Camera Network Topology Estimation Using Sparse Methods

Francisco Lourenço
Instituto Superior Técnico
Email: francisco.m.lourenco@tecnico.ulisboa.pt

***Abstract*—Knowledge of the topological structure of a camera network can help with tasks such as person re-identification by limiting the correspondence possibilities for each observed person. In situations where manually determining the camera layout is not practical, automated camera network topology estimation can be used to discover the connections between the different camera views.**

This work proposes two methods for discovering the topology of a non-overlapping camera network based on sparse inverse covariance estimation. Using sparse inverse covariance estimation an empirical covariance matrix obtained from observed feature data can be used to discover a sparse graphical model representing the conditional dependence structure between its nodes. The first method establishes direct correspondences between appearances of the same person in different cameras via color based thresholding, whereas the second method estimates a graph representing connections between events based on observed features without directly establishing correspondences.

Favorable experimental results are shown for both methods using randomly generated trajectory data obtained from an event simulator. Synthetic noise was added to the generated color descriptors to test the methods' robustness to noisy data.

I. Introduction

Video surveillance is increasingly common nowadays [1], being used in both public and private environments such as offices, hospitals, homes and even outdoor spaces like roads. In many large environments a network composed of multiple cameras is necessary for an effective surveillance [2]. One of the most important aspects of surveillance in a network of cameras is the task of following targets across multiple camera views.

In order to properly follow a target through the various cameras composing the network, especially cameras with non-overlapping fields of view, it can be helpful to know the spatial relations between them. While precise metric information of each camera's position can be useful, it is often enough to find a topological representation of the camera network map. This topological map can take the form of a graph where each vertex represents either a camera or an entry/exit zone of a camera's field of view. This latter option allows for a better understanding of the network, especially in more closed indoor areas where different hallways and doorways in the same camera view can lead to different places. The graph edges represent connections between these zones - if two vertices are connected by an edge in the graph, then there is an off-camera path between the two zones and it is possible for objects that disappear in an entry/exit zone associated with one of the two vertices to appear in the zone associated with the other vertex.

The problem of camera network topology estimation has been approached in several different ways. Kamenetsky [3] divides this problem into two main types: overlapping and non-overlapping. In overlapping camera networks the fields of view of the different cameras overlap, partially observing the same area. In the topology graphs for these problems two cameras are linked by an edge if their fields of view overlap, forming a *vision graph*. In non-overlapping problems different cameras observe different parts of the environment, without any overlap in their fields of view. In these cases the relationships between the different cameras are represented by a *communication graph*. These are undirected graphs where each node represents one camera and an edge connecting two nodes indicates that transitions between the two cameras are possible. These edges can also have weights corresponding to transition probabilities and times. Our work is focused on these non-overlapping camera networks, where the graph nodes represent specific zones through which people enter and exit the camera's field of view.

In addition to this distinction, Li et al. [4] state that non-overlapping problems can be approached with correspondence based methods and correspondence-free methods. Correspondence based methods require direct correspondences between people sighted in different cameras, that is, they need to identify the same person in different camera views. This can be done manually or through an automatic re-identification method. Correspondence-free methods have no such requirements. In [4] the authors propose a correspondence-free method for camera network topology estimation based on accumulated cross correlation and Gaussian fitting.

Another correspondence-free method is proposed by Makris et al. [5]. For this method the entry and exit zones of the different cameras are learned automatically from observed trajectories using an expectation-maximization algorithm. Visible connections are learned through a standard on-camera object tracking algorithm and the cross-correlation data for each pair of nodes is used to identify invisible connections as well as overlapping nodes and additional visible connections.

Cho et al. [6] propose a correspondence based iterative method to estimate camera network topology and perform automated person re-identification. Correspondence data obtained from automated re-identification is used to estimate the network topology and this information is used to improve the re-identification. This process is repeated until the topology estimate converges.

In this work we propose two camera network estimation methods focused on discovering an undirected graph of invisible connections between specific manually determined entry and exit zones belonging to non-overlapping cameras. An edge connecting two entry/exit zones indicates that a person disappearing from one of those zones may reappear in the other, and no transition probabilities or time distributions are computed. The first method is correspondence based, while the second is a correspondence-free method. Both methods use a sparse graph learning algorithm detailed in section III to solve the topology estimation problem by using observed feature data to discover the graph that represents the camera network.

II. Problem Definition

As people move through the area observed by the network of non-overlapping cameras, they will either be on-camera (within a camera's field of view) or off-camera (outside the field of view of every camera). The transition between the on-camera and off-camera areas is done through entry and exit zones, which are the parts of a camera's field of view from which people appear and disappear. All of these zones are simultaneously entry zones and exit zones, as any zone through which people can enter a camera's field of view can also be used to exit it. These entry/exit zones will serve as the nodes in the camera network graphs we want to estimate, and these graphs will represent the connections between the different entry/exit zones - an edge connecting two zones indicates that a person can travel between those two zones without passing through any other.

We can divide each trip made by a person into its on-camera (*visible*) and off-camera (*hidden* or *invisible*) portions, as can be seen in figure 1.

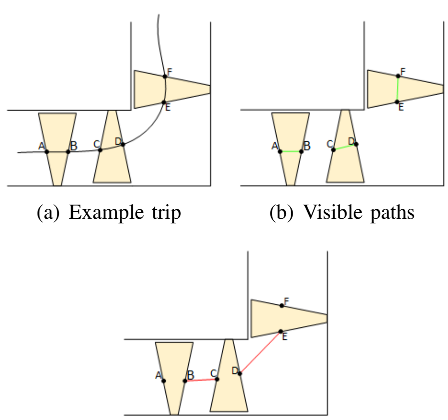

(a) Example trip (b) Visible paths

(c) Hidden paths

Fig. 1. Example of a path travelled by a person as well as its visible and hidden portions. The areas in yellow represent the fields of view of three different non-overlapping cameras in a camera network.

The points A, B, C, E and F are entry/exit zones of the fields of view of three cameras in a network. The example trip shown in figure 1(a) passes through all of these zones and can be divided into five sequential segments: $A \to B$, $B \to C$, $C \to D$, $D \to E$ and $E \to F$. Segments where both the beginning and the end nodes belong to the same camera's field of view ($A \to B$, $C \to D$ and $E \to F$) form visible paths, while segments where the two nodes are in the fields of view of different cameras ($B \to C$ and $D \to E$) form hidden paths. Our goal is to find the graph that represents all the hidden paths in a camera network, that is, all the paths that can be travelled off-camera. These graphs are undirected, as all paths can be travelled both ways - if a person can travel from node A to node B, they can also travel from node B to node A.

These hidden path graphs will be found based on entry and exit events. These events correspond to a person entering or exiting a camera's field of view. Associated with each entry/exit event is the entry/exit zone through which it happened, the instant when it happened and a descriptor of the person involved in the event. This information can be used to match different entry/exit events, forming trips travelled by people in the camera network and thereby identifying possible paths. Unlike with entry/exit zones, the distinction between entry and exit events is important, as entry events can not be matched with other entry events, or exit events with other exit events.

III. Graph Estimation Method

An undirected weighted graph can be fully represented by its Laplacian matrix [7]. The Laplacian matrix $\Delta = [\Delta_{ij}]$ of an undirected weighted graph G with n nodes is given by

$$\Delta_{ij} = \begin{cases} -w_{ij}, & i \neq j \\ \sum_k w_{ik}, & i = j \end{cases} \quad , \; i, j = 1, ..., n, \tag{1}$$

where w_{ij} is the weight of the graph edge connecting nodes i and j. These weights are nonnegative (that is, $w_{ij} \geq 0 \; \forall i, j = 1, ..., n$) and it is considered that if two nodes i and j are not directly connected by an edge in the graph, the weight w_{ij} between the two nodes is 0. As the diagonal of the Laplacian contains the weighted degree of the corresponding nodes and can be obtained from the non diagonal entries of Δ, the graph can be fully represented by these non diagonal entries alone. Therefore, it is sufficient to estimate the weighted adjacency matrix $W = [w_{ij}]$ with $i, j = 1, ..., n$. This matrix is symmetric as the graph is undirected, and it is considered that $w_{ii} = 0$ for all $i = 1, ...n$.

The weighted adjacency matrix W can be estimated as described by Lake & Tenenbaum [8] by defining the following *maximum a posteriori* (MAP) estimate:

$$\begin{aligned} \hat{W} &= \underset{W}{\operatorname{argmax}} \; p(W|D) \\ &= \underset{W}{\operatorname{argmax}} \; p(D|W) \cdot p(W). \end{aligned} \tag{2}$$

The matrix D is an $n \times m$ feature matrix with each of the m columns representing a feature vector $f^{(k)} = (f_1^{(k)}, ..., f_n^{(k)})^T$ where $f_i^{(k)}$ represents the value of feature k pertaining to node

i. These features are assumed to be independent and identically distributed draws from a normal distribution

$$p(f^{(k)}|W) = (2\pi)^{-\frac{n}{2}}|\tilde{\Delta}^{-1}|^{-\frac{1}{2}} \cdot \cdot \exp\left(-\frac{1}{4}\sum_{i,j} w_{ij}\left(f_i^{(k)} - f_j^{(k)}\right)^2 - \frac{1}{2}\sum_i \frac{\left(f_i^{(k)}\right)^2}{\sigma_i^2}\right). \tag{3}$$

This is equivalent to $p(f^{(k)}|W) \sim N(0, \tilde{\Delta}^{-1})$, where $\tilde{\Delta}$ is a regularized graph Laplacian given by $\tilde{\Delta} = \Delta + \text{diag}\left(1/\sigma_i{}^2\right)$, with $\text{diag}\left(1/\sigma_i{}^2\right)$ being a diagonal matrix containing n regularization terms. This regularization of the graph Laplacian matrix is important as the unregularized matrix is singular [9].

The term $-\frac{1}{4}\sum_{i,j} w_{ij}(f_i^{(k)} - f_j^{(k)})^2$ imposes a form of feature smoothness, as two nodes i and j connected by a large weight w_{ij} must have similar values of $f^{(k)}$ to maximize the value of the likelihood function $p(f^{(k)}|W)$. As the various features $f^{(k)}$ are assumed to be independent, the likelihood function pertaining to the full feature set D can be defined as

$$p(D|W) = \prod_{k=1}^{m} p(f^{(k)}|W). \tag{4}$$

Assuming each weight w_{ij} is independently drawn from an exponential distribution $p(w_{ij}) \sim \text{Exponential}(\beta)$, the graph weight prior can be defined as

$$p(W) = \prod_{1\leq i<j\leq n} p(w_{ij}) = \prod_{1\leq i<j\leq n} \beta e^{-\beta w_{ij}}. \tag{5}$$

This prior results in sparse graphs by encouraging weights close to zero due to the nature of the exponential distribution.

The MAP estimate of W can then be written as

$$\begin{aligned}\hat{W} &= \underset{W}{\text{argmax}}\ p(D|W)\cdot p(W)\\ &= \underset{W}{\text{argmax}}\ \prod_{k=1}^{m} p(f^{(k)}|W)\cdot \prod_{1\leq i<j\leq n} p(w_{ij}).\end{aligned} \tag{6}$$

Using the result from (3) we have

$$\prod_{k\,=1}^{m} p(f^{(k)}|W) = (2\pi)^{-\frac{nm}{2}}|\tilde{\Delta}^{-1}|^{-\frac{1}{2}m} \cdot \exp\left(\sum_{k=1}^{m}\left(-\frac{1}{4}\sum_{i,j}^{n} w_{ij}\left(f_i^{(k)} - f_j^{(k)}\right)^2 - \frac{1}{2}\sum_{i=1}^{n}\frac{\left(f_i^{(k)}\right)^2}{\sigma_i^2}\right)\right), \tag{7}$$

which gives us

$$\hat{W} = \underset{W}{\text{argmax}}\ (2\pi)^{-\frac{nm}{2}}|\tilde{\Delta}^{-1}|^{-\frac{1}{2}m} \cdot \exp\left(\sum_{k=1}^{m}\left(-\frac{1}{4}\sum_{i,j}^{n} w_{ij}\left(f_i^{(k)} - f_j^{(k)}\right)^2 - \frac{1}{2}\sum_{i=1}^{n}\frac{\left(f_i^{(k)}\right)^2}{\sigma_i^2}\right)\right) \cdot \prod_{1\leq i<j\leq n} \beta e^{-\beta w_{ij}}. \tag{8}$$

Maximizing the expression in (8) is equivalent to maximizing its logarithm, as the logarithm is a monotonically increasing function. Therefore, we can solve

$$\begin{aligned}\hat{W} = \underset{W}{\text{argmax}}\ & -\frac{nm}{2}\log 2\pi - \frac{m}{2}\log|\tilde{\Delta}^{-1}|\\ & -\frac{1}{2}\text{Tr}\left(\tilde{\Delta}DD^T\right) + \frac{n(n-1)}{2}\log\beta - \frac{\beta}{2}||W||_1,\end{aligned} \tag{9}$$

with $||W||_1 = \sum_{i,j=1}^{n}|w_{ij}|$.

As we are only interested in finding the matrix W that maximizes (9), we can ignore constants that do not depend on W, resulting in the final expression

$$\hat{W} = \underset{W}{\text{argmax}}\ \log|\tilde{\Delta}| - \frac{1}{m}\text{Tr}\left(\tilde{\Delta}DD^T\right) - \frac{\beta}{m}||W||_1. \tag{10}$$

This MAP estimation with a Gaussian likelihood function and an exponential prior can also be seen as a Gaussian maximum likelihood estimation with an ℓ_1-norm penalty term similar to the commonly used Lasso regularization [10]. The penalty parameter β controls the trade-off between the likelihood of the data in matrix D and the sparsity inducing penalty term, with larger values of β encouraging sparser graphs in detriment of the likelihood.

The relationship between the regularized graph Laplacian $\tilde{\Delta}$ and the weight matrix W can be written as $\tilde{\Delta} = -W + \text{diag}\,(q_i) + \text{diag}\left(1/\sigma_i{}^2\right)$, where $\text{diag}\,(q_i)$ is a diagonal matrix with each entry being $q_i = \sum_j w_{ij}$. As $\tilde{\Delta}$ can be easily obtained from W, discovering $\hat{W}$ is essentially equivalent to discovering $\hat{\tilde{\Delta}}$ and therefore solving (10) corresponds to estimating the precision or inverse covariance matrix $\tilde{\Delta}$ of the normal distribution $p(f^{(k)}|W) \sim N(0, \tilde{\Delta}^{-1})$. If this distribution describes a Gaussian Markov random field [11] then zero entries in its precision matrix indicate conditional independence between variables - if $\tilde{\Delta}_{ij} = 0$ for $i \neq j$ then $f_i^{(k)}$ and $f_j^{(k)}$ are conditionally independent given all other variables $\left\{f_a^{(k)} : a \neq i \wedge a \neq j\right\}$. Therefore by using the empirical covariance matrix $\frac{1}{m}DD^T$ obtained from m samples of $p(f^{(k)}|W) \sim N(0, \tilde{\Delta}^{-1})$ we can estimate a sparse precision matrix that defines a graphical model representing the dependencies between the different nodes. This graphical model is the graph we are estimating by solving (10), with the absence of an edge connecting two nodes indicating their conditional independence. Therefore this problem can also be seen as the estimation of a sparse Gaussian Markov random field by computing its precision matrix from empirical covariance data.

The final optimization problem (10) can be solved using a generic optimization solver. Different coordinate descent methods [12]–[15] have also been developed to solve optimization problems of this type.

IV. Topology Estimation Methods

The goal is to estimate a graph representing the hidden paths, i.e., the paths not directly observable by the camera network. This corresponds to paths travelled by a person between their exit from a camera's entry/exit node and their entry into another camera's entry/exit node.

In order to estimate the camera network topology two different approaches were developed, based on the graph estimation method detailed in section III. Subsection IV-A describes a pairwise method that identifies correspondences between pairs of events. In subsection IV-B a correspondence-free method that estimates a single graph representing the relations between all observed events is detailed.

A. Pairwise method

In the first method specific trips made by a person between two nodes were identified and the color descriptors of the person at the entry and exit from a camera's field of view were used as corresponding features for each of the two nodes. This is done by considering only entry and exit events, denoted by e_+ and e_- respectively. The algorithm starts by iterating over the various exit events, which are the start points of the hidden trips. For each exit event e_-^i, the corresponding entry event $a(e_-^i)$ that completes the trip must be found. First, the set $\hat{S}_+^i = \left\{ e_+^j : d_{BHATT}(f(e_-^i), f(e_+^j)) < \epsilon_{color},\ t(e_+^j) - t(e_-^i) > 0 \right\}$ containing all the entry events pertaining to what can be considered the same person and taking place after the chosen exit event is found. This is done by selecting only entry events where a distance metric $d_{BHATT}(f(e_-^i), f(e_+^j))$ based on a modified Bhattacharya coefficient [16] between the color descriptors $f(e_-^i)$ and $f(e_+^i)$ is below a certain similarity threshold ϵ_{color}. This accounts for small differences in color, but it is assumed that the color descriptors of the same person in different events will be mostly similar. Then, the end of the trip is chosen as

$$a(e_-^i) = \underset{e_+^j \in \hat{S}_+^i}{\operatorname{argmin}} \quad t(e_+^j) - t(e_-^i), \tag{11}$$

where $t(e^i)$ is the instant when event e^i takes place.

If no entry events are found within the acceptable threshold and $\hat{S}_+^i = \emptyset$, no path is considered and the chosen exit event e_-^i is ignored.

Once the trips are identified, a feature matrix D as described in section III is built for each pair of nodes using the color descriptors associated with the exit and entry events of all trips between the two nodes. As only two nodes are being considered, each feature matrix is a $2 \times pq$ matrix of the form $D = \begin{bmatrix} \mathbf{f_1^1} & \mathbf{f_1^2} & \dots & \mathbf{f_1^p} \\ \mathbf{f_2^1} & \mathbf{f_2^2} & \dots & \mathbf{f_2^p} \end{bmatrix}$, where $\mathbf{f_i^j}$ is a row vector of length q containing the color descriptor associated with the entry/exit event of the jth trip at the ith node of the trip, and p is the total number of trips between the two nodes.

These feature matrices are then used to compute the weight matrix W as described in section III. The optimization problem defined in (10) was solved using CVX, a package for specifying and solving convex programs [17], [18]. This process is memory intensive but can be applied in this case due to the reduced size of the graphs being estimated. As only two nodes are being considered at a time, this corresponds to computing a single weight w_{ij}, which is the weight of the connection between nodes i and j, for each pair of nodes. In effect this means that rather than using the event data to estimate a single graph representing the full camera network, one graph is inferred for each pair of nodes, which consists in determining whether or not the nodes are connected by an edge. The final camera network graph is the union of all the graphs obtained through this method.

While this approach does not make use of all the properties of the sparse inverse covariance estimation due to considering each pair of nodes individually, the sparsity penalty does provide an advantage over using a more typical distance metric to compare the color descriptors associated with two nodes by avoiding the need for further thresholding, as the penalty term $\frac{\beta}{m}||W||_1$ in expression (10) helps eliminate weights connecting dissimilar nodes.

B. Event graph method

Unlike the method described in subsection IV-A, the second approach does not need to establish one-to-one correspondences between events. Rather than identifying events belonging to the same trip, each entry or exit event is used as a node to estimate a graph representing connections between events. The data pertaining to all the observed events is stored in a $2n \times (q+1)$ matrix of the form

$$D = \begin{bmatrix} \mathbf{f}(e_+^1) & \tau t(e_+^1) \\ \mathbf{f}(e_-^1) & \tau t(e_-^1) \\ \mathbf{f}(e_+^2) & \tau t(e_+^2) \\ \mathbf{f}(e_-^2) & \tau t(e_-^2) \\ \dots & \\ \mathbf{f}(e_+^n) & \tau t(e_+^n) \\ \mathbf{f}(e_-^n) & \tau t(e_-^n) \end{bmatrix},$$

where $\mathbf{f}(e_+^i)$ is a row vector of length q containing the color descriptor associated with the ith entry event and $\mathbf{f}(e_-^i)$ is a row vector of length q containing the color descriptor associated with the ith exit event. It is important to note that two events e_+^i and e_-^i are not necessarily corresponding entry and exit events, they are simply the ith entry event to be recorded and the ith exit event to be recorded, respectively. Maintaining this ordering of the matrix rows is not necessary to obtain the event graph, but having a consistent order facilitates the treatment of the resulting data. The instant when each event takes place is used as an additional feature. These time values are multiplied by a parameter τ in order to control the trade-off between the weights of the color descriptors and the moments when the events take place, with larger values of τ placing a larger emphasis on the temporal factor over the color descriptors in the trade-off. Two corresponding entry and exit events (events belonging to the same trip) should have similar color descriptors, and it is considered that most trips should be relatively short. Therefore, the value of τ should be such that among events with similar color descriptors those that are closer in time should be more likely to be connected in the final graph but with the color similarity being the more important factor.

The feature matrix D is then used to estimate a weight matrix W^{ev} representing the graph of all events $G^{ev} = (V(G^{ev}), E(G^{ev}))$ as described in section III using the R package *dpglasso* [19], [13], as, unlike in subsection IV-A the size of the graph being estimated makes using CVX impossible due to memory constraints. In the resulting graph each vertex represents an event, with an edge between two vertices indicating that both events form a trip. In order to obtain the final graph $G^{net} = (V(G^{net}), E(G^{net}))$ representing the camera network, a new matrix W^{net} is initialized with every entry w^{net}_{ij} set to 0. For each edge $\{u, v\} \in E(G^{ev})$ the entry/exit zones z_a and z_b associated with the events represented by vertices u and v are identified, and the absolute value of the weight w^{ev}_{uv} of edge $\{u, v\}$ is added to entry w^{net}_{ab} of matrix W^{net}. The edge set of the final camera network graph is given by $E(G^{net}) = \left\{\{u, v\} \in (V(G^{net}))^2 : u \neq v, w^{net}_{uv} > \epsilon_{thr}\right\}$. The threshold ϵ_{thr} is important to remove false positives from the final graph. As the original graph being estimated is a graph of single events, a single false positive can result in a connection between unrelated events. It is assumed that for large enough data sets real paths will be travelled multiple times and large numbers of false positives are unlikely to occur in the same path. As the final weight values in matrix W^{net} are the sums of every edge weight associated with each possible path, applying the threshold ϵ_{thr} allows us to exclude most false positives while not affecting real paths.

Unlike the pairwise method described in subsection IV-A, this method makes no distinction between visible and hidden paths. As the goal is to find the hidden paths, after obtaining the full network graph we simply consider only connections between entry/exit zones belonging to different cameras. This ensures that visible paths are not included in the final graph.

V. Results

To test the performance of the developed methods multiple tests were run with different sets of simulated data. This data was generated with an event simulator originally developed for [20], adapted to generate Black-Value-Tint histograms [21] as the color descriptors associated with the different events. This simulator generates data according to the transition probabilities and temporal distributions obtained from the HDA+ data set [22]. The use of the simulator was necessary to generate additional data as the real camera footage data from the HDA+ data set contains a relatively small number of events for which favorable results could not be obtained.

The number of generated events can be controlled by adjusting two simulation parameters: t_{sim} defines the total run time of the simulation in seconds and p_{new} is the probability of a new person being added to the simulation in each second. It is also possible to add uniformly distributed random noise to the simulated color descriptors with a variable maximum amplitude defined by r_{max}. While different simulations can have different trajectories, they are based on the same environment and therefore the possible paths are the same for every test. The total number of existing edges in the ideal graph is 15.

We want to maximize the number of correctly detected edges while minimizing the number of false positives and missing connections. To evaluate these results we rely on precision and recall metrics [23] and we use the F-measure to optimize the trade-off between the two by computing the harmonic mean of the two ratios [24]. We also present the total number of edges included in each estimated graph to evaluate the sparsity of the obtained estimates.

A. Pairwise method

In this subsection the results obtained through the approach detailed in subsection IV-A are presented. The first test had no added noise, and the test parameters were $t_{sim} = 10800\ s$ and $p_{new} = 0.01$. The values of β used for this test were $\beta = \{0, 1000, 1500, 2000, 3000\}$. The results of this test can be seen in table I.

β	Number of detected edges	Precision	Recall	F-measure
0	15	1	1	1
1000	15	1	1	1
1500	15	1	1	1
2000	15	1	1	1
3000	12	1	0.8	0.889

TABLE I
Results of the pairwise method for $t_{sim} = 10800\ s$ without added noise.

The ideal simulated data without added noise results in perfect estimations for most lower values of the sparsity coefficient β. For higher values of β, the sparsity encouraging term grows large enough to exclude correct connections from the estimated graph, resulting in a lower recall value. The precision was 1 for all tested values of β as no false positive connections are identified. As there is no random noise added to the color descriptors, all events associated with the same person have the exact same descriptor. This means that in order for a false positive connection to be obtained, the random color descriptors generated for two different people must be within the chosen similarity threshold of $\epsilon_{color} = 0.1$, which is highly unlikely.

The results obtained for a second simulation with $r_{max} = 2 \times 10^{-3}$ using $t_{sim} = 86400\ s$, $p_{new} = 0.01$ and $\beta = \{0, 3500, 7000, 10000, 20000\}$ are presented in table II.

β	Number of detected edges	Precision	Recall	F-measure
0	84	0.179	1	0.303
3500	21	0.714	1	0.833
7000	15	1	1	1
10000	13	1	0.867	0.929
20000	10	1	0.667	0.8

TABLE II
Results of the pairwise method for $t_{sim} = 86400\ s$.

Unlike the results obtained with the noiseless data, in this test the precision value is no longer 1 for every value of β. This happens because as each instance of a person's color descriptor is altered with significant random additive noise, it is possible for the noisy data to result in false positives. As the sparsity coefficient β increases, the number of edges in the estimated graph decreases. Lower values of β result in low precision and high recall, with the opposite occurring for higher values of β. For $\beta = 7000$ it was possible to estimate a graph that perfectly represents all the hidden paths, as indicated by the fact that both precision and recall have a value of 1. This means that there are no false positives or false negatives - all existing paths are identified in the estimate, and no false paths are incorrectly identified. Figure 2 shows the resulting estimated graph for this value of β, which corresponds to the ideal graph of hidden paths.

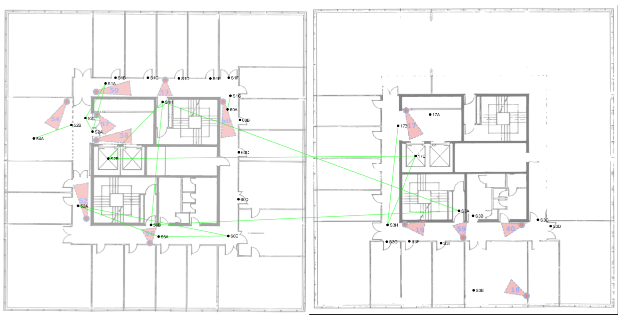

Fig. 2. Estimated graph of hidden paths obtained with the pairwise method for $t_{sim} = 86400$ s and $\beta = 7000$. Green lines represent correctly identified paths.

A third test was run with the parameters $p_{new} = 0.01$, $t_{sim} = 10800$ s and $r_{max} = 2 \times 10^{-3}$. The results for this test for $\beta = \{0, 1000, 1500, 2000, 3000\}$ can be seen in table III.

β	Number of detected edges	Precision	Recall	F-measure
0	38	0.395	1	0.566
1000	19	0.79	1	0.882
1500	15	0.933	0.933	0.933
2000	11	1	0.733	0.846
3000	9	1	0.6	0.75

TABLE III
RESULTS OF THE PAIRWISE METHOD FOR $t_{sim} = 10800$ s.

As with the previous test, increasing the value of the sparsity coefficient β decreases the number of edges in the estimated graph as well as the recall value, while increasing the precision. However, due to the smaller amount of data available for the estimation of the graph, in this test it was not possible to obtain a completely correct estimation, i.e. a graph with both precision and recall values of 1. In this case the results represent a trade-off between both metrics, with lower values of β resulting in a recall of 1 but lower precision values, higher values of β resulting in lower recall but a precision of 1, and $\beta = 1500$ resulting in a high value for both metrics but with neither of them being perfect. Figure 3 shows the resulting estimated graph for $\beta = 1500$, where both one false positive and one false negative can be seen.

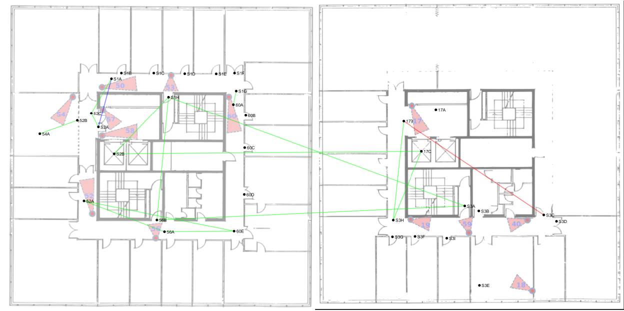

Fig. 3. Estimated graph of hidden paths obtained with the pairwise method for $t_{sim} = 10800$ s and $\beta = 1500$. Green lines represent correctly identified paths, blue lines represent false negatives and red lines represent false positives.

B. Event graph method

In this subsection the results obtained through the approach detailed in subsection IV-B are presented.

An initial test was run with noiseless data with a test duration of $t_{sim} = 10800$ s and $p_{new} = 0.01$ for $\beta = \{0, 0.025, 0.075, 0.125, 0.2\}$. The value used for the time coefficient was $\tau = 2 \times 10^{-4}$. The results of this test can be seen in table IV.

β	Number of detected edges	Precision	Recall	F-measure
0	175	0.086	1	0.158
0.025	29	0.517	1	0.681
0.075	14	1	0.933	0.966
0.125	13	1	0.867	0.929
0.2	10	1	0.667	0.8

TABLE IV
RESULTS OF THE EVENT GRAPH METHOD FOR $t_{sim} = 10800$ s WITHOUT ADDED NOISE.

For the lower values of β, the resulting graph is less sparse than those resulting from the pairwise method. This happens because while the pairwise method utilized a thresholding based on the color descriptors associated with each event prior to applying the graph estimation algorithm, this method relies only on the sparsity penalty to eliminate connections in the resulting graph. This means that when using $\beta = 0$ every possible hidden path is included in the estimated graph, resulting in an extremely low precision value. The graph obtained for this value of β simply corresponds to all paths between two entry/exit zones in different cameras.

It is also worth nothing that despite the fact that the color descriptors used for this test had no added noise, it was not possible to obtain a perfect estimate. This is due to the usage of the time of the event as an additional feature, which can result in real paths that take a particularly long time to walk being excluded from the estimate. Figure 4 shows the estimated graph obtained for $\beta = 0.075$, which was the best result obtained as it includes no false positives while missing only a single connection. This missing connection is a path between two different floors, which takes longer to walk than other paths in the network.

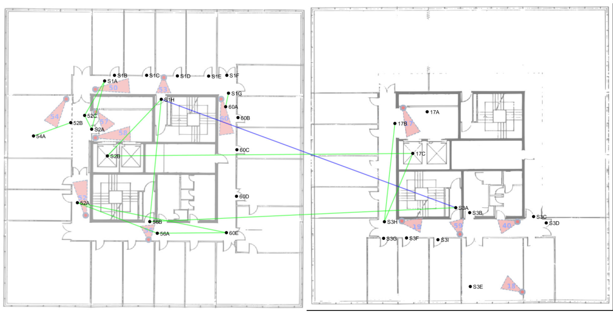

Fig. 4. Estimated graph of hidden paths obtained with the event graph method for $t_{sim} = 10800\ s$ and $\beta = 0.075$ without added noise. Green lines represent correctly identified paths and blue lines represent false negatives.

A second test was run, this time with uniform additive noise of amplitude $r = 2 \times 10^{-3}$ being added to the color descriptors. The test parameters used were $t_{sim} = 10800\ s$ and $p_{new} = 0.01$ and the test was run for $\beta = \{0, 0.025, 0.075, 0.125, 0.2\}$, once again with $\tau = 2 \times 10^{-4}$. The results of this test can be seen in table V.

β	Number of detected edges	Precision	Recall	F-measure
0	175	0.086	1	0.158
0.025	35	0.429	1	0.6
0.075	14	1	0.933	0.966
0.125	13	1	0.867	0.929
0.2	8	1	0.533	0.696

TABLE V
RESULTS OF THE EVENT GRAPH METHOD FOR $t_{sim} = 10800\ s$.

For $\beta = 0$ we obtain a graph containing every possible invisible path again. As before, increasing the value of β increases the sparsity of the obtained graphs while also increasing the precision and decreasing the recall value of the estimate. For $\beta = 0.075$ it was possible to achieve a precision of 1 and a recall of 0.933. In addition to having the same precision and recall values as the result obtained from the noiseless data, this constitutes an improvement over the results obtained with the pairwise method. The graph obtained for this value of β is represented in figure 5, where it can be seen only one existing path was not included in the estimated graph.

As with the previous test using noiseless data, the path that the algorithm failed to identify connects two different floors of the testing environment. As the moment when each event takes place is used as a feature in this method, paths that take a long time to complete (such as the path in question) are less likely to be included in the estimated graph, as the difference between the temporal features in two events that form a trip between two particularly distant nodes will be large. Reducing the value of the time coefficient τ can help include these paths in the estimated graph, but it also makes it more likely that incorrect paths will be included.

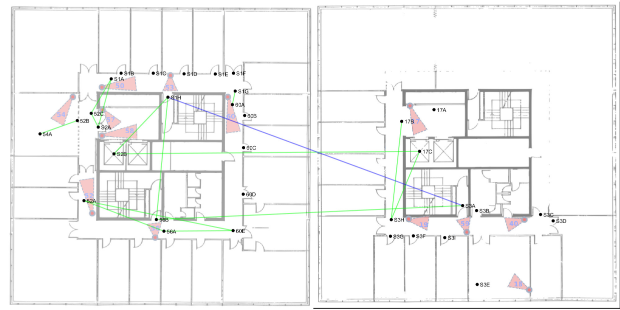

Fig. 5. Estimated graph of hidden paths obtained with the event graph method for $t_{sim} = 10800\ s$ and $\beta = 0.075$. Green lines represent correctly identified paths and blue lines represent false negatives.

In order to evaluate the effect of the time coefficient τ another set of tests was run with $\tau = 1 \times 10^{-4}$ for $t_{sim} = 10800\ s$ and $\beta = \{0, 0.025, 0.075, 0.125, 0.2\}$ with $r = 2 \times 10^{-3}$. The results of these tests can be seen in table VI.

β	Number of detected edges	Precision	Recall	F-measure
0	175	0.086	1	0.158
0.025	35	0.429	1	0.6
0.075	15	0.933	0.933	0.933
0.125	13	1	0.867	0.929
0.2	8	1	0.533	0.696

TABLE VI
RESULTS OF THE EVENT GRAPH METHOD FOR $t_{sim} = 10800\ s$ AND $\tau = 1 \times 10^{-4}$.

While the results are mostly identical to those obtained with $\tau = 2 \times 10^{-4}$, for $\beta = 0.075$ the precision value is lower. The graph obtained for this value of β can be seen in figure 6.

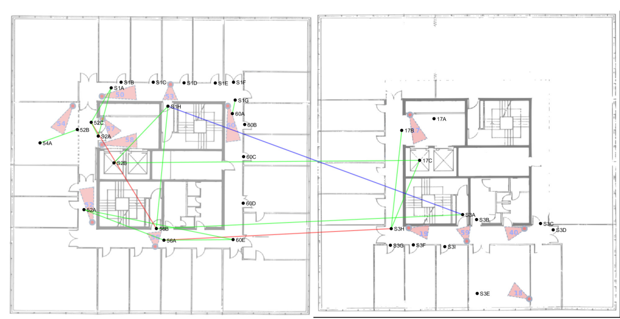

Fig. 6. Estimated graph of hidden paths obtained with the event graph method for $t_{sim} = 10800\ s$, $\tau = 1 \times 10^{-4}$ and $\beta = 0.075$. Green lines represent correctly identified paths, red lines represent false positives and blue lines represent false negatives.

In addition to the fact that the missing path between the two floors is still not included in the estimated graph, there are two new incorrect paths in the estimate. This means that even relaxing the time coefficient τ enough to match events that do not directly follow one another is still not enough to include paths with such a long travel time as the one connecting the two opposing stairwells on different floors.

These results demonstrate that while the event graph method can not be as easily used with large data sets as the pairwise method, for smaller amounts of data it can obtain more precise results. Moreover, while the pairwise method showed a better performance with the ideal noiseless color descriptors due to not using temporal features, the event graph method is more robust to noisy data, being able to obtain the same estimated graph when using noiseless data as well as color descriptors with uniform additive noise of amplitude $r = 2 \times 10^{-3}$.

VI. Conclusions and future work

In this work two automated methods for camera network topology estimation were tested with different sets of simulated data. The graphs obtained during the testing were compared with the ideal graphs containing all the off-camera trajectories within the network, and the test results were evaluated based on precision and recall metrics. For each test multiple values of the sparsity coefficient β were tested with the goal of finding the values of β that result in the best estimated graph. In general there is a trade-off between the two metrics we want to maximize so we used the F-measure to select the best value of β for each test.

We were able to obtain favorable results with both methods, with the event graph method demonstrating more robustness to noise in the color descriptor data but performing worse than the pairwise method with ideal data due to the effect of using the event time as a feature. This suggests that the pairwise method may be a better option for data with lower noise values, while the event graph method can be used when there is more noise in the data.

For the pairwise method, data sets with different numbers of events were used for testing. The best values of β depend on the number of events, with larger data sets requiring larger values of the sparsity coefficient to obtain similar results. As future work it would be important to test this method with data from a different environment to evaluate how the same values of β perform with different data sets. While we used different simulations with different parameters and different randomly generated trajectories, they were all based on the real data from the HDA+ data set and shared the same transition probabilities and temporal distributions, which means that the method was not tested with varied data containing different types of trajectories in different environments.

The event graph method was only tested with smaller sets of data due to its greater computational complexity making it impractical to use larger simulations for testing. The best results were consistently obtained for the same value of β across all the different tests, although as with the pairwise method it would be useful to use data from a different environment for future testing.

Both of the methods perform more poorly with smaller data sets, requiring a sizeable amount of data to provide favorable results. As the real data from the HDA+ data set amounts to a relatively small number of events it was impossible to obtain adequate estimates of the camera network topology from the real data. Testing with a larger amount of real data would be an important next step for future testing, as all the testing done so far used simulated color descriptors with additive uniform noise which do not properly reproduce all the challenges of using real color descriptors obtained from camera images.

References

[1] 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, Part C (Applications and Reviews)*, vol. 34, no. 3, pp. 334–352, 2004.

[2] X. Wang, "Intelligent multi-camera video surveillance: A review," *Pattern recognition letters*, vol. 34, no. 1, pp. 3–19, 2013.

[3] D. Kamenetsky, "Camera network topology discovery literature review," Defence Science and Technology Organisation, Tech. Rep. DSTO-GD-0667, 2011.

[4] X. Li, W. Dong, F. Chang, and P. Qu, "Topology learning of non-overlapping multi-camera network," *International Journal of Signal Processing, Image Processing and Pattern Recognition*, vol. 8, no. 11, pp. 243–254, 2015.

[5] D. Makris, T. Ellis, and J. Black, "Bridging the gaps between cameras," in *Proceedings of the 2004 IEEE computer society conference on Computer vision and pattern recognition*. IEEE Computer Society, 2004, pp. 205–210.

[6] Y.-J. Cho, J.-H. Park, S.-A. Kim, K. Lee, and K.-J. Yoon, "Unified framework for automated person re-identification and camera network topology inference in camera networks," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 2601–2607.

[7] F. Chung and R. M. Richardson, "Weighted laplacians and the sigma function of a graph," *Contemporary Mathematics*, vol. 415, pp. 93–107, 2006.

[8] B. Lake and J. Tenenbaum, "Discovering structure by learning sparse graphs," in *Proceedings of the Annual Meeting of the Cognitive Science Society, 32*, 2010, pp. 778–783.

[9] X. Zhu, J. Lafferty, and Z. Ghahramani, "Semi-supervised learning: From gaussian fields to gaussian processes," Carnegie Mellon University, Tech. Rep. CMU-CS-03-17, 2013.

[10] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.

[11] H. Rue and L. Held, *Gaussian Markov random fields: theory and applications*. CRC press, 2005.

[12] J. Friedman, T. Hastie, and R. Tibshirani, "Sparse inverse covariance estimation with the graphical lasso," *Biostatistics*, vol. 9, no. 3, pp. 432–441, 2008.

[13] R. Mazumder and T. Hastie, "The graphical lasso: New insights and alternatives," *Electronic journal of statistics*, vol. 6, p. 2125, 2012.

[14] J. Duchi, S. Gould, and D. Koller, "Projected subgradient methods for learning sparse gaussians," in *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 2008, pp. 153–160.

[15] O. Banerjee, L. El Ghaoui, and A. dAspremont, "Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data," *Journal of Machine Learning Research*, vol. 9, pp. 485–516, 2008.

[16] D. Comaniciu, V. Ramesh, and P. Meer, "Real-time tracking of non-rigid objects using mean shift," in *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on*, vol. 2. IEEE, 2000, pp. 142–149.

[17] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.1," http://cvxr.com/cvx, Mar. 2014.

[18] ——, "Graph implementations for nonsmooth convex programs," in *Recent Advances in Learning and Control*, ser. Lecture Notes in Control and Information Sciences, V. Blondel, S. Boyd, and H. Kimura, Eds. Springer-Verlag Limited, 2008, pp. 95–110, http://stanford.edu/~boyd/graph_dcp.html.

[19] R. Mazumder and T. Hastie, *dpglasso: Primal Graphical Lasso*, 2012, r package version 1.0. [Online]. Available: https://CRAN.R-project.org/package=dpglasso

[20] I. Silva, "Estimação de topologia de uma rede de câmaras usando métodos esparsos (in Portuguese)," Master's thesis, Instituto Superior Técnico, 2017.

[21] D. S. Cheng, M. Cristani, M. Stoppa, L. Bazzani, and V. Murino, "Custom pictorial structures for re-identification," in *Proceedings of the British Machine Vision Conference*. BMVA Press, 2011, pp. 68.1–68.11.

[22] D. Figueira, M. Taiana, A. Nambiar, J. Nascimento, and A. Bernardino, "The hda+ data set for research on fully automated re-identification systems," in *European Conference on Computer Vision*. Springer, 2014, pp. 241–255.

[23] M. Buckland and F. Gey, "The relationship between recall and precision," *Journal of the American society for information science*, vol. 45, no. 1, pp. 12–19, 1994.

[24] N. Chinchor, "Muc-4 evaluation metrics," in *Proceedings of the 4th conference on Message understanding*. Association for Computational Linguistics, 1992, pp. 22–29.