

This CVPR workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

TripletTrack: 3D Object Tracking using Triplet Embeddings and LSTM

Nicola Marinello¹, Marc Proesmans^{1,3}, Luc Van Gool^{1,2,3}, ¹ KU Leuven/ESAT-PSI, ² ETH Zurich/CVL, ³ TRACE vzw

{nicola.marinello,marc.proesmans,luc.vangool}@esat.kuleuven.be

Abstract

3D object tracking is a critical task in autonomous driving systems. It plays an essential role for the system's awareness about the surrounding environment. At the same time there is an increasing interest in algorithms for autonomous cars that solely rely on inexpensive sensors, such as cameras. In this paper we investigate the use of triplet embeddings in combination with motion representations for 3D object tracking. We start from an off-the-shelf 3D object detector, and apply a tracking mechanism where objects are matched by an affinity score computed on local object feature embeddings and motion descriptors. The feature embeddings are trained to include information about the visual appearance and monocular 3D object characteristics, while motion descriptors provide a strong representation of object trajectories. We will show that our approach effectively re-identifies objects, and also behaves reliably and accurately in case of occlusions, missed detections and can detect re-appearance across different field of views. Experimental evaluation shows that our approach outperforms state-of-the-art on nuScenes by a large margin. We also obtain competitive results on KITTI.

1. Introduction

3D Multi-Object Tracking (MOT) is a crucial task for autonomous cars since it plays a central role in providing surround situational awareness. At the same time there is an increasing interest in exploring perception algorithms that solely rely on cameras since they are relatively inexpensive, compared to other sensors such as LiDARs and radars. Nevertheless camera-only 3D MOT methods so far have not received as much attention as 2D methods [28]. Recent developments in monocular 3D object detection open new possibilities for 3D object tracking. The number of 3D object detectors in fact has been increasing recently, demonstrating promising results [4,7,8,16,29,41,42,49,50]. These detectors are capable of estimating the 3D location in camera frame coordinates of an object, along with its size, orientation and category. 3D object detectors allow to work



Figure 1. Illustration of re-identification with our tracking pipeline. The fuchsia colored car on the left is being occluded for a few frames by the blue car and the ID is picked up again.

in 3D and thus benefit from the additional information incorporated in the detections. In addition 3D MOT permits to follow each object in the real world space, a task that is simply not possible with 2D MOT.

3D tracking essentially consists of two sub-tasks: 3D detection and re-identification. 3D detection is the operation of recognizing objects in images and estimating their position, dimension and orientation. Re-identification allows to consistently follow each of them and to keep track of their position history. Tracked objects have to be associated with a unique ID and the same ID has to be correctly re-assigned to the same object each time it is detected. Furthermore, the tracking system has to be able to re-assign the same ID to the same object even after a temporary occlusion, a missed detection, or when the object reappears in a different camera.

In this work we introduce an online camera-only tracking mechanism that follows the Tracking-by-Detection paradigm. We operate under the traditional framework that consists of a detection step and a data association step, focusing on the design of the second one. Similarly to previous 2D approaches [20,30,36] we investigate how to model the motion of the objects in 3D space with a single neural network and how to extract powerful view-independent object visual features [30]. To extract visual features we borrow techniques from the re-identification [15, 25] and face recognition [38] literature. Our approach applies an independent detector to all images frames and it tries to link detected objects to trajectories of tracked objects, by exploiting motion and appearance information. Appearance features are condensed in triplet embeddings and the objects motion is modeled through a Long Short-Term Memory (LSTM). Identity association is then finally solved with the Hungarian algorithm [24] based on an affinity matrix that summarizes affinity scores between previously tracked objects and detected objects. The method can operate with any camera setup and with any off-the-shelf 3D object detector. Through experimental evaluation we show that a traditional approach consisting of a detection step and a data association step can be highly competitive in terms of tracking performance.

Some state-of-the-art methods perform joint detection and tracking but at the expense of weaker objects representations or lack of synergy between motion and appearance cues. Some also manually design similarity metrics for the data association step, not fully exploiting the power of Neural Networks [6,17,18,62]. We instead advocate for a data-driven method that focuses on strong appearance and motion representations. Exploiting both cues makes it robust in case of temporal occlusions and missed detections, being able to pick up again objects identities. An example is shown in Fig. 1. Our main contributions can be summarized as follows:

- We propose a camera-only tracking method that can be used with any off-the-shelf 3D object detector and any camera setup.
- We introduce an explicit embedding model that uses 2D and 3D appearance cues jointly with 3D motion.
- We set a new state-of-the-art on the nuScenes camera only tracking benchmark and we obtain competitive results on KITTI.

2. Related Work

The 3D tracking task aims at following each detected object and to keep track of their trajectories in the real world space. As for 2D MOT current research largely focuses on the tracking-by-detection paradigm [6, 32, 52, 57, 60], including state-of-the-art trackers [17, 18, 34, 54]. [52] sets a simple baseline using a combination of a 3D Kalman filter and the Hungarian algorithm [24]. QD-3DT [17, 18]

proposes an unified framework for joint detection, appearance feature extraction and tracking. It combines appearance and predicted motion information to match detections with tracklets. DEFT [6] also advocates for an unified detection and tracking framework while using a LSTM that captures motion constraints to filter out physically implausible matches. [1] integrates a Kalman filter for state estimation and motion prediction along with appearance features to associate detections with tracked objects. [47] makes use of the triplet loss to learn embeddings that improve data association in a self-supervised manner. On the other side others try to track objects without any appearance information [46]. [37] pairs up an object detector and a recursive filter that predicts and refine object positions and kinematical properties to track objects in 3D space. In [62] they keep track of objects by consistently following their centers in 2D space, however, they trade the ability to re-identify identities after long occlusions with speed. [45] proposes an approach in 2D that can directly be extended to 3D to hallucinate trajectories of fully occluded objects by extending the CenterTrack model with a recurrent memory module. [48] performs tracking given point cloud data with and end-to-end network that detects objects and directly assigns tracking IDs. Graph Neural Networks have also been used in the context of 3D MOT [53, 58]. OGR3MOT [58] models data association, and track management through a graph representation, while GNN3DMOT [53] extracts appearance and motion features from 2D and 3D spaces and with a Graph Neural Network (GNN) that models their interaction. [39] proposes shape pose and motion costs to improve data association and tracking. [26] presents a tracking pipeline that exploits 2D and 3D world space motion consistency to improve long-term tracking under occlusions. [2] proposes a confidence-based method for initialization and termination of tracklets. State-of-the-art performance can also be achieved without using the tracking-by-detection paradigm [55]. Finally, there is an entire line of works that take advantage of multiple sensors by combining their information flows [9, 19, 22, 40, 59, 60].

3. Method

3.1. TripletTrack

Let's formalize the re-identification task: we can define the state of a tracked object as $T_i = (x, y, z, \theta, w, h, l, \text{ID}, \mathbf{p})$, the state of a detected object $D_j = (x, y, z, \theta, w, h, l)$ where x, y and z are the center of the 3D bounding box in global coordinates, θ is the yaw angle, w, h, and l are the width, height and length respectively, ID is a unique id assigned to a tracked object and \mathbf{p} is the object path history. We can also define the set of tracked objects at time t - 1 as $T^{t-1} = \{T_1^{t-1}, T_2^{t-1}, ..., T_m^{t-1}\}$ and the set of detected objects $D^t = \{D_1^t, D_2^t, ..., D_n^t\}$.



Figure 2. **Proposed pipeline.** At every time frame t all pairs tracked-detected objects T_i^{t-1} , D_j^t are compared to compute an affinity matrix. The computed appearance embedding of a detected object is concatenated with the appearance embedding of a tracked object. The motion representation of a tracked object is concatenated with BEV 3D position of the detected object. The appearance and motion components are fed to an affinity network that computes the affinity score for a specific pair of objects. The affinity matrix is later used by the Hungarian algorithm to associate objects and re-assign the same ID. Lost objects are repeatedly compared with new detections up to a maximum amount of time. If they have not been re-identified after the maximum amount of time they are discarded. Unmatched detections are considered new objects.

The objective of a re-identification algorithm is to take as input the set of previously tracked objects T^{t-1} and the new detections set D^t and assign to the detected objects the same ID that they had in the previous frames. If new objects enter into the scene they have to be assigned with a new ID. Human drivers have the innate capacity to perform such a task by implicitly using two cues: *appearance* and *motion*. We therefore take inspiration from this to design an algorithm that exploits these cues.

Motion. Motion can be learned with a Neural Network. The purpose of learning a motion model is to be able to extract a compact representation of an object trajectory in space. A well suited model for this purpose is the LSTM. This is a particular type of Recurrent Neural Network (RNN) which can process a series of inputs and identify patterns within the input motion sequence. By training the model with objects' trajectories we believe that it can learn how objects typically move in space providing a meaningful representation. Our ablation study in Section 5.1 shows that this can be attained. In this work we aim to train a single LSTM that can be used to re-identify objects of any class. Since we want to model the objects motion in space, we believe that the most straightforward way is to feed the LSTM with the 3D position projected on the ground plane, which we will call "BEV 3D position", expressed in global coordinates. This can be achieved by using a 3D object detector and a GPS/IMU that determines the ego pose with respect to a global reference system.

2D Appearance. The other key attribute is the visual appearance of an object, such as an objects shape, or its reflectance properties observed through a camera or set of cameras. Appearance information is especially desirable in case of long lasting occlusions of dynamic objects. When an object is not in sight for a long time, the motion analysis may struggle to re-identify it. In these cases, appearance information provides a strong clue that allows to correctly assign the ID to the object that reappeared. In order to effectively extract robust object visual features, we employed a dedicated Convolutional Neural Network (CNN) that computes triplet embeddings. Ideally, the visual characteristics of an object should be invariant to the view point, scale, illumination, partial occlusions, truncation and background clutter.

3D Appearance. Each detection out of a 3D object detector comes with the estimated 3D bounding box size. This information can be regarded as an additional appearance feature and exploited to improve the embeddings discriminative ability. We therefore concatenate the estimated 3D bounding box size of the corresponding object at the last layer of the CNN.

Affinity network. Our re-identification method extracts appearance features with the CNN, and motion features with the LSTM from tracked and detected objects. A tiny affinity network is then responsible to combine them and to determine their similarity. The re-identification pipeline is shown in Fig. 2.

3.2. Training

The proposed pipeline is trained in two stages, and during the entire training process we make use of the annotations provided with the dataset. This makes the training process completely independent from the object detector that will be used during inference time.

In the first stage we train the CNN to learn objects embeddings. We use the *triplet loss*, first proposed in [38]. The triplet loss has been extensively used in the past, especially in the context of re-identification for surveillance purposes [15,25,43,51]. It is computed over a triplet of images: an *anchor*, a *positive* and a *negative*; the anchor and the positive images show the same object, while the negative shows an object with a different identity. It learns an embedding space by forcing embeddings of the same identity to be closer to each other than an embedding from any other identity by at least a margin m. Since the aforementioned datasets provide unique identity annotations for each object, they can be used as ground truth labels to train an embedding model with the triplet loss. The triplet loss is expressed as follows:

$$\mathcal{L}_{\rm tri} = \sum_{a,p,n} \left[m + D_{a,p} - D_{a,n} \right]_{+}$$
(1)

where D is the embedding distance. As input to the network we provide crops of projected ground truth 3D bounding boxes and the corresponding 3D bounding box sizes. A triplet of crops can contain objects from any scene at any frame in the dataset. We supervise the network with the triplet loss and the ground truth objects identity labels from the dataset. However, particularly in automotive datasets that provide continuous camera recordings, the similarity between the same object in succeeding frames is very high, hence it is crucial to train the CNN on the most informative triplets. Therefore we followed the online mining strategy proposed by [15]. The training batches are built in the following way: P identities are randomly sampled from the dataset and then K images of each identity to build batches of PK images. Over such a batch, a maximum of PK(PK-K)(K-1) triplets combinations can potentially be sampled. In order to sample the most informative triplets we use the *batch hard* online mining [15] to compute the triplet loss over the hardest triplets, resulting in PK triplet contributions in the loss computation.

In the second stage we jointly train the entire pipeline that includes the LSTM, the affinity network and the CNN. However the CNN weights are not updated in this stage. The pipeline is trained in batches and each batch element corresponds to a specific frame of a specific scene in the dataset. A single batch contains frames from any scene and in random order. For each frame a maximum amount of tracked objects and detected objects is randomly sampled. With a tracked object we refer to any object whose most recent annotation in the scene was there at most B frames in the past. With a detected object we refer to any annotated object in the current frame. We supervise the pipeline using the binary cross entropy loss and a ground truth affinity matrix. Let M denote the number of sampled tracked objects and N the number of sampled detected objects in a specific frame. We define a ground truth affinity matrix $Y \in \{0,1\}^{M \times N}$ using their identity labels. For every pair (T_i^{t-1}, D_i^t) the pipeline determines if they correspond to the same identity by computing a similarity score. The affinity network receives as input the appearance embedding¹ and the position history representation of T_i^{t-1} , along with the appearance embedding and the BEV 3D position of D_i^t . These cues are combined and used to compute the estimated affinity $\hat{y}_{ij} \in [0, 1]$. By repeating this process for every pair (T_i^{t-1}, D_i^t) we obtain an affinity matrix $\hat{Y} \in [0, 1]^{M \times N}$. To produce the feature embeddings the CNN is fed with the crop of the projected 3D bounding box and the 3D bounding box size. While the LSTM processes the BEV 3D position history of the tracked object. The binary cross entropy loss is then computed over corresponding elements of the two matrices as follows:

$$\mathcal{L}_e = \sum_{i}^{M} \sum_{j}^{N} (-p \, y_{ij} \log(\hat{y_{ij}}) + (1 - y_{ij}) \log(1 - \hat{y_{ij}}))$$
(2)

where $y_{ij} \in Y$, $\hat{y}_{ij} \in \hat{Y}$ and p is the weight for the positive pairs. The ground truth y_{ij} corresponds to 1 when the two identity labels are the same, 0 otherwise. Moreover, at each frame, the number of positive pairs is much lower than the number of negative pairs, hence, when computing the binary cross entropy loss, we balance the positive weight p in the loss accordingly.

However BEV 3D coordinates cannot be directly fed to the LSTM as they are. They can in fact become arbitrarily large, depending on the distance between the object and the origin of the reference system. The consequence is that the training process is numerically unstable. To overcome this problem we translate each motion sequence point w.r.t the first point of the input sequence. We also limit the number of last positions processed by the LSTM to a maximum of L. With these tricks the input values always span approximately the same range and let the training kick off. The BEV 3D position of D_j^t is also given w.r.t. the first T_i^{t-1} sequence point.

¹The appearance embedding of a tracked object is computed over its last visible annotation.

3.3. Inference

During inference time, the pipeline is run over sequences and it takes as input detections from the object detector. At each frame it computes the similarity score between previously tracked objects and detected objects in the current frame. The estimated affinity matrix is then solved by the Hungarian algorithm [24] to complete the assignment task. The lowest tolerated score for a match is set to 0.5. Solving the affinity matrix with the Hungarian algorithm implies to find the pairs that are likely to be the same object, based on their motion and their aspect. The process is repeated frame-by-frame to re-identify objects over time.

Tracking score. A tracking score is given to each detected object after the pipeline processing. To compute a value that takes into account detection and re-identification, we compute the tracking score as the product between the detection score and the affinity score. When a detected object is considered to be new we set the tracking score equal to the detection score.

4. Implementation and evaluation

4.1. Datasets and metrics

nuScenes. nuScenes [5] is a multimodal autonomous driving dataset released in 2019. It contains annotated sequences of approximately 20 seconds long captured from a vehicle driving in urban environments. The dataset provides data acquired from different sensors namely cameras, LiDAR and radars. The number of cameras is 6 and they capture the full 360 degrees view around the car. Different types of objects are annotated with 3D bounding boxes in global coordinates at a frame rate of 2 FPS. These frames are named keyframes. Each annotated object is provided with a unique ID across the entire dataset which in our method serves as ground truth supervision. nuScenes also provides frames in between annotated frames, making available a stream of 12 FPS. However, since our method relies on annotations during the training processes, we only used keyframes during the evaluation. nuScenes hosts a 3D tracking benchmark, making it suitable to evaluate our method.

KITTI. KITTI MOT [12] is a single-camera MOT benchmark. It includes 21 training sequences and 29 test sequences of footage recorded by a single camera, mounted on top of a car. Several object categories are annotated with 3D bounding boxes in camera coordinates and unique IDs, similarly to nuScenes. GPS/IMU data is also available, which allows objects to be expressed w.r.t. a global coordinates system.

Metrics. To evaluate the proposed tracking method we rely on the evaluation tools provided by nuScenes and KITTI.

The nuScenes 3D tracking benchmark uses evaluation

metrics introduced in [52], named Average Multi Object Tracking Accuracy (AMOTA) and Average Multi Object Tracking Precision (AMOTP). These modified metrics have been introduced to address the main limitation of the conventional CLEAR metrics Multi Object Tracking Accuracy (MOTA) and Multi Object Tracking Precision (MOTP) [3]. MOTA accounts for all object configuration errors made by the tracker: false positives, misses, mismatches, over all frames. MOTP shows the ability of the tracker to estimate precise object positions, independent of its skill at recognizing object configurations and keeping consistent trajectories. Conventional metrics are computed at the best recall value that yields the maximum value of MOTA. However in this way the metrics do not provide an exhaustive overview of the tracker performance at different recall values. To overcome this issue the AMOTA and AMOTP are computed as the integral of MOTA and MOTP across the spectrum of recall values, better summarizing the tracker performance.

KITTI MOT benchmark assesses the performance of a tracker using the recently proposed Higher Order Tracking Accuracy (HOTA) [27]. This metric is computed as the combination of three Intersection over Union (IoU) scores that account for localization, detection and association. For this reason a single number can be used to rank tracker performance. The HOTA metric allows to focus on single abilities of a tracker and to compare them in a more fine-grained manner e.g. it extends the concept of precision and recall to measure association performance.

4.2. Implementation details

To train the CNN we use a batch size of 128 with K = 4and P = 32. The margin value m in the triplet loss is set to 1. Using a hard mining strategy (in particular selecting the hardest negatives) to train the CNN could lead to a collapsed model i.e. the network predicts every embedding close to 0 [38]. This problem can be solved by first warming up the network by sampling semi-hard triplets as proposed in [38] or with the *batch all* online mining strategy [15] and then continue the training with the batch hard mining. Note also that constructing batches with two or more identical vehicles with different label identities, might potentially lead to distortions in the embedding space. In fact they would be pushed far apart by the triplet loss since they would be treated as negative examples. Nevertheless, we assume the probability of two or more identical vehicles, ending up in the same batch sufficiently low. For our experiments we use a ResNet-18 [14] with pre-trained weights on ImageNet [11] where we replaced the last layer with a fully-connected (FC) one of size 128. As optimizer we use Adam [23] with learning rate 10^{-5} on nuScenes and 10^{-6} when fine-tuning on KITTI.

Crops preparation. The quality of crops taken out from

3D bounding box projections can vary greatly. In fact, projected boxes would hardly include the target, in case of large occlusions, heavy truncation and projections that lie on the edges of the image. Low quality crops can potentially reduce the quality of the embeddings and introduce distortions in the embedding space. To mitigate this issue we apply filters on the selection of the projected bounding boxes, such that the CNN is trained on cleaner crops. Filtering is carried out based on bounding box resolution, their aspect ratio and visibility score. Crops are resized to a fixed size of 224×224 .

Crops and bounding box sizes augmentation. To make the network more robust under a variety of conditions, we applied different data augmentation techniques on the input crops. Color jittering, random horizontal flipping, random rotation and cutout [61]. To also make it more robust against fluctuations of estimated 3D bounding box sizes, typical of object detectors, we add noise on width, height and length. We apply a percentage of noise to each dimension of the bounding box drawn from a uniform distribution $X \sim \mathcal{U}(-a, a)$ with a = 0.2. Adding noise to the 3D bounding box size is necessary. This is due to the fact that the annotated size of an object is consistent over different annotations. In fact, without noise, the size would act as fingerprint that artificially differentiate objects, leading to a network that only inspects the bounding box size to distinguish identities.

To train the second stage we use a batch size of 32 on nuScenes and 16 on KITTI. The number of sampled tracked objects M and sampled detected objects N is set to 16. We set B to a value that approximates 5 seconds and L = 40. For our experiments we use a one-layer LSTM with hidden state size 128. The affinity network consists of 2 FC layers that concatenates the appearance and motion embeddings before passing them to the subsequent layers. In the last layer a sigmoid activation function computes the affinity score. We optimized the weights using Adam [23] with learning rate 5×10^{-4} on nuScenes and 10^{-5} when fine-tuning on KITTI.

Trajectories data augmentation. 3D object detectors reliably estimates the boundaries of detected objects. Nevertheless, in particular monocular object detectors, struggle to precisely estimate the object position, introducing more noise in the direction of the camera. For this reason we introduce noise on the position history of tracked objects during training. Even though more noise should be added in the direction of the camera, we add position noise uniformly in every direction for simplicity reasons. The noise added on each position point is drawn from $X \sim \mathcal{N}(\mu, \sigma^2)$ with $\mu = 0$ and $\sigma^2 = 1$. Note that for the aforementioned reason the noise solely affects the 3D objects position and it is not propagated in the projected 2D crop location.

Object detector. We use an existing 3D object detector

Method	AMOTA \uparrow	$AMOTP \downarrow$
3D center dist	0.208	1.501
Our motion model	0.251	1.489

Table 1. **Motion model efficacy.** Experiment on nuScenes validation set to demonstrate LSTM + affinity network effectiveness. In both cases we use the same set of detections from the Mapillary [42] 3D object detector.

Method	AMOTA \uparrow	$AMOTP \downarrow$
DEFT [6]	0.209	—
QD-3DT [18]	0.242	1.518
Our appearance model w/ QD-3DT det	0.235	1.512
Our motion model w/ QD-3DT det	0.285	1.485

Table 2. **Our motion model and appearance model vs SOTA.** Experiment on nuScenes validation set. We use QD-3DT [18] 3D object detector.

from the current leading camera-only 3D tracking method on nuScenes to fairly compare with them. We apply our tracking pipeline on raw detections from the QD-3DT [18] object detector which is based on Faster RCNN [35]. The object detector is trained on our target datasets and it performs predictions on monocular images. To facilitate future comparisons and to demonstrate that our method can work with any object detector we also evaluate our pipeline on the validation set of nuScenes using the monocular Mapillary 3D object detector [42], which is provided as baseline for the 3D tracking task.

Detections provided by the 3D object detector are filtered by a score threshold. The threshold is set to 0.8 for every object, except for the pedestrian category on KITTI set to 0.85. Moreover, since a single object might be visible from multiple cameras, at each time frame, multiple detections for that object might be available. We then apply a nonmaxima suppression on duplicate detections based on their 2D IoU. When an object is visible in multiple cameras, a feature embedding is extracted from all the views. The final object embedding is then computed as the weighted average of the views embeddings by their resolution.

5. Experiments

In this section we will provide ablation experiments and the results that we obtained on the public benchmarks nuScenes and KITTI.

5.1. Ablation study

LSTM motion model. We conducted an experiment to verify that the LSTM has actually learned how objects typically move in space and that it can produce a meaningful representation of object trajectories. We establish as baseline a simple algorithm that re-identifies objects solely

Method	AMOTA	↑ AMOTP	↓ MOTAR	MOTA	MOTP	RECALL	MT	ML	FAF	TP	FP	FN	IDS	FRAG	TID	LGD
CenterTrack-Vision [62]	0.046	1.543	0.046	0.043	0.753	0.233	573	5235	75.945	26544	17574	89214	3807	2645	2.057	3.819
PermaTrack [45]	0.066	1.491	0.321	0.060	0.724	0.189	652	5065	51.342	29662	16318	86305	3598	2656	2.163	4.248
DEFT [6]	0.177	1.564	0.484	0.156	0.770	0.338	1951	3232	67.741	52099	22163	60565	6901	3420	1.600	3.080
QD-3DT [18]	0.217	1.550	0.563	0.198	0.773	0.375	1893	2970	53.795	52553	16495	60156	6856	3001	1.620	2.961
TripletTrack (Ours)	0.268	1.504	0.605	0.245	0.800	0.400	2085	2922	60.0	62355	18517	56166	1044	3978	1.33	2.50

Table 3. nuScenes tracking leaderboard. Tracking leaderboard of nuScenes tracking benchmark for camera-only methods.

Method	HOTA \uparrow	DetA	AssA	DetRe	DetPr	AssRe	AssPr	LocA
MOTBeyondPixels [39]	63.75 %	72.87 %	56.40 %	76.58 %	85.38 %	59.05 %	86.70 %	86.90 %
JCSTD [44]	65.94 %	65.37 %	67.03 %	68.49 %	82.42 %	71.02 %	82.25 %	84.03 %
MASS [21]	68.25 %	72.92 %	64.46 %	76.83 %	85.14 %	72.12 %	81.46 %	86.80 %
Quasi-Dense [31]	68.45 %	72.44 %	65.49 %	76.01 %	85.37 %	68.28 %	88.53 %	86.50 %
SRK_ODESA [30]	68.51 %	75.40 %	63.08 %	78.89 %	86.00 %	65.89 %	87.47 %	86.88 %
IMMDP [56]	68.66 %	68.02 %	69.76 %	71.47 %	83.28 %	74.50 %	82.02 %	84.80 %
MOTSFusion [26]	68.74 %	72.19 %	66.16 %	76.05 %	84.88 %	69.57 %	85.49 %	86.56 %
TuSimple [10]	71.55 %	72.62 %	71.11 %	76.78 %	83.84 %	74.51 %	86.26 %	85.72 %
SMAT [13]	71.88 %	72.13 %	72.13 %	74.43 %	87.33 %	74.77 %	88.30 %	87.19 %
TrackMPNN [33]	72.30 %	74.69 %	70.63 %	80.02 %	83.11 %	73.58 %	87.14 %	86.14 %
QD-3DT [18]	72.77 %	74.09 %	72.19 %	78.13 %	85.48 %	74.87 %	89.21 %	87.16 %
CenterTrack-Vision [62]	73.02 %	75.62 %	71.20 %	80.10 %	84.56 %	73.84 %	89.00 %	86.52 %
LGM [46]	73.14 %	74.61 %	72.31 %	80.53 %	82.16 %	76.38 %	84.74 %	85.85 %
mono3DT [17]	73.16 %	72.73 %	74.18 %	76.51 %	85.28 %	77.18 %	87.77 %	86.88 %
DEFT [6]	74.23 %	75.33 %	73.79 %	79.96 %	83.97 %	78.30 %	85.19 %	86.14 %
Mono 3D KF [34]	75.47 %	74.10 %	77.63 %	78.86 %	82.98 %	80.23 %	88.88 %	85.48 %
PermaTrack [45]	78.03 %	78.29 %	78.41 %	81.71 %	86.54 %	81.14 %	89.49 %	87.10 %
TripletTrack (Ours)	73.58 %	73.18 %	74.66 %	76.18 %	86.81 %	77.31 %	89.55 %	87.37 %

Table 4. KITTI MOT (Car). Tracking leaderboard of KITTI MOT (car) tracking benchmark for camera-only methods.

Method	HOTA \uparrow	DetA	AssA	DetRe	DetPr	AssRe	AssPr	LocA
TrackMPNN [33]	39.40 %	44.24 %	35.45 %	50.78 %	64.58 %	38.98 %	69.80 %	77.56 %
JCSTD [44]	39.44 %	34.20 %	45.79 %	36.15 %	69.39 %	49.38 %	69.00 %	76.23 %
CenterTrack-Vision [62]	40.35 %	44.48 %	36.93 %	49.91 %	66.83 %	41.05 %	70.19 %	77.81 %
QD-3DT [18]	41.08 %	44.01 %	38.82 %	48.96 %	67.19 %	42.09 %	72.44 %	77.38 %
Quasi-Dense [31]	41.12 %	44.81 %	38.10 %	48.55 %	70.39 %	41.02 %	72.47 %	77.87 %
Mono 3D KF [34]	42.87 %	40.13 %	46.31 %	46.02 %	59.91 %	52.86 %	63.50 %	74.03 %
SRK_ODESA [30]	43.73 %	53.73 %	36.05 %	58.01 %	73.19 %	40.05 %	69.44 %	78.91 %
TuSimple [10]	45.88 %	44.66 %	47.62 %	47.92 %	69.51 %	52.04 %	69.88 %	76.43 %
PermaTrack [45]	48.63 %	52.28 %	45.61 %	57.40 %	71.03 %	49.63 %	73.28 %	78.57 %
TripletTrack (Ours)	42.77 %	39.54 %	46.54 %	41.97 %	71.91 %	50.86 %	71.26 %	77.93 %

Table 5. KITTI MOT (Pedestrian). Tracking leaderboard of KITTI MOT (pedestrian) tracking benchmark for camera-only methods.

based on their 3D bounding box center distance, matching objects by minimizing the sum of the distances between the set of detected objects and the set of tracked objects, with the Hungarian algorithm. We set the maximum tolerated distance for a match to 10 meters^2 . We compare in Table 1 this baseline with our method trained to exclusively make use of the LSTM motion model. In this case the affinity network has to be slightly modified since there are no appearance embeddings as input. As can be seen from the table, the LSTM jointly trained with the affinity network offers superior performance, demonstrating that it is producing a meaningful representation of the objects trajectories instead of merely providing an indication on the last trajectory position. We also compare in Table 2 our motion model with state-of-the-art methods on nuScenes. Surprisingly, our motion model is sufficient to outperform existing methods, providing an additional evidence of the internal LSTM representation effectiveness.

CNN appearance model. Similar to the motion cue we compare our appearance model against state-of-the-art methods on nuScenes. Also in this case a slightly modified affinity network has been trained jointly with the CNN. We show in Table 2 that our appearance model can compete with or even outperform recent methods.

To verify that our pipeline benefits from using appearance features and that the bounding box size increases the quality of the triplet embeddings, we summarize the effect of each cue in Table 6. We also perform the experiments with two different object detectors. Results suggest that each cue contributes to improve the tracking performance.

We experienced the aforementioned collapsed model issue when training the motion model on nuScenes without the bounding box size information. To overcome this, we first warmed up the network with the batch all online mining strategy and then we trained it with the batch hard strategy. We did not suffer from this issue when using the bounding box size as additional input, since the box size assists the network to discriminate too hard examples.

²10 meters is the value that maximizes the metrics performance.

Detector	Mot	App	Boxs	AMOTA \uparrow	AMOTP \downarrow
	-	\checkmark	\checkmark	0.235	1.512
OD 2DT [19]	\checkmark	-	-	0.285	1.485
QD-501 [10]	\checkmark	\checkmark	-	0.289	1.475
	\checkmark	\checkmark	\checkmark	0.295	1.468
	-	\checkmark	\checkmark	0.204	1.515
Manillany [42]	\checkmark	-	-	0.251	1.489
Mapinary [42]	\checkmark	\checkmark	-	0.284	1.474
	\checkmark	\checkmark	\checkmark	0.293	1.468

Table 6. **Effect of each cue.** The pipeline is evaluated on nuScenes validation set using Mapillary [42] and QD-3DT [18] object detectors.

5.2. Benchmark results

nuScenes. We report the quantitative results of our method in Table 3, obtained on the test set. The table shows that our method TripletTrack outperforms prior work by a large margin 23.5%. Compared to state-of-the-art methods our pipeline improves the tracking score relative to each object category and nearly all submetrics, with a huge reduction of identity switches. Contrarily to the current leading method QD-3DT [18], our pipeline makes use of keyframes only. This means that our method achieves higher performance while operating at a lower frame rate.

KITTI MOT. We report the quantitative results in Table 4 and Table 5 obtained on the test set. Since our method benefits from more data, we pre-train our models on nuScenes and we fine-tune them on KITTI. Our method is competitive with most existing ones. Note that some of those performing better need to jointly train the object detector with the tracking mechanism [6] or make use of simulation platforms to generate synthetic training data [45]. By looking at the HOTA metric components we can notice that the gap between our tracking pipeline and the best methods is higher on the detection submetrics compared to the ones related to tracking. This suggests that our pipeline compares well with top methods on data association and that if it were paired with a better object detector it would produce even higher overall tracking scores. Since the proposed pipeline gets the highest AssPr score, it sets a new Pareto optimal front on the submetrics pair AssPr-AssRe for camera-only tracking methods. This also confirms that our method is able to keep the number of identity switches low.

In Figure 3, we show two additional examples of our reidentification pipeline, dealing with occlusion between different objects and camera viewpoints.

6. Conclusions

We propose TripletTrack, a camera-only 3D multi object tracking method that combines appearance features and motion information. Appearance features are extracted through a CNN trained with the triplet loss while motion data is processed by an LSTM. The ID assignment step is performed by the Hungarian algorithm, based on an affinity matrix. The affinity matrix is computed by an affinity network that combines appearance features and motion representation. Our method can also work on top of any 3D object detector and camera setup. Finally it outperforms the camerabased state-of-the-art on the nuScenes tracking benchmark and obtains competitive results on KITTI MOT.



Figure 3. Top: illustration of pedestrian (yellow) and car (red) being occluded and picked up again. Bottom: Illustration of a car and truck re-identified across different cameras.

Acknowledgment. The authors thankfully acknowledge support by Toyota via the TRACE project. Furthermore, we would like to thank Jonas Heylen and Bruno Dawagne for their valuable insights and comments.

References

- Erkan Baser, Venkateshwaran Balasubramanian, Prarthana Bhattacharyya, and Krzysztof Czarnecki. Fantrack: 3d multi-object tracking with feature association network. In 2019 IEEE Intelligent Vehicles Symposium (IV), pages 1426– 1433, 2019. 2
- [2] Nuri Benbarka, Jona Schröder, and Andreas Zell. Score refinement for confidence-based 3d multi-object tracking. arXiv preprint arXiv:2107.04327, 2021. 2
- [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. 5
- [4] Garrick Brazil and Xiaoming Liu. M3d-rpn: Monocular 3d region proposal network for object detection. In 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pages 9286–9295, 2019. 1
- [5] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. arXiv preprint arXiv:1903.11027, 2019. 5
- [6] Mohamed Chaabane, Peter Zhang, Ross Beveridge, and Stephen O'Hara. Deft: Detection embeddings for tracking. arXiv preprint arXiv:2102.02267, 2021. 2, 6, 7, 8
- [7] Xiaozhi Chen, Kaustav Kundu, Ziyu Zhang, Huimin Ma, Sanja Fidler, and Raquel Urtasun. Monocular 3d object detection for autonomous driving. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 2147–2156, 2016. 1
- [8] Xiaozhi Chen, Kaustav Kundu, Yukun Zhu, Andrew G Berneshawi, Huimin Ma, Sanja Fidler, and Raquel Urtasun. 3d object proposals for accurate object class detection. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, Advances in Neural Information Processing Systems, volume 28. Curran Associates, Inc., 2015. 1
- [9] Hsu-Kuang Chiu, Jie Li, Rareş Ambruş, and Jeannette Bohg. Probabilistic 3d multi-modal, multi-object tracking for autonomous driving. In 2021 IEEE International Conference on Robotics and Automation (ICRA), pages 14227–14233, 2021. 2
- [10] Wongun Choi. Near-online multi-target tracking with aggregated local flow descriptor. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2015. 7
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. IEEE Computer Society, 2009. 5
- [12] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 5
- [13] Nicolas Franco Gonzalez, Andres Ospina, and Philippe Calvez. Smat: Smart multiple affinity metrics for multiple object tracking. In Aurélio Campilho, Fakhri Karray, and

Zhou Wang, editors, *Image Analysis and Recognition*, pages 48–62, Cham, 2020. Springer International Publishing. 7

- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 770–778, 2016. 5
- [15] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification, 2017. 2, 4, 5
- [16] Jonas Heylen, Mark De Wolf, Bruno Dawagne, Marc Proesmans, Luc Van Gool, Wim Abbeloos, Hazem Abdelkawy, and Daniel Olmeda Reino. Monocinis: Camera independent monocular 3d object detection using instance segmentation. In 2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW), pages 923–934, 2021. 1
- [17] Hou-Ning Hu, Qi-Zhi Cai, Dequan Wang, Ji Lin, Min Sun, Philipp Krahenbuhl, Trevor Darrell, and Fisher Yu. Joint monocular 3d vehicle detection and tracking. In *Proceedings* of the IEEE/CVF International Conference on Computer Vision (ICCV), October 2019. 2, 7
- [18] Hou-Ning Hu, Yung-Hsu Yang, Tobias Fischer, Fisher Yu, Trevor Darrell, and Min Sun. Monocular quasi-dense 3d object tracking. ArXiv:2103.07351, 2021. 2, 6, 7, 8
- [19] Kemiao Huang and Qi Hao. Joint multi-object detection and tracking with camera-lidar fusion for autonomous driving. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 6983–6989, 2021. 2
- [20] Xiaolong Jiang, Peizhao Li, Yanjing Li, and Xiantong Zhen. Graph neural based end-to-end data association framework for online multiple-object tracking, 2019. 2
- [21] Hasith Karunasekera, Han Wang, and Handuo Zhang. Multiple object tracking with attention to appearance, structure, motion and size. *IEEE Access*, 7:104423–104434, 2019. 7
- [22] Aleksandr Kim, Aljoša Ošep, and Laura Leal-Taixé. Eagermot: 3d multi-object tracking via sensor fusion. In 2021 IEEE International Conference on Robotics and Automation (ICRA), pages 11315–11321, 2021. 2
- [23] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. cite arxiv:1412.6980Comment: Published as a conference paper at the 3rd International Conference for Learning Representations, San Diego, 2015. 5, 6
- [24] H. W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, pages 83–97, 1955. 2, 5
- [25] Ratnesh Kumar, Edwin Weill, Farzin Aghdasi, and Parthsarathy Sriram. Vehicle re-identification: an efficient baseline using triplet embedding, 2019. 2, 4
- [26] Jonathon Luiten, Tobias Fischer, and Bastian Leibe. Track to reconstruct and reconstruct to track. *IEEE Robotics and Automation Letters*, 2020. 2, 7
- [27] 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 (IJCV)*, 2020. 5

- [28] Wenhan Luo, Junliang Xing, Anton Milan, Xiaoqin Zhang, Wei Liu, and Tae-Kyun Kim. Multiple object tracking: A literature review. *Artificial Intelligence*, 293:103448, Apr 2021. 1
- [29] Arsalan Mousavian, Dragomir Anguelov, John Flynn, and Jana Košecká. 3d bounding box estimation using deep learning and geometry. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 5632–5640, 2017. 1
- [30] Dmytro Mykheievskyi, Dmytro Borysenko, and Viktor Porokhonskyy. Learning local feature descriptors for multiple object tracking. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, November 2020. 2, 7
- [31] 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 (CVPR)*, pages 164–173, June 2021. 7
- [32] Ziqi Pang, Zhichao Li, and Naiyan Wang. Simpletrack: Understanding and rethinking 3d multi-object tracking. arXiv preprint arXiv:2111.09621, 2021. 2
- [33] Akshay Rangesh, Pranav Maheshwari, Mez Gebre, Siddhesh Mhatre, Vahid Ramezani, and Mohan M Trivedi. Trackmpnn: A message passing graph neural architecture for multi-object tracking. arXiv preprint arXiv:2101.04206. 7
- [34] Andreas Reich and Hans-Joachim Wuensche. Monocular 3d multi-object tracking with an ekf approach for long-term stable tracks. In 2021 IEEE 24th International Conference on Information Fusion (FUSION), pages 1–7, 2021. 2, 7
- [35] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. 6
- [36] Amir Sadeghian, Alexandre Alahi, and Silvio Savarese. Tracking the untrackable: Learning to track multiple cues with long-term dependencies. In 2017 IEEE International Conference on Computer Vision (ICCV), pages 300–311, 2017. 2
- [37] Samuel Scheidegger, Joachim Benjaminsson, Emil Rosenberg, Amrit Krishnan, and Karl Granström. Mono-camera 3d multi-object tracking using deep learning detections and pmbm filtering. In 2018 IEEE Intelligent Vehicles Symposium (IV), pages 433–440, 2018. 2
- [38] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 815–823, 2015. 2, 4, 5
- [39] Sarthak Sharma, Junaid Ahmed Ansari, J. Krishna Murthy, and K. Madhava Krishna. Beyond pixels: Leveraging geometry and shape cues for online multi-object tracking. In *Proceedings of the IEEE International Conference on Robotics* and Automation (ICRA), 2018. 2, 7
- [40] Abhijeet Shenoi, Mihir Patel, JunYoung Gwak, Patrick Goebel, Amir Sadeghian, Hamid Rezatofighi, Roberto Martín-Martín, and Silvio Savarese. Jrmot: A real-time 3d

multi-object tracker and a new large-scale dataset. In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 10335–10342, 2020. 2

- [41] Xuepeng Shi, Qi Ye, Xiaozhi Chen, Chuangrong Chen, Zhixiang Chen, and Tae-Kyun Kim. Geometry-based distance decomposition for monocular 3d object detection. In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pages 15172–15181, October 2021. 1
- [42] Andrea Simonelli, Samuel Rota Bulò, Lorenzo Porzi, Manuel López-Antequera, and Peter Kontschieder. Disentangling monocular 3d object detection. In *arXiv*, May 2019. 1, 6, 8
- [43] Chunfeng Song, Yan Huang, Wanli Ouyang, and Liang Wang. Mask-guided contrastive attention model for person re-identification. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 1179–1188, 2018. 4
- [44] Wei Tian, Martin Lauer, and Long Chen. Online multi-object tracking using joint domain information in traffic scenarios. *IEEE Transactions on Intelligent Transportation Systems*, 21(1):374–384, 2020. 7
- [45] 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 (ICCV), pages 10860–10869, October 2021. 2, 7, 8
- [46] Gaoang Wang, Renshu Gu, Zuozhu Liu, Weijie Hu, Mingli Song, and Jenq-Neng Hwang. Track without appearance: Learn box and tracklet embedding with local and global motion patterns for vehicle tracking. In *Proceedings* of the IEEE/CVF International Conference on Computer Vision (ICCV), pages 9876–9886, October 2021. 2, 7
- [47] Jianren Wang, Siddharth Ancha, Yi-Ting Chen, and David Held. Uncertainty-aware self-supervised 3d data association. In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 8125–8132, 2020. 2
- [48] Sukai Wang, Peide Cai, Lujia Wang, and Ming Liu. Ditnet: End-to-end 3d object detection and track id assignment in spatio-temporal world. *IEEE Robotics and Automation Letters*, 6(2):3397–3404, 2021. 2
- [49] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In 2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW), pages 913–922, 2021. 1
- [50] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Probabilistic and geometric depth: Detecting objects in perspective, 2021. 1
- [51] Yan Wang, Lequn Wang, Yurong You, Xu Zou, Vincent Chen, Serena Li, Gao Huang, Bharath Hariharan, and Kilian Q. Weinberger. Resource aware person re-identification across multiple resolutions. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8042– 8051, 2018. 4
- [52] Xinshuo Weng, Jianren Wang, David Held, and Kris Kitani. AB3DMOT: A Baseline for 3D Multi-Object Tracking and New Evaluation Metrics. *ECCVW*, 2020. 2, 5

- [53] Xinshuo Weng, Yongxin Wang, Yunze Man, and Kris M. Kitani. Gnn3dmot: Graph neural network for 3d multiobject tracking with 2d-3d multi-feature learning. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 6498–6507, 2020. 2
- [54] Hai Wu, Wenkai Han, Chenglu Wen, Xin Li, and Cheng Wang. 3d multi-object tracking in point clouds based on prediction confidence-guided data association. *IEEE Transactions on Intelligent Transportation Systems*, pages 1–10, 2021. 2
- [55] Hai Wu, Qing Li, Chenglu Wen, Xin Li, Xiaoliang Fan, and Cheng Wang. Tracklet proposal network for multi-object tracking on point clouds. In *IJCAI*, 2021. 2
- [56] Yu Xiang, Alexandre Alahi, and Silvio Savarese. Learning to track: Online multi-object tracking by decision making. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2015. 7
- [57] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Centerbased 3d object detection and tracking. In 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 11779–11788, 2021. 2
- [58] Jan-Nico Zaech, Alexander Liniger, Dengxin Dai, Martin Danelljan, and Luc Van Gool. Learnable online graph representations for 3d multi-object tracking. *IEEE Robotics and Automation Letters*, pages 1–1, 2022. 2
- [59] Yihan Zeng, Chao Ma, Ming Zhu, Zhiming Fan, and Xiaokang Yang. Cross-modal 3d object detection and tracking for auto-driving. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 3850– 3857, 2021. 2
- [60] Wenwei Zhang, Hui Zhou, Shuyang Sun, Zhe Wang, Jianping Shi, and Chen Change Loy. Robust multi-modality multi-object tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2
- [61] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation, 2017. 6
- [62] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. *ECCV*, 2020. 2, 7