

JOINT LEARNING ON THE HIERARCHY REPRESENTATION FOR FINE-GRAINED HUMAN ACTION RECOGNITION

Mei Chee Leong¹, Hui Li Tan¹, Haosong Zhang^{1,2}, Liyuan Li¹, Feng Lin², Joo Hwee Lim^{1,2}
{leong_mei_chee, hltan}@i2r.a-star.edu.sg, haosong001@e.ntu.edu.sg,
asflin@ntu.edu.sg, {lyli, jooHwee}@i2r.a-star.edu.sg

Institute for Infocomm Research (I²R), A*STAR¹
School of Computer Science and Engineering, Nanyang Technological University, Singapore²

ABSTRACT

Fine-grained human action recognition is a core research topic in computer vision. Inspired by the recently proposed hierarchy representation of fine-grained actions in FineGym and SlowFast network for action recognition, we propose a novel multi-task network which exploits the FineGym hierarchy representation to achieve effective joint learning and prediction for fine-grained human action recognition. The multi-task network consists of three pathways of SlowOnly networks with gradually increased frame rates for events, sets and elements of fine-grained actions, followed by our proposed integration layers for joint learning and prediction. It is a two-stage approach, where it first learns deep feature representation at each hierarchical level, and is followed by feature encoding and fusion for multi-task learning. Our empirical results on the FineGym dataset achieve a new state-of-the-art performance, with 91.80% Top-1 accuracy and 88.46% mean accuracy for element actions, which are 3.40% and 7.26% higher than the previous best results.

Index Terms— Action recognition, fine-grained action recognition, multi-task learning, joint representation

1. INTRODUCTION

As a fundamental research topic in computer vision, action recognition has always been an important research problem. Driven by advancements in deep learning models, such as 2D and 3D convolutional neural network (CNN) models, record breaking accuracies have been achieved on state-of-the-art action recognition datasets, such as ActivityNet [1], Kinetics [2] and Sports1M [3]. Nonetheless, fine-grained action recognition, where the appearances of action and background are very similar and the recognition relies more on the modeling of temporal dynamics, remains a challenging research problem. To address fine-grained action recognition, various

fine-grained datasets [4, 5, 6, 7] have been recently proposed. Particularly, FineGym [7] is a gymnastic video dataset with structured temporal and semantics hierarchy, decomposing an action into coarse-to-fine event-level, set-level, and element-level annotations. The scale and diversity of FineGym poses significant challenges for fine-grained action recognition, involving not only recognition of complex dynamics of intense motions, but also differentiation in subtle spatial semantics. Hierarchy representation of fine-grained actions provides discriminating and complementary semantic, contextual and motion information at different granularity levels. However, this information is not fully utilized for joint representation and recognition, as most of the existing methods still focus on learning and prediction of single level action class.

In this paper, inspired by the FineGym hierarchy representation, we propose a multi-task network for effective joint representation and learning for fine-grained action recognition. First, we introduce a new deep architecture according to the three-level hierarchy for effective learning of multi-level representation of fine-grained actions. The structure of the new architecture is designed to effectively learn and represent the associations and constraints across different levels of fine-grained action representation. Second, we propose a compact network of joint learning layers with multi-task learning to achieve joint learning and prediction of the three-level action class of the fine-grained action. Our experimental results on the FineGym benchmark dataset have shown a clear improvement on the state-of-the-art performance.

2. RELATED WORKS

There are two main frameworks for fine-grained action recognition. The first type is the 2D+1D framework [8, 9, 10, 11] where a 2D CNN model is used for extracting spatial features, and is followed by a 1D model for temporal aggregation. TSN [8] samples frames from segmented video, and then aggregates the scores of each segment for final prediction. TRN [9] introduces a temporal relational reasoning module that allows aggregation of multi-scale temporal relations between frames.

This research is supported by the Agency for Science, Technology and Research (A*STAR) under its AME Programmatic Funding Scheme (Project A18A2b0046).

TSM [10] proposes a temporal shift module that achieves 3D performance with 2D complexity by shifting some channels along the temporal dimension. ActionVLAD [11] aggregates spatio-temporal features from two-stream networks by pooling across space and time streams.

The second type of framework is based on the 3D CNN model [12, 13, 14, 15]. I3D [13] expands 2D CNN into 3D CNN by adding a temporal dimension on filters and pooling kernels to capture the spatio-temporal features. Non-local neural network [14] computes weighted sum of features in space-time to capture long-range dependencies. SlowFast [15] comprises a slow and fast pathway to capture spatial semantic at low frame rate and learn motion information at high frame rate. Lateral connections are used to aggregate the features from the two pathways. SlowOnly is a variant of SlowFast without the fast pathway. There are also hybrid framework that fuses 2D and 3D CNN in a network [16, 17, 18, 19, 20].

In existing works, investigations mainly focus on recognizing each level of the fine-grained actions separately. Different from existing works, we investigate joint spatio-temporal representation, learning and prediction of the hierarchical action class in fine-grained action recognition.

3. PROPOSED METHOD

To exploit the hierarchical action representation and make use of the complementary contexts and semantics across different levels, we propose a multi-task network that consists of three input pathways for each hierarchy level and a compact integration network to jointly learn and predict the coarse-to-fine action representations.

3.1. Hierarchy Representation of Fine-Grained Actions

In [7], a hierarchy representation format is proposed to represent fine-grained human actions. The hierarchy tree is composed of three levels, *i.e.*, *events*, *sets*, and *elements*. The events represent the coarsest level of action categories. In FineGym, it represents the actions of different gymnastic routines. The sets represent middle level sub-action categories under the corresponding event node in the hierarchy tree. Each sub-action category represents a cluster of several technically and visually similar element actions. The elements represent the sub-action categories of the finest granularity under the corresponding set node in the hierarchy tree.

The hierarchy representation provides a multi-level coarse-to-fine representation of fine-grained human actions. There are several advantages of such representation for both learning and prediction. First, it represents a large set of fine-grained actions as a multi-level semantic hierarchy tree. The semantic hierarchy provides a domain constraint and solid foundation for joint learning and perception for coarse-level

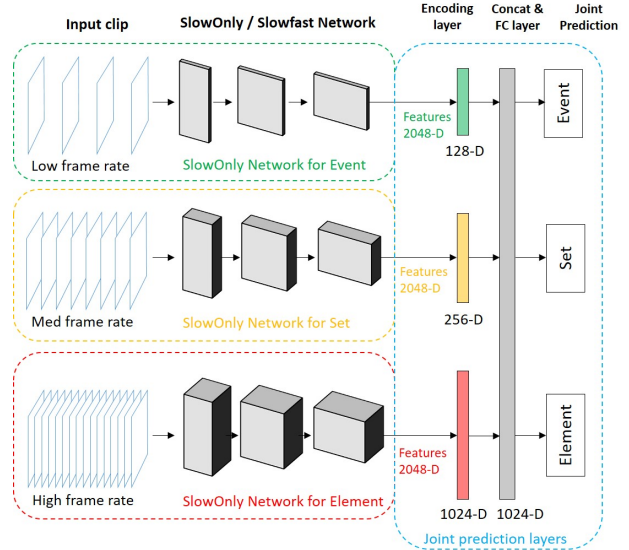


Fig. 1. The architecture of our multi-task network for effective representation and joint learning of three-level hierarchy of fine-grained human actions.

to fine-grained action understanding. Second, the hierarchy annotations are derived on a decision tree. Hence, the hierarchy representation encodes the distinctive features of semantic meaning, visual appearance and motion information on different levels. For example, some element actions can be distinguished from their background, objects and viewing angles, while some actions are mainly different in fine-grained motion features and duration. Training a single model on the element level actions may not be sufficient to learn the discriminative feature representation in a hierarchy structure. In addition, the hierarchy structure also encodes domain knowledge of action classes. Hence, exploiting multi-level action representations may help in better aggregation of spatial semantics and temporal dynamics, and learning of the domain class information implicitly. Such semantic information and domain knowledge might not be well exploited when training on a single level, especially the element-level for fine-grained action recognition.

3.2. Network Architecture for Multi-task Learning

Inspired by the SlowFast network that captures spatial and temporal semantics from its slow and fast pathways, we design a three-pathway multi-task network for joint learning of the hierarchical representation of fine-grained actions. Specifically, the network has Event pathway, Set pathway and Element pathway where each pathway samples the input clip at different frame rate and output features for our joint prediction layers. We first encode the features from each pathway to a lower dimensional vector, before fusing them for joint prediction. The encoded features from all pathways are con-

catenated for multi-task prediction on the hierarchical action representation. The architecture of the proposed network is illustrated in Fig. 1, where the base model for each pathway is a SlowOnly network trained for single level prediction.

3.2.1. Base model for individual pathways

The observations in [15] suggest that, action representation such as the event class at the coarsest granularity is more related to visual appearance and background context, and it relies less on motion information. On the other hand, the element class at the finest granularity relies more on detailed motion information and has subtle differences in visual appearance [7]. According to the hierarchy representation, we propose to use a slow pathway of low frame rate for event class which captures semantic information from a few sparsely sampled frames, a medium pathway of moderate frame rate for set class which captures increasing motion information from increased sample frames, and a fast pathway of high frame rate for element class which captures the detailed motion information from densely sampled frames. Following the notation in [15], we denote the feature shape of a pathway as $\{T, S^2, C\}$, where T is the number of the frames sampled, S represents the spatial size, and C denotes the number of channels.

3.2.2. Joint Prediction Layers

After training the base models for each hierarchy level, we combine them to form the three pathways of our multi-task network. The pre-trained features for Event, Set and Element pathways are encoded from the output vector (2048-D) to a lower dimension vector of size 128, 256 and 1024 respectively. Element has the largest dimension of feature vector, and those of Set and Event are reduced gradually since more bits are required to encode the details of low-level sub-actions. This is to involve more temporal dynamic features extracted from high frame rate input, and sufficient spatial information from lower frame rate input. The features are then concatenated and full-connected to a linear layer of dimension 1024. Finally, the linear layer is connected to three classifiers for joint prediction of the event, set and element labels. We adopt cross-entropy loss in (1) for the three categories, where c denotes category of event, set or element. s_i^c is the CNN score of class i of category c , t_i is the ground-truth value and only one element t_p is not zero for ground-truth positive class. N^c denotes the number of classes in category c .

$$L^c = - \sum_{i=1}^{N^c} t_i^c \log \left(\frac{e^{s_i^c}}{\sum_{j=1}^{N^c} e^{s_j^c}} \right) = -s_p^c + \log \sum_{j=1}^{N^c} e^{s_j^c} \quad (1)$$

The total loss, L^{total} in (2), for training the multi-task network is a weighted sum of all losses, where the ratio for the weights λ_1 , λ_2 and λ_3 are set as 1, 2 and 4 respectively.

$$L^{total} = \lambda_1 L^{Event} + \lambda_2 L^{Set} + \lambda_3 L^{Element} \quad (2)$$

4. EXPERIMENTS

4.1. Implementation Details

Our framework is implemented using MMAction2 [21] utilizing four GPUs GeForce GTX2080Ti. We adopt ResNet-50 as the backbone for SlowOnly base models. The base models and multi-task network are trained with learning rate 0.01, momentum 0.9, weight decay $1e-4$ and gradient clip 40. Following the original configurations, SlowOnly is trained with learning rate decay at fixed step 90 and 110.

Our experiment is conducted on FineGym dataset [7], specifically Gym99 which contains annotation on 4 events, 15 sets, and 99 elements. At the point where we download the dataset, some of the videos are no longer available.

4.2. Base Models

Recently, there is a study¹ that implements SlowOnly network on FineGym99. With the parameters of $T = 16$ and base $C = 64$, the model produced 79.3% Top-1 and 70.2% mean accuracies, which are still lower than the state-of-the-art performance. We first implement and investigate training of SlowOnly networks individually for Event, Set and Element actions as our base models with different configurations. For Event prediction, we sample 4 frames with 16 frames interval as low frame rate input, i.e., $T = 4$. For Set prediction, we experiment with two configurations, i.e., $T = 8$ with interval 8, and $T = 16$ with interval 4. For Element prediction, we sample 32 frames with interval 2, i.e., $T = 32$. The size of the input image is (224, 224). The first convolution kernel is set as (1, 7, 7) with base $C = 64$. Additionally, we also implement SlowFast network as a base model for Element prediction. The first convolution kernels are set as (1, 7, 7) with $C = 64$ and (5, 7, 7) with $C = 8$ for the slow and fast pathways respectively. SlowFast follows the implementation of SlowOnly, but adopts Cosine Annealing learning rate schedule. Both SlowOnly and SlowFast models are trained for 120 epochs and tested on six clips with center crop. The results for the individual models, evaluated on Top-1, Top-5, and mean accuracies are shown in Table 1.

As observed in Table 1, SlowOnly outperforms SlowFast network in Element prediction, with 8.7% gain in Top-1 accuracy. This is due to the reduced channel dimension in the fast pathway in SlowFast network which could not capture sufficient temporal context to distinguish the fine-grained motion. For Set and Element predictions, utilizing higher number of frame inputs leads to results improvement in Top-1 and mean accuracies. The results in Table 1 are also served as baselines for comparison with our multi-task network.

¹<https://github.com/open-mmlab/mmaaction2/tree/master/configs/recognition/slowonly>

Table 1. Individual model results for Event, Set and Element.

Class	Network	T	base C	Event			Set			Element		
				Top-1	Top-5	Mean	Top-1	Top-5	Mean	Top-1	Top-5	Mean
Event	SlowOnly	4	64	99.28	100	99.22	-	-	-	-	-	-
Set	SlowOnly	8	64	-	-	-	95.49	99.95	95.43	-	-	-
Set	SlowOnly	16	64	-	-	-	97.70	99.83	97.68	-	-	-
Element	SlowOnly ¹	16	64	-	-	-	-	-	-	79.30	-	70.2
Element	SlowOnly	32	64	-	-	-	-	-	-	91.05	97.82	88.40
Element	SlowFast	4 (slow) 32 (fast)	64 (slow) 8 (fast)	-	-	-	-	-	-	82.32	98.15	78.93

Table 2. Multi-task network results for joint recognition.

Combined pathways			Event			Set			Element		
Event model	Set model	Element model	Top-1	Top-5	Mean	Top-1	Top-5	Mean	Top-1	Top-5	Mean
SlowOnly	SlowOnly, $T = 8$	SlowOnly	99.50	100	99.40	98.42	100	98.16	91.58	99.64	87.50
SlowOnly	SlowOnly, $T = 16$	SlowOnly	99.54	100	99.37	98.94	100	98.87	91.80	99.69	88.46
SlowOnly	SlowOnly, $T = 8$	SlowFast	99.81	100	99.63	98.18	100	97.75	82.68	98.51	77.62
SlowOnly	SlowOnly, $T = 16$	SlowFast	99.81	100	99.49	98.78	99.98	98.47	82.95	98.73	77.76

Table 3. Comparison with State-of-the-Art performance on FineGym99 action recognition.

Model	Modality	Event		Set		Element	
		Top-1	Mean	Top-1	Mean	Top-1	Mean
TSN [8]	2Stream	99.86	98.47	97.69	91.97	86.0	76.4
TRNms [9]	2Stream	-	-	-	-	87.8	80.2
TSM [10]	2Stream	-	-	-	-	88.4	81.2
I3D [13]	2Stream	-	-	-	-	75.6	64.4
NL I3D [14]	2Stream	-	-	-	-	75.3	64.3
SlowFast	RGB	99.81	99.63	97.70	97.40	79.16	74.93
Multi-task (Ours)	RGB	99.54	99.37	98.94	98.87	91.80	88.46

4.3. Multi-task Network

In this section, we take one base model from each hierarchy class in Table 1 to form the pathways of our multi-task network. During training, the weights of the base models are frozen, where we utilize the output features to train our joint prediction layers. The network is trained for 60 epochs and tested on a single clip with center crop. We experiment with four different configurations and obtain the results in Table 2.

As observed in Table 2, multi-task networks with fused features perform better than their corresponding base model trained separately for each category in Table 1. This is due to the encoding and learning of complementary feature representation from multi-level spatio-temporal semantics. It is also apparent that multi-task networks with SlowOnly network performs better than SlowFast network in the Element prediction due to the difference in channels dimension.

4.4. Comparison with State-of-the-Art

We compare our multi-task learning results with state-of-the-art results of action recognition on FineGym99 presented in

[7]. Specifically, we compare with TSN [8], TRNms [9], TSM [10], I3D [13] and NL I3D [14] with 2-stream modalities. The results are presented in Table 3. The Event, Set and Element results from TSN are trained separately, while our results are joint prediction from our multi-task network. Additionally, we implemented SlowFast with multi-task learning by modifying its last fully-connected layer to three classifiers and adopt the same weighted loss function for training.

Our network outperforms all the action models in Element prediction, achieving 91.80% in Top-1 accuracy and 88.46% in mean accuracy, with improvement of 3.40% and 7.26% respectively over state-of-the-art performance. For Set prediction, our joint learning results performs better than TSN, with 1.25% and 6.90% increment in Top-1 and mean accuracies respectively. As compared to the baseline SlowFast multi-task network, our network with encoding and fusion layers learns better joint representation that leads to improved performance in Element and Set predictions. For Event prediction, our performance is comparable to TSN and SlowFast results.

5. CONCLUSION

In this paper, we presented a multi-task network for effective representation and joint learning of fine-grained human actions on the three-level hierarchy proposed for FineGym. Our experiment results show the effectiveness of exploiting and leveraging on the semantic and temporal context of parallel pathways with varying input sampling. We added integration layers to allow joint encoding and learning of complementary spatio-temporal features of hierarchical action categories. Our multi-task network outperforms networks with single task prediction. For future work, we will look into networks with end-to-end training to jointly learn and refine the hierarchical action representations for multi-task prediction.

6. REFERENCES

- [1] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles, “Activitynet: A large-scale video benchmark for human activity understanding,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 961–970.
- [2] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al., “The kinetics human action video dataset,” *arXiv preprint arXiv:1705.06950*, 2017.
- [3] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, “Large-scale video classification with convolutional neural networks,” in *2014 IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 1725–1732.
- [4] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzyńska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, Florian Hoppe, Christian Thureau, Ingo Bax, and Roland Memisevic, “The ”something something” video database for learning and evaluating visual common sense,” 2017.
- [5] Jihoon Chung, Cheng hsin Wu, Hsuan ru Yang, Yu-Wing Tai, and Chi-Keung Tang, “Haa500: Human-centric atomic action-dataset with curated videos,” 2020.
- [6] Manisha Verma, Sudhakar Kumawat, Yuta Nakashima, and Shanmuganathan Raman, “Yoga-82: a new dataset for fine-grained classification of human poses,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 1038–1039.
- [7] Bo Dai Dian Shao, Yue Zhao and Dahua Lin, “Finegym: A hierarchical video dataset for fine-grained action understanding,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2020.
- [8] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool, “Temporal segment networks for action recognition in videos,” 2017.
- [9] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba, “Temporal relational reasoning in videos,” 2018.
- [10] Ji Lin, Chuang Gan, and Song Han, “Tsm: Temporal shift module for efficient video understanding,” 2019.
- [11] Rohit Girdhar, Deva Ramanan, Abhinav Gupta, Josef Sivic, and Bryan Russell, “Actionvlad: Learning spatio-temporal aggregation for action classification,” 2017.
- [12] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri, “Learning spatiotemporal features with 3d convolutional networks,” 2015.
- [13] Joao Carreira and Andrew Zisserman, “Quo vadis, action recognition? a new model and the kinetics dataset,” 2017.
- [14] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He, “Non-local neural networks,” 2018.
- [15] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He, “Slowfast networks for video recognition,” in *Proceedings of the IEEE international conference on computer vision*, 2019, pp. 6202–6211.
- [16] Lin Sun, Kui Jia, Dit-Yan Yeung, and Bertram E Shi, “Human action recognition using factorized spatiotemporal convolutional networks,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 4597–4605.
- [17] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman, “Convolutional two-stream network fusion for video action recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1933–1941.
- [18] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri, “A closer look at spatiotemporal convolutions for action recognition,” in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018, pp. 6450–6459.
- [19] Mei Chee Leong, Dilip K Prasad, Yong Tsui Lee, and Feng Lin, “Semi-cnn architecture for effective spatiotemporal learning in action recognition,” *Applied Sciences*, vol. 10, no. 2, pp. 557, 2020.
- [20] Christoph Feichtenhofer, “X3d: Expanding architectures for efficient video recognition,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 203–213.
- [21] MMAction2 Contributors, “Openmmlab’s next generation video understanding toolbox and benchmark,” <https://github.com/open-mmlab/mmdetection>, 2020.