

MeNToS : Tracklets Association with a Space-Time Memory Network

Mehdi Miah, Guillaume-Alexandre Bilodeau and Nicolas Saunier
Polytechnique Montréal

{mehdi.miah, gabilodeau, nicolas.saunier}@polymtl.ca

Abstract

We propose a method for multi-object tracking and segmentation (MOTS) that does not require fine-tuning or per benchmark hyperparameter selection. The proposed method addresses particularly the data association problem. Indeed, the recently introduced HOTA metric, that has a better alignment with the human visual assessment by evenly balancing detections and associations quality, has shown that improvements are still needed for data association. After creating tracklets using instance segmentation and optical flow, the proposed method relies on a space-time memory network (STM) developed for one-shot video object segmentation to improve the association of tracklets with temporal gaps. To the best of our knowledge, our method, named MeNToS, is the first to use the STM network to track object masks for MOTS. We took the 4th place in the RobMOTS challenge. The project page is <https://mehdimiah.com/mentos.html>.

1. Introduction

Multi-object tracking (MOT) is a core problem in computer vision. Given a video, the objective is to detect all objects of interest then to track them throughout the video with consistent identities. Common difficulties are occlusions, small objects, fast motion (or equivalently low framerate) and deformations. Recently, the multi-object tracking and segmentation (MOTS) task [10] was introduced: instead of localizing objects with bounding boxes, they are described by their segmentation mask at the pixel level.

The MOTA metric has been commonly used to evaluate MOT but has a tendency to give more weight to detection errors compared to association errors. The newly introduced HOTA metric [5] balances these two aspects and provides further incentives to work on the association step. That is why we developed a method which relies first on an instance segmentation followed by two data association steps. The first one is applied between consecutive frames using optical flow to obtain mask location prediction and mask Intersection over Union (mIoU) for matching. The

second association step relies on a space-time memory network (STM). This is our main contribution. It is inspired by some results in one-shot video object segmentation (OSVOS), a task in computer vision that consists of tracking at a pixel level a mask provided only at the first frame. We use mask propagation with a STM network to associate tracklets separated by longer temporal gaps. Experiments show that the long-term data association significantly improves the HOTA score over the datasets used in the challenge.

2. Related works

MOTS Similarly to MOT where the “tracking-by-detection” paradigm is popular, MOTS is mainly solved by creating tracklets from segmentation masks then by building long-term tracks by merging the tracklets [3, 12, 13]. Usually, methods use an instance segmentation method to generate binary masks; Re-MOTS [12] used two advanced instance segmentation methods and self-supervision to refine masks. As for the association step, many methods require a re-identification (reID) step. For instance, Voigtlaender et al. [10] extended the Mask R-CNN by an association head to return an embedding for each detection, Yang et al. [12] associated two tracklets if they were temporally close, without any temporal overlap with similar appearance features based on all their observations and a hierarchical clustering, while Zhang et al. [13] used temporal attention to lower the weight of frames with occluded objects.

STM Closely related to MOTS, OSVOS requires tracking objects whose segmentation masks are only provided at the first frame. STM [6] was proposed to solve OSVOS by storing some previous frames and masks in a memory that is later read by an attention mechanism to predict the new mask in a target image. Such a network was recently used [2] to solve video instance segmentation, a problem in which no prior knowledge was given about the objects to track. However, it is unclear how STM behaves when multiple instances from the same class appear in the video. We show in this work they behave well and that they can help to solve a reID problem by taking advantage of the information at the pixel level and the presence of other objects.

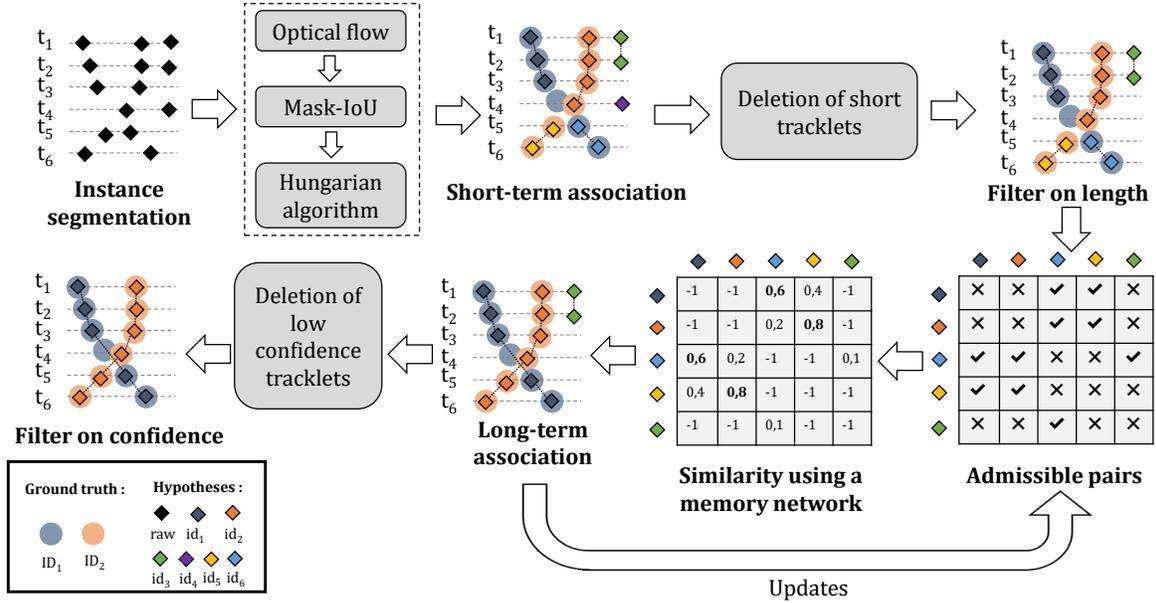


Figure 1. Illustration of our MeNToS method. Given an instance segmentation, binary masks are matched in consecutive frames to create tracklets. Very short tracklets are deleted. An appearance similarity, based on a memory network, is computed between two admissible tracklets. Then, tracklets are gradually merged starting with the pair having the highest similarity while respecting the updated constraints. Finally, low confidence tracks are deleted.

3. Method

As illustrated in Figure 1, our pipeline for tracking multiple objects is based on three main steps: detections of all objects of interest, a short-term association of segmentation masks in consecutive frames and a greedy long-term association of tracklets using a memory network.

3.1. Detections

Our method follows the “tracking-by-detection” paradigm. First, we used the public raw object masks provided by the challenge. They were obtained from a Mask R-CNN X-152 and Box2Seg Network. Objects with a detection score higher than θ_d and bigger than θ_a are extracted. Then, to avoid, for instance, that a car is simultaneously detected as a car and as a truck, segmentation masks having a mutual mIoU higher than θ_{mIoU} are merged to form a multi-class hypothesis.

3.2. Short-term association (STA)

We associate temporally close segmentation masks between consecutive frames by computing the Farneback optical flow [1] due to its simplicity. Masks from the previous frames are warped and a mIoU is computed between these warped masks and the masks from the next frame.

The Hungarian algorithm is used to associate masks where the cost matrix is computed based on the negative mIoU. Matched detections with a mIoU above a threshold

θ_s are connected to form a tracklet and the remaining detections form new tracklets. The class of the tracklet is the most dominant one among its detections. Then, a non-overlap algorithm is applied to avoid any spatial overlap between masks, giving the priority to the pixels of the most confident mask. Finally, tracklets with only one detection are deleted since they often correspond to a false positive.

3.3. Greedy long-term association (GLTA)

GLTA and the use of a memory network for re-identification are the novelties of our approach. Once tracklets have been created, it is necessary to link them in case of fragmentation caused, for example, by occlusion. In this long-term association, we used a memory network to propagate some masks of a tracklet in the past and the future. In case of a spatial overlap with another tracklet, the two tracklets are merged. Given the fact that this procedure is applied at the pixel-level on the whole image, the similarity is only computed on a selection of admissible tracklets pairs to reduce the computational cost. At this step, we point out that all tracklets have a length longer than or equal to two.

3.3.1 Measure of similarity between tracklets

Our similarity measure is based on the ability to match some parts of two different tracklets (say T^A and T^B) and can be interpreted as a pixel-level visual-spatial alignment rather than a patch-level visual alignment [12, 13]. For that, we

propagate some masks of tracklet T^A to other frames where the tracklet T^B is present and then compare the masks of T^B and the propagated version of the mask heatmaps, computed before the binarization, for T^A . The more they are spatially aligned, the higher the similarity is. In details, let us consider two tracklets $T^A = (M_1^A, M_2^A, \dots, M_N^A)$ and $T^B = (M_1^B, M_2^B, \dots, M_P^B)$ of length N and P respectively, such that T^A appears first and where M_1^A denotes the first segmentation mask of the tracklet T^A . We use a pre-trained STM network [6] to store two binary masks as references (and their corresponding frames): the closest ones (M_N^A for T^A and M_1^B for T^B) and a second mask a little farther (M_{N-n-1}^A for T^A and M_n^B for T^B). The farther masks are used because the first and last object masks of a tracklet are often incomplete due, for example, to occlusions. Then, the reference frames are used as queries to produce heatmaps with values between 0 and 1 ($H_N^A, H_{N-n-1}^A, H_1^B, H_n^B$). Finally, the average cosine similarity between these four heatmaps and the four masks ($M_N^A, M_{N-n-1}^A, M_1^B, M_n^B$) is the final similarity between the two tracklets. Figure 2 illustrates a one-frame version of this similarity measure between tracklets.

3.3.2 Selection of pairs of tracklets

Instead of estimating a similarity measure between all pairs of tracklets, a naive selection is made to reduce the computational cost. The selection is based on the following heuristic: two tracklets may belong to the same objects if they belong to the same class, are temporally close, spatially close and with a small temporal overlap.

In details, let us denote $f(M)$ the frame where the mask M is present, \bar{M} its center and fps, H and W respectively the number of frames per second, height and width of the video. The temporal ($C_t(T^A, T^B)$), spatial ($C_s(T^A, T^B)$) and temporal overlap ($C_o(T^A, T^B)$) costs between T^A and T^B are defined respectively as:

$$C_t(T^A, T^B) = \frac{|f(M_N^A) - f(M_1^B)|}{fps}, \quad (1)$$

$$C_s(T^A, T^B) = \frac{2}{H + W} \times \|\bar{M}_N^A - \bar{M}_1^B\|_1, \quad (2)$$

$$C_o(T^A, T^B) = |\{f(M) \forall M \in T^A\} \cap \{f(M) \forall M \in T^B\}| \quad (3)$$

A pair (T^A, T^B) is admissible if the tracklets belong to the same class, $C_t(T^A, T^B) \leq \tau_t$, $C_s(T^A, T^B) \leq \tau_s$ and $C_o(T^A, T^B) \leq \tau_o$.

3.3.3 Greedy association

Similarly to Singh et al. [8], we gradually merge the admissible pairs with the highest cosine similarity, if it is above a threshold θ_l , while continuously updating the admissible

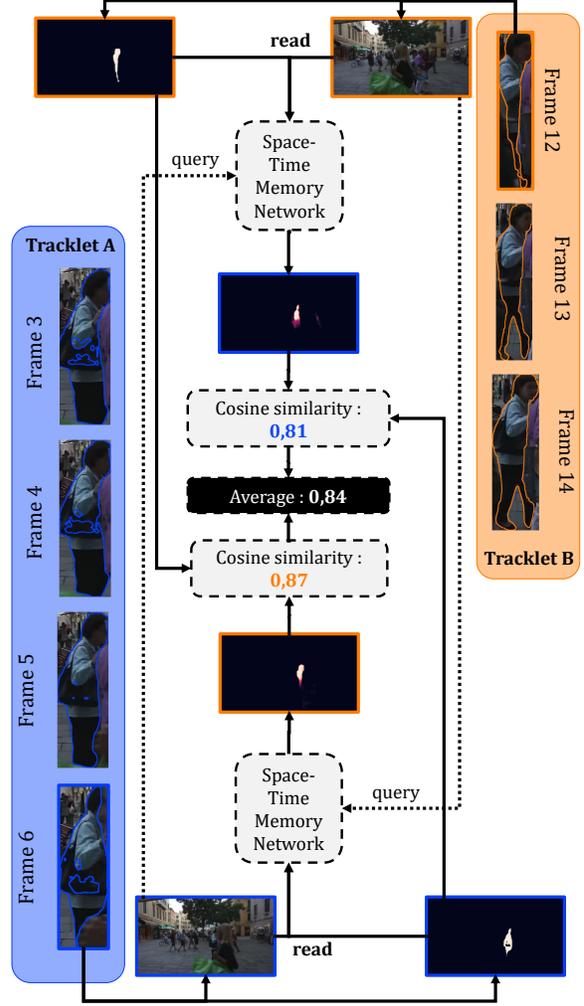


Figure 2. Similarity used at the long-term association step. For a matter of simplicity, only one mask and frame is used as reference and as target in the space-time memory network.

pairs using equation 3. A tracklet can therefore be repeatedly merged with other tracklets. Finally, tracks having their highest detection score lower than 90 % are deleted.

4. Experiments

4.1. Implementation details

At the detection step, θ_d is 0.5, small masks whose area is less than $\theta_a = 128$ pixels are removed and $\theta_{mIoU} = 0.5$. For the GLTA step, the selection is done with $(\tau_t, \tau_s, \tau_o) = (1.5, 0.2, 1)$. To measure similarity, the second frame is picked using $n = 5$. If that frame is not available, $n = 2$, is used instead. As for the thresholds at the STA and GLTA steps, we selected $\theta_s = 0.15$ and $\theta_l = 0.30$. These hyperparameters were selected through cross-validation and remain fixed regardless of the dataset and object classes.

Method	BDD	DAVIS	KITTI	MOTSCha.	OVIS	TAO	Waymo	YT-VIS	RobMOTS		
	HOTA	HOTA	HOTA	HOTA	HOTA	HOTA	HOTA	HOTA	HOTA	DetA	AssA
RobTrack [11]	57.9	56.9	71.6	61.0	61.6	55.0	57.2	68.3	61.2	59.4	64.8
SBT [9]	53.0	50.3	74.0	64.4	55.6	51.8	55.2	64.4	58.6	55.9	63.1
SIA [7]	53.4	47.4	70.8	62.2	54.8	49.6	54.1	62.7	56.9	55.8	59.8
MeNToS	52.3	49.6	69.7	60.2	55.6	39.2	53.4	64.2	55.5	52.4	60.8
STP [4]	49.4	48.2	66.4	60.4	52.8	43.8	51.8	62.3	54.4	55.8	55.0

Table 1. Results on the RobMOTS test set. The HOTA metrics on each benchmark is reported alongside with the overall DetA, AssA and HOTA. Red and blue indicate respectively the first and second best methods.

4.2. Datasets and performance evaluation

The tracking algorithms are applied on the benchmarks on the RobMOTS challenge [4]. It consists of eight tracking datasets with a high diversity in terms of framerate (ranges from 1 to 30 frames per second), objects of interest, duration and number of objects. Here, we considered the 80 categories of objects from COCO.

Recently, the HOTA metric was introduced to fairly balance the quality of detections and associations. It can be decomposed into the DetA and AssA metrics to measure the quality of these two components. The higher the HOTA is, the more the tracker is aligned with the human visual assessment. The final HOTA on RobMOTS is the average of the eight HOTA metrics.

4.3. Results

Results of Table 1 indicate that our method is competitive for MOTs. MeNToS performs well on all benchmarks except TAO. This benchmark is more difficult for MeNToS, and all the other methods, since it is composed of videos with a very low framerate (1 fps). Without this outlier, MeNToS would perform on par with SIA.

This issue comes from the second step of our method where the optical flow struggles to correctly associate consecutive masks of the same object. As a result, the deletion of very short tracklets leads to remove detections in this case, thus reducing DetA, the quality of detection (-3 percentage points on DetA compared to the baseline STP).

However, MeNToS is able to correctly associate tracklets in the long-term data association partially balances this drawback (+6 percentage points on AssA compared to STP).

5. Conclusion

In this work, we have developed a memory network-based tracker for multi-object tracking and segmentation. After creating tracklets, the STM network is used to compute a similarity score between tracklets. We can interpret this evaluation as a pixel-level visual-spatial alignment leveraging segmentation masks and the information of the whole image. Improving the creation of tracklets during

the short-term data association may lead to further improvements.

Acknowledgment

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), [DGDND-2020-04633 and DG individual 06115-2017].

References

- [1] Gunnar Farneback. Two-Frame Motion Estimation Based on Polynomial Expansion. In Josef Bigun and Tomas Gustavsson, editors, *Image Analysis, Lecture Notes in Computer Science*, pages 363–370, Berlin, Heidelberg, 2003. Springer. 2
- [2] Shubhika Garg and Vidit Goel. Mask Selection and Propagation for Unsupervised Video Object Segmentation. In *WACV*, 2021. 1
- [3] J. Luiten, T. Fischer, and B. Leibe. Track to Reconstruct and Reconstruct to Track. *IEEE Robotics and Automation Letters*, 5(2):1803–1810, Apr. 2020. 1
- [4] Jonathon Luiten, Arne Hoffhues, Blin Beqa, Paul Voigtlaender, István Sárándi, Patrick Dendorfer, Aljosa Osep, Achal Dave, Tarasha Khurana, Tobias Fischer, Xia Li, Yuchen Fan, Pavel Tokmakov, Song Bai, Linjie Yang, Federico Perazzi, Ning Xu, Alex Bewley, Jack Valmadre, Sergi Caelles, Jordi Pont-Tuset, Xinggang Wang, Andreas Geiger, Fisher Yu, Deva Ramanan, Laura Leal-Taixé, and Bastian Leibe. RobMOTS : A Benchmark and Simple Baselines for Robust Multi-Object Tracking and Segmentation. In *CVPR RVSU Workshop*, 2021. 4
- [5] Jonathon Luiten, Aljosa Osep, Patrick Dendorfer, Philip Torr, Andreas Geiger, Laura Leal-Taixé, and Bastian Leibe. HOTA: A Higher Order Metric for Evaluating Multi-Object Tracking. *International Journal of Computer Vision (IJCV)*, Oct. 2020. 1
- [6] Seoung Wug Oh, Joon-Young Lee, Ning Xu, and Seon Joo Kim. Video Object Segmentation Using Space-Time Memory Networks. In *ICCV*, 2019. 1, 3
- [7] Jeongwon Ryu and Kwangjin Yoon. SIA : Simple Re-Identification Association for Robust Multi-Object Tracking and Segmentation. In *CVPR RVSU Workshop*, 2021. 4
- [8] Gurinderbeer Singh, Sreeraman Rajan, and Shikharesh Majumdar. A Greedy Data Association Technique for Multiple Object Tracking. In *2017 IEEE Third International Conference on Multimedia Big Data (BigMM)*, pages 177–184, Apr. 2017. 3
- [9] Jiasheng Tang, Fei Du, Weihua Chen, Hao Luo, Fan Wang, and Hao Li. SBT : A Simple Baseline with Cascade Association for Robust Multi-Objects Tracking. In *CVPR RVSU Workshop*, 2021. 4
- [10] Paul Voigtlaender, Michael Krause, Aljosa Osep, Jonathon Luiten, Berin Balachandrar Gnana Sekar, Andreas Geiger, and Bastian Leibe. MOTs: Multi-Object Tracking and Segmentation. In *CVPR*, 2019. 1
- [11] Dongxu Wei, Jiashen Hua, Hualiang Wang, Baisheng Lai, Kejie Huang, Chang Zhou, Jianqiang Huang, and Xiansheng Hua. RobTrack : A Robust Tracker Baseline towards Real-World Robustness in Multi-Object Tracking and Segmentation. In *CVPR RVSU Workshop*, 2021. 4
- [12] Fan Yang, Xin Chang, Chenyu Dang, Ziqiang Zheng, Sakriani Sakti, Satoshi Nakamura, and Yang Wu. ReMOTS: Self-Supervised Refining Multi-Object Tracking and Segmentation. In *CVPR - Workshops*, 2020. 1, 2
- [13] Haotian Zhang, Yizhou Wang, Jiarui Cai, Hung-Min Hsu, Haorui Ji, and Jen-Neng Hwang. LIFTS: Lidar and monocular image fusion for multi-object tracking and segmentation. In *CVPR - Workshops*, 2020. 1, 2