

# Relation Modeling in Spatio-Temporal Action Localization

Yutong Feng<sup>1,2</sup> Jianwen Jiang<sup>2\*</sup> Ziyuan Huang<sup>2</sup> Zhiwu Qing<sup>2</sup>  
Xiang Wang<sup>2</sup> Shiwei Zhang<sup>2</sup> Mingqian Tang<sup>2</sup> Yue Gao<sup>1\*</sup>  
<sup>1</sup>THUICS, BNRist, School of Software, Tsinghua University  
<sup>2</sup>Alibaba Group

fytl19@mails.tsinghua.edu.cn , gaoyue@tsinghua.edu.cn  
{jianwen.jjw, pishi.hzy, qingzhiwu.qzw}@alibaba-inc.com  
{xiaolao.wx, zhangjin.zsw, mingqian.tmq}@alibaba-inc.com

## Abstract

*This paper presents our solution to the AVA-Kinetics Crossover Challenge of ActivityNet workshop at CVPR 2021. Our solution utilizes multiple types of relation modeling methods for spatio-temporal action detection and adopts a training strategy to integrate multiple relation modeling in end-to-end training over the two large-scale video datasets. Learning with memory bank and finetuning for long-tailed distribution are also investigated to further improve the performance. In this paper, we detail the implementations of our solution and provide experiments results and corresponding discussions. We finally achieve 40.67 mAP on the test set of AVA-Kinetics and obtain the 1st place of the AVA-Kinetics 2021 challenge.*

## 1. Introduction

Spatio-temporal action localization aims to localize atomic actions of people in videos with 3D bounding boxes, which has attract large efforts in recent years [5, 25, 20, 16, 4, 9]. Generally, there are two main factors showing fundamental influence on the performance of this task, *i.e.* video backbones and relation modeling. The design of video networks has been widely studied [4, 19, 3] and greatly enhance the performance of downstream tasks. Besides, pretraining such networks on large-scale networks is also demonstrated to be effective [20, 16], *e.g.* pretrain on Kinetics700 [2]. And there are also multiple ways to perform the pretraining, such as supervised pre-training [21, 4, 1] as used in [18, 17, 24], and unsupervised ones [8, 6]. For relation modeling, different approaches has been studied

in the fields of computer vision [23, 1], social networks [12, 10] and nature language processing [22]. Specifically, transformed-based relation modeling has been proved for improving the spatio-temporal localization task [25, 20, 16].

In this paper, we investigate multiple types of relation modeling methods for spatio-temporal action localization. Inspired by previous works, an off-the-shelf person detector are employed first to generate all human bounding boxes in the videos. Then we adopt a backbone model to extract visual features and build a relation module upon the feature maps of each person via roi align [7]. After relation module, an action predictor is used to generate score for each action category. The whole pipeline of our solution is shown in 1. In following sections, we first detail the implementation of our method. Then the experimental results and corresponding discussions are provided.

## 2. Method

In this section, we present our approach for relation modeling in spatio-temporal action localization. Firstly, we introduce our overall pipeline for this task. Then relation modeling module is presented with transformer-based architectures to capture the relations among persons in the spatial and temporal dimensions. Furthermore, we adapt memory bank for storing person features along the temporal context of video clips to model long-range relations. Different strategies of online or offline maintaining memory bank on the AVA-Kinetics Crossover are studied. Finally, we investigate learning approaches for the lone-tailed category distribution in the AVA-Kinetics dataset [5, 13].

### 2.1. Overall Pipeline

Our designed pipeline is illustrated in Figure 1. Given the input video clip, the key frame of this clip is extracted and fed into a 2D person detector to generate bounding boxes of persons inside this clip. The whole video clip is

\*Corresponding authors

This work was done when Yutong Feng (Tsinghua University), Ziyuan Huang (National University of Singapore), Zhiwu Qing and Xiang Wang (Huazhong University of Science and Technology) were interns at Alibaba Group.

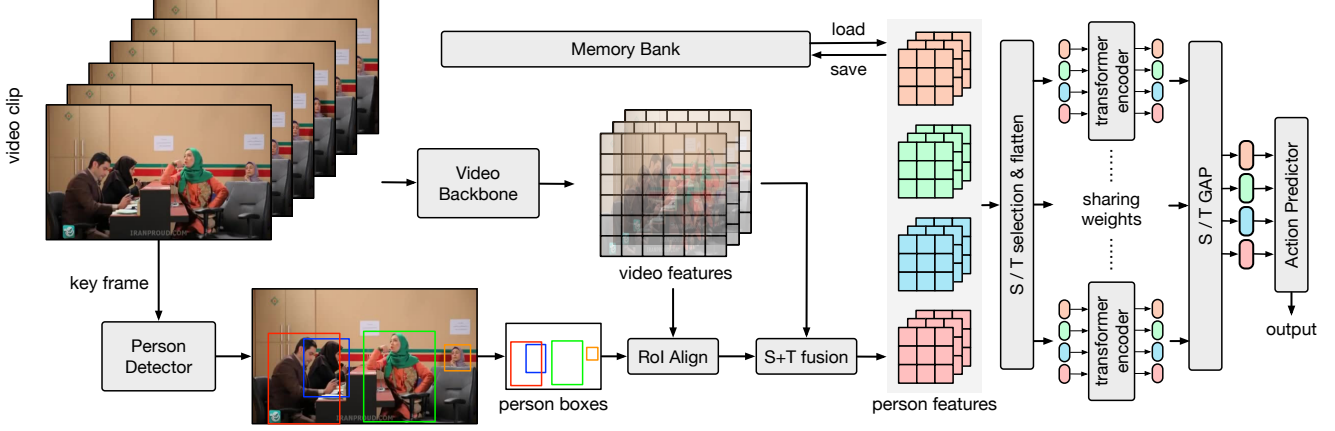


Figure 1. **Pipeline of our approach for spatio-temporal action localization.** ‘S’ and ‘T’ indicate operation along the spatial and temporal dimensions, respectively.

sampled into frames in specified interval and encoded with an video backbone, *e.g.* SlowFast [4] and CSN [21], to output a 3D video feature map. Then the 2D person boxes are inflated along the temporal dimension and used for extracting person features from the feature map with 3D RoI-Align [7]. The pooled person features are further fused together with the video feature map via channel-wise concatenation and convolution layers. To model the hidden relations among persons inside the same video clip for improving the effectiveness of action prediction, we fed the person features into our relation modeling module with transformer-based blocks. To specify the spatial and temporal relations, we select features along the same spatial or temporal dimension from different persons. The selected features are flattened to a sequence of tokens and fed into a transformer encoder block to model their relations via attention mechanism. Finally, the output tokens of all block in spatial or temporal dimensions are globally averaged and fed into full-connected layers to predict the action classes for each detected person.

## 2.2. Person Relation Modeling

The input person features to relation modeling module are 3D feature maps denoted as  $P^i \in \mathbb{R}^{T \times H \times W}$ , where  $i$  is the person index. Such person features are firstly transferred into sequential tokens as the input of transformer encoder block. To effectively model the relations along spatial and temporal dimension while maintain low computation cost, two types of relation head, *i.e.* S-only and T-only, are proposed to extract relations separately on each dimension. For S-only head, we generate  $HW$  sequences of input tokens. The sequence corresponding to spatial position  $(h, w)$  is  $\{agg_{t=1}^T(P_{t,h,w}^i) \mid i \in I\}$ , where  $agg(\cdot)$  indicates an aggregation function, *e.g.* average pooling or max pooling, and  $I$  is the index set of persons. Similarly, T-only head generate  $T$  sequences of input tokens,

where the sequence corresponding to temporal position  $t$  is  $\{agg_{h=1}^H agg_{w=1}^W(P_{t,h,w}^i) \mid i \in I\}$ .

Each sequence of tokens are then fed into a transformer encoder block for relation modeling. It is noted that transformer blocks for different spatial or temporal positions share the same weights. The output of transformer blocks are then averaged along different spatial or temporal positions into the final representation of persons.

## 2.3. Memory Bank

Maintaining feature memory bank for storing and utilizing representations along long term context has been demonstrated to be effective strategies for this task [25, 20, 16]. We also adapt the feature bank, which saves our pooled feature features and provides previously stored person features of timestamps within a long-range of current video clip. The loaded features are concatenated and fed into the relation modeling module.

Existing methods use online maintaining strategy of memory bank, which continually updates the stored features of current video during the training stage. However, for the AVA-Kinetics Crossover, such an online strategy is hard to be implemented since there is only one officially annotated clip of each video in Kinetics and the remaining clips will not be reached in the training stage. To address this issue, we design a two-stage training strategy. In the first stage, the non-annotated clips in Kinetics are not considered. We either train only on the AVA dataset or maintain an empty memory bank for the Kinetics dataset. In the second stage, we extract and store features of all clips in Kinetics, freeze the weights of backbone and finetune the relation head and classifier on both AVA and Kinetics. Besides, we also investigate the strategy that training without memory bank in the first stage, and retrain another head with memory bank in the second stage. Comparison of different strategies are

backbone	pretrain	input	head	+M	val mAP@0.5			test mAP@0.5
					AVA	Kinetics	A+K	A+K
SlowFast-152	K700	8x8+32x2	Linear	×	35.29	32.88	36.26	-
SlowFast-152	K700	8x8+32x2	S-only	×	36.16	33.82	36.92	-
SlowFast-152	K700	8x8+32x2	T-only	×	36.34	33.16	36.88	-
SlowFast-152	K700	8x8+32x2	T-only	✓	36.88	33.44	37.35	-
SlowFast-152	K700	16x4+64x1	T-only	✓	37.68	35.54	38.12	-
ir-CSN-152	K400+IG65M	32x2	Linear	×	35.38	32.58	36.20	-
ir-CSN-152	K400+IG65M	32x2	S-only	×	36.07	32.62	36.58	-
ir-CSN-152	K400+IG65M	32x2	T-only	×	36.29	32.81	36.58	-
ir-CSN-152	K400+IG65M	32x2	T-only	✓	36.99	33.44	37.34	-
ir-CSN-152	K700	32x2	Linear	×	35.67	33.20	36.22	-
ir-CSN-152	K700	32x2	S-only	×	36.69	34.12	37.16	-
ir-CSN-152	K700	32x2	T-only	×	36.41	34.80	36.98	-
ir-CSN-152	K700	32x2	T-only	✓	37.32	35.01	37.78	-
ir-CSN-152	K700	64x1	T-only	✓	37.95	35.26	38.43	-
ensemble	-	-	mixed	-	<b>40.52</b>	<b>37.04</b>	<b>40.97</b>	<b>40.67</b>

Table 1. **Main results on AVA-Kinetics v1.0.** All models are tested with 3 scales and horizontal flips. ”+M” indicates training with memory. ”A+K” indicates results on AVA-Kinetics.

shown in experiments.

## 2.4. Long-tailed Learning

There exists obvious long-tailed category distribution in the original AVA dataset, which leads to the challenge of learning those classes with less number of samples. With the join of Kinetics annotated data, the long-tailed distribution still remains a large problem for this task.

To perform a more suitable training, we consider the decoupling strategy from [11]. The training process is decoupled into two stages. Stage one follows the normal training strategy with randomly sampled data. While in stage two, we freeze all the models except the final classifier and train with class-balanced data sampling. Such a strategy helps to improve the performance on small classes.

## 3. Experiments

### 3.1. Experimental Settings

**Person Detector.** We adopt GFocalV2 [14] as the person detector. We first train the model on the subset of person category from COCO Dataset [15]. Then we continue training the model on the AVA-Kinetics dataset from COCO-Person pre-training.

**Backbone.** We adopt CSN152 [21], slowfast101 [4], slowfast152 [4] as the visual feature extractors. We first train the model on the training set of Kinetics700 [2] dataset and then use the weight to initial the backbone part of our

pipeline.

**Heads.** We train with four types of heads to predict actions. The linear head simply use full-connected layers as the baseline head. For relation heads, considering the size of features from S-only head will make it changeable to maintain memory bank, we train S-only without memory bank and T-only with/without memory bank.

**Training and Inference.** During the training process, we concatenate the data list of AVA and Kinetics for mixed learning. We train with SGD optimizer, initial learning rate as 1e-2, batch size as 64, weight decay as 1e-7 and total training iterations as 30k. All the batchnorm layers are frozen in training. We do linear warmup in the first 1.5k iterations, drop the learning rate by 0.66 at iteration 13.5k, 18k, 22.5k and 27k. The input videos are resized with minimum side of 256 and maximum side of 464. Color jitter and box jitter are used for data augmentation. During the inference process, we test with three scales {256, 288 and 320} and horizontal flips.

### 3.2. Main Results

Table 1 shows our main results on AVA-Kinetics. We report results of models using different backbones, pretrained datasets, input formats relation heads and memory banks. Our best single model achieves 37.95 mAP and 35.26 mAP on AVA and Kinetics, respectively, and 38.43 mAP on AVA-Kinetics Crossover. For backbones, ir-CSN-152 achieves better performance compared with SlowFast-152, which

Stage 1			Stage 2				val mAP@0.5	
dataset	+M(A)	+M(K)	dataset	+M(A)	+M(K)	finetune	AVA	Kinetics
A	✓	×	A+K	✓	✓	✓	35.67	28.59
A+K	×	×	A+K	✓	✓	×	35.91	33.35
A+K	✓	×	A+K	✓	✓	✓	<b>36.88</b>	<b>33.44</b>

Table 2. **Comparison of different training strategies for memory bank on AVA-Kinetics.** All models are trained with SlowFast-152 backbone and T-only head. Results are tested with 3 scales and horizontal flips. ”+M(A)” and ”+M(K)” indicate maintaining online memory bank of AVA and Kinetics during the training stage. ”A” indicates training with AVA only, and ”A+K” indicates training jointly with AVA-Kinetics.

may attribute to that ir-CSN-152 is pretrained with smaller downscale of 1/16, while SlowFast-152 is 1/32. For the pretrained dataset, pretraining on Kinetics700 dataset generally improves 1 mAP compared with those pretrained on Kinetics400. We also investigate results with different temporal resolutions, and it is shown that increasing the temporal resolution  $\times 2$  could improve around 0.6 mAP. For S-only and T-only relation heads, they both achieve great improvement compared with linear head and similar results to each other. T-only head with memory bank increase mAP of about 0.8 mAP. Among those models, we select 15 models and ensemble their results with average voting. The ensemble result achieves 40.97 mAP on the validation set of AVA-Kinetics, and 40.67 mAP on the test set. It is noted that even without training on the validation set, we only drop 0.3 mAP compared with the validation set.

### 3.3. Ablation Studies and Discussions

**Influence of memory bank strategies.** We train the three different strategies of maintain memory bank on AVA-Kinetics. As shown in Table 2, only training on AVA in the first stage and finetuning on AVA-Kinetics could achieve satisfying performance on AVA, but the results on Kinetics will be poor since the backbones are not fully trained on this dataset. Training head without memory bank on both AVA and Kinetics in the first stage could improve the result on Kinetics in a large margin. However, this strategy must re-initialize the weights of head in the second stage and may lose some valuable training efforts. The final strategy trains on AVA-Kinetics on two stages and maintain empty bank in the first stage, which achieve the best performance and and used in our final version.

**Influence of decoupled learning.** We compare the results of all classes before and after the decoupled class-balanced finetuning. The classes are ranked by their number of samples, and we list the averaged difference of the top-20 and bottom-20 classes. As shown in Table 3, there is an improvement of 1.05 mAP of those small classes. At the same time, the performance of the top-20 classes almost do not drop. We also report classes with largest changes after decoupled training. It is noted that some of the small classes,

classes	mAP@0.5 diff.
top-20	- 0.16
bottom-20	+ 1.05
all-60	+ 0.14
cut	+ 22.57
write	+ 4.93
text on cellphone	+ 3.66
get up	- 3.93
give/serve to	- 4.12
work on computer	- 8.85

Table 3. **Influence of applying decoupled training for lone-tail learning.** Results of ir-CSN-152+T-only on AVA dataset are used for comparison. Classes are ranked by their numbers of labeled samples. Averaged results of the top-20 and bottom-20 classes, top-3 improved classes and bottom-3 improved classes are listed.

detector	AVA		Kinetics	
	AP@0.5	mAP@0.5	AP@0.5	mAP@0.5
det-v1	94.5	35.72	77.0	32.70
det-v2	95.1	36.16	81.7	33.84
det-v3	95.1*	36.41	82.3	34.80
GT	-	43.53	-	48.49

Table 4. **Influence of person detector on AVA and Kinetics validation sets.** The column of AP@0.5 shows the detection results of different detectors. Results using GT bounding boxes are listed for reference. All the results are trained with ir-CSN-152 backbone and T-only head, and are tested with 3 scales and horizontal flips. ”\*” notes the results of det-v3, which has similar AP@0.5 to det-v2 but achieves better performance when measuring AP@0.5:0.95.

e.g. cut, could get great improvement, while there also exists classes with dropped performance. This may attribute to the overfitting of those classes with increased training times, which could be further improved.

**Influence of human detector.** Here we have also in-



investigate how much the improvement of the detector performance will gain the final action detection performance. Three detection boxes with 77.0%, 81.7%, 82.3% AP on Kinetics subset of AVA-Kinetics, respectively, are applied in our pipeline. As Table 4 shown, when the performance of the detector increases from 77% to 81.7%, the final performance can obtain significant improvement of about 1 mAP on both AVA and Kinetics parts of the dataset. If the GT boxes are applied, the performance of action detector will reach 43.5 mAP and 48.5 mAP respectively on the two subset. Moreover, even the improvements of human detector are similar (95.1 v.s. 95.1 and 81.7 v.s. 82.3 on AVA and Kinetics subset respectively), the final action detector also obtains corresponding gains of performance. The results demonstrate that human detector is one of bottleneck modules in the action detection pipeline.

## Acknowledgement

This work was supported by Alibaba Group through Alibaba Innovative Research Program.

## References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. *arXiv preprint arXiv:2103.15691*, 2021. 1
- [2] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 1, 3
- [3] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. *arXiv preprint arXiv:2104.11227*, 2021. 1
- [4] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 6201–6210. IEEE, 2019. 1, 2, 3
- [5] Chunhui Gu, Chen Sun, David A. Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, Cordelia Schmid, and Jitendra Malik. AVA: A video dataset of spatio-temporally localized atomic visual actions. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 6047–6056. IEEE Computer Society, 2018. 1
- [6] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised co-training for video representation learning. *arXiv preprint arXiv:2010.09709*, 2020. 1
- [7] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 2980–2988. IEEE Computer Society, 2017. 1, 2
- [8] Ziyuan Huang, Shiwei Zhang, Jianwen Jiang, Mingqian Tang, Rong Jin, and Marcelo Ang. Self-supervised motion learning from static images. *arXiv preprint arXiv:2104.00240*, 2021. 1
- [9] Jianwen Jiang, Yu Cao, Lin Song, Shiwei Zhang, Yunkai Li, Ziyao Xu, Qian Wu, Chuang Gan, Chi Zhang, and Gang Yu. Human centric spatio-temporal action localization. In *ActivityNet Workshop on CVPR*, 2018. 1
- [10] Jianwen Jiang, Yuxuan Wei, Yifan Feng, Jingxuan Cao, and Yue Gao. Dynamic hypergraph neural networks. In *IJCAI*, pages 2635–2641, 2019. 1
- [11] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217*, 2019. 3
- [12] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. 1
- [13] Ang Li, Meghana Thotakuri, David A. Ross, João Carreira, Alexander Vostrikov, and Andrew Zisserman. The ava-kinetics localized human actions video dataset. *CoRR*, abs/2005.00214, 2020. 1
- [14] Xiang Li, Wenhai Wang, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss v2: Learning reliable localization quality estimation for dense object detection. *arXiv preprint arXiv:2011.12885*, 2020. 3
- [15] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 3
- [16] Junting Pan, Siyu Chen, Zheng Shou, Jing Shao, and Hongsheng Li. Actor-context-actor relation network for spatio-temporal action localization. *CoRR*, abs/2006.07976, 2020. 1, 2
- [17] Zhiwu Qing, Haisheng Su, Weihao Gan, Dongliang Wang, Wei Wu, Xiang Wang, Yu Qiao, Junjie Yan, Changxin Gao, and Nong Sang. Temporal context aggregation network for temporal action proposal refinement. *arXiv preprint arXiv:2103.13141*, 2021. 1
- [18] Lin Song, Shiwei Zhang, Gang Yu, and Hongbin Sun. Tacnet: Transition-aware context network for spatio-temporal action detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11987–11995, 2019. 1
- [19] Jonathan Stroud, David Ross, Chen Sun, Jia Deng, and Rahul Sukthankar. D3d: Distilled 3d networks for video action recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 625–634, 2020. 1
- [20] Jiajun Tang, Jin Xia, Xinzhi Mu, Bo Pang, and Cewu Lu. Asynchronous interaction aggregation for action detection. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XV*, volume 12360 of *Lecture Notes in Computer Science*, pages 71–87. Springer, 2020. 1, 2

- [21] Du Tran, Heng Wang, Matt Feiszli, and Lorenzo Torresani. Video classification with channel-separated convolutional networks. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 5551–5560. IEEE, 2019. [1](#), [2](#), [3](#)
- [22] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017. [1](#)
- [23] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018. [1](#)
- [24] Xiang Wang, Baiteng Ma, Zhiwu Qing, Yongpeng Sang, Changxin Gao, Shiwei Zhang, and Nong Sang. Cbr-net: Cascade boundary refinement network for action detection: Submission to activitynet challenge 2020 (task 1). *arXiv preprint arXiv:2006.07526*, 2020. [1](#)
- [25] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krähenbühl, and Ross B. Girshick. Long-term feature banks for detailed video understanding. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 284–293. Computer Vision Foundation / IEEE, 2019. [1](#), [2](#)