

Multiple Object Tracking from appearance by hierarchically clustering tracklets

Andreu Girbau¹

agirbau@nii.ac.jp

Ferran Marqués²

ferran.marques@upc.edu

Shin'ichi Satoh¹

satoh@nii.ac.jp

¹ National Institute of Informatics

Tokyo, Japan

² Universitat Politècnica de Catalunya

Barcelona, Spain

Abstract

Current approaches in Multiple Object Tracking (MOT) rely on the spatio-temporal coherence between detections combined with object appearance to match objects from consecutive frames. In this work, we explore MOT using object appearances as the main source of association between objects in a video, using spatial and temporal priors as weighting factors. We form initial tracklets by leveraging on the idea that instances of an object that are close in time should be similar in appearance, and build the final object tracks by fusing the tracklets in a hierarchical fashion. We conduct extensive experiments that show the effectiveness of our method over three different MOT benchmarks, MOT17, MOT20, and DanceTrack, being competitive in MOT17 and MOT20 and establishing state-of-the-art results in DanceTrack. Code is available at https://github.com/NII-Satoh-Lab/MOT_FCG.

1 Introduction

The objective of the Multiple Object Tracking (MOT) task is to estimate the trajectory of a set of objects (e.g. pedestrians or vehicles) along a video sequence. The objects of interest must be represented with an accurate bounding box, and keep the associated identity over time. MOT can be useful in many applications, such as autonomous driving, robotics, or automatic production of events.

The dominant approach in MOT is tracking-by-detection, a two-step process where, first, the objects of interest are detected frame-by-frame, and then associated across frames to form tracks. This approach became feasible thanks to the latest advances in object detection [1, 3, 9, 16, 27], leading to the current dominant trackers, which use spatio-temporal priors combined with appearance features as the main source of association.

In this paper, we explore the association of objects by primarily focusing on their appearance, using spatial and temporal priors as a support to weight object relationships, and present a simple and adaptable MOT method based on tracklet generation and clustering. Our tracker, FCG (Feature Combinatorial Grouping), leverages on the idea that a specific object is prone to have a similar appearance in a temporal neighborhood. Based on this,

we form high quality tracklets, smaller parts of an object track, for each object in a local neighborhood, capturing small variations of the object appearance, and cluster the generated tracklets in a hierarchical way to form the final tracks, taking into account the sequential nature of a video, without any kind of post-processing or optimization when performing the tracking.

We conduct extensive experiments to show the validity of our method, and verify the quality of FCG in three MOT benchmarks, MOT17 [24], MOT20 [40], and DanceTrack [65], achieving comparable performance against current state-of-the-art trackers for MOT17 and MOT20, and setting state-of-the-art results for DanceTrack.

Our contributions in this paper are:

- We explore the effectiveness of focusing on appearance as the main source for associating objects through time for MOT, and present FCG, a tracker that relies on object appearance features to form tracks by hierarchically clustering tracklets.
- We show that a fairly simple tracklet association by clustering method, using off-the-shelf pre-trained models and without any kind of optimization or post-processing, can perform competitively to current, way more complex, state-of-the-art techniques.
- We achieve state-of-the-art results for the DanceTrack benchmark, while presenting competitive results in MOT17 and MOT20.

2 Related work

Advances in object detection [11, 13, 26, 27] allowed current tracking methods to rely on frame-by-frame detections to perform tracking of multiple objects. This led to the current predominant approach in MOT, tracking-by-detection, which addresses the tracking problem in two steps: (i) detection of the objects in the scene, and (ii) the association of the detections through time to form tracks. Simultaneously, object appearance features for Re-Identification have been highly improved, and have been proven very effective [15, 21, 28, 40].

2.1 Object appearance in MOT

Despite some state-of-the-art MOT methods do not use the object appearance at all [4, 6, 34, 51], relying on motion models to associate close objects between frames, object appearance is extensively used in MOT. Some MOT methods account indirectly for object appearance [36, 48, 54], where the motion of objects is inferred using previous detections and image features, train models to both learn to detect and represent objects at the same time with multiple output heads [47, 50], using recurrent neural networks [42], attention mechanisms [7] or correlation operations between convolutional feature maps [17, 41] to encode both motion and feature information, or even with end-to-end approaches where feature extraction, affinity estimation, and objects association are refined in a single network [8].

Explicit usage of object appearance is also extensively used in MOT, combining the spatio-temporal information with object appearance information as a weighting factor [43], optimizing a graph where the position and feature embedding information form the nodes and edges [5, 16], or to recover lost tracks by doing re-identification [4, 23].

The key difference between our proposed method and previous works is what is prioritized for doing object association, as previous works prioritize spatio-temporal priors or their

combination with object appearance to associate objects. In this work, we explore using appearance features to associate the objects as the main source of association, using spatial and temporal information as a support weighting factor.

2.2 Tracking by tracklet association

Tracklets, introduced in [10], are smaller parts of whole object tracks that, once fused, form the whole object track. Tracking by associating tracklets is a well established technique [11], that has had far less attention than other offline tracking techniques such as graph optimization considering individual detections. One common assumption in tracking by associating tracklets is to consider tracklets as a robust temporal representation of the object. Once the initial tracklets are built, they are associated to form the final tracks.

Some methods associate tracklets by using a MAP (Maximum A Posteriory) approach, conditioning the overall result every time a tracklet is fused [12], cluster tracklets based on color, spatial and temporal attributes [13], form tracks by clustering objects using spatial cues [14, 15], gradually add similar objects to the initial tracklets [16], or treat tracklets as the nodes of a graph, and the similarity between tracklets as the vertexes [9, 17, 18, 19].

Tracklet association methods are usually treated as an optimization problem, whether it is iteratively updating cluster assignments [14, 15, 16], using a graph formulation with the subsequent optimization [9, 17, 18, 19, 20], or re-training CNN networks to adapt ReID features to new sequences [21]. Our method hierarchically fuses tracklets based on the objects appearance considering the sequential nature of video object tracking *without requiring any type of optimization or post-processing* within the tracking pipeline.

3 Method

FCG leverages on the idea that instances of the same object have similar appearance in a temporal neighborhood. It consists on two stages: the first stage generates an initial set of short, but reliable, tracklets, and the second stage fuses these tracklets over time by clustering them in a sequentially-guided hierarchical way, naturally leading to the final object tracks. To generate tracklets and further fuse them, we use the implementation [22] of UPGMA (Unweighted Pair Group Method with Arithmetic mean) [23], which iteratively fuses pairs of clusters, forming a hierarchy. In Figure 1 we show FCG pipeline.

Tracklets

A tracklet T is defined as a set of instances of a specific object being tracked between two time instants. The union of the object tracklets through a whole video sequence will form the final object track. To form the initial set of tracklets, we cluster the object instances present in the frames contained in a non-overlapping temporal window W of size τ . For a video sequence, we will have N possible temporal windows.

Lifted frames

To hierarchically cluster tracklets, we introduce the concept of *lifted frames*, consisting on artificial time instants that contain tracklets instead of individual, discrete, detections. Formally, we define a lifted frame $F_{[n,m]}^i$ as the set of tracklets that contain detections from time $t = n \cdot \tau$ to $t = m \cdot \tau$, where $i \in [0, N]$ corresponds the current level of the hierarchy, and $n \in [0, N - 1]$ and $m \in [1, N]$, $\{n, m \in \mathbb{N} | m > n\}$ correspond to the lifted frame temporal indexes. Note that each tracklet $T^k = F_{[n,m]}^{i,k}$ present in a lifted frame $F_{[n,m]}^i$ has a length $l \in [1, (m - n)\tau]$.

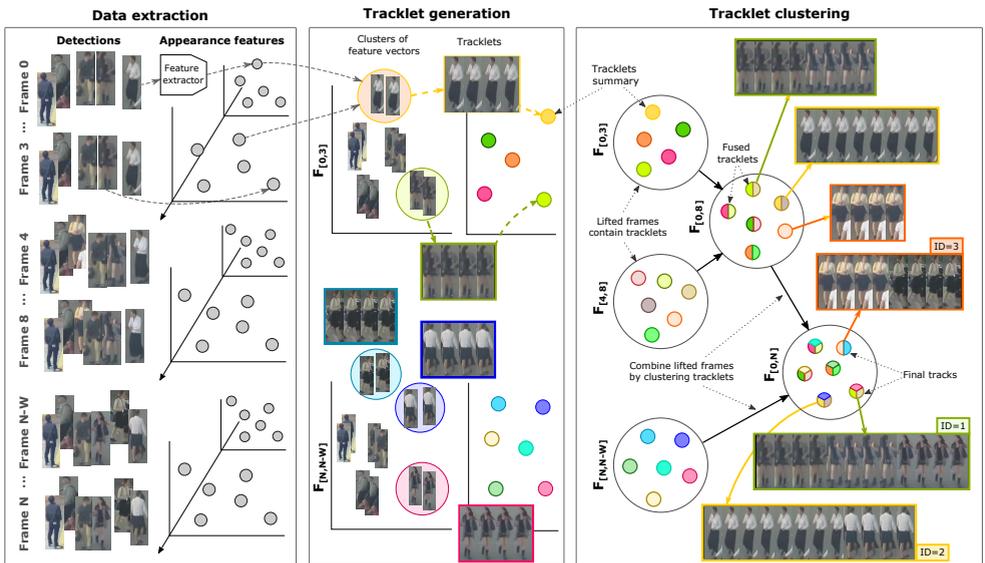


Figure 1: We track objects by clustering their appearance features. We build tracklets to have a robust representation of the objects in a temporal neighborhood, while capturing small variations of the object appearance, and hierarchically cluster the tracklets to form the final tracks.

3.1 Stage zero: extracting the data

As a pre-processing step we first detect all the objects in a video using a given object detector, in this work YOLOX (You Only Look Once X) [13], and extract their appearance features with a feature-extractor network, in this case SBS (Stronger Baseline of Specials) [12].

3.2 First stage: tracklet generation

Once we have all the detections in each frame and their associated appearance feature vectors, we proceed to generate the initial set of tracklets for the objects of interest. These tracklets will be a representation of an object in a temporal neighborhood, capturing small variations of the object appearance that will be key to form the final tracks.

To compare two object instances, O^k and O^p , we use the cosine distance over the corresponding appearance features, \mathbf{h}^k and $\mathbf{h}^p \in \mathbb{R}^{2048}$.

$$d(O^k, O^p) = 1 - \frac{\mathbf{h}^k \top \mathbf{h}^p}{\|\mathbf{h}^k\| \|\mathbf{h}^p\|} \quad (1)$$

To generate tracklets, we cluster object instances within a temporal window W of size τ . We constrain the fusion of objects to different frames, i.e., two objects from the same frame will not be able to be clustered in the same tracklet. In this work, we use $W = 6$, and set the cluster fusion threshold to 0.055.

3.3 Second stage: tracklet clustering

Once the object tracklets are formed for each lifted frame $F_{[n,n+1]}^1$, for $n \in [0, N-1]$, we fuse them in a hierarchical way. Whenever a set of lifted frames are fused, i.e., the tracklets in the lifted frames are clustered together, the current level of the hierarchy i is increased by 1.

To efficiently compare two tracklets, $T^k = F_{[n,m]}^{i,k}$ and $T^q = F_{[u,v]}^{i,q}$, we summarize the collection of appearance feature vectors present in each tracklet, $\mathcal{H}^k, \mathcal{H}^q$, as their element-wise median, and compute the cosine distance between both.

$$d(T^k, T^q) = 1 - \frac{\tilde{\mathbf{h}}^k \top \tilde{\mathbf{h}}^q}{\|\tilde{\mathbf{h}}^k\| \|\tilde{\mathbf{h}}^q\|}, \quad \tilde{\mathbf{h}}^k = \text{Med}(\mathcal{H}^k), \quad \tilde{\mathbf{h}}^q = \text{Med}(\mathcal{H}^q) \quad (2)$$

Generically, we define the fusion of two lifted frames as the clustering of the union of the tracklets contained in each lifted frame.

$$F_{[n,m] \cup [u,v]}^{i+1} = F_{[n,m]}^i \cup F_{[u,v]}^i \quad (3)$$

For the MOT setting, we constrain the fusion to consecutive lifted frames, following the timeline of the video sequence. This is:

$$F_{[n,u]}^{i+1} = F_{[n,m]}^i \cup F_{[m+1,u]}^i, \quad \{n < m < u\} \quad (4)$$

In the ablation study of Sec. 4.2 we further analyse the impact of this weighting.

By iteratively clustering tracklets by fusing lifted frames, we naturally get to a point where all tracklets are defined within a single lifted frame. This happens at $F_{[0,N]}^N$, where the clustering of tracklets is at the top level N of the hierarchy. We assign an ID to each tracklet, becoming the final track for each tracked object.

3.4 Multiple object tracking

We use FCG for the MOT task given the natural way the clustered tracklets form the final tracks of objects. Furthermore, we consider some additions with regards to the MOT challenge that can be embedded into the clustering scheme.

Temporal coherence. It is a common practice in MOT trackers to terminate the tracking of an object if no new detections are associated to that object within a specified time.

We incorporate a similar idea to FCG so that, if two tracklets T^k and T^q are further than a specified temporal distance K_T , given the last detection in T^k, B^k , and the first detection in T^q, B^q , the fusion between the nodes of the hierarchy will be harder, but not impossible, as FCG relies on appearance features to associate tracks, being able to recover from longer-term occlusions. We achieve this by multiplying the distance between two tracklets that are further than K_T by a constant $\lambda_T \geq 1$, making the two tracklets λ_T times harder to be associated, only allowing very similar tracklets to be fused.

$$d'(T^k, T^q) = d(T^k, T^q) \cdot \lambda_T, \quad \lambda_T = \begin{cases} 1 & \Delta t(B^k, B^q) \leq K_T \\ c_T & \Delta t(B^k, B^q) > K_T \end{cases} \quad (5)$$

In this work, we set the constant value to $c_T = 4$, and $K_T = 40$.

Spatial coherence. As many objects share similar features, we also take into account spatial information, making closer objects easier to fuse, while constraining the fusion of farther objects.

For two tracklets T^k and T^q , we compute both the IoU (Intersection over Union) and a L2 distance normalized w.r.t. bounding box size between the last detection in T^k , B^k , and the first detection in T^q , B^q . We multiply the original distance between tracklets $d(T^k, T^q)$ by λ_C and λ_F , corresponding to close and far objects respectively.

$$d'(T^k, T^q) = d(T^k, T^q) \cdot \lambda_C \cdot \lambda_F \quad (6)$$

$$\lambda_C = \min(1, d_{\text{IoU}} + \text{off}), \quad \lambda_F = \begin{cases} 1 & d_{\text{box}}(B^k, B^q) \leq K_F \\ c_F & d_{\text{box}}(B^k, B^q) > K_F \end{cases}$$

We set $\text{off} = 0.15$ to avoid multiplying by 0 the distance when two objects are completely overlapped, $c_F = 2$ to make spatially far detections twice as hard to be fused, and $K_F = 2$ as the normalized distance between bounding boxes. Note that $\lambda_C \in [\text{off}, 1]$, easing the fusion between tracklets, and $\lambda_C \in \{1, c_F\}$, making them harder to fuse, but not impossible. We convert the IoU between B^k and B^q to the distance d_{IoU} ,

$$d_{\text{IoU}}(B^k, B^q) = 1 - \frac{B^k \cap B^q}{B^k \cup B^q} \quad (7)$$

and compute the normalised displacement between B^k and B^q as the distance d_{box} between their extreme points (left, top, right, bottom),

$$d_{\text{box}}(B^k, B^q) = \frac{d_1 + d_2}{2} \quad (8)$$

where d_1 corresponds to the distance between the left-top corners of the bounding boxes B^k and B^q , and d_2 to the distance between the right-bottom corners of the bounding boxes B^k and B^q . Both distances are normalized by the average of the bounding boxes width and height to have in consideration the size of the bounding box when computing the displacement.

$$d_1 = \sqrt{\left(\frac{B_{\text{left}}^k - B_{\text{left}}^q}{\frac{B_{\text{width}}^k + B_{\text{width}}^q}{2}} \right)^2 + \left(\frac{B_{\text{top}}^k - B_{\text{top}}^q}{\frac{B_{\text{height}}^k + B_{\text{height}}^q}{2}} \right)^2}$$

$$d_2 = \sqrt{\left(\frac{B_{\text{right}}^k - B_{\text{right}}^q}{\frac{B_{\text{width}}^k + B_{\text{width}}^q}{2}} \right)^2 + \left(\frac{B_{\text{bot}}^k - B_{\text{bot}}^q}{\frac{B_{\text{height}}^k + B_{\text{height}}^q}{2}} \right)^2}$$

With this, the distance between tracklets T^k and T^q can be weighted by both temporal and spatial priors.

$$d'(T^k, T^q) = d(T^k, T^q) \cdot \lambda_T \cdot \lambda_C \cdot \lambda_F \quad (9)$$

Motion estimation. We integrate a simple motion estimation, consisting on a constant velocity assumption, to help in the spatial coherence. This motion is calculated as the difference between the previous $t - 1$ and current t positions of a bounding box B existing in a tracklet, and applied to displace B to the next time step $t + 1$.

4 Experimental evaluation

Datasets. Experiments are performed in three different MOT datasets: MOT17 [24], MOT20 [10], and DanceTrack [5]. MOT17 and MOT20 address the problem of tracking pedestrians in the wild, while DanceTrack focuses on a group of dancers on a stage. Following the work in MOT17, MOT20 has a much larger density of pedestrians in the scene, longer sequences, and the camera is set further away. In DanceTrack, where the setting is much more artistic, the main difficulties consist on the amount of very similarly dressed dancers, appearance changes due to zooms or moving cameras, the amount of occlusions and grouping, and the diverse motion, which, in contrast with MOT17 and MOT20, can be much more random.

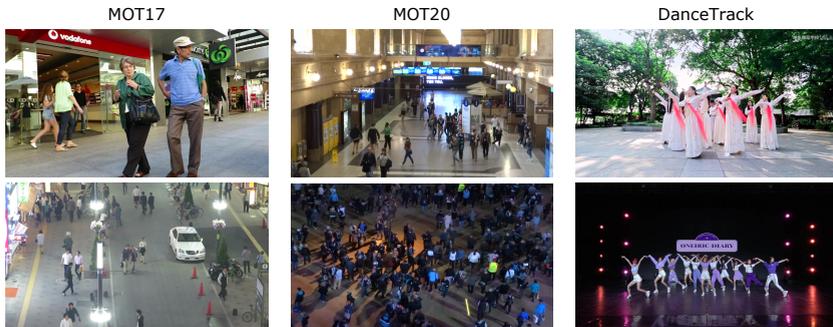


Figure 2: We test FCG on three different benchmarks. MOT17 and MOT20 focus on tracking pedestrians in the wild, and DanceTrack on tracking dancers on a stage.

Metrics. To evaluate our method, we use the main established metrics for MOT, which are HOTA [20], built as the combination of DetA (Detection Accuracy score) and AssA (Association Accuracy score), MOTA [3], and IDF1 [29]. For all metrics, higher is better.

Implementation details. To detect the objects in the scene we use YOLOX [13], pre-trained with COCO [18] and fine-tuned on each dataset. In this work, we use a detector threshold of 0.7. For the extraction of appearance features we use SBS [14] trained with Market1501 [53]. We conducted experiments using different detectors and feature extraction networks in Sec. 4.2. For all benchmarks, we build the initial tracklets from detections within a sliding window of 6 frames.

If not stated otherwise, for MOT17 and MOT20, where pedestrians may leave the scene to never reappear, we use the temporal constraint and motion estimation introduced in Sec. 3.4. For DanceTrack, where is usual for dancers to reappear in the scene, we don't use neither the temporal constraint nor motion estimation. For all benchmarks, we use the spatial coherence constraints, also introduced in Sec. 3.4, where two tracklets are easier or more difficult to fuse depending on the spatial distance between their corresponding detections.

4.1 Comparison to the state-of-the-art

In Table 1, we report results of the performance of FCG on the test set of each respective benchmark, and compare it to other state-of-the-art methods. All method use their own, private, detectors. The results for each dataset are presented as the overall results for all sequences.

Methods	MOT17 [14]					MOT20 [10]					DanceTrack [15]				
	HOTA	DetA	AssA	MOTA	IDF1	HOTA	DetA	AssA	MOTA	IDF1	HOTA	DetA	AssA	MOTA	IDF1
CenterTrack [14]	52.2	53.8	51.0	67.8	64.7	-	-	-	-	-	41.8	78.1	22.6	86.8	35.7
FairMOT [14]	59.3	60.9	58.0	73.7	72.3	54.6	54.7	54.7	61.8	67.3	39.7	66.7	23.8	82.2	40.8
QDTrack [15]	53.9	55.6	52.7	68.7	66.3	-	-	-	-	-	45.7	72.1	29.2	83.0	44.8
TransTrack [15]	54.1	61.6	47.9	75.2	63.5	48.9	53.3	45.2	65.0	59.4	45.5	75.9	27.5	88.4	45.2
TraDes [16]	52.7	55.2	50.8	69.1	63.9	-	-	-	-	-	43.3	74.5	25.4	86.2	41.2
MOTR [15]	55.1	56.2	54.2	67.4	67.0	43.6	51.8	37.0	59.1	49.9	48.4	71.8	32.7	79.2	46.1
Trackformer [15]	57.3	60.9	54.1	74.1	68.0	54.7	56.7	53.0	68.6	65.7	-	-	-	-	-
ReMOT [15]	59.7	62.8	57.1	77.0	72.0	61.2	63.9	58.7	77.4	73.1	-	-	-	-	-
CrowdTrack [15]	60.3	61.5	59.3	75.6	73.6	55.0	57.7	52.6	70.7	68.2	-	-	-	-	-
MAA [15]	62.0	64.2	60.2	79.4	75.9	57.3	59.7	55.1	73.9	71.2	-	-	-	-	-
BYTE [15]	63.1	64.5	62.0	80.3	77.3	61.3	63.4	59.6	77.8	75.2	47.7	71.0	32.1	89.6	53.9
FCG (Ours)	62.6	62.2	63.4	76.7	77.7	57.3	56.7	58.1	68.0	69.7	48.7	79.8	29.9	89.9	46.5

Table 1: Comparison between our method, *FCG*, to the state-of-the-art on the test set of MOT17, MOT20, and DanceTrack. Results of all methods over private detections.

By comparing to other state-of-the-art results, *FCG* shows that similar tracking performance can be achieved by using appearance between objects as the basic source of information to associate objects over time, with a much simpler technique and without the need of any kind of optimization or post-processing when performing the tracking. This raises the intuition that having strong object information, in this case detections and person re-identification features, can be as important as the tracking method itself.

Overall, *FCG* presents competitive results compared to other methods on MOT17 and MOT20 datasets, achieving state-of-the-art results for the DanceTrack benchmark.

4.2 Analysis

Design choices for tracklet generation

In Figure 3 we highlight the effect in performance when choosing different sizes of the temporal window W to build the initial tracklets. For all benchmarks, there is a similar performance when forming the tracklets within temporal windows of size $W = 2$ to $W = 6$ frames, starting to decrease from $W = 7$ for the DanceTrack dataset, and from $W = 15$ for MOT17 and MOT20. This is due to the appearance change in the person being tracked when considering longer tracklets, as many false positives can be added to the tracklet. For all benchmarks, tracking by directly clustering individual detections was computationally unfeasible. The same applies to window sizes of $W = 2$ and $W = 3$ in MOT20, where large amount of elements to be tracked are present.

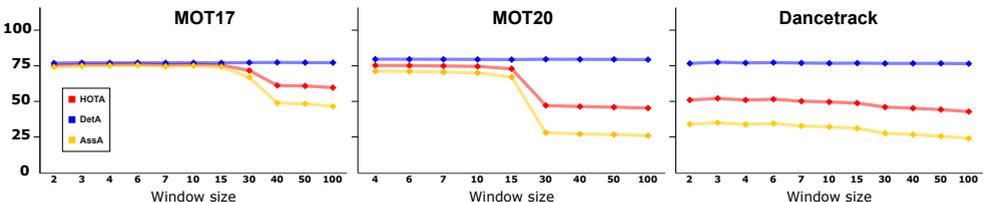


Figure 3: *FCG* tracking performance for tracklet construction over different window sizes.

C	S	M	T	MOT17					MOT20					DanceTrack				
				HOTA	DetA	AssA	MOTA	IDF1	HOTA	DetA	AssA	MOTA	IDF1	HOTA	DetA	AssA	MOTA	IDF1
				61.6	77.2	49.6	86.6	64.1	33.6	79.5	14.3	91.4	27.7	46.2	77.8	27.6	82.8	41.6
✓				67.6	77.2	59.6	86.7	71.9	43.3	79.5	23.7	91.8	40.1	41.8	77.9	22.5	82.8	37.0
✓	✓			71.9	77.3	67.2	87.3	77.4	58.4	79.8	42.9	92.4	57.0	50.0	77.2	32.5	87.5	51.8
✓	✓	✓		72.4	77.5	68.0	87.3	78.2	58.2	79.8	42.4	92.4	56.2	48.9	77.2	31.1	87.4	50.8
✓			✓	72.3	77.2	68.0	86.7	79.0	57.6	79.5	41.9	91.8	62.5	38.9	77.8	19.6	82.9	33.4
✓	✓		✓	72.4	77.5	68.0	87.3	78.2	74.8	79.8	70.1	92.4	85.4	48.9	77.2	31.1	87.4	50.8
✓	✓	✓	✓	76.3	77.5	75.4	87.3	85.1	74.6	79.8	69.8	92.4	85.2	49.1	77.3	31.4	87.5	48.9

Table 2: Ablations on FCG. C stands for Consecutive clustering of tracklets, S for Spatial coherence, M for Motion estimation, and T for Temporal coherence.

Ablation study

We perform a set of ablations to FCG, considering the enhancements presented in Sec. 3.4, and report the results in Table 2. Throughout the ablation, detection-based scores (DetA and MOTA) are very similar. This is due to be using all detections, without adding or discarding any, being the small variations related to ID switching. For this, we only highlight the overall score, HOTA, and the scores related to the association, AssA and IDF1.

First, comparing FCG baselines (first two rows in Table 2), we highlight the fact that following a consecutive (C) order when clustering tracklets is important for MOT17 and MOT20, given the sequential nature of the pedestrian movement, but not for DanceTrack, due to the random dance formation changes, combined with the resemblance of the dancers, producing many ID switches between tracks during transitions.

Spatial and temporal enhancements do not apply to the non-consecutive baseline due to the need of spatio-temporal coherence when comparing the detections in tracklets. To the FCG baseline (second row in Table 2), considering spatial (S) coherence is key due to the amount of frames where the same object overlaps within a video sequence. Combining spatial and a simple Motion (M) estimation, there is a small improvement in MOT17, and small decay for MOT20 and DanceTrack, probably due to the simplicity of the motion model.

On the other hand, temporal (T) coherence provides a considerable improvement in the pedestrian datasets, MOT17 and MOT20, and a small decay for DanceTrack. This is also due to the nature of the datasets as, in MOT17 and MOT20, pedestrians may leave the scene to never reappear, while in DanceTrack dancers are usually always present in the video, even if they are occluded for a long time. We highlight the fact that MOT20 has a staggering increase in 31.5 points in HOTA score when using spatial (S) and temporal (T) coherence during the clustering of tracklets. Weighting the object associations with spatio-temporal priors is specially needed in this dataset due to the high density of pedestrians in the scene, making the correct assignment difficult, combined with their small size due to camera positioning.

Low fps setting

A problem that current state-of-the-art trackers face comes when using them in low frames per second (fps) frameworks. We show this in Figure 4, where we tested three MOT trackers, BYTE [50], Trackformer [23], and FCG -this work-, in the MOT17 [24] training set. As the videos have different fps, we sub-sampled them by using a sampling ratio that ranges from 2, corresponding to half the original video fps -e.g. from 30 fps to 15 fps-, to 30, which corresponds to 1/30 of the original video -e.g. from 30 fps to 1 fps-.

From the results, FCG seems to be more robust to lower fps settings than BYTE or Trackformer. With a performance of 62.1 points of HOTA score on the most extreme sampling, 1 frame every 30 of the original video, FCG outperforms the other trackers, BYTE (44.6) and

Trackformer (50.3). This comes as a side effect of the object association by using appearance features, and only using spatial and temporal priors as a weighting factor.

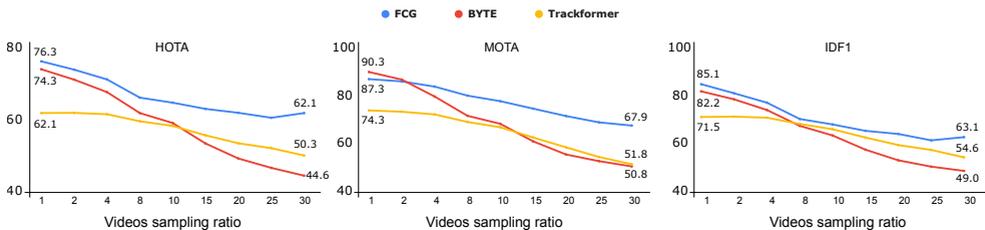


Figure 4: Low fps video setting experiment for state-of-the-art MOT trackers.

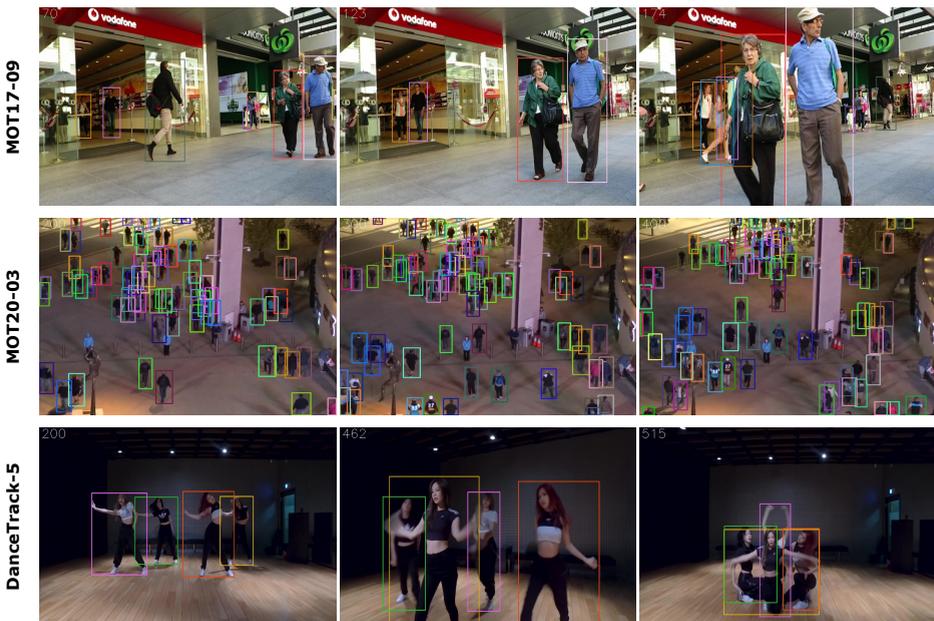


Figure 5: Qualitative results of FCG tracking for three sequences of the tested benchmarks, MOT17, MOT20 and DanceTrack. Same color corresponds to the same ID.

5 Conclusions

In this paper we explore multiple object tracking using object appearance as the principal element for the object association step, and present FCG, a MOT tracker that leverages on the idea that a specific object is similar in a temporal neighborhood to generate tracklets, and cluster them in a hierarchical way to form the final object tracks. While being a much simpler method than other current MOT trackers, FCG achieves state-of-the-art performance on the DanceTrack benchmark, and presents competitive results in both MOT17 and MOT20 datasets.

References

- [1] Jerome Berclaz, Francois Fleuret, and Pascal Fua. Robust people tracking with global trajectory optimization. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 1, pages 744–750. IEEE, 2006.
- [2] Philipp Bergmann, Tim Meinhardt, and Laura Leal-Taixe. Tracking without bells and whistles. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 941–951, 2019.
- [3] Keni Bernardin and Rainer Stiefelhagen. Evaluating multiple object tracking performance: the clear mot metrics. *EURASIP Journal on Image and Video Processing*, 2008: 1–10, 2008.
- [4] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Upcroft. Simple online and realtime tracking. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 3464–3468. IEEE, 2016.
- [5] Guillem Brasó and Laura Leal-Taixé. Learning a neural solver for multiple object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6247–6257, 2020.
- [6] Long Chen, Haizhou Ai, Zijie Zhuang, and Chong Shang. Real-time multiple people tracking with deeply learned candidate selection and person re-identification. In *2018 IEEE international conference on multimedia and expo (ICME)*, pages 1–6. IEEE, 2018.
- [7] Long Chen, Haizhou Ai, Rui Chen, and Zijie Zhuang. Aggregate tracklet appearance features for multi-object tracking. *IEEE Signal Processing Letters*, 26(11):1613–1617, 2019.
- [8] Peng Chu and Haibin Ling. Famnet: Joint learning of feature, affinity and multi-dimensional assignment for online multiple object tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6172–6181, 2019.
- [9] Peng Dai, Renliang Weng, Wongun Choi, Changshui Zhang, Zhangping He, and Wei Ding. Learning a proposal classifier for multiple object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2443–2452, 2021.
- [10] Patrick Dendorfer, Hamid Rezaatofighi, Anton Milan, Javen Shi, Daniel Cremers, Ian Reid, Stefan Roth, Konrad Schindler, and Laura Leal-Taixé. Mot20: A benchmark for multi object tracking in crowded scenes. *arXiv preprint arXiv:2003.09003*, 2020.
- [11] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6569–6578, 2019.
- [12] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Detect to track and track to detect. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3038–3046, 2017.

- [13] Zheng Ge, Songtao Liu, Feng Wang, Zeming Li, and Jian Sun. Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*, 2021.
- [14] Lingxiao He, Xingyu Liao, Wu Liu, Xinchen Liu, Peng Cheng, and Tao Mei. Fastreid: A pytorch toolbox for general instance re-identification. *arXiv preprint arXiv:2006.02631*, 2020.
- [15] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017.
- [16] Andrea Hornakova, Roberto Henschel, Bodo Rosenhahn, and Paul Swoboda. Lifted disjoint paths with application in multiple object tracking. In *International Conference on Machine Learning*, pages 4364–4375. PMLR, 2020.
- [17] Chang Huang, Bo Wu, and Ramakant Nevatia. Robust object tracking by hierarchical association of detection responses. In *European Conference on Computer Vision*, pages 788–801. Springer, 2008.
- [18] 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.
- [19] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [20] Jonathon Luiten, Aljosa Osep, Patrick Dendorfer, Philip Torr, Andreas Geiger, Laura Leal-Taixe, and Bastian Leibe. Hota: A higher order metric for evaluating multi-object tracking. *International Journal of Computer Vision*, 2020.
- [21] Hao Luo, Youzhi Gu, Xingyu Liao, Shenqi Lai, and Wei Jiang. Bag of tricks and a strong baseline for deep person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019.
- [22] Liqian Ma, Siyu Tang, Michael J Black, and Luc Van Gool. Customized multi-person tracker. In *Asian conference on computer vision*, pages 612–628. Springer, 2018.
- [23] Tim Meinhardt, Alexander Kirillov, Laura Leal-Taixe, and Christoph Feichtenhofer. Trackformer: Multi-object tracking with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8844–8854, 2022.
- [24] A. Milan, L. Leal-Taixé, I. Reid, S. Roth, and K. Schindler. MOT16: A benchmark for multi-object tracking. *arXiv:1603.00831 [cs]*, March 2016. URL <http://arxiv.org/abs/1603.00831>. arXiv: 1603.00831.
- [25] Jiangmiao Pang, Linlu Qiu, Xia Li, Haofeng Chen, Qi Li, Trevor Darrell, and Fisher Yu. Quasi-dense similarity learning for multiple object tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 164–173, 2021.

- [26] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017.
- [27] 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.
- [28] Ergys Ristani and Carlo Tomasi. Features for multi-target multi-camera tracking and re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6036–6046, 2018.
- [29] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *European Conference on Computer Vision*, pages 17–35. Springer, 2016.
- [30] Peter HA Sneath and Robert R Sokal. Unweighted pair group method with arithmetic mean. *Numerical Taxonomy*, pages 230–234, 1973.
- [31] Daniel Stadler and Jürgen Beyerer. Multi-pedestrian tracking with clusters. In *2021 17th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–10. IEEE, 2021.
- [32] Daniel Stadler and Jürgen Beyerer. On the performance of crowd-specific detectors in multi-pedestrian tracking. In *2021 17th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–12. IEEE, 2021.
- [33] Daniel Stadler and Jürgen Beyerer. Modelling ambiguous assignments for multi-person tracking in crowds. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 133–142, 2022.
- [34] Peize Sun, Jinkun Cao, Yi Jiang, Rufeng Zhang, Enze Xie, Zehuan Yuan, Changhu Wang, and Ping Luo. Transtrack: Multiple object tracking with transformer. *arXiv preprint arXiv:2012.15460*, 2020.
- [35] Peize Sun, Jinkun Cao, Yi Jiang, Zehuan Yuan, Song Bai, Kris Kitani, and Ping Luo. Dancetrack: Multi-object tracking in uniform appearance and diverse motion. *arXiv preprint arXiv:2111.14690*, 2021.
- [36] Pavel Tokmakov, Jie Li, Wolfram Burgard, and Adrien Gaidon. Learning to track with object permanence. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10860–10869, 2021.
- [37] Ibrahim Saygin Topkaya, Hakan Erdogan, and Fatih Porikli. Tracklet clustering for robust multiple object tracking using distance dependent chinese restaurant processes. *Signal, Image and Video Processing*, 10(5):795–802, 2016.
- [38] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert

- Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.
- [39] 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.
- [40] Guanshuo Wang, Yufeng Yuan, Xiong Chen, Jiwei Li, and Xi Zhou. Learning discriminative features with multiple granularities for person re-identification. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 274–282, 2018.
- [41] Qiang Wang, Yun Zheng, Pan Pan, and Yinghui Xu. Multiple object tracking with correlation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3876–3886, 2021.
- [42] Shuai Wang, Hao Sheng, Yang Zhang, Yubin Wu, and Zhang Xiong. A general recurrent tracking framework without real data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13219–13228, 2021.
- [43] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple online and realtime tracking with a deep association metric. In *2017 IEEE international conference on image processing (ICIP)*, pages 3645–3649. IEEE, 2017.
- [44] Jialian Wu, Jiale Cao, Liangchen Song, Yu Wang, Ming Yang, and Junsong Yuan. Track to detect and segment: An online multi-object tracker. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12352–12361, 2021.
- [45] Yu Wu, Yutian Lin, Xuanyi Dong, Yan Yan, Wanli Ouyang, and Yi Yang. Exploit the unknown gradually: One-shot video-based person re-identification by stepwise learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5177–5186, 2018.
- [46] Fan Yang, Xin Chang, Sakriani Sakti, Yang Wu, and Satoshi Nakamura. Remot: A model-agnostic refinement for multiple object tracking. *Image and Vision Computing*, 106:104091, 2021.
- [47] En Yu, Zhuoling Li, Shoudong Han, and Hongwei Wang. Relationtrack: Relation-aware multiple object tracking with decoupled representation. *arXiv preprint arXiv:2105.04322*, 2021.
- [48] Fangao Zeng, Bin Dong, Tiancai Wang, Xiangyu Zhang, and Yichen Wei. Motr: End-to-end multiple-object tracking with transformer. *arXiv preprint arXiv:2105.03247*, 2021.
- [49] Yang Zhang, Hao Sheng, Yubin Wu, Shuai Wang, Weifeng Lyu, Wei Ke, and Zhang Xiong. Long-term tracking with deep tracklet association. *IEEE Transactions on Image Processing*, 29:6694–6706, 2020.
- [50] Yifu Zhang, Chunyu Wang, Xinggong Wang, Wenjun Zeng, and Wenyu Liu. Fairmot: On the fairness of detection and re-identification in multiple object tracking. *International Journal of Computer Vision*, 129(11):3069–3087, 2021.

- [51] Yifu Zhang, Peize Sun, Yi Jiang, Dongdong Yu, Fucheng Weng, Zehuan Yuan, Ping Luo, Wenyu Liu, and Xinggang Wang. Bytetrack: Multi-object tracking by associating every detection box. *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022.
- [52] Zhimeng Zhang, Jianan Wu, Xuan Zhang, and Chi Zhang. Multi-target, multi-camera tracking by hierarchical clustering: Recent progress on dukemtmc project. *arXiv preprint arXiv:1712.09531*, 2017.
- [53] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE international conference on computer vision*, pages 1116–1124, 2015.
- [54] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. *arXiv preprint arXiv:2004.01177*, 2020.