Multi-hypothesis Motion Planning for Visual Object Tracking

Haifeng Gong, Jiwoong Sim and Jianbo Shi
hfgong@seas.upenn.edu

August 17, 2011

Explanation of Datasets

We have collected a video from a moving vehicle in an urban city.

- We have picked 7 sequences which contains multiple people, and have interesting interactions and occlusions.
- Details of all sequences are shown in Table.
- There are total of 48 people in all the sequences. Many of these people cannot be tracked through the entirety of the sequence, because of the high occlusion rates.

	# obj	# frames	# BB	# Occl. BB
seq #1	13	169	1139	471
seq #2	12	60	532	130
seq #3	7	35	210	125
seq #4	4	40	148	51
seq #5	5	112	211	46
seq #6	5	41	170	17
seq #7	2	27	54	16
Total	48	484	2464	856

They are divided into 3 difficulties according to the number of occluded bounding boxes. (BB = Bounding boxes.)

Example of crowd street scene

Motion planning as motion model for visual object tracking

- In crowded street scenes, frequent occlusions, lead to ambiguous data association or ‘drifting’ in tracking.

Motion planning as motion model for visual object tracking

- In crowded street scenes, frequent occlusions, lead to ambiguous data association or ‘drifting’ in tracking.
- Many of these occlusions could be dealt with using a long-term motion model.

Motion planning as motion model for visual object tracking

- In crowded street scenes, frequent occlusions, lead to ambiguous data association or ‘drifting’ in tracking.
- Many of these occlusions could be dealt with using a long-term motion model.
- We propose to construct a set of ‘plausible’ plans for each person.
  - multi-hypotheses,
  - no redundancy, no unnecessary loop,
  - no collisions with other objects.

Tracking with motion planning

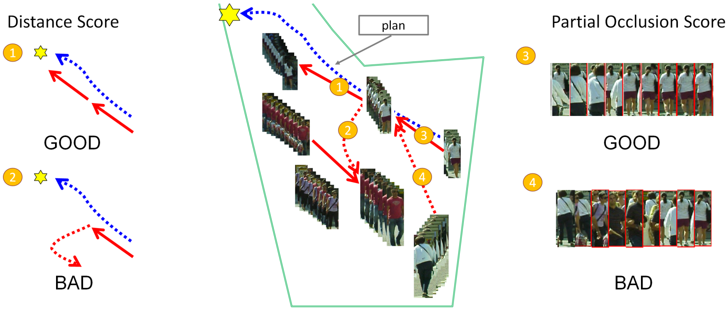

Tracking with multi-hypothesis motion planning

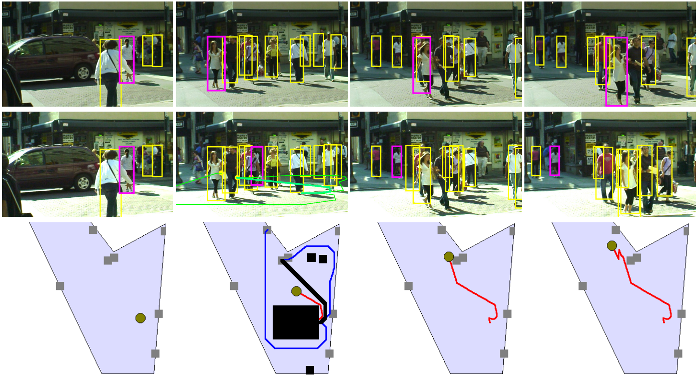

Top: tracking without planning. Middle: tracking with planning. Bottom, top view of tracking with planning. Note that we plan in advance, therefore, the obstacles are other objects a few frames ago.

Plausible plans for visual object tracking

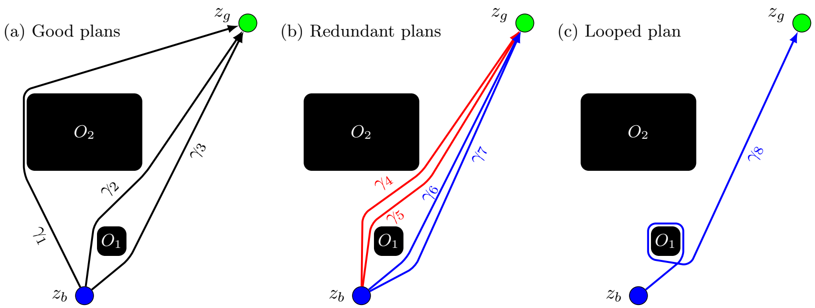

Examples of plausible plans and bad plans for visual object tracking. O_1 and O_2 are two obstacles. γ_i are possible paths. z_b and z_g are the start point and goal respectively.

Homotopy-class planning [Bhattacharya2010]

Notations

- z — a point in the complex plane,
- z_b — the start point and
- z_g — the goal of an agent (where it is intended to go).

A path $\gamma(s)$ is a complex function of arc length parameter $s \in [0, T]$, with constraints $\gamma(0) = z_b$ and $\gamma(T) = z_g$.

Homotopy-class planning [Bhattacharya2010] cont'

To distinguish different homotopy classes, a complex *obstacle marker function* is defined as

$$F(z) = \frac{f_0(z)}{(z-\zeta_1)(z-\zeta_2)\cdots(z-\zeta_N)} \tag{1}$$

where $f_0(z)$ is a complex Homomorphic function and ζ_i is a point in the area covered by obstacle i in the complex plane.

Homotopy-class planning [Bhattacharya2010] cont'

Cauchy Integral Theorem Two trajectories $\gamma_1(s)$ and $\gamma_2(s)$ connecting the same pair of points lie in the same homotopy class if and only if

$$\int_{\gamma_1} F(z)dz = \int_{\gamma_2} F(z)dz \quad (2)$$

given the assumption that $f_0(z)$ meets certain conditions. Therefore they use the L-value, defined as

$$L(\gamma) = \int_{\gamma} F(z)dz \quad (3)$$

to index homotopy classes.

Drawbacks of [Bhattacharya2010]

1. When obstacles differ greatly in size, [Bhattacharya2010] performs poorly.

Drawbacks of [Bhattacharya2010]

1. When obstacles differ greatly in size, [Bhattacharya2010] performs poorly.
   - It might loop around small obstacles before taking bigger obstacles into account.

Drawbacks of [Bhattacharya2010]

1. When obstacles differ greatly in size, [Bhattacharya2010] performs poorly.
   - It might loop around small obstacles before taking bigger obstacles into account.

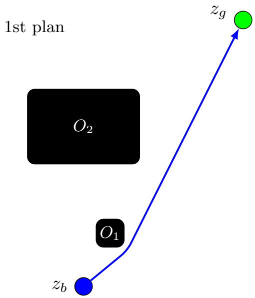

Drawbacks of [Bhattacharya2010]

1. When obstacles differ greatly in size, [Bhattacharya2010] performs poorly.
   - It might loop around small obstacles before taking bigger obstacles into account.

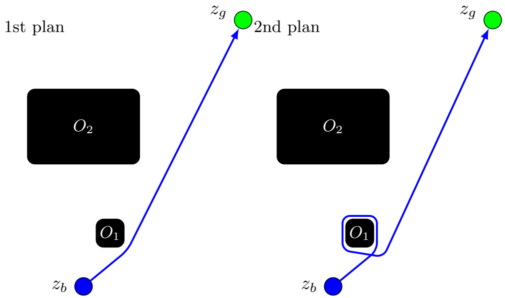

Drawbacks of [Bhattacharya2010]

1. When obstacles differ greatly in size, [Bhattacharya2010] performs poorly.
   - It might loop around small obstacles before taking bigger obstacles into account.

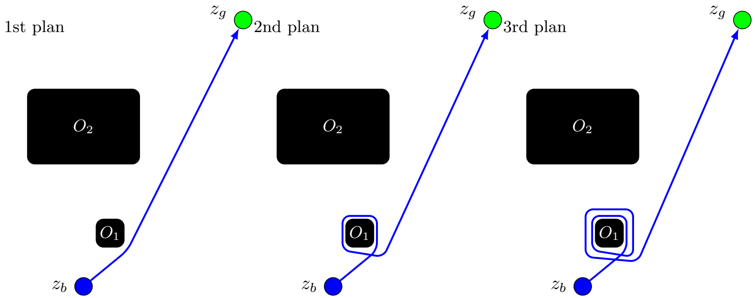

Drawbacks of [Bhattacharya2010] cont'

2. Obstacle marker function must be carefully chosen for numeric stability of *L*-values in real-world applications.

Drawbacks of [Bhattacharya2010] cont'

2. Obstacle marker function must be carefully chosen for numeric stability of *L*-values in real-world applications.
3. The representation of state space is an infinite augmented graph.

From L-value to winding numbers

- We propose replacing L-value with a more informative index, that incorporates the number of loops around obstacles.
- This allow us to screen out any paths with many loops, which are unlikely to be the paths that people actually take.

From L-value to winding numbers

- The L-value of a plan γ with respect to a single obstacle,

$$L = \int_{\gamma} \frac{f(z)}{z - z_0} dz \tag{4}$$

From L-value to winding numbers

- The L-value of a plan γ with respect to a single obstacle,

$$L = \int_{\gamma} \frac{f(z)}{z - z_0} dz \tag{4}$$

- L-values for a single obstacle must be in the discrete set of

$$\{k * 2\pi i f(z_0) + L_0 : k \in \mathbb{Z}\}. \tag{5}$$

From L-value to winding numbers

- The L-value of a plan γ with respect to a single obstacle,

$$L = \int_{\gamma} \frac{f(z)}{z - z_0} dz \tag{4}$$

- L-values for a single obstacle must be in the discrete set of

$$\{k * 2\pi i f(z_0) + L_0 : k \in \mathbb{Z}\}. \tag{5}$$

- Thus we can use k (winding number) to distinguish homotopy classes with respect to one obstacle which

Example of winding numbers

$L = L_0 - 4\pi i f(z_0)$	$L = L_0 - 2\pi i f(z_0)$	$L = L_0$	$L = L_0 + 2\pi i f(z_0)$
$k = -2$	$k = -1$	$k = 0$	$k = 1$
$\Delta\theta = -3\pi$	$\Delta\theta = -1\pi$	$\Delta\theta = \pi$	$\Delta\theta = 3\pi$

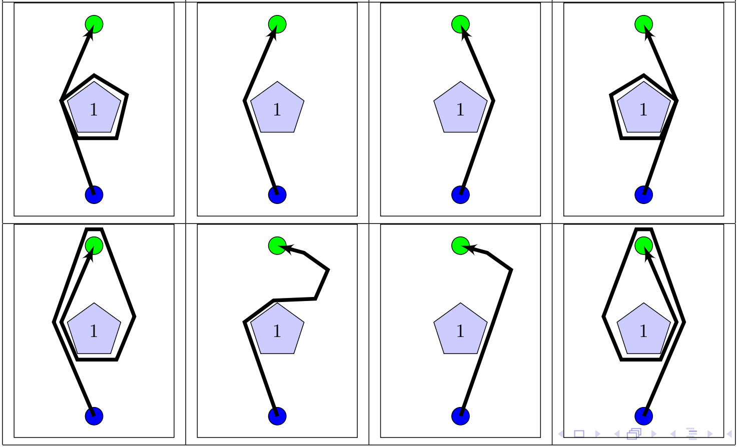

Winding numbers

- $k > 0$ indicates a path to the right of the obstacle that includes k loops around it.

Winding numbers

- $k > 0$ indicates a path to the right of the obstacle that includes k loops around it.
- $k < -1$ indicates a path to the left of the obstacle that includes $-k-1$ loops around it.

Winding numbers

- $k > 0$ indicates a path to the right of the obstacle that includes k loops around it.
- $k < -1$ indicates a path to the left of the obstacle that includes $-k - 1$ loops around it.
- For a plausible path, the values of k will likely be 0 or -1, meaning ‘go-right’ or ‘go-left’ around the obstacle.

Vector of winding numbers

Definition

By letting k_i be the k-value associated with the i-th obstacle, we can denote a homotopy class with respect to all obstacles as an integer vector (vector of winding numbers, or k**-vector**)

$$\mathbf{k} = (k_1, k_2, \cdots, k_N)^T. \quad (6)$$

Theorem

Two trajectories γ_1 and γ_2 with k-vectors $\mathbf{k}_1$ and $\mathbf{k}_2$ connecting the same points lie in the same homotopy class if and only if $\mathbf{k}_1 = \mathbf{k}_2$.

From winding numbers to winding angles

- A path γ can be written in parametric form, $\gamma(s) = z_0 + r(s)\exp[i\theta(s)]$.

From winding numbers to winding angles

- A path γ can be written in parametric form, $\gamma(s) = z_0 + r(s)\exp[i\theta(s)]$.
- The obstacle marker function can be a constant $f(z) = 1$.

From winding numbers to winding angles

- A path γ can be written in parametric form, $\gamma(s) = z_0 + r(s)\exp[i\theta(s)]$.
- The obstacle marker function can be a constant $f(z) = 1$.
- Then L-value can be computed in closed form as

$$L = \mathrm{Const} + i[\theta(T) - \theta(0)]. \tag{7}$$

From winding numbers to winding angles

- A path γ can be written in parametric form, $\gamma(s) = z_0 + r(s)\exp[i\theta(s)]$.
- The obstacle marker function can be a constant $f(z) = 1$.
- Then L-value can be computed in closed form as

$$L = \mathrm{Const} + i[\theta(T) - \theta(0)]. \tag{7}$$

- The imaginary part

$$\Delta\theta = \theta(T) - \theta(0) = \Delta\theta_0 + 2k\pi \tag{8}$$

may differ by $2k\pi$, where k is also a winding number.

From winding numbers to winding angles

- A path γ can be written in parametric form, $\gamma(s) = z_0 + r(s)\exp[i\theta(s)]$.
- The obstacle marker function can be a constant $f(z) = 1$.
- Then L-value can be computed in closed form as

$$L = \mathrm{Const} + i[\theta(T) - \theta(0)]. \tag{7}$$

- The imaginary part

$$\Delta\theta = \theta(T) - \theta(0) = \Delta\theta_0 + 2k\pi \tag{8}$$

  may differ by $2k\pi$, where k is also a winding number.
- We call $\Delta\theta$ the winding angle of γ w.r.t. obstacle z_0.

Augmented Graph

- Like [Bhattacharya2010], we use a graph based search algorithm, but we search on a finite graph.
- We begin with neighborhood graph G, in which each grid point on ground not occupied by an obstacle is a vertex, and each pair of neighboring points are connected by an edge.
- Each vertex in G is represented by its coordinate on ground z.
- We augment this graph with winding angle to create an augmented graph $\bar{G}$.
- We equip both vertices with winding angles and edges with increments of winding angles.

Augmented Graph

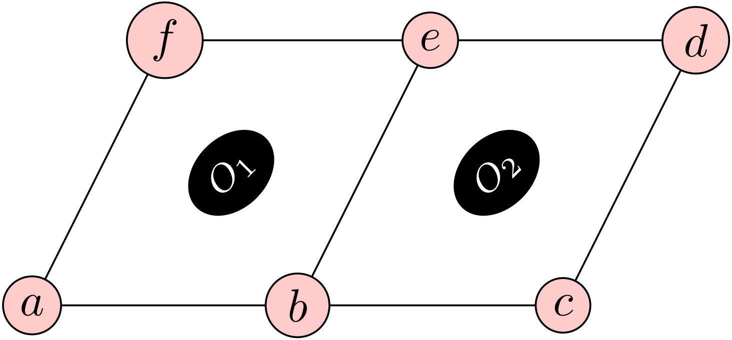

Augmented Graph

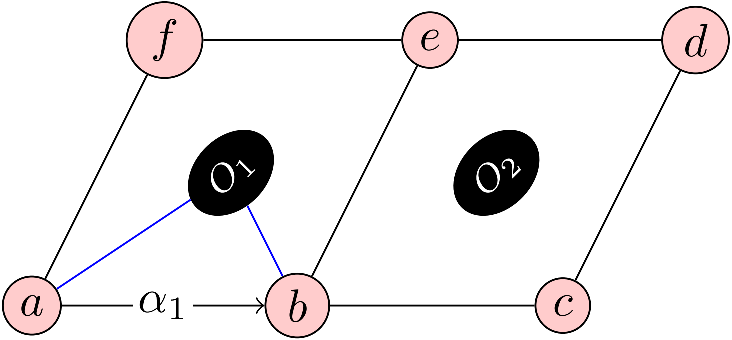

Augmented Graph

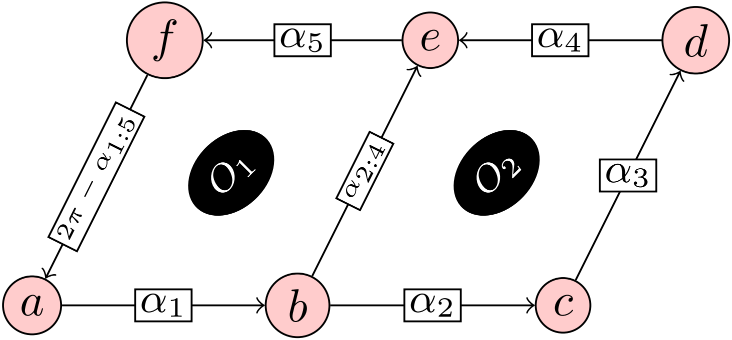

Augmented Graph

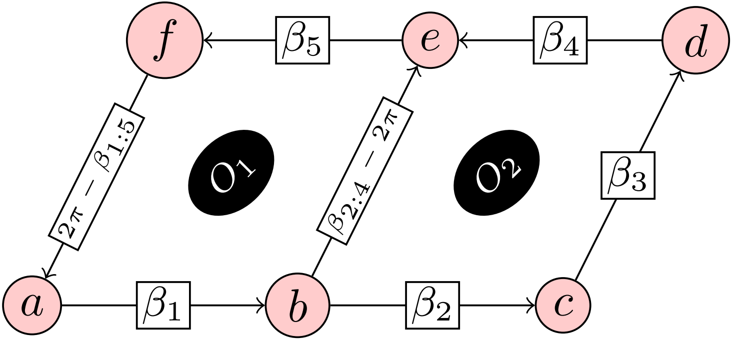

Augmented Graph

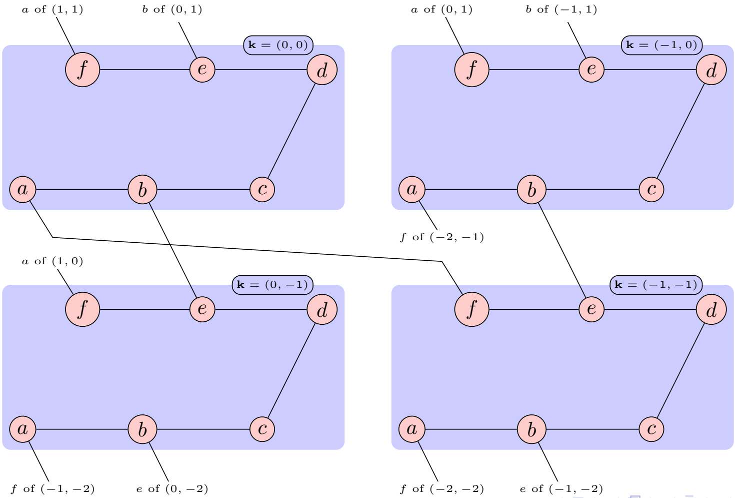

Augmented Graph

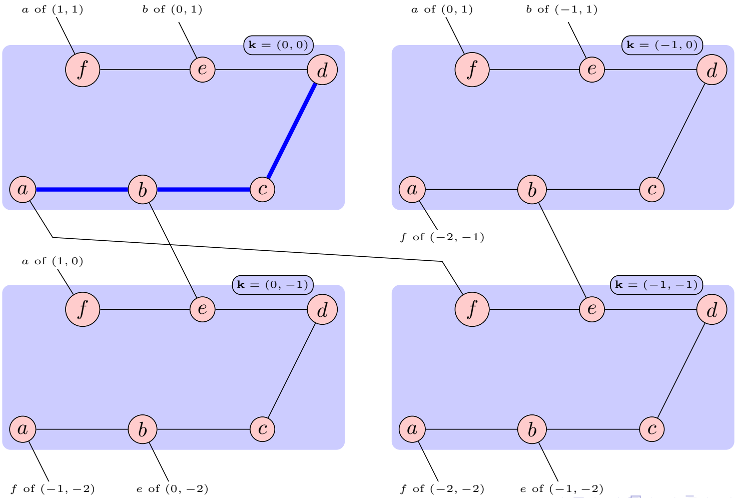

Augmented Graph

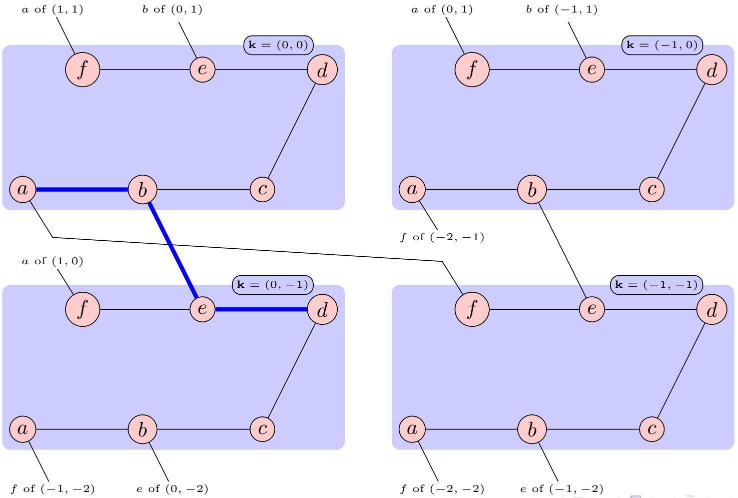

Augmented Graph

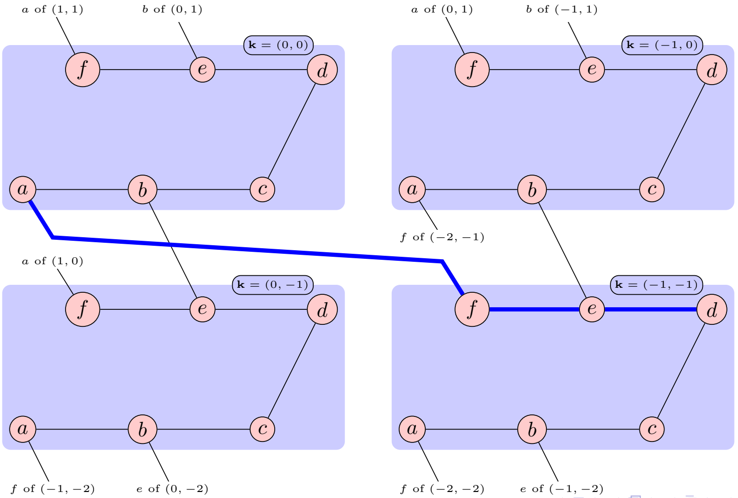

Augmented Graph

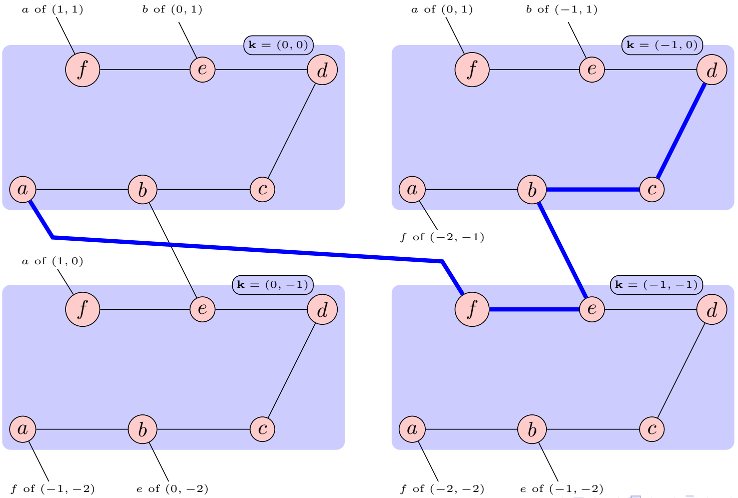

Tracking by Planning

- We test our motion model in a batchmode tracking by detection framework.
- Tracking a person in the visible state leads to a short trajectory that we call a tracklet.
- A conservative threshold is used to terminate the trajectory when the tracking score becomes too low.
- After termination, the same person may be picked up again by the detection algorithm, and tracked to produce associated tracklets.
- After tracklets are obtained, we can link them using both appearance and planning consistency.

Criteria for tracklets linking by planning

- Assume that we have a set of tracklets $\mathcal{T} = \{F_1, \cdots, F_{N_{\mathrm{Tr}}}\}$.

Criteria for tracklets linking by planning

- Assume that we have a set of tracklets $\mathcal{T} = \{F_1, \cdots, F_{N_{\mathrm{Tr}}}\}$.
- Each tracklet is described by 3D point series $F_i = (t_0^i, t_1^i, \mathbf{x}_{t_0^i}^i, \cdots, \mathbf{x}_{t_1^i}^i)$, where t_0^i is the start time of F_i, t_1^i is the end time of F_i and $\mathbf{x}_t^i$ is the object position at time t.

Criteria for tracklets linking by planning

- Assume that we have a set of tracklets $\mathcal{T} = \{F_1, \cdots, F_{N_{\text{Tr}}}\}$.
- Each tracklet is described by 3D point series $F_i = (t_0^i, t_1^i, \mathbf{x}_{t_0^i}^i, \cdots, \mathbf{x}_{t_1^i}^i)$, where t_0^i is the start time of F_i, t_1^i is the end time of F_i and $\mathbf{x}_t^i$ is the object position at time t.
- We then link and extend these tracklets, $\mathcal{T}$, into complete trajectories.

Criteria for tracklets linking by planning

- Assume that we have a set of tracklets $\mathcal{T} = \{F_1, \cdots, F_{N_{\mathrm{Tr}}}\}$.
- Each tracklet is described by 3D point series $F_i = (t_0^i, t_1^i, \mathbf{x}_{t_0^i}^i, \cdots, \mathbf{x}_{t_1^i}^i)$, where t_0^i is the start time of F_i, t_1^i is the end time of F_i and $\mathbf{x}_t^i$ is the object position at time t.
- We then link and extend these tracklets, $\mathcal{T}$, into complete trajectories.
- Let $L_{i,j}$ be the indicator of linking i-th and j-th tracklet:

$$L_{i,j} = \begin{cases} 1 & F_i \to F_j \\ 0 & \text{otherwise} \end{cases} . \tag{9}$$

Criteria for tracklets linking by planning

To link tracklets into plausible goal-directed obstacle-avoiding paths, we design the following criterion for tracking:

$$\max_{L} \epsilon(L) = \sum_{i,j:L_{i,j}=1} [S_{\mathrm{App}}(i,j) + \alpha S_{\mathrm{Plan}}(i,j)] - \beta |L| \tag{10}$$

where

- $S_{\mathrm{App}}(i,j)$ measures appearance similarity between tracklets F_i and F_j,
- $S_{\mathrm{Plan}}(i,j)$ measures 1) how consistent F_i and F_j are with a plausible goal directed path; and 2) how partial occlusion in the gap can be explained by appearance of F_i and F_j.

We seek an approximate solution using Linear Programming.

The criterion is subject to a set of other constraints.

Appearance and motion scores

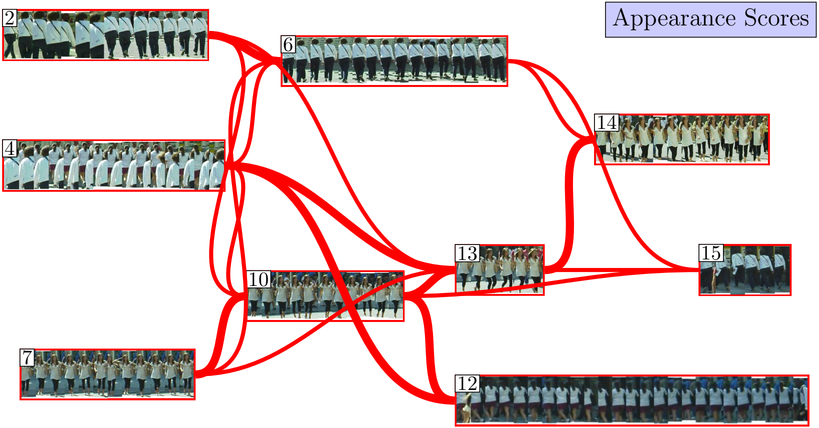

Appearance and motion scores

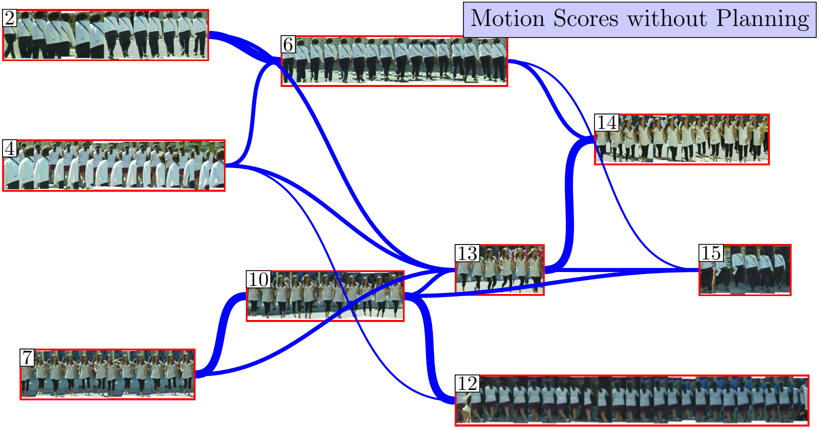

Appearance and motion scores

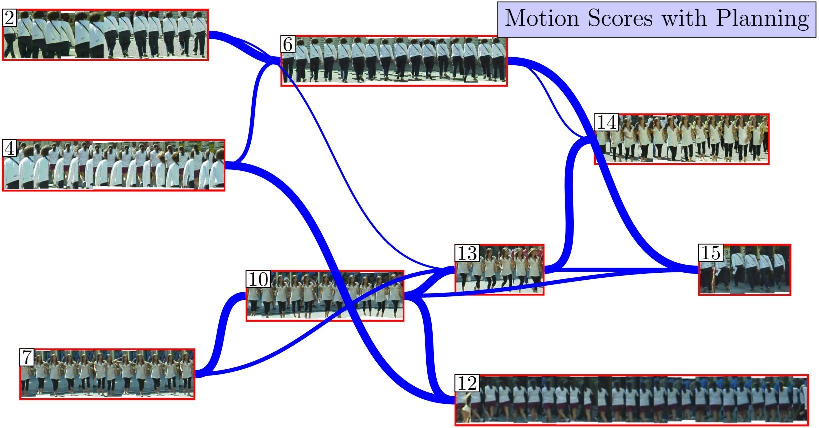

Appearance and motion scores

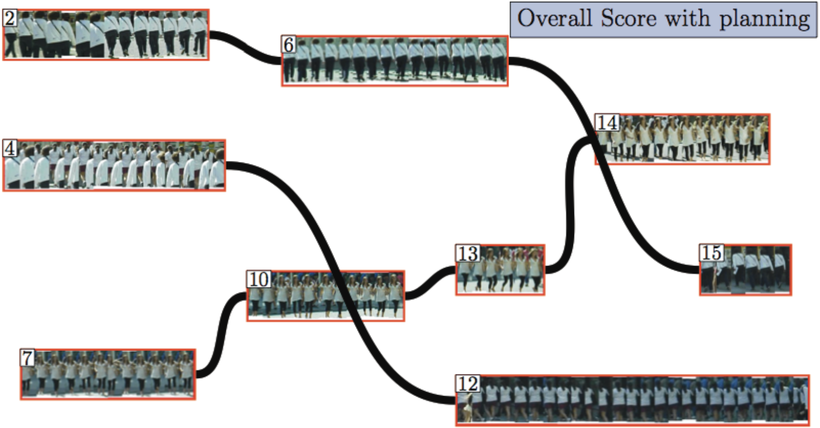

Appearance and motion scores

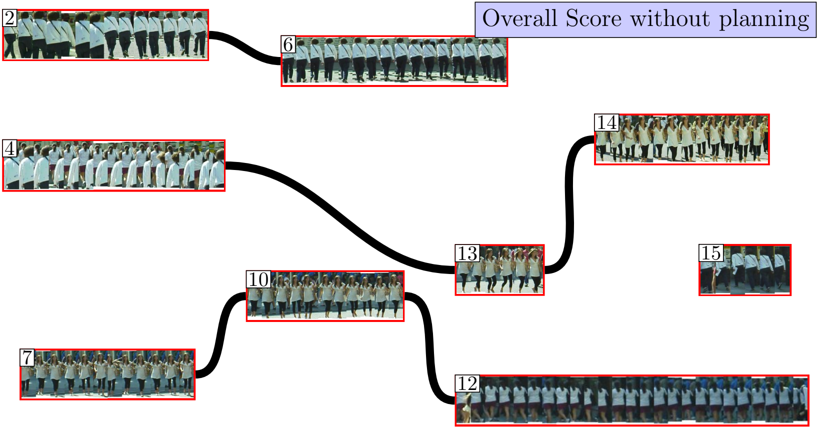

Additional constraints

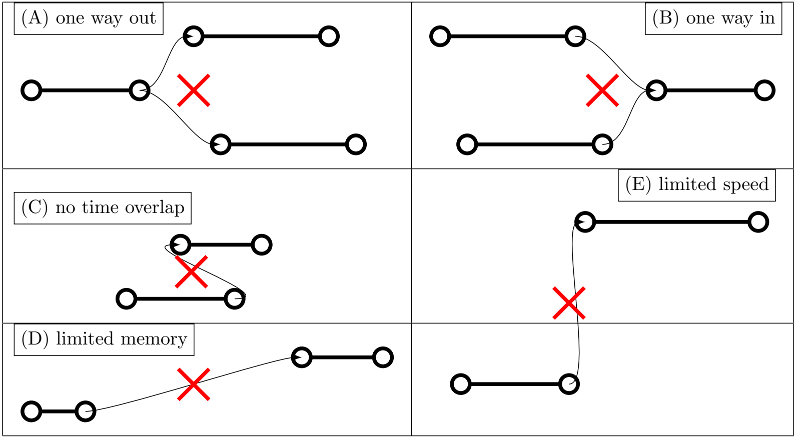

Planning score

- The planning score is given by finding the best planned path to fill the gap between tracklet i and j.
- The best path is compatible with tracklet i and tracklet j geometrically, and allows possible partial matches by appearance during occlusions.
- We use the following score:

$$S_{\mathrm{Plan}}(i,j) = \max_{r \in \mathrm{paths}} -\mathrm{Dist}(r, F_i) - \mathrm{Dist}(r, F_j) + S_{\mathrm{Occl}}(F_i, F_j, r),$$

  where $\mathrm{Dist}(r, F_i)$ is the distance between path r and tracklet F_i and $S_{\mathrm{Occl}}(F_i, F_j, r)$ is the score for picking up the partial occlusions along the gap.
- To reduce computation, we prune paths whose costs are higher than the minimal one above a threshold.

Matching Plans and Partial Occlusion Score

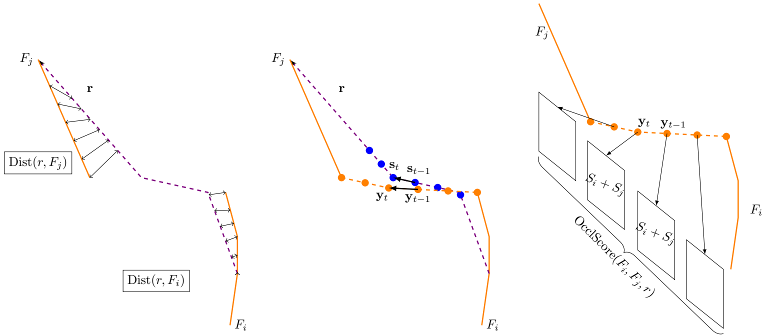

(a) $\text{Dist}(r, F_i)$ (b) Diffusion equation to fill the gap (c) partial occlusion score $\text{OcclScore}(F_i, F_j, r)$

Adaptive appearance model.

- For a pedestrian, we divide his image patch into three parts: head, torso and legs.
- Using a part based representation allows us to reason under partial occlusion.
- For each part k at time t, we collect the color histogram using $8 \times 8 \times 8$ bins, denoted by $\mathbf{p}_t(k)$, and we also collect the histogram of surrounding background, denoted by $\mathbf{q}_t(k)$.
- We use simple color feature instead of more advanced shape features for simplicity and computation efficiency.
- We maintain running means of the histograms as an object model.

Tracklet creation

- We use a detector based on [Felzenszwalb2008] to detect peoples and cars in the current frame.
- To track a person in frame $t+1$ given the models of previous frame, Model_t, we measure two scores
  - *Consistent Score (S1)* to ensure that it is similar to foreground appearance model $\mathbf{f}_t$ and different from $\mathbf{b}_t$,
  - and *Contrast Score (S2)* to ensure that the foreground is different from its surroundings in current frame.
- The appearance score in Linear Program criterion is obtained by testing the appearance model of tracklet i on model of tracklet j and vice-versa.

Experiment Setting

Street Scene To test our algorithm we have collected a video from a moving vehicle in an urban city.

Binocular Sensor The stereo images were collected at 1024×768 resolution and 6 FPS.

3D 3D scene layout/goal estimation, and camera ego-motion computation.

Detection 3D people detection (based on [Felzenszwalb2008]).

Goals We estimated building planes and ground plane in each frame and intersected them to get street side lines. The goals are estimated by intersecting the street side lines, plus infinity points along the street.

Obstacles We only track people, but detect cars as dynamic obstacles. When planning for a specified object, other objects are regarded as obstacles.

Data

	# obj	# frames	# BB	#Occl. BB
seq #1	13	169	1139	471
seq #2	12	60	532	130
seq #3	7	35	210	125
seq #4	4	40	148	51
seq #5	5	112	211	46
seq #6	5	41	170	17
seq #7	2	27	54	16
Total	48	484	2464	856

Test Videos with 3 difficulty levels according to the number of occluded bounding boxes. (BB = Bounding boxes.)

Comparison

- For comparison, we implement two baselines,
  - 1) (LINEAR) tracklet linking without planning, that is, using a straight line as a plan to try to link the gaps, and
  - 2) (LTA) Linear Trajectory Avoidance [Pellegrini2009].

Results — CLEAR Metrics

		miss rate	fa rate	id switch
seq #1	PLAN	**0.413**	0.089	9
	LINEAR	0.442	**0.070**	**8**
	LTA	0.488	0.214	8
seq #2	PLAN	**0.259**	**0.193**	**0**
	LINEAR	0.330	0.199	4
	LTA	0.366	0.310	6
seq #3	PLAN	**0.311**	0.223	**1**
	LINEAR	0.340	**0.200**	2
	LTA	0.476	0.445	6
seq #4	PLAN	**0.176**	**0.00**	**0**
	LINEAR	**0.176**	0.110	**0**
	LTA	0.270	0.212	**0**
seq #5	PLAN	0.137	0.032	**0**
	LINEAR	**0.123**	**0.016**	**0**
	LTA	0.189	0.090	0
seq #6	PLAN	**0.147**	0.194	**0**
	LINEAR	0.153	**0.152**	6
	LTA	0.211	0.394	5
seq #7	PLAN	**0.056**	**0.00**	**0**
	LINEAR	**0.056**	**0.00**	**0**
	LTA	0.203	0.157	**0**

Results

Results

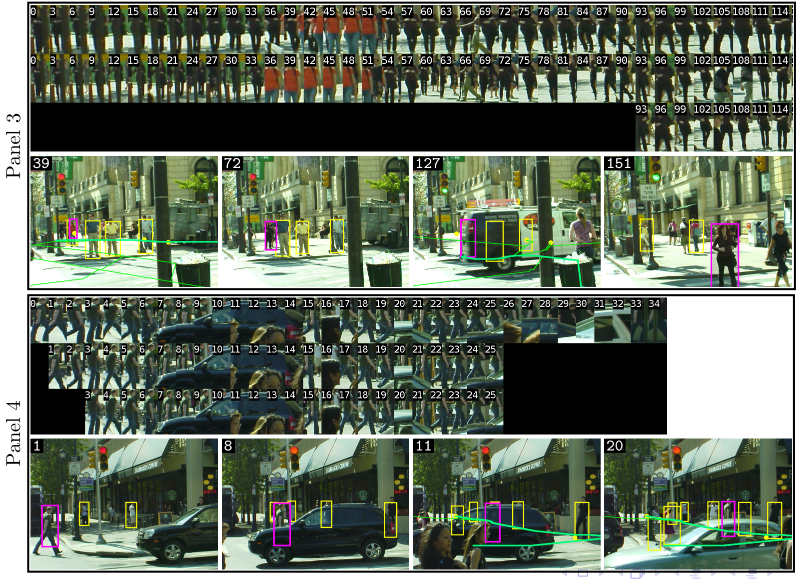
Panel 3
0 3 6 9 12 15 18 21 24 27 30 33 36 39 42 45 48 51 54 57 60 63 66 69 72 75 78 81 84 87 90 93 96 99 102 105 108 111 114
0 3 6 9 12 15 18 21 24 27 30 33 36 39 42 45 48 51 54 57 60 63 66 69 72 75 78 81 84 87 90 93 96 99 102 105 108 111 114
93 96 99 102 105 108 111 114
39
72
127
151
Panel 4
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25
3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25
1
8
11
20

End!

Thanks

hfgong@seas.upenn.edu