

IA-MOT: Instance-Aware Multi-Object Tracking with Motion Consistency

Jiarui Cai, Yizhou Wang, Haotian Zhang, Hung-Min Hsu, Chengqian Ma, Jenq-Neng Hwang
Department of Electrical and Computer Engineering
University of Washington, Seattle, WA, USA

{jrcai, ywang26, haotiz, hmhsu, cm74, hwang}@uw.edu

Abstract

Multiple object tracking (MOT) is a crucial task in computer vision society. However, most tracking-by-detection MOT methods, with available detected bounding boxes, cannot effectively handle static, slow-moving and fast-moving camera scenarios simultaneously due to ego-motion and frequent occlusion. In this work, we propose a novel tracking framework, called “instance-aware MOT” (IA-MOT), that can track multiple objects in either static or moving cameras by jointly considering the instance-level features and object motions. First, robust appearance features are extracted from a variant of Mask R-CNN detector with an additional embedding head, by sending the given detections as the region proposals. Meanwhile, the spatial attention, which focuses on the foreground within the bounding boxes, is generated from the given instance masks and applied to the extracted embedding features. In the tracking stage, object instance masks are aligned by feature similarity and motion consistency using the Hungarian association algorithm. Moreover, object re-identification (ReID) is incorporated to recover ID switches caused by long-term occlusion or missing detection. Overall, when evaluated on the MOTS20 and KITTI-MOTS dataset, our proposed method won the first place in Track 3 of the BMTT Challenge in CVPR2020 workshops.

1. Introduction

Multi-Object Tracking (MOT) leveraging detected bounding box locations has been the mainstream and well-studied, a.k.a., tracking-by-detection. Most tracking-by-detection methods firstly extract the objects’ deep features from the detected bounding boxes. With these extracted features, a tracking algorithm is followed to associate the detections along consecutive frames. However, the bounding-box-based features contain significant background noises, which are usually redundant and would downgrade the tracking performance. Recently, Multi-Object Tracking and Segmentation (MOTS) [10] is posed as a new task to track

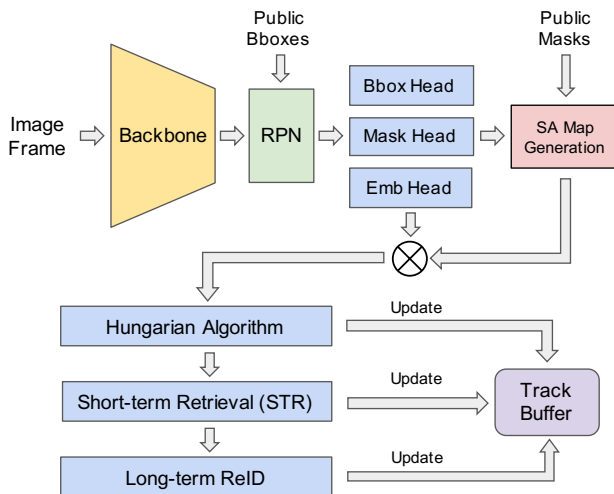


Figure 1. The framework of the proposed IA-MOT. First, embedding features are extracted using a variant of Mask R-CNN and spatial attention generation. Then, isolated detections are aligned by Hungarian algorithm based on the mask IOU and cosine feature similarity. Finally, missing detection and long-term occlusion are handled by the short-term retrieval and ReID modules.

multiple objects with the corresponding instance segmentation. MOTS is more informative and spatially precise, and the masks also introduce potential promotions to tracking accuracy.

Typically, there are three key components in MOTS: object detection and segmentation, feature extraction, and multi-object tracking. Some related works are briefly reviewed as follows. Recent studies on object detection and segmentation achieve impressive performance with various advanced structures [8, 1]. Especially, for mask generation in MOTS, Luiten et al. [5] present the Proposal-generation, Refinement and Merging for Video Object Segmentation algorithm (PREMVOS). First, a set of object segmentation mask proposals are generated from Mask R-CNN and refined by a refinement network frame-by-frame. Then, with the assistance of optical-flow and ReID information, these

selected proposals are merged into pixel-wise object tracks over a video sequence. Feature extraction with deep neural networks is widely used in the appearance similarity comparison in tracking. FaceNet [9] proposes a metric learning strategy. Moreover, a multi-scaled pedestrian detector [12] is proposed to allow target detection and appearance features to be learned in a shared model. On top of that, several MOT methods are studied to obtain the object trajectories. Wang et al. [11] propose the TrackletNet Tracking (TNT) framework that jointly considers appearance and geometric features to form a tracklet-based graph model. To incorporate segmentation into MOT, Voigtlaende et al. [10] propose a new baseline approach, Track R-CNN, which addresses detection, tracking, and segmentation via a unified convolutional neural network. Milan et al. [6] use super-pixel information and consider detection and segmentation jointly with a conditional random field model. CAMOT [7] exploits stereo information to perform mask-based tracking of generic objects on the KITTI dataset.

Instance segmentation is trained and inferred in parallel with the bounding box detection in the previous studies. In this paper, to better utilize the instance masks for embedding feature extraction, we propose a novel MOTS framework, called instance-aware multi-object tracking (IA-MOT) with motion consistency (IA-MOT) to integrate the segmentation information with instance features. With the given detection results, the sequences are sent into a variant of Mask R-CNN [1] to obtain the embedding features as region proposals of RPN. Here, to handle instance-aware features, we apply the spatial attention generated from the corresponding instance masks to weight more on the foreground of each bounding box. Second, the extracted instance-aware features are utilized in the following tracking module based on the Hungarian assignment algorithm. In this tracking module, object motion consistency, i.e., the similarity of object motion, including object sizes, moving direction, and moving speed, are jointly considered. Moreover, discontinuity caused by missing detection is recovered by the short-term retrieval (STR) module and object re-identification (ReID) is further incorporated to reduce ID switches caused by long-term occlusion.

2. The Proposed Method

The proposed IA-MOT includes three steps: embedding feature extraction, online object tracking with STR, and object re-identification with motion consistency for final refinement. The overall framework is shown in Figure 1.

2.1. Instance-Aware Embedding Features

First of all, embedding feature for each given bounding box is extracted from a variant of Mask R-CNN [2] with an extra embedding head. This network is trained by the joint optimization of bounding box classification and regression,

object mask prediction, and object feature extraction as a multi-task problem. The loss function is

$$\mathcal{L}_{total} = \alpha_1 \mathcal{L}_{box} + \alpha_2 \mathcal{L}_{cls} + \alpha_3 \mathcal{L}_{mask} + \alpha_4 \mathcal{L}_{emb}, \quad (1)$$

where \mathcal{L}_{box} , \mathcal{L}_{cls} and \mathcal{L}_{mask} are the original Mask R-CNN loss, \mathcal{L}_{emb} is the cross-entropy loss for object identity classification. Here, each public detection $D_n \in \mathcal{D}$ is treated as a region proposal and fed into the network to acquire its bounding box feature emb_n^{box} .

Then, the extracted embedding features emb_n^{box} are filtered by the spatial attention (SA) map w_n to form our instance-aware (IA) embedding features emb_n^{IA} , where

$$emb_n^{IA} = w_n \cdot emb_n^{box}. \quad (2)$$

Here, the SA map w_n is generated from the provided instance segmentation $mask_n$, where

$$w_n(i, j) = \begin{cases} 1, & \text{if } mask_n(i, j) \text{ is foreground,} \\ 0.5, & \text{if } mask_n(i, j) \text{ is background.} \end{cases} \quad (3)$$

2.2. Online Tracking with STR

After the IA embedding features are extracted, they are transmitted to the following tracking algorithm to associate the same identities among consecutive frames. With detections $\mathcal{D}^{(t)}$ and the corresponding feature embeddings $\mathcal{E}^{(t)}$ in frame t , the hypotheses $D_i^{(t)} \in \mathcal{D}^{(t)}$ are matched with live tracks $D_j^{(t-1)} \in \mathcal{D}^{(t-1)}$ in frame $t - 1$ using the Hungarian assignment algorithm [3]. The assignment cost between the j -th track and the i -th detection is computed by

$$\mathcal{C} = 2 - \text{maskIOU}(D_j^{(t-1)}, D_i^{(t)}) - \text{simi}(emb_j, emb_i^{(t)}), \quad (4)$$

where the j -th track's feature emb_j is the stack of its first 5 frames and the most recent 5 frames, $\text{simi}(\cdot)$ represents the cosine similarity. Considering the visibility of an object may vary due to camera movement, the features are not aggregated into one, but kept separated. Feature similarity is compared pair-wisely, and the maximum is taken as the result. Missing detections are compensated by the short-term retrieval (STR) module for lost objects. STR tries to match unassigned detections in frame t with the live tracks that without a detection in frame $t - 1$.

In addition to the feature similarity, the bounding boxes in current frame of the lost track are extrapolated by Huber regression. The distance between the regressed location and the track's last location is confined to within twice the object width. Tracklets will be marked as terminated if there are no alignments for the most recent N_1 frames and not be included in the Hungarian assignment or STR.



Figure 2. Qualitative results of the proposed IA-MOT. 1st row: MOTS20-01 (static camera); 2nd row: MOTS20-06 (stroller-mounted camera); 3rd row: KITTI-MOTS-0011 (car-mounted camera, up-slope view); 4th row: KITTI-MOTS-0012 (car-mounted camera, turning view). Red arrows indicate that the targets are tracked robustly even with frequent occlusion or turning.

2.3. ReID with Motion Consistency

Then, long-term occlusions are recovered by feature-based re-identification (ReID). In this stage, two tracklets ξ_u and ξ_v without overlapped frames in time (assuming ξ_u is earlier than ξ_v), within N_2 frames apart, and with feature similarity higher than β_1 , are considered as possible matched pairs. For static cameras, the interval in between is extrapolated from ξ_u and ξ_v , respectively. ξ_u and ξ_v are reconnected if the average bounding box IOU between two extrapolation are above β_2 . For moving cameras, tracklet motion vectors are estimated from the ξ_u 's last or ξ_v 's first N_3 frames. Thus, the motion of ξ_u can be defined as

$$M_u = [M_{ux}, M_{uy}] = \frac{1}{N_3 - 1} \left[\sum_{j=1}^{N_3-1} (x_u^{(j+1)} - x_u^{(j)}), \sum_{j=1}^{N_3-1} (y_u^{(j+1)} - y_u^{(j)}) \right], \quad (5)$$

where $[x_u^{(j)}, y_u^{(j)}]$ is the top-left point of the j -th detection in ξ_u . The motion M_v can be defined in the same manner. Then, ξ_u and ξ_v are recognized as the same object if the cosine similarity between M_u and M_v are positive and above a threshold β_3 .

3. Experiments

3.1. Dataset

The data of the BMTT Challenge consists of MOTS20 and KITTI-MOTS dataset [10]. There are 8 sequences in MOTS20 dataset for pedestrian tracking, and evenly split for training and testing. In testing set, the resolution varies from 640×480 to 1920×1080 with an average density of 10.6 targets per frame. KITTI-MOTS is a driving scenario dataset for both pedestrian and car tracking tasks, consisting of 21 training sequences and 29 testing sequences. The pre-computed detections are generated from Mask R-CNN X152 [1] and refined by the refinement net [5].

3.2. Implementation Details

The proposed modified Mask R-CNN uses ResNet50 [2] as the backbone, which is pretrained on COCO dataset [4] and fine-tuned on the MOTS20 dataset and KITTI-MOTS dataset. Based on the convergence speed and scaling of each loss component, the detection loss weight $\alpha_1 = \alpha_2 = 1$, mask weight $\alpha_3 = 0.8$ and embedding weight $\alpha_4 = 0.1$. The output feature is 1024-dim. Short-term memory interval to determine the state of a track is $N_1 = 0.2$ second, while the long-term interval for ReID is $N_2 = 1$ second for pedestrian, and $N_1 = 0.1$ second, $N_2 = 0.5$ second for cars. $N_3 = 5$ frames for both categories.

Moreover, due to the large number of false positives

Method	KITTI-MOTS		MOTS20	Total
	Car	Ped		
MCFPA	77.0	67.2	66.1	69.1
TPM-MOTS	75.8	67.3	66.6	69.1
ReMOTS	72.6	64.6	67.9	68.3
GMPHD_SAF	76.2	64.3	64.3	67.3
Lif_TS	77.5	55.8	65.3	66.0
SRF	71.4	60.9	60.0	63.1
KQD	74.4	61.8	57.3	62.7
USN	72.1	59.3	59.5	62.6
YLC	62.3	57.2	59.1	59.4
SI	68.5	55.5	56.2	59.1
FK	64.1	54.5	54.3	56.8
IA-MOT (Ours)	76.4	64.0	69.4	69.8

Table 1. Track 3 evaluation results of the BMTT Challenge, evaluated by sMOTSA. Best results are marked in **bold**.

in the provided detections, we create three different filters, including detection confidence, bounding box size, and bounding box aspect ratio, to select valid candidates. After tracking with the filtered detections, short tracks or tracks with low average confidence are discarded. In addition, the trajectory IOU is defined as the average mask IOU of two tracks over their co-exist frames. If the trajectory IOU is larger than 0.75, the shorter track will be discarded. These processing steps efficiently remove duplicated and non-target detections.

3.3. MOTSA Results

We evaluate our IA-MOT on Track 3 dataset of the BMTT Challenge (the combination of MOTSA20 and KITTI-MOTS) and get the first place, with 69.8 sMOTSA, out of 13 methods. The quantitative results are shown in Table 1 and some qualitative examples are shown in Figure 2.

Specifically, the proposed framework achieves 69.4 in MOTSA20 dataset with the public detections, which is nearly even with the leading method in the MOTSA20 leaderboard with private detectors. Although IA-MOT does not get the best performance on KITTI-MOTS, it still ranks the 3rd for KITTI-MOTS car and 5th for pedestrian, indicating its ability to generalize to different object categories and the potential to handle different complex MOTSA scenarios.

4. Conclusion

In this paper, a novel MOTSA framework – IA-MOT is proposed. A Mask R-CNN with an additional embedding head and spatial attention first generate discriminating features. The following MOT stage consists of online Hungarian assignment, short-term retrieve module and ReID. In addition, several implementation details are presented for the MOTSA20 and KITTI-MOTS dataset. The proposed

framework could effectively track both pedestrian and car, with static or moving cameras, and is flexible for different video resolution and scenarios. Our proposed IA-MOT achieved the winner in Track 3 of the BMTT Challenge in CVPR2020 workshops.

References

- [1] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 1, 2, 3
- [2] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2, 3
- [3] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 2
- [4] 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
- [5] B Leibe J Luiten, P Voigtlaender, and B Leibe. Premvos: Proposal-generation. *Refinement and Merging for the DAVIS Challenge on Video Object Segmentation*, 2018. 1, 3
- [6] Anton Milan, Laura Leal-Taixé, Konrad Schindler, and Ian Reid. Joint tracking and segmentation of multiple targets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5397–5406, 2015. 2
- [7] Aljoša Ošep, Wolfgang Mehner, Paul Voigtlaender, and Bastian Leibe. Track, then decide: Category-agnostic vision-based multi-object tracking. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8. IEEE, 2018. 2
- [8] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 1
- [9] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. 2
- [10] Paul Voigtlaender, Michael Krause, Aljoša Ošep, Jonathon Luiten, Berin Balachandar Gnana Sekar, Andreas Geiger, and Bastian Leibe. Mota: Multi-object tracking and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7942–7951, 2019. 1, 2, 3
- [11] Gaoang Wang, Yizhou Wang, Haotian Zhang, Renshu Gu, and Jenq-Neng Hwang. Exploit the connectivity: Multi-object tracking with trackletnet. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 482–490, 2019. 2
- [12] Zhongdao Wang, Liang Zheng, Yixuan Liu, and Shengjin Wang. Towards real-time multi-object tracking. *arXiv preprint arXiv:1909.12605*, 2019. 2