

# Technical Report: Disentangled Action Parsing Networks for Accurate Part-level Action Parsing

Xuanhan Wang<sup>1</sup>   Xiaojia Chen<sup>1</sup>   Lianli Gao<sup>1</sup>   Lechao Cheng<sup>2</sup>   Jingkuan Song<sup>1</sup>

<sup>1</sup> Center for Future Media, University of Electronic Science and Technology of China, Chengdu, China

<sup>2</sup>Zhejiang Lab, Hangzhou, China

wxuanhan@hotmail.com, josonchan1998@163.com, lianli.gao@uestc.edu.cn

chenglc@zhejianglab.com, jingkuan.song@gmail.com

## Abstract

*Part-level Action Parsing aims at part state parsing for boosting action recognition in videos. Despite of dramatic progresses in the area of video classification research, a severe problem faced by the community is that the detailed understanding of human actions is ignored. Our motivation is that parsing human actions needs to build models that focus on the specific problem.*

*We present a simple yet effective approach, named disentangled action parsing (DAP). Specifically, we divided the part-level action parsing into three stages: 1) person detection, where a person detector is adopted to detect all persons from videos as well as performs instance-level action recognition; 2) Part parsing, where a part-parsing model is proposed to recognize human parts from detected person images; and 3) Action parsing, where a multi-modal action parsing network is used to parse action category conditioning on all detection results that are obtained from previous stages. With these three major models applied, our approach of DAP records a global mean of 0.605 score in 2021 Kinetics-TPS Challenge.*

## 1. Introduction

The Part-level Action Parsing (PAP) task was firstly proposed as a part of the DeeperAction Challenge at ICCV 2021, which requires models to recognize a human action by compositional learning of body part state in videos. In this work, we present the disentangled action parsing (DAP), a modular framework for part-level action parsing research, which formed the basis for the submission to the Kinetics-TPS Challenge from **CFM-HAG** team.

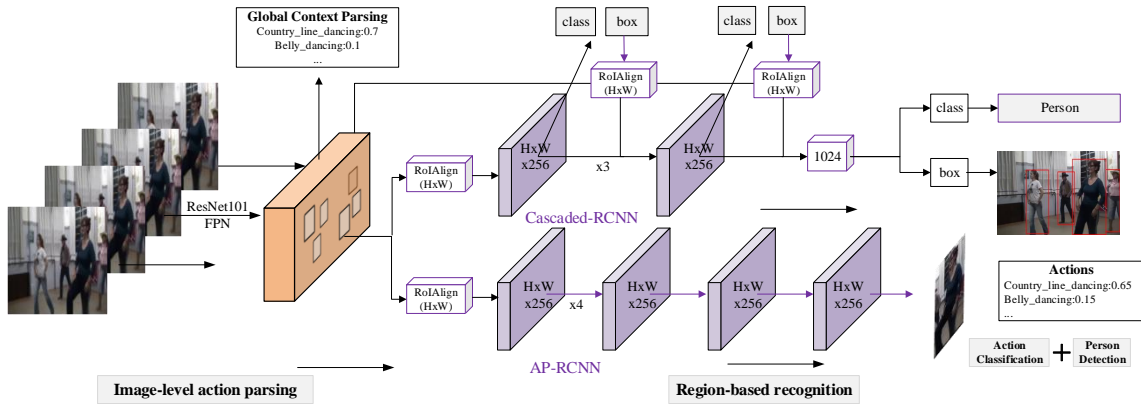
The motivation behind the proposed method comes from the following observations: First, the process of part-level action parsing (PAP) is interpreted as a three-stage pipeline,

which detects all person instances from all video frames, then solves the single-person part parsing problem, and finally predicts the action category conditioned on results that are obtained from previous stages. Secondly, the top-down strategy has become the dominant solution to various human-centric tasks, such as person detection [6, 4], key-point detection [9, 5, 11] and human parsing [1, 13, 10]. Generally, it divides a task into a sequence of sub-problems (e.g., detect-then-estimate), and orderly solves them by designing problem-specific neural networks (e.g., Faster-RCNN for detection and HRNet [9] for estimation). Third, we find out that bottleneck of the PAP lies in the “Part parsing” stage for boosting the final performance. In particular, we conduct a diagnostic experiment by investigating a simple approach, where an extended FPN detector [6] is adopted for instance-level recognition and a TPN [12] model is used for video classification. The experimental results are summarized in Tab. 1. From the results, we observe that the performance is dramatically improved when replacing the predictions of part with their ground truth. According to these observations, we believe that the precise human part-level parsing plays an important role for action parsing. In this technical report, we aim to study one problem: *how to design an effective pipeline for supporting distinct sub-tasks (e.g., person detection, part parsing and action parsing), in particular improving the performance of **part parsing**.*


## 2. Approach

The overall pipeline of the proposed method is shown in Fig. 1. We follow the design rule of “divide-and-conquer” and build our PAP model by designing three task-specific neural networks. Next, we orderly present the proposed models for person detection, part parsing and action parsing.

### Stage 1: Person Deteccion



### Stage 2: Part Parsing



### Stage 3: Action parsing




Figure 1. The overview of the proposed pipeline for part-level action parsing.

Table 1. Bottleneck analysis for part-level action parsing. “✓” means the predictions are replaced by corresponding ground truth.

Actor Detection	Frame-level		Video-level	Metric
	Part parsing			
	part_det	state_parsing	Action parsing	$Acc^p$
				33.32%
✓				35.33%
	✓			36.22%
	✓	✓		72.46%
			✓	42.20%
		✓		45.60%
✓	✓			38.10%
✓	✓	✓		77.30%
✓			✓	44.76%
	✓	✓	✓	93.41%
✓	✓	✓	✓	99.90%

## 2.1. Person Detection

To detect persons from video frames, we follow previous works and adopt the state-of-the-art two-stage object detection approach based on Cascaded-RCNN [3], which starts by region proposal generation, and then refines each proposal in a cascaded manner for predicting objects’ locations with their categories. In addition, human actions can be inferred either from person instances or visual context. Thus, we further extend the Cascaded-RCNN by adding one RCNN branch for instance-level action parsing and a global context module for frame-level action parsing.

**Instantiation:** To implement our person detector, we take the Cascaded-RCNN with ResNet-101 as the basic detection model. In parallel with the person detection branch, two sub-networks are built, involving the Action Parsing

RCNN (AP-RCNN) and the Global Context Parsing (GCP). Specifically, the AP-RCNN is constructed by four consecutive convolutional layers with 256 channels followed by a linear layer with  $C$  channels, where convolutional layers output the instance-level action features  $f_{ia} \in \mathbb{R}^{256}$  and  $C$  denotes the number of action categories. As for the GCP, we simple adopt a global average pooling layer followed by one fully connected layers, where the first layer outputs frame-level action features  $f_c \in \mathbb{R}^{2048}$  with 2048 dimensions and the second one is the  $C$ -way classifier. As a result, the person detector outputs a set of predictions, including the bounding box  $\{\mathcal{P}_{cls} \in \mathbb{R}^2, \mathcal{P}_{box} \in \mathbb{R}^8\}$ , the instance-level categorical distribution  $\mathcal{C}_{ins} \in \mathbb{R}^C$  and the frame-level categorical distribution  $\mathcal{C}_{img} \in \mathbb{R}^C$ .

**Learning objectives:** To enable the model to perform person detection, we design our learning objectives as follows:

$$\begin{aligned} \ell_{cls} &= Cross\_Entropy(\mathcal{P}_{cls}, \mathcal{P}_{cls}^*) \\ \ell_{box} &= SmoothL1(\mathcal{P}_{box}, \mathcal{P}_{box}^*) \\ \ell_{ins} &= Cross\_Entropy(\mathcal{C}_{ins}, \mathcal{C}_{ins}^*) \\ \ell_{img} &= Cross\_Entropy(\mathcal{C}_{img}, \mathcal{C}_{img}^*) \\ \ell_{det} &= \ell_{cls} + \ell_{box} + \ell_{ins} + \ell_{img} \end{aligned} \quad (1)$$

where  $\mathcal{P}_{cls}^*, \mathcal{P}_{box}^*, \mathcal{C}_{ins}^*$  and  $\mathcal{C}_{img}^*$  is the corresponding ground truth.

## 2.2. Part Parsing

In this section, we introduce our method to build part parsing network for facilitating part-level recognition. Given a person image  $I \in \mathbb{R}^{H \times W \times 3}$  of size  $H \times W$ , the goal of part-level recognition is to predict the location of each body part with action states, involving two sub-problems (i.e., part detection and state classification). In addition, the HRNet [9] has shown superiority on many human-centric visual tasks. Therefore, we build our model on the top of HRNet and extend it to the Parsing-HRNet for part-level recognition. The extended version consists of three major components: a visual feature network for visual feature generation, the part parsing modular (PPM) for part detection and a state parsing modular (SPM) for state recognition.

**Instantiation:** Our part parsing model is built on the top of HRNet-W48 [9]. Formally, let  $f \in \mathbb{R}^{d \times H_h \times W_h}$  be the visual features extracted from the visual feature network (i.e., HRNet-W48,  $d = 48$ ). Then, we separately feed visual features into the PPM for part detection and the SPM for state recognition. The PPM is constructed by three convolution layers, where the first two layers are used to generate part features  $f_{pa} \in \mathbb{R}^{d \times H_h \times W_h}$ , and the third one is the output layer that predicts  $K$  heatmaps  $\mathcal{O} \in \mathbb{R}^{K \times H_h \times W_h}$  for part detection. Furthermore, the SPM first applies a global average pooling on visual features for context feature generation. Then the context features are transformed into state features  $f_{sta} \in \mathbb{R}^{192}$  by using two fully connected

layers. Conditioning on the state features, the SPM outputs  $S$ -way probability distribution  $\mathcal{D}_{sta} \in \mathbb{R}^S$  by using one linear layer, where  $S$  is the number of action state classes.

**Learning objectives:** Two loss functions are used to enable the proposed part parsing model to perform part-level parsing. For part heatmap regression, we use the mean square error as the learning objective, and we use the standard focal loss [7] for state classification. The overall loss function is defined in Eq. 2:

$$\begin{aligned} \ell_p &= MSE(\mathcal{O}, \mathcal{O}^*) \\ \ell_s &= FocalLoss(\mathcal{D}_{sta}, \mathcal{D}_{sta}^*) \\ \ell_{part} &= \ell_p + \lambda \ell_s \end{aligned} \quad (2)$$

where  $\lambda$  is the hyperparameter and it is set to 0.5 for balancing training.  $\mathcal{O}^*$  is the ground-truth heatmaps and  $\mathcal{D}_{sta}^*$  is the ground-truth action state. Note that each heatmap is defined as a part-specific binary mask, where each pixel across part area is set to 1, and 0 otherwise.

## 2.3. Action Parsing

Video-level action feature is the key to precisely recognize human actions from videos. In this work, video-level action features are generated from multi-modal features, including non-video-level features and visual video-level features. In particular, non-video-level features involve frame-level action features, instance-level features, part features and part state features. After feature extraction, multiple MLP networks are used for final action recognition.

**Feature extraction:** For non-video-level feature extraction, we directly extract multiple features from the latent layers in the person detection network or the part parsing network, since these two models are fully trained in previous stages. Specifically, we feed each video frame into person detection network and extract outputs from latent layers, including frame-level features  $f_c$  from GCP module and multiple instance-level action features  $\{f_{ia}\}$  from AP-RCNN. As for part-level features, we feed each person image into part parsing network and extract features from two sub-modules, involving spatially pooled part features  $f'_{pa} \in \mathbb{R}^d$  from PPM and action state features  $f_{sta}$  from SPM. Formally, non-video-level features for each video are orderly denoted as: 1) frame-level action features  $f_c \in \mathbb{R}^{T \times 2048}$ ; 2) instance-level action features  $f_{ia} \in \mathbb{R}^{T \times P \times 256}$ ; 3) part features  $f'_{pa} \in \mathbb{R}^{T \times P \times 48}$  and 4) state features  $f_{sta} \in \mathbb{R}^{T \times P \times 192}$ , where  $T$  denotes the number of sampled frames and  $P$  is the number of person boxes predicted in each frame. In our implementation, 32 frames are randomly sampled from each video. For instance-level feature extraction, top-10 person boxes with the highest score in each frame are considered. Besides, we also extract video-level features from two trained action models, including TimeSformer [2] and SwimTransformer [8]. Formally, we denote the  $f_v^t \in \mathbb{R}^{768}$

is the video-level feature extracted from TimeSformer, and  $f_v^s \in \mathbb{R}^{1024}$  is the one from SwimTransformer.

**Inference:** Given non-video-level features, we first feed them into a MLP network, where it contains two fully-connected layers. Then we orderly perform max-pooling along with  $P$  dimension and  $T$  dimension, resulting in video-level features. Finally, we concatenate all video-level features together, and feed this feature into a linear layer to get action prediction.

### 3. Experiments

In this section, evaluation datasets and our implementation details are introduced at first. Then, we conduct detailed ablation studies to investigate the variants of each task-specific model. Finally, we provide our final version that forms the basis for the submission to the Kinetics-TPS Challenge.

#### 3.1. Experimental Settings

We evaluate our models on the 2021 KineticsTPS dataset. There are 3809 annotated videos in training set. Since the test set only provides video-level annotations, we randomly pick 30% of training set as the minival set in the validation phase, resulting in 2686 videos for training and 1123 videos for validation. In the test phase, all models are trained in whole training set with 3809 videos and tested on the official server<sup>1</sup>. We adopt the Average Precision ( $mAP$ ), mean class accuracy ( $Acc$ ) as well as video accuracy conditioning on Part State Correctness ( $Acc^p$ ) as the evaluation metrics. Our models are implemented based on OpenMMLab<sup>2</sup> on an Ubuntu server with eight Tesla V100 graphic cards. For optimization, the details of training configuration is summarized in Tab. 2.

Table 2. Training configurations.

Model	Example model	Training epochs	Learning rate	Drop steps	Solver
Person detection	Cascaded-RCNN(ResNet101)	12	0.02	Linear{8,11}	SGD
Part Parsing	HRNet-W48	40	0.0001	Linear{30,35}	Adam
Action Parsing	SwimTransformer-Base	30	0.001	CosineAnnealing	AdamW

#### 3.2. Ablation Study

In this section, we investigate the variants of proposed disentangled action parsing (DAP) method.

**Person detection.** We choose the Cascaded-RCNN with ResNet101 as the baseline model. Then we investigate the effect of the GCP module and the AP-RCNN by gradually incorporating them into baseline mode. The experimental results are summarized in Tab. 3. From the results, we ob-

serve that the GCP and AP-RCNN bring minor effects to the AP score.

Table 3. Ablation study. The investigation of the person detector variants.

Cascaded-RCNN	GCP	AP-RCNN	mAP
✓			76.6%
✓	✓		77.0%
✓	✓	✓	76.0%

**Part parsing.** In this section, we compare four settings: 1) Shared heatmaps, where the part parsing network outputs action state heatmaps which can be decoded to part location and action states. 2) Separated heatmaps, where the part parsing model outputs part heatmaps and state heatmaps for part decoding and state decoding, respectively. 3) The part parsing model outputs part heatmaps but predicts state onehot vector. 4) Focal loss that used for state classification during training. To evaluate these models, we use the TimeSformer as the video-level classification model that has video accuracy of 85.6% on minival set. We fix the video-level classification model throughout above experiments. The experimental results are summarized in Tab.4. From the results, we observe that separated version is better than shared version in terms of  $Acc^p$  score. The parsing model can be further improved when applying focal loss. This is because the annotation of state classes exists characteristic of imbalance, in particular the number of class 'none' is much higher than that of others.

**Action parsing.** In this section, we investigate the effect of action features. The experimental results are summarized in Tab. 5. From the results, we observe that the multiple features are critical. With non-video-level features, the accuracy score is  $\sim 74\%$ , which implies human actions can be well inferred from person states or frame-level context. In addition, ensemble multiple predictions as the final action scores can bring significant improvement. For example, the video-level accuracy can be increased to 93.2% when ensemble all predictions. With the improved accuracy, the  $Acc^p$  score has been increased to 57.5% which is the best performance so far. Similar performance can be observed in Tab.6.

**Predictions on test set.** With above experiences, we build the final version of DAP and apply it to obtain the predictions results on test set. Specifically, we use the extended Cascaded-RCNN trained on all non-split training set as the person detector, and we apply it to predict human boxes. With predicted human boxes, we choose the best variants of HRNet-W48 that explored in Tab.4 and apply it to predict part locations and part states. As for action parsing, we ensemble multiple predicted results as the final action scores which has a video-level accuracy of 94.1% on test set. With those predicted results, we have formed the submission to Kinetics-TPS challenge. The Tab. 7 summaries the methods

<sup>1</sup><https://competitions.codalab.org/competitions/2360>

<sup>2</sup><https://github.com/open-mmlab>

Table 4. Ablation study. The investigation of part parsing variants.

Shared-heatmaps	Part-heatmaps	State-heatmaps	State-vectors	Focal loss	$Acc^p$
✓					50.9%
	✓	✓			51.3%
	✓		✓		51.7%
	✓		✓	✓	53.2%

Table 5. Ablation study. The effect of video features on minival set.

Frame-level	Instance-level	Video-level	Metric	
$\hat{f}_c$	$\hat{f}_{ia}$ $\hat{f}'_{pa}$ $\hat{f}_{sta}$	$f_v^t$ $f_v^s$	$Acc$	$Acc^p$
✓			73.0%	46.1%
	✓ ✓ ✓		74.2%	46.4%
		✓	85.6%	53.2%
		✓	87.0%	53.6%
Model Ensemble				
Frame-level	Instance-level	Video-level	$Acc$	$Acc^p$
✓	✓		79.3%	49.7%
✓	✓	✓	93.2%	57.5%

Table 6. Ablation study. The effect of video features on test set.

Frame-level	Instance-level	Video-level	Metric	
$\hat{f}_c$	$\hat{f}_{ia}$ $\hat{f}'_{pa}$ $\hat{f}_{sta}$	$f_v^t$ $f_v^s$	$Acc$	
✓			73.8%	
	✓ ✓ ✓		79.0%	
		✓	85.0%	
		✓	87.5%	
Model Ensemble				
Frame-level	Instance-level	Video-level	$Acc$	
✓	✓		82.4%	
✓	✓	✓	94.1%	

from the leaderboard<sup>3</sup> of Kinetics-TPS Challenge. From the results, we see that the proposed DPA obtains 60.5% in  $Acc^p$ , which ranks third place on the test set.

#### 4. Conclusion

In this report, we present a simple yet effective approach for part-level action parsing, named disentangled action parsing (DAP). We divide the part-level action parsing into three stages, involving person detection, part parsing and action parsing. Following this, we design three models, in-

<sup>3</sup><https://competitions.codalab.org/competitions/32360#results>

Table 7. Kinetics-TPS Challenge results on test set.

Ranks	Method	$Acc^p$
(1)	yuzheming	0.630532
(2)	Sheldong	0.613722
(3)	<b>Ours</b>	<b>0.605059</b>
(4)	fangwudi	0.590167
(5)	uestc.wxh	0.536067
(6)	hubincsu	0.490984
(6)	scc1997	0.490984
(7)	KGH	0.486483
(8)	zhao_THU	0.434311
(9)	TerminusBazinga	0.396735
(10)	cjx_AILab	0.370753
(11)	xubo Cheng	0.358834
(12)	haifwu	0.247614
(13)	Aicity	0.189669
(14)	fog	0.188455

cluding the cascaded-rnn variant for person detection, the variant of HRNet for part parsing and multi-modal action parsing model for video-level action parsing. Extensive experiments on Kinetics-TPS dataset demonstrate the effectiveness of the proposed method, and it obtain 60.5%  $Acc^p$  score in 2021 Kinetics-TPS Challenge.

#### References

- [1] Riza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *CVPR*, pages 7297–7306, 2018.
- [2] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding?, 2021.
- [3] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: High quality object detection and instance segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, page 11, 2019.
- [4] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, pages 2980–2988, 2017.
- [5] Jiefeng Li, Can Wang, Hao Zhu, Yihuan Mao, Hao-Shu Fang, and Cewu Lu. Crowdpose: Efficient crowded scenes pose estimation and a new benchmark. In *CVPR*, 2019.
- [6] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 936–944, 2017.

- [7] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollr. Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(2):318–327, 2020.
- [8] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. *arXiv preprint arXiv:2106.13230*, 2021.
- [9] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *CVPR*, pages 5693–5703, 2019.
- [10] Xuanhan Wang, Lianli Gao, Jingkuan Song, and Heng Tao Shen. KTN: knowledge transfer network for multi-person densepose estimation. In *ACM MM*, pages 3780–3788, 2020.
- [11] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *ECCV*, 2018.
- [12] Ceyuan Yang, Yinghao Xu, Jianping Shi, Bo Dai, and Bolei Zhou. Temporal pyramid network for action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [13] Lu Yang, Qing Song, Zhihui Wang, and Ming Jiang. Parsing r-cnn for instance-level human analysis. *arXiv preprint arXiv:1811.12596*, 2018.