

Multi-Scene Generalized Trajectory Global Graph Solver with Composite Nodes for Multiple Object Tracking

Yan Gao¹, Haojun Xu¹, Jie Li^{1*}, Nannan Wang¹, Xinbo Gao^{1,2*}

¹Xidian University

²Chongqing University of Posts and Telecommunications

{yangao, haojunxu}@stu.xidian.edu.cn, nnwang@xidian.edu.cn, leejie@mail.xidian.edu.cn, gaodb@cqupt.edu.cn

Abstract

The global multi-object tracking (MOT) system can consider interaction, occlusion, and other “visual blur” scenarios to ensure effective object tracking in long videos. Among them, graph-based tracking-by-detection paradigms achieve surprising performance. However, their fully-connected nature poses storage space requirements that challenge algorithm handling long videos. Currently, commonly used methods are still generated trajectories by building one-forward associations across frames. Such matches produced under the guidance of first-order similarity information may not be optimal from a longer-time perspective. Moreover, they often lack an end-to-end scheme for correcting mismatches. This paper proposes the Composite Node Message Passing Network (CoNo-Link), a multi-scene generalized framework for modeling ultra-long frames information for association. CoNo-Link’s solution is a low-storage overhead method for building constrained connected graphs. In addition to the previous method of treating objects as nodes, the network innovatively treats object trajectories as nodes for information interaction, improving the graph neural network’s feature representation capability. Specifically, we formulate the graph-building problem as a top-k selection task for some reliable objects or trajectories. Our model can learn better predictions on longer-time scales by adding composite nodes. As a result, our method outperforms the state-of-the-art in several commonly used datasets.

Introduction

Multi-Object Tracking (MOT) aims to generate trajectories for all moving objects in a video stream. It is a fundamental module for video content analysis in application areas such as autonomous driving and intelligent robotics. In recent years, tracking-by-detection as the dominant paradigm in the field divides the task into (i) frame-by-frame object detection and (ii) data association, i.e., linking potential targets to the correct object trajectory whenever possible. Data association is performed mainly on neighboring frames, greedily matching trajectories with detection proposals through simple cues such as position and motion prediction (Zhang et al. 2022; Tokmakov et al. 2021; Zhou, Koltun, and Krähenbühl




Figure 1: Node connectivity relationship graphs for different graph structures. The 2nd bolded gray link in T indicates that the connection in each graph is the relationship between the blue node in the first frame and the nodes in the next frame.

2020; Bergmann, Meinhardt, and Leal-Taixé 2019; Bewley et al. 2016) or appearance features (Pang et al. 2021; Zhang et al. 2021). The trajectories formed by such local trackers can satisfy the accuracy requirements. However, severe object occlusion and appearance changes in crowded scenes can pose challenges for long-term identity preservation. Many studies have executed offline associations over the entire video frame to obtain long-range trajectories. Such global trackers typically require the construction of global appearance and motion cues or pairwise association between tracklets on all frames based on graphical representations.

To maximize the applicability to each scenario, recent graph-based MOT studies (Brasó and Leal-Taixé 2020; Hornáková et al. 2020, 2021; Cetintas, Brasó, and Leal-Taixé 2023) have adopted a generic approach to global tracking by using a unified module to process videos in a single framework. While the research has shown promising results, there are still some issues with the generalized graph-based approach. A significant challenge is the need to consider the time cues of video clips more. Generic methods generally use a uniform, fully-connected graph-building system, and due to storage constraints, the node information encoded is limited to neighboring frames, forming trajectories that may not be optimal in the temporal dimension. In contrast, when humans perform dynamic tracking, they must maintain the temporal consistency of the perceived object

*Corresponding Author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

(Hyvärinen et al. 2014). Another issue is that the fixed mode of building graphs suffers from an error accumulation effect. The incorrectly predicted features can ultimately affect the quality of feature learning. The above analysis raises the question, what does the graph structure needed for a generic MOT approach look like?

In this work, we propose CoNo-Link, a composite node message-passing network that combines improved graph formulation rules and effective learning strategies into a graph solver for unified learning. The framework comprises a NodeNet that generates nodes to build a graph and a graph neural network (GNN) link block that accomplishes reasoning about the entire domain. Specifically, NodeNet inputs long clips and outputs a one-piece partially connected graph G^{part} containing potential connections and composite nodes. We then reason by the link block to predict the edge score for G^{part} and output the final result.

G^{part} is the core of our CoNo-Link, inspired by the human visual system. Research (Hyvärinen et al. 2014) shows that infants can expand from tracking salient objects to arbitrary objects when they observe the world. That means observation and tracking is a continuous process in which there are priorities of interest. We summarize three tracking characteristics in the visual system and design two types of nodes and three types of connections accordingly. Among them, the detection node priority confirms the significantly moving objects, the trajectory node records the continuous motion of things, and the connections between composite nodes create pathways for the expression of relationships between arbitrary objects. We use a GNN to process the trajectory nodes and propose a complementary edge learning strategy for these nodes that allows potentially misclassified edges to resume learning.

Fig. 1 includes the node connectivity relations of different graph structures when generating a trajectory T . In the view of blue nodes, both the fully-connected frame-by-frame graph G^{fully} and the hierarchical neighboring frames graph G^{local_fully} have to be connected to all the nodes of the next frame, while our G^{part} only needs to be connected to meaningful object nodes. This scheme solves two limitations of the previous approaches: (i) Partially link significantly reduces memory consumption and can efficiently handle long clips. (ii) Distant nodes are reachable in G^{part} , with better feature learning potential. We evaluate the proposed CoNo-Link on several benchmarks: achieve a competitive 83.7 IDF1 and 82.7 MOTA on MOT17, 81.8 and 77.5 on MOT20, and 64.1 and 89.7 on DanceTrack.

In summary, we make the following main contributions.

- We propose a composite node messaging network that achieves globally optimal convergence by gradually aggregating meaningful trajectory nodes. It guarantees continuous solutions with significantly lower computational costs.
- We construct a partially connected graph based on composite nodes to build a perfect graph domain in which a generalized MOT graph solver can effectively utilize temporal information. It is a natural graph structure abstracted from human visual tracking.

- We achieve the state-of-the-art result in three public datasets and have multi-scene applicability.

Related Work

Local Short-Term Tracking

Many modern trackers run online using frame-by-frame or local frames (Bewley et al. 2016; Bergmann, Meinhardt, and Leal-Taixé 2019; Meinhardt et al. 2022; Pang et al. 2021; Tokmakov et al. 2021; Zeng et al. 2022; Zhang et al. 2022, 2021; Gao et al. 2022). Distance metrics for motion and spatial cues are often central to the design of trajectory formation for these methods. SORT (Bewley et al. 2016) and DeepSORT (Wojke, Bewley, and Paulus 2017; Wojke and Bewley 2018) have led many methods to use the Kalman filter as motion models (Bewley et al. 2016; Zhang et al. 2022, 2021; Dendorfer et al. 2022). CenterTrack (Zhou, Koltun, and Krähenbühl 2020) implements motion prediction and execution association through neural networks, which has led to the emergence of methods to improve motion prediction through improved model structures. Tracktor (Bergmann, Meinhardt, and Leal-Taixé 2019) inspired a frame-by-frame regression-based tracking framework (Liang et al. 2022; Meinhardt et al. 2022; Zhou, Koltun, and Krähenbühl 2020; Bergmann, Meinhardt, and Leal-Taixé 2019). Some trackers use appearance to identify the same object to increase robustness in low-quality video scenes (Pang et al. 2021; Zhang et al. 2021; Xu et al. 2019; Wojke, Bewley, and Paulus 2017), such as low frame rates or strong camera motion (Dendorfer et al. 2022). Although these pair-wise association-based trackers have good tracking stability, they do not focus on the long-term preservation of object identity. Here we do this by performing association on all objects throughout the temporal dimension.

Graph-Based Global Tracking

Graphs are a framework well suited for modeling data association, where each object trajectory can be considered a simple graph with entry and exit degrees of 1 for each node except for the start and end nodes. They use object detection as object nodes and represent edges as possible trajectory hypotheses. In contrast to trackers that use neighbor frame information, graph-based approaches define the cross-frame object association problem as a global combinatorial optimization problem (Koh et al. 2022; Dendorfer et al. 2020a; Brasó and Leal-Taixé 2020; Zeng et al. 2022). To this end, many studies have used different optimization strategies, including multi-cuts (Tang et al. 2017), minimal cliques (Zamir, Dehghan, and Shah 2012), network flow (Berclaz et al. 2011; Butt and Collins 2013), and disjoint path approaches (Hornáková et al. 2020, 2021). Following (Brasó and Leal-Taixé 2020), this paper relies on a simplified minimum cost flow formulation (Zhang, Li, and Nevatia 2008a) to ensure that graph-based network structures are manageable.

Early graph-based methods often used methods based on handcrafted models (Takala and Pietikäinen 2007) or conditional random fields (Yang and Nevatia 2012) to obtain association cues. More recently, many approaches employ GNNs (Brasó, Cetintas, and Leal-Taixé 2022; Brasó and Leal-Taixé

2020; Dai et al. 2021) or Transformers (Meinhardt et al. 2022; Zhou et al. 2022; Zeng et al. 2022) to learn features for the association. Lif_T (Hornáková et al. 2020) incorporates epistemic and pose features in optimizing disjoint path problems. MPNTrack (Brasó and Leal-Taixé 2020) proposes a neural solver to optimize a simplified graph. MOTR (Zeng et al. 2022) follows the DETR (Carion et al. 2020) structure to iteratively update the tracking queries in a propagated manner. The global association method GTR (Zhou et al. 2022) uses transformer queries to generate the entire trajectory simultaneously. These GNN and transformer-based works have achieved good performance. However, the current class of methods iteratively completes object building through all frames and still suffers from the limitations of local information association. In addition, out-of-memory (OOM) is prone to occur if the length of the processed video is considerable.

Recently multi-stage tracking methods (Chen et al. 2020; Wu et al. 2021) become popular. These methods first form short-range trajectories and then deals with lost tracklets and occlusions over long periods. Different optimization techniques (Wu et al. 2021) and cues (Gupta, Dollár, and Girshick 2019) are designed for multi-stage merging when forming long-range trajectories. Many modern trackers incorporate such techniques to improve error correction (Berclaz et al. 2011; Brasó and Leal-Taixé 2020; Meinhardt et al. 2022; Wojke and Bewley 2018). In this paper, we propose a learning strategy to gain the ability to reload missing edges during training.

Methodology

Preliminaries

Tracking-by-Detection (TBD) The TBD paradigm performs the task through frame-by-frame object detection and well-designed inter-target association. Let I be a set of images in a video; the detector first identifies and locates targets in all elements of I , producing a set of objects O with positions $\{p_i \mid p_i \in \mathbb{R}^4\}$. The object set O used for the association is determined in some way, and it is common practice to keep the proposals above a set confidence threshold as candidate objects. Each candidate object o_i contains the position p_i , the image patch p_i , and the time t . Then, the tracker associates each object in O to obtain its trajectory over time $T = \{\tau_1, \dots, \tau_K\}$, and $\tau_i = \{o_{i1}, \dots, o_{in_i}\}$ where n_i is the trajectory length of the object o_i . In this paper, our algorithm also follows this paradigm.

Our tracker is built on modeling inter-frame object relations using an undirected graph $G = (V, E)$ where each node V of G corresponds to an object detection o_i . The edge $E \rightarrow \mathbb{R}^{V \times V}$ denotes the possible interactions between objects on different frames. The association assumption E guarantees the connectivity of object pairs in different frames and facilitates trajectory error correction.

Graph-Based Tracking Baseline Based on G as the underlying representation of inter-object relations, the proposed model is performed on a message-passing graph network tracker (Brasó and Leal-Taixé 2020) based on the classical MOT network flow formulation (Zhang, Li, and Neva-

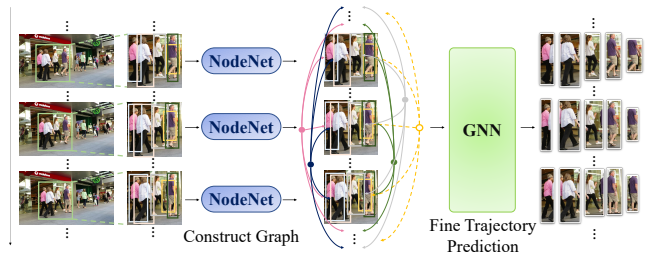


Figure 2: We propose the CoNo-Link framework. It comprises a node network (NodeNet) and a graph neural network (GNN). Given long-time video clips as inputs, the NodeNet first performs a graph-building operation to construct the objects appearing throughout the clips into graph nodes. Then, the constrained connection graph of the clip completes the prediction of delicate trajectories in the GNN.

tia 2008b). It is a framework that performs global reasoning directly over the entire graph domain and predicts whether object connectivity relations hold. Formally, it represents a trajectory τ_i as an object o_i in G linked time-by-time, i.e., $E(\tau_i) = (o_{i1}, o_{i2}), \dots, (o_{in_i-1}, o_{in_i})$. Then, the corresponding paths in G are subjected to binary decision through discrete optimization. Thus, if edge $(u, v) \rightarrow E(\tau_i)$ denotes a correct assumption, it means that the edge cost is $y(u, v) = 1$; otherwise, it is an incorrect assumption, i.e., $y(u, v) = 0$. The algorithm obtains the final trajectory by classifying the edges in the graph.

CoNo-Link

Tracking in a Gradual Way We provide an overview of CoNo-Link in Fig. 2. It comprises a node network (NodeNet) and message-passing graph neural networks (GNN). The NodeNet first performs the preparatory work of constructing the constrained connectivity graph in a video clip. Specifically, it constructs the objects appearing throughout the clip as graph nodes. In this, we define the graph nodes as composite node representations that contain the object’s detection nodes N^{det} and trajectory nodes N^{traj} (with an initial length of 1). Then, we establish the connections between composite nodes $\{N^{det}, N^{traj}\}$ according to the graph-building rules in the next section. Eventually, the complete partial connection graph of the entire input video clip will be fed into the GNN to predict delicate trajectories. It is worth noting that the NodeNet is trained separately. The trajectory nodes built by it are rough representations containing global information, which is intended to ensure the envelopment of edges included in the built graph to the ground truth edges. In addition, CoNo-Link can uniformly access all the information in long video clips, thus ensuring that object identities can be aligned over a wide range of time scales.

Our CoNo-Link is a tracking algorithm that gradually learns the connectivity between composite nodes through GNN, i.e., an incremental prediction process from detection nodes to sub-trajectory nodes and then from sub-trajectory nodes to the complete trajectory. This progressive learning reduces the learning difficulty of the GNN, which can effec-

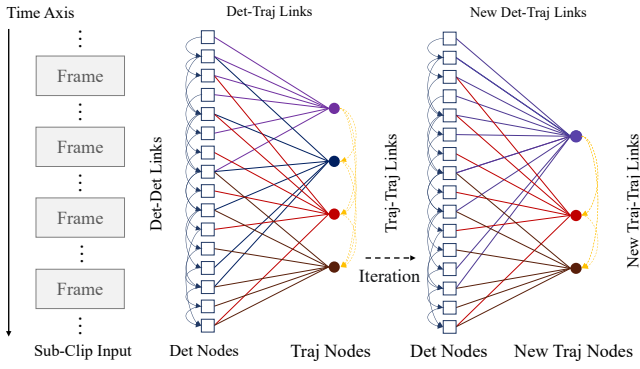


Figure 3: G^{part} structure. This graph contains potential links with a high ground-truth link share, which is the GNN input.

tively deal with ultra-long-time partially connected graphs and thus accomplish long-term tracking.

Constructing a Partially Connected Tracking Graph

This study proposes matching using a composite node-based partial connectivity graph G^{part} learning global information interactions over the time series of the entire video clip. Based on our experiments, we found that generating a fully connected graph and pruning it will always miss the edges. There are two types of reasons; one is that the same identity (ID) target is not in the video sampling window, and the other is that the connection is missing due to the feature learning problem of the model itself. Since missing links are unavoidable, we abandon the commonly used pruning method and instead construct the graph by identifying the “most promising” node relationships and backfilling the edges. Among them, the nodes can be the correct object detection or the shortest paths with the same ID that is easy to construct. The main idea is to follow a coarse-to-fine strategy: build an imperfect but relatively correct part of the node relationship graph and then determine new connections from the GNN to realize accurate tracking. Therefore, we propose NodeNet to establish the initial relationships between detections.

Establish the initial relationships. We employ a query representation to capture objects’ temporal and interaction relations. Specifically, we give each query node an initially constructed coarse trajectory in the semantic feature space by track search but keep only the candidate links for each node along the temporal dimension. Let $O = \{o_1^t, \dots, o_{N_t}^t\}$ be a set of detector output objects for image I^t with total number $N = \sum_t N_t$. Let $F^t = \{f_i^t \rightarrow R^D | f_1^t, \dots, f_{N_t}^t\}$ be the set of D -dimensional features extracted from their corresponding bounding boxes and $F = F^1 \cup \dots \cup F^T$ is the set of all features in time slice T . As shown in Fig. 4, all objects with $F \in R^{N \times D}$ and query $Q_k \in R^D$ are the inputs to our NodeNet. The association scores $s \in R^N$ that produce the coarse trajectories are the outputs. Formally, we use a softmax activation to model the likelihood of association between an object i and each trajectory k at time t as $P_M(o_i | Q_k, F) = \exp(s_i^t) / \sum_{j \in \{0, 1, \dots, N_t\}} \exp(s_j^t)$. Then, the distribution of all objects at time t corresponding to tra-

jectory k is $P^t(p | Q_k, F) = \sum_{i=1}^{N_t} l_{[p=p_i^t]} P_M(o_i | Q_k, F)$, where $l_{[\cdot]}$ assigns a bounding box p to each query. Ultimately, the distribution of trajectory k over the entire period is $P^T(\tau | Q_k, F) = \prod_{t=1}^T P^t(\tau^t | Q_k, F)$.

During training, we maximize the log-likelihood of the ground-truth trajectory, and features that are not matched are used as background queries and supervised empty set frames. Let \hat{o}_k be matched objects for a ground-truth trajectory τ , and $F_{\hat{o}_k}^m$ be the features of \hat{o}_k . For each trajectory τ , we optimize the training objective of the query assignment as:

$$L_{NodeNet} = L_{bg}(F) + \sum_{\tau} L_{matched}(F, \tau).$$

$$L_{bg}(F) = - \sum_{m=1}^T \sum_{j: \nexists \hat{o}_k^m = j} \sum_{t=1}^T \log P_M(o^t = \emptyset | F_j^m, F).$$

$$L_{matched}(F, \tau) = - \sum_{m \in \{1, \dots, T\} | \hat{o}_k^m \neq \phi} \sum_{t=1}^T \log P_M(\hat{o}_k^m | F_{\hat{o}_k^m}^m, F). \quad (1)$$

We first obtain the similarity matrix \mathcal{M} in the inference process. The method is to let NodeNet process the video stream as a sliding window. The length of the video stream is $T_{clip} = 512$, the window size is $T = 32$, and the step size is 16. We use a backbone network to extract the detected frames within the window to obtain F sequentially. Then, let $Q = K = V = F$ feed them into NodeNet, and output an association matrix of $N \times N$. Assign this matrix to the corresponding position in the large matrix \mathcal{M} , and if there is already a value in the corresponding place, take the average value in the overlapping part.

After obtaining the similarity matrix \mathcal{M} , we will perform a frame-by-frame association. For the first frame, we initialize all detections to trajectories. For the N_t bounding boxes in the subsequent frame I^t , the similarity sub-matrix M of the bounding boxes within the window of size $\min(t, T)$ is taken out in \mathcal{M} with this frame as the end of the window. The similarity between the query detections (trajectories) and the object detections within the window that have the same ID is summed to obtain the similarity matrix of each ID to the object of the current frame $\bar{M} \in R^{ID \times N_t}$. In addition, we compute the IoU matrix $\hat{M} \in R^{ID \times N_t}$ between the last occurrence of each ID box and the current frame box. Our final cost matrix is $C = -\max(\bar{M}, \hat{M})$. Further, we use the Hungarian algorithm (Kuhn 2010) to get the optimal association for each ID. We start a new trajectory if the average association score with any previous ID is below a threshold ϑ . Otherwise, we attach the generated trajectory’s underlying current detection (query) to the matching existing trajectory. In the similarity matrix \bar{M} , we connect the optimal associations for each ID and the Top- k bounding boxes with the highest similarity. These detections connections (*Det-Det links*) are the potential links we need.

Predict several short paths. We feed G^{part} containing T_{clip} information into the GNN for edge prediction. Then, we aggregate some natural short trajectories generated by the GNN into trajectory nodes N^{traj} . Precisely, we fill in

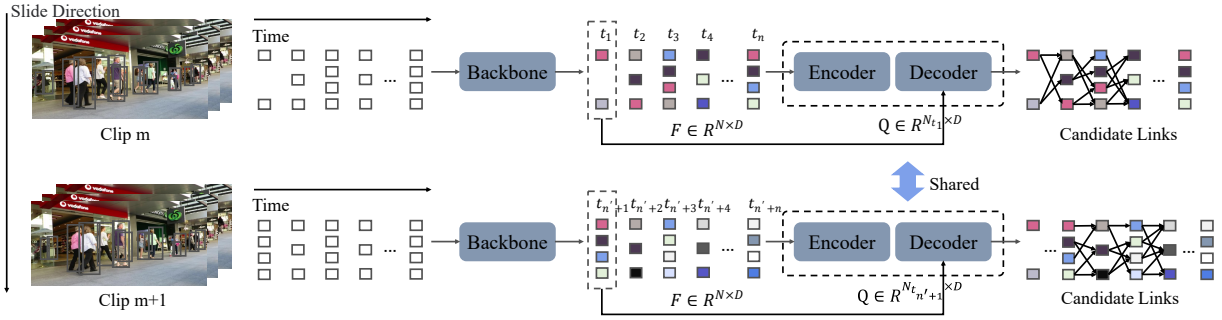


Figure 4: The detailed network architecture of NodeNet. It contains a layer of an encoder and a layer of a decoder following (Carion et al. 2020; Zhou et al. 2022). The object feature F and the query Q are the inputs to the encoder and decoder, respectively, and generate an affinity matrix between the query and the object \mathcal{M} as the basis for developing our candidate links.

the predicted scores of the GNN into \mathcal{M}_{GNN} as the similarity between the nodes and determine the trajectory IDs using the same association mode for generating Det-Det links. For each path (containing Det-Traj links), the detection features within the trajectory are averaged as the new trajectory feature representation. Finally, we build connections between N^{traj} (Traj-Traj links) and do not establish the relationship if temporal IoU exists between N^{traj} .

G^{part} is now the entire form that the complete set of candidate links established between pre-generated composite nodes. It contains Det-Det links, Det-Traj links, and Traj-Traj links, and we illustrate its construction in Fig. 3. These three links are significant because the Det-Det link generates new N^{traj} . Det-Traj link is responsible for building the information interaction between N^{det} and N^{traj} . The Traj-Traj link is responsible for fusing the sub-trajectories. In each iteration, the composite nodes undergo a directed transformation, and we transform the Traj-Traj links that satisfy the requirements into new N^{traj} , which are then merged with the existing N^{traj} to generate the final result.

The summary of G^{part} vs. G^{fully} . The pre-generated composite nodes retain meaningful object nodes and reduce invalid edges, alleviating the labeling imbalance. This strategy reduces storage occupancy, so the graph contains all the information in clip time, guaranteeing spatial information interaction over long distances. At the same time, establishing multiple rather than a single kind of edges makes some nodes unavailable to reachable, making mega-graph long-distance matching plausible. Through upper-bound analysis experiments, we determined that the restricted connected graph based on composite node learning performs better than the hierarchical neighboring frame fully connected graph learning. Also, proof that our method satisfies globally optimal matching is provided in the supplementary.

Learn an Effective Link Tracker We use a message-passing GNN (Brasó and Leal-Taixé 2020) to process a partially connected graph with composite nodes. Our main contribution is a learning strategy for GNN training. Specifically, the Traj-Traj links of our G^{part} are fully-connected edges constructed for N^{traj} . Trajectory nodes are relatively few at this time, from 10k down to around the 0.1k level (see

Algorithm 1: Graph aggregation algorithm

Input: current graph node V , current graph edge E^I

Parameter: GNN f_{θ_1} , f_{θ_2} , and $\theta_1 = \theta_2$, assign threshold ϵ

Output: id of every node K

- 1: Compute edge classification score $S_{E^I} = f_{\theta_1}(V, E^I)$
 - 2: Assign id to every node in V by S_{E^I} use the tracker described in Sec 3.2.1. The id of V is K_0 .
 - 3: Aggregate the same id node in V to get trajectory node V_{node}
 - 4: Get fully-connected edge E_{fc} of V_{node}
 - 5: Compute edge classification score $S_{E_{fc}} = f_{\theta_2}(V_{node}, E_{fc})$
 - 6: $E_+ = E_{fc}[S_{E_{fc}} > \epsilon]$
 - 7: Find the connected and nontime-nonoverlap node group
 - 8: Give the same id to every node in node group and map the id in V_{node} to V . The id of V now is \vec{K} .
 - 9: **return** K
-

Table 2), making it possible to build fully-connected graphs between them, and can minimize the problem of missing links when creating graphs. We obtain the final trajectory by performing edge classification to decide whether to keep the established connections. We show the learning process in algorithm 1. The key to making the GNN learn efficiently is maintaining the proportion of positive sample edges, which is why our method works. In the following, we detail a GNN link block that makes long-time node connections.

Let $G^I = (V^I, E^I)$ be a graph of our iteration I . $h_n^{(0)}$ is the node embedding for each $n \in V^I$ and $h_{(u,v)}^{(0)}$ is the edge embedding for each $(u, v) \in E^I$. Following the time-aware framework of (Brasó and Leal-Taixé 2020), the GNN aims to learn a function that encodes the higher-order semantic context contained in the node and edge feature vectors via information propagation. It feeds the edge embedding to a multi-layer perceptron (MLP) and outputs a score $y_{pred}^{(u,v)} = MLP(h_{(u,v)}^{(s)})$, where (s) denotes the number of message passes. We set this score to a similarity with a minimum threshold ϵ limitation for association. As mentioned,

we compute similarity scores based on association cues between nodes recorded in \mathcal{M}_{GNN} . The cues contain the similarity of the appearance embedding and the estimated score based on the closest temporal distance. Additional details are provided in the supplementary material.

Training CoNo-Link contains two GNNs (Brasó and Leal-Taixé 2020) which are trained jointly at all levels. For this purpose, we first train the first network and add the second network to train at a certain level. This ensures stability during the training process. Specifically, we unfreeze the second network after 5000 training iterations. We trained them using a focal loss (Lin et al. 2020) with a $\gamma = 1$ and summed all the losses as our final loss.

Experiments

Datasets and Metrics

We conducted experiments on three public benchmarks: MOT17 (Leal-Taixé et al. 2015; Milan et al. 2016), MOT20 (Dendorfer et al. 2020b), and DanceTrack (Sun et al. 2022). The MOT series is a dense pedestrian tracking dataset evaluated under public and private detection protocols. DanceTrack is a dataset of dance videos with similar appearance and complex motion patterns. We evaluate the performance using several widely used metrics: IDF1 (Ristani et al. 2016), MOTA (Kasturi et al. 2009), and HOTA (Luiten et al. 2021). IDF1 focuses on identity maintenance quality, while the official metric MOTA focuses on detection quality and tracking stability (IDS). HOTA combines two aspects to unify detection localization and association performance.

Implementation Details

Training parameters: we used the pre-trained ResNet50-IBN (Dai et al. 2021) as our ReID network and froze it during training. The GNNs were co-trained with a learning rate of 3×10^{-4} , weight decay of 10^{-4} , and a batch of 2 clips with 200 epochs. The optimizer was Adam (Kingma and Ba 2015). Inference: our model can handle sequences of arbitrary length, and for training and memory efficiency, we have T_{clip} of 512 frames. During inference, we fill the trajectory gaps by linear interpolation. The runtime on MOT17 is 19 FPS with given detections. Object detection: for DanceTrack and the private settings of MOT17 and MOT20, we follow the detections obtained in the YOLOX (Ge et al. 2021) trained in (Zhang et al. 2022). Our tracker was implemented on a single RTX 3090 GPU using Python 3.10 and Pytorch 2.0.0 (Paszke et al. 2019).

Ablation Study

Experimental Settings In this section, we use the MOT17 dataset for all experiments. To evaluate our model, we employ four video sequences (04, 05, 09, and 11) for training and the most challenging three sequences (02, 10, and 13) for validation. All ablation experiments were performed on the validation set.

Candidate Link Generation Structure We consider NodeNet’s three patterns, IoU, ReID, and Trans, for generating potential connections. IoU and ReID denote Hun-

Top- k	MOTA \uparrow			IDF1 \uparrow			Cover. \uparrow		
	IoU	ReID	Trans.	IoU	ReID	Trans.	IoU	ReID	Trans.
1	59.3	18.3	60.8	48.8	2.9	55.4	95.8	49.9	98.7
2	59.9	28.7	60.9	50.1	4.4	55.9	97.3	75.2	99.0
5	60.0	43.7	61.0	50.3	9.8	56.0	97.6	95.0	99.0
10	60.0	52.5	61.0	50.3	18.8	56.0	97.8	98.8	99.1

Table 1: Candidate link generation method selection.

garian matching (Kuhn 2010) based on IoU and ReID cues, and Trans. denotes query-based Transformer matching (i.e., for similarity associations considering global information). Based on these three models, we start the search for the relation-building range Top- k using MOTA, IDF1, and the ground truth (GT) links coverage (Cover.) as evaluation metrics. To obtain as high edge coverage as possible without being affected by environmental factors, we use GT boxes as inputs to NodeNet based on the purity assumption. Table 1 shows the results. We use query matching to establish relationships between the top 5 matches. Although $k = 10$ performs best, the gain is small, and the number of edges becomes large. We also show in Fig. 5 the effect of the window size for establishing potential connections on these metrics. There is the highest GT envelope at $T = 32$ and does not have memory overflow.

Temporally Connected Graph Domains We investigate the upper bounds on the performance of our proposed partially connected graph and the recent hierarchical fully connected graph. The experiment still uses GT detection as input for graph construction. We observe higher upper bounds for almost all metrics of G^{part} in Table 2, including a reduction of IDS to 18. We calculate the edge and node numbers in Table 4 for both graphs under the same upper bound. G^{local_fully} has an average of 256,095 edges and 7,482 nodes, whereas our G^{part} has only 143,603 edges and 151 nodes. The results show that the proposed G^{part} achieves 99.9% GT edge coverage before rounding than G^{local_fully} when using fewer edges to build the graph. As expected, our G^{part} is compelling in graph domain construction because it preserves temporal connectivity in addition to considering neighboring information. It allows the GNN to “see” nodes over time and make more credible decisions, thus generating significantly less IDS.

CoNo-Link’s Component Ablation Experiments We show the performance of each part on the MOT17 validation set in Table 3. The first row shows the results after feed-




Figure 5: The window size experiment of NodeNet.

# Graph	MOTA \uparrow	IDF1 \uparrow	IDS \downarrow	HOTA \uparrow	MT \uparrow	Cover. \uparrow
1 $G^{local-fully}$	61.7	76.2	38	61.4	315	99.9
2 G^{part}	61.8	76.3	18	61.5	316	99.9

Table 2: Experiments on the upper bound of graph structures.

# Method	MOTA \uparrow	IDF1 \uparrow	IDS \downarrow
1 G^{fully} + GNN	59.6	57.8	844
2 G^{part} + GNN	60.1 (0.5 \uparrow)	58.1 (0.3 \uparrow)	636 (208 \downarrow)
3 G^{part} + GNN + LS_1	59.9 (0.2 \downarrow)	64.4 (6.3 \uparrow)	553 (83 \downarrow)
4 G^{part} + GNN + LS_2	60.1 (0.0 -)	61.3 (3.2 \uparrow)	601 (35 \downarrow)
5 G^{part} + GNN + LS_3	60.2 (0.1 \uparrow)	65.2 (7.1 \uparrow)	534 (102 \downarrow)
6 G^{part} + GNN + LS_4	60.3 (0.2 \uparrow)	69.9 (11.8 \uparrow)	449 (187 \downarrow)

Table 3: Ablation study of CoNo-Link.

# Method	GFLOPs	FPS	VRAM (MiB)	E_{avg}	N_{avg}
1 $G^{local-fully}$	46.7	17.1	3,689	256,095	7,482
2 G^{part}	36.3	18.9	3,523	143,603	151

Table 4: Computational complexity experiments.

ing G^{part} directly into the GNN. Then, we add the proposed learning strategy (LS) to get the results of the Link blocks in rows 2-5. There are four settings for LS: (i) mapping the valid edges ($\epsilon > 0.5$) in Traj-Traj links to the corresponding Det-Det links and tracking the Det-Det links to get the final results; (ii) mapping the Traj-Traj links and their classification scores to the corresponding Det-Det links, and track the Det-Det links to get the final results. (iii) Map all the Traj-Traj links to the corresponding Det-Det links and use GNN to re-extract the scores of the corresponding edges and track to output the final results. (iv) Use Traj-Traj links to directly fuse N^{traj} to obtain the final results. Except for LS_1, which slightly decreases MOTA, the other strategies improve in all metrics. In particular, our LS_4 has 0.2 % of MOTA, 11.8 % of IDF1, and 187 IDS gains. Note that all results in Table 3 are without interpolation.

Benchmark Results

MOT We report the results of our model in Table 5 for MOT17 and MOT20 under the private detection protocol. Our approach achieves the desired results on both challenges. On MOT17, we outperform SUSHI based on hierarchical fully-connected graphs on all metrics. In the highly crowded scene of MOT20, it slightly underperforms SUSHI in identity switches. This result may be due to the large number of relational interactions in a short period that impact our approach. Still, again, the global information ensures that there are enough accurate tracklets in the results. Compared to ByteTrack, our model improves this by 6.4 IDF1, 4.0 HOTA, and 2.4 MOTA, reducing IDS by 50.3%. Our performance demonstrates the positive significance of the time-domain connectivity partially connected graph.

DanceTrack Table 6 demonstrates that in scenarios with multiple complex motion patterns, we have good improvements on all metrics and that our approach has a better bal-

Method	IDF1 \uparrow	HOTA \uparrow	MOTA \uparrow	IDS \downarrow
MOT17				
GTR (2022)	71.5	59.1	75.3	2859
FairMOT (2021)	72.3	59.3	73.7	3303
GRTU (2021b)	75.0	62.0	74.9	1812
CorrTracker (2021a)	73.6	60.7	76.5	3369
Unicorn (2022)	75.5	61.7	77.2	5379
ByteTrack (2022)	77.3	63.1	80.3	2196
SUSHI (2023)	83.1	66.5	81.1	1149
Ours	83.7	67.1	82.7	1092
MOT20				
MPNTrack* (2020)	59.1	46.8	57.6	1210
TrackFormer (2022)	65.7	54.7	68.6	1532
MeMOT (2022)	66.1	54.1	63.7	1938
FairMOT (2021)	67.3	54.6	61.8	5243
CorrTracker (2021a)	69.1	-	65.2	5183
ByteTrack (2022)	75.2	61.3	77.8	1223
SUSHI (2023)	79.8	64.3	74.3	706
Ours	81.8	65.9	77.5	956

Table 5: Test set results on MOT17 and MOT20 benchmark. MPNTrack* is the result under the public protocol.

Method	IDF1 \uparrow	HOTA \uparrow	MOTA \uparrow	AssA \uparrow	DetA \uparrow
CenterTrack (2020)	35.7	41.8	86.8	22.6	78.1
FairMOT (2021)	40.8	39.7	82.2	23.8	66.7
TraDes (2021)	41.2	43.3	86.2	25.4	74.5
GTR (2022)	50.3	48.0	84.7	31.9	72.5
QDTrack (2021)	50.4	54.2	87.7	36.8	80.1
MOTR (2022)	51.5	54.2	79.7	40.2	73.5
ByteTrack (2022)	53.9	47.7	89.6	32.1	71.0
SUSHI (2023)	63.4	63.3	88.7	50.1	80.1
Ours	64.1	63.8	89.7	50.7	80.2

Table 6: Test set results on DanceTrack benchmark.

ance between IDF1 and MOTA performance without excessive loss of accuracy. Results of this dataset show that an overall evaluation of trajectories in the temporal dimension can provide some correct clues for association.

Conclusion

This paper presents CoNo-Link, a feasible method for oversized graph tracking. Through ablation experiments, we demonstrate that partially connected graphs can reduce the video memory footprint and efficiently handle long video clips; the trajectory learning strategy can improve the upper bound of G^{part} performance. In addition, the proposed method outperforms the SOTA approaches on three benchmarks. In the future, we would like to design an approach with CLIP (Radford et al. 2021) to enhance the performance in tracking through multimodal tasks, e.g., image caption.

Acknowledgments

This work is supported in part by the National Natural Science Foundation of China under Grants 62176195, 62036007, U22A2096, U21A20514, and 62221005; in part by the Technology Innovation Leading Program of Shaanxi under Grant 2022QFY01-15; in part by Open Research Projects of Zhejiang Lab under Grant 2021KG0AB01.

References

- Berclaz, J.; Fleuret, F.; Türetken, E.; and Fua, P. 2011. Multiple Object Tracking Using K-Shortest Paths Optimization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(9): 1806–1819.
- Bergmann, P.; Meinhardt, T.; and Leal-Taixé, L. 2019. Tracking Without Bells and Whistles. In *IEEE/CVF International Conference on Computer Vision*, 941–951.
- Bewley, A.; Ge, Z.; Ott, L.; Ramos, F. T.; and Upcroft, B. 2016. Simple online and realtime tracking. In *IEEE International Conference on Image Processing*, 3464–3468.
- Brasó, G.; Cetintas, O.; and Leal-Taixé, L. 2022. Multi-Object Tracking and Segmentation Via Neural Message Passing. *Int. J. Comput. Vis.*, 130(12): 3035–3053.
- Brasó, G.; and Leal-Taixé, L. 2020. Learning a Neural Solver for Multiple Object Tracking. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6246–6256.
- Butt, A. A.; and Collins, R. T. 2013. Multi-target Tracking by Lagrangian Relaxation to Min-cost Network Flow. In *IEEE Conference on Computer Vision and Pattern Recognition*, 1846–1853.
- Cai, J.; Xu, M.; Li, W.; Xiong, Y.; Xia, W.; Tu, Z.; and Soatto, S. 2022. MeMOT: Multi-Object Tracking with Memory. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8080–8090.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-End Object Detection with Transformers. In *Computer Vision - ECCV*, volume 12346, 213–229.
- Cetintas, O.; Brasó, G.; and Leal-Taixé, L. 2023. Unifying Short and Long-Term Tracking With Graph Hierarchies. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22877–22887.
- Chen, Y.; Cao, Y.; Hu, H.; and Wang, L. 2020. Memory Enhanced Global-Local Aggregation for Video Object Detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10334–10343.
- Dai, P.; Weng, R.; Choi, W.; Zhang, C.; He, Z.; and Ding, W. 2021. Learning a Proposal Classifier for Multiple Object Tracking. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2443–2452.
- Dendorfer, P.; Rezatofighi, H.; Milan, A.; Shi, J.; Cremers, D.; Reid, I. D.; Roth, S.; Schindler, K.; and Leal-Taixé, L. 2020a. MOT20: A benchmark for multi object tracking in crowded scenes. *CoRR*, abs/2003.09003.
- Dendorfer, P.; Rezatofighi, H.; Milan, A.; Shi, J.; Cremers, D.; Reid, I. D.; Roth, S.; Schindler, K.; and Leal-Taixé, L. 2020b. MOT20: A benchmark for multi object tracking in crowded scenes. *CoRR*, abs/2003.09003.
- Dendorfer, P.; Yugay, V.; Osep, A.; and Leal-Taixé, L. 2022. Quo Vadis: Is Trajectory Forecasting the Key Towards Long-Term Multi-Object Tracking? In *Advances in Neural Information Processing Systems*.
- Gao, Y.; Xu, H.; Zheng, Y.; Li, J.; and Gao, X. 2022. An Object Point Set Inductive Tracker for Multi-Object Tracking and Segmentation. *IEEE Trans. Image Process.*, 31: 6083–6096.
- Ge, Z.; Liu, S.; Wang, F.; Li, Z.; and Sun, J. 2021. YOLOX: Exceeding YOLO Series in 2021. *CoRR*, abs/2107.08430.
- Gupta, A.; Dollár, P.; and Girshick, R. B. 2019. LVIS: A Dataset for Large Vocabulary Instance Segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5356–5364.
- Hornáková, A.; Henschel, R.; Rosenhahn, B.; and Swoboda, P. 2020. Lifted Disjoint Paths with Application in Multiple Object Tracking. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, 4364–4375. PMLR.
- Hornáková, A.; Kaiser, T.; Swoboda, P.; Rolínek, M.; Rosenhahn, B.; and Henschel, R. 2021. Making Higher Order MOT Scalable: An Efficient Approximate Solver for Lifted Disjoint Paths. In *IEEE/CVF International Conference on Computer Vision*, 6310–6320.
- Hyvärinen, L.; Walther, R.; Jacob, N.; Chaplin, K. N.; and Leonhardt, M. 2014. Current Understanding of What Infants See. *Current Ophthalmology Reports*, 2: 142 – 149.
- Kasturi, R.; Goldgof, D. B.; Soundararajan, P.; Manohar, V.; Garofolo, J. S.; Bowers, R.; Boonstra, M.; Korzhova, V. N.; and Zhang, J. 2009. Framework for Performance Evaluation of Face, Text, and Vehicle Detection and Tracking in Video: Data, Metrics, and Protocol. *IEEE Trans. Pattern Anal. Mach. Intell.*, 31(2): 319–336.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations*.
- Koh, J.; Kim, J.; Yoo, J. H.; Kim, Y.; Kum, D.; and Choi, J. W. 2022. Joint 3D Object Detection and Tracking Using Spatio-Temporal Representation of Camera Image and LiDAR Point Clouds. In *Thirty-Sixth AAAI Conference on Artificial Intelligence*, 1210–1218. AAAI Press.
- Kuhn, H. W. 2010. The Hungarian Method for the Assignment Problem. In *50 Years of Integer Programming 1958-2008 - From the Early Years to the State-of-the-Art*, 29–47. Springer.
- Leal-Taixé, L.; Milan, A.; Reid, I. D.; Roth, S.; and Schindler, K. 2015. MOTChallenge 2015: Towards a Benchmark for Multi-Target Tracking. *CoRR*, abs/1504.01942.
- Liang, C.; Zhang, Z.; Zhou, X.; Li, B.; and Hu, W. 2022. One More Check: Making "Fake Background" Be Tracked Again. In *Thirty-Sixth AAAI Conference on Artificial Intelligence*, 1546–1554. AAAI Press.
- Lin, T.; Goyal, P.; Girshick, R. B.; He, K.; and Dollár, P. 2020. Focal Loss for Dense Object Detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(2): 318–327.
- Luiten, J.; Osep, A.; Dendorfer, P.; Torr, P. H. S.; Geiger, A.; Leal-Taixé, L.; and Leibe, B. 2021. HOTA: A Higher Order Metric for Evaluating Multi-object Tracking. *Int. J. Comput. Vis.*, 129(2): 548–578.
- Meinhardt, T.; Kirillov, A.; Leal-Taixé, L.; and Feichtenhofer, C. 2022. TrackFormer: Multi-Object Tracking with Transformers. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8834–8844.

- Milan, A.; Leal-Taixé, L.; Reid, I. D.; Roth, S.; and Schindler, K. 2016. MOT16: A Benchmark for Multi-Object Tracking. *CoRR*, abs/1603.00831.
- Pang, J.; Qiu, L.; Li, X.; Chen, H.; Li, Q.; Darrell, T.; and Yu, F. 2021. Quasi-Dense Similarity Learning for Multiple Object Tracking. In *IEEE Conference on Computer Vision and Pattern Recognition*, 164–173.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Köpf, A.; Yang, E. Z.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*, 8024–8035.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 8748–8763. PMLR.
- Ristani, E.; Solera, F.; Zou, R. S.; Cucchiara, R.; and Tomasi, C. 2016. Performance Measures and a Data Set for Multi-target, Multi-camera Tracking. In *Computer Vision - ECCV*, volume 9914 of *Lecture Notes in Computer Science*, 17–35.
- Sun, P.; Cao, J.; Jiang, Y.; Yuan, Z.; Bai, S.; Kitani, K.; and Luo, P. 2022. DanceTrack: Multi-Object Tracking in Uniform Appearance and Diverse Motion. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20961–20970. IEEE.
- Takala, V.; and Pietikäinen, M. 2007. Multi-Object Tracking Using Color, Texture and Motion. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*.
- Tang, S.; Andriluka, M.; Andres, B.; and Schiele, B. 2017. Multiple People Tracking by Lifted Multicut and Person Re-identification. In *IEEE Conference on Computer Vision and Pattern Recognition*, 3701–3710.
- Tokmakov, P.; Li, J.; Burgard, W.; and Gaidon, A. 2021. Learning to Track with Object Permanence. In *IEEE/CVF International Conference on Computer Vision*, 10840–10849.
- Wang, Q.; Zheng, Y.; Pan, P.; and Xu, Y. 2021a. Multiple Object Tracking With Correlation Learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3876–3886.
- Wang, S.; Sheng, H.; Zhang, Y.; Wu, Y.; and Xiong, Z. 2021b. A General Recurrent Tracking Framework without Real Data. In *IEEE/CVF International Conference on Computer Vision*, 13199–13208.
- Wojke, N.; and Bewley, A. 2018. Deep Cosine Metric Learning for Person Re-identification. In *IEEE Winter Conference on Applications of Computer Vision*, 748–756.
- Wojke, N.; Bewley, A.; and Paulus, D. 2017. Simple online and realtime tracking with a deep association metric. In *IEEE International Conference on Image Processing*, 3645–3649.
- Wu, J.; Cao, J.; Song, L.; Wang, Y.; Yang, M.; and Yuan, J. 2021. Track to Detect and Segment: An Online Multi-Object Tracker. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12347–12356.
- Xu, J.; Cao, Y.; Zhang, Z.; and Hu, H. 2019. Spatial-Temporal Relation Networks for Multi-Object Tracking. In *IEEE/CVF International Conference on Computer Vision*, 3987–3997.
- Yan, B.; Jiang, Y.; Sun, P.; Wang, D.; Yuan, Z.; Luo, P.; and Lu, H. 2022. Towards Grand Unification of Object Tracking. In *Computer Vision - ECCV*, volume 13681, 733–751.
- Yang, B.; and Nevatia, R. 2012. An online learned CRF model for multi-target tracking. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2034–2041.
- Zamir, A. R.; Dehghan, A.; and Shah, M. 2012. GMCP-Tracker: Global Multi-object Tracking Using Generalized Minimum Clique Graphs. In *Computer Vision - ECCV*, volume 7573, 343–356.
- Zeng, F.; Dong, B.; Zhang, Y.; Wang, T.; Zhang, X.; and Wei, Y. 2022. MOTR: End-to-End Multiple-Object Tracking with Transformer. In *Computer Vision - ECCV*, 659–675. Cham: Springer Nature Switzerland.
- Zhang, L.; Li, Y.; and Nevatia, R. 2008a. Global data association for multi-object tracking using network flows. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*.
- Zhang, L.; Li, Y.; and Nevatia, R. 2008b. Global data association for multi-object tracking using network flows. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE Computer Society.
- Zhang, Y.; Sun, P.; Jiang, Y.; Yu, D.; Weng, F.; Yuan, Z.; Luo, P.; Liu, W.; and Wang, X. 2022. ByteTrack: Multi-object Tracking by Associating Every Detection Box. In *Computer Vision - ECCV*, volume 13682, 1–21.
- Zhang, Y.; Wang, C.; Wang, X.; Zeng, W.; and Liu, W. 2021. FairMOT: On the Fairness of Detection and Re-identification in Multiple Object Tracking. *Int. J. Comput. Vis.*, 129(11): 3069–3087.
- Zhou, X.; Koltun, V.; and Krähenbühl, P. 2020. Tracking Objects as Points. In Vedaldi, A.; Bischof, H.; Brox, T.; and Frahm, J., eds., *Computer Vision - ECCV*, volume 12349, 474–490.
- Zhou, X.; Yin, T.; Koltun, V.; and Krähenbühl, P. 2022. Global Tracking Transformers. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8761–8770.