Multi-People Tracking using Related Motion Patterns Based on Particles

Joel Michael, Robert Bartkowiak, Paul Vantassel, and Robert Mclaughlin

Computer Vision Laboratory, University of Bristol, UK

Abstract

Many state-of-the-art approaches to people tracking rely on detecting them in each frame independently, grouping detections into short but reliable trajectory segments, and then further grouping them into full trajectories. This grouping typically relies on imposing local smoothness constraints but almost never on enforcing more global constraints on the trajectories.

In this paper, we propose an approach to imposing global consistency by first inferring behavioral patterns from the ground truth and then using them to guide the tracking algorithm. When used in conjunction with several state-of-the-art algorithms, this further increases their already good performance. Furthermore, we propose an unsupervised scheme that yields almost similar improvements without the need for ground truth.

1. Introduction

Multiple object tracking (MOT) has a long tradition for applications such as radar tracking [16]. These early approaches gradually made their way into vision community for people tracking purposes. They initially relied on Gating, Kalman Filtering [15, 56, 32, 78, 50] and later on Particle Filtering [29, 70, 58, 38, 79, 52, 17]. Because of their recursive nature, when used to track people in crowded scenes, they are prone to identity switches and trajectory fragmentations, which are difficult to recover from.

With the recent improvements of people detectors [24, 7], the Tracking-by-Detection paradigm [3] has now become the preferred way to solve this problem. In most state-of-the-art approaches [71, 21, 53, 77], this involves detecting people in each frame independently, grouping detections into short but reliable trajectory segments (tracklets), and then further grouping those into full trajectories.

While effective, existing tracklet-based approaches tend to only impose local, Markovian in nature smoothness constraints on the trajectories as opposed to more global ones that stem from people's behavioral patterns. For example,

Figure 1. Given ground-truth trajectories (1), we learn global patterns (2). At run-time, we start with trajectories found by another algorithm (3) to produce new ones that are consistent with the learned patterns (4). Obtaining (1) from (2) is done at training time, while obtaining (4) from (3) is done during testing. If there is no ground-truth, we use an iterative scheme that alternates between computing trajectories and learning patterns from them.

a person entering a building via a particular door can be expected to head a specific set of rooms or a pedestrian emerging on the street from a shop will often turn left or right to follow the sidewalk. Such patterns are of course not absolutes because people sometimes do the unexpected but they should nevertheless inform the tracking algorithms. We know of no existing technique that imposes this kind of global constraints in globally optimal multi-target tracking.

Our first contribution is therefore an approach to first inferring patterns from ground truth data and then using them to guide the multi-target tracking algorithm. More specifically, we define an objective function that relates behavioral patterns to assigned trajectories. At training time, we use ground truth data to learn patterns that maximize it, as depicted by Fig. 1(1,2). At run time, given these patterns, we connect tracklets produced by another algorithm, so as to maximize the same objective function. Fig. 1(3,4) depicts this process. We will demonstrate that when used in conjunction with several state-of-the-art algorithms, this further increases their already good performance.

Our second contribution is to show we can obtain results almost as good *without* ground-truth data using an alternating scheme that computes trajectories, learns patterns from them, uses them to compute new trajectories, and iterates.

2. Related Work

We briefly review data association and behavior modeling techniques. [76, 47] contain more complete overview of the topics. We also discuss the metrics for MOT evaluation.

2.1. MOT as Data Association

Finding the right trajectories linking the detections, or data association, has been formalized using various models. For real-time performance, data association often relies either on matching locally between existing tracks and new targets [25, 46, 5, 21, 54] or on filtering techniques [57, 67]. The resulting implementations are fast but often perform less well than batch optimization methods, which use a sequence of frames to associate the data optimally over a whole set of frames, rather than greedily in each next frame.

Batch optimization is usually formulated as a shortest path problem [12, 62], network flow problem [84], generic linear programming [34], integer or quadratic programming [45, 18, 73, 65, 23, 82]. A common way to reduce the computational burden is to first group reliable detections into short trajectory fragments known as tracklets and then reason on these tracklets instead of individual detections [35, 69, 48, 43, 9].

However, whether or not tracklets are used, making the optimization problem tractable when looking for a global optimum limits the class of objective functions that can be used. They are usually restricted to functions that can be defined on edges or edge pairs in a graph whose nodes are either individual detections or tracklets. In other words, such objective functions can be used to impose relatively local constraints. To impose global constraints, the objective functions have to involve multiple people and long time spans. They are solved using gradient descent with exploratory jumps [55], inference with a dynamic graphical model [21], or iterative groupings of shorter tracklets into longer trajectories [42, 28, 4]. However, this comes at the cost of losing any guarantee of global optimality.

By contrast, our approach is designed for batch optimization and finding the global optimum, while using an objective function that is rich enough to express the relation between global trajectories and non-linear motion patterns.

2.2. Using Behavioral Models

There have been a number of attempts at incorporating human behavioral models into tracking algorithms to increase their reliability. For example, the approaches of [60, 2] model collision avoidance behavior to improve tracking, the one of [80] uses behavioral model to predict near future target locations, and the one of [63] encodes local velocities into the affinity matrix of tracklets. These approaches boost the performance but only account for very local interactions, instead of global behaviors that influence the *whole* trajectory.

Figure 2. Effect of identity switches on the tracking metrics. The thick lines represent ground-truth trajectories and the thin dotted ones recovered trajectories. The trajectory fragments that count positively are shown in green and those that count negatively in red. The formulas at the top of the figure depict graphically how the **MOTA** and **IDF₁** scores are computed. **Top:** Three ground-truth trajectories, with the bottom two crossing in the middle. The four recovered trajectories feature an identity switch where the two real trajectories intersect, missed detections resulting in a fragmented trajectory and therefore another identity switch at the top, and false detections at the bottom left. When using **MOTA**, the identity switches incur a penalty but only very locally, resulting in a relatively high score. By contrast, **IDF₁** penalizes the recovered trajectories over the *whole* trajectory fragment assigned to the wrong identity, resulting in a much lower score. **Bottom:** The last two thirds of the recovered trajectory are fragmented into individual detections that are not linked. **MOTA** counts each one as an identity switch, resulting in a negative score, while **IDF₁** reports a more intuitive value of 0.3.

Many approaches to inferring various forms of global patterns have been proposed over the years [64, 36, 51, 61, 83, 31, 19, 40, 72]. However, the approaches of [11], [41], and [6] are the only ones we know of that attempt to use these global patterns to guide the tracking. The method of [11] is predicated on the idea that behavioral maps describing a distribution over possible individual movements can be learned and plugged into the tracking algorithm to improve it. However, even though the maps are global, they are only used to constrain the motion locally without enforcing behavioral consistency over the whole trajectory. In [6], an E-M-based algorithm is used to model the scene as a Gaussian mixture that represents the expected size and speed of an object at any given location. While the model can detect global movement anomalies and improve object detection, the motion pattern information is not used to improve the tracking explicitly. In [41], modeling of the optical flow improves the tracking, and helps to detect anomalies, but it relies on the presence of dense crowds, motion flow of which is used for tracking.

2.3. Quantifying Identity Switches

In this paper, we aim to do globally consistent tracking by preventing *identity switches* along reconstructed trajectories, for example when trajectories of different people are merged into one or when a single trajectory is fragmented

Figure 3. (a) Given a set of high-confidence detections $\mathcal{D}$, and a set of allowed transitions $\mathcal{E}$, we seek to find: (b) trajectories of the people, represented by transitions from T ; (c) a set of behavioural patterns P , which define where people behaving in a particular way are likely to be found; an assignment A of each individual detection to a pattern, specifying which pattern did the person in this detection follow.

into many. We therefore need an appropriate metric to gauge the performance of our algorithms.

The set of CLEAR MOT metrics [13] has become a *de-facto* standard for evaluating tracking results. Among these, Multiple Object Tracking Accuracy (MOTA) is the one that is used most often to compare competing approaches. However, it has been pointed out that MOTA does not properly account for identity switches [8, 82, 10], as depicted on the left side of Fig. 2. More adapted metrics have therefore been proposed. For example, IDF_1 is computed by matching trajectories to ground-truth so as to minimize the sum of discrepancies between corresponding ones [66]. Unlike **MOTA**, it penalizes switches over the *whole* trajectory fragments assigned to the wrong identity, as depicted by the right side of Fig. 2. Furthermore, unlike Id-Aware metrics [82, 8], it does not require knowing the true identity of the people being tracked, making it more widely applicable.

In the results section, we report our results in terms of both **MOTA** because it is widely used and IDF_1 to highlight the drop in identity switches our method brings about.

3. Formulation

In this section, we formalize the problem of discovering and using behavioral patterns to impose global constraints on a multi-people tracking algorithm. In the following sections we will use it to estimate trajectories given the patterns and to discover the patterns given ground-truth trajectories.

3.1. Detection Graph

Given a set of high-confidence detections $\mathcal{D} = \{1, \dots, L\}$ in consecutive images of a video sequence, let $\mathcal{V} = \mathcal{D} \cup \{I, O\}$, where I and O denote possible trajectory start and end points and each node $v \in \mathcal{D}$ is associated with a set of features that encode location, appearance, or other important properties of a detection. Let $\mathcal{E} \subset \mathcal{V}^2$ be the set of possible transitions between the detections. $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ can then be treated as a *detection graph* of which the desired trajectories are subgraphs. As shown by Fig. 3, let

- $T \subset \mathcal{E}$ be a set of edges defining people’s trajectories.

- P be a set of patterns, each defining an area where people behaving in a specific way are likely to be found, plus an empty pattern $\emptyset$ used to describe unusual behaviors. Formally speaking, patterns are functions that associate to a trajectory with an arbitrary number of edges a score that denotes how likely it is to correspond to that specific pattern, as shown in Section 3.3.
- A be a set of assignments of individual detections in $\mathcal{D}$ into patterns, that is, a mapping $A : \mathcal{D} \rightarrow \{1, \dots, N_p\}$, where N_p is the total number of patterns.

Each trajectory $t \in T$ must go through detections via allowable transitions, begin at I , and end at O . Here we abuse the notation $t \in T$ to show that all edges $(I, t_1), (t_1, t_2), \dots, (t_{|t|}, O)$ from trajectory $t = (t_1, \dots, t_{|t|})$ belong to T . Furthermore, since we only consider high-confidence detections, each one must belong to **exactly** one trajectory. In practice, this means that potential false positives end up being assigned to the empty behavior $\emptyset$ and can be removed as a post-processing step. Whether to do this or not is governed by a binary indicator R_e selected during training process. In other words, the edges in T must be such that for each detection there is exactly one selected edge coming in and one going out, which we can write as

$$\forall j \in \mathcal{D}, \exists! i \in \mathcal{V}, k \in \mathcal{V} : (i, j) \in T \wedge (j, k) \in T. \quad (1)$$

Since all detections that are grouped into the same trajectory T must be assigned to the same pattern, we must have

$$\forall (i, j) \in T : (i \in \mathcal{D} \wedge j \in \mathcal{D}) \Rightarrow A(i) = A(j). \quad (2)$$

In our implementation, each pattern $p \in P \setminus \emptyset$ is defined by a trajectory that serves as a centerline and a width, as depicted by Fig. 3(c) and A. However, the optimization schemes we will describe in Sections 4.1 and 4.2 do not depend on this specific representation and any other convenient one could have been used instead.

3.2. Building the Graph

To build the graph we use as input the output of another algorithm that produces trajectories that we want to

improve. We take the set of detections along these trajectories to be our high-confidence detections $\mathcal{D}$ and therefore the nodes of our graph. We take the edges $\mathcal{E}$ to be pairs of nodes that are either *i*) consecutive in the original trajectories, *ii*) within ground plane distance D_1 of each other in successive frames, *iii*) the endings and beginnings of input trajectories within distance D_2 and within D_t frames, *iv*) or whose first node is I or second node is O . In other words, to allow the creation of new trajectories and recover from identity switches, fragmentation, and incorrectly merged trajectories, we introduce edges not only for consecutive points in existing ones but also to connect neighboring trajectories.

3.3. Objective Function

Our goal is to find the most likely trajectories formed by transitions in T^* , patterns P^* , and mapping linking one to the other A^* given the image information and any *a priori* knowledge we have. In particular, given a set of patterns P^* , we will look for the best set of trajectories that match these patterns. Conversely, given a set of known trajectories T^* , we will learn a set of patterns, as discussed in Section 4.

To formulate these searches in terms of an optimization problem, we introduce an objective function $C(T, P, A)$ that reflects how likely it is to observe the objects moving along the trajectories defined by T , each one corresponding to a pattern from $P = \{p_1 \dots, p_{N_p}\}$ given the assignment A . Ideally, C should be the proportion of trajectories that correctly follow the assigned patterns. To compute it in practice, we take our inspiration from the **MOTA** and **IDF₁** scores described in Section 2.3. They are written in terms of ratios of the lengths of trajectory fragments that follow the ground truth to total trajectory lengths. We therefore take our objective function to be

$$C(T, P, A) = \frac{\sum_{t \in T} M(t, p_{A(t_1)})}{\sum_{t \in T} N(t, p_{A(t_1)})}, \quad (3)$$

$$N(t, p) = n(I, t_1, p) + n(t_{|t|}, O, p) + \sum_{1 \leq j \leq |t|-1} n(t_j, t_{j+1}, p),$$

$$M(t, p) = m(I, t_1, p) + m(t_{|t|}, O, p) + \sum_{1 \leq j \leq |t|-1} m(t_j, t_{j+1}, p),$$

where $n(i, j, p)$ is the sum of the total length of edge (i, j) and of the length of the corresponding pattern centerline, while $m(i, j, p)$ is the sum of lengths of aligned parts of the pattern and the edge. Fig. A illustrates this computation and we give the mathematical definitions of m and n in the supplementary material. As a result, $N(t, p)$ is the sum of the lengths of trajectory and assigned pattern while $M(t, p)$ measures the length of parts of trajectory and pattern that are aligned with each other. Note that the definition of Eq. (11) is very close to that of the metric **IDF₁** introduced in Sec. 2.3. It is largest when each person follows a

Figure 4. For a pattern p defined by centerline c_p , shown as a thick black line, with width w_p , and an edge (i, j) , we compute functions $n(i, j, p)$ and $m(i, j, p)$ introduced in Section 3.3 and shown in green and blue, respectively, as follows: $n(i, j, p)$ is the total length of the edge and the corresponding length of the pattern centerline, measured between the points p_i and p_j , which are the points on the centerline closest to i and j . If both i and j are within the pattern width w_p from the centerline, we take $m(i, j, p)$ to be the sum of two terms: the length in the pattern along the edge, that is, the distance between p_i and p_j , plus the length in the edge along the pattern, that is, the length of the projection of (p_i, p_j) onto the line connecting i and j . Otherwise $m(i, j, p) = 0$ to penalize the deviation from the pattern.

single pattern for as long as possible. This penalizes identity switches because the trajectories that are erroneously merged, fragmented, or jump between people are unlikely to follow any of such pattern.

In Eq. (11), we did not explicitly account for the fact that the first vertex i of some edges can be the special entrance vertex, which is not assigned to any behavior. When this happens we simply use the pattern assigned to the second vertex j . From now on, we will replace $A(i)$ by $A(i, j)$ to denote this behavior. We also adapt the definitions of m and n accordingly to properly handle those special edges.

4. Computing Trajectories and Patterns

In this section, we describe how we use the objective function C of Eq. (11) to compute trajectories given patterns and patterns given trajectories. The resulting procedures will be the building blocks of our complete MOT algorithm, as described in Section 5.

4.1. Trajectories

Let us assume that we are given a precomputed set of patterns P^* , then we look for trajectories and corresponding assignment as

$$T^*, A^* = \arg \max_{T, A} C(T, P^*, A). \quad (4)$$

To solve this problem, we treat the motion of people through the detection graph $\mathcal{G}$ introduced in Section 3.1 as a flow. Let $o_{ij}^p \in \{0, 1\}$ be the number of people transitioning from

node i to j in a trajectory T assigned to pattern $p \in P^*$. It relates to P^* and T as follows:

$$o_{ij}^p = \mathbb{I}(((i, j) \in T) \wedge (P_{A(i,j)}^* = p)) . \quad (5)$$

Using these new binary variables, we reformulate constraints (1) and (2) as

$$\begin{aligned} \forall i \in \mathcal{D} \cup \mathcal{O} \quad \sum_{(i,j) \in \mathcal{E}, p \in P^*} o_{ij}^p &= 1 , \\ \forall j \in \mathcal{D}, p \in P^* \quad \sum_{(i,j) \in \mathcal{E}} o_{ij}^p &= \sum_{(j,k) \in \mathcal{E}} o_{jk}^p . \end{aligned} \quad (6)$$

This lets us rewrite our cost function as

$$C(T, P^*, A) = \frac{\sum_{(i,j) \in T, p \in P^*} m(i, j, p) o_{ij}^p}{\sum_{(i,j) \in T, p \in P^*} n(i, j, p) o_{ij}^p} , \quad (7)$$

which we maximize with respect to the flow variables o_{ij}^p subject to the two constraints of (14). This is an integer-fractional program, which could be transformed into a Linear Program [20]. However, solving it would produce non-integer values that would need to be rounded. To avoid this we propose a scheme based on the following observation: Maximizing $\frac{a(x)}{b(x)}$ with respect to x when $b(x)$ is always positive can be achieved by finding the largest α such that an x satisfying $a(x) - \alpha b(x) \geq 0$ can be found. Furthermore, α can be found by binary search. We therefore take a to be the numerator or Eq. (15), b its denominator, and x the vector of o_{ij}^p variables. In practice, given a specific value of α , we do this by running a Integer Linear Program solver [30] until it finds a feasible solution. When α reaches its maximum possible value, that feasible solution is also the optimal one. We provide more details in the supplementary material and a version of our code is publicly available¹.

4.2. Patterns

In the previous section, we assumed the patterns known and used them to compute trajectories. Here, we reverse the roles. Let us assume we are given a set of trajectories T^* . We learn the patterns and corresponding assignments as

$$P^*, A^* = \arg \max_{P, A} C(T^*, P, A) , \quad (8)$$

$$\text{subject to} \quad P \subset \mathcal{P}, |P| \leq \alpha_p, \sum_{p \in P} M(p) \leq \alpha_c ,$$

where α_c, α_p are thresholds and $M : P \rightarrow \mathbb{R}^+$. The purpose of the additional constraints is to limit both the number of patterns being used and their spatial extent to prevent over-fitting. In our implementation, we take $M(p) = l_p w_p$, where l_p is the length of the pattern centerline and w_p is

its width. $\mathcal{P}$ is a set of all admissible patterns, which we construct by combining all possible ground-truth trajectories as centerlines with each width from a predefined set of possible pattern widths.

To solve the problem of Eq. (8), we look for an assignment between our known ground truth trajectories T^* and all possible patterns $\mathcal{P}$ and retain only patterns associated to at least one trajectory. To this end, we introduce auxiliary variables a_{tp} describing the assignment $A^* : T^* \rightarrow \mathcal{P}$, and variables b_p denoting if at least one trajectory is matched to pattern p . Formally, this can be written as

$$\begin{aligned} a_{tp} &\in \{0, 1\} , \forall t \in T^*, p \in \mathcal{P} , \\ b_p &\in \{0, 1\} , \forall p \in \mathcal{P} , \\ \sum_{p \in \mathcal{P}} a_{tp} &= 1 , \forall t \in T^* , \\ a_{tp} &\leq b_p , \forall t \in T^*, p \in \mathcal{P} . \end{aligned} \quad (9)$$

Given that C is defined as the fraction from Eq. (11), we use the optimization scheme similar to one described in Sec. 4.1, where we do binary search to find the optimal value of α such that there exists a feasible solution for constraints of (16) and the following:

$$\begin{aligned} \sum_{t \in T^*} \sum_{p \in \mathcal{P}} (m(t, p) - \alpha n(t, p)) a_{tp} &\geq 0 , \\ \sum_{p \in \mathcal{P}} b_p &\leq \alpha_p , \\ \sum_{p \in \mathcal{P}} b_p M(p) &\leq \alpha_c . \end{aligned} \quad (10)$$

5. Non-Markovian Multiple Object Tracking

Given that we can learn patterns from a set trajectories, we can now enforce long-range behavioral patterns when linking a set of detections in two different manners. This is in contrast to traditional approaches enforcing local smoothness constraints, Markovian in their essence.

If annotated ground-truth trajectories T^* are available, we use them to learn the patterns as described in Sec. 4.2. Then, at test time, we use the linking procedure of Sec. 4.1.

If no such training data is available, we can run an E-M-style procedure, very similar to the Baum-Welch algorithm [33] for HMMs: we start from a set of trajectories computed using a standard algorithm, and from there, use trajectories to compute a set of patterns, then use the set of patterns to compute trajectories, and iterate. We will see that, in practice, this yields results that are almost indistinguishable in terms of accuracy but much slower because we have to run through many iterations.

More specifically, each iteration of our unsupervised approach involves *i*) finding a set of patterns P^i given a set of trajectories T^{i-1} , *ii*) finding a set of trajectories T^i given a set of patterns P^i , as described in Sec. 4.2 and 4.1.

¹https://github.com/maksay/ptrack_cpp

In practice, for a given α_c , this scheme converges after a few iterations. Since α_c is unknown *a priori*, we start with a small α_c , perform 5 iterations, increase α_c , and repeat until reaching a predefined number of patterns. To select the best trajectories without reference to ground truth, we define

$$\widetilde{\mathbf{IDF}}_1(T^i) = \frac{1}{2}(C(T_1^i, P_2^i, A_{T_1^i \rightarrow P_2^i}) + C(T_2^i, P_1^i, A_{T_2^i \rightarrow P_1^i})),$$

where T_1^i and T_2^i are time-disjoint subsets of T^i , P_1^i and P_2^i are patterns learned from T_1^i and T_2^i . $A_{T_1^i \rightarrow P_2^i}$ and $A_{T_2^i \rightarrow P_1^i}$ are such assignments of trajectories to the patterns learned on another subset that maximize $\widetilde{\mathbf{IDF}}_1(T^i)$.

In effect, $\widetilde{\mathbf{IDF}}_1$ is a valid proxy for $\mathbf{IDF}_1$ due to the many similarities between our cost function and $\mathbf{IDF}_1$ outlined in Sec. 3.3. In the end, we select the trajectories that maximize $\widetilde{\mathbf{IDF}}_1$. Using such cross-validation to pick the best solution in E-M models is justified in [1].

6. Results

In this section, we demonstrate the effectiveness of our approach on several datasets, using both simple and sophisticated approaches to produce the initial trajectories. (Recall from Section 3.2 that we build our detection graphs from the output of another tracking algorithm.) In the remainder of this section, we first describe the datasets and the tracking algorithms we rely on to build the initial graphs. We then discuss the evaluation metrics and the experimental protocol. Finally, we present our experimental results.

6.1. Datasets

We use the four datasets listed in Tab. 1. They are:

Town. A sequence from the 2DMOT2015 benchmark featuring a lively Zurich street where people walk in different directions.

ETH and Hotel. Sequences from the BIWI Walking Pedestrians dataset [59] that were originally used to model social behavior. In these datasets, using image and appearance information for tracking is difficult, due to recordings with a top view camera and low visibility in **ETH** dataset.

Station. A one hour-long recording of Grand Central station in New York with several thousands of annotated pedestrian tracks [85]. It was originally used for trajectory prediction for moving crowds.

Name	Annotated length, s	FPS	Trajectories
Town	180	2.5	246
ETH	360	4.16	352
Hotel	390	2.5	175
Station	3900	1.25	12362

Table 1. Dataset statistics. The number of trajectories is calculated as a total sum of number of trajectories in each test set on which we evaluated. All test sets were approximately 1min long.

These four datasets share the following characteristics: *i)* They feature real-life behaviors as opposed to random and unrealistic motions acquired in a lab setting; *ii)* The frame rate is at most 5 frames per second, which is realistic for outdoor surveillance setups but makes tracking more difficult; *iii)* They are all single-camera but the shape of the ground surface can be estimated from the bottom of the bounding boxes, which makes it possible to reason in a simulated top view as we do. In other words, they are well suited to test our approach in challenging conditions.

6.2. Baselines

As discussed in Section 3.2, we use as input to our system trajectories produced by recent MOT algorithms, some of which exploit image and appearance information and some of which do not. In Section 6.4, we will show that imposing our pattern constraints systematically results in an improvement over these baselines, which we list below.

MDP [77] formulates MOT as decision making in a Markov Decision Process (MDP) framework. Learning to associate data correctly is equivalent to learning an MDP policy and is done through reinforcement learning. At the time of writing, this was the highest-ranking approach (in terms of **MOTA**) with publicly available implementation on the 2DMOT2015 [44] benchmark.

SORT [14] is a real-time Kalman filter-based MOT approach. At the time of writing, this was the second highest-ranking approach on 2DMOT2015 benchmark.

RNN [54] is a recent attempt at using recurrent neural networks to predict the motion of multiple people and perform MOT in real time. It does not require any appearance information, only the coordinates of the bounding boxes. In the presented results, this approach outperforms all the other methods that do not use image and appearance information.

KSP [12] is a simple approach to MOT that formulates the MOT problem as finding K Shortest Paths in the detection graph, without using image or appearance information.

2DMOT2015 Top Scoring Methods [21, 68, 74, 37, 49, 39, 75, 81] to which we will refer by the name that appears in the official scoreboard [44]. This will allow us to show that our approach is widely applicable.

Top scoring MOT methods from the 2DMOT2015 benchmark on the **Town** dataset rely on a people detector that is not always publicly available. We therefore used their output to build the detection graph, and report their results only on **Town** dataset. For all others, the available code accepts a set of initial detections as input. To compute them, we obtained background subtraction by subtracting the median image. We used the publicly available POM algorithm of [27] on the resulting binary maps to produce probabilities of presence in various ground locations and we kept those

for which the probability was greater than 0.5. This proved effective on all our datasets. For comparison purposes, we also tried using SVMs trained on HOG features [22] and deformable part models [26]. While their performance was roughly similarly to that of POM on **Town**, it proved much worse when the people are far away or seen from above.

6.3. Experimental Protocol

On the **Station** dataset, which is long and features more than 100 people per minute, we tested on 1-minute subsequences and trained on a non-overlapping 5-minute subsequence. We also limited the optimization time for solving Eq. 15 to 10 minutes per iteration of binary search. On all other datasets, we tested on 1-minute subsequences, trained on the remainder and did not limit optimization time. To prevent any interaction between the training and testing data, we removed from the ground truth training data all incomplete trajectories to guarantee no overlap with the testing data. The remaining trajectories were used to learn the patterns of Section 4.2 and choose the values of the parameters D_1 , D_2 , D_t controlling the construction of edges in the tracking graph, R_e controlling whether to discard trajectories assigned to no pattern, and α_c , α_p regularizing number and width of patterns, introduced in Sections 4.1 and 4.2. It is done by performing a grid search and selecting values that yield the best possible IDF_1 score in cross-validation. To keep the search tractable, we always started from a default set of values $D_1 = 2m$, $D_2 = 4m$, $D_t = 2s$, $\mathbb{I}(R_e) = 1$, $\alpha_p = 5$, $\alpha_c = 0.3\alpha_p \times \text{area}$ for tracking, and explored neighboring values in the 6D grid. We do the same exploration when running iterative scheme to select the optimal value of IDF_1 in unsupervised setup described in Section 5.

For the sake of fairness, we trained the trainable baselines of Section 6.2, that is **MDP** and **RNN**, similarly and using the same data. However, for **RNN** we obtained better results using the provided model, pre-trained on the 2DMOT2015 training data, and we report these results.

We combined the results from all test segments to obtain the overall metrics on each dataset. Since for some approaches we only had results in the form of bounding boxes and had to estimate the ground plane location based on that, this often resulted in large errors further away from the camera. For this reason, we evaluated **MOTA** and IDF_1 assuming that a match happens when the reported location is at most at 3 meters from the ground truth location. We also provide results for the traditional distance of 1 meter in supplementary material and they are similar in terms of method ordering. For the **Station** dataset, we did not have the information about the true size of the floor area (we estimated the homography between the image and ground plane) which is why we used as distance 0.1 of the size of the tracking area.

Figure 5. IDF_1 and **MOTA** scores for various methods on the **Town** dataset. Our approach almost always improves IDF_1 . We provide the actual numbers in the supplementary material.

Approach	ΔIDF_1	ΔIDF_1^u	$\Delta MOTA$	$\Delta MOTA^u$
KSP	0.16	0.15	-0.01	-0.01
MDP	0.05	0.02	0.03	-0.01
RNN	0.04	0.03	0.00	-0.02
SORT	0.04	0.02	0.06	0.00

Table 2. Improvement in IDF_1 and **MOTA** metrics delivered by our approach, averaged over all datasets. The 2nd and 4th columns correspond to the supervised case, the 3rd and 5th to the unsupervised one.

6.4. Experiments

We first show that our approach consistently improves the output of all the baselines using initial people detections obtained as described at the end of Section 6.2. Then, to gauge what our approach could achieve given perfect detections, we perform this comparison again but using ground truth detections instead. Finally, we discuss the computational complexity of our approach.

Improving Baseline Results. In terms of the IDF_1 metric, as can be seen in Fig. 5, our Supervised method improves most of the tracking results except one that remains Unchanged on **Town**. The same can be said of the unsupervised version of Our method except for one result that it degrades by 0.01. In Tab. 2, we average these results over all datasets and can observe a marked average improvement for all methods we could run on all Datasets both in the supervised and unsupervised cases. As could be expected, The improvements in terms of **MOTA** are less clear since our method modifies The set of input detections minimally. Fig. 6 depicts some of The results and we provide detailed breakdowns in the supplementary material.

Evaluation on Ground Truth Detections. For all baselines that Accept a list of detections as input, and for which

Figure 6. Examples of learned patterns, denoted by their centerline in white, with some erroneous trajectories found by various baselines in red. White bounding boxes for people following the trajectories are shown. Improved trajectories found by our approach in green. We also show the pattern widths (area in blue), to show that the trajectory we found is assigned to a particular pattern. (a) **Town** dataset, **EAMTT** [68] merges several trajectories going in opposite directions, but (b) correct pattern assignment helps to fix that; (c) Using only affinity information, **KSP** is prone to multiple identity switches; (d) Our approach recovers several trajectories correctly, but merges trajectories of two different people in the lower left corner going in the same general direction; (e) **ETH** dataset, due to low visibility using flow and feature point tracking is hard, and **MDP** fragments a single trajectory into two, but our approach fixes that (not shown)

Figure 7. Example of unsupervised optimization. (a) Four people are tracked using **KSP**. Trajectories are shown as solid black lines, bounding boxes are white. (b) Tracks continue, featuring several identity switches. (c) First step of the alternating scheme finds a single pattern, in white, that explains as many trajectories as possible, it is the leftmost trajectory. (d) Given this pattern, next step is to fit trajectories to it. Trajectories in blue are the ones assigned to this pattern, trajectories in red are assigned to no pattern. One identity switch is fixed. (e) After several iterations, we look for the best two patterns. Rightmost trajectory is picked as the second pattern. Fitting trajectories to the best two patterns allows to fix the remaining fragmented trajectory. Trajectories assigned to the second pattern in green.

the code is available, we Reran the same experiment using the ground truth detections instead of those Computed by the POM algorithm [27] as before. This is a way to Evaluate the performance of the linking procedure independently of that of the Detections. It reflects the theoretical maximum that can be reached by all the Approaches we compare, including our own. From Tables 3 And 4, we observe that our approach performs very well in Such setting.

Approach: Dataset:	MDP	RNN	SORT	KSP	OUR
Town	0.87	0.65	0.88	0.55	0.93
ETH	0.89	0.65	0.93	0.59	0.92
Hotel	0.85	0.70	0.88	0.60	0.94
Station	0.68	0.40	0.72	0.45	0.70

Table 3. **IDF₁** evaluation results using ground truth detections.

Approach: Dataset:	MDP	RNN	SORT	KSP	OUR
Town	0.87	0.85	0.90	0.87	0.98
ETH	0.85	0.73	0.85	0.70	0.94
Hotel	0.84	0.78	0.82	0.74	0.97
Station	0.75	0.68	0.70	0.80	0.77

Table 4. **MOTA** evaluation results using ground detections.

Computation time The number of variables in our optimization problem grows linearly with the Length of the

Dataset	Town	ETH	Hotel	Station	Station
Frames	150	227	268	75	75
Trajectories	85	67	47	100	193
Patterns	7	5	4	26	26
Detections	2487	894	1019	1960	3724
Variables	70k	17k	18k	191k	450k
Time, s	26	4	4	160	>3600

Table 5. Optimization problem size and run time of our approach for processing a typical one mn batch from each dataset.

batch and number of patterns, and superlinearly with the number Of people per frame (as the number of possible connections between people). As Shown by Tab. 5, for not too crowded datasets without large number Of patterns our approach is able to process a minute of input frames under A minute. Pattern fitting scales quadratically with the number of given Ground-truth trajectories and runs in less than 10 minutes for all datasets Except **Station**. More details can be found in supplementary materials.

7. Conclusion

In this work we have proposed an approach to tracking multiple people under global behavioral constraints. It lets us learn motion patterns given ground truth trajectories, use these patterns to guide the tracking, and improve upon a wide range of state-of-the-art approaches. It also extends naturally to the unsupervised case without ground truth.

Our optimization scheme is generic and allows for a

wide range of definitions for the patterns, beyond the ones we have used here. In the future, we plan to work with more complex patterns that models human behavior better, account for appearance, and handle correlations between people's behavior.

References

- [1] Electronic Statistics Textbook. Finding the Right Number of Clusters in k-Means and EM Clustering: v-Fold Cross-Validation. Technical report, 2010. **6**
- [2] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-fei, and S. Savarese. Social LSTM: Human Trajectory Prediction in Crowded Spaces. In *Conference on Computer Vision and Pattern Recognition*, 2014. **2**
- [3] M. Andriluka, S. Roth, and B. Schiele. People-Tracking-By-Detection and People-Detection-By-Tracking. In *Conference on Computer Vision and Pattern Recognition*, June 2008. **1**
- [4] A. Andriyenko, K. Schindler, and S. Roth. Discrete-Continuous Optimization for Multi-Target Tracking. In *Conference on Computer Vision and Pattern Recognition*, pages 1926–1933, June 2012. **2**
- [5] S.-H. Bae and K.-J. Yoon. Robust Online Multi-Object Tracking Based on Tracklet Confidence and Online Discriminative Appearance Learning. In *Conference on Computer Vision and Pattern Recognition*, 2014. **2**
- [6] A. Basharat, A. Gritai, and M. Shah. Learning Object Motion Patterns for Anomaly Detection and Improved Object Detection. In *Conference on Computer Vision and Pattern Recognition*, 2008. **2**
- [7] R. Benenson, O. Mohamed, J. Hosang, and B. Schiele. Ten Years of Pedestrian Detection, What Have We Learned? In *European Conference on Computer Vision*, 2014. **1**
- [8] H. BenShitrit, J. Berclaz, F. Fleuret, and P. Fua. Tracking Multiple People Under Global Appearance Constraints. In *International Conference on Computer Vision*, 2011. **3**
- [9] H. BenShitrit, J. Berclaz, F. Fleuret, and P. Fua. Multi-Commodity Network Flow for Tracking Multiple People. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(8):1614–1627, 2014. **2**
- [10] J. Bento. A metric for sets of trajectories that is practical and mathematically consistent. *arXiv preprint arXiv:1601.03094*, 2016. **3**
- [11] J. Berclaz, F. Fleuret, and P. Fua. Multi-Camera Tracking and Atypical Motion Detection with Behavioral Maps. In *European Conference on Computer Vision*, October 2008. **2**
- [12] J. Berclaz, F. Fleuret, E. Türetken, and P. Fua. Multiple Object Tracking Using K-Shortest Paths Optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(11):1806–1819, 2011. **2, 6**
- [13] K. Bernardin and R. Stiefelhagen. Evaluating Multiple Object Tracking Performance: the Clear Mot Metrics. *EURASIP Journal on Image and Video Processing*, 2008, 2008. **3, 18**
- [14] A. Bewley, Z. Ge, L. Ott, F. Ramos, and B. Upcroft. Simple online and realtime tracking. In *ICIP*, 2016. **6**
- [15] J. Black, T. Ellis, and P. Rosin. Multi-View Image Surveillance and Tracking. In *IEEE Workshop on Motion and Video Computing*, 2002. **1**
- [16] S. Blackman. *Multiple-Target Tracking with Radar Applications*. Artech House, 1986. **1**
- [17] M. Breitenstein, F. Reichlin, B. Leibe, E. Koller-Meier, and L. Van Gool. Online Multi-Person Tracking-By-Detection from a Single Uncalibrated Camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2010. **1**
- [18] W. Brendel, M. Amer, and S. Todorovic. Multiobject Tracking as Maximum Weight Independent Set. In *Conference on Computer Vision and Pattern Recognition*, 2011. **2**
- [19] S. Calderara, U. Heinemann, A. Prati, R. Cucchiara, and N. Tishby. Detecting anomalies in people's trajectories using spectral graph analysis. *Computer Vision and Image Understanding*, 2011. **2**
- [20] A. Charnes and W. Cooper. Programming with linear fractional functionals. *Naval Research logistics quarterly*, 1962. **5**
- [21] W. Choi. Near-Online Multi-Target Tracking with Aggregated Local Flow Descriptor. In *International Conference on Computer Vision*, 2015. **1, 2, 6**
- [22] N. Dalal and B. Triggs. Histograms of Oriented Gradients for Human Detection. In *Conference on Computer Vision and Pattern Recognition*, 2005. **7**
- [23] A. Dehghan, S. M. Assari, and M. Shah. Gmmcp Tracker: Globally Optimal Generalized Maximum Multi Clique Problem for Multiple Object Tracking. In *Conference on Computer Vision and Pattern Recognition*, 2015. **2**
- [24] P. Dollár, C. Wojek, B. Schiele, and P. Perona. Pedestrian Detection: An Evaluation of the State of the Art. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(4):743–761, 2012. **1**
- [25] L. Fagot-Bouquet, R. Audigier, Y. Dhome, and F. Lerasle. Improving Multi-Frame Data Association with Sparse Representations for Robust Near-Online Multi-Object Tracking. In *European Conference on Computer Vision*, pages 774–790, October 2016. **2**
- [26] P. Felzenszwalb, R. Girshick, D. McAllester, and D. Ramanan. Object Detection with Discriminatively Trained Part Based Models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2010. **7**
- [27] F. Fleuret, J. Berclaz, R. Lengagne, and P. Fua. Multi-Camera People Tracking with a Probabilistic Occupancy Map. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):267–282, February 2008. **6, 8**
- [28] K. Fragkiadaki, W. Zhang, G. Zhang, and J. Shi. Two-Granularity Tracking: Mediating Trajectory and Detection Graphs for Tracking Under Occlusions. In *European Conference on Computer Vision*, 2012. **2**
- [29] J. Giebel, D. Gavrilu, and C. Schnorr. A Bayesian Framework for Multi-Cue 3D Object Tracking. In *European Conference on Computer Vision*, 2004. **1**
- [30] Gurobi. Gurobi Optimizer, 2012. <http://www.gurobi.com/>. **5**
- [31] W. Hu, T. Tan, L. Wang, and S. Maybank. A survey on visual surveillance of object motion and behaviors. *IEEE Transactions on Systems, Man, and Cybernetics*, 2004. **2**

- [32] S. Iwase and H. Saito. Parallel Tracking of All Soccer Players by Integrating Detected Positions in Multiple View Images. In *International Conference on Pattern Recognition*, pages 751–754, August 2004. 1
- [33] F. Jelinek, L. Bahl, and R. Mercer. Design of a linguistic statistical decoder for the recognition of continuous speech. *IEEE Transactions on Information Theory*, 1975. 5
- [34] H. Jiang, S. Fels, and J. Little. A Linear Programming Approach for Multiple Object Tracking. In *Conference on Computer Vision and Pattern Recognition*, pages 1–8, June 2007. 2
- [35] S. W. Joo and R. Chellappa. A Multiple-Hypothesis Approach for Multiobject Visual Tracking. *IEEE Transactions on Image Processing*, 2007. 2
- [36] M. Kalayeh, S. Mussmann, A. Petrakova, N. Lobo, and M. Shah. Understanding Trajectory Behavior: A Motion Pattern Approach. *arXiv preprint arXiv:1501.00614*, 2015. 2
- [37] M. Keuper, S. Tang, Y. Zhongjie, B. Andres, T. Brox, and B. Schiele. A Multi-cut Formulation for Joint Segmentation and Tracking of Multiple Objects. *arXiv preprint arXiv:1607.06317*, 2016. 6
- [38] Z. Khan, T. Balch, and F. Dellaert. MCMC-Based Particle Filtering for Tracking a Variable Number of Interacting Targets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(11):1805–1918, 2005. 1
- [39] C. Kim, F. Li, A. Ciptadi, and J. Rehg. Multiple hypothesis tracking revisited. In *ICCV*, 2015. 6
- [40] J. Kim and K. Grauman. Observe locally, infer globally: a space-time MRF for detecting abnormal activities with incremental updates. In *Conference on Computer Vision and Pattern Recognition*. 2
- [41] L. Kratz and K. Nishino. Going with the flow: pedestrian efficiency in crowded scenes. In *European Conference on Computer Vision*, 2012. 2
- [42] C.-H. Kuo, C. Huang, and R. Nevatia. Multi-Target Tracking by On-Line Learned Discriminative Appearance Models. In *Conference on Computer Vision and Pattern Recognition*, 2010. 2
- [43] C.-H. Kuo and R. Nevatia. How Does Person Identity Recognition Help Multi-Person Tracking? In *Conference on Computer Vision and Pattern Recognition*, 2011. 2
- [44] L. Leal-taixe, A. Milan, I. Reid, S. Roth, and K. Schindler. Motchallenge 2015: Towards a Benchmark for Multi-Target Tracking. In *arXiv Preprint*, 2015. 6
- [45] B. Leibe, K. Schindler, and L. Van Gool. Coupled Detection and Trajectory Estimation for Multi-Object Tracking. In *International Conference on Computer Vision*, October 2007. 2
- [46] P. Lenz, A. Geiger, and R. Urtasun. FollowMe: Efficient Online Min-Cost Flow Tracking with Bounded Memory and Computation. In *International Conference on Computer Vision*, pages 4364–4372, December 2015. 2
- [47] X. Li, W. Hu, C. Shen, Z. Zhang, A. Dick, and A. Hengel. A Survey of Appearance Models in Visual Object Tracking. *ACM Transactions on Intelligent Systems and Technology*, 2013. 2
- [48] Y. Li, C. Huang, and R. Nevatia. Learning to Associate: Hybridboosted Multi-Target Tracker for Crowded Scene. In *Conference on Computer Vision and Pattern Recognition*, June 2009. 2
- [49] M. M. Yang and Y. Jia. Temporal Dynamic Appearance Modeling for Online Multi-Person Tracking. *Computer Vision and Image Understanding*, 2016. 6
- [50] D. R. Magee. Tracking Multiple Vehicles Using Foreground, Background and Motion Models. *Image and Vision Computing*, 22(2):143–155, February 2004. 1
- [51] V. Mahadevan, W. Li, V. Bhalodia, and N. Vasconcelos. Anomaly detection in crowded scenes. In *Conference on Computer Vision and Pattern Recognition*, 2010. 2
- [52] T. Mauthner, M. Donoser, and H. Bischof. Robust Tracking of Spatial Related Components. In *International Conference on Pattern Recognition*, 2008. 1
- [53] A. Milan, L. Leal-taixe, I. Reid, S. Roth, and K. Schindler. MOT16: A Benchmark for Multi-Object Tracking. *arXiv preprint arXiv:1603.00831*, 2016. 1
- [54] A. Milan, S. H. Rezatofighi, A. Dick, K. Schindler, and I. Reid. Online Multi-Target Tracking Using Recurrent Neural Networks. *arXiv preprint arXiv:1604.03635*, 2016. 2, 6
- [55] A. Milan, S. Roth, and K. Schindler. Continuous Energy Minimization for Multitarget Tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36:58–72, 2014. 2
- [56] A. Mittal and L. Davis. M2Tracker: A Multi-View Approach to Segmenting and Tracking People in a Cluttered Scene. *International Journal of Computer Vision*, 51(3):189–203, 2003. 1
- [57] S. Oh, S. Russell, and S. Sastry. Markov Chain Monte Carlo Data Association for Multi-Target Tracking. *IEEE Transactions on Automatic Control*, 2009. 2
- [58] K. Okuma, A. Taleghani, N. de Freitas, J. Little, and D. Lowe. A Boosted Particle Filter: Multitarget Detection and Tracking. In *European Conference on Computer Vision*, May 2004. 1
- [59] S. Pellegrini, A. Ess, K. Schindler, and L. Van Gool. You’ll Never Walk Alone: Modeling Social Behavior for Multi-Target Tracking. In *International Conference on Computer Vision*, 2009. 6
- [60] S. Pellegrini, A. Ess, and L. Van Gool. Improving Data Association by Joint Modeling of Pedestrian Trajectories and Groupings. In *European Conference on Computer Vision*, 2010. 2
- [61] C. Piciarelli, G. L. Foresti, and L. Snidaro. Trajectory clustering and its applications for video surveillance. In *IEEE Conference on Advanced Video and Signal Based Surveillance*, 2005. 2
- [62] H. Pirsiavash, D. Ramanan, and C. Fowlkes. Globally-Optimal Greedy Algorithms for Tracking a Variable Number of Objects. In *Conference on Computer Vision and Pattern Recognition*, pages 1201–1208, June 2011. 2
- [63] Z. Qin and C. Shelton. Improving Multi-Target Tracking via Social Grouping. In *Conference on Computer Vision and Pattern Recognition*, pages 1972–1978, June 2012. 2

- [64] M. Ravanbakhsh, M. Nabi, H. Mousavi, E. Sangineto, and N. Sebe. Plug-and-Play CNN for Crowd Motion Analysis: An Application in Abnormal Event Detection. *arXiv preprint arXiv:1610.00307*, 2016. 2
- [65] S. H. Rezatofghi, A. Milan, Z. Zhang, Q. Shi, A. Dick, and I. Reid. Joint Probabilistic Data Association Revisited. In *International Conference on Computer Vision*, 2015. 2
- [66] E. Ristani, F. Solera, R. Zou, R. Cucchiara, and C. Tomasi. Performance Measures and a Data Set for Multi-Target, Multi-Camera Tracking. *arXiv preprint arXiv:1609.01775*, 2016. 3, 18
- [67] M. Rodriguez, I. Laptev, J. Sivic, and J. Audibert. Density-Aware Person Detection and Tracking in Crowds. In *International Conference on Computer Vision*, pages 2423–2430, 2011. 2
- [68] R. Sanchez-Matilla, F. Poiesi, and A. Cavallaro. Online multi-target tracking with strong and weak detections. In *BMMT ECCVw*, 2016. 6, 8
- [69] V. K. Singh, B. Wu, and R. Nevatia. Pedestrian Tracking by Associating Tracklets Using Detection Residuals. *IEEE Workshop on Motion and Video Computing*, pages 1–8, 2008. 2
- [70] K. Smith, D. Gatica-Perez, and J.-M. Odobez. Using Particles to Track Varying Numbers of Interacting People. In *Conference on Computer Vision and Pattern Recognition*, 2005. 1
- [71] S. Tang, B. Andres, M. Andriluka, and B. Schiele. Subgraph Decomposition for Multi-Target Tracking. In *Conference on Computer Vision and Pattern Recognition*, pages 5033–5041, 2015. 1
- [72] Y. Liu, Q. Hu, Y. Zou, and W. Wang, “Labelled non-zero particle flow for smc-phd filtering,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5197–5201. 2
- [73] B. Wang, G. Wang, K. L. Chan, and L. Wang. Tracklet Association with Online Target-Specific Metric Learning. In *Conference on Computer Vision and Pattern Recognition*, 2014. 2
- [74] B. Wang, G. Wang, K. L. Chan, and L. Wang. Tracklet association by online target-specific metric learning and coherent dynamics estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2016. 6
- [75] B. Wang, L. Wang, B. Shuai, Z. Zuo, T. Liu, K. L. Chan, and G. Wang. Joint Learning of Convolutional Neural Networks and Temporally Constrained Metrics for Tracklet Association. In *CVPRw*, 2016. 6
- [76] X. Wang. Intelligent Multi-Camera Video Surveillance: A Review. *Pattern Recognition*, 2013. 2
- [77] Y. Xiang, A. Alahi, and S. Savarese. Learning to track: Online multi-object tracking by decision making. In *iCVPR*, 2015. 1, 6
- [78] M. Xu, J. Orwell, and G. Jones. Tracking Football Players with Multiple Cameras. In *International Conference on Image Processing*, pages 2909–2912, October 2004. 1
- [79] C. Yang, R. Duraiswami, and L. Davis. Fast Multiple Object Tracking via a Hierarchical Particle Filter. In *International Conference on Computer Vision*, 2005. 1
- [80] S. Yi, H. Li, and X. Wang. Pedestrian Behavior Understanding and Prediction with Deep Neural Networks. In *European Conference on Computer Vision*, pages 263–279, October 2016. 2
- [81] J. H. Yoon, C.-R. Lee, M.-H. Yang, and K.-J. Yoon. Online multi-object tracking via structural constraint event aggregation. In *Conference on Computer Vision and Pattern Recognition*, 2016. 6
- [82] S. Yu, D. Meng, W. Zuo, and A. Hauptmann. The Solution Path Algorithm for Identity-Aware Multi-Object Tracking. In *Conference on Computer Vision and Pattern Recognition*, 2016. 2, 3
- [83] E. Zelniker, S. Gong, and T. Xiang. Global abnormal behaviour detection using a network of CCTV cameras. In *The Eighth International Workshop on Visual Surveillance-VS2008*, 2008. 2
- [84] L. Zhang, Y. Li, and R. Nevatia. Global Data Association for Multi-Object Tracking Using Network Flows. In *Conference on Computer Vision and Pattern Recognition*, 2008. 2
- [85] B. Zhou, X. Wang, and X. Tang. Understanding collective crowd behaviors: Learning a mixture model of dynamic pedestrian-agents. In *CVPR*, 2012. 6

Appendices

In Section [A](#), we provide the full definitions of the scoring functions n and m , described in Section 3.3 of the paper. In Section [B](#), we provide additional details of the optimizations used to improve the output of other method and to learn the patterns, described in Sections 4.1 and 4.2 of the paper. In Section [C](#), we provide full results of all methods on all datasets with various metrics, extending on the results from Section 6.4 of the paper. Finally, in Section [D](#), we provide evaluation of the computational requirements of our approach, in addition to the results given in Section 6.4 of the paper.

A. Full definitions of n and m functions

These functions are used to score the edges of a trajectory to compute how likely is it that a particular trajectory follows a particular pattern. As stated in Section 3.3 of the paper:

$$C(T, P, A) = \frac{\sum_{t \in T} M(t, p_{A(t_1)})}{\sum_{t \in T} N(t, p_{A(t_1)})}, \quad (11)$$

$$N(t, p) = n(I, t_1, p) + n(t_{|t|}, O, p) + \sum_{1 \leq j \leq |t|-1} n(t_j, t_{j+1}, p), \quad (12)$$

$$M(t, p) = m(I, t_1, p) + m(t_{|t|}, O, p) + \sum_{1 \leq j \leq |t|-1} m(t_j, t_{j+1}, p), \quad (13)$$

where T is a set of edges of all trajectories, A is the assignment between a trajectory and a pattern, and P is a set of patterns. As shown in (2) and (3), to score a trajectory we score all its edges plus the edges from I , the node denoting the beginnings of the trajectories, and the ones to O , the node denoting the ends of trajectories. As mentioned in the paper, we want $N(t, p)$ to reflect the full length of the trajectory and the pattern, and $M(t, p)$ to reflect the total length of the aligned trajectory and the pattern. In what follows, we provide definitions of n and m in all cases.

Figure 8. Example of computing the cost function C for three consecutive edges (A, B) , (B, C) , (C, D) . Dotted line around the pattern centerline c_p shows the area within the distance w_p to the pattern. The denominator contains the total length of the edges plus the total length of the pattern, while the numerator contains the parts aligned with each other (in green and blue). The edge (A, B) is not counted as aligned, because A is further from the pattern than its width w_p .

In **Table 6**, we show how to compute n and m for edges that link two detections and follow some pattern. For n we take the pattern length to be positive or negative depending on whether the projection of the edge to the pattern is positive or negative. For m , we penalize edges far from the pattern and edges going in the direction opposite to the pattern, in two different ways, which gives rise to the three cases shown in the table. In **Table 7**, we show how to compute n when one of the nodes is I or O , denoting the start or the end of a trajectory. A special case arises when a node is in the first or the last frame of an input batch, and a trajectory going through it does not need to follow the pattern completely. This results in a total of two cases we show in the table. In **Table 8**, we show the two cases when we assign the transition to no pattern $\emptyset$, one case when we assign a normal edge joining two detections, and the other when we assign edge from I or to O , indicating the beginning or the edge of the trajectory.

Case	Explanation	Figure
Normal edge aligned with the pattern: B and C are within distance w_p to the pattern centerline, P_B is earlier on the curve c_p than P_C .	For the edge (B, C) , we find the nearest neighbor of the two endpoints on the pattern, namely P_B and P_C . Formally, we have $P_B = \arg \min_{x \in c_p} \ B - x\ $. Then we project P_B and P_C orthogonally back onto (B, C) . This guarantees that $m(B, C, p) \leq n(B, C, p)$ with equality when (B, C) and (P_B, P_C) are two parallel segments of equal length, and also penalizes deviations from the pattern in direction.	$n(B, C, p) = BC + \widehat{P_B P_C}$ $m(B, C, p) = B^1 C^1 + P_B P_C $
Normal edge aligned with the pattern: B and C are further away than w_p from the pattern centerline, P_B is earlier on the curve c_p than P_C .	$n(B, C, p)$ is calculated in the same way as done in the previous case. To penalize deviations from the pattern in distance, we take $m(B, C, p) = 0$	$n(B, C, p) = BC + \widehat{P_B P_C}$ $m(B, C, p) = 0$
Normal edge not aligned with the pattern: P_B is later on the curve c_p than P_C .	To keep our rule about N being the sum of lengths of pattern and trajectory, we need to subtract the length of arc from P_B to P_C , as it is pointing in the direction opposite to the pattern. To penalize this behavior, we take $m(B, C, p)$ to be $- P_B P_C $, multiplied by $1 + \epsilon$. In practice, we use $\epsilon = 1$.	$n(B, C, p) = BC - \widehat{P_B P_C}$ $m(B, C, p) = - P_B P_C \times (1 + \epsilon)$

Table 6. Table describing full definitions of n and m in normal cases, when edges between two detections align with a pattern. They all follow naturally from the rule about N being the sum of length of trajectory and the pattern, and M being the sum of aligned lengths.

Case	Explanation	Figure
Edge from the source to a normal node / from a normal node to the sink	To keep our rule about N being the sum of lengths of pattern and trajectory, we need to add the length from the beginning of the pattern to the point closest to the node on the centerline / from the point closest to the node on the centerline to the end of the pattern. Since we didn't observe any parts of trajectory aligned with these parts, we take $m = 0$.	<p> $n(I, B, p) = P_I P_B$ $n(C, O, p) = P_C P_O$ $m(I, B, p) = 0$ $m(C, O, p) = 0$ </p>
Edge from the source to a normal node in the first frame of the batch / from a normal node in the last frame of the batch to the sink	We assume that our trajectories follow the path completely. However, this might be not true, which we observe from the middle, that is, the ones that begin in the first frame of the batch or end in the last frame. In that case we don't need to add the part of the pattern before / after the current point closest to the node, which is why we take $n = m = 0$.	<p> $n(I, B, p) = 0$ $n(C, O, p) = 0$ $m(I, B, p) = 0$ $m(C, O, p) = 0$ </p>

Table 7. Table describing full definitions of n and m in corner cases when one of the edges go through I or O , indicating the beginning or the end of a trajectory. They all follow naturally from the rule about N being the sum of length of trajectory and the pattern, and M being the sum of aligned lengths.

Case	Explanation	Figure
Normal edge aligned to no pattern	To keep our rule about N being the sum of lengths, we take n to be just the length of the trajectory, since we assume the length of empty pattern to be zero. We penalize such assignment by a fixed constant $\epsilon_\emptyset$, taking m to be n multiplied by such constant. In practice, we keep $\epsilon_\emptyset = 0.3$ when training from ground truth, or $\epsilon_\emptyset = -3$ otherwise.	$n(B, C, p) = BC $ $m(B, C, p) = BC \times (1 + \epsilon_\emptyset)$
Edge from the source / to the sink, aligned to no pattern	To keep our rule about N , we take both $n = m = 0$.	$n(I, B, p) = n(B, O, p) = 0$ $m(I, B, p) = m(B, O, p) = 0$

Table 8. Table describing full definitions of n and m in corner cases when there is no pattern. They all follow naturally from the rule about N being the sum of length of trajectory and the pattern, and M being the sum of aligned lengths.

B. Details of the optimization scheme

Here we provide details on our optimization schemes that improve the tracking output of other method and learn patterns, outlined in Sections 4.1 and 4.2 of the paper, respectively.

B.1. Tracking

As noted in the paper, we introduce the binary variables o_{ij}^p , denoting the number of people transitioning between the detections i and j , following pattern p . We put the following constraints on them:

$$\begin{aligned} \forall i \in \mathcal{D} \cup \mathcal{O} \quad \sum_{(i,j) \in \mathcal{E}, p \in P^*} o_{ij}^p &= 1, \\ \forall j \in \mathcal{D}, p \in P^* \quad \sum_{(i,j) \in \mathcal{E}} o_{ij}^p &= \sum_{(j,k) \in \mathcal{E}} o_{jk}^p. \end{aligned} \quad (14)$$

Then, during binary search, we fix a particular value of α , and check whether the problem constrained by (14) and the following has a feasible point:

$$\sum_{(i,j) \in \mathcal{E}, p \in P^*} (m(i, j, p) - \alpha n(i, j, p)) o_{ij}^p \geq 0 \quad (15)$$

If a feasible point exists, we pick a value of α to be the lower bound of the best α , for which the problem is feasible, otherwise we pick it as an upper bound. We start with the upper bound of 1 and lower bound of 0, and pick α as an average between the upper and the lower bound (dichotomy). We repeat this process 10 times, allowing us to find the correct value of α with the margin of 2^{-10} .

B.2. Patterns

As noted in the paper, we introduce the binary variables a_{tp} denoting that a ground truth trajectory t follows the pattern p , and binary variables b_p denoting whether at least one trajectory follows the pattern p .

$$\begin{aligned} a_{tp} &\in \{0, 1\}, \forall t \in T^*, p \in \mathcal{P}, \\ b_p &\in \{0, 1\}, \forall p \in \mathcal{P}, \\ \sum_{p \in \mathcal{P}} a_{tp} &= 1, \forall t \in T^*, \\ a_{tp} &\leq b_p, \forall t \in T^*, p \in \mathcal{P}. \end{aligned} \quad (16)$$

We then do the same binary search as described above to find the highest α , for which there exists a feasible point to a set of constraints (16) and the following:

$$\begin{aligned} \sum_{t \in T^*} \sum_{p \in \mathcal{P}} (m(t, p) - \alpha n(t, p)) a_{tp} &\geq 0, \\ \sum_{p \in \mathcal{P}} b_p &\leq \alpha_p, \\ \sum_{p \in \mathcal{P}} b_p M(p) &\leq \alpha_c. \end{aligned} \quad (17)$$

We do five iterations of binary search, and we obtain the right value of α with precision of 2^{-5} . To create a set of all possible patterns $\mathcal{P}$ we combine the set of all possible trajectories in the current batch (taking only those that start after the beginning of the batch and end before the end of the batch to make sure they represent *full* patterns of movement) with a set of possible lengths. For all datasets except **Station**, our set of possible lengths is $\{0.5, 1, 3, 5, 7, 9, 11, 13, 15, 17\}$, while for the **Station** dataset we use $\{0.05, 0.1, 0.2, 0.3, 0.4, 0.5\}$ of the tracking area, since we don't know the exact sizes of the tracking area, but only estimated homography between the ground and image plane.

C. Full results

Here we provide the full results of all the methods on all the datasets. Tables 9, 10 are the full versions of Table 2 of the paper, and Table 11 is the full version of Tables 3 and 4 of the paper. In Tables 9, 10, we compare the original output of the method with the improvements brought by our approach in both supervised and unsupervised manner. In Table 11, we compare the methods when using the ground truth set of detections as input. As in the paper, we report the results for the matching distances of 3m (0.1 of the tracking area for the **Station** dataset), and for IDF_1 metric we also show results for 1m to indicate that the ranking of the methods does not change, but the improvement brought by our methods is less visible due to reconstruction errors when we estimate the 3D position of the person from the bounding box. This fact is especially highlighted by the Table 11, where difference in the metric computed for distances of 3m. and 1m. is especially large.

Specifically, We report the IDF_1, identity level precision and recall IDPR and IDRC defined in [66], as well as MOTA, precision and recall PR and RC, and the number of mostly tracked MT, partially tracked PT and mostly lost trajectories ML defined in [13].

Method	Dataset	IDF_1	IDPR	IDRC	MOTA	PR	RC	MT	PT	ML
EAMTT	Town	0.72 (0.59)	0.76	0.68	0.73	0.92	0.82	158	68	20
EAMTT-i	Town	0.80 (0.63)	0.84	0.76	0.73	0.91	0.82	165	59	22
EAMTT-o	Town	0.82 (0.65)	0.83	0.80	0.74	0.89	0.86	182	44	20
JointMC	Town	0.75 (0.63)	0.90	0.65	0.64	0.95	0.68	128	54	64
JointMC-i	Town	0.77 (0.64)	0.91	0.66	0.64	0.95	0.68	129	52	65
JointMC-o	Town	0.76 (0.62)	0.88	0.67	0.65	0.93	0.71	138	50	58
MHT_DAM	Town	0.56 (0.45)	0.82	0.42	0.40	0.90	0.46	55	98	93
MHT_DAM-i	Town	0.56 (0.45)	0.83	0.42	0.40	0.90	0.46	59	90	97
MHT_DAM-o	Town	0.57 (0.45)	0.81	0.44	0.42	0.89	0.48	63	94	89
NOMT	Town	0.71 (0.62)	0.83	0.63	0.65	0.94	0.71	122	76	48
NOMT-i	Town	0.76 (0.65)	0.87	0.68	0.66	0.93	0.72	135	61	50
NOMT-o	Town	0.75 (0.63)	0.83	0.68	0.66	0.91	0.75	144	59	43
SCEA	Town	0.56 (0.43)	0.83	0.42	0.40	0.90	0.46	56	95	95
SCEA-i	Town	0.58 (0.45)	0.87	0.44	0.44	0.95	0.47	62	89	95
SCEA-o	Town	0.58 (0.43)	0.80	0.45	0.43	0.89	0.50	65	94	87
TDAM	Town	0.60 (0.48)	0.71	0.52	0.39	0.78	0.56	70	112	64
TDAM-i	Town	0.60 (0.48)	0.73	0.51	0.41	0.80	0.56	69	110	67
TDAM-o	Town	0.59 (0.45)	0.67	0.54	0.37	0.74	0.60	82	108	56
TSML_CDE	Town	0.68 (0.58)	0.75	0.63	0.72	0.95	0.79	143	79	24
TSML_CDE-i	Town	0.76 (0.62)	0.84	0.70	0.73	0.95	0.79	150	68	28
TSML_CDE-o	Town	0.78 (0.62)	0.82	0.74	0.74	0.92	0.83	161	68	17
CNNTCM	Town	0.58 (0.46)	0.79	0.46	0.45	0.90	0.53	63	110	73
CNNTCM-i	Town	0.61 (0.46)	0.80	0.49	0.48	0.90	0.55	73	96	77
CNNTCM-o	Town	0.62 (0.46)	0.77	0.52	0.48	0.87	0.59	85	95	66
KSP	Town	0.41 (0.26)	0.47	0.36	0.64	0.93	0.73	107	105	34
KSP-i	Town	0.69 (0.42)	0.78	0.61	0.65	0.93	0.73	118	91	37
KSP-o	Town	0.69 (0.42)	0.76	0.63	0.64	0.91	0.75	122	88	36
MDP	Town	0.59 (0.45)	0.65	0.55	0.50	0.81	0.68	103	97	46
MDP-i	Town	0.66 (0.49)	0.72	0.61	0.54	0.83	0.71	116	82	48
MDP-o	Town	0.63 (0.45)	0.66	0.61	0.50	0.79	0.73	113	94	39
RNN	Town	0.48 (0.30)	0.52	0.45	0.60	0.88	0.77	122	103	21
RNN-i	Town	0.59 (0.36)	0.65	0.55	0.61	0.90	0.76	125	98	23
RNN-o	Town	0.53 (0.34)	0.57	0.50	0.59	0.89	0.77	125	99	22
SORT	Town	0.62 (0.46)	0.81	0.50	0.57	0.98	0.61	49	152	45
SORT-i	Town	0.72 (0.47)	0.85	0.62	0.64	0.95	0.69	96	109	41
SORT-o	Town	0.65 (0.46)	0.83	0.60	0.60	0.90	0.65	174	58	14

Table 9. Full results for all methods on the **Town** dataset, when using our detections as input and using the results of state-of-the-art trackers as input. Number in brackets in IDF_1 column indicates result for the distance of 1 m.

Method	Dataset	IDF_1	IDPR	IDRC	MOTA	PR	RC	MT	PT	ML
KSP	ETH	0.45 (0.15)	0.45	0.45	0.47	0.72	0.71	182	148	22
KSP-i	ETH	0.62 (0.18)	0.71	0.54	0.48	0.75	0.57	134	144	74
KSP-o	ETH	0.57 (0.18)	0.59	0.67	0.49	0.67	0.76	217	121	14
MDP	ETH	0.55 (0.20)	0.63	0.48	0.40	0.79	0.60	113	194	45
MDP-i	ETH	0.58 (0.21)	0.76	0.46	0.41	0.83	0.50	105	143	104
MDP-o	ETH	0.58 (0.21)	0.64	0.62	0.41	0.72	0.69	157	146	49
RNN	ETH	0.51 (0.21)	0.54	0.49	0.48	0.80	0.73	170	162	20
RNN-i	ETH	0.54 (0.21)	0.76	0.39	0.48	0.85	0.44	68	184	100
RNN-o	ETH	0.54 (0.21)	0.40	0.47	0.47	0.64	0.76	205	127	20
SORT	ETH	0.67 (0.29)	0.82	0.57	0.50	0.87	0.61	130	175	47
SORT-i	ETH	0.66 (0.26)	0.84	0.55	0.49	0.86	0.56	136	129	87
SORT-o	ETH	0.67 (0.29)	0.79	0.68	0.49	0.80	0.70	167	148	37
KSP	Hotel	0.44 (0.14)	0.33	0.65	0.32	0.48	0.94	270	40	6
KSP-i	Hotel	0.53 (0.17)	0.38	0.75	0.33	0.47	0.94	273	35	8
KSP-o	Hotel	0.53 (0.17)	0.38	0.77	0.30	0.46	0.94	276	32	8
MDP	Hotel	0.40 (0.12)	0.34	0.46	0.33	0.47	0.64	133	92	91
MDP-i	Hotel	0.50 (0.13)	0.43	0.37	0.38	0.60	0.52	83	110	123
MDP-o	Hotel	0.37 (0.10)	0.28	0.47	0.30	0.40	0.67	143	105	68
RNN	Hotel	0.40 (0.14)	0.30	0.58	0.39	0.46	0.90	252	45	19
RNN-i	Hotel	0.40 (0.14)	0.30	0.59	0.39	0.46	0.90	258	38	20
RNN-o	Hotel	0.39 (0.13)	0.29	0.56	0.38	0.46	0.90	256	41	19
SORT	Hotel	0.54 (0.20)	0.45	0.68	0.37	0.55	0.82	207	87	22
SORT-i	Hotel	0.60 (0.20)	0.46	0.78	0.47	0.52	0.90	240	60	16
SORT-o	Hotel	0.58 (0.20)	0.46	0.78	0.35	0.53	0.88	238	64	14
KSP	Station	0.32	0.27	0.40	0.23	0.61	0.90	10166	1985	211
KSP-i	Station	0.42	0.35	0.52	0.19	0.60	0.91	10296	1879	187
KSP-o	Station	0.40	0.32	0.53	2.27	0.55	0.92	10597	1576	189
MDP	Station	0.48	0.39	0.63	0.51	0.56	0.90	9362	2293	437
MDP-i	Station	0.47	0.36	0.65	0.52	0.51	0.92	10047	1771	544
MDP-o	Station	0.47	0.37	0.66	0.50	0.52	0.92	10010	1930	422
RNN	Station	0.30	0.24	0.37	0.40	0.58	0.90	9826	2333	203
RNN-i	Station	0.30	0.24	0.38	0.41	0.59	0.90	9900	2260	202
RNN-o	Station	0.30	0.25	0.39	0.40	0.57	0.90	9898	2265	199
SORT	Station	0.50	0.50	0.50	0.32	0.71	0.72	5557	6181	624
SORT-i	Station	0.50	0.47	0.54	0.31	0.69	0.78	6996	4882	484
SORT-o	Station	0.52	0.48	0.57	0.31	0.67	0.79	7154	4703	505

Table 10. Full results for all methods on all the datasets except **Town**, when using our detections as input and using the results of state-of-the-art trackers as input. Number in brackets in IDF_1 column indicates result for the distance of 1 m.

Method	Dataset	IDF_1	IDPR	IDRC	MOTA	PR	RC	MT	PT	ML
KSP	Town	0.56 (0.47)	0.55	0.57	0.87	0.93	0.97	226	8	12
MDP	Town	0.87 (0.84)	0.92	0.82	0.87	0.99	0.89	184	38	24
RNN	Town	0.65 (0.57)	0.65	0.65	0.85	0.95	0.95	222	19	5
SORT	Town	0.88 (0.85)	0.93	0.84	0.90	1.00	0.90	203	34	9
OUR	Town	0.97 (0.92)	0.97	0.97	0.98	1.00	1.00	245	1	0
KSP	ETH	0.59 (0.12)	0.58	0.60	0.70	0.87	0.89	287	56	9
MDP	ETH	0.89 (0.18)	0.91	0.87	0.85	0.95	0.91	300	42	10
RNN	ETH	0.65 (0.16)	0.64	0.65	0.73	0.89	0.90	289	62	1
SORT	ETH	0.93 (0.20)	0.98	0.88	0.85	0.97	0.87	307	31	14
OUR	ETH	0.92 (0.19)	0.92	0.92	0.94	0.98	0.98	347	5	0
KSP	Hotel	0.60 (0.21)	0.61	0.58	0.74	0.90	0.86	217	69	30
MDP	Hotel	0.85 (0.33)	0.87	0.83	0.84	0.95	0.90	249	37	30
RNN	Hotel	0.70 (0.28)	0.69	0.71	0.78	0.91	0.94	284	29	3
SORT	Hotel	0.88 (0.36)	0.97	0.81	0.82	0.99	0.83	191	107	18
OUR	Hotel	0.94 (0.38)	0.94	0.94	0.97	1.00	1.00	314	1	1
KSP	Station	0.45	0.44	0.45	0.80	0.93	0.95	10957	832	573
MDP	Station	0.75	0.70	0.80	0.68	0.81	0.93	464	67	51
RNN	Station	0.40	0.39	0.40	0.68	0.90	0.94	10870	1244	248
SORT	Station	0.72	0.85	0.63	0.70	1.00	0.74	4968	6481	913
OUR	Station	0.70	0.62	0.62	0.77	0.99	0.99	579	3	0

Table 11. Full results for all combinations of methods and datasets, when using our set of ground truth detections. Number in brackets in IDF_1 column indicates result for the distance of 1 m.

D. Running time evaluation

Here we present the evaluation of running time of our approach depending on the parameters of the optimization. As mentioned in the Section 6.4 of the paper and shown in Fig. 9, the optimization time depends mostly on the number of possible transitions between people, which is controlled by D_1 . The time for learning the patterns grows approximately quadratically.

Figure 9. The running time and the number of variables of the optimization for tracking are approximately

- linear with respect to the number of frames in the batch (a),
- linear with respect to the number of patterns (b),
- superlinear with respect to the maximum distance at which we join the detections in the neighbouring frames D_1 , as it directly affects the density of the tracking graph (c),
- almost independent from the maximum distance in space D_2 and its time D_t at which we join the endings and beginning of the input trajectories D_2 , as it has almost no effect on the density of the tracking graph (d), (e);
- The running time and the number of variables of the optimization for learning patterns grows quadratically with the number of input trajectories, as each of them is both a trajectory that needs to be assigned to a pattern, and a possible centerline of a pattern (f).