

Human Action Recognition Without Human

Yun He, Soma Shirakabe, Yutaka Satoh, and Hirokatsu Kataoka^(✉)

National Institute of Advanced Industrial Science and Technology (AIST),
Tsukuba, Ibaraki, Japan
{yun.he,shirakabe-s,yu.satou,hirokatsu.kataoka}@aist.go.jp

Abstract. The objective of this paper is to evaluate “human action recognition without human”. Motion representation is frequently discussed in human action recognition. We have examined several sophisticated options, such as dense trajectories (DT) and the two-stream convolutional neural network (CNN). However, some features from the background could be too strong, as shown in some recent studies on human action recognition. Therefore, we considered whether a background sequence alone can classify human actions in current large-scale action datasets (e.g., UCF101).

In this paper, we propose a novel concept for human action analysis that is named “human action recognition without human”. An experiment clearly shows the effect of a background sequence for understanding an action label.

1 Introduction

An effective motion representation is in demand for action recognition, event recognition, and video understanding. In human action recognition especially, several survey papers have been published in the last two decades [1, 2, 10, 11]. We have investigated a more reliable and faster algorithm to put action recognition into practice. The target applications of action recognition can be easily imagined, for example, surveillance, robotics, augmented reality, and intelligent surgery. However, current vision-based video representations focus on the media to improve the recognition rate on UCF101 [16], HMDB51 [7], and ActivityNet [4].

Here we categorize action recognition into two types: direct and contextual approaches.

The direct approach, which is motion representation, has been studied in action recognition. Since Laptev *et al.* proposed space-time interest points (STIP) [8, 9], *xyt* keypoint acquisition has been well established in temporal representation. STIP is significantly improved with densely connected keypoints in the dense trajectories approach (DT) [17, 18]. The DT is a more natural approach for understanding whole body motions because it uses a large amount of tracked keypoints. Recently, two-stream CNN has been applied as a representative method in action recognition [15]. The two-stream convolutional neural network (CNN) uses spatial and temporal streams to extract appearance and motion

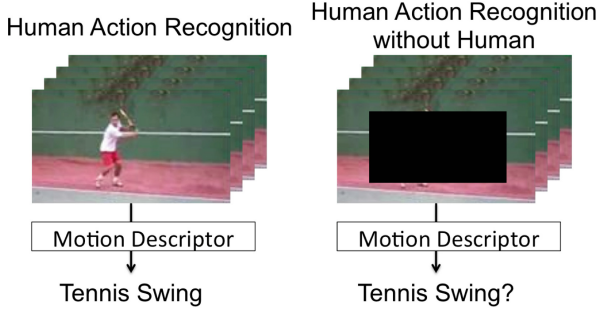


Fig. 1. Human action recognition (left) and human action recognition without human (right): We simply replace the center-around area with a black background in an image sequence. We evaluate the performance rate with only the limited background sequence as a contextual cue.

features from RGB and optical flow input. The classification scores at each stream are fused for evaluating an objective video. Other CNN-based approaches apply a dynamic scene descriptor such as a pooled time series (PoT) [14] and capture sensitive motion with a subtle motion descriptor (SMD) [6].

The contextual approach is focused around the region of a human and can provide an important cue to improve human action recognition. In related work, Jain *et al.* [5] and Zhou *et al.* [20] showed that object and scene context aid in the recognition of human actions. Jain *et al.* carried out an evaluation of how much object usage is needed for action recognition [5]. They combined object information with a classifier score into the improved DT (IDT) plus Fisher vectors (FVs) [12] as a motion feature from a human area. A large number of object labels (15,923 objects), e.g., computer and violin, are corresponded to an output function with AlexNet as an object prior. The response of CNN-based object information must be combined with a motion vector for a richer understanding of human actions. In their evaluation, motion + object vector allow us to obtain a better feature in an image sequence. When using object information, the performance rate rises by +3.9%, +9.9%, and 0.5% on UCF101, the THUMOS14 validation set, and KTH, respectively. According to experiments, the object vector improves recognition accuracy on a large-scale action database. Zhou *et al.* proposed a combination of a contextual human-object interaction and a motion feature for fine-grained action recognition [20]. Object proposals are captured by using BING [3]. However, some useless proposals are generated around a human area. The pruning of extra regions is executed by referring to dense trajectories around object proposals. The recognition rate can be improved with human-object interaction as a mid-level feature. The mid-level feature records an outstanding rate 72.4% on the MPII cooking dataset [13], which is known as a fine-grained action database. These two examples are convincing enough to integrate a mid-level feature into a motion vector.

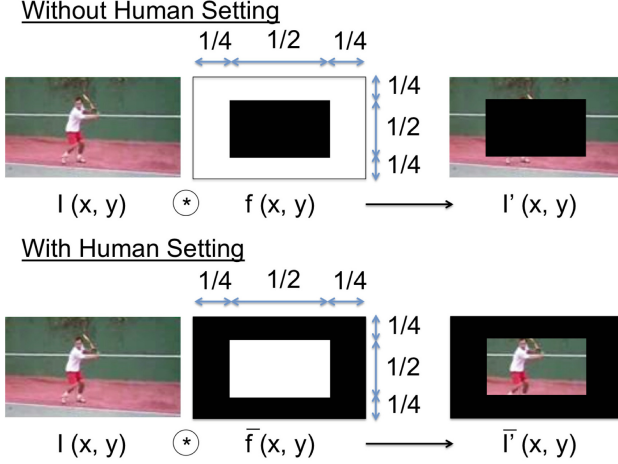


Fig. 3. Image filtering for human action recognition without human.

where \bar{I}' and filter \bar{f} are an inverse image and filter in the setting without a human. The background is eliminated with the inverse filter at the bottom of Fig. 3.

Training of Two-Stream CNN. The learning parameters of the spatial and temporal streams are based on [19]. Our goal is to predict the video label without additional training in the setting without a human (see Fig. 1). By using an original pre-trained model [19], we obtained the following results on UCF101 split 1: 74.86 % (spatial), 80.33 % (temporal), and 84.30 % (two-stream)¹, as shown in Table 1.

3 Experiment

Dataset. We apply the well-studied UCF101 dataset. This large-scale dataset was mainly collected from YouTube videos of sports and musical instrument performance scenes. The recognition task is to predict an action label from a given video. The dataset contains several computer vision difficulties, e.g., camera motion, scaling, posture change, and viewpoint difference. The mean average accuracy is calculated with three training and test splits. Here we calculate an average precision with training/test split 1.

Quantitative Evaluation. Table 2 shows the performance rate on the UCF101 dataset with or without a human. Surprisingly, the two-stream CNN performance was 47.42 % in the setting without a human. We understand that a motion

¹ Our implementation is different from the report of Wang [19]. The performance rate depends on the parameter tuning. They reported 79.8 % (spatial), 85.7 % (temporal) and 90.9 % (two-stream) on UCF101 split 1.

Table 1. Performance rate on the UCF101 dataset with baseline two-stream CNN


Stream	% on UCF101 (split 1)
Spatial stream	74.86
Temporal stream	80.33
Two-stream (S+T) [19]	84.30

Table 2. Performance rate of human action recognition with or without a human

With or without a human	Stream	% on UCF101 (split 1)
With human	Spatial stream	51.26
	Temporal stream	40.50
	Two-stream	56.91
Without human	Spatial stream	45.33
	Temporal stream	26.80
	Two-stream	47.42

recognition approach relies on a background sequence. The spatial stream is +18.53 % better than the temporal stream. Therefore, an appearance tends to classify between backgrounds. Motion features contribute slightly to the background classification; that is, the performance rate is increased +2.09 % with the temporal stream. The two-stream CNN recorded 56.91 % in the with human setting, which is +9.49 % higher than the setting without a human.

Qualitative Dataset Evaluation. Figure 4 shows examples without a human setting on the UCF101 dataset. Where we evaluated partial and complete images without a human (Figs. 4(a) and (b), respectively), the number of partial images without a human was 1,114 in 3,783 videos. The rate was 29.45 % in UCF101 split 1. The complete images without a human were not found on the videos.

**Fig. 4.** Qualitative evaluation of the setting without a human on the UCF101 dataset.

4 Conclusion

To the best of our knowledge, this is the first study of human action recognition without human. However, we should not have done that kind of thing. The motion representation from a background sequence is effective to classify videos in a human action database. We demonstrated human action recognition in with and without a human settings on the UCF101 dataset. The results show the setting without a human (47.42 %) was close to the setting with a human (56.91 %). We must accept this reality to realize better motion representation.

References

1. Aggarwal, J.K., Cai, Q.: Human motion analysis: a review. *Comput. Vis. Image Underst. (CVIU)* **73**, 428–440 (1999)
2. Aggarwal, J.K., Ryoo, M.S.: Human activity analysis: a review. *ACM Comput. Surv.* (2011)
3. Cheng, M.-M., Zhang, Z., Lin, W.-Y., Torr, P.: BING: binarized normed gradients for objectness estimation at 300fps. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2014)
4. Heilbron, F.C., Escorcia, V., Ghanem, B., Niebles, J.C.: Activitynet: a large-scale video benchmark for human activity understanding. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015)
5. Jain, M., van Gemert, J.C., Snoek, C.G.M.: What do 15,000 object categories tell us about classifying and localizing actions? In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015)
6. Kataoka, H., Miyashita, Y., Hayashi, M., Iwata, K., Satoh, Y.: Recognition of transitional action for short-term action prediction using discriminative temporal CNN feature. In: *British Machine Vision Conference (BMVC)* (2016)
7. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: HMDB: a large video database for human motion recognition. In: *International Conference on Computer Vision (ICCV)* (2011)
8. Laptev, I.: On space-time interest points. *Int. J. Comput. Vis. (IJCV)* **64**, 107–123 (2005)
9. Laptev, I., Lindeberg, T.: Space-time interest points. In: *International Conference of Computer Vision (ICCV)* (2003)
10. Moeslund, T.B., Hilton, A., Kruger, V.: A survey of advances in vision-based human motion capture and analysis. *Comput. Vis. Image Underst. (CVIU)* **104**, 90–126 (2006)
11. Moeslund, T.B., Hilton, A., Kruger, V., Sigal, L.: *Visual Analysis of Humans: Looking at People*. Springer, Heidelberg (2011)
12. Perronnin, F., Sánchez, J., Mensink, T.: Improving the Fisher kernel for large-scale image classification. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) *ECCV 2010. LNCS*, vol. 6314, pp. 143–156. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-15561-1_11](https://doi.org/10.1007/978-3-642-15561-1_11)
13. Rohrbach, M., Amin, S., Andriluka, M., Schiele, B.: A database for fine grained activity detection of cooking activities. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2012)

14. Ryoo, M.S., Rothrock, B., Matthies, L.: Pooled motion features for first-person videos. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015)
15. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition. In: Neural Information Processing Systems (NIPS) (2014)
16. Soomro, K., Zamir, A.R., Shah, M.: Ucf101: a dataset of 101 human action classes from videos in the wild. CRCV-TR-12-01 (2012)
17. Wang, H., Klaser, A., Schmid, C.: Dense trajectories and motion boundary descriptors for action recognition. *Int. J. Comput. Vis. (IJCV)* **103**, 60–79 (2013)
18. Wang, H., Klaser, A., Schmid, C., Cheng-Lin, L.: Action recognition by dense trajectories. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2011)
19. Wang, L., Xiong, Y., Wang, Z., Qiao, Y.: Towards good practices for very deep two-stream convnets. arXiv pre-print [1507.02159](https://arxiv.org/abs/1507.02159) (2015)
20. Zhou, Y., Ni, B., Hong, R., Wang, M., Tian, Q.: Interaction part mining: a mid-level approach for fine-grained action recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015)