primary reason for this failure has been the relatively small data-scale of video datasets for optimizing the large number of parameters in 3D CNNs, which can only be trained on video datasets. More recently, however, (Carreira and Zisserman 2017) achieved a significant breakthrough using the Kinetics dataset (Kay et al. 2017), which includes more than 300K annotated videos. It was created with the aim of being positioned as standard video dataset roughly equivalent to the position held by ImageNet in image domain. Thus, we now have the benefit of a 3D convolution that can directly extract spatio-temporal features, by virtue of the Kinetics dataset.

However, most of the previous studies have trained 3D CNNs using all labels in the Kinetics dataset. Therefore, we argue that developing a self-supervised method to train 3D CNNs is a worthwhile pursuit.

Self-Supervised Representation Learning

To overcome the inherent thirst for data in fully supervised training, a large body of literature have studied unsupervised feature learning. A recently emerging line of research is self-supervised feature learning where the supervision signal is obtained automatically from unlabeled images or videos.

Over the last few years, several self-supervised tasks have been introduced. For instance, methods that use context arrangement of image patches (Doersch, Gupta, and Efros 2015; Noroozi and Favaro 2016), image completion (Pathak et al. 2016), motion frame ordering (Misra, Zitnick, and Hebert 2016; Lee et al. 2017) and multi-task of many models (Kim et al. 2018) have been proposed. Our work is closely related to the context-based methods (Doersch, Gupta, and Efros 2015; Noroozi and Favaro 2016; Misra, Zitnick, and Hebert 2016; Lee et al. 2017; Kim et al. 2018). These are a popular approach, and work by creating an arrangement of image patches in either space or time. Each distinct arrangement is assigned a class label,

and the network then predicts the correct arrangement of these patches by solving a supervised classification problem. Context-based methods have the advantage of being easy to understand, network architecture agnostic, and frequently straightforward to implement. They also tend to perform well on standard measures of transfer learning. For instance, (Noroozi and Favaro 2016) and (Lee et al. 2017) are the top performers on PASCAL VOC 2007 detection and UCF101 action classification respectively, even among a large number of new arrivals.

However, these context-based methods only leverage either of spatial *or* temporal dimension. Furthermore, they are based on 2D CNNs which can only extract frame-level features that cannot detect scene dynamics by nature. Since the spatial appearances and temporal relations are both very important cues for video understanding, our work investigates the use of both spatial and temporal dimensions using 3D CNNs in videos. To our best knowledge, only few works (Zhao et al. 2017; Vondrick, Pirsivash, and Torralba 2016) exploit the 3D architectures in self-supervised feature learning. They use reconstruction/generation-based pretext tasks, and aim for a specific target task: anomaly detection and video generation respectively. In contrast, we mainly investigate the fine-tuning of the learned feature representations for the video action recognition tasks. Arguably, the action recognition is a hallmark problem in video understanding, so it can serve as a general task, similarly to object recognition in image understanding. In experiments, quantitative comparisons demonstrate our supervisory signals are able to generate much richer 3D feature representations than previous 3D CNN-based methods, as well as the 2D CNN-based competitors, even with fewer number of parameters.

Proposed Approach

Pretext Task: Space-Time Cubic Puzzles

Our goal is to learn spatio-temporal representations with 3D CNNs using *unlabeled* videos. We propose a 3D cubic puzzle problem called as *Space-Time Cubic Puzzles*; Given a randomly permuted sequence of 3D spatio-temporal pieces cropped from a video clip, we train a network to predict their original arrangement. Although this is a difficult task even for a human, it becomes easy once we identify the objects and their actions in the video crops. We hypothesize that the successful network in this task captures representative and discriminative features for each 3D crop by determining their spatio-temporal arrangement. Thus, our learned clip-level 3D representations are transferable to downstream tasks in videos as well.

To generate the puzzle pieces, we consider a spatio-temporal cuboid consisting of $2 \times 2 \times 4$ grid cells for each video, as shown in Fig. 2 and Fig. 4. Given 16 crops, there are 16! possible permutations. However, these include very similar permutations which make the puzzle task very ambiguous. For example, if the difference between two permutations lies only in two crops that are similar-looking, it will be impossible for the network to predict the right solution (Noroozi and Favaro 2016). To avoid such ambiguity, we sample 4 crops instead of 16, in either spatial or tempo-

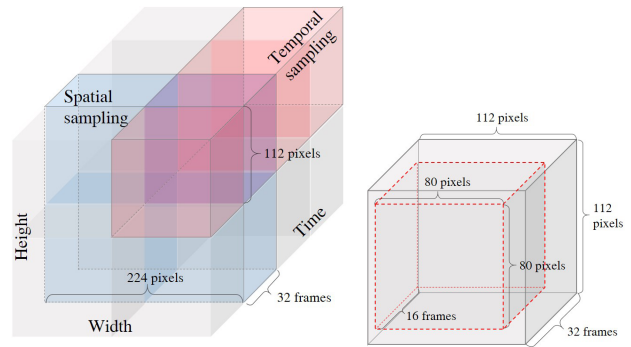


Figure 2: space-time cuboid (left) and spatio-temporal jittering (right).

ral dimension. More specifically, the 3D crops are extracted from a 4-cell grid of shape $2 \times 2 \times 1$ (colored in blue in Fig. 2-left) or $1 \times 1 \times 4$ (colored in red in Fig. 2-left) along the spatial or temporal dimension respectively. Finally, we randomly permute them to make our input. The network must feed the 4 input crops through several convolutional layers, and produce an output probability to each of the possible permutations that might have been sampled. Note, however, that we ultimately wish to learn spatio-temporal features for the *individual* 3D crop.

Network Architecture

To achieve this, we use a late-fusion architecture shown in Fig. 3-(a). It is a 4-tower siamese network, where the towers share the same parameters, and follow the 3D ResNet (Hara, Kataoka, and Satoh 2018) architecture to provide a comparison with the Kinetics-pretraining. Each 3D crops are processed separately until the fully-connected layer, so that the network cannot "cheat" by viewing low-level statistics such as edge boundaries without having to understand the global scene dynamics. Since only two last fully-connected layers receive input from all 4 crops, we expect the network to perform the most semantic reasoning for each crops separately. Furthermore, each towers are agnostic of whether it was spatial or temporal dimension the input crops had been sampled from. That is, each tower must encode the spatial and temporal structures in a given video crop simultaneously, because it does not know if the problem to solve in the last layers is a spatial puzzle or a temporal puzzle. Similar to the jigsaw puzzle problem (Noroozi and Favaro 2016), we formulate the rearrangement problem as a multi-class classification task. In practice, for each tuple of four crops, we flip all the frames upside-down with 50% probability, doubling the number of classes to 48 (that is, $2 \times 4!$) to further boost our performance, as suggested in (Mundhenk, Ho, and Chen 2018).

Avoiding Trivial Learning

When designing a pretext task, it is crucial to ensure that the task forces the network to learn the desired semantic structure, without bypassing the understanding by finding

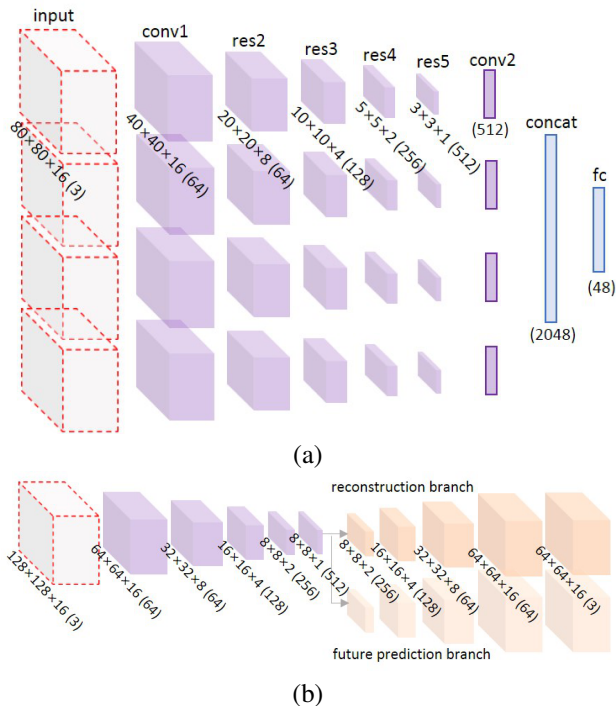


Figure 3: The overall architecture.

low-level clues that reveal the location of a video crop. As pointed out by (Doersch, Gupta, and Efros 2015), an example of this is *chromatic aberration*, which occurs naturally as a result of camera lensing. A common remedy for this is to partially drop color channels (Doersch, Gupta, and Efros 2015), replicate one channel (Lee et al. 2017), use grayscale inputs (Noroozi and Favaro 2016). We choose channel replication as our data preprocessing.

Another often-cited worry in all context-based works relates to trivial low-level boundary pattern completion (Doersch, Gupta, and Efros 2015; Noroozi and Favaro 2016; Lee et al. 2017). The network may learn the alignment between video crops not based on the semantics, but instead by matching the volume boundaries. Thus, we apply spatio-temporal jittering when extracting each video crops from the grid cells to avoid the trivial cases, as shown in the right side of Fig. 2.

Implementation Details

Network and dataset. We implement our method and conduct all experiments mostly using the 3D ResNet (Hara, Kataoka, and Satoh 2018) as a backbone architecture, since its performances on the random initialization and Kinetics-pretraining are well studied in their work. We can immediately compare the performance of our pretraining method to those scores. The training uses Kinetics datasets, which includes 400 human action classes, and consists of more than 400 videos for each class. The videos were temporally trimmed and last around 10 seconds. During the pretraining, we use the training split which has total 240K videos.

Pretraining. We use video clips with 224×224 pixel frames and convert every video file into PNG images in our experiments. We sample 128 consecutive frames from each clip, and split them into $2 \times 2 \times 4$ -cell grid; That is, one grid cell consists of $112 \times 112 \times 32$ pixels, and for each cell, we sample $80 \times 80 \times 16$ pixels with random jittering to generate a 3D video crop. We set the mini-batch size as 128 and the initial learning rate as 0.01. We use stochastic gradient descent with a momentum of 0.9 on two GTX-1080Ti GPUs. All the pre-trained models and the source codes will be available soon.

Experimental Results

In this section, we evaluate the effectiveness of our space-time cubic puzzle as a pretext task for self-supervised pre-training of 3D CNNs. As in prior works on self-supervised learning, we use the learned 3D CNNs features as the initialization for a fine-tuning stage for video recognition tasks. Better results indicate better qualities and generalization abilities of the learned video representations. We organize our experimental results as follows: 1) comparison with the random initialization and Kinetics-pretraining (supervised), 2) comparison with our alternative strategies, 3) ablation analysis, 4) comparison with the state-of-the-art methods, and 5) Visualization of the low-level filters and high-level activations. The followings are the datasets and fine-tuning details in all our experiments.

Datasets. We conduct video recognition experiments on two benchmark action recognition datasets, namely UCF101 (Soomro, Zamir, and Shah 2012) and HMDB51 (Kuehne et al. 2011). UCF101 contains 101 actions classes, 13K videos, and 27 hours of video data in total. The HMDB51 dataset consists of realistic videos captured from movies and Web videos, and contains 6,766 videos from 51 action classes. To be noted, all the experiments follow the training/test splits of UCF101 and HMDB51, and we mostly report the average classification accuracy over the three splits for UCF101, as done in (Hara, Kataoka, and Satoh 2018).

Fine-tuning for action recognition. Once we finish the pretraining stage, we use our learned parameters to initialize the 3D CNNs for action recognition, while the last fully-connected layer is initialized randomly. During the fine-tuning and testing, we follow the same protocol in (Hara, Kataoka, and Satoh 2018) to provide a fair comparison. Specifically, for each clip, we randomly sample 16 consecutive frames, and spatially resize the frames at 112×112 pixels. During the fine-tuning, we apply random spatial cropping, scaling and horizontal flipping to perform data augmentation. We start from a learning rate of 0.05, and assign a weight decay of $5e-4$. In testing, we adopt the sliding window manner to generate input clips, so that each video is split into non-overlapped 16-frame clips. The clip class scores are averaged over all the clips of the video.

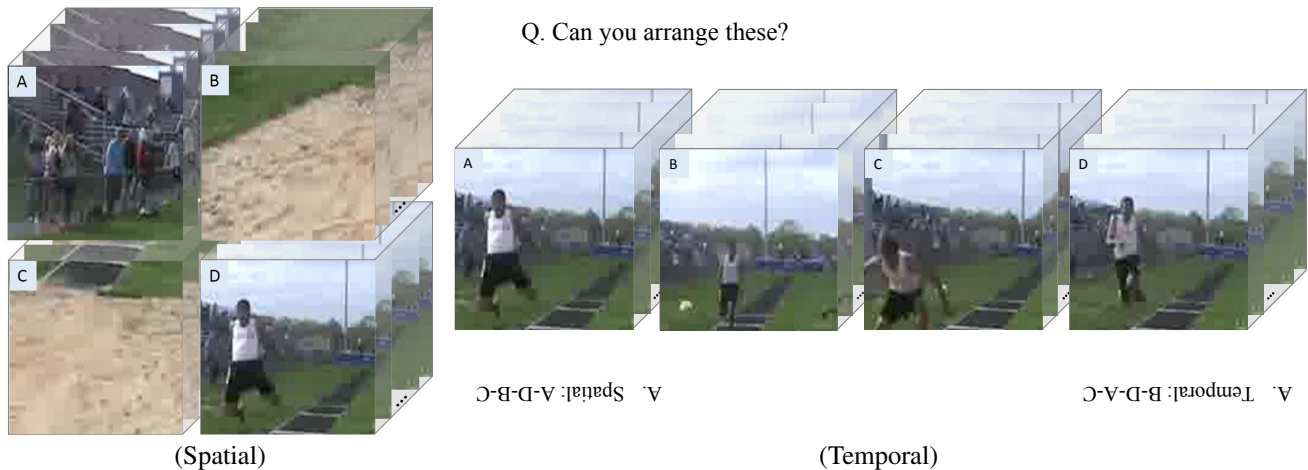


Figure 4: Example spatial and temporal tuples.

Initialization	UCF101(%)	HMDB51(%)
Random init.	42.4	17.1
3D ST-puzzle (ours)	65.8	33.7
Kinetics 1/8	64.2	33.2
Kinetics 1/4	71.1	41.1
Kinetics 1/2	78.0	48.6
Kinetics full	84.4	56.4
ImageNet-inflated	60.3	30.7

Table 1: **Comparison with random initialization / fully-supervised pretraining.** Top-1 accuracies on UCF101 and HMDB51. All methods use 3D ResNet-18, and the accuracies are averaged over three splits.

Comparison with Random Initialization and Fully-Supervised Pretraining

In these experiments we study the advantage of our self-supervised pretraining for action recognition in comparison to training from the scratch and several fully-supervised pretraining methods. We report the performances in Table. 1. Our self-supervised pretraining shows a dramatic **improvement of +23.4%** over training from scratch in UCF101 and a significant gain of +16.6% in HMDB51. This impressive gain demonstrates the effectiveness of our self-supervised cubic puzzle task.

Also, to quantitatively assess the effectiveness of our method in comparison to fully supervised methods, we gradually reduce the number of class labels in Kinetics dataset (full, 1/2, 1/4, and 1/8), and evaluate the pretraining results. Still having gap with the full Kinetics-pretraining, our method performs slightly better than the pretraining with *one eighth* of the Kinetics labels (that is, 50 out of 400 classes). In addition, to provide a comparison with ImageNet-pretraining, we import the existing ImageNet supervised 2D filters and inflate them into 3D, as suggested in (Carreira and Zisserman 2017). Our self-supervised pretraining results through *Space-Time Cubic Puzzles* outperform ImageNet-pretraining by +5.5% and +3.0% in each

Method	UCF101(%)
3D AE	48.7
3D AE + future	50.1
3D inpainting	50.9
3D S-puzzle	58.5
3D T-puzzle	59.3
3D ST score ensemble	61.3
3D ST-puzzle (full)	65.8

Table 2: **comparison with alternative methods.** Top-1 accuracies on UCF10. All methods use 3D ResNet-18, and the accuracies are averaged over three splits.

benchmark datasets. This implies that our video representations learned from the spatio-temporal context reasoning can be more powerful than the massively supervised 2D image-based representations in video recognition.

Alternative Pretraining Strategies

Since there are few prior works on self-supervised representation learning using 3D CNNs, we enumerate several alternative self-supervision tasks to provide our own reference levels and validate the effectiveness of our method. While we mainly focus on the context-based approaches, we also explore the reconstruction-based methods: spatio-temporal autoencoders (Zhao et al. 2017) and 3D inpainting (Pathak et al. 2016) as well. All the methods and experiments use the same 3D ResNet-18 as a backbone architecture, and use Kinetics dataset (without labels). To itemize, they are:

Context-based methods. Refer to Fig. 3-(a) for network architecture. We use cross entropy loss to train the networks.

- **3D ST-puzzle (spatio-temporal, Our full method):** The *Space-Time Cubic Puzzles*, where the tuple of 4 video crops is sampled in the spatial dimension with 50% probability, and in the temporal dimension otherwise. Due to

Method	Backbone	UCF101(%)	HMDB51(%)
Random initialization	3D ResNet-18	42.4	17.1
Random initialization	AlexNet	38.4	13.4
Temporal Coherency (Mobahi, Collobert, and Weston 2009)	AlexNet	45.4	15.9
Object Patch (Wang and Gupta 2015)	AlexNet	42.7	15.6
Sequence Verification (Misra, Zitnick, and Hebert 2016)	AlexNet	50.9	19.8
OPN (Lee et al. 2017)	AlexNet	<u>56.3</u>	22.1
Geometry (Gan et al. 2018)	AlexNet	54.1	<u>22.6</u>
Time Arrow (Wei et al. 2018)	AlexNet	55.3	-
Video Generation (Vondrick, Pirsivash, and Torralba 2016)	C3D	52.1	-
3D ST-puzzle (ours)	C3D	60.6	28.3
	3D ResNet-10	63.4	30.8
	3D ResNet-18	65.8	33.7

Table 3: **Comparison with the state-of-the-art methods.** Top-1 accuracies on UCF101 and HMDB51. All methods use 3D ResNet-18, and the accuracies are averaged over three splits.

this randomness, the network is forced to learn both spatial and temporal structures simultaneously.

- **3D S-puzzle (spatial only):** 3D extension of (Doersch, Gupta, and Efros 2015; Noroozi and Favaro 2016). Same as above, with the input tuple always generated from the spatial dimension.
- **3D T-puzzle (temporal only):** 3D extension of (Misra, Zitnick, and Hebert 2016). Same as above, with the input tuple always generated from the temporal dimension.
- **3D ST score ensemble:** Score ensemble of the classification scores of the S-puzzle and T-puzzle tasks. We average the softmax probabilities from both puzzle tasks.

Reconstruction-based methods. Refer to Fig. 3-(b) for network architecture. We use MSE loss to train the networks

- **3D AE (reconstruction):** The network is trained to reconstruct input stack of 16 frames. We use four 3D deconvolution layers with stride $2 \times 2 \times 2$ in the decoder. We followed the same decoder structure and the training protocol in (Zhao et al. 2017).
- **3D AE + future (recon. + future prediction):** Same as above, with one more decoder branch for joint future prediction of additional 16 frames, as in (Zhao et al. 2017).
- **3D inpainting:** 3D extension of (Pathak et al. 2016). The network is trained to recover the missing center region ($64 \times 64 \times 16$) in the input 16-frame stack.

We compare these methods in Table. 2. The context-based methods consistently outperform the reconstruction-based baselines. Also, we can see that the score ensemble gives better scores than the single-dimension baselines, implying that the knowledge from the spatial appearance are indeed complementary with those from the temporal relations. Our full method brings additional 3% performance gain on top of the score ensemble. This implies that our proposed method effectively aggregates spatio-temporal video features, and these features are much more discriminative and representative than those from the single-dimension baselines or their late fusion ensemble.

Method	UCF101(%)
with no regularizations	58.7
+ channel replication	61.5
+ random jittering	63.9
+ rotation with classification	65.8

Table 4: **Ablation studies.** Top-1 accuracies on UCF101. Each methods are accumulated down from the top and use 3D ResNet-18. The accuracies are averaged over three splits.

Ablation Studies

In order to validate various regularization techniques in our pretraining method, we evaluate the effect of each design choices on the UCF-101 dataset.

Channel replication. As mentioned earlier, chromatic aberration is one of the often-cited worries in context-based self-supervised learning because it leads to learning trivial color features. To prevent such issue, we first use grayscale images. We further experimented with channel replication where we randomly choose one representative channel and replicate its values to the other two channels. Table. 4 shows that channel replication improves the performance.

Spatio-temporal jittering. Analogous to the random gap used in the puzzle-solving task (Noroozi and Favaro 2016), we apply spatio-temporal jittering to each video crops to prevent the network from learning low-level statistics. In practice, we crop $80 \times 80 \times 16$ pixels from a $112 \times 112 \times 32$ -pixel cuboid with random shifts in all horizontal, vertical and temporal directions. Table. 4 shows that applying random jittering does help the network to learn better video features.

Rotation with classification. Recently, Mundhenk, Ho, and Chen (2018) developed a set of methods to improve on the results of self-supervised learning using context. To see if our model can benefit from these technologies, we apply one of their methods: *rotation with classification* (RWC) which encourages the network to identify if the inputs are right-side-up or upside-down. We do this by flipping all




Figure 5: **Learned filters with self-supervision vs. fully supervised-pretraining.** Visualization of the learned 64 filters in conv1 layer: (a) the resulting 3D $7 \times 7 \times 7$ filters of our self-supervised learning, (b) 3D $7 \times 7 \times 7$ filters from the Kinetics-pretrained network, and (c) 2D 7×7 filters from ImageNet-pretrained network. Note that our representation incorporate temporal dynamics and have rich temporal structure, without requiring massive human labels.

video crops in a tuple upside-down and doubling the number of classes ($24 \times 2 = 48$ in our work). Table. 4 shows that RWC does prevent *learning to bypass* and improves over the baseline. This implies that other off-the-shelf techniques for context-based self-supervised learning would further boost the performance of our pretraining method.

Comparison with the State-of-the-art Methods

We show a comparison of our results and the state-of-the-art self-supervised methods in Table. 3. In particular, we compare with (Mobahi, Collobert, and Weston 2009; Wang and Gupta 2015; Misra, Zitnick, and Hebert 2016; Lee et al. 2017; Gan et al. 2018; Wei et al. 2018) using the RGB video data. We quote the numbers directly from the published papers. It should be noted that the direct comparison with these 2D CNN-based methods is difficult due to the fundamental architectural difference. To complement this discrepancy, we conduct experiments with 3D architectures with different number of parameters and layers: C3D, 3D ResNet-10 and 3D ResNet-18. These networks have fewer parameters (11M, 14M and 33M respectively) compared with the AlexNet (58M) which is the backbone architecture in the 2D CNN-based methods. However, our approach outperforms other recent self-supervised methods. (Fernando et al. 2017) utilizes temporal order verification as a supervisory signal and can be used as a baseline as well. The minor difference is that this baseline uses stacks of frame differences (15 channels) as inputs. To use a similar setting, we use frame difference as inputs during fine-tuning and testing. With 3D ResNet-18, we achieved 75.3% on UCF101, that is outperforming Odd-One-Out method by a margin of +15.0%. (Vondrick, Pirsivash, and Torralba 2016) used C3D architecture for video generation and tested their learned representations on action recognition. Our results with the same C3D backbone network brings +8.5% performance gain over this, showing the informativeness of our self-supervised task.

Visualization of Learned Filters

All the learned conv1 filters from our self-supervised learning, Kinetics-pretraining, and ImageNet-pretraining are visualized in Fig. 5. We observe that: 1) All our filters change in the time dimension, meaning each encodes temporal information; 2) For most of the ImageNet-pretrained 2D filters, we can find a 3D filter with a similar appearance pattern mostly at the center slice, 4th out of 7, both in our filters and Kinetics-pretrained ones. These observations may imply that our learned 3D representations are able to not only cover the appearance information in 2D filters, but can also capture useful temporal motion simultaneously, like the Kinetics-pretrained representations do.

Conclusion

In this study, we examined the self-supervised feature learning for spatio-temporal 3D CNNs. We propose *Space-Time Cubic Puzzles* as our pretext task, and train with unlabeled Kinetics dataset. Our method enables learning both spatial appearances and temporal relations in video, which has been hardly achieved by previous 2D CNN-based self-supervisions. Our self-supervised pretraining performs slightly better than supervised pretraining on *one eighth* of the Kinetics labels on UCF101 and HMDB51 datasets. The visualization shows that our learned 3D representations indeed encode spatial and temporal information jointly.

We believe that the results of this study will facilitate further advances in self-supervised representation learning for spatio-temporal 3D CNNs. In recent years, significant progress has been made in self-supervised learning has narrowed the gap with ImageNet-pretraining in image domain. Similar to these, our self-supervised learning with 3D CNNs also shows promising results towards our ultimate goal of reducing human supervision in video domain. In our future work, we will investigate transfer learning not only for action recognition but also for other such tasks.

References

- Bengio, Y.; Courville, A.; and Vincent, P. 2013. Representation learning: A review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)* 35(8):1798–1828.
- Carreira, J., and Zisserman, A. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*.
- Doersch, C.; Gupta, A.; and Efros, A. A. 2015. Unsupervised visual representation learning by context prediction. In *Proc. of Int'l Conf. on Computer Vision (ICCV)*.
- Everingham, M.; Van Gool, L.; Williams, C. K. I.; Winn, J.; and Zisserman, A. 2007. The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results. <http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html>.
- Everingham, M.; Van Gool, L.; Williams, C. K. I.; Winn, J.; and Zisserman, A. 2012. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>.
- Fernando, B.; Bilen, H.; Gavves, E.; and Gould, S. 2017. Self-supervised video representation learning with odd-one-out networks. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 5729–5738.
- Gan, C.; Gong, B.; Liu, K.; Su, H.; and Guibas, L. J. 2018. Geometry guided convolutional neural networks for self-supervised video representation learning. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 5589–5597.
- Hara, K.; Kataoka, H.; and Satoh, Y. 2018. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*.
- Ji, S.; Xu, W.; Yang, M.; and Yu, K. 2013. 3d convolutional neural networks for human action recognition. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)* 35(1):221–231.
- Karpathy, A.; Toderici, G.; Shetty, S.; Leung, T.; Sukthankar, R.; and Fei-Fei, L. 2014. Large-scale video classification with convolutional neural networks. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 1725–1732.
- Kay, W.; Carreira, J.; Simonyan, K.; Zhang, B.; Hillier, C.; Vijayanarasimhan, S.; Viola, F.; Green, T.; Back, T.; Natsev, P.; Suleyman, M.; and Zisserman, A. 2017. The kinetics human action video dataset. *CoRR* abs/1705.06950.
- Kim, D.; Cho, D.; Yoo, D.; and Kweon, I. S. 2017. Two-phase learning for weakly supervised object localization. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 3554–3563. IEEE.
- Kim, D.; Cho, D.; Yoo, D.; and Kweon, I. S. 2018. Learning image representations by completing damaged jigsaw puzzles. In *Proc. of Winter Conference on Applications of Computer Vision (WACV)*, 793–802.
- Kuehne, H.; Jhuang, H.; Garrote, E.; Poggio, T.; and Serre, T. 2011. HMDB: a large video database for human motion recognition. In *Proc. of Int'l Conf. on Computer Vision (ICCV)*.
- Lee, H.-Y.; Huang, J.-B.; Singh, M.; and Yang, M.-H. 2017. Unsupervised representation learning by sorting sequences. In *Proc. of Int'l Conf. on Computer Vision (ICCV)*.
- Misra, I.; Zitnick, C. L.; and Hebert, M. 2016. Shuffle and learn: unsupervised learning using temporal order verification. In *Proc. of European Conf. on Computer Vision (ECCV)*, 527–544.
- Mobahi, H.; Collobert, R.; and Weston, J. 2009. Deep learning from temporal coherence in video. In *Proceedings of the 26th Annual International Conference on Machine Learning*, 737–744. ACM.
- Mundhenk, T. N.; Ho, D.; and Chen, B. Y. 2018. Improvements to context based self-supervised learning. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*.
- Noroozi, M., and Favaro, P. 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In *Proc. of European Conf. on Computer Vision (ECCV)*, 69–84.
- Oquab, M.; Bottou, L.; Laptev, I.; and Sivic, J. 2015. Is object localization for free?—weakly-supervised learning with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Pathak, D.; Krähenbühl, P.; Donahue, J.; Darrell, T.; and Efros, A. A. 2016. Context encoders: Feature learning by inpainting. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2536–2544.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *Int'l Journal of Computer Vision (IJCV)* 115(3).
- Simonyan, K., and Zisserman, A. 2014. Two-stream convolutional networks for action recognition in videos. In *Proc. of Neural Information Processing Systems (NIPS)*, 568–576.
- Soomro, K.; Zamir, A. R.; and Shah, M. 2012. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*.
- Tran, D.; Bourdev, L.; Fergus, R.; Torresani, L.; and Paluri, M. 2015. Learning spatiotemporal features with 3d convolutional networks. In *Proc. of Int'l Conf. on Computer Vision (ICCV)*, 4489–4497.
- Vondrick, C.; Pirsivash, H.; and Torralba, A. 2016. Generating videos with scene dynamics. In *Advances In Neural Information Processing Systems*, 613–621.
- Wang, X., and Gupta, A. 2015. Unsupervised learning of visual representations using videos. In *Proc. of Int'l Conf. on Computer Vision (ICCV)*, 2794–2802.
- Wang, L.; Xiong, Y.; Wang, Z.; Qiao, Y.; Lin, D.; Tang, X.; and Gool, L. V. 2016. Temporal segment networks: Towards good practices for deep action recognition. In *Proc. of European Conf. on Computer Vision (ECCV)*, 20–36.
- Wei, D.; Lim, J.; Zisserman, A.; and Freeman, W. T. 2018. Learning and using the arrow of time. In *CVPR*.
- Zhao, Y.; Deng, B.; Shen, C.; Liu, Y.; Lu, H.; and Hua, X.-S. 2017. Spatio-temporal autoencoder for video anomaly detection. In *Proc. of Multimedia Conference (MM)*.