

Simultaneous Detection and Tracking with Motion Modelling for Multiple Object Tracking

ShiJie Sun¹, Naveed Akhtar², XiangYu Song³, HuanSheng Song¹, Ajmal Mian², and Mubarak Shah⁴

¹ Chang'an University, Xi'an, Shaanxi, China
{shijieSun, hshsong}@chd.edu.cn

² University of Western Australia, 35 Stirling Highway, Crawley, WA, Australia
{naveed.akhtar, ajmal.mian}@uwa.edu.au

³ Deakin University, RWAurn Ponds, Victoria 3216, Melbourne, Australia
xiangyu.song@deakin.edu.au

⁴ University of Central Florida, Orlando, FL, America
shah@crcv.ucf.edu

Abstract. Deep learning based Multiple Object Tracking (MOT) currently relies on off-the-shelf detectors for tracking-by-detection. This results in deep models that are detector biased and evaluations that are detector influenced. To resolve this issue, we introduce Deep Motion Modeling Network (DMM-Net) that can estimate multiple objects' motion parameters to perform joint detection and association in an end-to-end manner. DMM-Net models object features over multiple frames and simultaneously infers object classes, visibility and their motion parameters. These outputs are readily used to update the tracklets for efficient MOT. DMM-Net achieves PR-MOTA score of 12.80 @ 120+ fps for the popular UA-DETRAC challenge - which is better performance and orders of magnitude faster. We also contribute a synthetic large-scale public dataset Omni-MOT for vehicle tracking that provides precise ground-truth annotations to eliminate the detector influence in MOT evaluation. This 14M+ frames dataset is extendable with our public script (Code at [Dataset](#), [Dataset Recorder](#), [Omni-MOT Source](#)). We demonstrate the suitability of Omni-MOT for deep learning with DMM-Net, and also make the source code of our network public.

Keywords: Multiple Object Tracking, Tracking-by-detection, Deep Learning, Simultaneous detection and tracking.

1 Introduction

Multiple Object Tracking (MOT) [10], [25], [44], [46] is a longstanding problem in Computer Vision [27]. Contemporary deep learning based MOT has widely adopted the tracking-by-detection paradigm [45], that capitalizes on the natural division of *detection* and *data association* tasks for the problem. In standard MOT evaluation protocol, the object detections are assumed to be known and public detections are provided on evaluation sequences, and MOT algorithms are expected to output object tracks by solving the data association problem.

Although attractive, using off-the-shelf detectors for MOT also has undesired ramifications. For instance, a deep model employed for the subsequent data association task

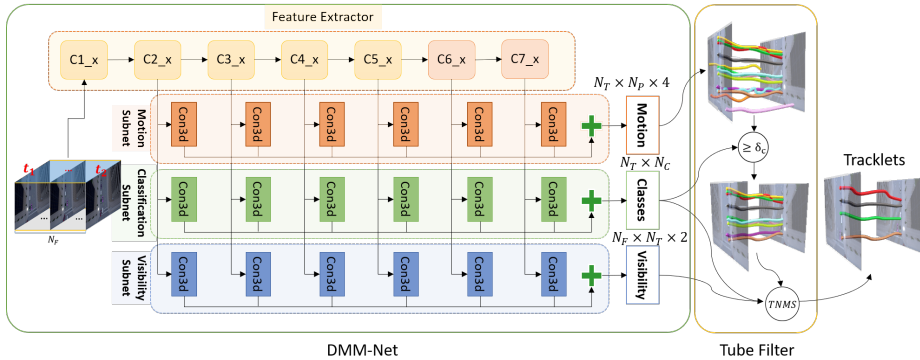


Fig. 1. Schematics of the proposed end-to-end trainable DMM-Net: N_F frames from time stamp $t_1 : t_2$ are input to the network. The frame sequence is first processed with a *Feature Extractor* comprising 3D ResNet-like [20] convolutional groups. Outputs of the last groups 2 to 7 are processed by Motion Subnet, Classifier Subnet, and Visibility Subnet. Each sub-network uses 3D convolutions to learn features that are concatenated to predict anchor tubes’ motion parameters ($O_M \in \mathbb{R}^{N_T \times N_P \times 4}$), object classes ($O_C \in \mathbb{R}^{N_T \times N_C}$), and visibility ($O_V \in \mathbb{R}^{N_F \times N_T \times 2}$), where N_T , N_P and N_C denote the number of anchor tubes, motion parameters and object classes (including background). We explain anchor tubes and motion parameters in Section 4. DMM-Net is trained with its specialized loss. For deployment, the anchor tubes predicted by the DMM-Net are filtered to compute tracklets defined over multiple frames. These tracklets are later combined to form a complete track.

(or a constituent sub-task) gets biased to the detector. The detector performance can also become a bottle-neck for the overall tracker. Additionally, composite techniques resulting from independent detectors fail to harness the true representation power of deep learning by sacrificing end-to-end training etc. It also seems unnatural that a MOT tracker must be evaluated on different detectors to interpret its tracking performance. All these problems are potentially solvable if trackers can implicitly detect the target objects, and detector bias can be removed from the ground-truth labeling of the tracks. This work makes strides towards these solutions.

We make two major contributions towards setting the tracking-by-detection paradigm free from off-the-shelf detectors in deep learning settings. Our first contribution comes as the first-of-its-kind deep network that performs object detections and data association by estimating multiple object motion parameters in an end-to-end manner. Our network, called Deep Motion Modeling Network (DMM-Net), predicts object motion parameters, their classes and their visibility with the help of three sub-networks, see Fig. 1. These sub-networks exploit feature maps of frames in a video sequence that are learned with a Feature Extractor comprising seven 3D ResNet-like [20] convolutional groups. Instead of individual frames, DMM-Net simultaneously processes multiple (i.e. 16) frames. To handle multiple tracks in those frames, we introduce the notion of anchor tubes that extends the concept of anchor boxes in object detection [26] along the temporal dimension for MOT. Similar to [26], these anchor tubes are pre-defined to reduce the computation and improve the network speed. The predicted motion parameters can describe the shape offset and scale of each pre-defined anchor tube in

the spatio-temporal space. We propose individual losses over the comprehensive representations of the sub-networks to predict object motion parameters, object classes and visibility. The DMM-Net output is readily usable to update the tracks.

As our second major contribution, we propose a realistic large-scale dataset with accurate and extensive ground-truth annotations. The proposed dataset, termed Omni-MOT for its comprehensive coverage of the MOT scenarios, is generated with CARLA simulator [13] for vehicle tracking. The dataset comprises 14M+ frames, 250K tracks, 110 million bounding boxes, three weather conditions, three crowd levels and three camera views in five simulated towns. By eliminating the use of off-the-shelf detectors in ground-truth labeling, it removes any detector bias in evaluating the techniques.

The Omni-MOT dataset and DMM-Net source code are both publicly available for the broader research community. For the former, we also provide software tools to freely extend the dataset. We demonstrate the suitability of the Omni-MOT for deep models by training and evaluating DMM-Net on its videos. We also augment DMM-Net with Omni-MOT and evaluate our technique on the popular UA-DETRAC challenge [47]. The remote server computed results show that DMM-Net is able to achieve a very competitive 12.80 score for the comprehensive PR-MOTA metric with the overall speed of 123 fps. The orders of magnitude increase in the speed is directly attributed to the intrinsic detections in our tracking-by-detection technique.

2 Related Work

Multiple Object Tracking (MOT) is a fundamental problem in Computer Vision that has attracted significant interest of researchers in recent years. For a general review of the related literature, we refer to Luo et al. [27] and Emami et al. [14]. Here, we focus on the key contributions that relate to this work more closely.

With the recent advances in object detectors, tracking-by-detection is fast becoming the common contemporary paradigm for MOT [39], [40], [45], [48]. In this scheme, objects are first detected frame-wise and later associated with each other across multiple frames. Relying on off-the-shelf detectors, techniques following this paradigm mainly focus on object association, which can make them inherently detector biased. These methods can be broadly categorized as local [34], [42] and global [10], [37], [50] approaches. The former use two frames for data association, while the latter associate objects over a larger number of frames.

More recent global techniques cast data association into a network flow problem [4], [7], [33], [41]. For instance, Berclaz et al. [4] solved a constrained flow optimization problem for MOT and used the k-shortest paths algorithm for associating the tracks. Chari et al. [8] added a pairwise cost to the min-cost network flow framework and proposed a convex relaxation of the problem. However, such methods rely on object detectors rather strongly, which makes them less attractive in the presence of occlusions and misdetections. To mitigate the problems resulting from occlusions, Milan et al. [29] employed a continuous energy minimization framework for MOT that incorporates explicit occlusion reasoning and appearance modeling. Wen et al. [49] also proposed a data association technique based on undirected hierarchical relation hyper-graph.

Recent years have witnessed numerous deep learning based techniques for MOT that use pre-trained networks to build a feature model. Object features computed by the model are later used for data association. Improving this scheme, Schulter et al. [38]

proposed a network that is trained for data association in the context of MOT. Son et al. [43] proposed a quadruplet CNN that learns to associate objects detected in different video frames. Sun et al. [44] proposed a deep affinity network to model features of pre-detected objects and compute object affinities by the same network. Bea et al. [3] modified the Siamese Network to learn deep representations for MOT with object association. There are few instances of deep learning techniques that aim at removing the reliance of tracking on independent detectors. For instance, Feichtenhofer et al [15] proposed an R-FCN based network [9] that performs object detection and jointly builds an object model to be used for data association. However, their method is limited to frame-wise detections. Consequently, it only allows frame-by-frame association, requiring manual adjustment of temporal stride. Our technique is inherently different, as it directly computes tracklets over multiple frames by motion modeling, enabling realtime solutions while considering all the frames.

Besides the development of novel techniques, the role of datasets is central to the advancement of deep learning based MOT. Currently, a few large datasets for this task are available, e.g. PETS2009 [17], KITTI [18], DukeMTMC [36], PoseTrack [23], and MOT Challenge datasets [25], [28]. These datasets are recorded in the real world with pedestrians and vehicles as the objects of interest. We refer to the original works for more details on these datasets. Below, we briefly discuss UA-DETRAC [47] for its high relevance to our contribution.

UA-DETRAC [47] is a large dataset for traffic scenes MOT. It provides object bounding boxes, their IDs and information on the overlapped ratio of the objects. However, the provided detections are individually generated by detectors and hence are prone to errors. This results in an undesired detector-bias in tracker evaluation. Besides, different pre-processing procedures of the detectors employed by the dataset also cause problems in fair evaluation. Although UA-DETRAC has served a great purpose in advancing the state-of-the-art in MOT for vehicles, the aforementioned issues call for a more transparent dataset that does not rely on off-the-shelf detectors for evaluating trackers. This work provides such a dataset with realistic settings and complete control over the environment conditions and camera viewpoints. The proposed large-scale dataset is also extendable and can further be used to augment the real-world datasets for improved tracking performance.

3 Omni-MOT Dataset

We term the proposed dataset as **Omni-MOT** (OMOT) dataset for its comprehensive coverage of the conditions and scenarios possible in MOT. The dataset is publicly available for download. Moreover, we also make the recording script for the dataset public that will enable the community to further extend the data. The provided script has the ability to generate multi-camera videos. The dataset is recorded using virtual cameras in the CARLA simulator [13].

In the proposed dataset, five different simulated cities are considered. For each city, we use up to 39 cameras. The cameras are placed with viewpoints that have three levels of difficulty for the MOT scenarios. Namely, (a) Easy view: which results in no occlusion of the vehicles. (b) Ordinary view: that allows temporary occlusions but forbids continuous occlusions. (c) Hard view: that allows continuous occlusions in the videos. Collectively, we provide 90 scenes in the dataset that result in 3,510 videos. There

are 14.04M frames of size 1920×1080 in the OMOT dataset that are recorded in the ‘XDIV’ format to provide high-quality videos with acceptable memory size.

Besides controlling the camera viewpoints, we include three different weather conditions; clear, cloudy and rainy. We also include three levels of traffic congestion in the OMOT dataset i.e. low, medium and severe congestion. We provide further details of the dataset in the supplementary material, where we also include a video showing different scenes, camera viewpoints and weather conditions. CARLA simulator allows us to capture comprehensive information on the target objects with high precision. Hence, besides being accurate, our ground truth annotations are detailed enough to be used for other related problems, e.g. velocity estimation, 3D localization, camera calibration.

To the best of our knowledge, Omni-MOT is the first realistic large dataset for tracking vehicles that completely relinquishes off-the-shelf detectors in ground truth generation, and provides comprehensive annotations. Moreover, with the provided scripts, the dataset can easily be extended for future research. To put the scale of Omni-MOT into perspective, the provided number of frames is almost 1, 200 times larger than MOT17. The number of provided tracks and boxes respectively are 210 and 30 times larger than UA-DETRAC. Not to mention, all the boxes and tracks are simulator generated that avoids any labeling error. Considering that OMOT can also be used for data augmentation, we make the ground truth for the test videos public as well.

4 Deep Motion Modeling Network

To absolve deep learning based tracking-by-detection from independently pre-trained off-the-shelf detectors, we propose **Deep Motion Modeling Network** (DMM-Net) for online MOT (see Fig. 1). Our network enables MOT by jointly performing object detection, tracking, and classification across multiple video frames without requiring pre-detections and subsequent data association. For a given input video, it outputs objects’ motion parameters, classes, and their visibility across the input frames. We provide a detailed discussion of our network below. However, we first introduce notations and conventions for a concise description of the text and figures.

Notations:

- N_F, N_C, N_P, N_T respectively denote the number of input frames, object classes (0 for ‘background’), time-related motion parameters, and anchor tubes.
- W, H are the frame width, and frame height.
- I_t denotes the video frame at time t . Subsequently, a 4-D tensor $I_{t_1:t_2:N_F} \in \mathbb{R}^{3 \times N_F \times W \times H}$ denotes N_F video frames from time t_1 to $t_2 - 1$. For simplicity, we often ignore the subscript “: N_F ”.
- $B_{t_1:t_2:N_F}, C_{t_1:t_2:N_F}, V_{t_1:t_2:N_F}$ respectively denote the ground truth boxes, classes and visibilities in the selected N_F video frames from time t_1 to $t_2 - 1$. The text also ignores “: N_F ” for these notations.
- $O_{M,t_1:t_2:N_F}, O_{C,t_1:t_2:N_F}, O_{V,t_1:t_2:N_F}$ denote the estimated motion parameters, classes, and visibilities. With time stamps and frames clear from the context, we simplify these notations as O_M, O_C, O_V . In Fig. 2, we illustrate the notions of object visibility, classes and motion parameters with a simple example for clarity.

Conventions: The shape of a network output tensor is considered to be Batch \times Channels \times Duration \times Width \times Height, where Duration accounts for the number of frames. For brevity, we often ignore the Batch dimension.




Fig. 2. (Best viewed in color) Illustration of motion visibility, class, and motion parameters: The visibility of the vehicle goes from 1 (fully visible) to 0 (invisible) and back to 0.5 (partially visible). Classes are predefined as 1 for vehicle and 0 for everything else. Motion parameters $\{p_{11}, \dots, p_{43}\}$ are used to locate the center x (cx), center y (cy), width (w) and height (h) of the tube/tracklet at anytime. We employ a quadratic motion model leveraging 4×3 matrices.

Table 1. Inputs for training and testing DMM-Net. The tensor dimensions are given as Channels \times Duration \times Width \times Height. $n_{t_1:t_2}$ is the number of tracks in $I_{t_1:t_2}$.

Input	Dimensions/Size	Train	Test
$I_{t_1:t_2}$	$3 \times N_F \times W \times H$	✓	✓
$t_{t_1:t_2}$	N_F	✓	✓
$B_{t_1:t_2}$	$N_F \times n_{t_1:t_2} \times 4$	✓	✗
$C_{t_1:t_2}$	$n_{t_1:t_2}$	✓	✗
$V_{t_1:t_2}$	$N_F \times n_{t_1:t_2}$	✓	✗

4.1 Data Preparation

Appropriate data preparation is important for training (and later testing) our network. In this process, we also avail the opportunity of augmenting the data to induce robustness in our model against camera photometric distortions, background scene variations, and other practical factors. The inputs expected by our network are summarized in Table 1. For training a robust network, we perform the following data pre-processing:

1. *Stochastic frame sequence* is introduced by randomly choosing N_F frames from $2N_F$ frames.
2. *Static scene emulation* is performed by duplicating selected frames N_F number of times.
3. *Photometric camera distortions* are introduced in the frames by scaling each pixel by a random value in the range $[0.7, 1.5]$. The resulting frames are converted to HSV format, and their saturation channel is again scaled by a random value in $[0.7, 1.5]$. A frame is then converted back to RGB format and re-scaled by a random value in the same range. This process of photometric distortion is inspired by [44].
4. *Frame expansion* is performed by enlarging video frames by a random factor in the range $[1, 1.2]$. To that end, we pad the original frames with extra pixels whose value is set to the mean pixel value of the training data.

In our pre-processing, the above-mentioned steps 2-4 are applied to the frames with probability 0.01, 0.5, 0.5, respectively. The resulting frame are then resized to the fixed dimension $H \times W \times 3$, and horizontally flipped with a probability of 0.5. We simultaneously process the selected N_F frames by applying the above mentioned transformations. These RGB video frames are then arranged as 4D-tensors $I_{t_1:t_2} \in \mathbb{R}^{3 \times N_F \times W \times H}$.

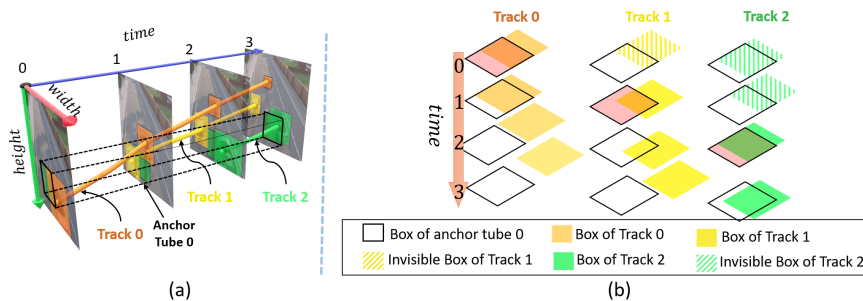


Fig. 3. (a) One anchor tube and three tracks are shown to illustrate the search for the best track of the anchor tube. (b) is the simplified demonstration of (a). Based on the largest overlap between the first tube box and the first *visible* track box, Track 2 is selected as the best match.

We fill the ground truth data matrices $\mathbf{B}_{t_1:t_2}$, $\mathbf{C}_{t_1:t_2}$, $\mathbf{V}_{t_1:t_2}$ with the detected boxes, and visibilities in the video frames from t_1 to $t_2 - 1$. We set the labeled boxes, classes, and visibilities of fully occluded boxes to 0. We also remove the tracks whose ratio of fully occluded boxes are greater than δ_v . In our experiments, $N_F = 16$ proved a reasonable bound for the used datasets.

4.2 Anchor Tubes

Inspired by the well-established and effective Single Shot Detector (SSD) [26] for object detection, we extend the core concept of this technique for MOT. Analogous to the anchor boxes of SSD, we introduce *anchor tubes* for DMM-Net. In our setup, an anchor tube is a set of anchor boxes (and associated object class and visibility), that share the same location in multiple frames along the temporal dimension (see supplementary material for further visualization). The information of anchor tubes is encoded with the tensors $\mathbf{B}_{t_1:t_2}$, $\mathbf{C}_{t_1:t_2}$, and $\mathbf{V}_{t_1:t_2}$. The DMM-Net is designed to predict the tube shape offset parameters along the temporal dimension, the confidence for each object class, and the visibility of each box in the tube.

Computing anchor tubes for network training can also be interpreted as further data preparation for DMM-Net. We first employ a search strategy to find the best-matched track for each anchor tube, and subsequently encode the tubes by their best-matched tracks. For the former, we specify the overlap between an anchor tube and a track as the overlap ratio between the first visible box of the track and the corresponding box in the anchor tube. A simplified illustration of this concept is presented in Fig. 3 with one anchor tube and three tracks. The anchor tube 0 iterates over all the tracks to find the largest overlap ratio. The first box of Track 1 and the first two boxes of Track 2 are occluded. Therefore, the overlap ratio of the anchor box filled red is used for selecting the best-matched track, i.e. Track 2 for the largest overlap.

To encode the anchor tubes, we employ their best-matched tracks ($\mathbf{B}_{t_1:t_2}$) along with the classes ($\mathbf{C}_{t_1:t_2}$) and visibility ($\mathbf{V}_{t_1:t_2}$). We denote a box of the i^{th} anchor tube at the t^{th} frame as $a_{i,t} = (a_{i,t}^{cx}, a_{i,t}^{cy}, a_{i,t}^w, a_{i,t}^h)$, where the superscripts cx and cy indicate the (x, y) location of the center, and w, h denote the width and height of the box. We use the same notation in the following text as well. Each anchor box is encoded by the




Fig. 4. Used ResNeXt block: (conv, x^3, F) denote ‘F’ convolutional kernels of size $x \times x \times x$. *BN* is Batch Normalization [22] and shortcut connection is summation.

box of its best-matched track as follows:

$$\begin{cases} g_{i,t}^w = \log(b_{i,t}^w/a_{i,t}^w) \\ g_{i,t}^h = \log(b_{i,t}^h/a_{i,t}^h) \\ g_{i,t}^{cx} = (b_{i,t}^{cx} - a_{i,t}^{cx})/a_{i,t}^w \\ g_{i,t}^{cy} = (b_{i,t}^{cy} - a_{i,t}^{cy})/a_{i,t}^h \end{cases} \quad (1)$$

where $(b_{i,t}^{cx}, b_{i,t}^{cy}, b_{i,t}^w, b_{i,t}^h)$ describe the box of the best-matched track at the t^{th} frame, and $(g_{i,t}^{cx}, g_{i,t}^{cy}, g_{i,t}^w, g_{i,t}^h)$ represent the resulting encoded box for the best-matched track. Following this encoding, we replace each original box, class, and visibility by its newly encoded counterpart. In the above-mentioned encoding, an anchor tube that has its best-matched track overlap ratio less than δ_o , is identified as the ‘background’ class.

4.3 Motion Model

For motion modeling, we leverage a quadratic model in time. One of the outputs of DMM-Net is a tensor of motion parameters $\mathbf{O}_M \in \mathbb{R}^{N_T \times 4 \times N_P}$, where $N_P = 3$ in our experiments and N_T indicates the number of anchor tubes. We estimate an encoded anchor tube for a track as:

$$\begin{cases} \hat{g}_{i,t}^w = p_{11}t^2 + p_{12}t + p_{13} \\ \hat{g}_{i,t}^h = p_{21}t^2 + p_{22}t + p_{23} \\ \hat{g}_{i,t}^{cx} = p_{31}t^2 + p_{32}t + p_{33} \\ \hat{g}_{i,t}^{cy} = p_{41}t^2 + p_{42}t + p_{43} \end{cases} \quad (2)$$

where $\hat{g}_{i,t}$ indicates an estimated box descriptor of the i^{th} encoded anchor tube at the t^{th} frame, the superscripts cx , cy , w and h indicate the center x, center y, height, width of the box respectively, and $\{p_{11}, \dots, p_{43}\}$ are the motion parameters. Each encoded ground-truth anchor tube can be decoded into a ground truth track by Eq. (1). We further elaborate on the relationship between the motion parameters and the ground truth tracks in the supplementary material. Note that the used motion function is replaceable by any differential function in our technique. We choose quadratic motion modeling considering vehicle tracking as our primary objective.

4.4 Architecture

As shown in Fig. 1, our network comprises a Feature Extractor and three sub-network. The network simultaneously processes N_F video frames for which features are first extracted. The feature maps of six intermediate layers of the Feature Extractor are fed to the sub-networks. These networks predict tensors containing motion parameters ($\mathbf{O}_M \in \mathbb{R}^{N_T \times N_P \times 4}$), object classes ($\mathbf{O}_C \in \mathbb{R}^{N_T \times N_C}$), and visibility information ($\mathbf{O}_V \in \mathbb{R}^{N_F \times N_T \times 2}$).

Table 2. Convolutional groups of the Feature Extractor. Conv is abbreviated as ‘C’, ‘F’ is the number of feature maps, as in Fig. 4, and ‘N’ is the number of bottleneck blocks in each layer.

C1_x	C2_x		C3_x		C4_x		C5_x		C6_x		C7_x	
conv, 7^3 , 64	F	N	F	N	F	N	F	N	F	N	F	N
	64	3	128	4	256	23	512	3	32	3	32	3

Table 3. Details of the Motion Subnet (M.S.), Classifier Subnet (C.S.), and Visibility Subnet (V.S.): K.S., P.S., S.S., and N.K. respectively denote the kernel shape, padding size, stride size, and the number of kernels. The row ‘Output’ reports output tensor shape. $\mathcal{K} = \{10, 8, 8, 5, 4, 4\}$ is the number of anchor tubes for each of the six input feature maps for each sub-network.

	M.S.	C.S.	V.S.
K.S.	$\{8, 4, 2, 1, 1, 1\} \times 3 \times 3$		
P.S.	$0 \times 1 \times 1$		
S.S.	$1 \times 1 \times 1$		
N.K.	$4\mathcal{K}N_P$	$\mathcal{K}N_C$	$2\mathcal{K}$
Output	$N_T \times N_P \times 4$	$N_T \times N_C$	$N_F \times N_T \times 2$

Feature Extractor: We construct our Feature Extractor based on the ResNeXt blocks of the 3D ResNet [20]⁵. The architectural details of the blocks are provided in Fig. 4. The blocks accept ‘F’ channel input that allows us to simultaneously model spatio-temporal features in N_F frames. We enhance the 3D ResNet architecture by removing the *fully-connected* and *softmax* layers of the network and appending two extra convolutional groups denoted as Conv6_x and Conv7_x. We perform this enhancement because the additional convolutional layers are still able to encode further higher level spatio-temporal features of the frames. Details on the convolutional groups used by the Feature Extractor are provided in Table 2. Here, Conv1_x (abbreviated as C1_x) contains 64 convolutional kernels of size $7 \times 7 \times 7$, stride $1 \times 2 \times 2$. A $3 \times 3 \times 3$ max-pooling layer with stride 2 is inserted before Conv2_x for down sampling. Each convolutional layer is followed by batch normalization [21] and ReLU [31]. Spatio-temporal down-sampling is performed by Conv3_1, Conv4_1 and Conv5_1 with a stride of 2.

Sub-networks: The DMM-Net has three sub-networks to compute motion parameters, object classes, and their visibility. These networks use the outputs of Conv2_x to Conv7_x groups of the Feature Extractor. Each sub-network processes the input with six convolutional layers. The architectural details of all three sub-networks are summarized in Table 3. The kernel shape, padding size and stride are the same for each network, whereas the employed numbers of kernels are different. We fix the temporal kernel size for each network to $\{8, 4, 2, 1, 1, 1\}$. Denoting the numbers of anchor tubes defined for the six input feature maps of a sub-network by \mathcal{K} , we let $\mathcal{K} = \{10, 8, 8, 5, 4, 4\}$ in our experiments. For each sub-network, we concatenate the output feature maps of each convolutional layer. We reshape the concatenated feature according to the dimensions mentioned in the last row of the table. These outputs are subsequently used to compute the network loss.

⁵ Notation are adopted from the original work.

Network Loss: Based on the architecture, the overall loss of DMM-Net is defined as a weighted sum of three sub-losses, including the motion loss (\mathcal{L}_M), the classification loss (\mathcal{L}_C) and the visibility loss (\mathcal{L}_V), given as:

$$\mathcal{L} = \frac{1}{N}(\alpha\mathcal{L}_M + \beta\mathcal{L}_C + \mathcal{L}_V), \quad (3)$$

where N is the number of *positive* anchor tubes, and α, β are both hyper-parameters. The positive anchor tubes exclude those tubes that get assigned to the background.

The motion loss \mathcal{L}_M is defined as the sum of *Smooth-L1* losses between the ground truth encoded tracks $g_{i,t}$ and their predicted encoded tracks $\hat{g}_{i,t}$, formally given as:

$$\mathcal{L}_M = \sum_{t=0}^{N_F-1} \sum_{i \in Pos} \sum_{m \in \{cx, cy, h, w\}} \|g_{i,t}^m - \hat{g}_{i,t}^m\|_1, \quad (4)$$

where Pos is the set of positive encoded anchor tube indices. We compute $\hat{g}_{i,t}^m$ using Eq. (2) discussed in Section 4.3.

For the classification loss \mathcal{L}_C , we employ the hard negative mining strategy [26] to balance the positive and negative anchor tubes, where the negative tubes correspond to the background. We denote $x_{i,j,t}^p \in \{1, 0\}$ to be an indicator for matching the i^{th} anchor tube and the j^{th} ground-truth track of class ‘ p ’ at the t^{th} frame. Each anchor tube at least has one best-matched ground-truth track, implying $\sum_i x_{i,j,t}^p \geq 1$ for any t . The classification loss of the DMM-Net is defined as:

$$\mathcal{L}_C = \sum_{t=0}^{N_F} \left(- \sum_{i \in Pos} x_{i,j,t}^p \log(\hat{c}_{i,t}^p) - \sum_{i \in Neg} \log(\hat{c}_{i,t}^0) \right), \quad (5)$$

where $\hat{c}_{i,t}^p = \frac{\exp c_{i,t}^p}{\sum_p \exp c_{i,t}^p}$ such that $c_{i,t}^p$ is the predicted confidence of the object being for class $p : p \in \{1, \dots, N_C\}$, Neg is the set of negative encoded anchor tube indices.

We also consider the visibility estimation task fro classification viewpoint and classify each positive box into invisible ‘0’ or visible ‘1’ box. Based on that, the loss is computed as:

$$\mathcal{L}_V = \sum_{t=0}^{N_F} \left(- \sum_{i \in Pos} y_{i,j,t}^q \log(\hat{v}_{i,t}^q) \right), \quad (6)$$

where $\hat{v}_{i,t}^q = \frac{\exp v_{i,t}^q}{\sum_q \exp v_{i,t}^q}$, and $y_{i,j,t}^q$ is an indicator for matching the i^{th} anchor tube and the j^{th} ground-truth track of visibility ‘ q ’ at the t^{th} frame, such that $q \in \{0, 1\}$.

4.5 Deployment

The trained DMM-Net is readily deployable for tracking as shown in Fig. 5. The overall tracker processes $2N_F$ frames $\mathbf{I}_{t_i:t_{i+1}:2N_F}$, where the DMM-Net first selects N_F consecutive frames and outputs the predicted tubes encoded with object motion parameters, classes and visibility using Eq. (2). the tubes are decoded with Eq. (1). A filtration is then performed to remove the undesired tubes and we get tracklets (details to follow).




Fig. 5. Deployment of DMM-Net: For the $2N_F$ frames $I_{t_i:t_{i+1}:2N_F}$ from t_i to t_{i+1} , the trained DMM-Net selects N_F frames as its input and outputs predicted tubes encoded by object motion parameter matrix (O_M), classification matrix (O_C) and visibility matrix (O_V). These matrices are filtered and the track set \mathcal{T}_{t_i} is updated by associating the filtered tracklets by their IOU with the previous track set $\mathcal{T}_{t_{i-1}}$.




Fig. 6. Example of calculating Tube IOU. There are two tracklets (yellow, cyan) and their intersection is colored green. The tracklet IOU is defined as the maximum IOU of the visible box pair. Although IOU at $t = 2$ is the largest, the lined yellow object is invisible, hence, we select the second largest IOU at $t = 1$ as the tracklet IOU due to visibility of both objects.

We compute updated trajectory set \mathcal{T}_{t_i} by associating the tracklets with the previous trajectory set $\mathcal{T}_{t_{i-1}}$ using their IOUs.

To filter, we first remove tubes with confidence lower than a threshold δ_c . We subsequently perform a Tube None Maximum Suppression (TNMS). To that end, we cluster the detected tubes of the same category into multiple groups by their IOUs with a threshold δ_{nms} . Then, we only keep one tube for each group that has the maximum confidence of being positive. The kept tubes, namely ‘‘tracklets’’, are employed to update trajectory set.

We initialize our track set \mathcal{T}_{t_i} with as many trajectories as the number of tracklets. The track set is updated from t_i^{th} to t_{i+1}^{th} stamp using the Hungarian algorithm [30] applied to an IOU matrix $\Psi \in \mathbb{R}^{n'_{t_{i-1}} \times n_{t_i}}$, where $n'_{t_{i-1}}$ is the number of element in the previous track set, and n_{t_i} is the number of tracklets. Notice that we perform association of tracklets and not the individual objects. The object association remains implicit and is done inside the DMM-Net which lets us define the tubes. The association of tracklets, defined across multiple frames, leads to significant computational gain. This is one of the core reasons of the orders of magnitude gain in the tracking speed of our technique over existing methods, as will be clear in Section 5.

Table 4. Results on Omni-MOT: The symbol \uparrow indicates higher values are better, and \downarrow implies lower values are favored.

Type	Camera	IDF1 \uparrow	IDP \uparrow	IDR \uparrow	RcII \uparrow	Pren \uparrow	GT \uparrow	MT \uparrow	PT \uparrow	ML \downarrow	FP \downarrow	FN \downarrow	IDs \downarrow	FM \downarrow	MOTA \uparrow	MOTP \uparrow
Test	Camera_0	72.2%	69.3%	75.4%	90.6%	83.2%	41	38	3	0	1762	911	18	98	72.1%	79.3%
	Camera_1	59.3%	56.2%	62.7%	81.9%	73.4%	35	19	15	1	1199	730	12	61	52.0%	75.5%
	Camera_5	40.2%	61.3%	29.9%	34.7%	71.1%	44	9	20	15	2779	12855	38	80	20.4%	76.1%
	Camera_7	77.8%	79.6%	76.1%	88.3%	92.4%	37	29	8	0	1518	2457	20	99	80.9%	80.3%
	Camera_8	30.5%	47.3%	22.5%	33.5%	70.7%	38	0	28	10	1789	8548	36	110	19.3%	70.3%
Train	Camera_0	59.2%	58.6%	59.7%	87.7%	86.0%	48	40	8	0	2517	2174	43	172	73.2%	79.5%
	Camera_1	47.6%	44.2%	51.6%	80.0%	68.6%	46	29	16	1	3681	2005	50	137	42.9%	76.9%
	Camera_5	45.0%	52.7%	39.2%	57.0%	76.7%	49	6	37	6	4844	12000	122	234	39.2%	77.3%
	Camera_7	74.7%	71.9%	77.8%	90.1%	83.3%	43	31	12	0	1804	985	16	91	71.8%	79.7%
	Camera_8	41.7%	52.2%	34.7%	44.1%	66.3%	43	0	41	2	1884	4705	18	117	21.4%	68.3%
Average		55.5%	61.1%	50.9%	66.5%	79.8%	424	201	188	35	23777	47370	373	1199	49.4%	77.8%

To form the matrix Ψ , we must compute IOU between the existing tracklets and the new tracklets. Fig. 6 shows the procedure adopted for calculating the IOU between two tracklets with a simplified example. Overall, our tracker is an on-line technique in the sense that it does not use future frames to predict object trajectories. Concurrently, DMM-Net can perform tracking in real-time that makes it a highly desirable technique for practical applications.

5 Experiments

We evaluate our technique using the proposed Omni-MOT dataset and the popular UA-DETRAC challenge [47]. The former is to demonstrate both the effectiveness of our network for MOT and trainability of deep models for MOT with the proposed Omni-MOT dataset. It is emphasized that we mainly focus on ‘vehicle’ tracking in this work. The proposed dataset is also for the vehicle tracking problem. Whereas DMM-Net is trained here for vehicle tracking, it is possible to train it for e.g. pedestrian tracking. However, due to the complex (detector + tracker) modeling performed by our network, it requires much larger amount of training data that commensurate with the modeling task complexity. Existing pedestrian tracking challenges e.g. MOT15, 17 provide relatively less training data. Moreover, the evaluation servers compute results that are influenced by off-the-shelf detectors that automatically account for the covariate shift between the training and testing data. These issues do not allow a fair benchmarking of our technique on these online challenges. Hence, we focus on the vehicle tracking problem in our experiments, for which our dataset can further help in fair evaluation.

5.1 Implementation Details

We implement the DMM-Net using Pytorch [32] and train it using a single NVIDIA GeForce GTX Titan GPU. The hyper-parameter values are selected with cross-validation to maximize the MOTA metric on the validation set of UA-DETRAC. The finally chosen values of the hyperparameters are as follows. Batch size $B = 8$, number of training epochs per model = 20, number of input frames = 16, input frame size = 168×168 , We use Adam Optimizer [24] for training. Other hyper-parameter values for the overall technique are $\delta_c = 0.4$, $\delta_{nms} = 0.3$, and $\delta_o = 0.5$.

5.2 Omni-MOT dataset

To demonstrated the efficacy of DMM-Net and trainability of deep models on the proposed dataset, we select five scenes from the Omni-MOT and perform tracking. The

Table 5. Results on UA-DETRAC: ‘T.S.’ is the tracker speed (fps), ‘D.S.’ is the detector speed (fps), and ‘A.S.’ is the overall speed (fps). The proposed DMM-Net does not need explicit detector. DMM-Net+ uses a subset of Omni-MOT for data augmentation.

Tracker	Detector	PR-MOTA \uparrow	PR-MT \uparrow	PR-ML \downarrow	PR-IDS \downarrow	PR-FRAG \downarrow	PR-FP \downarrow	PR-FN \downarrow	T.S. \uparrow	D.S. \uparrow	A.S. \uparrow
CEM [1]	ACF [12]	4.50%	2.90%	37.10%	265.4	366.0	15180.3	270643.2	3.74	0.67	0.57
CMOT [51]	ACF [12]	7.80%	14.30%	20.70%	418.3	2161.7	81401.4	183400.2	3.12	0.67	0.55
DCT [2]	ACF [12]	7.90%	4.80%	34.40%	108.1	101.4	13059.7	251166.4	1.29	0.67	0.44
GOG [33]	ACF [12]	10.80%	12.20%	22.30%	3950.8	3987.3	45201.5	197094.2	319.29	0.67	0.67
H2T [49]	ACF [12]	8.20%	13.10%	21.30%	1122.8	1445.8	71567.4	189649.1	1.08	0.67	0.41
IHTLS [11]	ACF [12]	6.60%	11.50%	22.40%	1243.1	4723.0	72757.5	198673.5	5.09	0.67	0.59
CEM [1]	DPM [16]	3.30%	1.30%	37.80%	265	317.1	13888.7	270718.5	4.49	0.17	0.16
DCT [2]	DPM [16]	2.70%	0.50%	42.70%	72.2	68.8	7785.8	280762.2	2.85	0.17	0.16
GOG [33]	DPM [16]	5.50%	4.10%	27.70%	1873.9	1988.5	38957.6	230126.6	476.52	0.17	0.17
IOUT [6]	DPM [16]	1.92%	0.84%	43.70%	61.4	106.0	3111.5	290412.2	100842	0.17	0.17
CEM [1]	R-CNN [19]	2.70%	2.30%	34.10%	778.9	1080.4	34768.9	269043.8	5.4	0.1	0.10
DCT [2]	R-CNN [19]	11.70%	10.10%	22.80%	758.7	742.9	336561.2	210855.6	0.71	0.1	0.09
CMOT [51]	R-CNN [19]	11.00%	15.70%	19.00%	506.2	22551.1	74253.6	177532.6	3.59	0.1	0.10
GOG [33]	R-CNN [19]	10.00%	13.50%	20.10%	7834.5	7401.0	58378.5	192302.7	352.8	0.1	0.10
H2T [49]	R-CNN [19]	11.10%	14.60%	19.80%	1481.9	1820.8	66137.2	184358.2	2.78	0.1	0.10
IHTLS [11]	R-CNN [19]	8.30%	12.00%	21.40%	1536.4	5954.9	68662.6	199268.8	11.96	0.1	0.10
DMM-Net	-	11.80%	10.30%	15.20%	230.3	658.0	36238.8	194886.4	-	-	123.25
DMM-Net+	-	12.20%	10.80%	14.90%	228.2	674.1	36355.8	192289.6	-	-	123.25

scenes are intentionally chosen with Easy camera viewpoint with clear weather conditions. The selected scenes are from Town 02 (with 50 vehicles) that are indexed 0, 1, 5, 7 and 8 in the dataset. In this experiment, we train and test DMM-Net using the train and test partitions of the selected scene, as provided by Omni-MOT. Our training required 59.2 hours for 22 epochs.

We use both CLEAR MOT [5], and MT/ML [35] metrics and summarize the results in Table 4. For the definitions of metrics we refer to the original works. We refer to [25] for the details on the comprehensive metrics MOTA and MOTAP. The table provides metric values for both training and test partitions. The variation between these values is a clear indication that despite the Easy camera view, the dataset provides reasonable challenges for the MOT task. We provide results with additional view points in the supplementary material to further put the reported values into better perspective.

An illustration of tracking for Camera_0 is also provided in Fig. 7 (top). Our technique is able to easily track vehicles that are stationary, (e.g. 1), moving along a straight path (e.g. 2), and moving along a curved path (e.g. 3), which justifies the selection of our motion model. Additionally, our tracker has the power to deal with even full occlusions between the frames. We show such a case for a hard camera view point in the middle row of Fig. 7. In this case, vehicle-1 is totally occluded at the 86th frame which can be detected and tracked by our tracker. At the 99th frame, vehicle-1 is partially occluded, but the tracklets association performs correctly. Note that, our tracker does not need to be evaluated with (detector-based) UA-DETRAC metrics [47] as the Omni-MOT dataset provides ground truth detections without using off-the-shelf detectors.

5.3 UA-DETRAC

The UA-DETRAC challenge [47] is arguably the most widely used large-scale benchmark for MOT of vehicles. It comprises 100 videos @ 25 fps, recorded in 24 locations with frame size 540 × 960. For this challenge, results are computed by a remote server using CLEAR MOT and MT/ML metrics with Precision-Recall curve of the detection. We refer to [47] for further details on the metrics.




Fig. 7. Tracking illustration of DMM-Net on Omni-MOT (first and second row) and UA-DETRAC (third row). The mentioned IDs are for reference in the text only.

In Table 5, we show our results on the UA-DETRAC challenge. The pre-fix PR for the metrics indicates the use of Precision-Recall curve. It can be seen that DMM-Net achieves excellent MOTA score, which is widely considered as the most comprehensive metric for MOT. We also provide results for DMM-Net+ which augments the UA-DETRAC training set with a subset of Omni-MOT. This subset contained 8 random videos from Omni-MOT. We can see an overall performance gain with this augmentation, exemplifying the utility of Omni-MOT in data augmentation.

Notice that our technique does not require an external ‘detector’ and achieves highly promising PR-MOTA score for UA-DETRAC @ 120+ fps - orders of magnitude faster than the existing methods. In the third row of Fig. 7, we show an example of tracking result. The box color indicates the predicted object identity. Our tracker leverages information from previous frames to detect partially occluded objects, e.g. vehicle-3 in frame 518. The figure also shows the generalization power of our in-built detection mechanism. For instance, vehicle-1 is a very rare vehicle that is consistently assigned the correct identity. In the supplementary material, we provide further example tracking videos that clearly demonstrate successful tracking by DMM-Net for UA-DETRAC.

6 Conclusion

In the context of tracking-by-detection, we proposed a deep network DMM-Net that removes the need for an explicit external detector and performs tracking at 120+ fps to achieve 12.80% PR-MOTA value on the UA-DETRAC challenge. Our network meticulously models object motions, classes and their visibility that are subsequently used for efficiently associating the object tracks. The proposed network provides an end-to-end solution for detection and tracklet generation across multiple video frames. We also propose a realistic CARLA simulator based large-scale dataset with over 14M frames for vehicle tracking. The dataset provides precise and comprehensive ground truth with full control over data parameters, which allows for the much needed transparency in evaluation. We also provide scripts to generate more data under our framework and we make the code of DMM-Net public.

References

1. Andriyenko, A., Schindler, K.: Multi-target tracking by continuous energy minimization. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 1265–1272 (2011). <https://doi.org/10.1109/CVPR.2011.5995311> 13
2. Andriyenko, A., Schindler, K., Roth, S.: Discrete-continuous optimization for multi-target tracking. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 1926–1933 (2012). <https://doi.org/10.1109/CVPR.2012.6247893> 13
3. Bae, S.H., Yoon, K.J.: Confidence-Based Data Association and Discriminative Deep Appearance Learning for Robust Online Multi-Object Tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **40**(3), 595–610 (2018). <https://doi.org/10.1109/TPAMI.2017.2691769> 4
4. Berclaz, J., Fleuret, F., Türetken, E., Fua, P.: Multiple object tracking using k-shortest paths optimization. *IEEE TPAMI* **33**(9), 1806–1819 (2011). <https://doi.org/10.1109/TPAMI.2011.21> 3
5. Bernardin, K., Stiefelhagen, R.: Evaluating multiple object tracking performance: The CLEAR MOT metrics. *Eurasip Journal on Image and Video Processing* **2008** (2008). <https://doi.org/10.1155/2008/246309> 13
6. Bochinski, E., Eiselein, V., Sikora, T.: High-Speed tracking-by-detection without using image information. 2017 14th IEEE International Conference on Advanced Video and Signal Based Surveillance, AVSS 2017 (2017). <https://doi.org/10.1109/AVSS.2017.8078516> 13
7. Butt, A.A., Collins, R.T.: Multi-target tracking by lagrangian relaxation to min-cost network flow. In: Proc. CVPR. pp. 1846–1853 (2013) 3
8. Chari, V., Lacoste-Julien, S., Laptev, I., Sivic, J.: On pairwise costs for network flow multi-object tracking. *Proc. CVPR* **07-12-June**, 5537–5545 (2015). <https://doi.org/10.1109/CVPR.2015.7299193> 3
9. Dai, J., Li, Y., He, K., Sun, J.: R-fcn: object detection via region-based fully convolutional networks. In: NIPS’16 Proceedings of the 30th International Conference on Neural Information Processing Systems. pp. 379–387 (2016), <https://academic.microsoft.com/paper/2407521645> 4
10. Dehghan, A., Modiri Assari, S., Shah, M.: Gmmcp tracker: Globally optimal generalized maximum multi clique problem for multiple object tracking. In: Proc. CVPR. pp. 4091–4099 (2015) 1, 3
11. Dicle, C., Camps, O.I., Sznaiier, M.: The way they move: Tracking multiple targets with similar appearance. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2304–2311 (2013). <https://doi.org/10.1109/ICCV.2013.286> 13
12. Dollar, P., Appel, R., Belongie, S., Perona, P.: Fast feature pyramids for object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2014). <https://doi.org/10.1109/TPAMI.2014.2300479> 13
13. Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., Koltun, V.: CARLA: An Open Urban Driving Simulator. In: Proceedings of the 1st Annual Conference on Robot Learning. pp. 1–16 (2017), <http://arxiv.org/abs/1711.03938> 3, 4
14. Emami, P., Pardalos, P.M., Eleftheriadou, L., Ranka, S.: Machine Learning Methods for Solving Assignment Problems in Multi-Target Tracking. *arXiv:1802.06897* **1**(1), 1–35 (2018) 3
15. Feichtenhofer, C., Pinz, A., Zisserman, A.: Detect to Track and Track to Detect. Proceedings of the IEEE International Conference on Computer Vision **2017-October**, 3057–3065 (2017). <https://doi.org/10.1109/ICCV.2017.330> 4
16. Felzenszwalb, P.F., Girshick, R.B., McAllester, D., Ramanan, D.: Object detection with discriminatively trained part-based models. *IEEE TPAMI* **32**(9), 1627–1645 (2010) 13

17. Ferryman, J., Shahrokni, A.: PETS2009: Dataset and challenge. In: Proceedings of the 12th IEEE International Workshop on Performance Evaluation of Tracking and Surveillance, PETS-Winter 2009 (2009). <https://doi.org/10.1109/PETS-WINTER.2009.5399556> 4
18. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the KITTI vision benchmark suite. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 3354–3361 (2012). <https://doi.org/10.1109/CVPR.2012.6248074> 4
19. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (2014). <https://doi.org/10.1109/CVPR.2014.81> 13
20. Hara, K., Kataoka, H., Satoh, Y.: Can Spatiotemporal 3D CNNs Retrace the History of 2D CNNs and ImageNet? CVPR pp. 6546–6555 (2018). <https://doi.org/10.1109/CVPR.2018.00685>, <http://arxiv.org/abs/1711.09577> 2, 9
21. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (2016). <https://doi.org/10.1109/CVPR.2016.90> 9
22. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv (2015) 8
23. Iqbal, U., Milan, A., Gall, J.: PoseTrack: Joint multi-person pose estimation and tracking. In: Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017 (2017). <https://doi.org/10.1109/CVPR.2017.495> 4
24. Kingma, D.P., Ba, J.L.: Adam: A Method for Stochastic Optimization. In: ICLR 2015 : International Conference on Learning Representations 2015 (2015), <https://academic.microsoft.com/paper/2964121744> 12
25. Leal-Taixé, L., Milan, A., Reid, I., Roth, S., Schindler, K.: MOTChallenge 2015: Towards a Benchmark for Multi-Target Tracking. arXiv:1504.01942 [cs] pp. 1–15 (2015), <http://arxiv.org/abs/1504.01942> 1, 4, 13
26. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: SSD: Single shot multibox detector. ECCV 9905 LNCS, 21–37 (2016). https://doi.org/10.1007/978-3-319-46448-0_2 2, 7, 10
27. Luo, W., Xing, J., Milan, A., Zhang, X., Liu, W., Zhao, X., Kim, T.K.: Multiple Object Tracking: A Literature Review. arXiv:1409.7618v4 pp. 1–18 (2017). <https://doi.org/10.1145/0000000.0000000> 1, 3
28. Milan, A., Leal-Taixé, L., Reid, I., Roth, S., Schindler, K.: MOT16: A Benchmark for Multi-Object Tracking. CoRR **abs/1603.0** (2016), <http://arxiv.org/abs/1603.00831> 4
29. Milan, A., Roth, S., Schindler, K.: Continuous energy minimization for multitarget tracking. IEEE Transactions on Pattern Analysis and Machine Intelligence **36**(1), 58–72 (2014). <https://doi.org/10.1109/TPAMI.2013.103> 3
30. Munkres, J.: Algorithms for the Assignment and Transportation Problems. Journal of the Society for Industrial and Applied Mathematics **5**(1), 32–38 (1957). <https://doi.org/10.1137/0105003> 11
31. Nair, V., Hinton, G.: Rectified Linear Units Improve Restricted Boltzmann Machines. In: Proceedings of the 27th International Conference on Machine Learning (2010) 9
32. Paszke, A., Chanan, G., Lin, Z., Gross, S., Yang, E., Antiga, L., Devito, Z.: Automatic differentiation in PyTorch. Advances in Neural Information Processing Systems 30 (Nips), 1–4 (2017) 12

33. Pirsiavash, H., Ramanan, D., Fowlkes, C.C.: Globally-optimal greedy algorithms for tracking a variable number of objects. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* pp. 1201–1208 (2011). <https://doi.org/10.1109/CVPR.2011.5995604> 3, 13
34. Reid, D., et al.: An algorithm for tracking multiple targets. *IEEE transactions on Automatic Control* **24**(6), 843–854 (1979) 3
35. Ristani, E., Solera, F., Zou, R., Cucchiara, R., Tomasi, C.: Performance measures and a data set for multi-target, multi-camera tracking. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. vol. 9914 LNCS, pp. 17–35 (2016). https://doi.org/10.1007/978-3-319-48881-3_2 13
36. Ristani, E., Tomasi, C.: Tracking Multiple People Online and in Real Time. In: *Asian Conference on Computer Vision*. pp. 444–459. Springer (2014) 4
37. Roshan Zamir, A., Dehghan, A., Shah, M.: GMCP-tracker: Global multi-object tracking using generalized minimum clique graphs. In: *Lecture Notes in Computer Science* (2012). https://doi.org/10.1007/978-3-642-33709-3_25 3
38. Schuster, S., Vernaza, P., Choi, W., Chandraker, M.: Deep Network Flow for Multi-Object Tracking. *Cvpr* pp. 6951–6960 (2017). <https://doi.org/10.1109/CVPR.2017.292> 3
39. Shafique, K., Shah, M.: A noniterative greedy algorithm for multiframe point correspondence. *IEEE transactions on pattern analysis and machine intelligence* **27**(1), 51–65 (2005) 3
40. Sheng, H., Zhang, Y., Chen, J., Xiong, Z., Zhang, J.: Heterogeneous association graph fusion for target association in multiple object tracking. *IEEE Transactions on Circuits and Systems for Video Technology* (2018) 3
41. Shitrit, H.B., Berclaz, J., Fleuret, F., Fua, P.: Multi-commodity network flow for tracking multiple people. *IEEE TPAMI* **36**(8), 1614–1627 (2014) 3
42. Shu, G., Dehghan, A., Oreifej, O., Hand, E., Shah, M.: Part-based multiple-person tracking with partial occlusion handling. In: *Proc. CVPR*. pp. 1815–1821. IEEE (2012) 3
43. Son, J., Baek, M., Cho, M., Han, B.: Multi-Object Tracking with Quadruplet Convolutional Neural Networks. *Cvpr* pp. 5620–5629 (2017). <https://doi.org/10.1109/CVPR.2017.403> 4
44. Sun, S., Akhtar, N., Song, H., Mian, A., Shah, M.: Deep Affinity Network for Multiple Object Tracking **13**(9), 1–15 (2018), <http://arxiv.org/abs/1810.11780> 1, 4, 6
45. Tian, Y., Dehghan, A., Shah, M.: On detection, data association and segmentation for multi-target tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence* pp. 1–1 (2018) 1, 3
46. Voigtlaender, P., Krause, M., Osep, A., Luiten, J., Sekar, B.B.G., Geiger, A., Leibe, B.: Mots: Multi-object tracking and segmentation. In: *Proc. CVPR*. pp. 7942–7951 (2019) 1
47. Wen, L., Du, D., Cai, Z., Lei, Z., Chang, M.C., Qi, H., Lim, J., Yang, M.H., Lyu, S.: UA-DETRAC: A New Benchmark and Protocol for Multi-Object Detection and Tracking (2015), <http://arxiv.org/abs/1511.04136> 3, 4, 12, 13
48. Wen, L., Du, D., Li, S., Bian, X., Lyu, S.: Learning non-uniform hypergraph for multi-object tracking. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 33, pp. 8981–8988 (2019) 3
49. Wen, L., Li, W., Yan, J., Lei, Z., Yi, D., Li, S.Z.: Multiple target tracking based on undirected hierarchical relation hypergraph. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* **1**, 1282–1289 (2014). <https://doi.org/10.1109/CVPR.2014.167> 3, 13
50. Wu, B., Nevatia, R.: Detection and tracking of multiple, partially occluded humans by bayesian combination of edgelet based part detectors. *International Journal of Computer Vision* **75**(2), 247–266 (2007) 3
51. Zhu, J., Yang, H., Liu, N., Kim, M.: Online Multi-Object Tracking with Dual Matching Attention Networks pp. 1–17 13