

The Multi-Strand Graph for a PTZ Tracker

Shachaf Melman
Tel Aviv University
shachafme@gmail.com

Yael Moses
The Interdisciplinary Center
yael@idc.ac.il

Gérard Medioni Yinghao Cai
University of Southern California
medioni@usc.edu, caiyinghao@gmail.com

Abstract

High-resolution images can be used to resolve matching ambiguities between trajectory fragments (tracklets), which is one of the main challenges in multiple target tracking. A PTZ camera, which can pan, tilt and zoom, is a powerful and efficient tool that offers both close-up views and wide area coverage on demand. The wide-area view makes it possible to track many targets while the close-up view allows individuals to be identified from high-resolution images of their faces. A central component of a PTZ tracking system is a scheduling algorithm that determines which target to zoom in on.

In this paper we study this scheduling problem from a theoretical perspective, where the high resolution images are also used for tracklet matching. We propose a novel data structure, the Multi-Strand Tracking Graph (MSG), which represents the set of tracklets computed by a tracker and the possible associations between them. The MSG allows efficient scheduling as well as resolving – directly or by elimination – matching ambiguities between tracklets. The main feature of the MSG is the auxiliary data saved in each vertex, which allows efficient computation while avoiding time-consuming graph traversal. Synthetic data simulations are used to evaluate our scheduling algorithm and to demonstrate its superiority over a naïve one.

1. Introduction

We consider a system consisting of a single PTZ camera (which can pan, tilt and zoom) to solve the problem of tracking multiple pedestrians while also capturing their faces. A necessary component of such a system is a scheduling algorithm that determines at any time step whether to remain in zoom-out mode or to zoom in on a face. This paper presents an efficient new data structure, the multi-strand graph (MSG), for multiple target tracking using a single PTZ system, and a novel scheduling algorithm based on it.

Our method aims to overcome one of the main challenges of a multiple target tracker: trajectory fragmentation. Such fragmentation is caused by occlusions, by the joining

of two or more targets who then split (e.g., Fig. 1(a)), or when the PTZ camera zooms in on another target, creating what we call a *blind gap* (Fig. 1(e)). Matching trajectory fragments (tracklets) is complicated due to ambiguities caused by similarity in appearance and location of different targets. We propose using the faces captured in zoom-in mode together with the available information of the system state, to resolve such ambiguities.

The objective of the proposed system is to maximize the total length of the *labeled tracklets*, that is, trajectory fragments with associated high resolution face images captured in zoom-in mode. Formally, let Z be the set of targets, and let $\{\tau_i(z)\}_{i=1}^{n_z}$ be the set of its detected tracklets. For each target $z \in Z$ we define $\tau_L(z)$ to be the union of the labeled tracklets of z . The objective is given by

$$M = \frac{\sum_{z \in Z} |\tau_L(z)|}{\sum_{z \in Z} \sum_{i=1}^{n_z} |\tau_i(z)|}, \quad (1)$$

where $|\cdot|$ denotes the tracklet length. Our scheduler selects the target with the highest probability to maximize M to be the next target that will be zoomed in on. Note that resolving ambiguities (e.g., matching v_2 and v_9 in Fig. 1(a)) can significantly increase M , as desired.

We introduce the Multi-Strand Tracking Graph (MSG), which represents the tracklets computed by a tracker and their possible associations (its basic structure is similar to [15]). We show that a straightforward use of the graph for the abovementioned task requires a graph traversal. The main contribution of this paper is the proposed auxiliary data stored in each vertex. We use this data to efficiently compute the system state information without traversing the graph. The graph is constructed online and the auxiliary data is recursively computed based only on the vertex itself and on its direct parents. Hence, all the required information is available when scheduling decisions are made. Other contributions of this paper are the use of high-resolution images to resolve matching ambiguities of tracklets and the design of an efficient scheduling algorithm that uses the MSG.

System overview: The tracking system considered in this paper consists of a single PTZ camera, and several components, described below, are assumed to be available.

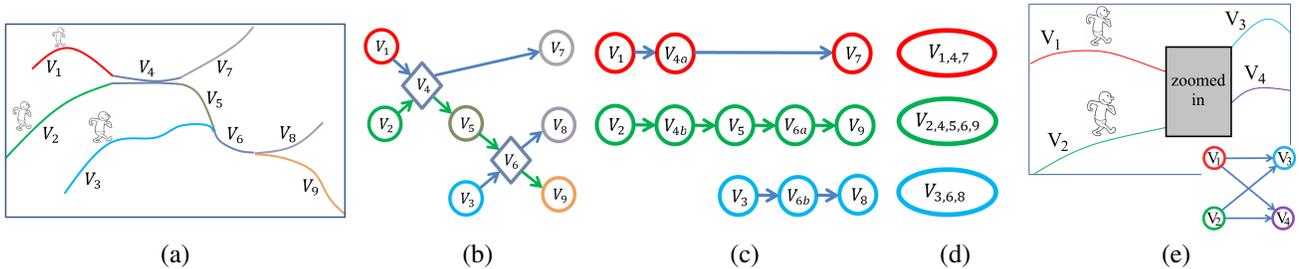


Figure 1. (a) Three targets walk in a scene, with two join and split events. (b) The corresponding MSG. Circle nodes represent solo vertices, and diamond nodes – compound vertices. (c) First untangling step after matching v_9 to v_2 using high-resolution images. (d) The final MSG after untangling: each solo vertex represents a full target trajectory. (e) A blind-gap scene and its MSG: targets who walk separately but move out of sight when the camera zooms in on another target, and then become visible again in the next zoom-out mode.

These include a tracker that detects and tracks pedestrians in zoom-out mode. It also detects joining and splitting events of two or more targets moving together (as in [15]). According to the proposed scheduler, the system selects a person to zoom in on using a camera control algorithm. The control algorithm chooses the FOV that makes it possible to zoom in on selected target (e.g., [2]). In the zoom-in mode, a face image is acquired and a face-to-face and a face-to-person matchings are computed. The system then zooms out, to the same wide view, to continue tracking. A person-to-person matching module associates tracklets when returning from zoom-in mode or after targets split from a group. Fig. 2 summarizes the system components. Our contribution to the system is the graph representation (MSG) and the efficient scheduling algorithm.

2. Previous Work

Scheduling of a single PTZ camera was considered in [14, 2, 12, 13]. Scenarios of joining/splitting targets were considered in [14, 2]. The goal of [14] was to minimize the slew time of an aerial camera tracking cars. High-resolution images were used to remove incorrect prediction hypothesis (stored as a tree). The greedy policy in [2] aimed to maximize the number of captured faces, considering the predicted time of each target to exit the scene and its movement angle w.r.t. the camera. An information-theoretic approach [12, 13] aims to decrease location uncertainty while capturing high-resolution images. A distributed game-theoretic approach for scheduling multiple PTZ cameras [8] aims to maximize the targets’ image quality and to capture their faces. None of the above scheduling algorithms considered the goal of resolving tracklet-matching ambiguities.

Other systems considered setups with both fixed and PTZ cameras, in a master-slave configuration. Matching in such setups is less challenging since a fixed camera continuously views the entire region. They vary from a single master and a single slave [1, 3] to multiple masters and multiple slaves [4, 6, 11, 17]. The considered objectives are to ac-

quire once [4], or as many times as possible [3], the face of each target, or to minimize camera motion [1]. The scheduling methods consider the expected distance from the camera [11], the viewing angle [1, 6, 14, 17], and expected occlusions [6, 14]. In addition to these objectives, our algorithm also considers how zooming in contributes to the resolution of past and future ambiguities of tracklet matching.

Graphs were previously used to represent relations between tracklets [5, 10, 16, 18, 19], where the weighted edges reflect the appearance similarity and the consistency of location with respect to the computed motion direction and sometimes speed. A graph with a similar structure to the MSG [5, 9, 15] was used to associate isolated tracklets of targets with indistinct appearance as well as tracklets of a set of targets that cannot be separated. The joins/splits of targets were computed by a tracker. The association of single target tracklets is solved by finding the most probable set of paths. All these papers use the target’s location and only one appearance descriptor level for matching while we use both low- and high-resolution images. Moreover, they do not use auxiliary data, which allows efficient scheduling and online graph updating in our method.

3. Method

We describe the basic structure of the MSG graph and its extension with auxiliary data for efficient matching by elimination. Then we present our scheduling algorithm.

3.1. Graph Definition

The basic structure of the MSG is a dynamic augmented graph, $G = (V, E)$, where V represents the set of tracklets computed by the tracker, and E the candidate associations of different tracklets computed by some available matching algorithm (similar to [15]). Each vertex is associated with the information regarding its tracklet. We consider two types of vertices that represent two types of tracklets. A *solo vertex* represents the tracklet of a single target and a *compound vertex* represents the shared tracklet of *joined*

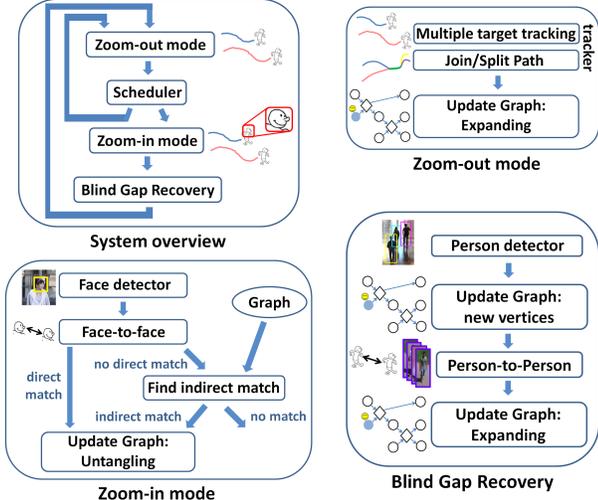


Figure 2. An overview of the system and its operation modes

targets, that is, a set of targets that walk together (Fig. 1(b)).

A directed edge, $e = (v_i, v_j) \in E$, represents the case in which at least one of the targets associated with v_i may also be associated with v_j , and v_i and v_j correspond to consecutive time intervals (ignoring zoom-in time). A compound vertex is generated as the child of other vertices when the tracker detects that the targets’ trajectories are joined into indistinguishable tracklets (see Fig. 1(a), 1(b)). A new solo vertex is generated when a new target enters the scene, a target trajectory splits from those of others (as a child of the compound vertex), or a target reappears when the camera returns to zoom-out (after a blind gap).

Edges between solo vertices are generated at consecutive layers (e.g., when returning from zoom-in mode), according to a matching algorithm that is based on the targets’ low-resolution images captured in zoom-out mode and on their locations. When the matching of a new solo vertex is ambiguous, edges are set between the vertex and all the matching candidates, forming an *X-type ambiguity* (see Fig. 1(e)). Additional edges are generated between compound and solo vertices, based on the tracker’s detection of splitting and joining targets.

3.2. Untangling

When no ambiguities are present, the trajectory of each target can be fully recovered, and the graph contains only unconnected solo vertices. We wish to reduce as much as possible the number of vertices by concatenating consecutive tracklets to a single tracklet when possible. When a univocal matching exists for a consecutive set of tracklets (all of the same target), their corresponding vertices form a *solo chain* in the graph – in which the out degree and the in degree of a vertex’s parent and child, respectively, are one (Fig. 1(c)). A solo chain, of any length, can be merged to

a single solo vertex (Fig. 1(d)). A compound chain can be defined and merged similarly.

A univocal matching of a pair of non-consecutive solo vertices, $\{u, v_L\}$, can be computed using an available face-to-face matching algorithm. In addition, an *indirect match* can also be obtained by elimination (see Sec. 3.3).

Such a univocal matching of $\{u, v_L\}$ can be used for untangling the graph as long as the connected component of u and v_L is a DAG (a graph that does not contain cycles). In this case, a breadth-first-search (BFS) algorithm is used to recover the graph path, $\ell(u, v_L)$ (e.g., $\ell(v_2, v_9)$ in Fig. 1(b)). All vertices of $\ell(u, v_L)$ are guaranteed to represent consecutive tracklets of the same target. Hence, edges ‘to’ and ‘from’ solo vertices of $\ell(u, v_L) \setminus u$ (representing X-type ambiguities) are removed except those that are part of the path. Each compound vertex $v_{comp} \in \ell(u, v_L)$ is split into two vertices. One solo vertex represents only the labeled target and is linked only to the solo chain. The second vertex, v_{split} , represents the remaining targets of the compound vertex and is disconnected from the chain (Fig. 1(c)). As a result, a solo chain of the labeled target, and possibly additional solo chains of other targets, are obtained. Each chain can be merged into a single solo vertex (Fig. 1(d)). Note that no information is lost in the untangling process.

3.3. Matching by Elimination

When there is sufficient confidence that a labeled vertex cannot be matched to any of the previously labeled vertices, the vertex can sometimes be indirectly matched to an unlabeled vertex by elimination. For example, assume that v_1, v_2 and v_9 in Fig. 1(b) were labeled and no match was found for the faces of either v_1 or v_2 with v_9 . It is possible to deduce that v_3 is the correct match to v_9 . Similarly, if only v_3 and v_9 were labeled, the non-source v_5 is deduced to be this match. We next define when an indirect match can be found in the general case, and how to compute it efficiently. Let V_L be the set of labeled solo vertices. We define an *unlabeled path* between w and v , $\tilde{\ell}(w, v)$, to be a path that does not contain any labeled vertex except possibly w and v , that is, $\forall u \in \tilde{\ell}(w, v) \setminus \{w, v\}, u \notin V_L$.

Claim 1: Sufficient and necessary conditions for a solo vertex $w \notin V_L$ to be an indirect match to $v \in V_L$ are (i) v cannot be matched to a previously labeled vertex; (ii) there exists an unlabeled path, $\tilde{\ell}(w, v)$; (iii) if an unlabeled solo vertex w' satisfies (ii) then $w' \in \tilde{\ell}(w, v)$. (See proof in [7].)

When (i) holds, an indirect match to a labeled solo vertex v_L can be computed in a straightforward manner by traversing the graph backwards from v_L and checking whether a vertex w that satisfies (ii) and (iii) exists. This is clearly time consuming. Instead, we propose to store auxiliary data in each vertex; this data, which can be efficiently computed online from the vertex itself and its

parents, makes it possible to directly compute an indirect match, if one exists. It is also used for scheduling (Sec. 3.4).

Auxiliary data for matching by elimination: We define w to be an *origin* of v if (i) w is a solo vertex; (ii) there exists an unlabeled path $\tilde{\ell}(w, v)$ and (iii) w is either a source of the graph (*unlabeled origin*) or a labeled vertex (*labeled origin*). A labeled vertex is the origin of itself and has no unlabeled origins. The set of origins of v consists of the set of vertices – each associated with a distinct target ID – that may represent the same target as v . Note that only a labeled origin of v may be directly matched to v .

We observe that a solo vertex v may have an indirect match only if it has at least one unlabeled origin (otherwise it can only be directly matched). Furthermore, v may have an indirect match only if just one of its parents has unlabeled origins (otherwise, the unlabeled origins, one from each parent, do not satisfy (iii) of Claim 1). Hence, to compute whether an indirect match exists, it is sufficient to store in each vertex the number of its unlabeled origins, denoted by $n_{-L}(v)$, and the single parent that has unlabeled sources, if one exists, $p_{\leftarrow}(v)$ (set to zero if one does not exist). Let $P(v)$ be the set of parents of v . Both $n_{-L}(v)$ (given by summing the number of unlabeled origins of $P(v)$) and $p_{\leftarrow}(v)$ can be recursively defined as follows:

$$n_{-L}(v) = \begin{cases} 1 & v \notin V_L \ \& \ |P(v)| = 0 \\ 0 & v \in V_L \\ \sum_{v_i \in P(v)} n_{-L}(v_i) & \text{otherwise.} \end{cases} \quad (2)$$

$$p_{\leftarrow}(v) = \begin{cases} u & \exists! u \mid u \in P(v) \ \& \ n_{-L}(u) > 0 \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

If a vertex w is the indirect match of a solo vertex $p_{\leftarrow}(v)$, it is also the indirect match of v . Therefore, we can efficiently and recursively compute the single candidate of an indirect match of v , $C(v)$:

$$C(v) = \begin{cases} 0 & n_{-L}(v) = 0 \\ C(p_{\leftarrow}(v)) & p_{\leftarrow}(v) \neq 0 \ \& \ C(p_{\leftarrow}(v)) \neq 0 \\ solo(v) & \text{otherwise,} \end{cases} \quad (4)$$

where $solo(v)$ holds v if v is a solo vertex and 0 otherwise. Note that a compound vertex, v , cannot have an indirect match; however, the value $C(v)$ contains the candidate indirect match for its descendants.

After labeling v and untangling the MSG (when possible), the auxiliary data is recalculated to be $n_{-L}(v) = 0$. This reflects that ambiguities of this target prior to the labeling are no longer relevant for future ambiguities. After each labeling and once the untangling is complete, C must also be recalculated for all the vertices that were disconnected from the solo chain during the untangling process. Each of

these vertices then propagates the updated value to all its descendants, who recalculate their own values accordingly.

3.4. Scheduling

The scheduler selects the tracklet whose target’s face will be acquired in the next zoom-in mode or selects to stay in zoom-out mode. The score it provides reflects the expected contribution of a tracklet labeling to maximize the system’s objective, M (Eq. 1). A prerequisite for choosing to zoom in on an unlabeled target is that the acquisition of its face is expected to be successful. A Boolean value that indicates the expected success, $E_s(v)$, can be computed by the tracker in a similar manner to previous studies (e.g., [2]). For example, the motion direction can be used for predicting occlusions and time to exit, and whether the face will be visible to the camera.

Ideally, the minimal number of required labelings for a full trajectory retrieval of N targets is N , one per target. The upper bound of the required number of ideal labelings is $\sum_i (2n_i - 1) \leq 2N - 1$, where $n_i = |G_i|$ and $G_i \subseteq G$ is a connected component. This sum includes the labeling of the first and last solo-walking tracklet of each target z . Thus, the full trajectory of z is recovered and labeled (under the no-cycle assumption). Note that each untangling may further reduce the required number of labelings.

In practice, an ideal labeling set is often impossible to obtain: the online algorithm leaves limited time for zooming in, and each labeling may cause additional ambiguities due to a blind gap. Moreover, the target identity and hence its contribution to M is unknown before zooming in.

Therefore, we propose a scheduling algorithm that approximates the estimated contribution of labeling each of the targets or staying in zoom-out mode to maximize M . A *zoom-out score*, S_{zo} , can reflect global properties of the scene, such as the number of new targets expected to enter it, and the prevention of X-type ambiguities caused by a blind gap in zoom-in mode. Here we set it to be a constant. A *labeling score*, $S_L(v)$, is set for each vertex and reflects the expected contribution to M if v is chosen to be labeled. The v to be selected for labeling is the one with the highest S_L as long as $S_L(v) > S_{zo}$. Otherwise, the system remains in zoom-out mode.

3.5. Labeling Score & Auxiliary Data

The score $S_L(v)$ is a weighted sum of two terms that estimate the expected resolution of future (S_F) and past (S_P) ambiguities:

$$S_L(v) = E_s(v) (\alpha(v)S_F(v) + \beta(v)S_P(v)), \quad (5)$$

where the weights $\alpha(v)$ and $\beta(v)$ are higher for source and expected sink vertices, respectively.

Future ambiguities: The probability that v was not labeled before is given by $(n_{-L}(v)/n_o(v))$, where $n_o(v)$

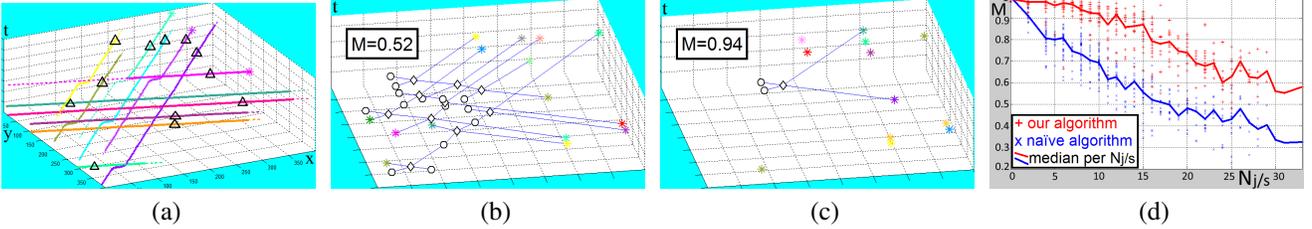


Figure 3. Simulation screen shots. (a) Trajectory trails of all the targets that entered or already exited the scene, using (X,Y,t) coordinates. Solid lines: labeled tracklets; dashed lines: unlabeled tracklets; black triangle: a face acquisition; asterisk: a labeled target. (b) The final MSG of the naïve scheduler. (c) The final MSG obtained by our method. Asterisk: a labeled vertex. (d) M as a function of $N_{j/s}$.

and $n_{-L}(v)$ are the number of origins and unlabeled origins of v , respectively. The score $S_F(v)$ is defined to be $S_F(v) = Join(v)n_{-L}(v)/n_o(v)$, where $Join(v)$ is a Boolean value that reflects that the target of v is expected to join another target with a similar appearance (computed by the tracker). The value of $n_{-L}(v)$ can be recursively computed (see Eq. 2). In a similar manner, $n_o(v)$ can also be recursively computed (as specified in [7]).

Past ambiguities: The score $S_P(v)$ reflects the expected increase in the length of the labeled tracklets, $L = \sum_{z \in Z} |\tau_L(z)| = \sum_{v \in V_L} |\tau(v)|$, if v is chosen for labeling. Labeling a vertex v increases the length of L by the length of the tracklet $\tau(v)$. In addition, if v is matched to u , directly or indirectly, then L is extended by the sum of $|\tau(w)|$ over all $w \in \ell(u, v)$ such that $w \notin V_L$. The identity of v 's target, and hence the origin to which v will be matched, is unknown prior to its labeling. Hence, we average over all possible increases of L with respect to the $n_o(v)$ possible origins that may be the match of v . Let $L_{-dir}(v)$ and $L_{dir}(v)$ be the increased sum of L over the sets of unlabeled origins and labeled origins, respectively. Then,

$$S_P(v) = \frac{1}{n_o(v)} (\Delta L_{dir}(v) + L_{-dir}(v)). \quad (6)$$

A straightforward computation of $S_P(v)$ is by graph traversal. To avoid such a computationally expensive operation for each candidate vertex, we store in each vertex the auxiliary data fields, $\Delta L_{dir}(v)$, $L_{-dir}(v)$, and $n_o(v)$. These values are recursively computed. We next describe the computation of $L_{-dir}(v)$. (For $\Delta L_{dir}(v)$, see [7]).

When no match is found (either direct or indirect), the contribution of labeling v to $L_{-dir}(v)$ is $|\tau(v)|$ for each unlabeled origin. If an indirect candidate match is given by $C(v)$, that is, $C(v) \notin \{0, v\}$, then the additional contribution of labeling v is given by $|\tau(\ell(C(v), p_{\leftarrow}(v)))|$. To compute it, $L_{-dir}(p_{\leftarrow}(v))$ is computed recursively as follows:

$$L_{-dir}(v) = \begin{cases} |\tau(v)| \cdot n_{-L}(v) + L_{-dir}(p_{\leftarrow}(v)) & v \notin V_L \ \& \ p_{\leftarrow}(v) \neq 0 \\ |\tau(v)| \cdot n_{-L}(v) & v \notin V_L \ \& \ p_{\leftarrow}(v) = 0 \\ 0 & v \in V_L. \end{cases} \quad (7)$$

Complexity: The computation of the score is linear with $|P(v)|$ – which is expected to be small for each visible target – instead of $O(|E| + |V|)$ for the necessary graph traversal without the auxiliary data. Note that without untangling, the graph is expected to grow very fast when more targets enter the scene and many tracklets are detected, hence making the alternative $O(|E| + |V|)$ even worse. Due to the overhead incurred by untangling, the auxiliary data of all the descendant vertices must be updated. In the worst case, it will require updating $O(|V|)$ vertices. However, this operation is rarely performed. Moreover, each time it takes place, the size of the graph is greatly reduced. Hence, the amortized complexity of updating the graph is expected to be $O(1)$ for each new vertex. A formal proof of this conjecture is left for future study.

4. Experiments

We used simulated data as an input to our method to evaluate the scheduler's performance independently from that of the other modules. Simulated data also makes it possible to bypass the limitations of comparing online algorithms on the same real data; each algorithm dictates different zoom-in operations, thus changing the data. We implemented our method as well as the simulated data using Matlab.

The simulated scene consists of a set of targets walking on a grid of intersecting diagonal roads. The targets' velocities (speed and direction), entrance time and location, and the probability that meeting targets start walking together, are determined randomly. All targets have the same low-resolution appearance to increase ambiguity, and low-resolution images are not used for matching.

We present the objective score of our algorithm, M (Eq. 1), as a function of the expected ambiguities in the scene. It is computed when the simulation ends and is based on the MSG's tracklets and on the ground truth. We estimate the ambiguities of the scene by $N_{j/s} = \sum_{z \in Z} n_{j/s}(z)$, where $n_{j/s}(z)$ is the number of times a solo-walking target, z , joins a group and then splits to walk alone again. Note that in practice blind gaps may cause additional ambiguities. For comparison we consider a naïve scheduler [2] that selects the unlabeled target predicted to leave the

scene first. In both cases, we assume that the tracker provides the necessary available information (e.g., whether the face is expected to be captured successfully). Examples of the simulation results of our and the naïve schedulers are shown in Fig. 3 (additional examples can be found in [7]).

The results on 414 simulations are presented in Fig. 3(d). When $N_{j/s}$ is small, the performance of our and the naïve algorithms is close to perfect. When $N_{j/s}$ increases, the performance of our method decreases, mainly due to the limited time available to label all the desired targets. However, for moderate ambiguity of $N_{j/s} = 15$, our method still performs well: $M > 0.85$. The superiority of our algorithm over the naïve one is apparent both for moderate and high $N_{j/s}$. For example, the score of the naïve algorithm obtained for $N_{j/s} = 15$ is $M = 0.55$, which is lower than the worst score of our method, for $N_{j/s} > 30$.

Two components of our algorithm contribute to its superiority over the naïve one. The global view we keep of the system state allows us to associate one or more labeled tracklets of the same target with additional tracklets of that target. Using graph terminology, this corresponds to the untangling and merging of vertices, either by direct labeling or as a byproduct of labeling other targets. In addition, our scheduling method explicitly considers the task of disambiguating tracklet associations, and uses global information of the current state of the system efficiently.

5. Discussion & Future Work

We proposed a method for tracking multiple pedestrians and capturing their faces using a single PTZ camera. The goal of the system is to maximize the length of the labeled trajectories recovered by the tracker. Our main contribution is a novel data structure, MSG, that efficiently utilizes all the available global information of a tracking system. The auxiliary data of the MSG is used for an efficient scheduling algorithm that resolves or prevents tracklet ambiguities and matches tracklets directly or indirectly via target labeling. The MSG may be modified for various applications that use several cameras, with or without overlapping fields of view, when two distinct resolution levels can be used.

Our method aims to represent and efficiently use the data available from basic components of trackers and recognition systems, most of which are assumed to be deterministic for ease of exposition. It is clearly prone to the expected errors of each of these components. Our method can be extended to handle a probabilistic setting where each component provides a degree of confidence for its output. This can be integrated into the graph by, for example, associating a weight with each edge. In the current system, X-type edges can represent the output of a probabilistic person-to-person matching algorithm. A threshold on the face-to-face matching confidence may be used for deciding whether to untangle the graph or wait for additional information.

Acknowledgment: This research was supported by the Israeli Ministry of Science, grant number 3-8700 and by award No. 2011-IJ-CX-K054, awarded by the National Institute of Justice, Office of Justice Programs, U. S. Department of Justice.

References

- [1] A. Bagdanov, A. del Bimbo, and F. Pernici. Acquisition of high-resolution images through on-line saccade sequence planning. In *VSSN*, 2005. 2
- [2] Y. Cai, G. Medioni, and T. Dinh. Towards a practical PTZ face detection and tracking system. In *WACV*, 2013. 2, 4, 5
- [3] C. Costello, C. Diehl, A. Banerjee, and H. Fisher. Scheduling an active camera to observe people. In *VSSN*, 2004. 2
- [4] C. Costello and I. Wang. Surveillance camera coordination through distributed scheduling. In *CDCECC*, 2005. 2
- [5] J. Henriques, R. Caseiro, and J. Batista. Globally optimal solution to multi-object tracking with merged measurements. In *ICCV*, 2011. 2
- [6] S. Lim, L. Davis, and A. Mittal. Constructing task visibility intervals for video surveillance. *MS*, 12(3), 2006. 2
- [7] S. Melman, Y. Moses, G. Medioni, and Y. Cai. The Multi-Strand Graph for a PTZ Tracker. *ArXiv e-papers: 1506.08485*, 2015. 3, 5, 6
- [8] A. Morye, C. Ding, A. Roy-Chowdhury, and J. Farrell. Distributed constrained optimization for bayesian opportunistic visual sensing. *TCST*, 2014. 2
- [9] P. Nillius, J. Sullivan, and S. Carlsson. Multi-target tracking-linking identities using bayesian network inference. In *CVPR*, 2006. 2
- [10] J. Prokaj, M. Duchaineau, and G. Medioni. Inferring tracklets for multi-object tracking. In *CVPRW*, 2011. 2
- [11] F. Qureshi and D. Terzopoulos. Surveillance in virtual reality: System design and multi-camera control. In *CVPR*, 2007. 2
- [12] P. Salvagnini, F. Pernici, M. Cristani, G. Lisanti, I. Masi, A. Del Bimbo, and V. Murino. Information theoretic sensor management for multi-target tracking with a single pan-tilt-zoom camera. In *WACV*, 2014. 2
- [13] E. Sommerlade and I. Reid. Information-theoretic active scene exploration. In *CVPR*, 2008. 2
- [14] T. Strat, P. Arambel, M. Antone, C. Rago, and H. Landan. A multiple-hypothesis tracking of multiple ground targets from aerial video with dynamic sensor control. In *SPIE*. 2004. 2
- [15] J. Sullivan and S. Carlsson. Tracking and labelling of interacting multiple targets. In *ECCV*. 2006. 1, 2
- [16] X. Wang, E. Türetken, F. Fleuret, and P. Fua. Tracking interacting objects optimally using integer programming. In *ECCV*. 2014. 2
- [17] C. Ward and M. Naish. Scheduling active camera resources for multiple moving targets. In *CCECE*, 2009. 2
- [18] Z. Wu, T. Kunz, and M. Betke. Efficient track linking methods for track graphs using network-flow and set-cover techniques. In *CVPR*, 2011. 2
- [19] B. Yang and R. Nevatia. An online learned CRF model for multi-target tracking. In *CVPR*, 2012. 2