Computer Vision II - Lecture 13

Multi-Object Tracking III

03.07.2014

Bastian Leibe
RWTH Aachen
http://www.vision.rwth-aachen.de

leibe@vision.rwth-aachen.de

Course Outline

- Single-Object Tracking
- Bayesian Filtering
  - Kalman filters
  - Particle filters
  - Case studies
- Multi-Object Tracking
  - Introduction
  - MHT, JPDAF
  - Network Flow Optimization
- Articulated Tracking

Topics of This Lecture

- Recap: MHT
- Data Association as Linear Assignment Problem
  - LAP formulation
  - Greedy algorithm
  - Hungarian algorithm
- Tracking as Network Flow Optimization
  - Min-cost network flow
  - Generalizing to multiple frames
  - Complications
  - Formulation

Recap: Multi-Hypothesis Tracking (MHT)

- Ideas
  - Instead of forming a track tree, keep a set of hypotheses that generate child hypotheses based on the associations.
  - Enforce exclusion constraints between tracks and measurements in the assignment.
  - Integrate track generation into the assignment process.
  - After hypothesis generation, merge and prune the current hypothesis set.

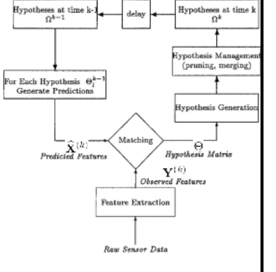

D. Reid, An Algorithm for Tracking Multiple Targets, IEEE Trans. Automatic Control, Vol. 24(6), pp. 843-854, 1979.

Image source: [Cox, IJCV'93]

Recap: Hypothesis Generation

- Create hypothesis matrix of the feasible associations

$$\Theta = \begin{array}{c} \mathbf{x}_1\ \mathbf{x}_2\ \mathbf{x}_{fa}\ \mathbf{x}_{nt} \\ \begin{bmatrix} 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix} \end{array} \begin{array}{c} \mathbf{y}_1 \\ \mathbf{y}_2 \\ \mathbf{y}_3 \\ \mathbf{y}_4 \end{array}$$

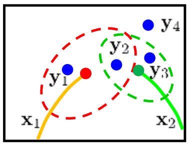

- Interpretation
  - Columns represent tracked objects, rows encode measurements
  - A non-zero element at matrix position (i,j) denotes that measurement $\mathbf{y}_i$ is contained in the validation region of track $\mathbf{x}_j$.
  - Extra column $\mathbf{x}_{fa}$ for association as *false alarm*.
  - Extra column $\mathbf{x}_{nt}$ for association as *new track*.
  - Turn this hypothesis matrix

Recap: Assignments

Z_j	$\mathbf{x}_1$	$\mathbf{x}_2$	$\mathbf{x}_{fa}$	$\mathbf{x}_{nt}$
$\mathbf{y}_1$	0	0	1	0
$\mathbf{y}_2$	1	0	0	0
$\mathbf{y}_3$	0	1	0	0
$\mathbf{y}_4$	0	0	0	1

- Impose constraints
  - A measurement can originate from only one object.

  $\Rightarrow$ Any row has only a single non-zero value.
  - An object can have at most one associated measurement per time step.

  $\Rightarrow$ Any column has only a single non-zero value, except for $\mathbf{x}_{fa}$, $\mathbf{x}_{nt}$

Recap: Calculating Hypothesis Probabilities

- **Probabilistic formulation**
  - It is straightforward to enumerate all possible assignments.
  - However, we also need to calculate the probability of each child hypothesis.
  - This is done recursively:

$$p(\Omega_j^{(k)}|\mathbf{Y}^{(k)}) = p(Z_j^{(k)}, \Omega_{p(j)}^{(k-1)}|\mathbf{Y}^{(k)})$$

$$\stackrel{Bayes}{=} \eta p(\mathbf{Y}^{(k)}|Z_j^{(k)}, \Omega_{p(j)}^{(k-1)}) p(Z_j^{(k)}, \Omega_{p(j)}^{(k-1)})$$

$$= \eta p(\mathbf{Y}^{(k)}|Z_j^{(k)}, \Omega_{p(j)}^{(k-1)}) p(Z_j^{(k)}|\Omega_{p(j)}^{(k-1)}) p(\Omega_{p(j)}^{(k-1)})$$

Normalization factor | Measurement likelihood | Prob. of assignment set | Prob. of parent

Recap: Measurement Likelihood

- **Use KF prediction**
  - Assume that a measurement $\mathbf{y}_i^{(k)}$ associated to a track $\mathbf{x}_j$ has a Gaussian pdf centered around the measurement prediction $\hat{\mathbf{x}}_j^{(k)}$ with innovation covariance $\widehat{\mathbf{\Sigma}}_j^{(k)}$.
  - Further assume that the pdf of a measurement belonging to a new track or false alarm is uniform in the observation volume W (the sensor's field-of-view) with probability W^{-1}.
  - Thus, the measurement likelihood can be expressed as

$$p\left(\mathbf{Y}^{(k)}|Z_j^{(k)}, \Omega_{p(j)}^{(k-1)}\right) = \prod_{i=1}^{M_k} \mathcal{N}\left(\mathbf{y}_i^{(k)}; \hat{\mathbf{x}}_j, \widehat{\mathbf{\Sigma}}_j^{(k)}\right)^{\delta_i} W^{-(1-\delta_i)}$$

$$= W^{-(N_{fal}+N_{new})} \prod_{i=1}^{M_k} \mathcal{N}\left(\mathbf{y}_i^{(k)}; \hat{\mathbf{x}}_j, \widehat{\mathbf{\Sigma}}_j^{(k)}\right)^{\delta_i}$$

Recap: Probability of an Assignment Set

$$p(Z_j^{(k)}|\Omega_{p(j)}^{(k-1)})$$

- **Composed of three terms**
  1. Probability of the number of tracks N_{det}, N_{fal}, N_{new}
     - Assumption 1: N_{det} follows a binomial distribution

$$p(N_{det}|\Omega_{p(j)}^{(k-1)}) = \binom{N}{N_{det}} p_{det}^{N_{det}} (1-p_{det})^{(N-N_{det})}$$

where N is the number of tracks in the parent hypothesis

   - Assumption 2: N_{fal} and N_{new} both follow a Poisson distribution with expected number of events $\lambda_{fal}W$ and $\lambda_{new}W$

$$p(N_{det}, N_{fal}, N_{new}|\Omega_{p(j)}^{(k-1)}) = \binom{N}{N_{det}} p_{det}^{N_{det}} (1-p_{det})^{(N-N_{det})} \cdot \mu(N_{fal}; \lambda_{fal}W) \cdot \mu(N_{new}; \lambda_{new}W)$$

Recap: Probability of an Assignment Set

2. Probability of a specific assignment of measurements
   - Such that $M_k = N_{det} + N_{fal} + N_{new}$ holds.
   - This is determined as 1 over the number of combinations

$$\binom{M_k}{N_{det}} \binom{M_k - N_{det}}{N_{fal}} \binom{M_k - N_{det} - N_{fal}}{N_{new}}$$

3. Probability of a specific assignment of tracks
   - Given that a track can be either *detected* or not *detected*.
   - This is determined as 1 over the number of assignments

$$\frac{N!}{(N-N_{det})!} \binom{N-N_{det}}{N_{det}}$$

$\Rightarrow$ When combining the different parts, many terms cancel out!

Laser-based Leg Tracking using MHT

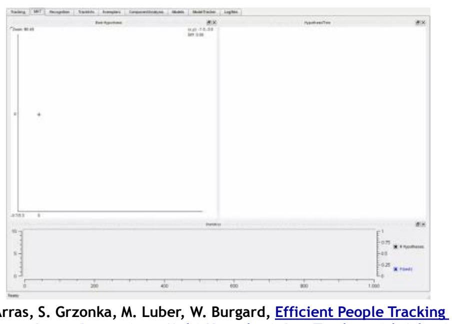

K. Arras, S. Grzonka, M. Luber, W. Burgard, Efficient People Tracking in Laser Range Data using a Multi-Hypothesis Leg-Tracker with Adaptive Occlusion Probabilities, ICRA'08.

video source: Social Robotics Lab, Univ. Freiburg

Laser-based People Tracking using MHT

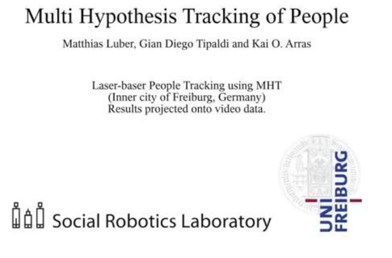

video source: Social Robotics Lab, Univ. Freiburg

Topics of This Lecture

- Recap: MHT
- **Data Association as Linear Assignment Problem**
  - LAP formulation
  - Greedy algorithm
  - Hungarian algorithm
- Tracking as Network Flow Optimization
  - Min-cost network flow
  - Generalizing to multiple frames
  - Complications
  - Formulation

Back to Data Association...

- Goal: Match detections across frames

Data Association

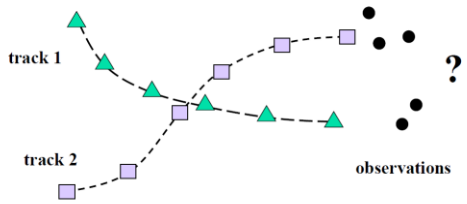

- Main question here
  - How to determine which measurements to add to which track?
  - Today: consider this as a matching problem

Linear Assignment Formulation

- Form a matrix of pairwise similarity scores
- Similarity could be
  - based on motion prediction
  - based on appearance
  - based on both

Frame $t+1$ (columns), Frame t (rows)

			
	0.11	**0.95**	0.23
	0.85	0.25	**0.89**
	**0.90**	0.12	0.81

- Goal
  - Choose one match from each row and column to maximize the sum of scores

Linear Assignment Formulation

- Example: Similarity based on motion prediction
  - Predict motion for each trajectory and assign scores for each measurement based on inverse (Mahalanobis) distance, such that closer measurements get higher scores.

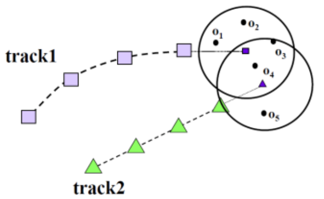

	ai1	ai2
1	3.0	
2	5.0	
3	**6.0**	1.0
4	9.0	**8.0**
5		3.0

  - Choose at most one match in each row and column to maximize sum of scores

Linear Assignment Problem

- Formal definition
  - Maximize $\sum_{i=1}^{N}\sum_{j=1}^{M} w_{ij} z_{ij}$

    subject to
    $$\sum_{j=1} z_{ij} = 1;\ i = 1,2,\ldots,N$$
    $$\sum_{i=1} z_{ij} = 1;\ j = 1,2,\ldots,M$$
    $$z_{ij} \in \{0,1\}$$
    Those constraints ensure that Z is a permutation matrix
  - The permutation matrix constraint ensures that we can only match up one object from each row and column.
  - Note: Alternatively, we can minimize cost rather than maximizing weights. $\arg\min_{z_{ij}} \sum_{i=1}^{N}\sum_{j=1}^{M} c_{ij} z_{ij}$

Greedy Solution to LAP

	1	2	3	4	5
1	0.95	0.76	0.62	0.41	0.06
2	0.23	0.46	0.79	0.94	0.35
3	0.61	0.02	0.92	0.92	0.81
4	0.49	0.82	0.74	0.41	0.01
5	0.89	0.44	0.18	0.89	0.14

- Greedy algorithm
  - Find the largest score
  - Remove scores in same row and column from consideration
  - Repeat
- Result: score =

Greedy Solution to LAP

	1	2	3	4	5
1	0.95	0.76	0.62	0.41	0.06
2	0.23	0.46	0.79	0.94	0.35
3	0.61	0.02	0.92	0.92	0.81
4	0.49	0.82	0.74	0.41	0.01
5	0.89	0.44	0.18	0.89	0.14

- Greedy algorithm
  - Find the largest score
  - Remove scores in same row and column from consideration
  - Repeat
- Result: score = 0.95 + 0.94 + 0.92 + 0.82 + 0.14 = 3.77

Is this the best we can do?

Greedy Solution to LAP

	1	2	3	4	5
1	0.95	0.76	0.62	0.41	0.06
2	0.23	0.46	0.79	0.94	0.35
3	0.61	0.02	0.92	0.92	0.81
4	0.49	0.82	0.74	0.41	0.01
5	0.89	0.44	0.18	0.89	0.14

Greedy solution
score = 3.77

	1	2	3	4	5
1	0.95	0.76	0.62	0.41	0.06
2	0.23	0.46	0.79	0.94	0.35
3	0.61	0.02	0.92	0.92	0.81
4	0.49	0.82	0.74	0.41	0.01
5	0.89	0.44	0.18	0.89	0.14

Optimal solution
score = 4.26

- Discussion
  - Greedy method is easy to program, quick to run, and yields "pretty good" solutions in practice.
  - But it often does not yield the optimal solution.

Optimal Solution

- Hungarian Algorithm
  - There is an algorithm called Kuhn-Munkres or "Hungarian" algorithm specifically developed to efficiently solve the linear assignment problem.
  - Reduces assignment problem to bipartite graph matching.
  - When starting from an $N \times N$ matrix, it runs in $\mathcal{O}(N^3)$.

  ⇒ If you need LAP, you should use it.
- In the following
  - Look at other algorithms that generalize to multi-frame (>2 frames) problems.

  ⇒ Min-Cost Network Flow

Topics of This Lecture

- Recap: MHT
- Data Association as Linear Assignment Problem
  - LAP formulation
  - Greedy algorithm
  - Hungarian algorithm
- Tracking as Network Flow Optimization
  - Min-cost network flow
  - Generalizing to multiple frames
  - Complications
  - Formulation

Min-Cost Flow

- Small example

	1	2	3
1	3	2	3
2	2	1	3
3	4	5	1

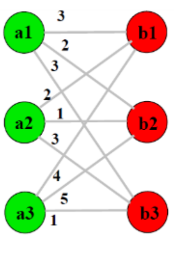

RWTH AACHEN UNIVERSITY

Computer Vision II, Summer'14

Min-Cost Flow

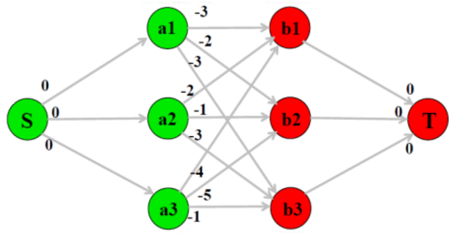

- Conversion into flow graph
  - Transform weights into costs $c_{ij} \quad \alpha \quad w_{ij}$
  - Add source/sink nodes with 0 cost.
  - Directed edges with a capacity of 1.

Slide credit: Robert Collins

B. Leibe

RWTH AACHEN UNIVERSITY

Computer Vision II, Summer'14

Min-Cost Flow

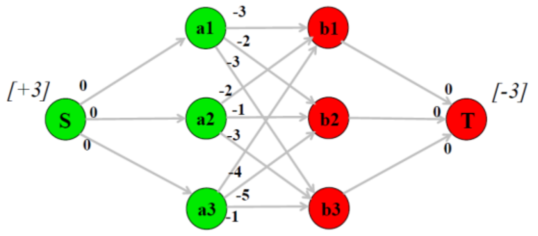

- Conversion into flow graph
  - Pump N units of flow from source to sink.
  - Internal nodes pass on flow ($\sum$ flow in $= \sum$ flow out).
  - $\Rightarrow$ Find the optimal paths along which to ship the flow.

Slide credit: Robert Collins

B. Leibe

RWTH AACHEN UNIVERSITY

Computer Vision II, Summer'14

Min-Cost Flow

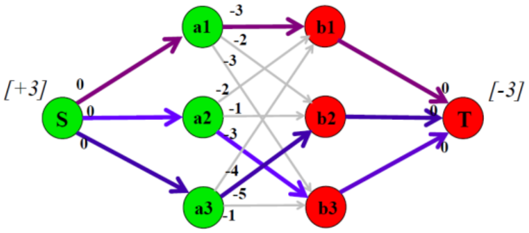

- Conversion into flow graph
  - Pump N units of flow from source to sink.
  - Internal nodes pass on flow ($\sum$ flow in $= \sum$ flow out).
  - $\Rightarrow$ Find the optimal paths along which to ship the flow.

Slide credit: Robert Collins

B. Leibe

RWTH AACHEN UNIVERSITY

Computer Vision II, Summer'14

Min-Cost Flow

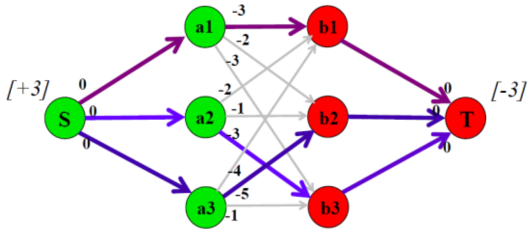

- Solving the Min-Cost Flow problem
  - There are standard algorithms for efficiently solving min-cost network flow
  - E.g., push-relabel or successive shortest path algorithms

Slide credit: Robert Collins

B. Leibe

RWTH AACHEN UNIVERSITY

Computer Vision II, Summer'14

Min-Cost Flow

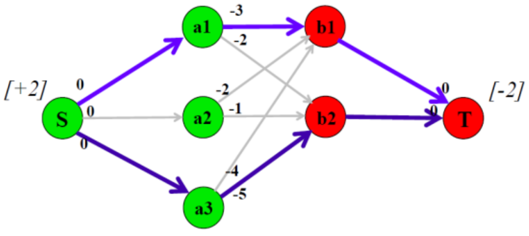

- Nice property
  - Min-cost formalism readily generalizes to matching sets with unequal sizes.

Slide credit: Robert Collins

B. Leibe

RWTH AACHEN UNIVERSITY

Computer Vision II, Summer'14

Using Network Flow for Tracking

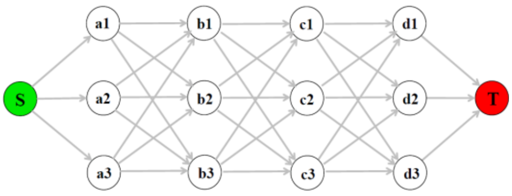

- Approach
  - Seek a globally optimal solution by considering observations over all frames in "batch mode".
  - $\Rightarrow$ Extend two-frame min-cost formulation by adding observations from all frames into the network.

Slide credit: Robert Collins

B. Leibe

Using Network Flow for Tracking

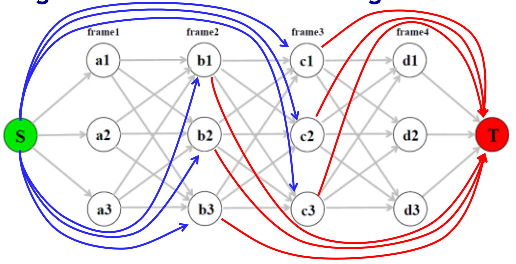

- **Complication 1**
  - Tracks can start later than frame1 (and end earlier than frame4)

  ⇒ Connect the source and sink nodes to all intermediate nodes.

Using Network Flow for Tracking

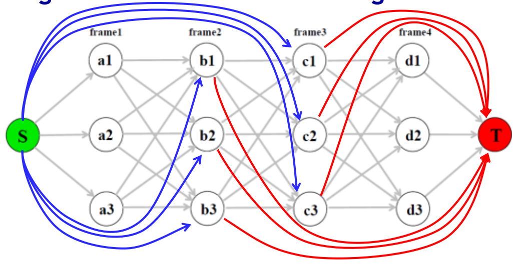

- **Complication 2**
  - Trivial solution: zero cost!

Using Network Flow for Tracking

- **Solution**
  - Divide each detection into 2 nodes

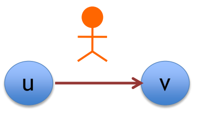

Detection edge

$$C_i = \log \frac{\beta_i}{1-\beta_i} \longleftarrow$$ Probability that detection i is a false alarm

Zhang, Li, Nevatia, Global Data Association for Multi-Object Tracking using Network Flows, CVPR'08.

Network Flow Approach

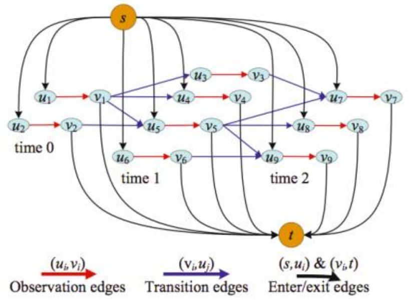

Zhang, Li, Nevatia, Global Data Association for Multi-Object Tracking using Network Flows, CVPR'08.

Network Flow Approach: Illustration

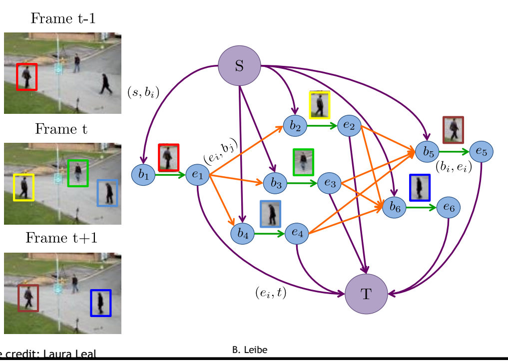

Min-Cost Formulation

- **Objective Function**

$$\mathcal{T}* = \underset{\mathcal{T}}{\operatorname{argmin}} \sum_i C_{in,i} f_{in,i} + \sum_i C_{i,out} f_{i,out}$$
$$+ \sum_{i,j} C_{i,j} f_{i,j} + \sum_i C_i f_i$$

- **subject to**
  - Flow conservation at all nodes

$$f_{in,i} + \sum_j f_{j,i} = f_i = f_{out,i} + \sum_j f_{i,j} \; \forall i$$

  - Edge capacities

$$f_i \leq 1$$

Min-Cost Formulation

- Objective Function

$$\mathcal{T}* = \underset{\mathcal{T}}{\text{argmin}} \sum_i C_{in,i} f_{in,i} + \sum_i C_{i,out} f_{i,out}$$

$$+ \sum_{i,j} C_{i,j} f_{i,j} + \sum_i C_i f_i$$

$$C_i = -log(P_i)$$

- Equivalent to Maximum A-Posteriori formulation

$$\mathcal{T}* = \underset{\mathcal{T}}{\text{argmax}} \prod_i P(\mathbf{o}_i|\mathcal{T})P(\mathcal{T})$$

Min-Cost Formulation

- Objective Function

IN OUT

$$\mathcal{T}* = \underset{\mathcal{T}}{\text{argmin}} \sum_i C_{in,i} f_{in,i} + \sum_i C_{i,out} f_{i,out}$$

TRANSITION

$$+ \sum_{i,j} C_{i,j} f_{i,j} + \sum_i C_i f_i$$

Likelihood of the detection

$$C_i = -log(P_i)$$

- Equivalent to Maximum A-Posteriori formulation

$$\mathcal{T}* = \underset{\mathcal{T}}{\text{argmax}} \prod_i P(\mathbf{o}_i|\mathcal{T})P(\mathcal{T})$$

$$P(\mathcal{T}) = \prod_{T_k \in \mathcal{T}} P(T_k)$$

Independence assumption + Markov

Network Flow Solutions

- Push-relabel method
  - Zhang, Li and Nevatia, **"Global Data Association for Multi-Object Tracking Using Network Flows,"** CVPR 2008.
- Successive shortest path algorithm
  - Berclaz, Fleuret, Turetken and Fua, **"Multiple Object Tracking using K-shortest Paths Optimization,"** IEEE PAMI, Sep 2011.
  - Pirsiavash, Ramanan, Fowlkes, **"Globally Optimal Greedy Algorithms for Tracking a Variable Number of Objects"**, CVPR'11.
  - **These both include approximate dynamic programming solutions**

References and Further Reading

- The original network flow tracking paper
  - **Zhang, Li, Nevatia, Global Data Association for Multi-Object Tracking using Network Flows, CVPR'08.**
- Extensions and improvements
  - **Berclaz, Fleuret, Turetken, Fua, Multiple Object Tracking using K-shortest Paths Optimization, IEEE PAMI, Sep 2011. (code)**
  - **Pirsiavash, Ramanan, Fowlkes, Globally Optimal Greedy Algorithms for Tracking a Variable Number of Objects, CVPR'11.**
- A recent extension to incorporate social walking models
  - **L. Leal-Taixe, G. Pons-Moll, B. Rosenhahn, Everybody Needs Somebody: Modeling Social and Grouping Behavior on a Linear Programming Multiple People Tracker, ICCV Workshops 2011.**