

TITLE: Video Action Transformer Network
http://openaccess.thecvf.com/content_CVPR_2019/html/Girdhar_Video_Action_Transformer_Network_CVPR_2019_paper.html
AUTHORS: Rohit Girdhar, Joao Carreira, Carl Doersch, Andrew Zisserman
HIGHLIGHT: We introduce the Action Transformer model for recognizing and localizing human actions in video clips.

TITLE: Timeception for Complex Action Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Hussein_Timeception_for_Complex_Action_Recognition_CVPR_2019_paper.html
AUTHORS: Nouredien Hussein, Efstratios Gavves, Arnold W.M. Smeulders
HIGHLIGHT: This paper focuses on the temporal aspect for recognizing human activities in videos; an important visual cue that has long been undervalued.

TITLE: STEP: Spatio-Temporal Progressive Learning for Video Action Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Yang_STEP_Spatio-Temporal_Progressive_Learning_for_Video_Action_Detection_CVPR_2019_paper.html
AUTHORS: Xitong Yang, Xiaodong Yang, Ming-Yu Liu, Fanyi Xiao, Larry S. Davis, Jan Kautz
HIGHLIGHT: In this paper, we propose Spatio-Temporal Progressive (STEP) action detector--a progressive learning framework for spatio-temporal action detection in videos.

TITLE: Relational Action Forecasting
http://openaccess.thecvf.com/content_CVPR_2019/html/Sun_Relational_Action_Forecasting_CVPR_2019_paper.html
AUTHORS: Chen Sun, Abhinav Shrivastava, Carl Vondrick, Rahul Sukthankar, Kevin Murphy, Cordelia Schmid
HIGHLIGHT: This paper focuses on multi-person action forecasting in videos.

TITLE: Long-Term Feature Banks for Detailed Video Understanding
http://openaccess.thecvf.com/content_CVPR_2019/html/Wu_Long-Term_Feature_Banks_for_Detailed_Video_Understanding_CVPR_2019_paper.html
AUTHORS: Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, Ross Girshick
HIGHLIGHT: In this paper, we enable existing video models to do the same.

TITLE: Which Way Are You Going? Imitative Decision Learning for Path Forecasting in Dynamic Scenes
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Which_Way_Are_You_Going_Imitative_Decision_Learning_for_Path_CVPR_2019_paper.html
AUTHORS: Yuke Li
HIGHLIGHT: To this end, we propose a novel Imitative Decision Learning (IDL) approach.

TITLE: What and How Well You Performed? A Multitask Learning Approach to Action Quality Assessment
http://openaccess.thecvf.com/content_CVPR_2019/html/Parmar_What_and_How_Well_You_Performed_A_Multitask_Learning_Approach_CVPR_2019_paper.html
AUTHORS: Paritosh Parmar, Brendan Tran Morris
HIGHLIGHT: In this paper, we propose to learn spatio-temporal features that explain three related tasks - fine-grained action recognition, commentary generation, and estimating the AQA score.

TITLE: MHP-VOS: Multiple Hypotheses Propagation for Video Object Segmentation
http://openaccess.thecvf.com/content_CVPR_2019/html/Xu_MHP-VOS_Multiple_Hypotheses_Propagation_for_Video_Object_Segmentation_CVPR_2019_paper.html
AUTHORS: Shuangjie Xu, Daizong Liu, Linchao Bao, Wei Liu, Pan Zhou
HIGHLIGHT: In this paper, we propose a novel approach to defer the decision making for a target object in each frame, until a global view can be established with the entire video being taken into consideration.

TITLE: Language-Driven Temporal Activity Localization: A Semantic Matching Reinforcement Learning Model
http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_Language-Driven_Temporal_Activity_Localization_A_Semantic_Matching_Reinforcement_Learning_Model_CVPR_2019_paper.html
AUTHORS: Weining Wang, Yan Huang, Liang Wang
HIGHLIGHT: Considering that current methods are generally time-consuming due to the dense frame-processing manner, we propose a recurrent neural network based reinforcement learning model which selectively observes a sequence of frames and associates the given sentence with video content in a matching-based manner.

TITLE: Gaussian Temporal Awareness Networks for Action Localization
http://openaccess.thecvf.com/content_CVPR_2019/html/Long_Gaussian_Temporal_Awareness_Networks_for_Action_Localization_CVPR_2019_paper.html

AUTHORS: Fuchen Long, Ting Yao, Zhaofan Qiu, Xinmei Tian, Jiebo Luo, Tao Mei
HIGHLIGHT: In this paper, we propose to address the problem by introducing Gaussian kernels to dynamically optimize temporal scale of each action proposal.

TITLE: Efficient Video Classification Using Fewer Frames
http://openaccess.thecvf.com/content_CVPR_2019/html/Bhardwaj_Efficient_Video_Classification_Using_Fewer_Frames_CVPR_2019_paper.html

AUTHORS: Shweta Bhardwaj, Mukundhan Srinivasan, Mitesh M. Khapra
HIGHLIGHT: In this work, we focus on building compute-efficient video classification models which process fewer frames and hence have less number of FLOPs.

TITLE: A Perceptual Prediction Framework for Self Supervised Event Segmentation
http://openaccess.thecvf.com/content_CVPR_2019/html/Aakur_A_Perceptual_Prediction_Framework_for_Self_Supervised_Event_Segmentation_CVPR_2019_paper.html

AUTHORS: Sathyanarayanan N. Aakur, Sudeep Sarkar
HIGHLIGHT: In this paper, we tackle the problem of self-supervised temporal segmentation that alleviates the need for any supervision in the form of labels (full supervision) or temporal ordering (weak supervision).

TITLE: COIN: A Large-Scale Dataset for Comprehensive Instructional Video Analysis

http://openaccess.thecvf.com/content_CVPR_2019/html/Tang_COIN_A_Large-Scale_Dataset_for_Comprehensive_Instructional_Video_Analysis_CVPR_2019_paper.html

AUTHORS: Yansong Tang, Dajun Ding, Yongming Rao, Yu Zheng, Danyang Zhang, Lili Zhao, Jiwen Lu, Jie Zhou
HIGHLIGHT: To address these problems, we introduce a large-scale dataset called "COIN" for COmprehensive INstruction video analysis.

TITLE: Recurrent Attentive Zooming for Joint Crowd Counting and Precise Localization

http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Recurrent_Attentive_Zooming_for_Joint_Crowd_Counting_and_Precise_Localization_CVPR_2019_paper.html

AUTHORS: Chenchen Liu, Xinyu Weng, Yadong Mu
HIGHLIGHT: To address this issue, this work proposes a novel framework that simultaneously solving two inherently related tasks - crowd counting and localization.

TITLE: An Attention Enhanced Graph Convolutional LSTM Network for Skeleton-Based Action Recognition

http://openaccess.thecvf.com/content_CVPR_2019/html/Si_An_Attention_Enhanced_Graph_Convolutional_LSTM_Network_for_Skeleton-Based_Action_CVPR_2019_paper.html

AUTHORS: Chenyang Si, Wentao Chen, Wei Wang, Liang Wang, Tieniu Tan
HIGHLIGHT: In this paper, we propose a novel Attention Enhanced Graph Convolutional LSTM Network (AGC-LSTM) for human action recognition from skeleton data.

TITLE: Graph Convolutional Label Noise Cleaner: Train a Plug-And-Play Action Classifier for Anomaly Detection

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhong_Graph_Convolutional_Label_Noise_Cleaner_Train_a_Plug-And-Play_Action_Classifier_CVPR_2019_paper.html

AUTHORS: Jia-Xing Zhong, Nannan Li, Weijie Kong, Shan Liu, Thomas H. Li, Ge Li
HIGHLIGHT: In this paper, we provide a new perspective, i.e., a supervised learning task under noisy labels.

TITLE: MAN: Moment Alignment Network for Natural Language Moment Retrieval via Iterative Graph Adjustment

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_MAN_Moment_Alignment_Network_for_Natural_Language_Moment_Retrieval_via_CVPR_2019_paper.html

AUTHORS: Da Zhang, Xiyang Dai, Xin Wang, Yuan-Fang Wang, Larry S. Davis
HIGHLIGHT: In this paper, we present Moment Alignment Network (MAN), a novel framework that unifies the candidate moment encoding and temporal structural reasoning in a single-shot feed-forward network.

TITLE: Less Is More: Learning Highlight Detection From Video Duration

http://openaccess.thecvf.com/content_CVPR_2019/html/Xiong_Less_Is_More_Learning_Highlight_Detection_From_Video_Duration_CVPR_2019_paper.html

AUTHORS: Bo Xiong, Yannis Kalantidis, Deepti Ghadiyaram, Kristen Grauman
HIGHLIGHT: We propose a scalable unsupervised solution that exploits video duration as an implicit supervision signal.

TITLE: DMC-Net: Generating Discriminative Motion Cues for Fast Compressed Video Action Recognition

http://openaccess.thecvf.com/content_CVPR_2019/html/Shou_DMC-Net_Generating_Discriminative_Motion_Cues_for_Fast_Compressed_Video_Action_CVPR_2019_paper.html

AUTHORS: Zheng Shou, Xudong Lin, Yannis Kalantidis, Laura Sevilla-Lara, Marcus Rohrbach, Shih-Fu Chang, Zhicheng Yan

HIGHLIGHT: To remedy these issues, we propose a lightweight generator network, which reduces noises in motion vectors and captures fine motion details, achieving a more Discriminative Motion Cue (DMC) representation.

TITLE: AdaFrame: Adaptive Frame Selection for Fast Video Recognition

http://openaccess.thecvf.com/content_CVPR_2019/html/Wu_AdaFrame_Adaptive_Frame_Selection_for_Fast_Video_Recognition_CVPR_2019_paper.html

AUTHORS: Zuxuan Wu, Caiming Xiong, Chih-Yao Ma, Richard Socher, Larry S. Davis

HIGHLIGHT: We present AdaFrame, a framework that adaptively selects relevant frames on a per-input basis for fast video recognition.

TITLE: Spatio-Temporal Video Re-Localization by Warp LSTM

http://openaccess.thecvf.com/content_CVPR_2019/html/Feng_Spatio-Temporal_Video_Re-Localization_by_Warp_LSTM_CVPR_2019_paper.html

AUTHORS: Yang Feng, Lin Ma, Wei Liu, Jiebo Luo

HIGHLIGHT: In this paper, we make an answer to the question of when and where by formulating a new task, namely spatio-temporal video re-localization.

TITLE: Completeness Modeling and Context Separation for Weakly Supervised Temporal Action Localization

http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Completeness_Modeling_and_Context_Separation_for_Weakly_Supervised_Temporal_Action_CVPR_2019_paper.html

AUTHORS: Daochang Liu, Tingting Jiang, Yizhou Wang

HIGHLIGHT: Specifically, to model the completeness of actions, we propose a multi-branch neural network in which branches are enforced to discover distinctive action parts.

TITLE: Cross-Task Weakly Supervised Learning From Instructional Videos

http://openaccess.thecvf.com/content_CVPR_2019/html/Zhukov_Cross-Task_Weakly_Supervised_Learning_From_Instructional_Videos_CVPR_2019_paper.html

AUTHORS: Dimitri Zhukov, Jean-Baptiste Alayrac, Ramazan Gokberk Cimbis, David Fouhey, Ivan Laptev, Josef Sivic

HIGHLIGHT: In this paper we investigate learning visual models for the steps of ordinary tasks using weak supervision via instructional narrations and an ordered list of steps instead of strong supervision via temporal annotations.

TITLE: D3TW: Discriminative Differentiable Dynamic Time Warping for Weakly Supervised Action Alignment and Segmentation

http://openaccess.thecvf.com/content_CVPR_2019/html/Chang_D3TW_Discriminative_Differentiable_Dynamic_Time_Warping_for_Weakly_Supervised_Action_CVPR_2019_paper.html

AUTHORS: Chien-Yi Chang, De-An Huang, Yanan Sui, Li Fei-Fei, Juan Carlos Niebles

HIGHLIGHT: We propose Discriminative Differentiable Dynamic Time Warping (D3TW), the first discriminative model using weak ordering supervision.

TITLE: Progressive Teacher-Student Learning for Early Action Prediction

http://openaccess.thecvf.com/content_CVPR_2019/html/Wang_Progressive_Teacher-Student_Learning_for_Early_Action_Prediction_CVPR_2019_paper.html

AUTHORS: Xionghui Wang, Jian-Fang Hu, Jian-Huang Lai, Jianguo Zhang, Wei-Shi Zheng

HIGHLIGHT: In this paper, we aim at improving early action prediction by proposing a novel teacher-student learning framework.

TITLE: Social Relation Recognition From Videos via Multi-Scale Spatial-Temporal Reasoning

http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Social_Relation_Recognition_From_Videos_via_Multi-Scale_Spatial-Temporal_Reasoning_CVPR_2019_paper.html

AUTHORS: Xinchun Liu, Wu Liu, Meng Zhang, Jingwen Chen, Lianli Gao, Chenggang Yan, Tao Mei

HIGHLIGHT: To overcome these challenges, we propose a Multi-scale Spatial-Temporal Reasoning (MSTR) framework to recognize social relations from videos.

TITLE: MS-TCN: Multi-Stage Temporal Convolutional Network for Action Segmentation

http://openaccess.thecvf.com/content_CVPR_2019/html/Abu_Farha_MS-TCN_Multi-Stage_Temporal_Convolutional_Network_for_Action_Segmentation_CVPR_2019_paper.html

AUTHORS: Yazan Abu Farha, Jurgen Gall

HIGHLIGHT: In this paper, we introduce a multi-stage architecture for the temporal action segmentation task.

TITLE: Transferable Interactiveness Knowledge for Human-Object Interaction Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Transferable_Interactiveness_Knowledge_for_Human-Object_Interaction_Detection_CVPR_2019_paper.html
AUTHORS: Yong-Lu Li, Siyuan Zhou, Xijie Huang, Liang Xu, Ze Ma, Hao-Shu Fang, Yanfeng Wang, Cewu Lu
HIGHLIGHT: In this paper, we explore Interactiveness Knowledge which indicates whether human and object interact with each other or not.

TITLE: Actional-Structural Graph Convolutional Networks for Skeleton-Based Action Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Actional-Structural_Graph_Convolutional_Networks_for_Skeleton-Based_Action_Recognition_CVPR_2019_paper.html
AUTHORS: Maosen Li, Siheng Chen, Xu Chen, Ya Zhang, Yanfeng Wang, Qi Tian
HIGHLIGHT: To capture richer dependencies, we introduce an encoder-decoder structure, called A-link inference module, to capture action-specific latent dependencies, i.e. actional links, directly from actions.

TITLE: Multi-Granularity Generator for Temporal Action Proposal
http://openaccess.thecvf.com/content_CVPR_2019/html/Liu_Multi-Granularity_Generator_for_Temporal_Action_Proposal_CVPR_2019_paper.html
AUTHORS: Yuan Liu, Lin Ma, Yifeng Zhang, Wei Liu, Shih-Fu Chang
HIGHLIGHT: In this paper, we propose a multi-granularity generator (MGG) to perform the temporal action proposal from different granularity perspectives, relying on the video visual features equipped with the position embedding information.

TITLE: Peeking Into the Future: Predicting Future Person Activities and Locations in Videos
http://openaccess.thecvf.com/content_CVPR_2019/html/Liang_Peeking_Into_the_Future_Predicting_Future_Person_Activities_and_Locations_CVPR_2019_paper.html
AUTHORS: Junwei Liang, Lu Jiang, Juan Carlos Niebles, Alexander G. Hauptmann, Li Fei-Fei
HIGHLIGHT: We propose an end-to-end, multi-task learning system utilizing rich visual features about human behavioral information and interaction with their surroundings.

TITLE: Re-Identification With Consistent Attentive Siamese Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Zheng_Re-Identification_With_Consistent_Attentive_Siamese_Networks_CVPR_2019_paper.html
AUTHORS: Meng Zheng, Srikrishna Karanam, Ziyang Wu, Richard J. Radke
HIGHLIGHT: We propose a new deep architecture for person re-identification (re-id).

TITLE: Object-Centric Auto-Encoders and Dummy Anomalies for Abnormal Event Detection in Video
http://openaccess.thecvf.com/content_CVPR_2019/html/Ionescu_Object-Centric_Auto-Encoders_and_Dummy_Anomalies_for_Abnormal_Event_Detection_in_CVPR_2019_paper.html
AUTHORS: Radu Tudor Ionescu, Fahad Shahbaz Khan, Mariana-Iuliana Georgescu, Ling Shao
HIGHLIGHT: In this work, we formalize abnormal event detection as a one-versus-rest binary classification problem.

TITLE: DDLSTM: Dual-Domain LSTM for Cross-Dataset Action Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Perrett_DDLSTM_Dual-Domain_LSTM_for_Cross-Dataset_Action_Recognition_CVPR_2019_paper.html
AUTHORS: Toby Perrett, Dima Damen
HIGHLIGHT: In this paper we introduce Dual-Domain LSTM (DDLSTM), an architecture that is able to learn temporal dependencies from two domains concurrently.

TITLE: The Pros and Cons: Rank-Aware Temporal Attention for Skill Determination in Long Videos
http://openaccess.thecvf.com/content_CVPR_2019/html/Doughty_The_Pros_and_Cons_Rank-Aware_Temporal_Attention_for_Skill_Determination_CVPR_2019_paper.html
AUTHORS: Hazel Doughty, Walterio Mayol-Cuevas, Dima Damen
HIGHLIGHT: We present a new model to determine relative skill from long videos, through learnable temporal attention modules.

TITLE: Collaborative Spatiotemporal Feature Learning for Video Action Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Collaborative_Spatiotemporal_Feature_Learning_for_Video_Action_Recognition_CVPR_2019_paper.html
AUTHORS: Chao Li, Qiaoyong Zhong, Di Xie, Shiliang Pu
HIGHLIGHT: In this paper, we propose a novel neural operation which encodes spatiotemporal features collaboratively by imposing a weight-sharing constraint on the learnable parameters.

TITLE: MARS: Motion-Augmented RGB Stream for Action Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Crasto_MARS_Motion-Augmented_RGB_Stream_for_Action_Recognition_CVPR_2019_paper.html
AUTHORS: Nieves Crasto, Philippe Weinzaepfel, Karteek Alahari, Cordelia Schmid
HIGHLIGHT: In this paper, we introduce two learning approaches to train a standard 3D CNN, operating on RGB frames, that mimics the motion stream, and as a result avoids flow computation at test time.

TITLE: Convolutional Relational Machine for Group Activity Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Azar_Convolutional_Relational_Machine_for_Group_Activity_Recognition_CVPR_2019_paper.html
AUTHORS: Sina Mokhtarzadeh Azar, Mina Ghadimi Atigh, Ahmad Nickabadi, Alexandre Alahi
HIGHLIGHT: We present an end-to-end deep Convolutional Neural Network called Convolutional Relational Machine (CRM) for recognizing group activities that utilizes the information in spatial relations between individual persons in image or video.

TITLE: Video Summarization by Learning From Unpaired Data
http://openaccess.thecvf.com/content_CVPR_2019/html/Rochan_Video_Summarization_by_Learning_From_Unpaired_Data_CVPR_2019_paper.html
AUTHORS: Mrigank Rochan, Yang Wang
HIGHLIGHT: We present an approach that learns to generate optimal video summaries using a set of raw videos (V) and a set of summary videos (S), where there exists no correspondence between V and S.

TITLE: Skeleton-Based Action Recognition With Directed Graph Neural Networks
http://openaccess.thecvf.com/content_CVPR_2019/html/Shi_Skeleton-Based_Action_Recognition_With_Directed_Graph_Neural_Networks_CVPR_2019_paper.html
AUTHORS: Lei Shi, Yifan Zhang, Jian Cheng, Hanqing Lu
HIGHLIGHT: In this work, we represent the skeleton data as a directed acyclic graph based on the kinematic dependency between the joints and bones in the natural human body.

TITLE: PA3D: Pose-Action 3D Machine for Video Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Yan_PA3D_Pose-Action_3D_Machine_for_Video_Recognition_CVPR_2019_paper.html
AUTHORS: An Yan, Yali Wang, Zhifeng Li, Yu Qiao
HIGHLIGHT: To fill this gap, we propose a concise Pose-Action 3D Machine (PA3D), which can effectively encode multiple pose modalities within a unified 3D framework, and consequently learn spatio-temporal pose representations for action recognition.

TITLE: Deep Dual Relation Modeling for Egocentric Interaction Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Li_Deep_Dual_Relation_Modeling_for_Egocentric_Interaction_Recognition_CVPR_2019_paper.html
AUTHORS: Haoxin Li, Yijun Cai, Wei-Shi Zheng
HIGHLIGHT: To exploit the strong relations for egocentric interaction recognition, we introduce a dual relation modeling framework which learns to model the relations between the camera wearer and the interactor based on the individual action representations of the two persons.

TITLE: Action Recognition From Single Timestamp Supervision in Untrimmed Videos
http://openaccess.thecvf.com/content_CVPR_2019/html/Moltisanti_Action_Recognition_From_Single_Timestamp_Supervision_in_Untrimmed_Videos_CVPR_2019_paper.html
AUTHORS: Davide Moltisanti, Sanja Fidler, Dima Damen
HIGHLIGHT: We propose a method that is supervised by single timestamps located around each action instance, in untrimmed videos.

TITLE: Time-Conditioned Action Anticipation in One Shot
http://openaccess.thecvf.com/content_CVPR_2019/html/Ke_Time-Conditioned_Action_Anticipation_in_One_Shot_CVPR_2019_paper.html
AUTHORS: Qiuhong Ke, Mario Fritz, Bernt Schiele
HIGHLIGHT: In this paper, we propose a novel time-conditioned method for efficient and effective long-term action anticipation.

TITLE: Dance With Flow: Two-In-One Stream Action Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhao_Dance_With_Flow_Two-In-One_Stream_Action_Detection_CVPR_2019_paper.html
AUTHORS: Jiaojiao Zhao, Cees G. M. Snoek
HIGHLIGHT: The goal of this paper is to detect the spatio-temporal extent of an action.

TITLE: Representation Flow for Action Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Piergiovanni_Representation_Flow_for_Action_Recognition_CVPR_2019_paper.html
AUTHORS: AJ Piergiovanni, Michael S. Ryoo
HIGHLIGHT: In this paper, we propose a convolutional layer inspired by optical flow algorithms to learn motion representations.

TITLE: LSTA: Long Short-Term Attention for Egocentric Action Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Sudhakaran_LSTA_Long_Short-Term_Attention_for_Egocentric_Action_Recognition_CVPR_2019_paper.html
AUTHORS: Swathikiran Sudhakaran, Sergio Escalera, Oswald Lanz
HIGHLIGHT: In this paper we propose LSTA as a mechanism to focus on features from spatial relevant parts while attention is being tracked smoothly across the video sequence.

TITLE: Learning Actor Relation Graphs for Group Activity Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Wu_Learning_Actor_Relation_Graphs_for_Group_Activity_Recognition_CVPR_2019_paper.html
AUTHORS: Jianchao Wu, Limin Wang, Li Wang, Jie Guo, Gangshan Wu
HIGHLIGHT: This paper aims at learning discriminative relation between actors efficiently using deep models.

TITLE: A Structured Model for Action Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_A_Structured_Model_for_Action_Detection_CVPR_2019_paper.html
AUTHORS: Yubo Zhang, Pavel Tokmakov, Martial Hebert, Cordelia Schmid
HIGHLIGHT: To address this limitation, we propose to incorporate domain knowledge into the structure of the model, simplifying optimization.

TITLE: Out-Of-Distribution Detection for Generalized Zero-Shot Action Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Mandal_Out-Of-Distribution_Detection_for_Generalized_Zero-Shot_Action_Recognition_CVPR_2019_paper.html
AUTHORS: Devraj Mandal, Sanath Narayan, Sai Kumar Dwivedi, Vikram Gupta, Shuaib Ahmed, Fahad Shahbaz Khan, Ling Shao
HIGHLIGHT: In this paper, we set out to tackle this issue by arguing for a separate treatment of seen and unseen action categories in generalized zero-shot action recognition.

TITLE: TACNet: Transition-Aware Context Network for Spatio-Temporal Action Detection
http://openaccess.thecvf.com/content_CVPR_2019/html/Song_TACNet_Transition-Aware_Context_Network_for_Spatio-Temporal_Action_Detection_CVPR_2019_paper.html
AUTHORS: Lin Song, Shiwei Zhang, Gang Yu, Hongbin Sun
HIGHLIGHT: In this paper, we define these ambiguous samples as "transitional states", and propose a Transition-Aware Context Network (TACNet) to distinguish transitional states.

TITLE: Learning Regularity in Skeleton Trajectories for Anomaly Detection in Videos
http://openaccess.thecvf.com/content_CVPR_2019/html/Morais_Learning_Regularity_in_Skeleton_Trajectories_for_Anomaly_Detection_in_Videos_CVPR_2019_paper.html
AUTHORS: Romero Morais, Vuong Le, Truyen Tran, Budhaditya Saha, Moussa Mansour, Svetha Venkatesh
HIGHLIGHT: We propose a new method to model the normal patterns of human movements in surveillance video for anomaly detection using dynamic skeleton features.

TITLE: Local Temporal Bilinear Pooling for Fine-Grained Action Parsing
http://openaccess.thecvf.com/content_CVPR_2019/html/Zhang_Local_Temporal_Bilinear_Pooling_for_Fine-Grained_Action_Parsing_CVPR_2019_paper.html
AUTHORS: Yan Zhang, Siyu Tang, Krikamol Muandet, Christian Jarvers, Heiko Neumann
HIGHLIGHT: In this paper we propose a novel bilinear pooling operation, which is used in intermediate layers of a temporal convolutional encoder-decoder net.

TITLE: Improving Action Localization by Progressive Cross-Stream Cooperation
http://openaccess.thecvf.com/content_CVPR_2019/html/Su_Improving_Action_Localization_by_Progressive_Cross-Stream_Cooperation_CVPR_2019_paper.html
AUTHORS: Rui Su, Wanli Ouyang, Luping Zhou, Dong Xu

HIGHLIGHT: In this work, we propose a new Progressive Cross-stream Cooperation (PCSC) framework to iterative improve action localization results and generate better bounding boxes for one stream (i.e., Flow/RGB) by leveraging both region proposals and features from another stream (i.e., RGB/Flow) in an iterative fashion.

TITLE: Two-Stream Adaptive Graph Convolutional Networks for Skeleton-Based Action Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Shi_Two-Stream_Adaptive_Graph_Convolutional_Networks_for_Skeleton-Based_Action_Recognition_CVPR_2019_paper.html
AUTHORS: Lei Shi, Yifan Zhang, Jian Cheng, Hanqing Lu
HIGHLIGHT: In this work, we propose a novel two-stream adaptive graph convolutional network (2s-AGCN) for skeleton-based action recognition.

TITLE: A Neural Network Based on SPD Manifold Learning for Skeleton-Based Hand Gesture Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Nguyen_A_Neural_Network_Based_on_SPD_Manifold_Learning_for_Skeleton-Based_CVPR_2019_paper.html
AUTHORS: Xuan Son Nguyen, Luc Brun, Olivier Lezoray, Sebastien Bougleux
HIGHLIGHT: This paper proposes a new neural network based on SPD manifold learning for skeleton-based hand gesture recognition.

TITLE: Large-Scale Weakly-Supervised Pre-Training for Video Action Recognition
http://openaccess.thecvf.com/content_CVPR_2019/html/Ghadiyaram_Large-Scale_Weakly-Supervised_Pre-Training_for_Video_Action_Recognition_CVPR_2019_paper.html
AUTHORS: Deepti Ghadiyaram, Du Tran, Dhruv Mahajan
HIGHLIGHT: This paper presents an in-depth study of using large volumes of web videos for pre-training video models for the task of action recognition.

TITLE: Learning Spatio-Temporal Representation With Local and Global Diffusion
http://openaccess.thecvf.com/content_CVPR_2019/html/Qiu_Learning_Spatio-Temporal_Representation_With_Local_and_Global_Diffusion_CVPR_2019_paper.html
AUTHORS: Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, Xinmei Tian, Tao Mei
HIGHLIGHT: In this paper, we present a novel framework to boost the spatio-temporal representation learning by Local and Global Diffusion (LGD).

TITLE: Unsupervised Learning of Action Classes With Continuous Temporal Embedding
http://openaccess.thecvf.com/content_CVPR_2019/html/Kukleva_Unsupervised_Learning_of_Action_Classes_With_Continuous_Temporal_Embedding_CVPR_2019_paper.html
AUTHORS: Anna Kukleva, Hilde Kuehne, Fadime Sener, Jurgen Gall
HIGHLIGHT: To address this issue, we propose an unsupervised approach for learning action classes from untrimmed video sequences.

TITLE: Double Nuclear Norm Based Low Rank Representation on Grassmann Manifolds for Clustering
http://openaccess.thecvf.com/content_CVPR_2019/html/Piao_Double_Nuclear_Norm_Based_Low_Rank_Representation_on_Grassmann_Manifolds_CVPR_2019_paper.html
AUTHORS: Xinglin Piao, Yongli Hu, Junbin Gao, Yanfeng Sun, Baocai Yin
HIGHLIGHT: In this paper, we propose a new low rank model for high-dimension data clustering task on Grassmann manifold based on the Double Nuclear norm which is used to better approximate the rank minimization of matrix.

TITLE: 2.5D Visual Sound
http://openaccess.thecvf.com/content_CVPR_2019/html/Gao_2.5D_Visual_Sound_CVPR_2019_paper.html
AUTHORS: Ruohan Gao, Kristen Grauman
HIGHLIGHT: We propose to convert common monaural audio into binaural audio by leveraging video.