

# SSP-GNN: Learning to Track via Bilevel Optimization

Griffin Golias  
Applied Physics Laboratory  
University of Washington  
Seattle, WA 98105 USA  
goliagri@uw.edu

Masa Nakura-Fan  
Paul G. Allen School of Computer Science  
University of Washington  
Seattle, WA 98105 USA  
mnakura@uw.edu

Vitaly Ablavsky\*  
Applied Physics Laboratory  
University of Washington  
Seattle, WA 98105 USA  
vxa@uw.edu

**Abstract**—We propose a graph-based tracking formulation for multi-object tracking (MOT) where target detections contain kinematic information and re-identification features (attributes). Our method applies a successive shortest paths (SSP) algorithm to a tracking graph defined over a batch of frames. The edge costs in this tracking graph are computed via a message-passing network, a graph neural network (GNN) variant. The parameters of the GNN, and hence, the tracker, are learned end-to-end on a training set of example ground-truth tracks and detections. Specifically, learning takes the form of bilevel optimization guided by our novel loss function. We evaluate our algorithm on simulated scenarios to understand its sensitivity to scenario aspects and model hyperparameters. Across varied scenario complexities, our method compares favorably to a strong baseline.

**Index Terms**—feature-aided tracking, data association, GNN

## I. INTRODUCTION

Multi-target and multi-object tracking (MOT) is often a necessary step in larger systems that address real-world challenges. Examples include tracking pedestrians in the context of autonomous driving, tracking animals and birds to understand environmental factors, and tracking players in team sports to analyze plays.

MOT remains a challenging problem even in scenarios where detections contain features (attributes) that could potentially help associate across target identities and distinguish targets from clutter. One of the challenges in feature-aided tracking is to exploit feature vectors of arbitrary (possibly high) dimensions. Another challenge is to reason jointly and optimally about attributes and target dynamics.

Learning-based approaches to tracking [26], [31] learn key components of a tracker (e.g., data association) directly from annotated training data. Recently, approaches to neural message passing [11] have been adopted by [5] to formulate tracking on a *detection graph* (where a detection graph is instantiated for a temporal window given all the measurements/detections). However, predicting which pairs of detections should be linked without inferring complete tracks, while computationally efficient, may be suboptimal.

We propose an end-to-end bilevel formulation for learning MOT, where the inner optimization (tracking) is accomplished via a successive shortest paths (SSP) algorithm on a *tracking*

*graph* [6], [32] with edge costs. Given these costs, SSP finds a solution that is guaranteed to be optimal. In our outer optimization problem we learn parameters of a function that computes the cost of each edge in the tracking graph given detections that include attributes. The edge-cost prediction function takes the form of a graph neural network (GNN).

Learning GNN parameters requires solving a bilevel optimization problem that is different from those proposed earlier, e.g. [16], [26]. Specifically, we derive a novel loss function defined with respect to SSP-computed tracks and the ground-truth tracks and an algorithm to learn the GNN parameters by gradient descent. At each iteration, GNN parameters are updated to increase the cost of incorrect tracks, effectively learning from (tracking) mistakes.

The contributions of our paper are as follows:

- A novel end-to-end learnable approach to graph-based tracking with lower computational cost than [16], [26].
- A novel use of SSP for inner optimization, guaranteeing an optimal solution satisfying tracking constraints.
- Quantitative analysis of our algorithm on a diverse set of synthetic scenarios and comparison with a strong baseline that relies on GNN but does not employ global path optimization.

## II. RELATED WORK

**Learning approaches to tracking** In many domains, using machine-learning tools to find optimal parameters for a tracking algorithm from data have been shown to outperform hand-crafting and hand-tuning an algorithm. In the computer vision domain, learning-based approaches to multi-object tracking have taken many different forms. [31] Demonstrated a Markov decision process (MDP) capable of initiating and terminating target tracks; parameters of this MDP were learned from data. In [20] recurrent neural networks were used in an end-to-end framework. Specifically, in [20] a target’s state was maintained by a dedicated recurrent neural network (RNN); assignment of measurements to target tracks and track creation was estimated via a long-short-term memory (LSTM) neural network. A *transformer* neural network was used in [19] as a unified formulation for detecting (new) targets, extending tracks, and creating new tracks.

\*Corresponding author

**Graph-based tracking via classical methods** Application of graph-based and linear-programming (LP) methods to multi-target tracking has a long history, e.g., [6], and includes recent attempts to combine graph-based tracking with multi-hypothesis tracking (MHT) [9]. Graph-based tracking has shown to be effective in computer vision where it is often employed with a tracking-by-detection paradigm where tracking is reduced to the grouping of detections (bounding boxes) proposed by a (target class specific) object detector.

An early graph-based method of [32] proposed the general form of a *tracking graph* and solved for an unknown number of targets via repeated invocation of min-cost flow. The computational complexity was reduced in [24] by introducing *augmenting paths*, thus extending an existing tracking solution to include an additional target instead of starting from scratch. While [32] and [24] performed tracking in the image plane, tracking on the ground plane from multiple camera views via K-shortest-paths (KSP) algorithm was proposed in [3]. Approaches mentioned thus far assumed that targets dynamics are Markovian and targets move independently of each other. An approach to extend LP-based tracking to take into account long-term (group) behavior was proposed in [18]. During training, their method alternated between learning behavior patterns and fitting trajectories to data given the observations and priors given by the learned patterns.

**Graph-based tracking via neural methods** A significant limitation of early graph-based approaches [3], [24], [32] is the requirement to hand-craft scalar costs on tracking graph edges. However, in many real-world scenarios where we wish to apply feature-aided tracking, measurements (detections) in addition to containing positional information also include (high-dimensional) vectors of attributes, such as re-identification (ReID) features. Neural networks could learn a mapping from these high-dimensional attributes to scalar edge weights, but deriving an end-to-end learnable method is nontrivial. Among the first attempts to apply neural methods to graph-based tracking, the approach of [26] stands out for formulating learning and tracking as bilevel [33] optimization. Specifically, in [26] the high-level optimization problem solved for the neural network parameters, while the lower-level problem was a constrained LP that solved a tracking-by-detection problem over a temporal window of pre-defined length. The bilevel approach of [26] was revisited in [16], but instead of relying on the implicit function theorem to deal with the inner LP problem, the latter used KKT conditions to compute the gradient at the lower problem's optimum.

Recent advances in neural methods for graph-structured problems, such as graph neural networks (GNNs) [10] led to advances in graph-based tracking. Inspired by the neural message-passing on graphs method of [11], the authors of [5] proposed to learn a function to estimate edge probabilities in a tracking graph. This approach was extended in [7] to jointly solve for short-term and long-term tracking. While the approaches of [5], [7] require a batch of frames, and thus operate with a constant lag, approaches of [19], [30] require only the current frame to manage tracks. The ideas of [11]




Fig. 1. System diagram of our learnable method.  $f_\theta$  corresponds to Eq.1, explained in detail by Algorithm 1. In our bi-level formulation (Eq.3), learning  $\theta$  is the outer objective, while SSP is the graph-based inner optimization that explicitly produces globally optimal tracks  $\mathcal{P}^*$  in  $\mathcal{G}^{\text{trk}}$ . GNN parameters  $\theta$  are updated in the backward pass through Eq.4.

were adopted in [17] to develop a neural-enhanced message-passing data association algorithm.

**Neural approaches to optimization** Learning-based approaches to MOT mentioned thus far solve the underlying assignment problem using classical methods e.g., LP. An approach to solve quadratic optimization as a layer in a neural network was proposed in [2]. In [23], the matrix of association costs was updated using Sinkhorn-Knopp's algorithm [27] to satisfy MOT constraints. A differentiable graph matching layer was proposed in [12] and used to associate graphs formed on detections across time frames.

**Our contribution** We draw on [10], [11] and [5] to formulate a GNN that infers edge costs on a graph structure, but unlike [5] our inner optimization with SSP ensures that track predictions are globally optimal and satisfy constraints. We demonstrate in Section IV that our method compares favorably to a GNN-based method without such inner optimization. Our method is similar to that of [16] in forming a bilevel problem during training, but we show in Section III-D that our inner optimization with SSP has lower computational complexity.

### III. APPROACH

Let  $\mathbf{x}$  denote target state and let  $\mathbf{y}$  denote measurement (detection). A detection  $\mathbf{y}_{j,t}$  at time  $t$  may originate from one of the targets (if present) or from clutter (i.e., it's a false alarm).

We adopt a sliding-window approach, where tracking is performed on a batch of data frames over a temporal window of length  $T$ . In a complete system, separate logic is used to produce an output that takes into account the tracking solution from the previous temporal window. Figure 1 shows the system diagram of our approach.

#### A. Graph-based formulation

Within the batch of  $T$  data frames, we solve the tracking problem by casting it in a graph-based formulation [24], [32]. Given a temporal window  $t \in 1, \dots, T$  our objective is to group  $\{\mathbf{y}_{j,t}\}$  into a set  $\mathcal{P} \doteq \{p_k\}$  of paths in a *tracking graph*  $\mathcal{G}^{\text{trk}}$  with positive and negative edge costs. The set  $\mathcal{P}$  may be empty if the tracker doesn't predict any tracks.

### B. The detection graph and the tracking graph

Given detections  $\mathcal{Y} \doteq \{\mathbf{y}_{j,t}\}_{t=1}^T$  in a temporal window with  $T$  frames, we construct a detection graph  $\mathcal{G}^{\text{det}}$ . Nodes in  $\mathcal{G}^{\text{det}}$  are in one-to-one correspondence with detections in  $\mathcal{Y}$ . An edge  $(j, j')$  exists between  $\mathbf{y}_{j,t}$  and  $\mathbf{y}_{j',t'}$  if  $t < t' \leq t + \Delta$  and the location components of  $\mathbf{y}_{j,t}$  and  $\mathbf{y}_{j',t'}$  satisfy *gating* constraints. The use of  $\Delta > 1$  is standard [16], [26] to handle short-duration obscuration (occlusion) of a target. In Fig. 2(a) we show an example of  $\mathcal{G}^{\text{det}}$  for  $T = 3$  and  $\Delta = 2$ . Two paths that satisfy node-disjoint constraints are shown.

Given  $\mathcal{G}^{\text{det}}$ , we define its companion *tracking graph*  $\mathcal{G}^{\text{trk}}$ , whose topology follows [6], [32]. In particular, every node in  $\mathcal{G}^{\text{det}}$  is represented by a pair of “twin” nodes in  $\mathcal{G}^{\text{trk}}$ . Furthermore,  $\mathcal{G}^{\text{trk}}$  contains source  $\mathbf{s}$  and terminal  $\mathbf{t}$  nodes with edges from  $\mathbf{s}$  and to  $\mathbf{t}$  defined in a way as to allow every detection to either correspond to a track’s birth or its termination. Fig. 2(b) shows the tracking graph  $\mathcal{G}^{\text{trk}}$  derived from its  $\mathcal{G}^{\text{det}}$ ; the paths (tracks) in  $\mathcal{G}^{\text{trk}}$  originate in  $\mathbf{s}$ , pass through “twin” nodes and terminate in  $\mathbf{t}$ .




Fig. 2. Our formulation employs two types of graphs: (a) *detection graph*  $\mathcal{G}^{\text{det}}$  constructed from measurements in a given temporal window; (b) its corresponding *tracking graph*  $\mathcal{G}^{\text{trk}}$ .

**Computing edge-costs of a tracking graph via MPN** Let  $f_\theta$  be a function, parameterized by  $\theta$  for computing edge costs in a graph  $\mathcal{G}^{\text{det}} = (\mathcal{V}, \mathcal{E}, \mathcal{Y})$

$$f_\theta : \mathcal{G}^{\text{det}} \rightarrow \{c_e\}_{e=1}^{|\mathcal{E}|}. \quad (1)$$

Our  $f_\theta$  takes the form of a message-passing network (MPN) [11], a specialization of a graph neural net (GNN) [10]. We adopt the ideas of [11] and [14] as implemented in [5].

Let  $L$  be an integer representing the number of message passing layers (iterations),  $\mathcal{E}$  be the edge set, and  $\mathcal{V}$  be the vertex set of  $\mathcal{G}^{\text{det}}$ . Let  $\mathbf{h}_{(i,j)}^{(l)}$  be the embedding of edge  $(i, j) \in \mathcal{E}$  at iteration  $l$  and similarly  $\mathbf{h}_z^{(l)}$  be the embedding of the vertex  $z \in \mathcal{V}$  at iteration  $l$ . We assume without loss of generality that node  $j$  is ahead of node  $i$  in time.

Messages are computed using the functions  $g_e(\cdot, \cdot, \cdot)$  and  $g_v(\cdot, \cdot, \cdot)$  which take the form of learnable multi-layer perceptions. In our implementation, as suggested by [5],  $g_v(\cdot, \cdot, \cdot)$  consists of two independently parameterized functions  $g_v^{\text{past}}(\cdot, \cdot, \cdot)$ ,  $g_v^{\text{fut}}(\cdot, \cdot, \cdot)$  which are specialized for forward and backward in time messages respectively. The readout function  $g_e^{\text{readout}}(\cdot, \cdot, \cdot)$ , also a learnable MLP, then computes the edge cost by aggregating the final attributes of the forward and backward in time edges between each pair of nodes. The parameters of these four functions comprise the vector  $\theta$ .

### Algorithm 1: Computing edge costs in $\mathcal{G}^{\text{det}}$

---

```

1 Initialize  $\mathbf{h}_{(i,j)}^{(0)}$  and  $\mathbf{h}_z^{(0)}$  for all  $(i, j) \in \mathcal{E}, z \in \mathcal{V}$ 
2  $l \leftarrow 1$ 
3 while  $l \leq L$  do
4   for  $(i, j) \in \mathcal{E}$  do
5      $\mathbf{h}_{(i,j)}^{(l)} = g_e([\mathbf{h}_i^{(l-1)}, \mathbf{h}_j^{(l-1)}, \mathbf{h}_{(i,j)}^{(l-1)}])$ 
6      $m_{j,(i,j)}^{(l)} = g_v^{\text{past}}([\mathbf{h}_j^{(l-1)}, \mathbf{h}_{(i,j)}^{(l)}])$ 
7      $m_{i,(i,j)}^{(l)} = g_v^{\text{fut}}([\mathbf{h}_i^{(l-1)}, \mathbf{h}_{(i,j)}^{(l)}])$ 
8   for  $z \in \mathcal{V}$  do
9      $\mathbf{h}_{z,\text{past}}^{(l)} = \sum_{(i,z) \in \mathcal{E}} m_{z,(i,j)}^{(l)}$ 
10     $\mathbf{h}_{z,\text{fut}}^{(l)} = \sum_{(z,j) \in \mathcal{E}} m_{z,(z,j)}^{(l)}$ 
11     $\mathbf{h}_z^{(l)} = g_v([\mathbf{h}_{z,\text{past}}^{(l)}, \mathbf{h}_{z,\text{fut}}^{(l)}])$ 
12   $l \leftarrow l + 1$ 
13 for  $(i, j) \in \mathcal{E}$  do
14    $c_{(i,j)} = g_e^{\text{readout}}(\mathbf{h}_{(i,j)}^{(L)}, \mathbf{h}_{(j,i)}^{(L)})$ 
15 Output:  $\{c_e\}_{e=1}^{|\mathcal{E}|}$ 

```

---

### C. End-to-end learning for tracking

**Optimization criterion** To derive our algorithm to learn  $f_\theta$  we need a differentiable optimization criterion (a loss function). The loss function should be defined with respect to the ground-truth tracks (paths)  $\mathcal{P}^+$  and the predicted ones  $\mathcal{P}^*$ .

Given  $\mathcal{G}^{\text{trk}} = (\mathcal{V}, \mathcal{E}, \mathcal{Y})$  with edge costs  $\{c_e\}_{e=1}^{|\mathcal{E}|}$  we rewrite the edge costs as a real-valued function  $c(\cdot)$ . For a track (path)  $p \in \mathcal{G}^{\text{trk}}$ , we define (with slight abuse of notation)  $c(p) = \sum_{e \in p} c(e) + c^{\text{en}} + c^{\text{ex}}$ , where  $c^{\text{en}} > 0$ ,  $c^{\text{ex}} > 0$  are entrance and exit costs. For a set of paths  $\mathcal{P}$ , we define  $c(\mathcal{P}) = \sum_{p \in \mathcal{P}} c(p)$ .

We define the *set loss*  $\mathcal{L}$  for two sets of paths  $\mathcal{P}^+, \mathcal{P}^*$  as

$$\mathcal{L}(\mathcal{P}^+, \mathcal{P}^*) = \ell_1(\mathcal{P}^+, \mathcal{P}^*) + \ell_2(\mathcal{P}^+). \quad (2)$$

The first term is defined  $\ell_1(\mathcal{P}^+, \mathcal{P}^*) = \max(0, c(\mathcal{P}^*) - c(\mathcal{P}^+))$  and the second  $\ell_2(\mathcal{P}^+) = \sum_{p \in \mathcal{P}^+} \max(0, c(p))$ .

The following claim follows immediately from the properties of an SSP algorithm.

**Claim 1.** Let  $\mathcal{P}^*$  be the estimated set of paths. Then  $c(\mathcal{P}^*) \leq 0$  and for all  $p \in \mathcal{P}^*$ ,  $c(p) \leq 0$ .

The next claim confirms that optimizing  $\mathcal{L}$  converges to  $\mathcal{P}^+$ .

**Claim 2.** Suppose that the cost of each unique set of mutually disjoint paths in the tracking graph is unique. Then  $\ell_1(\mathcal{P}^+, \mathcal{P}^*) = 0$  if and only if  $\mathcal{P}^+ = \mathcal{P}^*$  as sets. If these equivalent conditions hold,  $\mathcal{L}(\mathcal{P}^+, \mathcal{P}^*) = 0$  as well.

The first part follows from the fact that  $\mathcal{P}^*$  is the global optimum, but so is  $\mathcal{P}^+$ , by definition. The second part follows by showing that  $\ell_2(\mathcal{P}^+) = 0$ , and hence  $\mathcal{L}$  is zero.

**Bilevel optimization to learn MPN parameters** We formulate the following bilevel optimization problem [8], [21], [33]

to find the optimal set of GNN parameters given ground-truth paths  $\mathcal{P}^+$  :

$$\begin{aligned} \theta^* &= \underset{\theta}{\operatorname{argmin}} \mathcal{L}(\mathcal{P}^+, \mathcal{P}^*) \\ \text{s.t., } \mathcal{P}^* &= \text{track\_by\_ssp}(\mathcal{G}, \{c_e\}_{e=1}^{|\mathcal{E}|}) \end{aligned} \quad (3)$$

The problem defined in Equation 3 is non-trivial: the outer objective is non-linear, while the inner objective is stated in terms of a constrained search for minimum-cost paths in a graph. Applying the chain rule to our loss, yields

$$\frac{\partial \mathcal{L}^{\text{SSP-GNN}}}{\partial \theta} = \frac{\partial \ell}{\partial \mathbf{c}} \bigg|_{\mathcal{P}^*} \frac{\partial \mathbf{c}}{\partial \theta}, \quad (4)$$

where  $\mathbf{c} \doteq \{c_e\}_{e=1}^{|\mathcal{E}|}$ .

In practice, we implement optimization of  $\theta$  in Equation 3 via a two-stage learning algorithm. In Stage I, inspired by [22] which introduced MCMC data association *moves*, we “perturb”  $\mathcal{P}^+$  to pre-compute a fixed-size set of negative examples  $\mathcal{P}^-$ . Generating  $\mathcal{P}^-$  does not require running the SSP algorithm.

Let  $p^- \in \mathcal{P}^-$  be one of the paths (tracks) obtained via perturbation, and let  $p^+ \in \mathcal{P}^+$  be a ground-truth path. In general, given  $\theta^*$ , it is not the case that every  $p^-$  has a higher cost than every  $p^+$ . Nevertheless for Stage I we define

$$\ell^{+-}(p^+, p^-) = \max(0, \text{cost}(p^+) - \text{cost}(p^-) + \gamma) \quad (5)$$

for some margin  $\gamma \geq 0$ . And we can extend this to the entire set  $\mathcal{P}^-$  of all the paths obtained by perturbing  $\mathcal{P}^+$ . Optimizing  $\ell^{+-}$  leads to  $\hat{\theta}$  which serves as a “warm start” for Stage II which implements Equation 3.

Our end-to-end learning approach is defined in Algorithm 2; extending it to multiple training graphs is straightforward. We will demonstrate in the experiments section that learning  $f_\theta$  this way yields a stable training algorithm.

#### D. Computational Complexity

In our analysis, to provide a fair comparison between our proposed algorithm and related methods we focus only on forward inference of the inner optimization problem used to compute  $\mathcal{P}^*$  from  $\mathcal{G}^{\text{trk}}$ . In our implementation, this is solved using SSP, specifically the  $k$ -shortest-paths algorithm. As shown in [3], this algorithm has a worst-case time complexity of  $\mathcal{O}(k(m+n \cdot \log n))$ , where  $k$  is the number of paths/tracks found by the algorithm and  $m$  and  $n$  refer to the number of nodes and edges, respectively, contained in the graph  $\mathcal{G}^{\text{trk}}$ . The size and topology of  $\mathcal{G}^{\text{trk}}$  depends on  $\mathcal{Y}$ , and therefore is a function of both the size of the temporal window and the number of targets and false alarms.

In contrast, linear programming (LP)-based approaches to the min-cost flow problem have an average-case polynomial time complexity in the size of  $\mathcal{G}^{\text{trk}}$  [1]. This also applies to more recent approaches presented in [26] and [16], which embed the LP solution as a differentiable layer within a multi-layer perceptron.<sup>1</sup> The proposed SSP-GNN algorithm blends

<sup>1</sup>Note that [16] transforms the LP problem to a regularized quadratic program; however, the transformed problem has the same time complexity.

---

#### Algorithm 2: Training SSP-GNN

---

```

1 Inputs :
2  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{Y})$  (training) graph with node features  $\mathcal{Y}$ 
3  $\mathcal{P}^+ \triangleq \mathcal{P}^{\text{GT}}$  ground-truth paths (tracks) in  $\mathcal{G}$ 
4 Stage I parameters  $M, N_1^{\text{max}}$ ; Stage II parameter  $\epsilon \geq 0$ 
5  $f_\theta : \mathcal{G} \rightarrow \{c_e\}_{e=1}^{|\mathcal{E}|}, c_e \in \mathbb{R}$ 
6 // Stage I (bootstrapping)
7  $\mathcal{P}^- \triangleq \{p_1^-, \dots, p_M^-\}$ 
8  $\mathcal{P}^- \leftarrow \text{perturb\_paths}(\mathcal{G}, \mathcal{P}^+, M)$ 
9  $\theta \leftarrow \text{init\_gnn\_weights}()$ 
10 while  $\text{itr} < N_1^{\text{max}}$  do
11    $\{c_e\}_{e=1}^{|\mathcal{E}|} \leftarrow f_\theta(\mathcal{G})$ 
12   // compute loss  $\ell^{+-}(\mathcal{P}^+, \mathcal{P}^-)$ ; update  $\theta$ 
13 // Stage II
14 while True do
15    $\{c_e\}_{e=1}^{|\mathcal{E}|} \leftarrow f_\theta(\mathcal{G})$ 
16    $\mathcal{P}^* \leftarrow \text{track\_by\_ssp}(\mathcal{G}, \{c_e\}_{e=1}^{|\mathcal{E}|})$ 
17   if  $\mathcal{L}(\mathcal{P}^+, \mathcal{P}^*) < \epsilon$  then
18     break
19   update  $\theta$ 
20 Output:  $f_\theta$ 

```

---

the edge-cost expressivity of learning-based methods with the reduced time complexity of SSP methods.

## IV. EXPERIMENTS

We conduct extensive experiments to demonstrate the advantages of our method and to assess the impact of a variety of model and scenario attributes on performance.

#### A. Train and Test Scenarios

Our synthetic data is generated using Stone Soup<sup>2</sup>, an open-source framework. We generate diverse scenarios with multiple targets using constant velocity and Ornstein-Uhlenbeck motion models. When generating detections, we fix the detection probability ( $P_D$ ) at 1.0 except in Table VI which investigates the effects of an imperfect detector, and we fix measurement error standard deviation at 0.2. We refer to target features as re-identification (ReID) as is common in computer vision. However, our approach can generalize to other features and sensing modalities, e.g., [15]. Our ReID features are drawn from a mixture of Gaussians (MoG) distributions. We control the “strength” of ReID features by changing the amount of overlap between the foreground and background distributions and define categories in terms of KL divergence: *very weak* (0.02 nats/0.029 bits), *weak* (0.5 nats/0.721 bits), *moderate* (3.125 nats/4.508 bits), and *strong* (12.5 nats/18.034 bits).

#### B. Implementation Details

Given a pair of detections  $\mathbf{y}_{i,t}, \mathbf{y}_{i',t'}$  with  $t \neq t'$  we follow [16] to encode their spatio-temporal constraint into our

<sup>2</sup><https://github.com/dstl/Stone-Soup/>

latent representation. We denote the kinematic (spatial) component of a detection as  $\mathbf{y}_{i,t}^{(s)}$ ,  $\mathbf{y}_{i',t'}^{(s)}$ , the ReID component as  $\mathbf{y}_{i,t}^{(id)}$ ,  $\mathbf{y}_{i',t'}^{(id)}$ . For a node  $v$  in  $\mathcal{G}$ , we define  $\mathbf{h}_v^{(0)} = [\mathbf{y}_i^{(s)}, \mathbf{y}_i^{(id)}]$ . For an edge  $e$  between the two detections in  $\mathcal{G}$ , we define  $\mathbf{h}_e^{(0)} = [t' - t, \mathbf{y}_{i'}^{(s)} - \mathbf{y}_i^{(s)}, \|\mathbf{y}_{i'}^{(id)} - \mathbf{y}_i^{(id)}\|]$ . We use 2D scenarios and 2D ReID features, i.e.,  $\mathbf{y}^{(s)} \in \mathbb{R}^2$  and  $\mathbf{y}^{(id)} \in \mathbb{R}^2$ .

**Tracking graph and inference** We instantiate  $f_\theta$  and learn parameters  $\theta$  on the detection graph  $\mathcal{G}^{\text{det}}$  as shown in Fig. 2(a). Then,  $f_\theta(\mathcal{G}^{\text{det}})$  assigns scalar costs to the edges in  $\mathcal{G}^{\text{det}}$ . We copy these edge costs from  $\mathcal{G}^{\text{det}}$  to  $\mathcal{G}^{\text{trk}}$  by associating each edge  $e_{i,j} \in \mathcal{G}^{\text{det}}$  with the edge between the pair of twin nodes associated with  $i$  and the pair of twin nodes associated with  $j$  in  $\mathcal{G}^{\text{trk}}$ . We assign a cost of 0 to the edges within the twin nodes, shown as red arrows in Figure 2(b). This formulation allows us to use the GNN-computed  $\{c_e\}_{e=1}^{|\mathcal{E}|}$  as edge costs for the SSP algorithm to optimize over  $\mathcal{G}^{\text{trk}}$ . Detection and tracking graphs are represented via NetworkX<sup>3</sup> except for GNN learning and inference the detection graph is converted to a Deep Graph Library (DGL) [29] object.

**Training** We use the Adam optimizer [13] with a learning rate of 0.001. Fig. 3 demonstrates a typical loss curve across 20 re-initialized training trials on the same train set. In practice we observe that  $\ell_1$  tends to dominate total loss  $\mathcal{L}$  because (a) our Stage I bootstrapping guarantees that  $\ell_2 \approx 0$  at epoch 0 and (b)  $\ell_2$  captures the more trivial problem of simply keeping the ground truth paths sufficiently negative in cost, while  $\ell_1$  addresses the more difficult problem of predicting the ground truth paths over competing sets of paths.




Fig. 3. The average Stage II loss values across 20 SSP-GNN models with randomly initialized weights. As the number of epochs increases, the average loss converges to zero with decreasing standard deviation (shaded).

### C. Evaluation Metrics

Our method is inspired by [5], [16], etc., which were developed for tracking objects in video sequences. These methods tend to be evaluated with respect to Multiple Object Tracking Accuracy (MOTA) [4]. While our method is applicable to tracking extended targets in video sequences, we generate point observations in our Stone Soup simulation. We therefore report the MOTA metrics calculated with respect to the point

observations for all of our experiments. We also report GOSPA [25] and SIAP [28] metrics when comparing our model’s performance to a filter-based tracker provided by Stone Soup.

Following [4], MOTA is defined as:

$$\text{MOTA} = 1 - \frac{\sum_t |\text{FP}_t| + |\text{FN}_t| + |\text{IDS}_t|}{\sum_t |\text{gtDets}_t|}, \quad (6)$$

where  $t$  is the time step,  $|\text{FP}_t|$  is the number of false positives,  $|\text{FN}_t|$  is the number of false negatives, and  $|\text{IDS}_t|$  is the number of identity switches. We report MOTA, as well as each constituent metric as ratios with the number of ground truth detection’s (rates). Since we control data generation, we can utilize ID-based matching between tracks and true detections instead of the traditional Hungarian matching algorithm typically used for computer vision applications. The ability to perform more straightforward ID-based matching for evaluation is one advantage of using synthetic data.

Both GOSPA and SIAP metrics are popular evaluation methods for comparing predicted tracks to continuous ground-truth object paths. GOSPA quantifies localization and cardinality errors, and SIAP completeness (C), ambiguity (A), spuriousness (S), and positional error, provide a more complete insight into the accuracy of associations between predicted tracks and ground truth objects. We report metrics summed over all time frames for GOSPA and averaged for SIAP.

### D. Qualitative Demonstration

For demonstration, in Figure 4 we show a sample SSP-GNN result trained on one scenario. Both the training and test scenarios last 100 timesteps and include three targets with randomly-generated lifespans following Ornstein-Uhlenbeck dynamics. Each scenario uses moderate strength ReID features and 20 uniformly distributed false alarms per frame. We observe that our algorithm generalizes well to the test scenario (the circled predicted tracks follow ground truth paths). The MOTA score is 0.993 for train scenario and 0.987 for the test scenario.

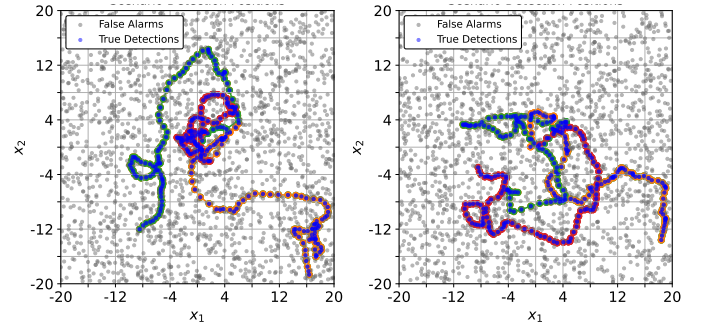


Fig. 4. Predicted tracks for the train (left) and test (right) scenario.

### E. Quantitative Evaluations

Since we are working with synthetic data, we can generate datasets of varying sizes and complexity. This enables isolating specific aspects of train and test scenarios to examine the effect on our model. However, since not all published methods report such analysis, and there is no standardized evaluation protocol,

<sup>3</sup><https://networkx.org/>



our results cannot be immediately compared to others reported in literature.

**Effect of GNN hyperparameters** We first determine the optimal value of two hyperparameters used in configuring the graph neural network: number of message passing layers and size of the hidden layers  $h$  in all used MLP's. Figure 5 shows the MOTA performance of the model on fixed train and test sets with variable number of GNN layers and hidden dimension size. The scenarios involve multiple ground truth tracks with constant velocity motion models, moderate strength ReID features, and 8 false alarms per frame in expectation.

As the number of message passing layers in the graph neural network increases model performance also tends to increase until saturation is reached, at which point it largely levels out. Using a hidden dimension size of 64 produced the best results overall, and it achieves saturation at 4 GNN layers. Therefore, we use these hyperparameter values in all experiments going forward. When working with substantially different environments from our synthetic data, the optimal values for these hyperparameters would need to be adapted.




Fig. 5. Effect of the number of GNN message passing layers and hidden dimension size. Note that as the model capacity increases so does the tracking accuracy, until saturation.

**Comparison to baselines** We compare our SSP-GNN tracker to an Extended Kalman Filter (EKF) based multi-object tracker and Edge-Belief GNN (based on [5]). The EKF tracker implementation is provided by the Stone Soup library. Although EKF only processes the kinematic component of measurements and is an online method (while ours is offline/batch-based), it establishes a reference point in our comparison.

The Edge-Belief GNN algorithm [5] uses the same MPN architecture as our SSP-GNN, but instead of predicting edge costs, Edge-Belief GNN individually estimates the probability of an edge in the detection graph being part of a track or not. This is accomplished by modifying  $g_e^{\text{readout}}(\cdot, \cdot)$  to output sigmoid rather than logits. During inference, edge beliefs are rounded up or down to yield binary predictions, and a follow-on edge traversal heuristic ensures that the computed paths are node-disjoint. We observed that the Edge-Belief GNN tends to produce short (spurious) tracks; we added a post processing step which removes tracks shorter than a set length in order to strengthen the baseline (empirically we found minimum path

threshold = 3 was optimal).

Figure 6 shows tracks produced by an SSP-GNN model and Edge-Belief GNN model on a test scenario. In this scenario, there are several instances where two tracks intersect. Our model correctly maintains target tracks on the left, while it swaps paths causing two identity switches on the right. Edge-belief GNN tends to fragment or deviate from tracks to nearby false alarms much more frequently than the SSP-GNN.

Table I compares Edge-Belief GNN, SSP-GNN, and EKF on similar scenarios. As inferred from the visual demonstration, the false negative rate is nearly twice that of SSP-GNN. Edge-Belief also has a substantially higher identity switch rate, indicating track fragmenting, but it does achieve a comparable false positive rate to our model.

Table II compares Edge-Belief GNN and SSP-GNN on similarly structured scenarios, but with increased difficulty due to false alarm rate and ReID feature weakness, (the EKF tracker is omitted as ignoring ReID features makes it uncompetitive). SSP-GNN consistently outperforms Edge-Belief by a sizable margin, which is even more pronounced in harder regimes.

Because the Edge-Belief GNN baseline uses the same underlying message passing algorithm as SSP-GNN, we ascribe the observed performance increase to bilevel optimization. The SSP solver allows our model to properly weigh various trade-offs in the context of the full scenario. For instance it can weigh the cost of including a series of positive cost edges that would link otherwise disparate tracks versus the extra entrance/exit cost of inferring more tracks.

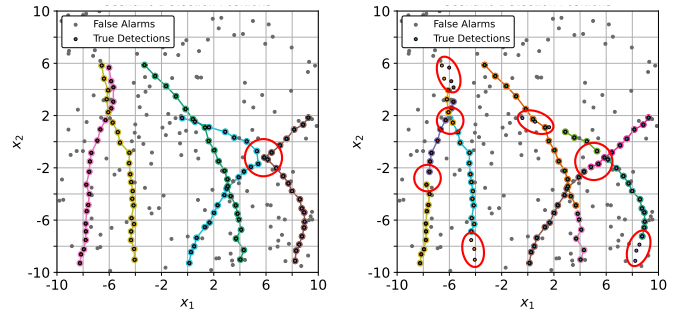


Fig. 6. SSP-GNN (left) and Edge-Belief GNN (right) algorithms trained and tested on identical sets. Small, colored dots with black outline are detection of true paths, while the larger circles linked together are algorithm generated tracks. Edge-Belief GNN tends to fragment tracks more often. Errors in the tracks are marked with red circles. Note false alarm density is 60% higher than in Figure 4.

**Effect of ReID feature strength and noisy dimensions** We use one trajectory following an Ornstein-Uhlenbeck motion model over a 50 timestep period with a high false alarm rate (expected 30 per frame). Table III shows a clear positive correlation between model performance and ReID feature strength. The two extremes of the table also indicate that the model leverages strong ReID features effectively, while also allowing for good motion-based predictions when ReID features are highly unreliable.

While it is clear that our method can handle higher-dimensional ReID features, it is interesting to consider a case

TABLE I  
COMPARISON WITH BASELINE TRACKERS.  $\uparrow$  INDICATES HIGHER IS BETTER. PARENTHESES NEXT TO METRICS INDICATE A TARGET VALUE.

Tracker	Metrics											
	MOTA				SIAP				GOSPA			
	MOTA $\uparrow$	FPR $\downarrow$	FNR $\downarrow$	IDS Rate $\downarrow$	C (1)	A (1)	S (0)	Pos. Err. (0)	Dist. $\downarrow$	Local. $\downarrow$	Missed $\downarrow$	False $\downarrow$
EKF					0.945	1.065	0.055	0.164	23.130	14.130	2.750	6.250
Edge-Belief	0.915 $\pm$ 0.008	<b>0.0175</b>	0.041	0.027	0.9605	1.009	<b>0.008</b>	0.129	8.772	<b>2.073</b>	1.940	<b>0.710</b>
SSP-GNN (ours)	<b>0.942<math>\pm</math>0.050</b>	0.022	<b>0.022</b>	<b>0.013</b>	<b>0.977</b>	<b>1.000</b>	0.023	<b>0.128</b>	<b>7.776</b>	2.441	<b>0.720</b>	0.720

TABLE II  
COMPARISON OF SSP-GNN VS EDGE BELIEF ON SCENARIOS OF DIVERSE COMBINATIONS OF FALSE ALARMS (FA) AND REID STRENGTH

Tracker	FA Rate	ReID Strength	MOTA $\uparrow$
Edge-Belief	10	Moderate	0.862 $\pm$ 0.012
SSP-GNN	10	Moderate	<b>0.950 <math>\pm</math> 0.019</b>
Edge-Belief	10	Weak	0.826 $\pm$ 0.008
SSP-GNN	10	Weak	<b>0.952 <math>\pm</math> 0.006</b>
Edge-Belief	20	Moderate	0.656 $\pm$ 0.032
SSP-GNN	20	Moderate	<b>0.891 <math>\pm</math> 0.034</b>
Edge-Belief	20	Weak	0.633 $\pm$ 0.018
SSP-GNN	20	Weak	<b>0.873 <math>\pm</math> 0.023</b>

where no dimensions are equally informative. We therefore constructed scenarios with up to ten additional ReID feature dimensions filled with i.i.d. noise. While training took longer, our learned tracker did not suffer any noticeable degradation in accuracy, indicating that the algorithm is able to find informative components of ReID and ignore the rest.

TABLE III  
EFFECT OF DECREASING THE STRENGTH OF REID FEATURES

ReID Str.	MOTA $\uparrow$	FP Rate $\downarrow$	FN Rate $\downarrow$	IDS Rate $\downarrow$
Strong	0.961 $\pm$ 0.031	0.015	0.025	0.000
Moderate	0.918 $\pm$ 0.023	0.041	0.041	0.000
Weak	0.909 $\pm$ 0.009	0.045	0.045	0.000
Very weak	0.893 $\pm$ 0.021	0.048	0.048	0.000

TABLE IV  
EFFECT OF INCREASING THE RATE OF FALSE ALARMS

FA per frame	MOTA $\uparrow$	FP Rate $\downarrow$	FN Rate $\downarrow$	IDS Rate $\downarrow$
2	0.994 $\pm$ 0.004	0.003	0.003	0.001
5	0.979 $\pm$ 0.007	0.009	0.009	0.004
10	0.938 $\pm$ 0.047	0.029	0.029	0.005
20	0.900 $\pm$ 0.040	0.048	0.046	0.006
30	0.907 $\pm$ 0.017	0.045	0.046	0.003

**Effect of false alarm rate** Table IV displays the performance of our method at a variety of expected numbers of false alarms per frame. The performance is nearly optimal when false alarms are sparse, and deteriorates gracefully as false alarms are added, however it seems to level off and not drop below a MOTA score of about 0.9. This indicates that our algorithm is capable of capitalizing on easier regimes with relatively few false positives, while also adapting to hard environments where

false alarms significantly outnumber ground truth tracks.

**Effect of the training set size** Table V test our model with a variety of training set sizes. Our experiments use limited size training sets so as to run multiple trials. Real world applications would likely train on more data than we use, so it is important that model accuracy scales with training set size.

This experiment uses a diverse array of scenarios using constant velocity and Ornstein–Uhlenbeck motion models with variable ReID strength (moderate to weak) and up to 5 ground truth tracks. Table V shows a positive correlation between number of graphs in the training set and model performance. Learned tracker consistency, represented by standard deviation, also improves substantially with training set size.

TABLE V  
EFFECT OF INCREASING TRAINING SET SIZE

Train Set Size	MOTA $\uparrow$	FP Rate $\downarrow$	FN Rate $\downarrow$	IDS Rate $\downarrow$
1	0.894 $\pm$ 0.046	0.045	0.053	0.007
5	0.937 $\pm$ 0.032	0.024	0.032	0.007
10	0.954 $\pm$ 0.006	0.020	0.020	0.007
20	0.958 $\pm$ 0.019	0.016	0.016	0.011
30	0.963 $\pm$ 0.007	0.014	0.014	0.009

**Effect of lower detection probability** Finally, we consider the effect of an imperfect detector i.e.,  $P_D < 1.0$ . To allow the model to select trajectories that skip time steps, we modify the tracking graph to contain edges between non-adjacent timesteps up to three apart.

In Table VI we observe a drop in MOTA of around 0.04 when detection probability decreases from 1.0 to 0.95, as adding missed detections increases the complexity of the representation that must be learned. Our model performance is resilient between detection probabilities of 0.95 to 0.85. From 0.85 to 0.8, there is another drop in performance, which may be due to the increasing prevalence of enduring missed detections lasting beyond what the model is capable of considering.

## V. CONCLUSIONS

We presented a novel trainable approach for graph-based multi-object tracking. Our approach predicts edge costs using a message passing GNN, then relies on a globally optimal and computationally-efficient path-finding algorithm to form tracks that satisfy the required constraints. The message-passing network can be learned end-to-end in a bilevel optimization

TABLE VI  
EFFECT OF DECREASING  $P_D$

$P_D$	MOTA $\uparrow$	FP Rate $\downarrow$	FN Rate $\downarrow$	IDS Rate $\downarrow$
1.0	0.978 $\pm$ 0.001	0.010	0.010	0.003
0.95	0.957 $\pm$ 0.013	0.022	0.014	0.007
0.9	0.945 $\pm$ 0.020	0.024	0.026	0.004
0.85	0.944 $\pm$ 0.012	0.031	0.022	0.002
0.8	0.913 $\pm$ 0.017	0.048	0.034	0.006

framework that is different from prior work. We thoroughly investigated model performance on a variety of scenarios. Our learned tracker compares favorably to a strong baseline. In our future work, we plan to focus on computationally-efficient extensions to our graph formulation and the message-passing network to optimally handle long-term occlusions.

## VI. ACKNOWLEDGEMENTS

We thank the anonymous reviewers for their feedback and Samuel C. Buckner for helpful discussions. Earlier work funded by U.S. Office of Naval Research Code 321.

## REFERENCES

- [1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. Network flows: Theory, algorithms, and applications. 1993.
- [2] B. Amos and J. Z. Kolter. OptNet: Differentiable optimization as a layer in neural networks. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 136–145. PMLR, 06–11 Aug 2017.
- [3] J. Berclaz, F. Fleuret, E. Turetken, and P. Fua. Multiple object tracking using K-shortest paths optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(9):1806–1819, 2011.
- [4] K. Bernardin and R. Stiefelhausen. Evaluating multiple object tracking performance: The clear mot metrics. *EURASIP Journal on Image and Video Processing*, 2008, 01 2008.
- [5] G. Brasó and L. Leal-Taixé. Learning a neural solver for multiple object tracking. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [6] D. Castañón. Efficient algorithms for finding the K best paths through a trellis. *IEEE Transactions on Aerospace and Electronic Systems*, 26(2):405–410, 1990.
- [7] O. Cetintas, G. Brasó, and L. Leal-Taixé. Unifying short and long-term tracking with graph hierarchies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22877–22887, June 2023.
- [8] B. Colson, P. Marcotte, and G. Savard. An overview of bilevel optimization. *Annals of Operations Research*, 153:235–256, Sept. 2007.
- [9] S. P. Coraluppi, C. A. Carthel, and A. S. Willsky. Multiple-hypothesis tracking and graph-based tracking extensions. *Journal of Advances in Information Fusion*, 14(2), 2019.
- [10] M. Defferrard, X. Bresson, and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NeurIPS’16, page 3844–3852, Red Hook, NY, USA, 2016. Curran Associates Inc.
- [11] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1263–1272, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- [12] J. He, Z. Huang, N. Wang, and Z. Zhang. Learnable graph matching: Incorporating graph partitioning with deep feature learning for multiple object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5299–5309, June 2021.
- [13] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2017.
- [14] T. Kipf, E. Fetaya, K.-C. Wang, M. Welling, and R. Zemel. Neural relational inference for interacting systems. In J. Dy and A. Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2688–2697. PMLR, 10–15 Jul 2018.
- [15] J. LeNoach, M. Lexa, and S. Coraluppi. Feature-aided tracking techniques for active sonar applications. In *IEEE 24th International Conference on Information Fusion (FUSION)*, pages 1–7. IEEE, 2021.
- [16] S. Li, Y. Kong, and H. Rezatofighi. Learning of global objective for network flow in multi-object tracking. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8845–8855. IEEE, 2022.
- [17] M. Liang and F. Meyer. Neural enhanced belief propagation for data association in multiobject tracking. In *2022 25th International Conference on Information Fusion (FUSION)*, pages 1–7, 2022.
- [18] A. Maksai, X. Wang, F. Fleuret, and P. Fua. Non-Markovian globally consistent multi-object tracking. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [19] T. Meinhardt, A. Kirillov, L. Leal-Taixé, and C. Feichtenhofer. Trackformer: Multi-object tracking with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8844–8854, June 2022.
- [20] A. Milan, S. H. Rezatofighi, A. Dick, I. Reid, and K. Schindler. Online multi-target tracking using recurrent neural networks. In *Proceedings of the AAAI conference on Artificial Intelligence*, volume 31, 2017.
- [21] P. Ochs, R. Ranftl, T. Brox, and T. Pock. Bilevel optimization with nonsmooth lower level problems. In J.-F. Aujol, M. Nikolova, and N. Papadakis, editors, *Scale Space and Variational Methods in Computer Vision*, pages 654–665, Cham, 2015. Springer International Publishing.
- [22] S. Oh, S. Russell, and S. Sastry. Markov chain monte carlo data association for multi-target tracking. *IEEE Transactions on Automatic Control*, 54(3):481–497, 2009.
- [23] I. Papakis, A. Sarkar, and A. Karpatne. GCNNMatch: Graph convolutional neural networks for multi-object tracking via sinkhorn normalization. *arXiv preprint arXiv:2010.00067*, 2020.
- [24] H. Pirsiavash, D. Ramanan, and C. C. Fowlkes. Globally-optimal greedy algorithms for tracking a variable number of objects. In *CVPR 2011*, pages 1201–1208, 2011.
- [25] A. S. Rahmathullah, A. F. Garcia-Fernandez, and L. Svensson. Generalized optimal sub-pattern assignment metric. In *2017 20th International Conference on Information Fusion (Fusion)*. IEEE, July 2017.
- [26] S. Schuster, P. Vernaza, W. Choi, and M. Chandraker. Deep network flow for multi-object tracking. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2730–2739, 2017.
- [27] R. Sinkhorn and P. Knopp. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21(2):343–348, 1967.
- [28] P. Votruba, R. Nisley, R. Rothrock, and B. Zombro. Single integrated air picture (SIAP) metrics implementation. In *Tech. Rep. Single Integrated Air Picture (SIAP) System Engineering Task Force*, Arlington, VA, 2001.
- [29] M. Wang, D. Zheng, Z. Ye, Q. Gan, M. Li, X. Song, J. Zhou, C. Ma, L. Yu, Y. Gai, T. Xiao, T. He, G. Karypis, J. Li, and Z. Zhang. Deep graph library: A graph-centric, highly-performant package for graph neural networks. *arXiv preprint arXiv:1909.01315*, 2019.
- [30] Y. Wang, K. Kitani, and X. Weng. Joint Object Detection and Multi-Object Tracking with Graph Neural Networks. *arXiv:2006.13164*, 2020.
- [31] Y. Xiang, A. Alahi, and S. Savarese. Learning to track: Online multi-object tracking by decision making. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [32] L. Zhang, Y. Li, and R. Nevatia. Global data association for multi-object tracking using network flows. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.
- [33] Y. Zhang, P. Khanduri, I. Tsaknakis, Y. Yao, M. Hong, and S. Liu. An introduction to bi-level optimization: Foundations and applications in signal processing and machine learning, 2023.