

Disentangling Motion, Foreground and Background Features in Videos

Xunyu Lin*
Beihang University
Beijing, China
xunyu.lin2017@outlook.com

Víctor Campos
Barcelona Supercomputing Center
Barcelona, Catalonia/Spain
victor.campos@bsc.es

Xavier Giro-i-Nieto
Universitat Politecnica de Catalunya
Barcelona, Catalonia/Spain
xavier.giro@upc.edu

Jordi Torres
Barcelona Supercomputing Center
Barcelona, Catalonia/Spain
jordi.torres@bsc.es

Cristian Canton Ferrer
Microsoft
Redmond (WA), USA
cristian.canton@microsoft.com

Abstract

This paper introduces an unsupervised framework to extract semantically rich features for video representation. Inspired by how the human visual system groups objects based on motion cues, we propose a deep convolutional neural network that disentangles motion, foreground and background information. The proposed architecture consists of a 3D convolutional feature encoder for blocks of 16 frames, which is trained for reconstruction tasks over the first and last frames of the sequence. A preliminary supervised experiment was conducted to verify the feasibility of proposed method by training the model with a fraction of videos from the UCF-101 dataset taking as ground truth the bounding boxes around the activity regions. Qualitative results indicate that the network can successfully segment foreground and background in videos as well as update the foreground appearance based on disentangled motion features. The benefits of these learned features are shown in a discriminative classification task, where initializing the network with the proposed pretraining method outperforms both random initialization and autoencoder pretraining. Our model and source code are publicly available at https://allenovo.github.io/cvprw17_webpage/.

1. Introduction

Unsupervised learning has long been an intriguing field in artificial intelligence. Human and animal learning is largely unsupervised: we discover the structure of the world mostly by observing it, not by being told the name of every object, which would correspond to supervised learning [3]. A sys-

tem capable of predicting what is going to happen by just watching large collections of unlabeled video data needs to build an internal representation of the world and its dynamics [5]. When considering the vast amount of unlabeled data generated every day, unsupervised learning becomes one of the key challenges to solve in the road towards general artificial intelligence.

Based on how a human would provide a high level summary of a video, we hypothesize that there are three key components to understand such content: namely *foreground*, *motion* and *background*. These three elements would tell us, respectively, what the main objects in the video are, what they are doing and where their location. We propose a framework that explicitly disentangles these three components in order to build strong features for action recognition, where the supervision signals can be generated without requiring from expensive and time consuming human annotations. The proposal is inspired by how infants who have no prior visual knowledge tend to group things that move as connected wholes and also move separately from one another [8]. Based on this intuition, we can build a similar unsupervised pipeline to segment foreground and background with global motion, i.e. the rough moving directions of objects. Such segmented foregrounds across the video can be used to model both the global motion (e.g. transition or stretch) and local motion (i.e. transformation of detailed appearance) from a pair of foregrounds at different time steps. Since background motion is mostly given by camera movements, we restrict the use of motion to the foreground and rely on appearance to model the former.

The contributions of this work are two-fold: (1) disentangling motion, foreground and background features in videos by human alike motion aware mechanism and (2) learning strong video features that improve the performance of action recognition task.

*This work was developed during an exchange period of Xunyu Lin at the Universitat Politecnica de Catalunya.

2. Related Work

Leveraging large collections of unlabeled videos has proven beneficial for unsupervised training of image models thanks to the implicit properties they exhibit in the temporal domain, e.g. visual similarity between patches in consecutive frames [14] and temporal coherence and order [6]. Since learning to predict future frames forces the model to construct an internal representation of the world dynamics, several works have addressed such task by predicting global features of future frames with Recurrent Neural Networks (RNN) [11] or pixel level predictions by means of multi-scale Convolutional Neural Networks (CNN) trained with an adversarial loss [5]. The key role played by motion has been exploited for future frame prediction tasks by explicitly decomposing content and motion [12] and for unsupervised training of video-level models [4]. Similarly in spirit, separate foreground and background streams have been found to increase the quality of generative video models [13].

Techniques exploiting explicit foreground and background segmentations in video generally require from expensive annotation methods, limiting their application to labeled data. However, the findings by Pathak et al. [7] show how models trained on noisy annotations learn to generalize and perform well when finetuned for other tasks. Such noisy annotations can be generated by unsupervised methods, thus alleviating the cost of annotating data for the target task. In this work we study our proposed method by using manual annotations, whereas evaluating the performance drop when replacing such annotations with segmentations generated in an unsupervised manner remains as future work.

3. Methodology

We adopt an autoencoder-styled architecture to learn features in an unsupervised manner. The encoder maps input clips to feature tensors by applying a series of 3D convolutions and max-pooling operations [1]. Unlike traditional autoencoder architectures, the bottleneck features are partitioned into three splits which are then used as input for three different reconstruction tasks, as depicted in Figure 1.

Disentangling of foreground and background: depending on the nature of the training data, reconstruction of frames may become dominated either by the foreground or background. We explicitly split the reconstruction task to guarantee that none of the parts dominates over the other. Partitioned foreground and background features will be passed into two different decoders for reconstruction. While segmentation masks are often obtained by manual labeling, it is worth noting they can be obtained without supervision as well, e.g. by using methods based on motion perceptual grouping such as uNLC [7]. The latter approach has proven beneficial for unsupervised pre-training of CNNs [7].

Disentangling of foreground motion: leveraging motion

information can provide a boost in action recognition performance when paired with appearance models [9]. We encourage the model to learn motion-related representations by solving a predictive learning task where the foreground in the last frame needs to be reconstructed from the foreground in the first frame. Given a pair of foregrounds at timesteps t_1 and t_2 , namely (f_{t_1}, f_{t_2}) , we aim to estimate a function M from motion features $m_{t_1 \rightarrow t_2}$ throughout t_1 and t_2 that maps f_{t_1} to f_{t_2} in deep feature space G :

$$G(f_{t_2}) = M(G(f_{t_1}), m_{t_1 \rightarrow t_2}) \quad (1)$$

Throughout this work, the space of encoded features is used for G , and M is parametrized by a deterministic version of cross convolution [15]. The foreground decoder weights are shared among all foreground reconstruction task. Gradients coming from the reconstruction of f_{t_2} are blocked from backpropagating through $G(f_{t_1})$ during training to prevent $G(f_{t_1})$ from storing information about f_{t_2} .

Frame selection: assuming that the background semantics stay close throughout the short clips, only the background in the first frame is reconstructed. First and last frames are chosen to perform foreground reconstruction, since they represent the most challenging pair in the clip.

Loss function: the model is optimized to minimize the L1 loss between the original frames and their reconstruction. In particular, the loss function is defined from a decomposition of the input video volume x of T frames into the foreground x_{fg} and background x_{bg} volumes:

$$\begin{aligned} x_{fg} &= x \cdot b_{fg} \\ x_{bg} &= x \cdot (1 - b_{fg}) \end{aligned} \quad (2)$$

where b_{fg} corresponds to a volume of binary values, so that 1 correspond to foreground pixels and 0 to the background ones.

This decomposition allows defining the reconstruction loss $L_{rec}(x)$ over the video volume x as the sum of three terms:

$$L_{rec}(x) = L_{fg}^1(x) + L_{bg}^1(x) + L_{fg}^T(x) \quad (3)$$

where the components L_{fg}^1 , L_{bg}^1 and L_{fg}^T represent the reconstruction loss for the first foreground and first background, and last foreground, respectively. These three terms are particularizations at the first ($t = 1$) and last ($t = T$) frames of the generic foreground $L_{fg}^t(x)$ and background $L_{bg}^t(x)$ reconstructions losses:

$$L_{fg}^t(x) = \frac{1}{A^t} \sum_{i,j} W^t[i,j] \cdot |\hat{x}_{fg}^t[i,j] - x_{fg}[i,j,t]| \quad (4)$$

$$L_{bg}^t(x) = \frac{1}{A^t} \sum_{i,j} |\hat{x}_{bg}^t[i,j] - x_{bg}[i,j,t]| \quad (5)$$




Figure 1: System architecture. Please note that in this work, the masks used to generate ground truth are from manual annotations while uNLC will be utilized in our future work.

where \hat{x}^t denotes a reconstructed foreground/ background at time t , A^t is the area of the reconstructed frame at time t , and W^t is an element-wise weighting mask at time t designed to leverage the focus between the foreground and background pixels:

$$W^t[i, j] = \begin{cases} 1 & \text{if } (i, j) \in \text{background} \\ \max\left[1, \frac{A_{bg}^t}{A_{fg}^t}\right] & \text{if } (i, j) \in \text{foreground} \end{cases} \quad (6)$$

During preliminary experiments, we observed that the reconstruction of the first foreground always outperformed the reconstruction of the last one by a large margin, given the increased difficulty of the latter task. In order to get finer reconstruction of the last foreground, we introduce an L2 loss L_{feat} on $G(ft_2)$. The pseudo ground truth for this task is obtained by getting first foreground features from the encoder fed with the temporally reversed clip. The final loss to optimize is the following:

$$L_{total}(x) = L_{rec}(x) + L_{feat}(x) \quad (7)$$

4. Experimental setup

Please note again we are showing results trained with ground truth masks to check the feasibility of our proposal and the pure unsupervised framework generating masks from uNLC [7] remains as future work.

Dataset: there are 24 classes out of 101 in UCF-101 with localization annotations [10, 2]. Following [7], we first evaluate the proposed framework with supervised annotations and use the bounding boxes in the subset of UCF-101 for such purpose. Evaluating the proposal in weak annotations collected by means of unsupervised methods remains as

future work. We follow the original splits of training and test set and also split 10% videos out of the training set as validation set in order to perform early stopping and prevent the network from overfitting the training data.

Training details: videos are split into clips of 16 frames each. These clips are then resized to 128×128 and their pixel values are scaled and shift to $[-1, 1]$. The clips are randomly temporally or horizontally flipped for data augmentation. Weight decay with rate of 10^{-3} is added as regularization. The network is trained for 125 epochs with Adam optimizer and a learning rate of 10^{-4} on batches of 40 clips.

5. Results

We tested our model on test set for reconstruction task. For better demonstrating the efficiency of our proposed pre-training pipeline, we also trained the network to do action recognition with pretrained features.

Reconstruction task: reconstruction results on test set are shown in Figure 2. From these results, we can clearly see that the network already can predict similar foreground segmentation as ground truth. However, the image reconstructions are still blurry. We argue that this is due to the properties of the L1 loss we are adopting [5]. One interesting fact is that the network has learned to generalize foreground to some other moving objects in the scene even though they are not included in the annotations. For example, the result shown in the top-right corner: instead of only segmenting the person, the dog walking beside the person is also included. This fact suggests that the network has successfully learned to identify foreground from motion cues.

Besides from foreground and background features, these results also demonstrate a good extraction of motion features. The learned motion features contain both global motions,




Figure 2: Reconstruction results on the test set. For each example, the top row shows the reconstruction while the bottom one contains the ground truth. Each column shows the segmentation of foreground in first frame, background in first frame and foreground in last frame, respectively.

e.g. transition of foreground, and local motions, e.g. change of human pose. In the bottom-center result, the generated kernels from motion feature successfully shift the object from right to the middle and change its gesture.

Action recognition: a good pretraining pipeline should show better performance on some typical discriminative tasks than random initialization, especially when training data is scarce [6, 4, 7, 13, 14]. We also conducted comparative experiments on the task of action recognition. By discarding the decoders in our framework and training a linear softmax layer on top of the disentangled features, we can obtain a simple network for action recognition. For the first experiment, we first pretrain our encoder on the subset of UCF-101 with the settings discussed above and then fine-tune the whole action recognition network with added softmax layer on the same subset. As baselines, we trained another two action recognition networks, one with all weights initialized randomly and another one pretrained with an unsupervised autoencoder architecture. This autoencoder shared the same 3D convolutional encoder architecture with ours, while its decoder was the mirrored version of the encoder but replacing the pooling operations with convolutions.

During training, we observed that our pretrained model reached 90% accuracy on training set immediately after one epoch while the randomly initialized network took 130 epochs to achieve it. All three models reached around 96% accuracy at the end of training and encountered severe overfitting problems. The accuracy of different methods on the

Method	Accuracy
Random initialization	52.2%
Pretrained (autoencoder)	56.8%
Pretrained (ours)	62.5%

Table 1: Action recognition accuracy of different methods on the test subset of UCF-101.

validation set during training time is shown in Figure 3. The best accuracy obtained on the test set with our pretrained model is 62.5%, while it drops to 52.2% and 56.8% respectively when using a random initialization and autoencoder as pretraining scheme, as shown in Table 1. We observe a margin of more than 10% on accuracy between our proposed method and random initialization on both validation set and test set. This further demonstrates that with our proposal, the network can learn features that generalize better. These results are specially promising given the small amount of data used during pretraining, which is just a fraction of UCF-101. While this demonstrates the efficiency of the approach, using a larger dataset for pretraining should provide additional gains and better generalization capabilities.

6. Conclusions

This work has proposed a novel framework towards an unsupervised learning of video features capable of disentangling of motion, foreground and background. Our method




Figure 3: Action recognition results on validation set. This figure shows the accuracy of each method on validation set during the training time.

mostly exploits motion in videos and is inspired by human perceptual grouping with motion cues. Our experiments using ground truth boxes render convincing results on both frame reconstruction and action recognition, showing the potential of the proposed architecture.

However, multiple aspects still need to be explored in our work. As our plans for the future work, we decide to (1) introduce unsupervised learning for foreground segmentation as well, as proposed in uNLC [7]; (2) train with a larger amount of unlabeled data; (3) introduce adversarial loss to improve the sharpness of the reconstructed frames [5]; and (4) fill the gap of absent motion features between the first frame and the last frame by reconstructing any random frame in the clip.

Our model and source code are publicly available at https://allenovo.github.io/cvprw17_webpage/.

References

- [1] T. Du, B. Lubomir, F. Rob, T. Lorenzo, and P. Manohar. C3D: generic features for video analysis. In *ICCV*, 2015.
- [2] A. Gorban, H. Idrees, Y.-G. Jiang, A. Roshan Zamir, I. Laptev, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes. <http://www.thumos.info/>, 2015.
- [3] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. In *Nature*, 2015.
- [4] Z. Luo, B. Peng, D.-A. Huang, A. A. Alahi, and L. Fei-Fei. Unsupervised learning of long-term motion dynamics for videos. In *CVPR*, 2017.
- [5] M. Mathieu, C. Couprie, and Y. LeCun. Deep multi-scale video prediction beyond mean square error. In *ICLR*, 2016.
- [6] I. Misra, C. L. Zitnick, and M. Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *ECCV*, 2016.
- [7] D. Pathak, R. Girshick, P. Dollár, T. Darrell, and B. Hariharan. Learning features by watching objects move. In *CVPR*, 2017.
- [8] E. S. Spelke. Principles of object perception. In *Cognitive Science*, 1990.
- [9] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 2014.
- [10] K. Soomro, A. Roshan Zamir, and M. Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. In *CRCV-TR-12-01*, 2012.
- [11] N. Srivastava, E. Mansimov, and R. Salakhutdinov. Unsupervised learning of video representations using lstms. In *ICML*, 2015.
- [12] R. Villegas, J. Yang, S. Hong, X. Lin, and H. Lee. Decomposing motion and content for natural video sequence prediction. In *ICLR*, 2017.
- [13] C. Vondrick, H. Pirsaviash, and A. Torralba. Generating videos with scene dynamics. In *NIPS*, 2016.
- [14] X. Wang and A. Gupta. Unsupervised learning of visual representations using videos. In *ICCV*, 2015.
- [15] T. Xue, J. Wu, K. L. Bouman, and W. T. Freeman. Visual dynamics: Probabilistic future frame synthesis via cross convolutional networks. In *NIPS*, 2016.