

Cross-Entropy Based Data Association for Multi-Target Tracking

Daniel Sigalov Nahum Shimkin

Technion - Israel Institute of Technology

Efficient Monte Carlo:
A Conference on the Occasion of Reuven Rubinstein's 70th Birthday
Thiele Center, July 2008

Outline

- Target Tracking and Data Association
- Model and Problem Formulation
- Cross Entropy Based Data Association
- Experimental Results and Discussion

Target Tracking

The problem:

- **Track multiple targets moving in 2D/3D space,** under conditions of
 - Noisy position (and/or velocity) measurement
 - Clutter ("False alarms")
 - Missed detections
 - Unknown number of targets

Major applications:

- Air traffic control
- Surveillance
- Microscopic object tracking

Illustration (1)

Scan 1: 10 target detections, ≈ 90 false observations.

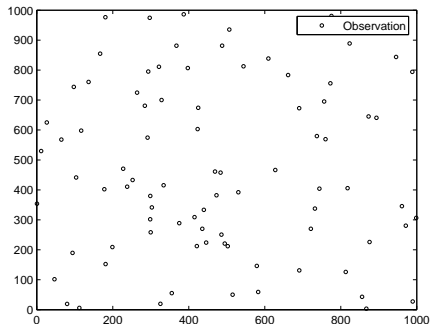


Illustration (2)

Scans 1 + 2

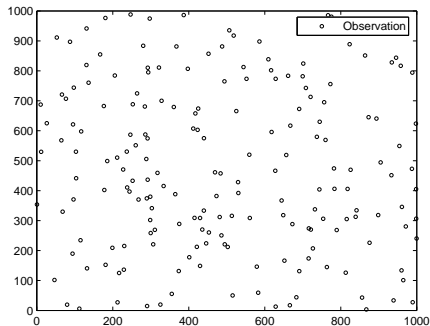


Illustration (3)

Scans 1 to 10


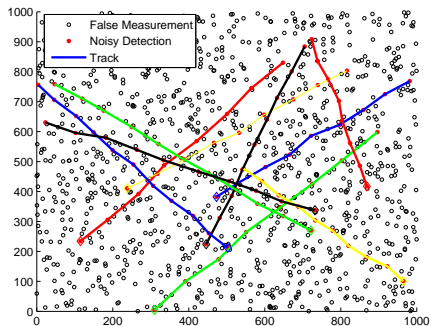


Illustration (4)


Actual Scenario: Tracks and Clutter



- Our goal: Identify the actual targets and their paths, by exploiting regularities in the target motion.

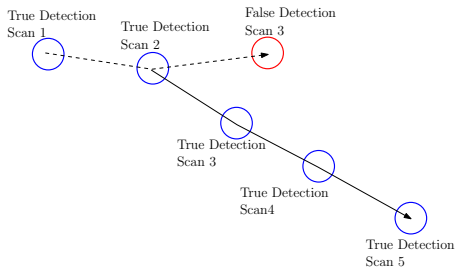
Data Association

- Data association (matching measurements to tracks) is the primary task in any Multi-Target Tracking system.
- The following (true scenario) illustrates the possible consequences of simple-minded data association.



Multi-Scan Data Association

- In **single scan** Data association, each scan is handled separately and sequentially when acquired.
- In **multi-scan** data association, several scans are accumulated and processed together to obtain a better fit.



General Setup

- Consider an unknown number (K) of targets moving across a surveillance region \mathcal{R} , scanned by a single sensor.
- T consecutive scans are considered at a time.
- Target generation and disappearance:
 - At each scan period, the number of *new* targets that appear is distributed as $\text{Pois}(\lambda_b V)$.
 - New targets appears at a random positions in the surveillance region.
 - An existing target disappears at each sampling time with probability p_z , and persists with probability $1 - p_z$.

Target-Motion Model

- Each target moves according to a known dynamic model, assumed to be linear:

$$x_{t_{i+1}} = A_i x_{t_i} + w_{t_i}$$

Here t_i is the time of the i -th scan, and w is a discrete-time zero mean white Gaussian process noise with known covariance matrix Q .

Observation and Sensor Models

The sensor –

- Detects each target with probability P_d .
- Generates n_t clutter measurements (false alarms) at scan t uniformly distributed across \mathcal{R} , where $n_t \sim \text{PoiSS}(\lambda_f V)$.

Observation and Sensor Models(cont.)

Observations:

- If y_t^j is the j -th observation at time t , then

$$y_t^j = \begin{cases} Cx_t^k + v_t & j\text{-th observation is from target } k \\ \text{random position} & \text{otherwise (clutter)} \end{cases}$$

v_t is a discrete-time zero mean white Gaussian measurement noise with known covariance matrix R independent of w_t .


- Each observation carries a cartesian position and time-tag - (x, y, t) .
- Let $Y_t = \{y_t^j : j = 1, \dots, n_t\}$ be the set of observations at time t and $Y_{1:T} = \bigcup_{t=1}^T Y_t$ be the set of all observations.

Solution Space

Let $\Omega = \{\omega\}$ be the set of **partitions** of Y such that

- $\omega = \{\tau_0, \tau_1, \dots, \tau_K\}$, with $K \geq 0$.
- τ_0 is the set of false alarms, and each τ_k , $k \geq 1$ is a track, namely a sequence of measurements from *different* scans. It is further required that $|\tau_k| \geq L_{\min} \geq 2$.

A simple example:



Optimality Criterion

We will be looking for an optimal solution in the MAP sense

$$\omega^* = \arg \max_{\omega \in \Omega} \mathbb{P} \{ \omega \mid Y_{1:T} \} = \arg \max_{\omega \in \Omega} p(Y_{1:T} \mid \omega) \cdot \mathbb{P} \{ \omega \}$$

- $\mathbb{P} \{ \omega \}$ is the prior probability to obtain a given partition $\omega = \{ \tau_0, \tau_1, \dots, \tau_K \}$. It may be computed (as function of $P_d, p_z, \lambda_b, \lambda_f$) independently of the measurement positions.
- $p(Y_{1:T} \mid \omega)$ is the likelihood to obtain a set of observations given the partition ω . It may be computed using Kalman filtering equations.


The Posterior Probability – Summary

The posterior probability of a partition ω computes as

$$\mathbb{P}\{\omega \mid Y_{1:T}\} = \frac{1}{Z_1} \prod_{t=1}^T p_z^{z_t} (1 - p_z)^{e_t - z_t} P_d^{d_t} (1 - P_d)^{u_t} \lambda_b^{a_t} \lambda_f^{f_t} \\ \times \prod_{\tau \in \omega \setminus \{\tau_0\}} \prod_{i=2}^{|\tau|} \mathcal{N}(\tau(t_i); \hat{y}_{t_i}(\tau), Q_{t_i}(\tau))$$

Here d_t, u_t, z_t, e_t, a_t are the numbers of detected, undetected, terminated, existing, and new targets at time t , and f_t is the number of false alarms. Further, \hat{y} is the Kalman prediction of the next measurement position along a track, and Q its associated covariance.

Complexity of Data Association



- Our data association problem may be viewed as a multidimensional assignment problem, which is
 - Solvable for $T = 2$ scans in $\mathcal{O}(n^3)$ (Hungarian, Auction).
 - NP-hard for $T \geq 3$ scans.
- Therefore, for $T \geq 3$ we cannot find an exact solution in reasonable time.
- \implies **Heuristic/approximate solutions are needed!**

Previous Approaches

- Multi Hypotheses Tracker (Reid, 1979)
- Joint Probabilistic Data Association (Fortmann, Bar-Shalom, Scheffé, 1980)
- Lagrangian Relaxation: (pseudo)-polynomial approximation (Deb, Yeddanapudi, Pattipati, Bar-Shalom, 1997)
- Greedy Tracker (Oh, Russel and Sastry, 2004)
- Markov Chain Monte Carlo (Oh, Russel, Sastry, 2004)
- Semigreedy Track Selection Algorithm (Capponi, 2004)
- Cross Entropy Data Association (Current Work)

The Cross Entropy Optimization Method

The basic CE Algorithm for the combinatorial optimization problem $\max_{\mathbf{x} \in X} S(\mathbf{x})$ may be summarized as follows:

- 1 Initialize - Choose some \mathbf{v}_0 . Set $t = 1$.
- 2 Sample $\mathbf{X}_1, \dots, \mathbf{X}_N \sim f(\cdot; \mathbf{v}_{t-1})$
- 3 Retain the upper ρ -percentile (the *the elite sample*) $\tilde{\mathbf{X}}_1, \dots, \tilde{\mathbf{X}}_{\rho N}$ with respect to the cost $S(\mathbf{x})$.
- 4 Compute $\tilde{\mathbf{v}}_t$ as the MLE based on the elite samples.
If $f(\cdot; \mathbf{v})$ is a NEF with mean \mathbf{v} , this obtains

$$\hat{\mathbf{v}}_t = \frac{1}{\rho N} \sum_{i=1}^{\rho N} \tilde{\mathbf{X}}_i = \frac{\sum_{i=1}^N \mathbf{X}_i \mathbb{1}_{\{S(\mathbf{x}_i) \geq \gamma_t\}}}{\sum_{i=1}^N \mathbb{1}_{\{S(\mathbf{x}_i) \geq \gamma_t\}}}$$

- 5 Relax the update: $\hat{\mathbf{v}}_t = \alpha \tilde{\mathbf{v}}_t + (1 - \alpha) \hat{\mathbf{v}}_{t-1}$, $0 \leq \alpha \leq 1$.

The Parametric MinxEnt Optimization Method

PME Algorithm for Combinatorial Optimization

- 1 Initialize - Choose some \mathbf{v}_0 . Set $t = 1$.
- 2 Sample $\mathbf{X}_1, \dots, \mathbf{X}_N \sim f(\cdot; \mathbf{v}_{t-1})$.
- 3 Compute $\lambda_t = \frac{1}{\rho N} \sum_{i=1}^{\rho N} S(\tilde{\mathbf{X}}_i)$, the mean performance over the elite ρ -percentile.
- 4 Solve $\frac{\sum_{i=1}^N S(\mathbf{X}_i) \exp\{-S(\mathbf{X}_i)\lambda\}}{\sum_{i=1}^N \exp\{-S(\mathbf{X}_i)\lambda\}} = \gamma_t$ to obtain λ_t
- 5 Update $\hat{\mathbf{v}}_t$ componentwise as the tilted average

$$\hat{\mathbf{v}}_t = \frac{\sum_{i=1}^N \mathbf{X}_i \exp\{-S(\mathbf{X}_i)\lambda_t\}}{\sum_{i=1}^N \exp\{-S(\mathbf{X}_i)\lambda_t\}}$$

- 6 Relax the update $\hat{\mathbf{v}}_t = \alpha \tilde{\mathbf{v}}_t + (1 - \alpha)\hat{\mathbf{v}}_{t-1}$, $0 \leq \alpha \leq 1$.
- 7 Increment t and repeat unless stopping criteria are met.

Implementing the CE and MinxEnt Methods

To implement these optimization schemes, we will need to

- Define the parametric family $f(\cdot, \mathbf{v})$ over the set of feasible solutions, namely the feasible partitions $\omega \in \Omega$.
- Devise an efficient scheme for sampling from these distributions.


It will be convenient to use graph notations for that purpose.

Basic Connectivity Graph

Given the measurements $Y_{1:T}$, we construct the basic **connectivity graph** $G = (V, E)$, in which:


- The nodes are the measurements $\{y_t^j\}$.
- Node i is connected to node j if it can be an immediate predecessor to v in some legal path (subject to velocity and time separation limitations).
- Evidently, $G = (V, E)$ is a directed acyclic graph.

Sampling Paths



- We now need to sample a random number K of tracks (i.e., directed paths) over the connectivity graph, which can start and terminate in arbitrary positions.
- Note that these paths must be vertex-disjoint.
- For that purpose we augment the graph with appropriate start and sink nodes.

The Augmented Connectivity Graph



The parameter vector now is: $\{p_b(i), p_f(i), p_{ij}\}$, where $p_b(i)$ is the probability that a path starts in node i , $p_f(i)$ is the termination probability at node i , and p_{ij} are the transition probabilities, re-normalized so that $\sum_j p_{ij} + p_f(i) = 1$.

Sampling Procedure

- To avoid excessive sample rejection due to path intersection, we use a standard **edge elimination** procedure.
- Thus, a node that is selected and all edges leading to it are removed from the graph, and the transition probabilities are re-normalized to sum to 1.

Parameter Update

The update rule for all transition probabilities is

$$\hat{p}_{ij} = \frac{\sum_{k=1}^N \mathbb{1}_{\{S(\omega_k) \geq \gamma\}} \mathbb{1}_{\{\omega_k \in \Omega_{ij}\}}}{\sum_{k=1}^N \mathbb{1}_{\{S(\omega_k) \geq \gamma\}}}$$

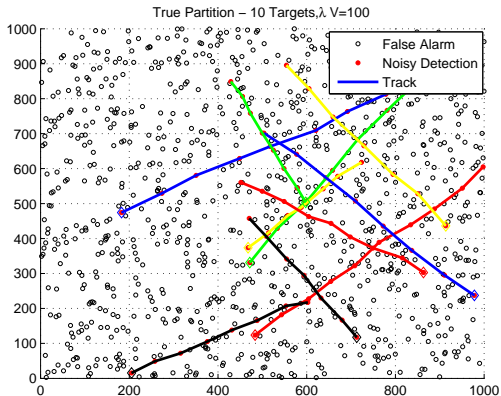
Simply put: \hat{p}_{ij} is set to its empirical frequency n_{ij}/n_i in the elite sample.

Additional Enhancements

The following enhancements were incorporated into the final algorithm, and greatly affected its performance:

- **Multi-node states:** In the basic algorithm, the transition probabilities depend only on the last node, namely on position information only. To allow encoding direction and velocity information, we considered **node pairs** as the basic state of the sampling graph. This increases the number of parameters in proportion to the average graph connectivity, but greatly enhances performance.
- **Two-directional sampling:** This greatly increases the path integrity and avoids sampling paths that start in the middle of an actual track. Again, the number of parameters approximately doubles.
- **Smart initialization:** The initial parameters \mathbf{v}_0 are selected using a local (3-point) Kalman filter based scheme.

Simulation Setup



$T = 10$ scans

Performance Measures



In addition to the posterior probability that serves as the cost function to be minimized, the following performance measures are natural and intuitively appealing for the problem at hand:

- **Normalized Correct Associations (NCA)** – the number of correctly recovered associations normalized by the total number of associations in the test scenario.
- **Incorrect-to-Correct Association Ratio (ICAR)** – the ratio between the number of incorrectly recovered association and the number of correctly recovered ones.

Simulation Results – Dense Scenarios

Setup



$$K = 10 \div 75, P_d = 0.999, \lambda_f V = 1.$$



Simulation Results – Cluttered Scenarios

Setup



$K = 10$, $P_d = 0.999$, $\lambda_f V = 1 \div 100$.



Simulation Results – Poor Detection

Setup

$K = 10$, $P_d = 0.3 \div 0.9$, $\lambda_f V = 1$.



Conclusion

Summary

- We have proposed two polynomial algorithms for the multi-scan data association problem.
- The algorithms are robust against dense scenarios, heavy clutter and low detection probabilities.
- The algorithms are efficient in the sense that each iteration is polynomial in the problem size. Empirically the number of iterations to convergence never exceeded 10.
- The performance obtained represent the state of the art relative to algorithms known to us.

We note in closing that these algorithms may be extended in many directions of interest within target tracking – including “online” tracking problems, multi-sensor tracking, and large single-scan associations.