Label-Free Estimation and Tracking

by Peter Willett
UConn ECE Department

with lots of help from (and most of the ideas and work by)

Lennart & Daniel Svensson, Marco Guerriero, Marcus Baum and David Crouse

June 24^{th}, 2014

What Is Measurement Origin Uncertainty (MOU)? Why Do Engineers Care?

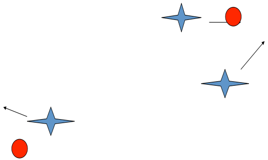

- Which measurements come from which targets?
- How many targets are there?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD

What Is Measurement Origin Uncertainty (MOU)? Why Do Engineers Care?

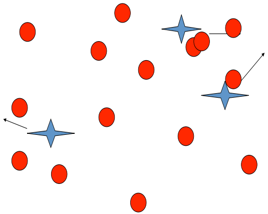

- Which measurements come from which targets?
- How many targets are there?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD

What Is Measurement Origin Uncertainty (MOU)? Why Do Engineers Care?

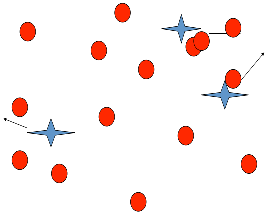

- Which measurements come from which targets?
- How many targets are there?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD

What Is Measurement Origin Uncertainty (MOU)? Why Do Engineers Care?

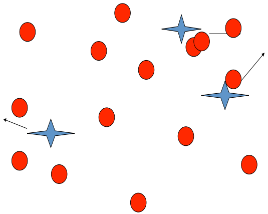

- Which measurements come from which targets?
- How many targets are there?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD

What Is Measurement Origin Uncertainty (MOU)? Why Do Engineers Care?

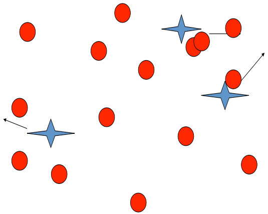

- Which measurements come from which targets?
- How many targets are there?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD

What Is Measurement Origin Uncertainty (MOU)? Why Do Engineers Care?

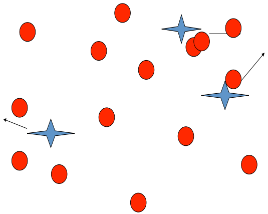

- Which measurements come from which targets?
- How many targets are there?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD

What Is Measurement Origin Uncertainty (MOU)? Why Do Engineers Care?

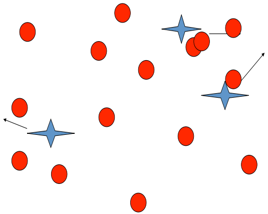

- Which measurements come from which targets?
- How many targets are there?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD

What Is Measurement Origin Uncertainty (MOU)? Why Do Engineers Care?

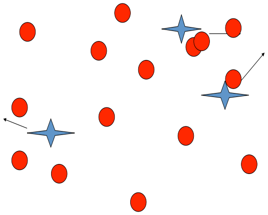

- Which measurements come from which targets?
- How many targets are there?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD

What Is Measurement Origin Uncertainty (MOU)? Why Do Engineers Care?

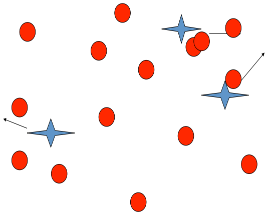

- Which measurements come from which targets?
- How many targets are there?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD

What Is Measurement Origin Uncertainty (MOU)? Why Do Engineers Care?

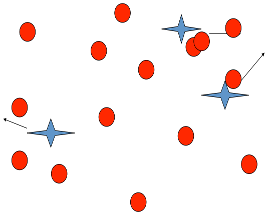

- Which measurements come from which targets?
- How many targets are there?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD

The Probabilistic Data Association Filter

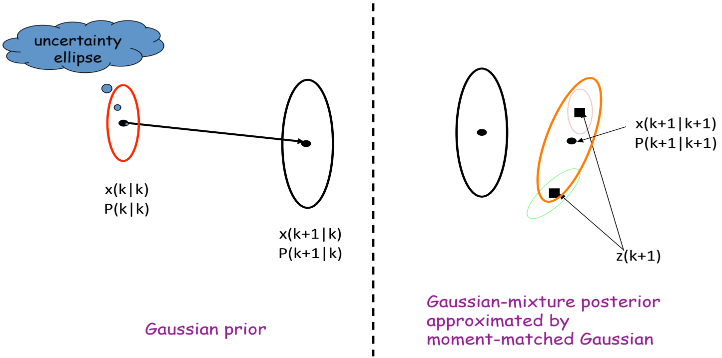

The PDAF is perhaps the only consistent tracker.
Joint (J)PDAF: for multiple targets, weights are computed from joint association events, and tracks are updated individually.
Here the key here is that the surviving estimates are the MMSEEs.

Lennart Svensson: Peter, how about a **Set** JPDAF?

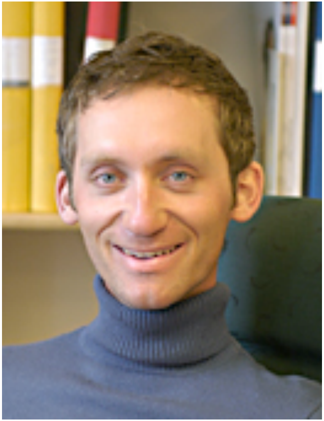

Lennart Svensson: Peter, how about a **Set** JPDAF?

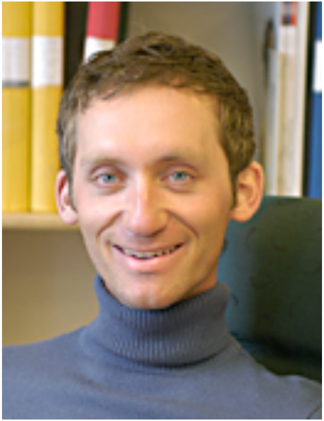

The usual intelligent reaction from me.

Marco Guerriero: Professor, you are wrong!

Marco Guerriero: Professor, you are wrong!

Most folks seemed to agree that I was wrong.

Outline

- Label-Free Estimation
  - Motivation
  - OSPA, RFS & MMOSPA
- Tracking Closely-Spaced Objects
  - Soft Association: Track Merging
  - Hard Association: Track Repulsion
- Label-Free Tracking
  - SJPDAF
- Labeled label-free tracking

Rockets & torpedoes

There are some circumstances in which multi-object estimation ought to be “label-free”. One wants to track all objects, but one is not especially concerned which one is which.

MMSE Estimation

Assuming that MMSE estimation is the goal

$$\{\hat{x}_i\}_{i=1}^n = \arg\min \left\{ \int \sum_{i=1}^n |x_i - \hat{x}_i|^2 \, p(x_1, \ldots, x_n) dx_1 \ldots dx_n \right\}$$

This will be a shorthand notation: in fact, for estimation, we would really have

$$\begin{aligned} &\{\hat{x}_i(z)\}_{i=1}^n \\ &= \arg\min \left\{ \int \sum_{i=1}^n |x_i - \hat{x}_i(z)|^2 \, p(x_1, \ldots, x_n | z) dx_1 \ldots dx_n \right\} \end{aligned}$$

in which z is an observation – that is, we are doing a-posteriori estimation.

Label-Constrained MMSE Estimation[1]

One does not usually write this, but when one estimates in the multi-object situation, one usually has the "labeling" constraints

$$x_1 \sim p_1, \quad x_2 \sim p_2, \quad \ldots \quad x_n \sim p_n$$

as opposed to, say

$$x_3 \sim p_1, \quad x_n \sim p_2, \quad \ldots \quad x_1 \sim p_n$$

Recalling that this is a minimization problem, the constraints can be doing no good as regards our minimization criterion (MSE). Are they doing any harm?

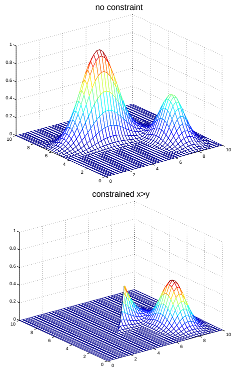

[1] Redner and Walker, "Mixture densities, maximum likelihood and the em algorithm," *SIAM Review*, 1984.

No Labeling? OK, so we need Random Finite Sets

RFS[2]: When permutation-symmetry holds (i.e., objects are exchangeable), Janossy densities[3] have the following interpretation:

$$j^n_{k|k}(x_1, x_2, \ldots, x_n)$$

$\triangleq$ Probability that there are exactly n unlabeled objects one in each of the n distinct infinitesimal regions $(x_i, x_i + dx_i)$, divided by $dx_1 dx_2 \ldots dx_n$

The Probability Hypothesis Density[4] (PHD) is

$$\text{PHD}_{k|k}(x) = \text{P}\{\text{There exists a target in } (x, x+dx)\} \cdot \frac{1}{dx}$$

$$\triangleq \sum_{n=0}^{\infty} \frac{1}{n!} \int j^{n+1}_{k|k}(x, y_1, \ldots, y_n) dy_1 dy_2 \ldots dy_n$$

---

[2] R. Mahler, "Statistics 101 for Multisensor, Multitarget Data Fusion," *IEEE AES Magazine*, 2004.

[3] Daley and Vere-Jones, *An Introduction to the Theory of Point Processes*, 2002.

[4] R. Mahler, *Statistical Multisource Multitarget Information Fusion*, Artech House, 2007.

More genuflection toward RFS

Janossy densities have been used for multi-target tracking[5]. The multi-target set posterior

$$
\begin{aligned}
&\xi_{k|k}(\{x_1, x_2, \ldots, x_n\}) \\
&\quad = \quad j^n_{k|k}(x_1, x_2, \ldots, x_n) \\
&\quad = \quad n!\, p_{k|k}(n)\, f^n_{k|k}(x_1, x_2, \ldots, x_n) \\
&\quad = \quad p_{k|k}(n)\, f^{jmp,n}_{k|k}(x_1, x_2, \ldots, x_n)
\end{aligned}
$$

where $f^n_{k|k}(x_1, x_2, \ldots, x_n)$ is the posterior of the n labeled objects, $p_{k|k}(n)$ is the posterior that there *are* n objects and $f^{jmp,n}_{k|k}(x_1, x_2, \ldots, x_n)$ is the joint multitarget posterior[6].

---

[5] S. Mori, C. Chong, "Tracking of Groups of Targets Using Generalized Janossy Measure Density Function", *Proc. International Conference on Information Fusion*, 2006.

[6] K. Kastella, "Joint Multitarget Probabilities for Detection and Tracking", *Proc. International Conference on Information Fusion*, 2005.

Back to MMSE: What's wrong with it?

The problem, of course, is that the MMSE criterion

$$\{\hat{x}_i\}_{i=1}^{n} \;=\; \arg\min \left\{ \int \sum_{i=1}^{n} |x_i - \hat{x}_i|^2 \, p(x_1, \ldots, x_n) dx_1 \ldots dx_n \right\}$$

has little meaning *without* labeling. What can one do?

Multi-Object Estimation Criteria

Pretty good estimation?

Multi-Object Estimation Criteria

Multi-Object Estimation Criteria

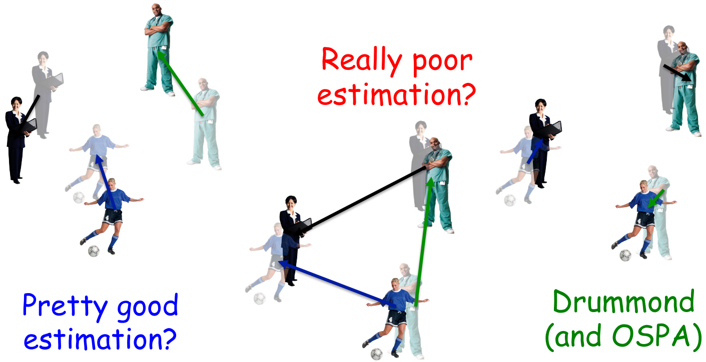

Drummond[7] discussed measures that do not overly penalize labeling mistakes by calculating the SE of the best assignment.

[7] Drummond & Fridling, “Ambiguities in evaluating performance of multiple target tracking algorithms,” *SPIE Conference on Signal and Data Processing of Small Targets*, 1992.

The OSPA Metric

The Optimal Mass Transfer (OMAT[8]) *metric* is less ad-hoc than assignment, but there are weaknesses. Here we'll refer to Optimal Subpattern Assignment (OSPA[9]) that is more general and encompassing.

$$\tilde{d}_p^{(c)}(\widehat{\mathbf{X}}, \mathbf{X}) = \left(\frac{1}{n} \left(\min_{\pi \in \Pi_n} \sum_{i=1}^{n} d^{(c)}(\mathbf{x}^{(i)}, \hat{\mathbf{x}}^{\pi(i)})^p \right) \right)^{1/p}$$

Here, $d^{(c)}(\mathbf{x}, \hat{\mathbf{x}}) \triangleq \min(c, d(\mathbf{x}, \hat{\mathbf{x}}))$ is the distance d between $\mathbf{x}$ and $\hat{\mathbf{x}}$, cut-off at c, and $\hat{\mathbf{x}}^{\pi(i)}$ describes the i^{th} permutation (target-reordering) of the vector $\hat{\mathbf{x}}$.
For now, we will make $c = \infty$, $p = 2$ and assume that object cardinality is known.

[8] Hoffman & Mahler, "Multitarget Miss Distance via Optimal Assignment," *IEEE T-SMC:A*, 2000.

[9] Schuhmacher, Vo & Vo, "A Consistent Metric for Evaluation of Multi-Object Filters," *IEEE T-SP*, 2008.

MMOSPA

- Since one can calculate the squared error (SE), and then minimize its mean value (MSE) to derive the MMSE, in obvious parallel here we are interested in minimizing the *mean* OSPA (MOSPA) measure (the MMOSPA estimate):

$$\text{MOSPA}_p^{(c)}(\widehat{\mathbf{X}}, p(\mathbf{X})) \triangleq \mathbb{E}_{p(\mathbf{X})}\{\tilde{d}_p^{(c)}\}.$$

- Nice theory about the combination of MMOSPA and RFS[10] and nice algorithms to find the approximate[11] and exact[12] MMOSPA estimates too.

---

[10] Guerriero, Svensson, Svensson and Willett, "Shooting Two Birds with Two Bullets: How to Find Minimum Mean OSPA," FUSION 2010.

[11] Crouse, Willett, Guerriero and Svensson, "An Approximate Miminum MOSPA Estimator," ICASSP 2011.

[12] Baum, Willett and Hanebeck, "Calculating Some Exact MMOSPA Estimates for Particle Densities," FUSION 2012.

Since this isn't going to be easy . . .

Applications where MOSPA is reasonable

- two threats are attacking
- radar cueing: there is no interest in which target is which, but where there are targets is very important
- collision avoidance: it is not of interest to know which car is which, the goal is to avoid all cars
- maritime surveillance: send helicopters as close as possible to a set of vessels with anomalous behavior

But it will turn out that MMOSPA has application in a wider context, even in situations where target labeling is very much desired.

Simple Gaussian Example, 2 Objects in 1-Space

□ MMOSPA X MMSE   + ML

MMOSPA is Difficult Even with Known Cardinality

- Except in the scalar case, MMOSPA estimation requires optimization over a difficult integral.

$$\hat{\mathbf{x}}^{\text{MMOSPA}} = \arg\min_{\hat{\mathbf{x}}} \frac{1}{N_T} \int \min_{\mathbf{a}} \|\mathbf{x}_{\mathbf{a}} - \hat{\mathbf{x}}\|^n p(\mathbf{x})\, d\mathbf{x}$$

(2010) Guerriero, Svensson, Svensson, Willett: Shooting Two Birds with Two Bullets: How to Find Minimum Mean OSPA Estimates

Proved that the MMOSPA estimate is the mean of a specific RFS density.

Which Specific RFS density?

Theorem

For any given $\hat{\mathrm{x}}$, there exists exactly one ordered density $\tilde{p}(\mathbf{x})$ within the RFS family, such that

$$\mathit{MSE}(\hat{\mathrm{x}}, \tilde{p}(\mathbf{x})) = \mathit{MOSPA}(\hat{\mathrm{x}})$$

It turns out that

$$\tilde{p}(\mathbf{x}) = \begin{cases} f(\{\mathbf{x}_1, \ldots, \mathbf{x}_n\}) & \text{if } \mathbf{x} \in \mathcal{A}(\hat{\mathrm{x}}) \\ 0 & \text{otherwise} \end{cases}$$

where

$$\mathcal{A}(\hat{\mathrm{x}}) = \left\{ \mathbf{x} : \frac{1}{n} \sum_{i=1}^{n} d(\hat{\mathrm{x}}_i, \mathbf{x}_i) = \tilde{d}_p(\hat{\mathrm{x}}, \mathbf{x}) \right\}$$

and, unfortunately, $\hat{x}$ is the MMSEE of $\tilde{p}$ – the MMOSPA estimate for p. An iterative solution does seem to work, though.

Cartoon of Iterated Algorithm in Scalar Case

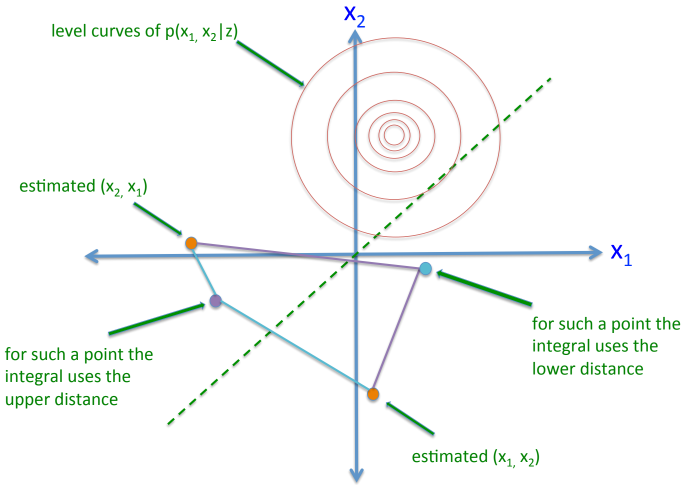

. . . and after it is done

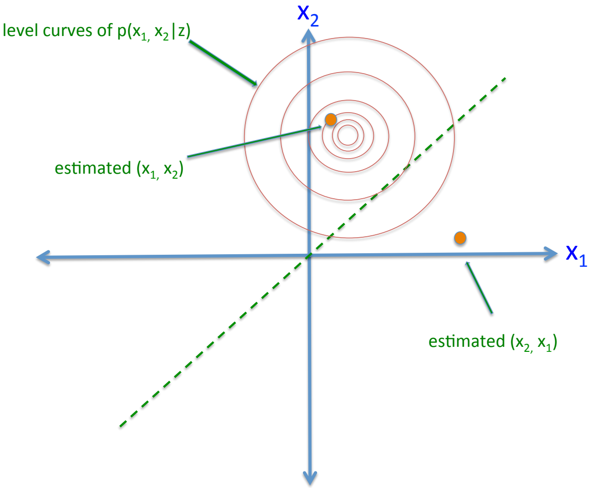

Iterative Scheme for 2 Objects in 2-Space, Particle pdf

Focus on 10 Particles to Watch Them Switch

Iterative Scheme for 3 Objects in 2-Space, Particle pdf

Focus on 10 Particles to Watch Them Switch

Exact (Non-Iterative) Algorithms for Particle pdfs

$$p(x) = \sum_{i=1}^{N_p} w_i \delta(x = x^{(i)})$$

Number of targets n

Dimension d

	2	3	4	$\dots$
1		$\mathcal{O}(N_p\, n \log n)$ Section VI		
2	$\mathcal{O}(N_p \log N_p)$	$\mathcal{O}\left((N_p (n!)^2)^{nd-1}\right)$ Section IV		
3				
$\vdots$	$\mathcal{O}((N_p)^{d-1})$ Section V			

Scalar Case

In 1d: Optimal permutation obvious!

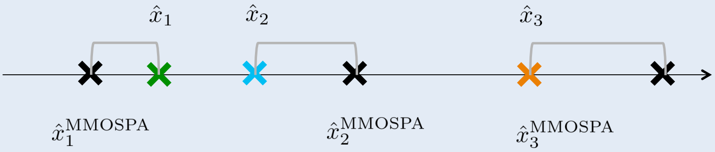

MMOSPA and Order Statistics

MMOSPA estimate = mean of the order statistics

$$x^{\mathrm{MMOSPA}} = \int \mathrm{sort}(x)\ p(x)\ dx$$

where $\mathrm{sort}(x)$ sorts the single target states

Scalar Case: Particles

Basic Algorithm

MMOSPA estimate for $p(x) = \sum_{i=1}^{N} w_i \cdot \delta(x - x^{(i)})$

- Sort the single target states in each particle

$$\hat{x}^{\text{MMOSPA}} := \sum_{i=1}^{N} w_i \, \text{sort}(x^{(i)})$$

Runtime Complexity

Linear in the number of particles, i.e., $\mathcal{O}(N \; n \log n)$

Two-Target Case

$$\hat{x}^{\text{MMOSPA}} = \text{argmin}_{\hat{x}} \frac{1}{2} \sum_{i=1}^{N} w_i \cdot \min_{\pi_i} ||\hat{x} - P_{\pi_i}(x^{(i)})||^2$$

$$= \frac{1}{2} \sum_{i=1}^{N} w_i \cdot P_{\pi_i^{\text{opt}}}(x^{(i)})$$

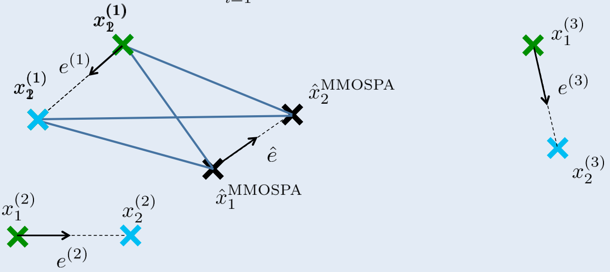

$$\pi_i^{\text{opt}} = (1, 2) \quad \Leftrightarrow \quad \text{sign}(\langle \hat{e}, e^{(i)} \rangle) = +1$$

Two-Target Case: the Regions that Need Checking

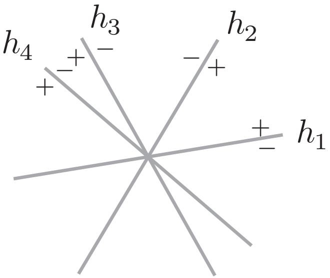

It's easy in 2-space.

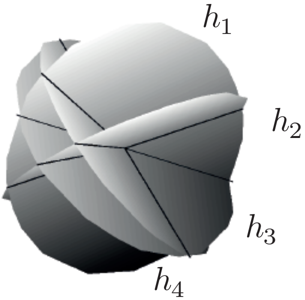

It's a little harder in 3-space.

Example of Exact Calculation

The MMSEE Suffers at the Hands of MOU

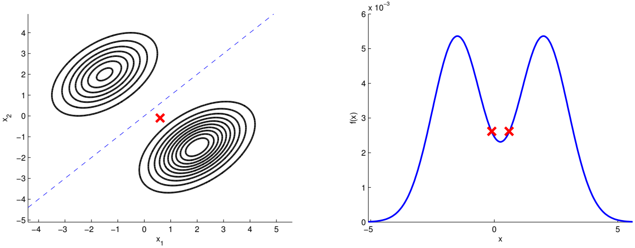

✗ marks the mean.

- Left: full PDF. Right: marginal RFS Density.
- The mean is located between the estimates, a general area of low likelihood.

The MMOSPA Estimate Does Not

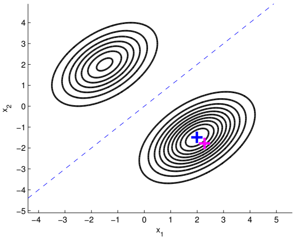

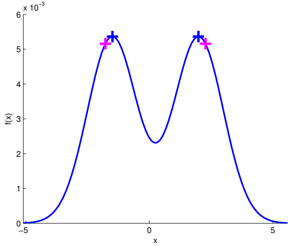

\+ marks the MMOSPA estimate[13] with $n = 2$.

\+ marks the MMOSPA estimate with $n = 4$.

- n is the exponent in the OSPA metric.
- Left: full PDF. Right: marginal RFS Density.

[13] Guess why this MMOSPA estimate is offset *away* from the middle?

The Chernov Bound Approximation

For two targets, the Chernov Bound tells us that

$$\frac{1}{N_T}\int_{\mathbf{x}\in\mathbb{R}^{2d}} \min\left[\|\mathbf{x}-\hat{\mathbf{x}}\|^n, \|\mathbf{x}-\chi\hat{\mathbf{x}}\|^n\right] p(\mathbf{x})\, d\mathbf{x}$$

$$\leq \frac{1}{N_T}\int_{\mathbf{x}\in\mathbb{R}^{2d}} \|\mathbf{x}-\hat{\mathbf{x}}\|^{n\beta} \|\mathbf{x}-\chi\hat{\mathbf{x}}\|^{n(1-\beta)} p(\mathbf{x})\, d\mathbf{x}$$

Normally, $\beta \in (0,1)$ is chosen to minimize the bound.

An Instantiation of the Bound

For $n = 4$ and $\beta = \frac{1}{2}$ – that is, Bhattacharyya – we have

$$\arg\min_{\hat{\mathbf{x}}} \frac{1}{N_T} \int_{\mathbf{x} \in \mathbb{R}^{2d}} \|\mathbf{x} - \hat{\mathbf{x}}\|^2 \|\mathbf{x} - \chi\hat{\mathbf{x}}\|^2 p(\mathbf{x})\, d\mathbf{x}$$

$$= \arg\min_{\hat{\mathbf{x}}} \mathrm{E}\Big[\left(\mathbf{x}^T\mathbf{x}\right)^2 - 2\left(\mathbf{x}^T\mathbf{x}\right)\mathbf{x}^T A\hat{\mathbf{x}} + 2(\mathbf{x}^T\mathbf{x})\left(\hat{\mathbf{x}}^T\hat{\mathbf{x}}\right)$$

$$-2(\hat{\mathbf{x}}^T\hat{\mathbf{x}})\left(\mathbf{x}^T A\hat{\mathbf{x}}\right) + 4\left(\mathbf{x}^T\hat{\mathbf{x}}\right)\left(\mathbf{x}^T\chi\hat{\mathbf{x}}\right) + (\hat{\mathbf{x}}^T\hat{\mathbf{x}})^2 \Big]$$

- $A \triangleq \chi + \mathbf{I}$
- χ is the skew-identity matrix
- We can minimize the bound for any PDF with explicit moments.
  - We have explicit solutions for Gaussian mixtures.
  - The optimization over $\hat{\mathbf{x}}$ can be performed using Newton's method; the mean can be used as the initial estimate.

The Chernov Bound Minimization Can Work Well

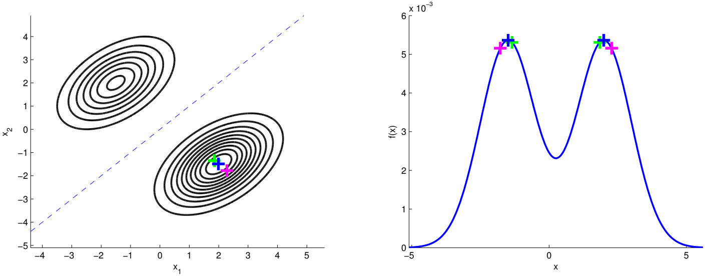

+ marks the estimate from minimizing the Chernov bound.

+ marks the MMOSPA estimate with exponent $n = 2$.

+ marks the MMOSPA estimate with exponent $n = 4$.

Left: full PDF. Right: marginal RFS PDF.

Example

Two Gaussian-prior targets, ballistic trajectory, separated by unit distance. We assume $P_d = 1$ – two unlabeled hits – and Gaussian observation noise at each scan:

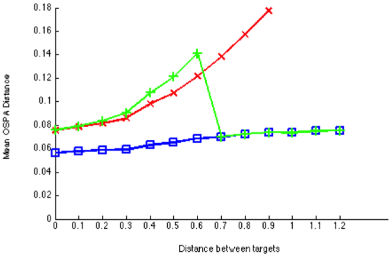

Big deal, Willett. You’re the fastest horse in your own race.

Does President Obama care about MOSPA?

Probably not . . . yet.

So let us define the kill probability as the probability that the estimates of the target locations are within a ball of radius ϵ centered around the true target locations

$$P_{kill}(\hat{\mathbf{x}}, \epsilon) = P(\|\hat{\mathbf{x}} - \mathbf{x}\| \leq \epsilon)$$

Another Example: CS versus MAP versus MMOSPA for Multiple AOAs

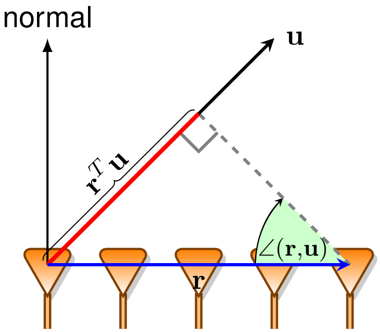

The element-level observations $\mathbf{z}$ with complex Gaussian amplitudes $\mathbf{b}$ and noise. This is the usual conversion of a direction finding problem to one of frequency estimation.

$$\mathbf{z} = \mathbf{A}\mathbf{b} + \mathbf{w}$$

The PDFs

The prior distribution on the complex amplitudes will be complex Gaussian (Swerling 1, 2). An uninformative (uniform) prior will be used on the directions of arrival.

$$
\begin{aligned}
p(\mathbf{U}, \mathbf{b}|\mathbf{z}) &= \frac{p(\mathbf{z}|\mathbf{U}, \mathbf{b})p(\mathbf{U}, \mathbf{b})}{p(\mathbf{z})} \\
p(\mathbf{z}|\mathbf{U}, \mathbf{b}) &= \pi^{-N_{\text{el}}}|\mathbf{Q}|^{-1}e^{-(\mathbf{z}-\mathbf{A}\mathbf{b})^H\mathbf{Q}^{-1}(\mathbf{z}-\mathbf{A}\mathbf{b})} \\
p(\mathbf{b}) &= \pi^{-M}|\mathbf{Q}_b|^{-1}e^{-\mathbf{b}^H\mathbf{Q}_b^{-1}\mathbf{b}} \\
p(\mathbf{U}) &= \prod_{m=1}^{M} p(\mathbf{u}_m)
\end{aligned}
$$

Solutions

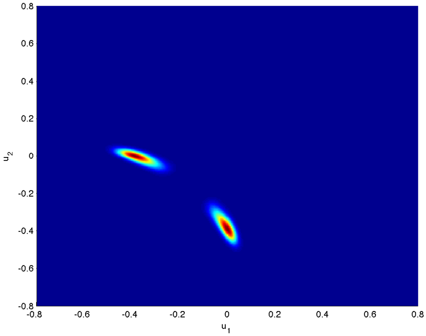

- The MMSE solution is between the two peaks.
- Compressed sensing.
- MAP (or ML).
- MMOSPA?

A Simulation

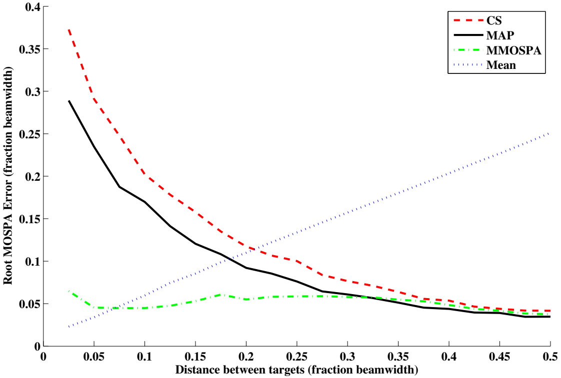

5-element linear array. One target was fixed at the boresight (0°) and the other was moved across the array. SNR $\approx 20\,$dB.

Same Thing in Two Dimensions

- 9 antennas arranged in a 3×3 grid
- Two targets
- Posterior density for DOAs: Particle approximation using importance sampling (2000 samples)
- Exact algorithm for MMOSPA estimate[3]

Radiation pattern and trajectory of target 2.

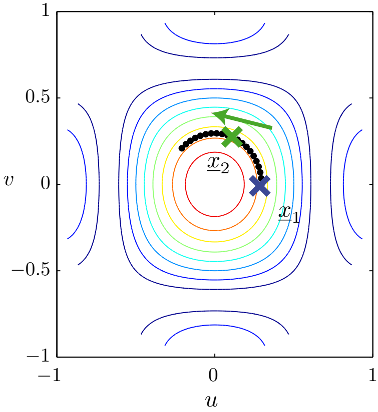

The Planar Array Results

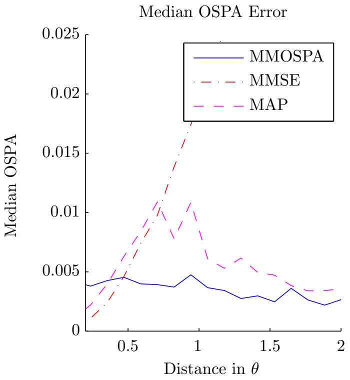

SNR 13dB.

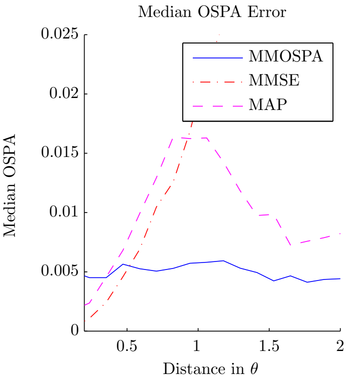

SNR 10dB.

The Problem with Closely-Spaced Objects

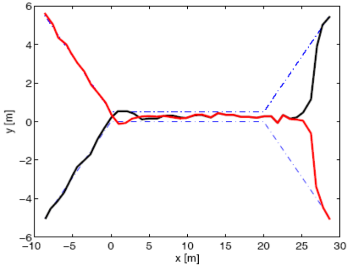

Top: two targets being tracked by the JPDAF – and believe it or not the *coupled* JPDAF is worse!

Close tracks tend to merge with soft-association trackers and to repel in hard-association ones.

Soft Assignment Trackers (PDA, JPDA) Merge CSOs

- MMSE target state update
- It is fairly well-known[14] that soft association trackers tend to "coalesce" (merge) nearby targets.

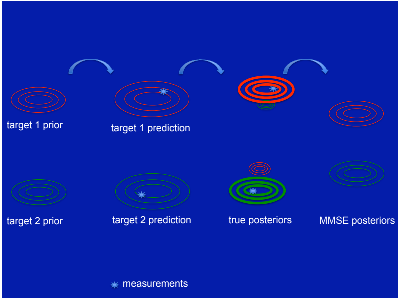

Two targets' pdf contours are notionally represented in red & green. For the "true posterior," the thickness of the lines indicates the posterior weight in the Gaussian-mixture mode. The MMSE estimates are the soft-association trackers' updated pdfs – note that due to the association uncertainty they have tended toward one another.

[14] R. Fitzgerald, "Track Biases and Coalescence with Probabilistic Data Association," IEEE T-AES, 1985.

Hard-Association Trackers (GNN, MHT) Repel CSOs!

- Choose only the association that fits best
- It is less well-known[15] that with hard association trackers nearby targets have difficulty crossing.
- Below: right happens more often than left.

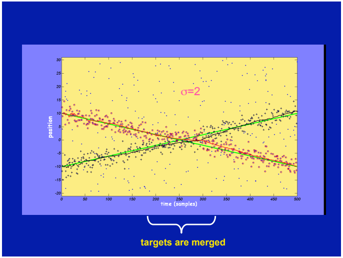

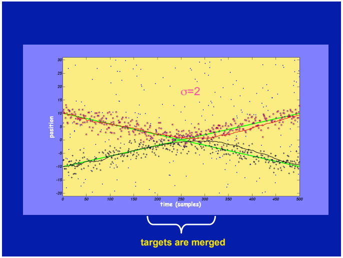

15 S. Coraluppi, C. Carthel, P. Willett, M. Dingboe, O. O'Neill, and T. Luginbuhl, "The Track Repulsion Effect in Automatic Tracking," *FUSION 2009*.

Intuition Behind Track Repulsion with Hard Association Trackers

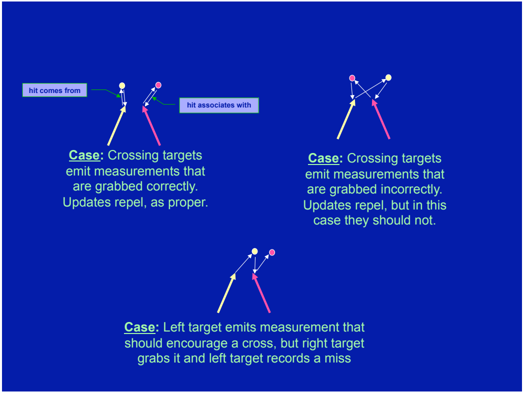

Although with high probability the association is correct, there is an *unbalanced* probability that the update is in the opposite direction from the other target.

The SJPDA

- With a multimodal pdf the MMSEE “hedges its bets” and produces an estimate somewhere among the modes
  - This is track coalescence.
- MMOSPA estimation eliminates some of the extra pdf modes, the ones that correspond to re-orderings of the items being estimated
- Since its genesis was in random finite set (RFS) theory, this was given the monicker *Set Joint Probabilistic Data Association*, or SJPDA

The SJPDA and MMOSPA

In the SJPDA context (Gaussian priors, Gaussian measurement noise, measurement-origin uncertainty) the MOSPA error is proportional to the trace of an *unordered* joint covariance matrix,

$$\begin{aligned}
\mathbf{P}(k) &\equiv \mathrm{E}\left[\min_{\mathbf{a}} \left(\mathbf{x}_{\mathbf{a}}(k) - \hat{\mathbf{x}}(k|k)\right)\left(\mathbf{x}_{\mathbf{a}}(k) - \hat{\mathbf{x}}(k|k)\right)^T \middle| Z^k\right] \\
&= \sum_{i=1}^{N_H} \Pr\{\theta_i\}\, \mathrm{E}\left[\min_{\mathbf{a}} (\mathbf{x}_{\mathbf{a}}(k) - \hat{\mathbf{x}}(k|k))(\mathbf{x}_{\mathbf{a}}(k) - \hat{\mathbf{x}}(k|k))^T \middle| \theta_i, Z^k\right] \\
&\approx \sum_{i=1}^{N_H} \Pr\{\theta_i\} \min_{\mathbf{a}_i} \mathrm{E}\left[\left(\mathbf{x}_{\mathbf{a}_i}(k) - \hat{\mathbf{x}}(k|k)\right)\left(\mathbf{x}_{\mathbf{a}_i}(k) - \hat{\mathbf{x}}(k|k)\right)^T \middle| \theta_i, Z^k\right]
\end{aligned}$$

The SJPDAF uses the a_i's chosen to minimize this (approximation's) trace, and then updates the covariance in

$$\sum_{i=1}^{N_H} \Pr\{\theta_i\} \left(\mathbf{P}_{\mathbf{a}_i}(k|k) + (\hat{\mathbf{x}}_{i,\mathbf{a}_i}(k|k) - \hat{\mathbf{x}}(k|k))(\hat{\mathbf{x}}_{i,\mathbf{a}_i}(k|k) - \hat{\mathbf{x}}(k|k))^T\right)$$

The SJPDA Intuition

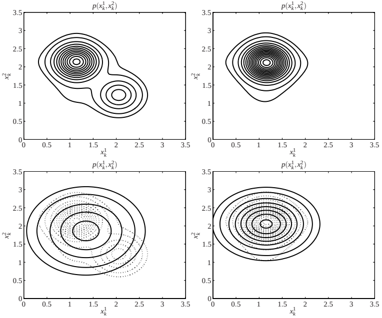

setup Gaussian prior, 2 targets, $P_d = 0.7$, 2 measurements

top left true posterior

top right MMOSPA posterior

bottom left JPDA posterior (the MMSE estimate)

bottom right SJPDA posterior (the approximate MMOSPA estimate)

The SJPDA Solution

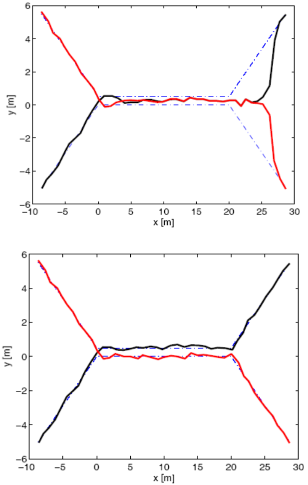

Top: two targets being tracked by the JPDAF – and believe it or not the *coupled* JPDAF is worse!

Below: the label-free SJPDAF.

But Prime Minister Rajoy *Wants* Labels . . .

An operator might want to:

- Know how many targets there are.
- Know the states of the targets.
- Know what the targets are. This could mean knowing:
  - Their identity relative to a specific time designated by the operator.
  - Their physical nature according to classification (feature) information.
- Have a probabilistic model for the uncertainties.
  - Localization accuracy (covariance)
  - Identity accuracy

The SJPDA With Identity Information

In a particle filter, each particle might represent a probability and a stacked set of states for each target.

$$p_i = \{w_i, \mathbf{x}_i\}$$

For this exposition, consider that a "particle" can be the same as a Gaussian mode in the SJPDAF.

Suppose that we want to be able to change the ordering of the states for the targets without losing or corrupting any information. We can do this by adding an extra component:

$$\tilde{p}_i = \{w_i, \mathbf{x}_i, \mathbf{o}_i\}$$

The vector $\mathbf{o}$ can be an $N_T! \times 1$ vector that lists the probability of a particular ordering of states, initially all zeros and a single 1.

- Suppose that the ordering is known exactly (i.e. $\mathbf{o}_i$ contains a single one and all zeros.
- Then we can represent $\mathbf{o}_i$ as a $N_T \times N_T$ permutation matrix, χ_i (containing exactly one 1 in each row and column).
- If we mix the permutation matrices, the result can be interpreted as a matrix of marginal probabilities, e.g.,

$$\chi_M = \begin{bmatrix} & \text{State 1} & \text{State 2} & \text{State 3} \\ \text{Track 1} & 2/3 & 1/6 & 1/6 \\ \text{Track 2} & 1/3 & 1/3 & 1/3 \\ \text{Track 3} & 0 & 1/2 & 1/2 \end{bmatrix}$$

so

$$\Pr\{\text{State 1 is Track 1}\} = 2/3$$

If we make an independence assumption, then we can say

$$\Pr\{\text{State 1 is Track 1} \bigcap \text{State 2 is Track 2} \bigcap \text{State 3 is Track 3}\} \approx \frac{\Pr\{\text{State 1 is Track 1}\} \Pr\{\text{State 2 is Track 2}\} \Pr\{\text{State 2 is Track 2}\}}{\text{perm}[\chi_M]}$$

This formulation makes mixing simpler than using the full $\mathbf{o}_i$.
Note that a label-free tracker with labeling is not the same as a labeled tracker. The difference is that target locations are estimated without a labeling constraint; and the labeling identity is an "overlay".
It is worth reiterating that traditional (labeled) trackers *do* have uncertain labels, but keep quiet about them and deliver estimates as if there were no uncertainty.

Simulation: Crossing & Recrossing Targets

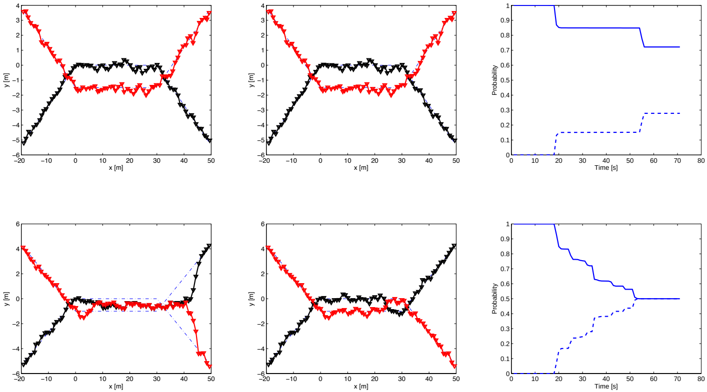

Separation cases 1.5 unit (upper) and 1.0 unit (lower).
Left: JPDAF. Middle: SJPDAF. Right: identity probabilities.

Monte Carlo Performance

Cross ?	min. sep.	SJPDA % Track loss	SJPDA P_{label}^{act}	SJPDA P_{label}^{est}	JPDA % Track loss	JPDA P_{label}^{act}
No	0.5	0	48	50	73	73
	1.0	0.1	54	54	0	56
	1.5	0	97	93	0	86
Yes	0.5	0	46	50	97	26
	1.0	0.1	60	51	3	38
	1.5	0.1	84	60	0.1	94

P_{label}^{act} is probability that both tracks have the same identity at the end as at the beginning. P_{label}^{est} is SJPDA's (averaged) estimated probability that that both tracks have the same identity at the end as at the beginning.

Unknown Object Cardinality

- Recall that

$$\tilde{d}_p^{(c)}(\widehat{\mathbf{X}}, \mathbf{X}) = \left(\frac{1}{n} \left(\min_{\pi \in \Pi_n} \sum_{i=1}^{n} d^{(c)}(\mathbf{x}^{(i)}, \hat{\mathbf{x}}^{\pi(i)})^p \right) \right)^{1/p}$$

  where $d^{(c)}(\mathbf{x}, \hat{\mathbf{x}}) \triangleq \min(c, d(\mathbf{x}, \hat{\mathbf{x}}))$ is the distance d between $\mathbf{x}$ and $\hat{\mathbf{x}}$, cut-off at c and $\hat{\mathbf{x}}^{\pi(i)}$ describes the i^{th} permutation (target-reordered) of the vector $\hat{\mathbf{x}}$.
- Results so far show that minimizing this over target number as well as (unlabeled) location works just fine.
- Reason is that cardinality errors decrease exponentially in number of observations, whereas "distance" errors decrease inversely: not knowing the number of objects is (asymptotically) a non-event.

Summary

I hope I have planted the seed of the idea that we can track better than with what we had previously thought “best”.

- Multi-object estimation without labels is *unconstrained*, and hence can be better
- MMOSPA estimation (useful and interesting in its own right)
- Label-free tracking: the SJPDA
- Labeling the label-free tracker via an *overlay*