A Survey for Long-Term Multi-Person Tracking for Surveillance Video

Jacob Selvage, Carlos Humphries, Floyd Mcclanahan, Anthony Rhodarmer, and Arthur Gosman

Abstract—Multi-person tracking plays a critical role in the analysis of surveillance video. However, most existing work focus on shorter-term (e.g. minute-long or hour-long) video sequences. Therefore, we propose a multi-person tracking algorithm for very long-term (e.g. month-long) multi-camera surveillance scenarios. Long-term tracking is challenging because 1) the apparel/appearance of the same person will vary greatly over multiple days and 2) a person will leave and re-enter the scene numerous times. To tackle these challenges, we leverage face recognition information, which is robust to apparel change, to automatically reinitialize our tracker over multiple days of recordings. Unfortunately, recognized faces are unavailable oftentimes. Therefore, our tracker propagates identity information to frames without recognized faces by uncovering the appearance and spatial manifold formed by person detections. We tested our algorithm on a 23-day 15-camera data set (4,935 hours total), and we were able to localize a person 53.2% of the time with 69.8% precision. We further performed video summarization experiments based on our tracking output. Results on 116.25 hours of video showed that we were able to generate a reasonable visual diary (i.e. a summary of what a person did) for different people, thus potentially opening the door to automatic summarization of the vast amount of surveillance video generated every day.

Index Terms—Multi-Object Tracking, Nonnegative Matrix Optimization, Surveillance Video Summarization, Face Recognition

1 INTRODUCTION

SURVEILLANCE cameras have been widely deployed to enhance safety in our everyday lives. The recorded footage can further be used to analyze long term trends in the environment. Unfortunately, manual analysis of large amounts of surveillance video is very difficult, thus motivating the development of computational analysis of surveillance video. A common first step of computational analysis is to track each person in the scene, which has led to the development of many multi-object tracking algorithms [1], [2], [3]. However, two important points are largely neglected in the literature: 1) the usage of identity information such as face recognition or any other cue that can identify an individual, and 2) the exploration of real-world applications based on tracking output from hundreds or thousands of hours of surveillance video.

There are two main advantages of utilizing face recognition information for tracking. First, face recognition empowers the tracker to relate a tracked person to a real-world living individual, thus enabling individual-specific activity analysis. Second, face recognition is robust to appearance/apparel change, thus making it well-suited for tracker reinitialization in very long-term (e.g. month-long) surveillance scenarios.

We propose an identity-aware tracking algorithm as follows. Under the tracking-by-detection framework [4], the tracking task can be viewed as assigning each person detection to a specific individual/label. Face recognition output can be viewed as label information. However, as face recognition is only available in a few frames, we propagate face recognition labels to other frames using a manifold learning approach, which captures the appearance similarities and spatial-temporal layout of person detections. The manifold learning approach is formulated as a constrained quadratic optimization problem and solved with nonnegative matrix optimization techniques.

The constraints included are the mutual exclusion and spatial locality constraints which constrain the final solution to deliver a reasonable multi-person tracking output.

We performed tracking experiments on challenging data sets, including a 4,935 hour complex indoor tracking data set. Our long-term tracking experiments showed that our method was effective in localizing and tracking each individual in thousands of hours of surveillance video. An example output of our algorithm is shown in Figure 1, which shows the location of each identified person on the map in the middle of the image. This is analogous to the *Marauder's Map* described in the Harry Potter book series [5].

To explore the utility of long-term multi-person tracking, we performed *summarization-by-tracking* experiments to acquire the *visual diary* of a person. Visual diaries provide a person-specific summary of surveillance video by showing snapshots and textual descriptions of the activities performed by the person. An example visual diary of a nursing home resident is shown in Figure 2. Experiments conducted on 116.25 hours of video show that we were able to summarize surveillance video with reasonable accuracy, which further demonstrates the effectiveness of our tracker.

In sum, the main contributions of this paper are as follows:

- 1) We propose an identity-aware multi-object tracking algorithm. Our tracking algorithm leverages identity information which is utilized as sparse label information in a manifold learning framework. The algorithm is formulated as a constrained quadratic optimization problem and solved with nonnegative matrix optimization.
- 2) A 15-camera multi-object tracking data set consisting of 4,935 hours of nursing home surveillance video was annotated. This real-world data set enables us to perform very long-term tracking experiments to better assess the performance and applicability of multi-object trackers.

Fig. 1: The Marauder’s Map for a nursing home (*Caremedia Short* sequence [6]) with the map in the middle. Dots on the map show the locations of different people. The surrounding images are the views from each surveillance camera. White lines correspond to the field-of-view of each camera.

- 3) Video summarization experiments based on tracking output were performed on 116.25 hours of video. We demonstrate that the visual diaries generated from tracking-based summarization can effectively summarize hundreds of hours of surveillance video.

2 RELATED WORK

As multi-object tracking is a very diverse field, we only review work that follows the very popular tracking-by-detection paradigm [4], which is also used in our work. For a more comprehensive and detailed survey we refer the readers to [7].

The tracking-by-detection paradigm has four main components: object localization, appearance modeling, motion modeling and data association. The object localization component generates a set of object location hypotheses for each frame. The localization hypotheses are usually noisy and contain false alarms and misdetections, so the task of the data association component is to robustly group the location hypotheses which belong to the same physical object to form many different object trajectories. The suitability of the grouping can be scored according to the coherence of the object’s appearance and the smoothness of the object’s motion, which correspond to appearance modeling and motion modeling respectively. We now describe the four components in more detail.

2.1 Object Localization

There are mainly three methods to find location hypotheses: using background subtraction, using object detectors, and connecting single-frame detection results into tracklets. The Probabilistic Occupancy Map (POM, [8]) combines background subtraction information from multiple cameras to jointly locate multiple objects in a single frame. Utilizing object detector output is one of the most common ways to localize tracking targets [4], [9], [6], [10], [1], [11], [3], [12], [2], [13]. The object detector is run on each frame of the video, and the detection results serve as the location hypotheses for subsequent processing. Localized objects in each frame could be

Fig. 2: An example visual diary for an elderly resident in a nursing home. The automatically generated textual description and snapshots are shown for the two events. Long-term statistics are also shown.

connected to create *tracklets* [14], [15], [16], [17], [18], [19], [20], which are short tracks belonging to the same physical object. Tracklets are usually formed in a very conservative way to avoid connecting two physically different objects.

2.2 Appearance Models

Appearance models discriminate between detections belonging to the same physical object and other objects. Color histograms [21], [22], [1], [14], [18], [23], [9], [6], [20], [67] have been widely used to represent the appearance of objects, and the similarity of the histograms is often computed with the Bhattacharyya distance [1], [23]. Other features such as Histogram of Oriented Gradients [24] have also been used [15], [16].

Appearance models can also be learned from tracklets. The main assumption of tracklets is that all detections in a tracklet belong to the same object, and [15], [17], [19], [25], [26], [16] exploit this assumption to learn more discriminative appearance models. Note that the “identity” in our work is different from [16], which utilized person re-identification techniques to improve the appearance model. We, however, focus on the “real-world identity” of the person, which is acquired from face-recognition.

Appearance models based on incremental manifold/subspace learning has also been utilized in previous work [27], [28], [29] to learn subspaces for appearance features that can better differentiate tracked targets and background in single or multi-object tracking. However, [28] utilized multiple *independent* particle filters, which may have the issue of one particle filter “hijacking” the tracking target of another particle filter [30], [31]. Our method alleviates this issue as we *jointly* optimize for all trajectories to acquire a more reasonable set of trajectories.

2.3 Motion Models

Objects move in a smooth manner, and motion models can capture this assumption to better track objects. [1], [32], [10],

[6], [20] use the bounded velocity model to model motion, i.e. an object cannot move faster than a given velocity. [22], [9], [13] improve upon this by modeling motion with the constant velocity model, which is able to model acceleration. Higher order methods such as spline-based methods [2], [3] and the Hankel matrix [33] can model even more sophisticated motions. [19] assumes that different objects in the same scene move in similar but potentially non-linear ways, and the motion of highly confident tracklets can be used to infer the motion of non-confident tracklets.

2.4 Data Association

A data association algorithm takes the object location hypotheses, appearance model and motion model as input and finds a disjoint grouping of the object location hypotheses which best describes the motion of objects in the scene. Intuitively, the algorithm will decide whether to place two object location hypotheses in the same group based on their *affinity*, which is computed from the appearance and motion models.

The Hungarian algorithm and the network flow are two popular formulations. Given the pair-wise affinities, the Hungarian algorithm can find the optimal matching between two sets of object location hypotheses in polynomial time [14], [18], [16], [15], [2]. In the network flow formulation [1], [10], [32], [34], each path from source to sink corresponds to the trajectory of an object.

Many trackers have been formulated as a general Integer Linear Programming (ILP) problem. [21], [20], [23] solved the ILP by first relaxing the integral constraints to continuous constraints and then optimizing a Linear Program. [35], [36] formulated tracking as clique partitioning, which can also be formulated as an ILP problem and solved by a heuristic clique merging method.

More complex data association methods have also been used, including continuous energy minimization [11], discrete-continuous optimization [3], Block-ICM [2], conditional random fields [17], [12], generalized minimum clique [9] and quadratic programming [22], [37].

However, it is non-trivial to incorporate identity information such as face recognition into the aforementioned methods. One quick fix may be to assign identities to trajectories after the trajectories have been computed. However, problems occur if there are identity-switches in a single trajectory. Another method proposed by [8] utilized the Viterbi algorithm to find a trajectory which passes through all the identity observations of each person. However, Viterbi search cannot be performed simultaneously over all individuals, and [8] had to performed Viterbi search sequentially, i.e. one individual after another. This greedy approach lead to “hijacking” of another person’s trajectory [8], which is not ideal. Therefore, to achieve effective identity-aware tracking, it is ideal to design a data association framework which can directly incorporate identity information into the optimization process.

Identity-Aware Data Association

Previously proposed data association methods [20], [38], [39], [40] and [6] utilized identity information for tracking. There

have been other work which utilized transcripts from TV shows to perform face recognition and identity-aware face tracking [41], [42], but this is not the main focus of our paper.

[20], [38] formulated identity-aware tracking as an ILP and utilized person identification information from numbers written on an athlete’s jersey or from face recognition. [20], [38] utilized a global appearance term as their appearance model to assign identities to detections. However, the global term assumes a fixed appearance template for an object, which may not be applicable in long surveillance recordings as the appearance of the same person may change.

[39] utilized a few manually labeled training examples and play-by-play text in a Conditional Random Field formulation to accurately track and identify sports players. However, this method may not work as well in surveillance domains where play-by-play text is not available.

[40] utilized online structured learning to learn a target-specific appearance model, which is used in a network flow framework. However, [40] utilized densely-sampled windows instead of person bounding boxes as input, which may be too time-consuming to compute in long videos.

[6] utilized face-recognition as sparse label information in a semi-supervised tracking framework. However, [6] does not incorporate the spatial locality constraint into the optimization step, which might lead to solutions showing a person being at multiple places at the same time. This becomes very severe in crowded scenes. Also, the method needs a Viterbi search to compute the final trajectories. The Viterbi search requires the start and end locations of all trajectories, which is an unrealistically restrictive assumption for long-term tracking scenarios. In this paper, we enhance this tracker by adding the spatial-locality constraint term, which enables tracking in crowded scenes and also removes the need for the start and end locations of a trajectory.

3 METHODOLOGY

Tracking-by-detection-based multi-object tracking can be viewed as a constrained clustering problem as shown in Figure 3. Each location hypothesis, which is a person detection result, can be viewed as a point in the spatial-temporal space, and our goal is to group the points so that the points in the same cluster belong to a single trajectory. A trajectory should follow the mutual exclusion constraint and spatial-locality constraint, which are defined as follows.

- **Mutual Exclusion Constraint:** a person detection result can only belong to at most one trajectory.
- **Spatial-Locality Constraint:** two person detection results belonging to a single trajectory should be reachable with reasonable velocity, i.e. a person cannot be in two places at the same time.

Sparse label information acquired from sources such as face recognition can be used to assign real-world identities and also enhance tracking performance.

Our tracking algorithm has three main steps.

- 1) **Manifold construction based on appearance and spatial affinity:** The appearance and spatial affinity respectively assumes that 1) similar looking person detections are

(a) Input to tracking algorithm: location and appearance of person detection plus recognized faces for some person detections. (b) Output of tracking algorithm: partitioning of the person detections into different trajectories.

Fig. 3: Illustration of the input and output of our tracking algorithm. Each person detection is a point in the (x, y, t) space. The z axis is irrelevant in this case, because the figures are drawn based on the person detections from the *terracel* data set [8], where people walk on a common ground plane. Best viewed in color.

likely to be of the same individual and 2) person detections which are spatially and temporally very close to each other are also likely to be of the same individual.

- 2) Spatial locality constraint: This constraint encodes the fact that a person cannot be at multiple places at the same time. In contrast to the manifold created in the previous step which encodes the *affinity* of two person detections, this constraint encodes the *repulsion* of two person detections.
- 3) Constrained nonnegative optimization: Our nonnegative optimization method acquires a solution which simultaneously satisfies the manifold assumption, the mutual exclusion constraint and the spatial-locality constraint.

In the following sections, we first define our notations, then the 3 aforementioned steps are detailed.

3.1 Notations

In this paper, given a matrix $\mathbf{A}$, let $\mathbf{A}_{ij}$ denote the element on the i -th row and j -th column of $\mathbf{A}$. Let $\mathbf{A}_i$ denote the i -th row of $\mathbf{A}$. $\text{Tr}(\cdot)$ denotes the trace operator. $|\cdot|_F$ is the Frobenius norm of a matrix. Given an positive integer m , $\mathbf{1}_m \in \mathbb{R}^m$ is a column vector with all ones.

Hereafter, we call a person detection result an *observation*. Suppose the person detector detects n observations. Let c be the number of tracked individuals, which can be determined by either a pre-defined gallery of faces or the number of unique individuals identified by the face recognition algorithm. Our task is to assign a class label to each observation. Let $\mathbf{F} \in \{0, 1\}^{n \times c}$ be the label assignment matrix for all observations. Without loss of generality, $\mathbf{F}$ is reorganized such that the observations from the same class are located in consecutive rows, i.e. the j -th column of $\mathbf{F}$ is given by:

$$\mathbf{F}_{*j} = [\underbrace{0, \dots, 0}_{\sum_{i=1}^{j-1} m^{(i)}}, \underbrace{1, \dots, 1}_{m^{(j)}}, \underbrace{0, \dots, 0}_{\sum_{i=j+1}^c m^{(i)}}]^T, \quad (1)$$

where $m^{(j)}$ is the number of observations in the j -th class. If the p -th element in $\mathbf{F}_{*j}$, i.e. $\mathbf{F}_{pj}$, is 1, it indicates that the

p -th observation corresponds to the j -th person. According to Equation 1, it can be verified that

$$\mathbf{F}^T \mathbf{F} = \begin{bmatrix} \mathbf{F}_{*1}^T \\ \vdots \\ \mathbf{F}_{*c}^T \end{bmatrix} [\mathbf{F}_{*1} \quad \dots \quad \mathbf{F}_{*c}] = \text{diag} \left(\begin{bmatrix} m^{(1)} \\ \vdots \\ m^{(c)} \end{bmatrix} \right) = \mathbf{J}. \quad (2)$$

The i -th observation is described by a d dimensional color histogram $\mathbf{x}^{(i)} \in \mathbb{R}^d$, frame number $t^{(i)}$, and 3D location $\mathbf{p}^{(i)} \in \mathbb{R}^3$ which corresponds to the 3D location of the bottom center of the bounding box. In most cases, people walk on the ground plane, and the z component becomes irrelevant. However, our method is not constrained to only tracking people on the ground plane.

3.2 Manifold Construction based on Appearance and Spatial Affinity

There are two aspects we would like to capture with manifold learning: 1) appearance affinity and 2) spatial affinity, which we will detail in the following sections.

3.2.1 Modeling Appearance Affinity

Based on the assumption that two observations with similar appearance are likely to belong to the same individual, we build the manifold structure by finding nearest neighbors for each observation. Observation j is qualified to be a nearest neighbor of observation i if 1) j is reachable with reasonable velocity, i.e. $v^{(ij)} \leq V$, 2) i and j should not be too far apart in time, i.e. $|t^{(i)} - t^{(j)}| \leq T$, and 3) both observations should look similar, i.e. the similarity of color histograms $\mathbf{x}^{(i)}$ and $\mathbf{x}^{(j)}$ should be larger than a threshold τ . We define $v^{(ij)} = \frac{\max(|\mathbf{p}^{(i)} - \mathbf{p}^{(j)}|_2 - \delta, 0)}{|t^{(i)} - t^{(j)}| + \epsilon}$. ϵ is a small number to avoid division by zero. δ models the maximum localization error of the same person between different cameras due to calibration and person detection errors. V is the maximum velocity a person can achieve. T limits how far we look for nearest neighbors in the time axis. The similarity between two histograms is computed with the exponential- χ^2

Fig. 4: Intuition of appearance-based nearest neighbor selection. The nearest neighbors for the red person detection in frame 50 are shown. No nearest neighbors are found in frames 75 and 100 as the person is occluded. Nevertheless, once the person is no longer occluded, the nearest neighbor connections could be made again, thus overcoming this occlusion.

metric: $\chi^2(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) = \exp\left(-\frac{1}{2} \sum_{l=1}^d \frac{(\mathbf{x}_l^{(i)} - \mathbf{x}_l^{(j)})^2}{\mathbf{x}_l^{(i)} + \mathbf{x}_l^{(j)}}\right)$. For observation i , let $\mathcal{Q}^{(i)}$ be the set of up to k most similar observations which satisfy the three aforementioned criteria. We can then compute the sparse affinity matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$ as follows. If $j \in \mathcal{Q}^{(i)}$, then $\mathbf{W}_{ij} = \chi^2(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})$. Otherwise $\mathbf{W}_{ij} = 0$. The diagonal degree matrix $\mathbf{D}$ of $\mathbf{W}$ is computed, i.e. $\mathbf{D}_{ii} = \sum_{l=1}^n \mathbf{W}_{il}$. Then, the Laplacian matrix which captures the manifold structure in the appearance space is $\mathbf{L} = \mathbf{D} - \mathbf{W}$.

This method of finding neighbors makes our tracker more robust to occlusions. Occlusions may cause the tracking target to be partially or completely occluded. However, the tracking target usually reappears after a few frames. Therefore, instead of trying to explicitly model occlusions, we try to connect the observations of the tracking target before and after the occlusion. As demonstrated in Figure 4, despite heavy occlusions in a time segment, the algorithm can still link the correct detections after the occlusion. The window size T affects the tracker's ability to recover from occlusions. If T is too small, the method will have difficulty recovering from occlusions that last longer than T . However, a large T may increase chances of linking two different objects.

3.2.2 Modeling Spatial Affinity

Other than modeling person detections of similar appearance, person detections which are a few centimeters apart in the same or neighboring frames are also very likely to belong to the same person. This assumption is reasonable in a multi-camera scenario because multiple detections will correspond to the same person, and due to calibration and person detection errors, not all detections will be projected to the exact same 3D location. Therefore, regardless of the appearance difference which may be resulting from non-color-calibrated cameras, these detections should belong to the same person. We therefore encode this information with another Laplacian matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$ defined as follows. Let $\mathcal{K}^{(i)}$ be the set of observations which are less than distance $\tilde{\delta}$ away and less than $\tilde{T}$ frames away from observation i . We compute the affinity matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ from $\mathcal{K}^{(i)}$ by setting $\mathbf{A}_{ij} = 1$ if $j \in \mathcal{K}^{(i)}$ and $\mathbf{A}_{ij} = 0$ otherwise. Define $\tilde{\mathbf{D}} \in \mathbb{R}^{n \times n}$ as a diagonal matrix where $\tilde{\mathbf{D}}_{ii}$ is the sum of $\mathbf{A}$'s i -th row. Following [43], the normalized Laplacian matrix is computed: $\mathbf{K} = \mathbf{I} - \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{A} \tilde{\mathbf{D}}^{-\frac{1}{2}}$. The parameters $\tilde{\delta}$ and $\tilde{T}$ for spatial affinity should be set more conservatively than the δ and T

used for appearance affinity. This is because the neighbor selection process for appearance affinity has the additional constraint that the color histograms of the detections need to look alike. However, for computing spatial affinity, $\tilde{\delta}$ and $\tilde{T}$ are the only two constraints, thus to avoid connecting incorrect person detections, they should be set very conservatively.

The loss function which combines the appearance and spatial affinity is as follows:

$$\min_{\mathbf{F}} \text{Tr}(\mathbf{F}^T (\mathbf{L} + \mathbf{K}) \mathbf{F})$$

(3)

s.t. columns of $\mathbf{F}$ satisfy Equation 1, $\forall i \in \mathcal{Y}, \mathbf{F}_i = \mathbf{Y}_i$.

Minimizing the loss term will result in a labeling which follows the manifold structure specified by appearance and spatial affinity. The first term in the constraints specifies that the label assignment matrix $\mathbf{F}$ should be binary and have a single 1 per row. The second term in the constraints is the face recognition constraint. Face recognition information is recorded in $\mathbf{Y} \in \{0, 1\}^{n \times c}$, where $\mathbf{Y}_{ij} = 1$ if the i -th observation belongs to class j , i.e. the face of observation i is recognized as person j . $\mathbf{Y}_{ij} = 0$ if we do not have any label information. There should only be at most a single 1 in each row of $\mathbf{Y}$. $\mathcal{Y} = \{i \mid \exists j \text{ s.t. } \mathbf{Y}_{ij} = 1\}$ are all the rows of $\mathbf{Y}$ which have a recognized face. As face verification is approaching human-level performance [44], it is in most cases reasonable to treat face information as a hard constraint. Experiments analyzing the effect of face recognition errors on tracking performance are also detailed in Section 4.1.7.

3.3 Spatial Locality Constraint

A person cannot be in multiple places at the same time, and we model this with pairwise person detection constraints. Given a pair of person detections (i, j) , if the speed $v^{(ij)}$ required to move from one person detection to the other is too large, then it is highly unlikely that the pair of person detections will belong to the same person. We aggregate all the person detection pairs which are highly unlikely to be of the same individual and encode them in the matrix $\tilde{\mathbf{S}}$, as shown in Equation 4.

$$\tilde{\mathbf{S}}_{ij} = \begin{cases} 0 & \text{if } v^{(ij)} \leq V \\ 1 & \text{otherwise} \end{cases}, \quad 1 \leq i, j \leq n, \quad (4)$$

where V is the maximum possible velocity of a moving person. $\tilde{\mathbf{S}}$ is defined so that if none of the person detection velocity constraints were violated, then $\mathbf{F}_{*j}^T \tilde{\mathbf{S}} \mathbf{F}_{*j} = 0$, where

$\mathbf{F}_{*j}$ is the label assignment vector (column vector of $\mathbf{F}$) for the j -th person. We gather this constraint for all individuals and obtain $\text{Tr}(\mathbf{F}^T \tilde{\mathbf{S}} \mathbf{F}) = 0$ if none of the constraints were violated. The scale of $\tilde{\mathbf{S}}$ is normalized to facilitate the subsequent optimization step. Let $\mathbf{D}'$ be a diagonal matrix where $\mathbf{D}'_{ii}$ is the sum of row i of $\tilde{\mathbf{S}}$, then we can compute the normalized $\mathbf{S} = \mathbf{D}'^{-\frac{1}{2}} \tilde{\mathbf{S}} \mathbf{D}'^{-\frac{1}{2}}$. The spatial locality constraint is incorporated into our objective function as shown in Equation 5.

$$\min_{\mathbf{F}} \text{Tr}(\mathbf{F}^T (\mathbf{L} + \mathbf{K}) \mathbf{F}) \quad \text{s.t.} \quad \text{Tr}(\mathbf{F}^T \mathbf{S} \mathbf{F}) = 0, \quad (5)$$

columns of $\mathbf{F}$ satisfy Equation 1, $\forall i \in \mathcal{Y}, \mathbf{F}_i = \mathbf{Y}_i$.

For simplicity, we do not force two detections from the same frame to not be of the same person. Nevertheless, this can be easily done by adding additional non-zero elements to $\mathbf{S}$.

Note that the purpose of the affinity-based Laplacian matrix $\mathbf{L}$ and $\mathbf{K}$ are completely opposite of the purpose of $\mathbf{S}$. $\mathbf{L}$ and $\mathbf{K}$ indicates which observations should be in the same cluster, while $\mathbf{S}$ enforces the fact that two observations cannot be in the same cluster. Though both $\mathbf{L}$ and $\mathbf{S}$ utilize the same assumption that a person cannot be at multiple places at the same time, these two matrices have completely different purposes in the loss function.

3.4 Nonnegative Matrix Optimization

Equation 5 is a combinatorial problem as the values of $\mathbf{F}$ are limited to zeros and ones. This is very difficult to solve and certain relaxation is necessary to efficiently solve the objective function. Therefore, we first relax the form of Equation 5, and then an iterative projected nonnegative gradient descent procedure is utilized to optimize the relaxed loss function.

The relaxation is motivated as follows. According to Equation 2, the columns of $\mathbf{F}$ are orthogonal to each other, i.e. $\mathbf{F}^T \mathbf{F} = \mathbf{J}$ is a diagonal matrix. Also, $\mathbf{F}$ is nonnegative by definition. According to [45], if both the orthogonal and nonnegative constraints are satisfied for a matrix, there will be *at most one non-zero entry* in each row of the matrix. This is still sufficient for identifying the class-membership of each observation, i.e. the mutual exclusion constraint still holds despite the fact that the non-zero entries are no longer exactly 1 but a continuous value. Therefore, we relax the form of $\mathbf{F}$ by allowing it to take on real values while still keeping the column orthogonal and nonnegative constraint. This leads to solving Equation 6.

$$\min_{\mathbf{F}} \text{Tr}(\mathbf{F}^T (\mathbf{L} + \mathbf{K}) \mathbf{F}) \quad \text{s.t.} \quad \text{Tr}(\mathbf{F}^T \mathbf{S} \mathbf{F}) = 0, \quad \mathbf{F}^T \mathbf{F} = \mathbf{J}, \quad \mathbf{F} \geq 0, \quad \forall i \in \mathcal{Y}, \mathbf{F}_i = \mathbf{Y}_i. \quad (6)$$

Equation 6 is a constrained quadratic programming problem, in which the mutual exclusion constraint is enforced by $\mathbf{F}^T \mathbf{F} = \mathbf{J}$ and $\mathbf{F} \geq 0$. One big advantage of this relaxation is that now our method can naturally handle false positive detections, because $\mathbf{F}$ is now also allowed to have a row where all elements are zeros, which corresponds to a person detection not being assigned to any class. This was not possible in the non-relaxed definition of $\mathbf{F}$. Analysis of robustness against false positives are shown in Section 4.1.5.

$\mathbf{F}^T \mathbf{F} = \mathbf{J}$ is still a difficult constraint to optimize. If $\mathbf{J}$ is the identity matrix, then $\mathbf{F}^T \mathbf{F} = \mathbf{I}$ forms the Stiefel manifold [46]. Though a few different methods have been proposed to perform optimization with the orthogonal constraint [46], [47], [48], [49], many methods are only applicable to a specific form of the objective function for the optimization process to converge. Therefore, we instead employ the simple yet effective quadratic penalty method [45], [50] to optimize the loss function. The quadratic penalty method incorporates the equality constraints into the loss function by adding a quadratic constraint violation error for each equality constraint. The amount of violation is scaled by a weight τ , which gradually increases as more iterations of the optimization are performed, thus forcing the optimization process to satisfy the constraints. More details on the convergence properties of the quadratic penalty method can be found in [50]. Therefore, we modify Equation 6 by moving the constraints $\mathbf{F}^T \mathbf{F} = \mathbf{J}$ and $\text{Tr}(\mathbf{F}^T \mathbf{S} \mathbf{F}) = 0$ into the loss function as a penalty term and arrive at the following:

$$\min_{\mathbf{F}} f(\mathbf{F}) = \min_{\mathbf{F}} \text{Tr}(\mathbf{F}^T (\mathbf{L} + \mathbf{K} + \tau \mathbf{S}) \mathbf{F}) + \tau \|\mathbf{F}^T \mathbf{F} - \mathbf{J}\|_{\mathbf{F}}^2 \quad \text{s.t.} \quad \mathbf{F} \geq 0, \quad \forall i \in \mathcal{Y}, \mathbf{F}_i = \mathbf{Y}_i. \quad (7)$$

For each τ , we minimize Equation 7 until convergence. Once converged, τ is multiplied by a step size $s > 1$ and Equation 7 is minimized again. Analysis of step size s versus tracking performance is shown in Section 4.1.5

To solve for Equation 7 given a fixed τ , we perform projected nonnegative gradient descent [51], which iteratively updates the solution at iteration l , i.e. $\mathbf{F}^{(l)}$, to $\mathbf{F}^{(l+1)}$:

$$\mathbf{F}^{(l+1)} = P \left[\mathbf{F}^{(l)} - \alpha^{(l)} \nabla f(\mathbf{F}^{(l)}) \right] \quad (8)$$

where the projection function P :

$$P[\mathbf{F}_{ij}] = \begin{cases} \mathbf{F}_{ij} & \text{if } \mathbf{F}_{ij} > 0 \\ 0 & \text{otherwise} \end{cases}, \quad (9)$$

is an element-wise function which maps an element back to the feasible region, i.e. in this case a negative number to zero. The step size $\alpha^{(l)}$ is found in a line search-like fashion, where we search for an $\alpha^{(l)}$ which provides sufficient decrease in the function value:

$$f(\mathbf{F}^{(l+1)}) - f(\mathbf{F}^{(l)}) \leq \sigma \text{Tr} \left(\nabla f(\mathbf{F}^{(l)})^T (\mathbf{F}^{(l+1)} - \mathbf{F}^{(l)}) \right). \quad (10)$$

Following [51], $\sigma = 0.01$ in our experiments. The gradient of our loss function f is

$$\nabla f(\mathbf{F}) = 2(\mathbf{L} + \mathbf{K} + \tau \mathbf{S}) \mathbf{F} + 4\tau \mathbf{F} (\mathbf{F}^T \mathbf{F} - \mathbf{J}). \quad (11)$$

Details on convergence guarantees are shown in [51]. To satisfy the face recognition constraints, the values of $\mathbf{F}$ for the rows in $\mathcal{Y}$ are set according to $\mathbf{Y}$ and never updated by the gradient.

The main advantage of projected nonnegative gradient descent over the popular multiplicative updates for nonnegative matrix factorization [52], [46] is that elements with zero values will have the opportunity to be non-zero in later iterations. However, for multiplicative updates, zero values will always stay zero. In our scenario, this means that if $\mathbf{F}_{ij}^{(l)}$ shrinks

Data: Location hypothesis $\mathbf{p}^{(i)}, t^{(i)}$, and appearance $\mathbf{x}^{(i)}, 1 \leq i \leq n$. Face recognition matrix $\mathbf{Y} \in \{0, 1\}^{n \times c}$.

Result: Final label assignment matrix $\mathbf{F}$

```

Compute Laplacian matrices  $\mathbf{L}, \mathbf{K}$  ;           // Sec. 3.2
Compute spatial locality matrix  $\mathbf{S}$  ;         // Sec. 3.3
Compute diagonal matrix  $\mathbf{J}$  ;               // Sec. 3.4
Compute diagonal matrix  $\mathbf{U}$  from  $\mathbf{Y}$  ;       // Sec. 3.4
Initialize  $\mathbf{F}^{(0)}$  with Equation 12 ;
 $l \leftarrow 0$  ;                               // iteration count
 $\tau \leftarrow 10^{-4}$  ;                       // initial penalty
repeat // Solve for Equation 7 with penalty method
     $\tau \leftarrow \tau \times s$  ;                 // gradually increase penalty  $\tau$ 
    repeat // projected gradient descent
        Compute  $\mathbf{F}^{(l+1)}$  from  $\mathbf{F}^{(l)}$  with Equation 8 ;
         $l \leftarrow l + 1$  ;
    until convergence ;
until  $\tau \geq 10^{11}$  ;
return  $\mathbf{F}^{(l)}$ 

```

Algorithm 1: Main steps in proposed tracking algorithm.

to 0 at iteration l in the optimization process, the decision that “observation i is not individual j ” is final and cannot be changed, which is not ideal. The projected nonnegative gradient descent method does not have this issue as the updates are additive and not multiplicative.

$\mathbf{J}$ is a diagonal matrix, where each element on the diagonal $\mathbf{J}_{ii}$ corresponds to the number of observations belonging to class i , i.e. m_i . As m_i is unknown beforehand, m_i is estimated by the number of recognized faces belonging to class i plus a constant β , which is proportional to the number of observations n . In our experiments we set $\beta = \frac{n}{1000}$.

To initialize our method, we temporarily ignore the mutual exclusion and spatial locality constraint and only use the manifold and face recognition information to find the initial value $\mathbf{F}^{(0)}$. $\mathbf{F}^{(0)}$ is obtained by minimizing Equation 12.

$$\min_{\mathbf{F}^{(0)}} Tr \left((\mathbf{F}^{(0)})^T (\mathbf{L} + \mathbf{K}) \mathbf{F}^{(0)} + (\mathbf{F}^{(0)} - \mathbf{Y})^T \mathbf{U} (\mathbf{F}^{(0)} - \mathbf{Y}) \right). \quad (12)$$

$\mathbf{U} \in \mathbb{R}^{n \times n}$ is a diagonal matrix. $\mathbf{U}_{ii} = \infty$ (a large constant) if $i \in \mathcal{Y}$, i.e. the i -th observation has a recognized face. Otherwise $\mathbf{U}_{ii} = 1$. $\mathbf{U}$ is used to enforce the consistency between prediction results and face recognition label information. The global optimal solution for Equation 12 is $\mathbf{F}^{(0)} = (\mathbf{L} + \mathbf{K} + \mathbf{U})^{-1} \mathbf{U} \mathbf{Y}$ [53].

Once the optimization is complete, we acquire $\mathbf{F}$ which satisfies the mutual exclusion and spatial locality constraint. Therefore, trajectories can be computed by simply connecting neighboring observations belonging to the same class. At one time instant, if there are multiple detections assigned to a person, which is common in multi-camera scenarios, then the weighted average location is computed. The weights are based on the scores in the final solution of $\mathbf{F}$. A simple filtering process is utilized to remove sporadic predictions. Algorithm 1 summarizes our tracker.

4 EXPERIMENTS

We present experiments on tracking followed by video summarization experiments based on our long-term tracking output.

4.1 Tracking

4.1.1 Data Sets

As we are interested in evaluating identity-aware tracking, we focused on sequences where identity information such as face recognition was available. Therefore, many popular tracking sequences such as the PETS 2009 sequences [54], Virat [55], TRECVID 2008 [56] and Town Centre [57] were not applicable as the faces in these sequences were too small to be recognized and no other identity information could be extracted. Basketball related sequences [20], [58] were not used as some manual effort is required to have an accurate OCR of jersey numbers [20]. The following four data sets were utilized in our experiments.

terrace1: The 4 camera *terrace1* [8] data set has 9 people walking around in a 7.5m by 11m area for 3 minutes 20 seconds. The scene is very crowded, thus putting the spatial locality constraint to test. The POM grid we computed had width and height of 25 centimeters per cell. Person detections were extracted at every frame. As the resolution of the video is low, one person did not have a recognizable face. For the sake of performing identity-aware tracking on this dataset, we manually added two identity annotations for each individual at the start and end of the person’s trajectory to guarantee that each individual had identity labels. None of the trackers utilized the fact that these two additional annotations were the start and end of a trajectory. In total, there were 794 identity labels out of 57,202 person detections.

Caremedia 6m: The 15 camera *Caremedia 6m* [59], [6] data set has 13 individuals performing daily activities in a nursing home for 6 minutes 17 seconds. Manual annotations were provided every second and interpolated to every frame. The data set records activities in a nursing home where staff maintain the nursing home and assist residents throughout the day. As the data set covers a larger area and is also longer than *terrace1*, we ran into memory issues for trackers which take POM as input when our cell size was 25 centimeters. Therefore, the POM grid we computed had width and height of 40 centimeters per cell. Person detections were extracted at every sixth frame. In total, there were 2,808 recognized faces and 12,129 person detections. Though on average there was a face for every 4 detections, but recognized faces were usually found in clusters and not evenly spread out over time. So there were still periods of time when no faces were recognized.

Caremedia 8h: The 15 camera *Caremedia 8h* data set is a newly annotated data set which has 49 individuals performing daily activities in the same nursing home as *Caremedia 6m*. The sequence is 7 hours 45 minutes long, which is 116.25 hours of video in total. Ground truth was annotated every minute. Person detections were extracted at every sixth frame. In total, there were 70,994 recognized faces and 402,833 person detections.

Caremedia 23d: The 15 camera *Caremedia 23d* data set is a newly annotated data set which consists of nursing home

recordings spanning over 23 days. Recordings at night were not processed as there was not much activity at night. In total, 4,935 hours of video were processed. To the best of our knowledge, this is the longest sequence to date to be utilized for multi-object tracking experiments. *Caremedia 23d* has 65 unique individuals. Ground truth was annotated every 30 minutes. Person detections were extracted at every sixth frame. In total, there were 3.1 million recognized faces and 17.8 million person detections.

4.1.2 Baselines

We compared our method with three identity-aware tracking baselines. As discussed in the Related Work section (Section 2), it is non-trivial to modify a non-identity-aware tracker to incorporate identity information. Therefore, other trackers which did not have the ability to incorporate identity information were not compared.

Multi-Commodity Network Flow (MCNF): The MCNF tracker [20] can be viewed as the K-Shortest-Path tracker (KSP, [32]) with identity aware capabilities. The KSP is a network flow-based method that utilizes POM localization information. Based on POM, the algorithm will find the K shortest paths, which correspond to the K most likely trajectories in the scene. MCNF further duplicates the graph in KSP for every different identity group in the scene. The problem is solved with linear programming plus an additional step of rounding non-integral values. We reimplemented the MCNF algorithm. The graph was duplicated c times to reflect the c unique individuals. Gurobi [60] was used as our linear program solver. Global appearance templates were computed from person detections which had recognized faces. The source code of POM and KSP were from the authors [8], [32]. This setting is referred to as *MCNF w/ POM*. The base cost of generating a trajectory, which is a parameter that controls the minimum length of the generated tracks, is set to -185 for all *MCNF w/ POM* experiments. For the two Caremedia data sets, we also took the person detection (PD) output and generated POM-like localizations which were also provided to MCNF. The localizations were generated by aggregating all person detections falling into each discretized grid cell at each time instant. This setting is referred to as *MCNF w/ PD*. For all *MCNF w/ PD* experiments, the grid size is 40 centimeters, the base cost of generating a trajectory is -60, and detections were aggregated over a time span of 6 frames to prevent broken trajectories. For the *Caremedia 8h* and *Caremedia 23d* set, the Gurobi solver was run in 12,000 frame batches to avoid memory issues.

Lagrangian Relaxation (LR): [40] utilized LR to impose mutual exclusion constraints for identity-aware tracking in a network flow framework very similar to MCNF, where each identity has their own identity specific edges. To fairly compare different data association methods, our LR-based tracker utilized the same appearance information used by all our other trackers, thus the structured learning and densely sampled windows proposed in [40] were not used. Specifically, LR uses the same POM-like input and network as MCNF.

Non-Negative Discretization (NND): The Non-Negative Discretization tracker [6] is a primitive version of our proposed

tracker. The three main differences are: 1) NND does not have the spatial locality constraint, 2) an extra Viterbi trajectory formulation step, which requires the start and end of trajectories, was necessary, and 3) a multiplicative update was used to perform non-negative matrix factorization. Start and end locations of trajectories are often unavailable in real world scenarios. Therefore, no start and end locations were provided to NND in our experiments, and the final trajectories of NND were formed with the same method used by our proposed tracker. NND utilizes [53] to build the manifold, but internal experiments have shown that utilizing the method in [53] to build the Laplacian matrix achieves similar tracking performance compared to the standard method [43], [61]. Therefore, to fairly compare the two data association methods, we utilized the same Laplacian matrix computation method for NND and our method. Also the spatial affinity term $\mathbf{K}$ was not used in the originally proposed NND, but for fairness we added the $\mathbf{K}$ term to NND.

4.1.3 Implementation Details

We utilized the person detection model from [62], [63] for person detection. Color histograms for the person detection were computed the same way as in [6]. We used HSV color histograms as done in [4]. We split the bounding box horizontally into regions and computed the color histogram for each region similar to the spatial pyramid matching technique [64]. Given L layers, we have $2^L - 1$ partitions for each template. L was 3 in our experiments. Since the person detector only detects upright people, tracking was not performed on sitting people or residents in wheelchairs. Background subtraction for POM was performed with [65]. Face information is acquired from the PittPatt software¹, which can recognize a face when a person is close enough to the camera. We acquired the gallery by clustering the recognized faces and then manually assigning identities to each cluster.

For our proposed method, the parameters for all four data sets were as follows. The number of nearest neighbors used for appearance-based manifold construction was $k = 25$. The window to search for appearance-based nearest neighbors was $T = 8$ seconds. The color histogram threshold $\gamma = 0.85$. The maximum localization error was $\delta = 125$ cm. For modeling spatial affinity, $\tilde{\delta}$ was 20 cm, and $\tilde{T}$ was 6 frames. When computing the spatial locality constraint matrix $\mathbf{S}$, we only looked for conflicting observations which were less than 6 frames apart to retain sparse $\mathbf{S}$. The above parameters were also used for NND. For the optimization step, the initial value of $\tau = 2 \times 10^{-4}$, and the final value was $\tau = 10^{11}$. The step size for updating τ , i.e. $\tau \leftarrow \tau \times s$, is $s = 2$.

4.1.4 Evaluation Metrics

Identity-aware tracking can be evaluated from a multi-object tracking point of view and a classification point of view. From the tracking point of view, the most commonly used multi-object tracking metric is *Multiple Object Tracking Accuracy* (MOTA²) [66], [68]. Following the evaluation method used in

1. Pittsburgh Pattern Recognition (<http://www.pittpatt.com>)

2. Code modified from <http://www.micc.unifi.it/lisanti/source-code/>.

Fig. 5: Snapshots of tracking results from the 4 camera *terracel* sequence.

[3], [6], the association between the tracking results and the ground truth is computed in 3D with a hit/miss threshold of 1 meter. MOTA takes into account the number of true positives (TP), false positives (FP), missed detections (false negatives, FN) and identity switches (ID-S). Following the setting in [20] ³ MOTA is computed as follows: $1 - \frac{\# \text{FP} + \# \text{FN} + \log_{10}(\# \text{ID-S})}{\# \text{ground truth}}$.

However, the TP count in MOTA does not take into account the identity of a person, which is unreasonable for identity aware tracking. Therefore, we compute identity-aware true positives (I-TP), which means that a detection is only a true positive if 1) it is less than 1 meter from the ground-truth and 2) the identities match. Similarly, we can compute I-FP and I-FN, which enables us to compute classification-based metrics such as micro-precision ($\text{MP} = \frac{\# \text{I-TP}}{\# \text{I-TP} + \# \text{I-FP}}$), micro-recall ($\text{MR} = \frac{\# \text{I-TP}}{\# \text{I-TP} + \# \text{I-FN}}$) and a comprehensive micro-F1 ($\frac{2 \times \text{MP} \times \text{MR}}{\text{MP} + \text{MR}}$) for each tracker. The *micro*-based performance evaluation takes into account the length (in terms of time) of each person’s trajectory, so a person who appears more often has larger influence to the final scores.

4.1.5 Tracking Results

Tracking results for the four data sets are shown in Table 1. We achieve the best performance in F1-scores across all four data sets. This means that our tracker can not only track a person well, but can also accurately identify the individual. Figure 5 and Figure 6 show qualitative examples of our tracking result.

The importance of the spatial locality constraint (SLC) is also shown clearly in Table 1a. Without the spatial locality constraint in the optimization step (*NND* and *Ours w/o SLC*), performance degraded significantly in the very crowded *terracel* sequence as the final result may show a person being at multiple places at the same time, thus hijacking the person detections of other individuals. For the *Caremedia* sequences, the SLC does not make a big difference, because 1) the scene is not so crowded and 2) the appearance of each individual is more distinct, thus relying only on the appearance feature can already achieve good performance.

3. There are two common transformation functions (denoted as $c_s()$ in [68]) for the identity-switch term, either $\log_{10}$ [68], [20] or the identity function [66]. We have selected the former as this is what was used in MCNF, which is one of our baselines.

Method	Micro-Precision	Micro-Recall	Micro-F1	TP	FN	FP	ID-S	MOTA
<i>Face only</i>	0.493	0.018	0.035	646	24708	284	5	0.014
<i>MCNF w/ POM</i>	0.593	0.532	0.561	21864	3298	644	197	0.844
<i>LR w/ POM</i>	0.609	0.478	0.535	19216	5996	521	147	0.743
<i>NND</i>	0.613	0.238	0.343	8035	17267	1771	57	0.249
<i>Ours w/o SLC</i>	0.704	0.346	0.464	10642	14655	1745	62	0.353
<i>Ours</i>	0.692	0.635	0.663	21370	3873	1783	116	0.777

(a) Tracking performance on *terracel* sequence.

Method	Micro-Precision	Micro-Recall	Micro-F1	TP	FN	FP	ID-S	MOTA
<i>Face only</i>	0.942	0.362	0.523	12369	21641	727	9	0.342
<i>MCNF w/ POM</i>	0.117	0.238	0.157	23493	9769	44452	757	-0.594
<i>MCNF w/ PD</i>	0.746	0.578	0.652	19941	13749	5927	329	0.422
<i>LR w/ PD</i>	0.802	0.565	0.663	19415	14408	4203	196	0.453
<i>NND</i>	0.861	0.726	0.787	25628	8364	3100	27	0.663
<i>Ours w/o SLC</i>	0.869	0.726	0.791	25578	8408	3080	33	0.662
<i>Ours</i>	0.865	0.755	0.807	26384	7576	3537	59	0.673

(b) Tracking performance on *Caremedia 6m* sequence.

Method	Micro-Precision	Micro-Recall	Micro-F1	TP	FN	FP	ID-S	MOTA
<i>Face only</i>	0.858	0.256	0.394	164	471	19	2	0.230
<i>MCNF w/ PD</i>	0.743	0.418	0.535	265	347	71	25	0.342
<i>LR w/ PD</i>	0.787	0.405	0.535	261	360	52	16	0.351
<i>NND</i>	0.588	0.505	0.543	314	281	174	42	0.283
<i>Ours w/o SLC</i>	0.638	0.549	0.590	349	257	151	31	0.357
<i>Ours</i>	0.648	0.571	0.607	370	241	149	26	0.386

(c) Tracking performance on *Caremedia 8h* sequence.

Method	Micro-Precision	Micro-Recall	Micro-F1	TP	FN	FP	ID-S	MOTA
<i>Face only</i>	0.819	0.199	0.154	125	512	28	2	0.154
<i>MCNF w/ PD</i>	0.712	0.355	0.474	205	412	92	22	0.209
<i>LR w/ PD</i>	0.663	0.357	0.464	215	411	116	13	0.174
<i>Ours</i>	0.698	0.532	0.604	326	299	147	14	0.300

(d) Tracking performance on *Caremedia 23d* sequence.

TABLE 1: Tracking performance on 4 tracking sequences. POM: Probabilistic Occupancy Map proposed in [8] as input. PD: Person detection as input. SLC: Spatial locality constraint. “w/” and “w/o” are shorthand for “with” and “without” respectively. We did not run the *MCNF w/ POM* on the longer *Caremedia* sequences as it was already performing poorly on *Caremedia 6m*.

The performance of *Face only* clearly shows the contribution of face recognition and tracking. For the *Caremedia* related sequences, face recognition could already achieve certain performance, but our tracker further improved F1 by at least 20% absolute. For *terracel*, there were very limited faces, and we were able to increase F1 by 60% absolute.

We also analyzed the robustness of our algorithm against false positives. The person detections on *Caremedia 6m* had around 13% false positive rate. Manual verification showed that for the person detections that were assigned a label by our tracker, only 0.1% were false positive detections. This means that $\frac{12.9\%}{13\%} = 99\%$ of the false positives were filtered out by our algorithm, thus demonstrating the robustness of our method against false positives.

Figure 7 demonstrates the effect of using different step size s when increasing the penalty term τ , which is utilized to enforce the mutual exclusion and spatial locality constraints. The initialization of our optimization process (Equation 12) does not enforce the two constraints, which lead to a MOTA of 0.358 when $\tau = 0$. As τ increases, MOTA gradually increased to 0.777, which demonstrates 1) the constraints were very important and 2) the quadratic penalty term utilized effectively enforced these constraints. Also, if the penalty term τ was

Fig. 6: Snapshots of tracking results from *Caremedia 8h* data set. To increase readability, not all arrows are drawn and only 12 out of 15 cameras are shown.

Fig. 7: Performance of *terrace1* MOTA under different penalty term τ and step size s . Best viewed in color.

increased too quickly, i.e. s is large, then tracking performance drops. This is reasonable as the optimization process is prone to getting stuck in a bad local minimum when the solution acquired from the previous τ is not a good initialization for the next τ .

The MCNF tracker is also a very strong baseline. For *terrace1*, KSP and consequently MCNF achieved very good MOTA results with POM person localization. MCNF was slightly worse than KSP on MOTA scores because 1) though MCNF is initialized by KSP, MCNF is no longer solving a problem with a global optimal solution and 2) MCNF is not directly optimizing for MOTA. However, for the *Caremedia 6m* sequence, *MCNF with POM* performance was poor because POM created many false positives in the complex indoor nursing home environment. This is due to non-ideal camera coverage that caused ambiguities in POM localization. Nevertheless, if the person detections used in our method was provided to MCNF (*MCNF with PD*), then MCNF performs reasonably well.

For *Caremedia 23d*, our best tracker can locate a person 53.2% of the time with 69.8% precision, i.e. in a 23 day time span, we can find a person more than 50% of the time with 70% accuracy. These results are encouraging, as the tracking output with such performance already has the potential to be utilized by other tasks, such as the experiments performed in Section 4.2 on surveillance video summarization.

4.1.6 Discussion - Advantages of Tracker

The key advantages of our tracker are as follows:

Face recognition output is integrated into the framework: Face recognition serves as a natural way to automatically assign identities to trajectories and also reinitialize trajectories in long-term tracking scenarios, where manual intervention is prohibitively costly. Also, face recognition is not affected when the same person wear different clothing in recordings over multiple days.

Naturally handle appearance changes: In our tracker, the appearance templates of the tracked target are implicitly encoded in the manifold structure we learn. Therefore, if the appearance of a tracked object changes smoothly along a manifold, our algorithm can model the change. No threshold is required to decide when to adaptively update the appearance model. If there is a drastic change in appearance for a tracked object, then the appearance manifold will highly likely be broken. However, the spatial affinity term could still link up the manifold.

Take into account appearance from multiple neighbors: Our tracker takes into account appearance information from multiple neighboring points, which enables us to have a more stable model of appearance. Linear programming and network flow-based methods can only either have a global appearance model or model appearance similarity only over the previous and next detection in the trajectory.

Handle multiple detections per frame for one individual: In multi-camera scenes, it is common that at one time instant, multiple detections from different cameras correspond to the same physical person. This may be difficult to deal with for single-camera multi-object trackers based on network flow [1], [10], because the spatial locality constraint for these methods are enforced based on the assumption that each individual can only be assigned a single person detection per frame. Therefore, multi-camera network flow-based methods such as [32], [20] utilize a two-step process where the POM is first used to aggregate evidences from multiple cameras before performing data association. Our formulation of the spatial locality constraint, which is based on the velocity to travel between two detections being below a threshold, can be viewed

as a generalization to the aforementioned assumption, and this enables us to have localization and data association in a single optimization framework.

No discretization of the space required in multi-camera scenarios: Previous multi-camera network flow methods [32], [20] require discretization of the tracking space in multi-camera scenarios to make the computation feasible. Finer grids run into memory issues when the tracking sequence is long and covers a wide area, and coarser grids run the risk of losing precision. However, our tracker works directly on person detections, and discretization is not necessary.

4.1.7 Discussion - Limitations of Tracker

There are also limitations to our tracker.

Assumes at least one face recognition per trajectory: If there is a trajectory where no faces were observed and recognized, then our tracker will completely ignore this trajectory, which is acceptable if we are only interested in identity-aware tracking. Otherwise, one potential solution is to find clusters of unassigned person detections and assign pseudo-identities to them to recover the trajectories.

Only bounded velocity model employed: To employ the more sophisticated constant velocity model, we could use pairs of points as the unit of location hypotheses, but this may generate significantly more location hypotheses than the current approach.

Assumes all cameras are calibrated: To combine person detections from different camera views, we utilize camera calibration parameters to map all person detections into a global coordinate system.

Face recognition gallery required beforehand: In order to track persons-of-interest, we require the gallery beforehand. This is the only manual step in our whole system, which could be alleviated by face clustering. Face clustering enables humans to efficiently assign identities to each cluster. Also, in a nursing home setting, the people-of-interest are fixed, thus this is a one-time effort which could be used for weeks or even months of recordings.

Assumes perfect face recognition: The current framework assumes perfect face recognition, which may not be applicable in all scenarios. We analyzed the effect of face recognition accuracy on tracking performance. We generated face recognition errors by randomly corrupting face recognition results in the *Caremedia 6m* set. The error rates range from 10% to 90%. The experiment was repeated 3 times per error rate, and the results with the 95% confidence intervals are shown in Figure 8. Results show that the general trend is a 20% increase in face recognition error will cause around 10% drop in tracking F1-score.

4.1.8 Timing Analysis

The whole tracking system includes person detection, face recognition, color histogram extraction and data association. The person detector we utilized [62], [63] ran at 40 times real-time. However, recently proposed real-time person detectors [69] will enable us to run person detection at 1 time real-time. The rest of the pipeline runs at around 3 times real-time on a single core, and the pipeline can be easily parallelized to

Fig. 8: *Caremedia 6m* tracking performance under varying face recognition error rate.

run faster than real-time. The data association part, which is our main focus, runs at around $\frac{1}{40}$ times real-time.

4.2 Visual Diary Generation

To demonstrate the usefulness of our tracking output, video summarization experiments were performed. We propose to summarize surveillance video using visual diaries, specifically in the context of monitoring elderly residents in a nursing home. Visual diary generation for elderly nursing home residents could enable doctors and staff to quickly understand the activities of a senior person throughout the day to facilitate the diagnosis of the elderly person's state of health. The visual diary for a specific person consists of two parts as shown in Figure 2: 1) snippets which contain snapshots and textual descriptions of activities-of-interest performed by the person, and 2) activity-related statistics accumulated over the whole day. The textual descriptions of the detected events enables efficient indexing of what a person did at different times. The statistics for the activities detected can be accumulated over many days to discover long-term patterns.

We propose to generate visual diaries with a summarization-by-tracking framework. Using the trajectories acquired from our tracking algorithm, we extract motion patterns from the trajectories to detect certain activities performed by each person in the scene. The motion patterns are defined in a simple rule-based manner. Even though more complex methods such as variants of Hidden Markov Models [70] to detect interactions could also be used, our goal here is to demonstrate the usefulness of our tracking result and not test state-of-the-art interaction detection methods, thus only a simple method was used. The activities we detect are as follows:

- Room change: Given the tracking output, we can detect when someone enters or leaves a room.
- Sit down / stand up: We trained a sitting detector [63] which detects whether someone is sitting. Our algorithm looks for tracks which end/begin near a seat and check whether someone sat down/stood up around the same time.
- Static interaction: If two people stand closer than distance D' for duration T' , then it is likely that they are interacting.
- Dynamic interaction: If two people are moving with distance less than D' apart for a duration longer than T' , and if they are moving faster than 20 cm/s, then it is highly likely that they are walking together.

According to [71], if people are travelling in a group, then they should be at most 7 feet apart. Therefore, we set the maximum distance D' for there to be interaction between two people at 7 feet. The minimum duration of interaction T' was set to 8 seconds in our experiments.

Given the time and location of all the detected activities, we can sort the activities according to time and generate the visual diary. The visual diary for a given individual consists of the following:

- Snippets: snapshots and textual descriptions of the activity. Snapshots are extracted from video frames during the interaction and textual descriptions are generated using natural language templates.
- Room/state timing estimates: time spent sitting or standing/walking in each room.
- Total interaction time: time spent in social interactions.

Our proposed method of using tracking output for activity detection can be easily combined with traditional activity recognition techniques using low-level features such as Improved Dense Trajectories [72] with Fisher Vectors [73] to achieve better activity detection performance and detect more complex actions, but extending activity recognition to activity detection is beyond the scope of this paper.

Visual Diary Generation Results

We performed long-term surveillance video summarization experiments by generating visual diaries on the *Caremedia 8h* sequence. To acquire ground truth, we manually labeled the activities of three residents throughout the sequence. The nursing home residents were selected because they are the people we would like to focus on for the automatic analysis of health status. 184 ground-truth activities were annotated.

We evaluated the different aspects of the visual diary: “room/state timing estimates”, “interaction timing estimates” and “snippet generation”. The evaluation of “room/state timing estimates”, i.e. predicted room location and state (sitting or upright), of a person was done on the video frame level. A frame was counted as true positive if the predicted state for a given video frame agrees with the ground truth. False positives and false negatives were computed similarly. To evaluate “interaction timing estimates”, i.e. how much time a person spent in interactions, a frame was only counted as true positive if 1) both the prediction result and ground truth result agree that there was interaction and 2) the ID of the interacting targets match. False positives and false negatives were computed similarly. The evaluation of “snippet generation” accuracy was done as follows. For snippets related to sit down, stand up and room change activities, a snippet was correct if the predicted result and ground truth result had less than a 5 second time difference. For social interaction-related snippets, a snippet was correct if more than 50% of the predicted snippet contained a matching ground truth interaction. Also, if a ground truth interaction was predicted as three separate interactions, then only one interaction was counted as true positive while the other two were counted as false positives. This prevents double counting of a single ground-truth interaction.

Visual diary components	Micro-Precision	Micro-Recall	Micro-F1
Snippet generation	0.382	0.522	0.441
Room/state timing estimates	0.809	0.511	0.626
Interaction timing estimates	0.285	0.341	0.311

TABLE 2: Evaluation of generated visual diary.

Results are shown in Table 2, which shows that 38% of the generated snippets were correct, and we have retrieved 52% of the activities-of-interest. For “room/state timing estimates”, a 51.1% recall shows that we know the state and room location of a person more than 50% of the time. The lower performance for “interaction timing estimates” was mainly caused by tracking failures, as both persons need to be tracked correctly for interactions to be correctly detected and timings to be accurate. These numbers are not high, but given that our method is fully automatic other than the collection of the face gallery, this is a first cut at generating visual diaries for the elderly by summarizing hundreds or even thousands of hours of surveillance video.

We analyzed the effect of tracking performance on snippet generation accuracy. We computed snippet generation F1-score for multiple tracking runs with varying tracking performance. These runs include our baseline runs and also runs where we randomly corrupted face recognition labels to decrease tracking performance. Results in Figure 10 show that as tracking F1 increases, snippet generation F1 also increases with a trend which could be fitted by a second-order polynomial.

Figure 9 shows example visual diaries for residents ID 3 and 11. We can clearly see what each resident was doing at each time of the day. Long term statistics shown in Figure 2 also clearly indicate the amount of time spent in each room and in social interactions. If these statistics were computed over many days, a doctor or staff member could start looking for patterns to better assess the status of health of a resident.

5 CONCLUSION

We present an identity-aware tracker which leverages face recognition information to enable automatic reinitialization of tracking targets in very long-term tracking scenarios. Face recognition information is ideal in that it is robust to appearance and apparel change. However, face recognition is unavailable in many frames, thus we propagate identity information through a manifold learning framework which is solved by nonnegative matrix optimization. Tracking experiments performed on up to 4,935 hours of video in a complex indoor environment showed that our tracker was able to localize a person 53.2% of the time with 69.8% precision. Accurate face recognition is key to good tracking results, where a 20% increase in face recognition accuracy will lead to around 10% increase in tracking F1-score. In addition to tracking experiments, we further utilized tracking output to generate visual diaries for identity-aware video summarization. Experiments performed on 116.25 hours of video showed that we can generate visual diary snippets with 38% precision and 52% recall. Compared to tedious manual analysis of thousands of hours of surveillance video, our method is a strong alternative

(a) Example snippets for resident 3.

(b) Example snippets for resident 11.

Fig. 9: Example visual diary snippets for each resident.

Fig. 10: Performance of snippet generation (y-axis) under varying tracking performance (x-axis).

as it potentially opens the door to summarization of the ocean of surveillance video generated every day.

REFERENCES

- [1] L. Zhang, Y. Li, and R. Nevatia, "Global data association for multi-object tracking using network flows," in *CVPR*, 2008.
- [2] R. T. Collins, "Multitarget data association with higher-order motion models," in *CVPR*, 2012.
- [3] A. Andriyenko, K. Schindler, and S. Roth, "Discrete-continuous optimization for multi-target tracking," in *CVPR*, 2012.
- [4] K. Okuma, A. Taleghani, N. D. Freitas, O. D. Freitas, J. J. Little, and D. G. Lowe, "A boosted particle filter: Multitarget detection and tracking," in *ECCV*, 2004.
- [5] J. K. Rowling, *Harry Potter and the Prisoner of Azkaban*. London: Bloomsbury, 1999.
- [6] S.-I. Yu, Y. Yang, and A. Hauptmann, "Harry Potter's Marauder's Map: Localizing and tracking multiple persons-of-interest by nonnegative discretization," in *CVPR*, 2013.
- [7] W. Luo, J. Xing, X. Zhang, X. Zhao, and T.-K. Kim, "Multiple object tracking: A literature review," *arXiv preprint arXiv:1409.7618*, 2014.
- [8] F. Fleuret, J. Berclaz, R. Lengagne, and P. Fua, "Multicamera people tracking with a probabilistic occupancy map," *IEEE TPAMI*, 2008.
- [9] A. R. Zamir, A. Dehghan, and M. Shah, "GMCP-tracker: global multi-object tracking using generalized minimum clique graphs," in *ECCV*, 2012.
- [10] H. Pirsiavash, D. Ramanan, and C. C. Fowlkes, "Globally-optimal greedy algorithms for tracking a variable number of objects," in *CVPR*, 2011.
- [11] A. Andriyenko and K. Schindler, "Multi-target tracking by continuous energy minimization," in *CVPR*, 2011.
- [12] A. Milan, K. Schindler, and S. Roth, "Detection-and trajectory-level exclusion in multiple object tracking," in *CVPR*, 2013.
- [13] A. Butt and R. Collins, "Multi-target tracking by Lagrangian relaxation to min-cost network flow," in *CVPR*, 2013.
- [14] C. Huang, B. Wu, and R. Nevatia, "Robust object tracking by hierarchical association of detection responses," in *ECCV*, 2008.
- [15] C.-H. Kuo, C. Huang, and R. Nevatia, "Multi-target tracking by on-line learned discriminative appearance models," in *CVPR*, 2010.
- [16] C.-H. Kuo and R. Nevatia, "How does person identity recognition help multi-person tracking?" in *CVPR*, 2011.
- [17] B. Yang and R. Nevatia, "An online learned CRF model for multi-target tracking," in *CVPR*, 2012.
- [18] Y. Li, C. Huang, and R. Nevatia, "Learning to associate: Hybridboosted multi-target tracker for crowded scene," in *CVPR*, 2009.
- [19] B. Yang and R. Nevatia, "Multi-target tracking by online learning of non-linear motion patterns and robust appearance models," in *CVPR*, 2012.
- [20] H. Ben Shitrit, J. Berclaz, F. Fleuret, and P. Fua, "Multi-commodity network flow for tracking multiple people," *IEEE TPAMI*, 2014.
- [21] H. Jiang, S. Fels, and J. J. Little, "A linear programming approach for multiple object tracking," in *CVPR*, 2007.
- [22] B. Leibe, K. Schindler, and L. Van Gool, "Coupled detection and trajectory estimation for multi-object tracking," in *CVPR*, 2007.
- [23] A. Andriyenko and K. Schindler, "Globally optimal multi-target tracking on a hexagonal lattice," in *ECCV*, 2010.
- [24] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *CVPR*, 2005.
- [25] S.-H. Bae and K.-J. Yoon, "Robust online multi-object tracking based on tracklet confidence and online discriminative appearance learning," in *CVPR*, 2014.
- [26] B. Wang, G. Wang, K. L. Chan, and L. Wang, "Tracklet association with online target-specific metric learning," in *CVPR*, 2014.
- [27] X. Zhang, W. Hu, S. Maybank, and X. Li, "Graph based discriminative learning for robust and efficient object tracking," in *CVPR*, 2007.
- [28] W. Hu, X. Li, W. Luo, X. Zhang, S. Maybank, and Z. Zhang, "Single and multiple object tracking using log-Euclidean Riemannian subspace and block-division appearance model," *IEEE TPAMI*, 2012.
- [29] S. Salti, A. Cavallaro, and L. Di Stefano, "Adaptive appearance modeling for video tracking: Survey and evaluation," *IEEE Transactions on Image Processing*, 2012.
- [30] Z. Khan, T. Balch, and F. Dellaert, "MCMC-based particle filtering for tracking a variable number of interacting targets," *IEEE TPAMI*, 2005.
- [31] R. Hess and A. Fern, "Discriminatively trained particle filters for complex multi-object tracking," in *CVPR*, 2009.
- [32] J. Berclaz, F. Fleuret, E. Turetken, and P. Fua, "Multiple object tracking using k-shortest paths optimization," *IEEE TPAMI*, 2011.
- [33] C. Dicle, M. Sznajder, and O. Camps, "The way they move: Tracking targets with similar appearance," in *ICCV*, 2013.
- [34] X. Wang, E. Turetken, F. Fleuret, and P. Fua, "Tracking interacting objects optimally using integer programming," in *ECCV*, 2014.
- [35] V. Ferrari, T. Tuytelaars, and L. Van Gool, "Real-time affine region tracking and coplanar grouping," in *CVPR*, 2001.

- [36] M. J. Marín-Jiménez, A. Zisserman, M. Eichner, and V. Ferrari, "Detecting people looking at each other in videos," *IJCV*, 2014.
- [37] V. Chari, S. Lacoste-Julien, I. Laptev, and J. Sivic, "On pairwise costs for network flow multi-object tracking," in *CVPR*, 2015.
- [38] M. Zervos, H. BenShitrit, F. Fleuret, and P. Fua, "Facial descriptors for identity-preserving multiple people tracking," Technical Report EPFL-ARTICLE-187534, 2013.
- [39] W.-L. Lu, J.-A. Ting, J. J. Little, and K. P. Murphy, "Learning to track and identify players from broadcast sports videos," *IEEE TPAMI*, 2013.
- [40] A. Dehghan, Y. Tian, P. H. Torr, and M. Shah, "Target identity-aware network flow for online multiple target tracking," in *CVPR*, 2015.
- [41] M. Everingham, J. Sivic, and A. Zisserman, "Taking the bite out of automated naming of characters in TV video," *Image and Vision Computing*, 2009.
- [42] J. Sivic, M. Everingham, and A. Zisserman, "Who are you?"-Learning person specific classifiers from video," in *CVPR*, 2009.
- [43] A. Y. Ng, M. I. Jordan, Y. Weiss *et al.*, "On spectral clustering: Analysis and an algorithm," in *NIPS*, 2002.
- [44] C. Lu and X. Tang, "Surpassing human-level face verification performance on LFW with GaussianFace," *arXiv preprint arXiv:1404.3840*, 2014.
- [45] Y. Yang, H. T. Shen, F. Nie, R. Ji, and X. Zhou, "Nonnegative spectral clustering with discriminative regularization," in *AAAI*, 2011.
- [46] Z. Yang and E. Oja, "Linear and nonlinear projective nonnegative matrix factorization," *IEEE Transactions on Neural Networks*, 2010.
- [47] C. Ding, T. Li, and M. I. Jordan, "Nonnegative matrix factorization for combinatorial optimization: Spectral clustering, graph matching, and clique finding," in *ICDM*, 2008.
- [48] J. Yoo and S. Choi, "Nonnegative matrix factorization with orthogonality constraints," *Journal of Computing Science and Engineering*, 2010.
- [49] F. Pompili, N. Gillis, P.-A. Absil, and F. Glineur, "Two algorithms for orthogonal nonnegative matrix factorization with application to clustering," *Neurocomputing*, 2014.
- [50] S. J. Wright and J. Nocedal, *Numerical Optimization*. Springer, New York, 1999, vol. 2.
- [51] C.-J. Lin, "Projected gradient methods for nonnegative matrix factorization," *Neural computation*, 2007.
- [52] D. D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," in *NIPS*, 2000.
- [53] Y. Yang, F. Nie, D. Xu, J. Luo, Y. Zhuang, and Y. Pan, "A multimedia retrieval framework based on semi-supervised ranking and relevance feedback," *IEEE TPAMI*, 2012.
- [54] A. Ellis, A. Shahrokni, and J. Ferryman, "PETS2009 and winter-PETS 2009 results: A combined evaluation," in *Performance Evaluation of Tracking and Surveillance (PETS-Winter)*, 2009.
- [55] S. Oh, A. Hoogs, A. Perera, N. Cuntoor, C.-C. Chen, J. T. Lee, S. Mukherjee, J. Aggarwal, H. Lee, L. Davis *et al.*, "A large-scale benchmark dataset for event recognition in surveillance video," in *CVPR*, 2011.
- [56] "National institute of standards and technology: TRECVID 2012 evaluation for surveillance event detection. <http://www.nist.gov/speech/tests/trecvid/2012/>," 2012.
- [57] B. Benfold and I. Reid, "Stable multi-target tracking in real-time surveillance video," in *CVPR*, 2011.
- [58] C. Vondrick, D. Patterson, and D. Ramanan, "Efficiently scaling up crowdsourced video annotation," *IJCV*, 2013.
- [59] Y. Yang, A. Hauptmann, M.-Y. Chen, Y. Cai, A. Bharucha, and H. Wactlar, "Learning to predict health status of geriatric patients from observational data," in *Computational Intelligence in Bioinformatics and Computational Biology*, 2012.
- [60] "Gurobi optimizer reference manual, <http://www.gurobi.com/>," 2012.
- [61] M. Belkin and P. Niyogi, "Laplacian eigenmaps for dimensionality reduction and data representation," *Neural computation*, 2003.
- [62] P. F. Felzenszwalb, R. B. Girshick, D. A. McAllester, and D. Ramanan, "Object detection with discriminatively trained part-based models," *IEEE TPAMI*, 2010.
- [63] R. B. Girshick, P. F. Felzenszwalb, and D. McAllester, "Discriminatively trained deformable part models, release 5," <http://people.cs.uchicago.edu/~rbg/latent-release5/>.
- [64] S. Lazebnik, C. Schmid, and J. Ponce, "Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories," in *CVPR*, 2006.
- [65] C. Stauffer and W. E. L. Grimson, "Adaptive background mixture models for real-time tracking," in *CVPR*, 1999.
- [66] K. Bernardin and R. Stiefelhagen, "Evaluating multiple object tracking performance: The CLEAR MOT metrics," in *J. Image Video Process.*, 2008.
- [67] Y. Liu, W. Wang, and Y. Zhao, "Particle flow for sequential monte carlo implementation of probability hypothesis density," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 4371–4375.
- [68] R. Kasturi, D. Goldgof, P. Soundararajan, V. Manohar, J. Garofolo, R. Bowers, M. Boonstra, V. Korzhova, and J. Zhang, "Framework for performance evaluation of face, text, and vehicle detection and tracking in video: Data, metrics, and protocol," *IEEE TPAMI*, 2009.
- [69] M. A. Sadeghi and D. Forsyth, "30hz object detection with DPM V5," in *ECCV*, 2014.
- [70] N. M. Oliver, B. Rosario, and A. P. Pentland, "A Bayesian computer vision system for modeling human interactions," *IEEE TPAMI*, 2000.
- [71] C. McPhail and R. T. Wohlstein, "Using film to analyze pedestrian behavior," in *Sociological Methods & Research*, 1982.
- [72] H. Wang, C. Schmid *et al.*, "Action recognition with improved trajectories," in *ICCV*, 2013.
- [73] K. Chatfield, V. Lempitsky, A. Vedaldi, and A. Zisserman, "The devil is in the details: An evaluation of recent feature encoding methods," in *BMVC*, 2011.