Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Discrete-Continuous Optimization for Multi-Target Tracking

Anton Andriyenko[1] Konrad Schindler[2] Stefan Roth[1]

[1]Department of Computer Science, TU Darmstadt, Germany [2]Photogrammetry and Remote Sensing Group, ETH Zürich, Switzerland

TECHNISCHE UNIVERSITÄT DARMSTADT

Motivation and Overview

Following the *tracking-by-detection* approach, multi-target tracking involves two tightly coupled challenges:

1. Data Association

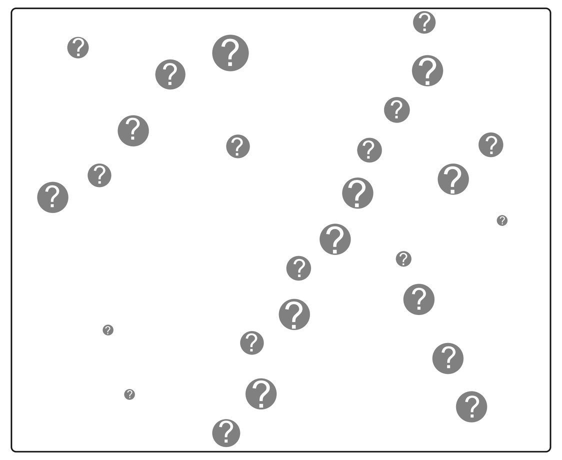

What is the source of each observation? (discrete problem)

2. Trajectory Estimation

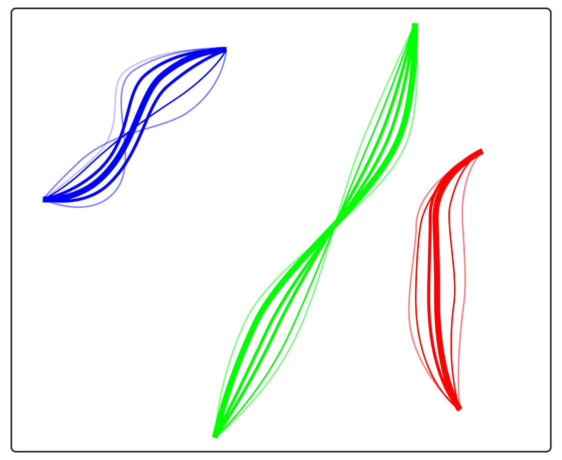

What are the actual spatio-temporal motion patterns of targets? (continuous problem)

Most previous work focused mainly on one aspect.

Our approach: Combine both problems into a single discrete-continuous energy; solve each aspect efficiently in its natural domain.

Setting

Given:

- a set of target hypotheses (detections) $\mathbf{D}$
- a set of trajectory hypotheses (models) $\mathcal{T} = \{\mathcal{T}_1, \ldots, \mathcal{T}_N\}$
- a set of labels $\mathbf{L} = \{1, \ldots, N\} \cup \varnothing$ ($\varnothing \equiv$ false alarms)

Goal:
Assign detections to models and improve trajectories by alternating between discrete and continuous optimization.

Continuous Trajectory Model

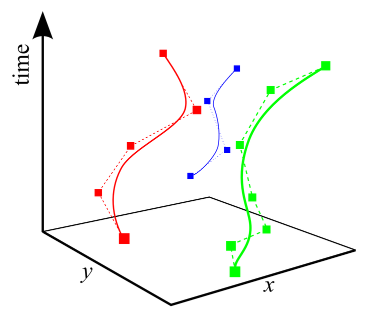

2D cubic B-splines with explicit start (s) and end (e) points.

$$\mathcal{T}_i : t \in [\mathrm{s}_i, \mathrm{e}_i] \to (x, y)^\mathsf{T} \in \mathbb{R}^2$$

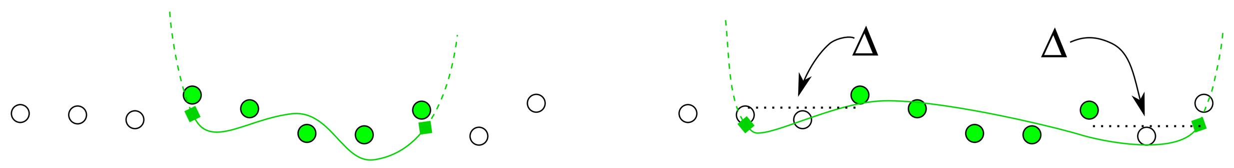

A support margin Δ is added for better spline behavior.

Convex Continuous Optimization

Subproblem 2: **Given data association**, fit parametrized trajectory models; perform **weighted least squares** on each active trajectory:

$$E_\mathrm{f}^{\text{te}}(\mathcal{T}_i) = \sum_t \sum_j c_j^t \cdot \left\| p_j^t - \mathcal{T}_i(t) \right\|^2 \to \min . \qquad (1)$$

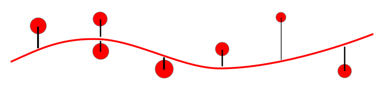

Graph-Cuts-Based Discrete Optimization

Subproblem 1: **Given trajectories**, perform data association (usually more challenging); we cast it as a **multi-labeling problem**

$$E_\mathcal{T}^{\text{da}}(\mathrm{f}) = \sum_d U(\mathrm{f}_d, \mathcal{T}) + \sum_{(d,d')} S(\mathrm{f}_d, \mathrm{f}_{d'}) \to \min \qquad (2)$$

with the data term

$$U(1, \mathcal{T}) = c_j^t \cdot \|p_j^t - \mathcal{T}_1(t)\|^2 \qquad (3)$$

and the generalized Potts smoothness term

$$S(\mathrm{f}_{d_t}, \mathrm{f}_{d'_{t+1}}) = \eta \cdot \delta[\mathrm{f}_{d_t} - \mathrm{f}_{d'_{t+1}}]. \qquad (4)$$

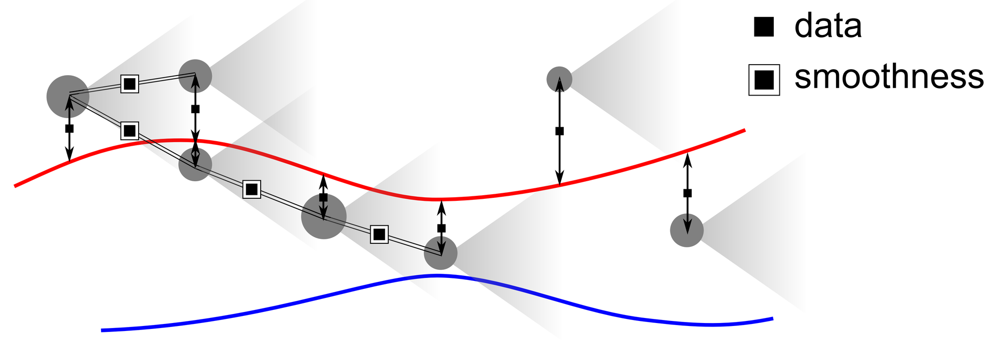

The energy (2) is **submodular** and can be minimized efficiently by **α-expansion**.

Discrete-Continuous Energy

A naive combination of (1) and (2) will not work well.
Challenge: How to incorporate a regularizer and higher-order terms?

Formulate the problem with a single discrete-continuous energy with label cost:

$$E(\mathcal{T}, \mathrm{f}) = \sum_d U(\mathrm{f}_d, \mathcal{T}) + \sum_{(d,d')} S(\mathrm{f}_d, \mathrm{f}_{d'}) + \kappa \cdot h_\mathrm{f}(\mathcal{T}). \qquad (5)$$

- For $\kappa = 0$ minimizing E *w.r.t.* $\mathcal{T}$ amounts to **convex** least squares optimization, *cf.* Eq. (1).
- Minimizing E *w.r.t.* f amounts to solving a multi-label **pairwise MRF**.

Multi-Target Tracking with Label Cost

Following the recent work of Delong *et al.* [1] we integrate the **label cost** $h(\mathcal{T})$ into our energy. It naturally handles

- Regularization: Avoids overfitting by enforcing a constant penalty on each active model.
- Model assessment: "Good" trajectories are favored while implausible ones are penalized.

Trajectory Assessment

The full label cost consists of five components:

$$h_\mathrm{f}(\mathcal{T}) = \sum_i \lambda h_i^{\text{dyn}} + \nu h_i^{\text{per}} + \xi h_i^{\text{fid}} + \zeta h_i^{\text{col}} + h_i^{\text{reg}}. \qquad (6)$$

Dynamics	Persistence & Length	Higher-Order Data Fidelity	Mutual Exclusion

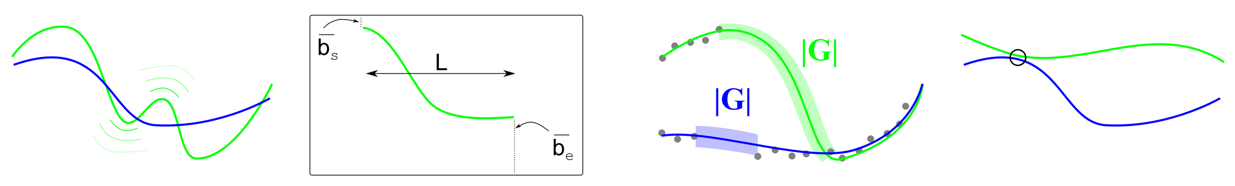

$\max_r C(r, 1)$	$\overline{b_s} + \overline{b_e} + L^{-1}$	$\sum_k \lvert\mathbf{G}_k\rvert^3$	min dist

Limitations. The label cost $h(\mathcal{T})$ depends on trajectory $\mathcal{T}$ $\to$ continuous optimization *w.r.t.* $\mathcal{T}$ is non-trivial.
Perform sanity check after each continuous optimization (1):

if $E_{\text{new}} \leq E_{\text{old}}$ accept new trajectory,
else discard this step.

Algorithm in a Nutshell

Detections → Initial Models → Relabeling → Refitting

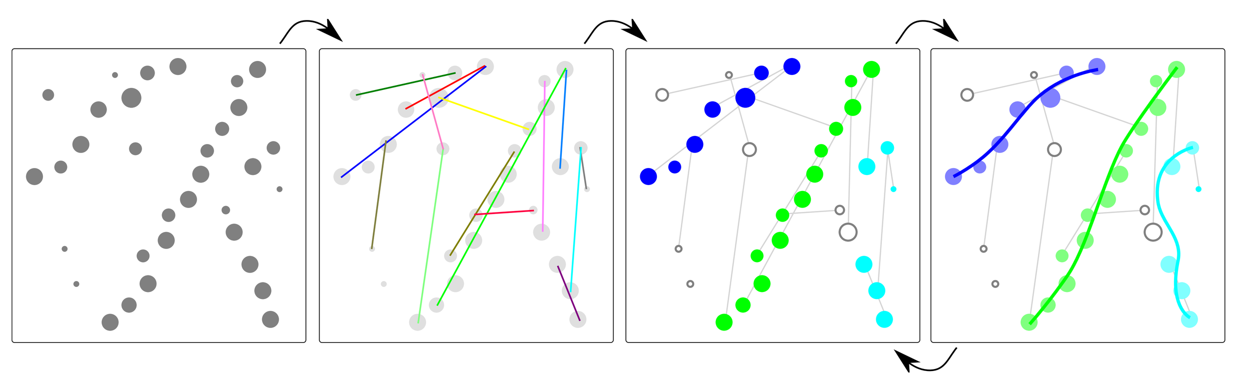

Hypotheses Maintenance

At each iteration new hypotheses are generated

- randomly
- by extending or merging existing ones

Unused models are removed after a few iterations.

Experiments

- Publicly available datasets: TUD and PETS'09 [2].
- Standard CLEAR MOT metrics.
- Evaluated fairly by the PETS organizers.

Averaged Baseline Comparison

TUD+PETS	MOTA↑	MOTP↑	FPR↓	FNR↓	ID Sw.↓
Detector	–	–	39	27.2	–
RANSAC	39.8	76.0	2.4	57.6	12.8
RANSAC w/ GT	62.0	76.7	9.7	28.2	7.0
Our method	71.4	74.7	4.4	24.1	7.0

Comparison to Other Methods

PETS S2.L1	MOTA	MOTP	MODA	MODP
ILP [3]	82 %	56 %	85 %	57 %
Particle Filtering [4]	75 %	60 %	89 %	60 %
Our method	89 %	56 %	91 %	57 %

Qualitative Results

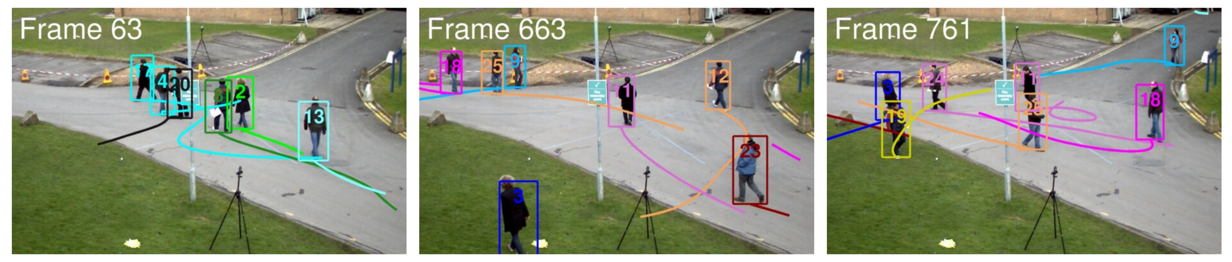

Conclusion

- We presented a **discrete-continuous energy** that combines data association and trajectory estimation.
- Our formulation captures many **desirable properties** of multi-target tracking.
- High-order terms are integrated through the **label cost**.
- By keeping the continuous part **convex** and the discrete part **submodular**, strong minima are found efficiently.
- Our **source code** is freely available at: `goo.gl/rkKXN`.

References

[1] A. Delong, A. Osokin, H. Isack, and Y. Boykov. Fast approximate energy minimization with label costs. *IJCV*, 2011.

[2] J.M. Ferryman and A. Shahrokni. PETS2009. In *Winter-PETS*, 2009.

[3] J. Berclaz, F. Fleuret, and P. Fua. Multiple object tracking using flow linear programming. In *Winter-PETS*, 2009.

[4] M. Breitenstein, F. Reichlin, B. Leibe, E. Koller-Meier, and L. Van Gool. Online multiperson tracking-by-detection from a single, uncalibrated camera. *PAMI'09*.

visual inference

Acknowledgements: This work was gratefully supported by the German Research Foundation (DFG) within the GRK 1362 "Cooperative, Adaptive and Responsive Monitoring of Mixed Mode Environments".
Further we thank Olga Veksler, Andrew Delong and Anton Osokin for publicly releasing GCO, Stefan Walk for providing detector code and Anna Ellis for evaluating our results.

/visinf @visinf goo.gl/rkKXN