

# Discriminative Appearance Modeling with Multi-track Pooling for Real-time Multi-object Tracking

Chanho Kim<sup>1\*</sup>   Li Fuxin<sup>2</sup>   Mazen Alotaibi<sup>2</sup>   James M. Rehg<sup>1</sup>

<sup>1</sup>Georgia Institute of Technology   <sup>2</sup>Oregon State University

## Abstract

*In multi-object tracking, the tracker maintains in its memory the appearance and motion information for each object in the scene. This memory is utilized for finding matches between tracks and detections, and is updated based on the matching. Many approaches model each target in isolation and lack the ability to use all the targets in the scene to jointly update the memory. This can be problematic when there are similarly looking objects in the scene. In this paper, we solve the problem of simultaneously considering all tracks during memory updating, with only a small spatial overhead, via a novel multi-track pooling module. We additionally propose a training strategy adapted to multi-track pooling which generates hard tracking episodes online. We show that the combination of these innovations results in a strong discriminative appearance model under the bilinear LSTM tracking framework, enabling the use of greedy data association to achieve online tracking performance. Our experiments demonstrate real-time, state-of-the-art online tracking performance on public multi-object tracking (MOT) datasets. The code and trained models are available at <https://github.com/chkim403/blstm-mtp>.*

## 1. Introduction

In the typical tracking-by-detection setting of multi-object tracking, a trained target detector is assumed to exist, and the goal of the tracking algorithm is to solve the *data association* problem: associating detections from different frames into tracks. The standard approach involves building an appearance model and a motion model for each target being tracked. The appearance model stores features from past appearances, which are compared against detections in each new frame, and the motion model predicts where the target will be located in the next frame. Because of the ambiguity in appearance, many previous trackers utilized sophisticated

data association schemes that relied on evidence from multiple frames to determine the correct track. These choices result in a tracker which is not capable of *real-time online* performance, i.e. it cannot generate a result *without delay* after processing a new frame in the video. This is unfortunate, because real-time online tracking is critically important for numerous practical applications in AI. For example, in robotics applications such as autonomous driving, online processing is critical to support real-time decision-making.

One issue that has necessitated sophisticated data association schemes is that the appearance models are usually built to match tracks and detections individually, without considering other tracks and detections in the matching process. This approach does not have sufficient discriminative power to address the ambiguities arising when several similar targets move adjacent to each other in the video, e.g. multiple men in black suits. In this situation, either more subtle features need to be utilized for target discrimination, or the likelihood of the matching needs to be decreased to accommodate this uncertainty. This in turn requires the matching approach to utilize the appearance information *from other nearby tracks* in making a determination.

This paper introduces a novel approach to appearance modeling that takes *all* tracked objects into account when matching detections. Suppose each track has a stored online memory about all of its past appearances, we propose a novel *multi-track pooling module* which stores a max-pooled version of the memory from all other tracks. This extension allows the appearance model to take into account **online** negative examples from other objects within the same video (see Fig. 1). We show that multi-track pooling greatly enhances the discriminative power of the appearance model for track scoring and improves overall performance.

We leverage the additional discriminative power of our matching approach to achieve real-time online tracking by means of a simple, greedy data association method: Each track is matched with the detection that has the highest likelihood of belonging to the track. Greedy association is extremely efficient, resulting in fast performance. In this work we test the hypothesis that our discriminative appearance and motion modeling can enable this simple data associa-

\* Part of this work was conducted while Chanho Kim was visiting Oregon State University.

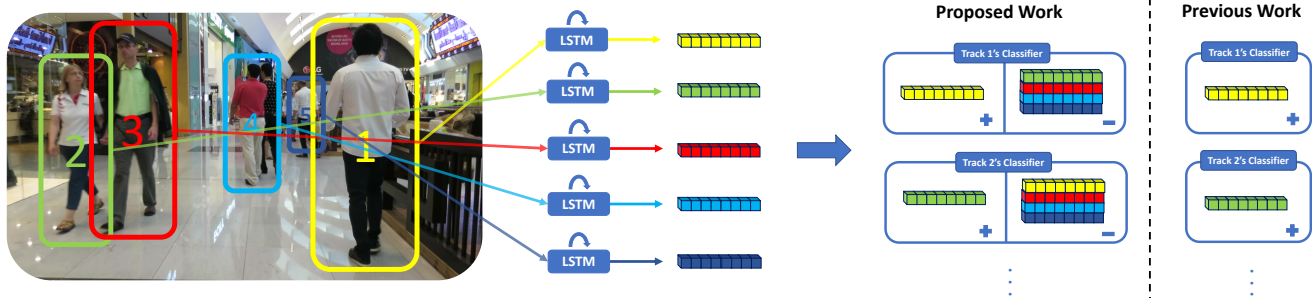


Figure 1. Existing recurrent neural network-based track classifiers used only matched detections for updating its appearance memory during tracking. This does not consider other objects in the scene (i.e. negative examples), which may have similar appearances. We propose to improve the predicted likelihood of such a classifier by augmenting its memory with appearance information about other tracks in the scene with multi-track pooling, leveraging the appearance information from the full set of tracks in the scene. The resulting classifier learns to adapt its prediction based on the information from other tracks in the scene.

tion mechanism to achieve strong tracking performance.

Our online tracking framework additionally incorporates components from recent tracking works. First, we utilize the Bilinear LSTM framework [17] as the basis for our matching approach, due to its effective memory construction. Second, we incorporate training strategies that handle long tracks using truncated backpropagation through time [32], in contrast to prior works [31, 17] in which the models were trained with short tracks. Third, we extend object tracks to frames without detections using a motion model, thereby compensating for missing detections. Fourth, we trained a bounding box corrector to correct bounding box coordinates, which is especially helpful for the extended tracks.

In summary, this paper makes the following contributions:

1. A novel multi-track pooling module which enables a track classifier to take online negative examples into account during testing and thus adjust its prediction adaptively depending on the objects in the scene.
2. A training strategy that is adapted to train the proposed track pooling module by utilizing within-batch dependencies among tracks and that enables long sequence training with the original full tracks instead of short tracks used in [31, 17].
3. A real-time, online multi-object tracking algorithm that achieves state-of-the-art online tracking performance on standard tracking benchmarks.

## 2. Related Work

Two groups of prior works have explored the incorporation of positive and negative samples during on-the-fly testing, and these are the closest related works to this paper. One line of work incorporates these samples by fine-tuning

a pretrained CNN using positive examples (target object) and negative examples (background in [27], other objects in [23]) during testing. While these approaches share our interest in utilizing scene-specific information during tracking, the need to fine-tune during testing adds an additional source of complexity and is a barrier to efficient online performance. In contrast, our model automatically adjusts its prediction based on scene-specific information without the need for fine-tuning. The second line of work uses relational networks or graph convolutional networks (GCNs) to incorporate the appearance of other detections in the same frame [35] and in neighboring frames [3] when computing the appearance features of each detection. However, [3] operates in a batch setting whereas our approach is online. In addition, our multi-track pooling method is significantly simpler and faster than graph convolutional networks which require multiple iterations of message passing.

We utilize the Bilinear LSTM architecture of [17] in developing our matching approach. We extend beyond this work in multiple ways, the primary difference being the introduction of a novel multi-object pooling approach which utilizes appearance information across tracks to significantly improve data association performance. We demonstrate that this makes it feasible to use a much simpler and more cost-effective matching algorithm following track scoring, achieving real-time multi-object tracking.

[31, 26] presented an LSTM-based track proposal classifier that integrates motion, appearance, and interaction cues. [40] presented a new attention network where attentions are computed both spatially and temporally to predict binary labels for track proposals. [11] proposed recurrent autoregressive networks where they learn generative models for both motion and appearance, and these generative models are used in data association to determine how likely it is that a detection belongs to a track. The primary difference between our work and these prior approaches is that our model makes a prediction for a detection-track pair by tak-

ing appearances of all tracks into account.

[28] and [6] exploited interactions between tracks in solving the multi-object tracking problem. [28] incorporated a social behavior model into their tracking algorithm. The social behavior model is based on the assumption that each person moves in such a way as to avoid collisions with other people. [6] incorporated high-level human activity cues into their tracking algorithm by exploiting the fact that human activities can influence how people move in the scene. In contrast to these works, our work focuses on incorporating multiple appearances from all tracks into our model, in order to make it more discriminative.

### 3. Track Proposal Classifier

Tracking-by-detection approaches in multi-object tracking evaluate multiple track proposals [5, 16, 33, 31, 17] when finding correspondences between tracks and detections. Track proposals are typically constructed by extending existing tracks from the previous frame with new detections in the current frame. Let us denote  $T_l(t)$  as the  $l^{\text{th}}$  track at time  $t$ . Let  $s(l)$  be the starting frame of the  $l^{\text{th}}$  track and  $d_t^l$  be the detection selected by the  $l^{\text{th}}$  track at time  $t$ . We then write the  $l^{\text{th}}$  track at time  $t$  as  $T_l(t) = \{d_{s(l)}^l, d_{s(l)+1}^l, \dots, d_{t-1}^l, d_t^l\}$ . Recurrent networks are trained to output the following conditional probability:

$$f(d_t, T_l(t-1); \theta) = p(d_t \in T_l(t) | T_l(t-1)) \quad (1)$$

where  $f(\cdot)$  and  $\theta$  represent a neural network and its learnable parameters respectively. The input to the neural network is the detection-track pair  $(d_t, T_l(t-1))$ .

#### 3.1. Bilinear LSTM

Vanilla LSTM equations are defined as follows [15]:

$$\begin{aligned} \mathbf{f}_t &= \sigma(\mathbf{W}_f[\mathbf{h}_{t-1}; \mathbf{x}_t]), & \mathbf{i}_t &= \sigma(\mathbf{W}_i[\mathbf{h}_{t-1}; \mathbf{x}_t]), \\ \mathbf{g}_t &= \sigma(\mathbf{W}_g[\mathbf{h}_{t-1}; \mathbf{x}_t]), & \mathbf{o}_t &= \tanh(\mathbf{W}_o[\mathbf{h}_{t-1}; \mathbf{x}_t]), \\ \mathbf{c}_t &= \mathbf{f}_t \circ \mathbf{c}_{t-1} + \mathbf{i}_t \circ \mathbf{g}_t, & \mathbf{h}_t &= \mathbf{o}_t \circ \tanh(\mathbf{c}_t) \end{aligned} \quad (2)$$

where  $[\cdot]$  denotes concatenation of 2 column vectors. In the sequence labeling problem,  $\mathbf{h}_t$  stores the information about the sequence and is fed into additional fully-connected or convolutional layers to generate the final prediction. The input  $\mathbf{x}_t$  and the LSTM memory  $\mathbf{h}_{t-1}$  are combined by additive interactions in the above equations.

There are two issues with this formulation. First, the *matching* operation is usually more easily represented by multiplicative relationships instead of additive, an intuitive example being the inner product as a correlation metric. Second, it is difficult to store and clearly distinguish multiple different appearances in the same LSTM memory vector, but a track exhibiting multiple different appearances in different frames is very common in multi-target tracking.

Bilinear LSTM [17] solves these issues by introducing a new representation based on the multiplicative interaction between the input and the LSTM memory:

$$\begin{aligned} \mathbf{h}_{t-1} &= [\mathbf{h}_{t-1,1}^\top, \mathbf{h}_{t-1,2}^\top, \dots, \mathbf{h}_{t-1,r}^\top]^\top \\ \mathbf{H}_{t-1}^{\text{reshaped}} &= [\mathbf{h}_{t-1,1}, \mathbf{h}_{t-1,2}, \dots, \mathbf{h}_{t-1,r}]^\top \\ \mathbf{m}_t &= g(\mathbf{H}_{t-1}^{\text{reshaped}} \mathbf{x}_t) \end{aligned} \quad (3)$$

where  $g(\cdot)$  is a non-linear activation function. A long vector  $\mathbf{h}_{t-1}$  from LSTM is reshaped into a matrix  $\mathbf{H}_{t-1}^{\text{reshaped}}$  before it being multiplied to the input  $\mathbf{x}_t$ . Hence, multiple memory vectors can be matched with the feature from the detection with an inner product. The new matching vector  $\mathbf{m}_t$  is then fed into additional fully-connected layers to generate the final prediction. Note that the way that the LSTM memory  $\mathbf{h}_{t-1}$  interacts with the input  $\mathbf{x}_t$  is changed in Bilinear LSTM, the output gate is not used, while standard LSTM memory updates in Eq. (2) are used since the memory vector  $\mathbf{h}_{t-1}$  is retained as in Eq. (3).

Bilinear LSTM bears some resemblance to the popular transformers model [34] in natural language processing in that both utilize a multiplicative relationship between the sequence and a new token, but they have some important differences. In a transformer model, the inner product is taken between the key of a new token with *all* previous tokens in the sequence. This makes them very memory-inefficient and unsuitable for online operations that span hundreds of frames in multi-object tracking. In Bilinear LSTM, the memory size is fixed, and the memory is updated automatically using the LSTM updating rule. This has the effect of automatically grouping similar appearances into the same row of the memory  $\mathbf{H}_{t-1}^{\text{reshaped}}$  in Eq. (3), so that the memory size does not grow linearly with respect to the sequence length. Hence, we believe that Bilinear LSTM is a suitable choice for online tracking in which multiple appearances for each track need to be stored.

#### 3.2. Application to Multi-object Tracking

When Bilinear LSTM is used in multi-object tracking, each track  $T_l(t-1)$  will have its own LSTM memory  $\mathbf{h}_{t-1}^l$  which is stored during the tracking process. All new detections at frame  $t$  go through a CNN to generate their corresponding  $\mathbf{x}_t$ , which are then used to compare with  $\mathbf{h}_{t-1}^l$  for all the tracks. When each detection has been scored with each track, the detections will be assigned to the existing tracks by either greedy assignment, or multiple hypothesis tracking (MHT) assignment [16]. Finally, the features from the assigned bounding box are used as  $\mathbf{x}_t^l$  to update the track memory with Eq. (2). The updated memory  $\mathbf{h}_t^l$  will be stored and then the same process will be repeated in the next frame. All tracks share the same LSTM network as their appearances will be dynamically updated in the memory, hence there is no re-training needed for any new object.




Figure 2. Bilinear LSTM and the proposed improvements by multi-track pooling module. Bilinear LSTM stores multiple memory vectors for each track (in blue) so that a new detection  $x_t$  can be matched with multiple templates via inner products. In this work, we improve on it by concatenating the memory with multiple stored memory vectors from other simultaneously tracked targets (in red) to serve as negative examples, so that the template matching process can take into account more subtle differences between the positive track and negative tracks. (best viewed in color)

### 3.3. Multi-Track Pooling Module

A limitation of the previous work [31, 17] is that only past appearances of the same track were considered, as only matched detections were inputted as  $x_t^l$  to update the LSTM memory at time  $t$ . However, during tracking, different targets can have similar appearances. For example, in pedestrian tracking, there could always be many people wearing similar white shirt and black pants, or black suits and black pants. These people need to be distinguished via more detailed features, such as shoes, objects in hand, etc. Thus, simple appearance matching may not be discriminative enough.

We propose to extend the track scoring process to consider **all** the tracked objects, other than only the current object of concern. Instead of simply growing the memory to store more templates, which can hit a memory bottleneck, we jointly consider all the objects that have been tracked. Tracked objects are usually the ones most easily confused with the target (if a detection does not come from a track, then its appearance is likely significantly different from any tracked objects (e.g. pedestrians)). Hence, taking the appearance information about these objects into consideration could greatly improve the discriminative power of the tracker for each target (see Fig. 2).

In order to consider other tracked objects as well, we propose to modify Eq. (1) to:

$$f(d_t, T_{1:M}(t-1); \theta) = p(d_t \in T_l(t) | T_{1:M}(t-1)) \quad (4)$$

where  $T_{1:M}(t-1) = \{T_1(t-1), T_2(t-1), \dots, T_M(t-1)\}$  represents the existing tracks in the previous frame. In this paper, we denote  $M$  as the number of tracks in the previous frame  $t-1$  and denote  $N$  as the number of detections in the current frame  $t$ . [31] also trained recurrent neural networks to output this kind of conditional probability by using track

interaction cues, but their interaction cues were the most recent locations of other tracks, not appearances of other tracks.

In Bilinear LSTM, each track in the previous frame is associated with a unique LSTM memory  $h_{t-1}^l$ . When new detections  $x_t$  arrive in the current frame, we compute  $m_t$  in Eq. (3) between each of the  $M$  existing tracks and each of the  $N$  new detections. We denote  $m_t^+$  as the matching vector between  $h_{t-1}^l$  computed using the target track (i.e. the current object of concern) and  $x_t$  (as computed by Eq. (3)). For the other  $M-1$  tracks in the scene, we denote  $M_t^-$  as a matrix in which each row represents  $m_t$  computed by a non-target object track (i.e. other objects in the scene) and the new detection ( $x_t$ ):

$$M_t^- = [m_{t,1}^-, m_{t,2}^-, \dots, m_{t,M-1}^-]^\top \quad (5)$$

where  $m_{t,i}^-$  represents the matching between the  $i$ -th non-target object track and the new detection.

Here the main difficulty is that we have one positive track and an indefinite number of negative tracks. We propose to compress  $M-1$   $m_t$ s using max pooling in order to obtain a fixed size matching representation  $m_t^-$  regardless of the number of the tracks in the previous frame:

$$m_t^-(j) = \max_i M_t^-(i, j) \quad (6)$$

where  $m_t^-(j)$  represents the  $j$ -th element in  $m_t^-$ , and  $M_t^-(i, j)$  represents an element located in the  $i$ -th row and  $j$ -th column of  $M_t^-$ . Thus, the  $j$ -th element in  $m_t^-$  is computed by taking the maximum of the  $j$ -th column of  $M_t^-$ .

Since the Bilinear LSTM matching vector stores correlation responses between multiple templates and the input, applying max pooling to  $M_t^-$  allows us to detect high correlation responses generated by non-target object tracks for the detection which is currently considered. Thus, values in  $m_t^-$  can show whether any of the non-target object tracks has a similar appearance to that of the current detection. We obtain the final matching representation for a track classifier by concatenating  $m_t^+$  and  $m_t^-$  as:

$$m_t^{\text{all}} = [(m_t^+); (m_t^-)] \quad (7)$$

The network runs a fully-connected layer after  $m_t^{\text{all}}$  followed by softmax, that outputs the binary decision that determines whether the new detection belongs to the target track ( $m_t^+$ ).

Table 1 shows the proposed network architecture. In practice we compute the Bilinear LSTM representation  $m_t$  for all the tracks and then construct  $M_t^-$  by simply stacking the precomputed vectors into a matrix, which can be done efficiently. We adopt ResNet 50 [14] as our convolutional neural network and use Bilinear LSTM [17] as our recurrent network. In addition to the proposed appearance



		Soft-max	2		
		FC	2		
		Concatenation ( $m_1^2$ )	16		
		Matrix-vector Multiplication-relu ( $m_1^2$ )	8	Max-pooling ( $m_1^2$ )	8
				Matrix-vector Multiplication-relu ( $M_1$ )	$(M-1) \times 8$
Reshape	$8 \times 256$	Reshape	$256 \times 1$	Reshape	$(M-1) \times 8 \times 256$
LSTM	2048			LSTM	$(M-1) \times 2048$
FC-relu	256	FC-relu	256	FC-relu	$(M-1) \times 256$
ResNet-50	2048	ResNet50	2048	ResNet-50	$(M-1) \times 2048$
$x_{t-1}$	$128 \times 64 \times 3$	$x_t$	$128 \times 64 \times 3$	$x_{t-1:t}$	$(M-1) \times 128 \times 64 \times 3$
				$x_t$	$128 \times 64 \times 3$

Table 1. Proposed Network Architecture for the Track Proposal Classifier. The right two columns represent the multi-track pooling module where  $(M - 1)$  other tracks are processed.

model, we adopt the motion model in [17] as the learned motion model. It is a binary classifier based on a traditional LSTM that receives the bounding box coordinates as input. We combine the appearance model and motion model to form a joint model that utilizes both appearance and motion cues. We used this joint model to generate the results in Table 6, 7, and 8. More details about the joint model are included in the supplementary.

## 4. Training

In this section, we describe the training strategy for the proposed neural network architecture by sampling mini-batches so that tracks within the mini-batch are correlated, in order for the multi-track pooling module to train effectively. We explain the method that generates our training data consisting of both *actual* multi-object tracking episodes and *random* tracking episodes from public multi-object tracking datasets [18, 25] and then present our choice of loss function.

### 4.1. Actual Tracking Episodes as Training Data

We generate actual tracking episodes by processing ground truth track labels sequentially and generating track proposals every frame in an online manner. Specifically, when we have  $M$  tracks in the previous frame and  $N$  detections in the current frame, we generate  $MN$  track-detection pairs in the current frame by considering all possible pairs between the existing tracks and the new detections. Each proposal of a track-detection pair is associated with a binary label where positive label indicates a match and negative represents that the track and detection belong to different objects. We use these  $MN$  proposals as a mini-batch for each training iteration and repeat this process until we process all the frames in video. Then we move on to the next training video. This process is repeated over all training videos during training.

**Truncated backpropagation through time.** As we process training videos sequentially, tracks become too long to fit into GPU memory. One way to avoid this is to split tracks into shorter ones, which is similar to the training data used in [31, 17] for recurrent models. However, this is sub-optimal because a recurrent model would only be able to learn temporal structures with an extent limited by the max-

imum sequence length that the model sees during training.

In order to address this issue, we adopt truncated back-propagation through time (BPTT) [32]. In truncated BPTT, the maximum number of time steps where the loss is back-propagated is limited by a fixed-size time window. However, entire sequences are processed by a recurrent model through multiple training iterations by a moving time window. Since long sequences are processed through multiple training iterations, the recurrent memories that store the sequence information need to be shared across the training iterations. This allows the model to access the temporal context beyond the temporal extent which is determined by the size of the time window. By training a recurrent model with the original long tracks, the model has a better chance to learn how to exploit long-term information when solving the track proposal classification problem.

### 4.2. Random Tracking Episodes as Training Data

We also use short random track proposals, which is similar to the training data used in [31, 17], as our additional training data. We generate short random track proposals as follows. At each training iteration, we first pick a video where we generate tracks and randomly select the start frame and end frame for obtaining track proposals. We take all the tracks that exist in the end frame. Denote  $N_{\max}$  as the maximum number of tracks that we use to construct random track proposals. If the number of the tracks in the end frame is greater than  $N_{\max}$ , we randomly pick  $N_{\max}$  tracks among them. In order for each mini-batch to have tracks of different lengths, we randomly clip the selected tracks such that each track in the batch start in different frames.

We take  $N$  detections from the selected tracks in the end frame and take  $M$  tracks from the previous frame to form track proposals ( $M < N$  when new tracks are born in the selected end frame. Otherwise,  $M = N$ ). We generate  $MN$  track-detection pairs by considering all possible matchings between these two. Thus,  $M$  track-detection pairs are associated with a positive label (same object), and  $MN - M$  pairs are associated with a negative label.

### 4.3. Loss Function

For each mini-batch, there are  $MN$  track-detection pairs, each of which incurs a loss value. Thus, our cross-entropy loss is written as:

$$L(t) = \frac{1}{MN} \sum_{i=1}^M \sum_{j=1}^N \alpha_{ij}(t) L_{ij}(t) \quad (8)$$

where  $t$  represents a frame number in video, and  $\alpha_{ij}$  is a weighting factor. The cross-entropy loss term for each train-

ing example  $L_{ij}(t)$  is defined as:

$$L_{ij}(t) = \begin{cases} -\log p(d_{jt} \in T_i(t)|T_{1:M}(t-1)), & \text{if } y_{ij}(t) = 1 \\ -\log(1 - p(d_{jt} \in T_i(t)|T_{1:M}(t-1))), & \text{otherwise} \end{cases} \quad (9)$$

where  $d_{jt}$  represents a  $j$ th detection in  $t$ , and  $y_{ij}(t)$  is a ground truth binary label representing whether  $d_{jt}$  belongs to a track  $T_i(t)$  or not. For the weighting factor  $\alpha_{ij}(t)$ , we adopted the weighting factor of Focal loss [20] to address the class imbalance of our training data (i.e. there are much more negative examples than positive examples). The weighting factor is then written as:

$$\alpha_{ij}(t) = \begin{cases} \beta_+(1 - p(d_{jt} \in T_i(t)|T_{1:M}(t-1)))^2, & \text{if } y_{ij}(t) = 1 \\ \beta_-(p(d_{jt} \in T_i(t)|T_{1:M}(t-1)))^2, & \text{otherwise} \end{cases} \quad (10)$$

where  $\beta_+$  and  $\beta_-$  are class-specific constant weights which are found empirically as suggested in [20]. For positive labels, we used  $\beta_+ = 4$ , and, for negative labels (i.e.  $d_{jt} \notin T_i(t)$ ), we used  $\beta_- = 1$ .

## 5. Tracking Algorithm

In this work, we choose to use the greedy association because it is the fastest association algorithm, and our training setting simulates such a greedy association algorithm (i.e. the loss value was calculated for each track independently and then averaged across all tracks in a mini-batch). The greedy data association algorithm runs in  $O(MN)$  time. A more complicated scheme would require additional computational complexity (e.g. the Hungarian algorithm runs in  $O(MN^2)$ ) and would require more careful parameter tuning for the tracker.

### 5.1. Greedy Data Association

Our greedy data-association works as follows. It initializes new tracks with the detections in the first frame. The tracker generates and stores the LSTM memory  $\mathbf{h}_{t-1}$  for each of the tracks. When new detections  $\mathbf{x}_t$  arrive in the next frame, we compute the association likelihood for every possible track-detection pair using the track proposal classifier. We set the threshold for the association likelihood to 0.5. The data association problem is then solved in a greedy manner starting from the highest matching likelihood. The tracks are updated with the newly assigned detections, and new tracks are born from the detections which are not associated with any of the existing tracks. The tracker updates the LSTM memory from  $\mathbf{h}_{t-1}$  to  $\mathbf{h}_t$  for every track according to the data association result. This process is repeated until all the video frames are processed.

We adopt a simple track termination strategy which terminates tracks if they have more missing detections than actual detections or the number of consecutive missing detections becomes larger than a threshold  $N_{\text{miss}}$ . The terminated tracks will not be used in the data association process anymore and will not be used for computing non-target object memories either.

### 5.2. Track Extension and Bounding Box Correction

In online tracking, we attempt to extend a track by generating additional detections. When there is no detection in a predicted target location in the current frame, we generate a new detection bounding box using the location predicted by a Kalman filter and then use our track proposal classifier to decide whether or not the newly generated detection belongs to the current target. Since we rely on motion cues when generating these additional detections, the bounding box coordinates of the newly generated detections might not be accurate. Thus, we train a bounding box corrector which predicts the correct bounding box coordinates based on the CNN features, which is similar to the bounding box regression module presented in [13]. More details are included in the supplementary material.

## 6. Experiments

We tested the proposed method on the MOT 16 and MOT 17 (Multiple Object Tracking) Challenge Benchmarks [25]. We trained on the MOT 17 training sequences, as well as additional public training sequences including 7 MOT 15 training sequences [18] which were not included in MOT 17 training/testing sequences to our training data. In contrast to several recent work [33, 31, 17], we did not utilize the Person Re-identification datasets [19, 38] to pretrain our CNN.

For our ablation study, we used two validation sets. The first one is the MOT 19 Challenge training sequences [10] which have 2,390 annotated object tracks. In contrast to the MOT 17 sequences, this new challenge dataset provides heavily crowded scenes captured in new environments, which makes it good for validating the proposed appearance model. Note that this dataset was only used for our ablation study and thus was not used to train our model in Table 6, 7, and 8. The second validation set is the MOT 17 Challenge training sequences which have 512 annotated object tracks. When the second validation set is used, we used the MOT 15 training sequences as our training data. See the supplementary document for the video sequence names used in the training, validation, and test sets.

We used the standard MOT metrics such as IDF1 [29], IDS, MOTA [2], Mostly Tracked (MT), Mostly Lost (ML), and Fragmentation (Frag) [25] for performance comparison in our experiments.

Method	MOTA	IDF1	IDS	Rcll	Prcn	MT	ML	Frag
B-LSTM	44.8	31.3	15,367	46.1	99.8	12.6	27.0	38,182
Ours	44.9	<b>35.0</b>	<b>11,940</b>	45.9	99.8	12.7	27.5	<b>37,017</b>

Table 2. Performance comparison on MOT 19 train sequences (val1) when motion gating is not used

Method	MOTA	IDF1	IDS	Rcll	Prcn	MT	ML	Frag
B-LSTM	49.1	52.5	1,112	51.7	96.9	21.1	31.9	1,066
Ours	49.4	<b>53.9</b>	<b>809</b>	51.6	97.1	21.1	31.5	1,070

Table 3. Performance comparison on MOT 17 train sequences (val2) when motion gating is not used

Method	MOTA	IDF1	IDS	Rcll	Prcn	MT	ML	Frag
B-LSTM	45.1	39.6	9,137	45.9	99.8	12.6	27.6	36,379
Ours	45.0	<b>40.5</b>	<b>7,873</b>	45.7	99.8	12.6	28.1	<b>35,169</b>

Table 4. Performance comparison on MOT 19 train sequences (val1) when the simple motion gating strategy is used

Method	MOTA	IDF1	IDS	Rcll	Prcn	MT	ML	Frag
B-LSTM	49.3	56.7	847	51.6	97.1	21.1	32.1	1,038
Ours	49.6	56.8	<b>616</b>	51.6	97.3	21.1	32.1	1,040

Table 5. Performance comparison on MOT 17 train sequences (val2) when the simple motion gating strategy is used

## 6.1. Ablation Study

In the ablation study, we show the effectiveness of the proposed multi-track pooling module by comparing its performance to the original Bilinear LSTM. Firstly, we tested their performance without using any motion cues during tracking. Secondly, we used motion cues during tracking by adopting a simple motion gating strategy that allows detections that are close to the current track to be considered as a possible matching.

**Data Association.** We ran the greedy data association algorithm described in the previous section with the following hyperparameter setting: 0.5 as the association threshold and  $N_{\text{miss}} = 60$ . In the ablation study, we did not interpolate missing detections in the final tracks using our track extension module in order to make the MOTA scores (also, Recall, Precision, MT, and ML) close across different methods. This allowed us to compare other tracking metrics in a fairer setting.

**Comparison with Bilinear LSTM.** We examined the effect of the proposed multi-track pooling module on the tracking performance. For this ablation study to be fair, both Bilinear LSTM and our method were trained on the same training set (see the supplementary document) with the same training setup. Table 2 and 3 show the tracking results when no motion cues were used during tracking. In this case, the appearance model needed to do the heavy lifting. Table 4 and 5 show the tracking results when a simple motion gating strategy was applied. Bilinear LSTM with the multi-track pooling module consistently outperformed the original Bilinear LSTM on IDF1, IDS and Fragmenta-

tions, showing its effectiveness in multi-object tracking on both of our validation sets.

In Table 2 - 5, the differences on IDF1 between the B-LSTM baseline and the proposed approach were relatively small compared to the differences on IDS. We suspect that this is a result of a mismatch between the binary classification loss and the IDF1 metric. As pointed out by [24], minimizing the binary classification loss does not always guarantee the model to perform better on IDF1.

## 6.2. MOT Challenges

We evaluated both the online and near-online versions of our tracker for the MOT 17/16 Benchmarks. In the online version, we utilized the track extension module described in Sec. 5.2 to recover missing detections (except for the online tracker in Table 6 in which we turned off both the extension module and the bounding box corrector). In the near-online version, we performed local track smoothing to recover missing detections instead of using the track extension module. For the second case, we denote the method as near-online in Table 7 and 8 since local track smoothing requires lookahead frames.

Recent approaches [36, 21] utilized Tracktor [1] to first refine public detections, which resulted in higher scores due to more accurate detections. In order to compare with these recent approaches, we also used public detections processed by Tracktor as input to our tracker and presented the comparison in Table 6 separately. For this experiment, we used the public detections provided by [3]. In Table 6, it can be seen that our approach significantly improves the IDF1 score and identity switches over other online tracktor-based approaches. We also compared our performance with [39] which also refined public detections using CenterNet. Note that CenterNet offers a significantly better refinement module than Tracktor, shown by their high detection accuracy in MOTA with significantly less false negatives, despite higher identity switches. In this case, a direct comparison cannot be made since a different method was used to refine the public detections. Note that our performance on IDF1 and ID switches still remains strong compared to [39].

In Table 7 and 8, we did not utilize Tracktor and compared our results with other online and near-online trackers which did not utilize Tracktor. Again our greedy tracker is the fastest among the top performing trackers and our performance is comparable with the best trackers. Considering its simplicity and speed, we believe our method demonstrates strong state-of-the-art performance on the MOT Challenge.

In near-online trackers, we significantly improve over our baseline MHT-bLSTM on both IDF1 (by 7.5%) and MOTA (by 12.8%), obtaining the best performance in near-online tracking. We also have the smallest amount of ID switches, fragmentations and the highest amount of mostly

Method	Type	IDF1	MOTA	IDS	MT	ML	Frag	FP	FN	Hz
CTTrackPub [39]	online	<b>59.6</b>	<b>61.5</b>	2,583	<b>26.4</b>	<b>31.9</b>	4,965	14,076	<b>200,672</b>	<b>17.0</b>
GSM-Tracktor [21]	online	57.8	<b>56.4</b>	<b>1,485</b>	<b>22.2</b>	<b>34.5</b>	<b>2,763</b>	14,379	<b>230,174</b>	<b>8.7</b>
Tracktor++v2 [1]	online	55.1	56.3	1,987	21.1	35.3	<b>3,763</b>	<b>8,866</b>	235,449	1.5
TrctrD17 [36]	online	53.8	53.7	1,947	19.4	36.6	4,792	11,731	247,447	4.9
Tracktor++ [1]	online	52.3	53.5	2,072	19.5	36.6	4,611	12,201	248,047	1.5
<b>BLSTM-MTP-T</b>	online	<b>60.5</b>	55.9	<b>1,188</b>	20.5	36.7	4,185	<b>8,663</b>	238,863	6.0 (24.8)

Table 6. MOT 17 Challenge (with trackers that utilized public detections + Tracktor [1] or public detections + CenterTrack [39]). We used the public detections provided by [3] as input to our tracker. Our tracker runs at 24.8 Hz when the detections processed by Tracktor were used as input. When we consider the processing time spent by Tracktor (8 Hz, calculated by [3]), our tracker runs at 6 Hz. Note that we used **bold** for the best number and **blue** color for the second-best number.

Method	Type	IDF1	MOTA	IDS	MT	ML	Frag	FP	FN	Hz
STRN-MOT17 [35]	online	<b>56.0</b>	50.9	2,397	18.9	<b>33.8</b>	9,363	25,295	<b>249,365</b>	13.8
DMAN [40]	online	<b>55.7</b>	48.2	<b>2,194</b>	<b>19.3</b>	38.3	5,378	26,218	263,608	0.3
MOTDT17 [22]	online	52.7	50.9	2,474	17.5	35.7	5,317	24,069	250,768	<b>18.3</b>
AM-ADM17 [30]	online	52.1	48.1	2,214	13.4	39.7	<b>5,027</b>	25,061	265,495	5.7
HAM-SADF17 [37]	online	51.1	48.3	<b>1,871</b>	17.1	41.7	<b>3,020</b>	<b>20,967</b>	269,038	5.0
PHD-GSDL17 [12]	online	49.6	48.0	3,998	17.1	35.6	8,886	23,199	265,954	6.7
FAMNet [8]	online	48.7	<b>52.0</b>	3,072	19.1	<b>33.4</b>	5,318	<b>14,138</b>	253,616	0.0
<b>BLSTM-MTP-O</b>	online	54.9	<b>51.5</b>	2,566	<b>20.4</b>	35.5	7,748	29,616	<b>241,619</b>	<b>20.1</b>
MHT-blSTM [17]	near-online	<b>51.9</b>	47.5	<b>2,069</b>	18.2	41.7	3,124	25,981	268,042	<b>1.9</b>
EDMT17 [4]	near-online	51.3	50.0	2,264	<b>21.6</b>	<b>36.3</b>	3,260	32,279	<b>247,297</b>	0.6
MHT-DAM [16]	near-online	47.2	<b>50.7</b>	2,314	20.8	36.9	<b>2,865</b>	<b>22,875</b>	252,889	0.9
<b>BLSTM-MTP-N</b>	near-online	<b>55.8</b>	<b>53.6</b>	<b>1,845</b>	<b>23.5</b>	<b>34.4</b>	<b>2,294</b>	<b>23,583</b>	<b>236,185</b>	<b>22.7</b>

Table 7. MOT 17 Challenge (Published online and near-online methods using public detections)

Method	Type	IDF1	MOTA	IDS	MT	ML	Frag	FP	FN	Hz
STRN-MOT16 [35]	online	<b>53.9</b>	<b>48.5</b>	747	<b>17.0</b>	<b>34.9</b>	2,919	9,038	<b>84,178</b>	13.5
DMAN [40]	online	<b>54.8</b>	46.1	<b>532</b>	<b>17.4</b>	42.7	1,616	7,909	89,874	0.3
MOTDT [22]	online	50.9	47.6	792	15.2	38.3	1,858	9,253	85,431	<b>20.6</b>
STAM16 [9]	online	50.0	46.0	<b>473</b>	14.6	43.6	<b>1,422</b>	6,895	91,117	0.2
RARI6pub [11]	online	48.8	45.9	648	13.2	41.9	1,992	6,871	91,173	0.9
KCF16 [7]	online	47.2	<b>48.8</b>	648	15.8	<b>38.1</b>	<b>1,116</b>	<b>5,875</b>	86,567	0.1
AMIR [31]	online	46.3	47.2	774	14.0	41.6	1,675	<b>2,681</b>	92,856	1.0
<b>BLSTM-MTP-O</b>	online	53.5	48.3	735	<b>17.0</b>	38.7	2,349	9,792	<b>83,707</b>	<b>21.0</b>
NOMT [5]	near-online	<b>53.3</b>	<b>46.4</b>	<b>359</b>	<b>18.3</b>	41.4	<b>504</b>	9,753	<b>87,565</b>	<b>2.6</b>
EDMT [4]	near-online	47.9	45.3	639	17.0	<b>39.9</b>	946	11,122	87,890	1.8
MHT-blSTM [17]	near-online	47.8	42.1	753	14.9	44.4	1,156	11,637	93,172	1.8
MHT-DAM [16]	near-online	46.1	45.8	590	16.2	43.2	781	<b>6,412</b>	91,758	0.8
<b>BLSTM-MTP-N</b>	near-online	<b>52.6</b>	<b>49.9</b>	<b>573</b>	<b>19.6</b>	<b>38.9</b>	<b>680</b>	<b>7,098</b>	<b>83,657</b>	<b>23.8</b>

Table 8. MOT 16 Challenge (Published online and near-online methods using public detections)

tracked objects in MOT 17. Note that these are obtained with greedy data association and only local smoothing is added on top of the online tracker performance. The speed is also faster than the online version since the track extension module used for the online tracker is replaced with simple local track smoothing in the near-online version.

## 7. Conclusion

In this paper, we introduce a novel multi-track pooling module that enables joint updating of appearance models using all tracks, thereby improving matching reliability when targets are similar in appearance. We propose a novel training strategy for track pooling that utilizes within-batch dependencies among tracks and supports training over

long sequences. The resulting tracker is based on a Bilinear LSTM architecture and performs greedy data association. With this simple approach, it achieves real-time tracking performance with an accuracy equivalent to state-of-the-art online trackers.

## Acknowledgements

This work was supported in part by NIH award 1R24O D020174-01A1 and DARPA contract N66001-19-2-4035. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Defense Advanced Research Projects Agency (DARPA).



## References

- [1] Philipp Bergmann, Tim Meinhardt, and Laura Leal-Taixé. Tracking without bells and whistles. In *ICCV 2019*, 2019. 7, 8
- [2] Keni Bernardin and Rainer Stiefelhagen. Evaluating multiple object tracking performance: the CLEAR MOT metrics. *Image and Video Processing*, 2008. 6
- [3] Guillem Brasó and Laura Leal-Taixé. Learning a neural solver for multiple object tracking. In *CVPR*, June 2020. 2, 7, 8
- [4] Jiahui Chen, Hao Sheng, Yonghui Zhang, and Zhang Xiong. Enhancing detection model for multiple hypothesis tracking. In *CVPR Workshops*, 2017. 8
- [5] Wongun Choi. Near-online multi-target tracking with aggregated local flow descriptor. In *ICCV*, 2015. 3, 8
- [6] W. Choi and S. Savarese. A unified framework for multi-target tracking and collective activity recognition. In *ECCV*, 2012. 3
- [7] Peng Chu, Heng Fan, Chiu Tan, and Haibin Ling. Online multi-object tracking with instance-aware tracker and dynamic model refreshment. In *WACV*, 2019. 8
- [8] Peng Chu and Haibin Ling. Famnet: Joint learning of feature, affinity and multi-dimensional assignment for online multiple object tracking. In *ICCV 2019*, 2019. 8
- [9] Qi Chu, Wanli Ouyang, Hongsheng Li, Xiaogang Wang, Bin Liu, and Nenghai Yu. Online multi-object tracking using cnn-based single object tracker with spatial-temporal attention mechanism. In *ICCV*, 2017. 8
- [10] P. Dendorfer, H. Rezatofighi, A. Milan, J. Shi, D. Cremers, I. Reid, S. Roth, K. Schindler, and L. Leal-Taixé. CVPR19 tracking and detection challenge: How crowded can it get? June 2019. arXiv: 1906.04567. 6
- [11] Kuan Fang, Yu Xiang, Xiaocheng Li, and Silvio Savarese. Recurrent autoregressive networks for online multi-object tracking. In *WACV*, 2018. 2, 8
- [12] Zeyu Fu, Pengming Feng, Federico Angelini, Jonathon A. Chambers, and Syed Mohsen Naqvi. Particle phd filter based multiple human tracking using online group-structured dictionary learning. *IEEE Access*, 6:14764–14778, 2018. 8
- [13] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 6
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4
- [15] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9:1735–80, 12 1997. 3
- [16] C. Kim, F. Li, A. Ciptadi, and J. Rehg. Multiple hypothesis tracking revisited. In *ICCV*, 2015. 3, 8
- [17] C. Kim, F. Li, and J. Rehg. Multi-object tracking with neural gating using bilinear lstm. In *ECCV*, 2018. 2, 3, 4, 5, 6, 8
- [18] L. Leal-Taixé, A. Milan, I. Reid, S. Roth, and K. Schindler. MOTChallenge 2015: Towards a benchmark for multi-target tracking. arXiv:1504.01942, 2015. 5, 6
- [19] Wei Li, Rui Zhao, Tong Xiao, and Xiaogang Wang. Deepreid: Deep filter pairing neural network for person re-identification. In *CVPR*, 2014. 6
- [20] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection (best student paper award). In *ICCV*, 2017. 6
- [21] Qiankun Liu, Qi Chu, Bin Liu, and Nenghai Yu. Gsm: Graph similarity model for multi-object tracking. In *IJCAI*, 2020. 7, 8
- [22] Chen Long, Ai Haizhou, Zhuang Zijie, and Shang Chong. Real-time multiple people tracking with deeply learned candidate selection and person re-identification. In *ICME*, 2018. 8
- [23] Liqian Ma, Siyu Tang, Michael J. Black, and Luc Van Gool. Customized multi-person tracker. In *ACCV 2018*, 2018. 2
- [24] Andrii Maksai and Pascal Fua. Eliminating exposure bias and metric mismatch in multiple object tracking. In *CVPR*, 2019. 7
- [25] A. Milan, L. Leal-Taixé, I. Reid, S. Roth, and K. Schindler. MOT16: A benchmark for multi-object tracking. arXiv:1603.00831, 2016. 5, 6
- [26] A. Milan, S. H. Rezatofighi, A. Dick, I. Reid, and K. Schindler. Online multi-target tracking using recurrent neural networks. In *AAAI*, 2017. 2
- [27] Hyeonseob Nam and Bohyung Han. Learning multi-domain convolutional neural networks for visual tracking. In *CVPR*, 2016. 2
- [28] Stefano Pellegrini, Andreas Ess, Konrad Schindler, and Luc J. Van Gool. You'll never walk alone: Modeling social behavior for multi-target tracking. In *ICCV*, 2009. 3
- [29] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *ECCV Workshop*, 2016. 6
- [30] S. Bae S. Lee, M. Kim. Learning discriminative appearance models for online multi-object tracking with appearance discriminability measures. *IEEE Access*, 2018. 8
- [31] A. Sadeghian, A. Alahi, and S. Savarese. Tracking the untrackable: Learning to track multiple cues with long-term dependencies. In *ICCV*, 2017. 2, 3, 4, 5, 6, 8
- [32] Ilya Sutskever. Training recurrent neural networks. *PhD thesis*, 2013. 2, 5
- [33] S. Tang, M. Andriluka, B. Andres, and B. Schiele. Multiple people tracking with lifted multicut and person re-identification. In *CVPR*, 2017. 3, 6
- [34] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc., 2017. 3
- [35] Jiarui Xu, Yue Cao, Zheng Zhang, and Han Hu. Spatial-temporal relation networks for multi-object tracking. In *ICCV 2019*, 2019. 2, 8
- [36] Yihong Xu, Aljosa Osep, Yutong Ban, Radu Horaud, Laura Leal-Taixé, and Xavier Alameda-Pineda. How to train your deep multi-object tracker. In *CVPR*, pages 6787–6796, 2020. 7, 8

- [37] Youngchul Yoon, Abhijeet Boragule, Youngmin Song, Kwangjin Yoon, and Moongu Jeon. Online multi-object tracking with historical appearance matching and scene adaptive detection filtering. In *AVSS*, 2018. 8
- [38] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, 2015. 6
- [39] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. *ECCV*, 2020. 7, 8
- [40] Ji Zhu, Hua Yang, Nian Liu, Minyoung Kim, Wenjun Zhang, and Ming-Hsuan Yang. Online multi-object tracking with dual matching attention networks. In *ECCV*, 2018. 2, 8