THE UNIVERSITY *of* EDINBURGH
School of Engineering

Scalable opportunistic calibration of fusion networks

Dr Murat Uney
Visiting scholar, University of Edinburgh
Scientist, CMRE

20th February 2019

THE UNIVERSITY *of* EDINBURGH
School of Engineering

Disclaimer: This presentation is based on Dr Murat Uney's research that was carried out at the University of Edinburgh under EPSRC/MoD University Defence Research Collaboration (UDRC) in signal processing Phase 2 programme.

Parts of the work presented is done in collaboration mainly with Prof Bernard Mulgrew (UoE), Dr Daniel Clark (Telecom SudParis), Dr Keith Copsey (Cubica Tech)

ERPem
Edinburgh Research Partnership in Engineering and Mathematics

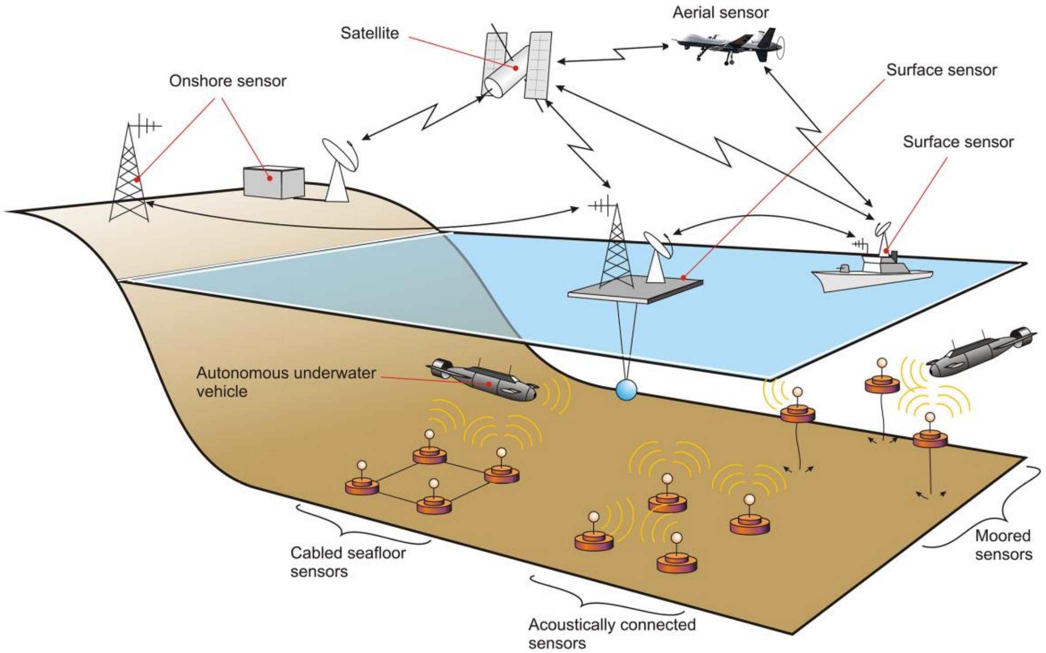

An underwater sensor network (inspired from [Akyildiz, Ad-hoc networks, 2005] and [Heideman, PToRSL, 2012]).

An autonomous sensor system for base and perimeter protection [Thomas, et al, Proc. SPIE, 2016]).

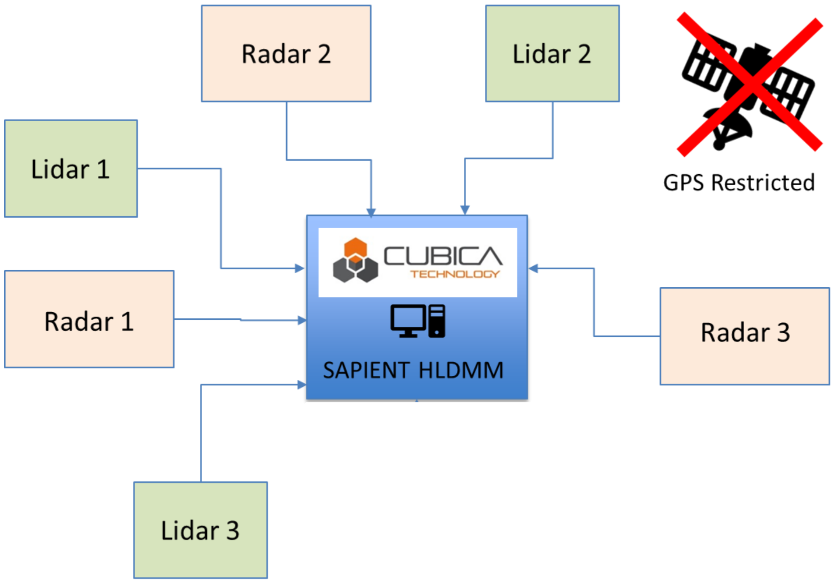

Functional view

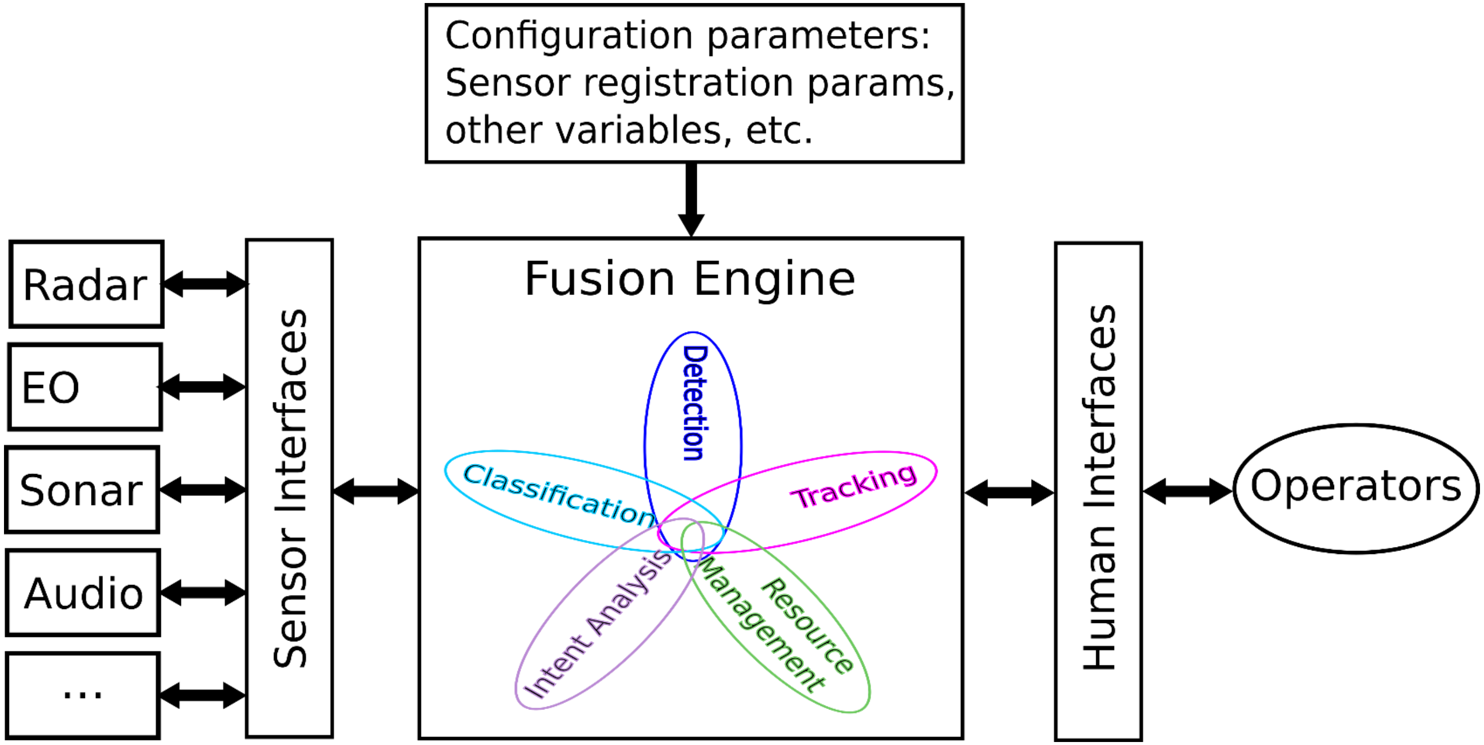

Functional view

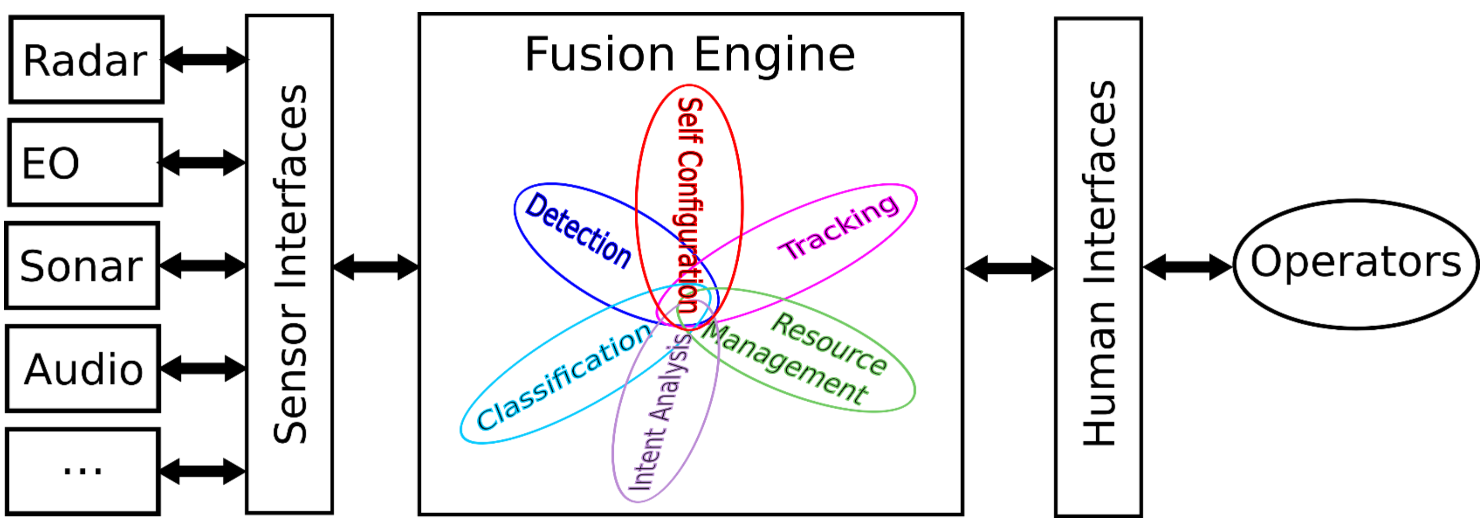

Functional view

$$p\left(X_{1:k},\theta\middle|Z_{1:k}^{1},\ldots,Z_{1:k}^{N}\right)=p\left(X_{1:k}\middle|Z_{1:k}^{1},\ldots,Z_{1:k}^{N},\theta\right)\underbrace{p\left(\theta\middle|Z_{1:k}^{1},\ldots,Z_{1:k}^{N}\right)}_{\propto\, l\left(Z_{1:k}^{1},\ldots,Z_{1:k}^{N}|\theta\right)p(\theta)}$$

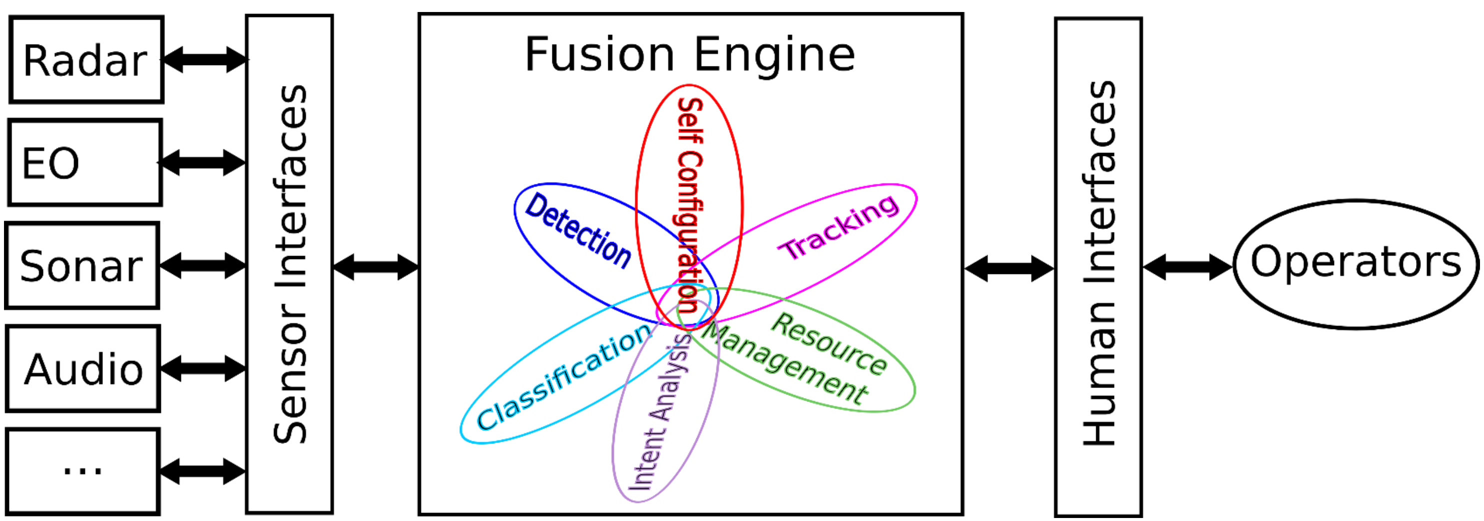

Contents

- Problem definition: Opportunistic self calibration
- Separable pseudo-likelihoods
- Markov random fields with separable likelihoods
- Self-calibration via belief propagation (BP)
- Demonstration on a SAPIENT network
- Conclusions

Problem Definition

- The concatenation of all configuration variables of a respective nature

$$\theta = [\theta_1, \ldots, \theta_N]$$

Example: Sensor locations and orientations with respect to a selected reference frame

$$\theta_j = [s_j, \alpha_j]$$

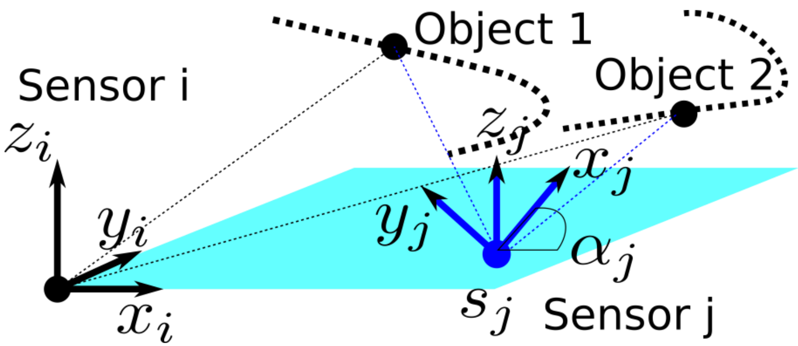

- The network-wide collected sensor data in the time window of k=1,…,t

$$Z \triangleq [Z^1_{1:t}, \ldots, Z^N_{1:t}]$$

Problem Definition

Multi-sensor multi-object state space model

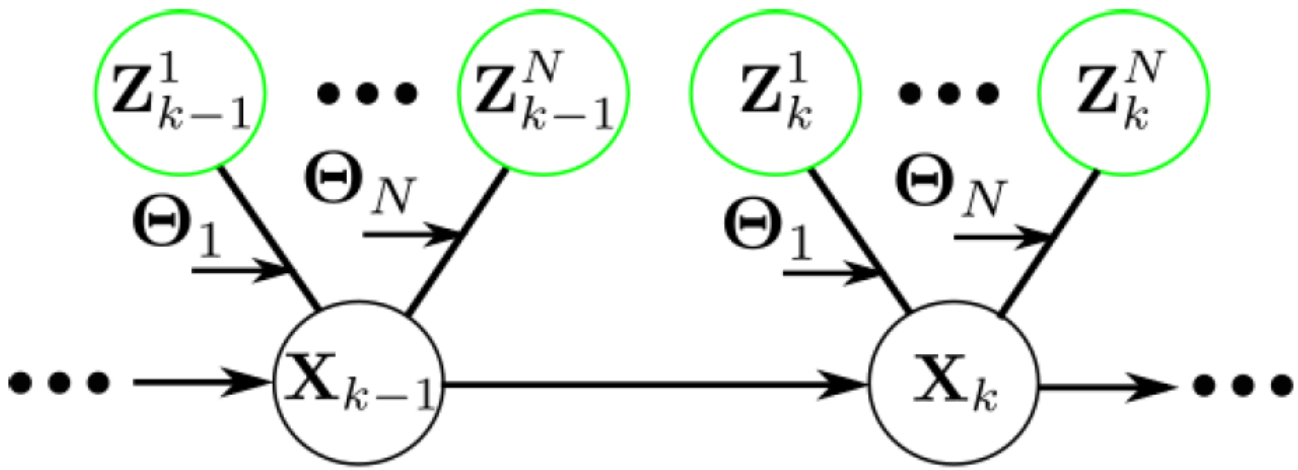

- $\mathbf{X}_k = \{\mathbf{x}_k^1, \ldots, \mathbf{x}_k^{M_k}\}$: Multi-object process evolving with a Markov shift $\pi(\boldsymbol{X}_k|\boldsymbol{X}_{k-1})$
- $\mathbf{Z}_k^i = \{\mathbf{z}_{k,1}^i, \ldots, \mathbf{z}_{k,L_k}^i\}$: Set valued measurement process characterised by $l_\mathrm{i}(\boldsymbol{Z}_k^i|\boldsymbol{X}_k, \theta)$
- Association uncertainty inherent in $l_i(.|.)$.

Problem Definition

Multi-sensor multi-object state space model

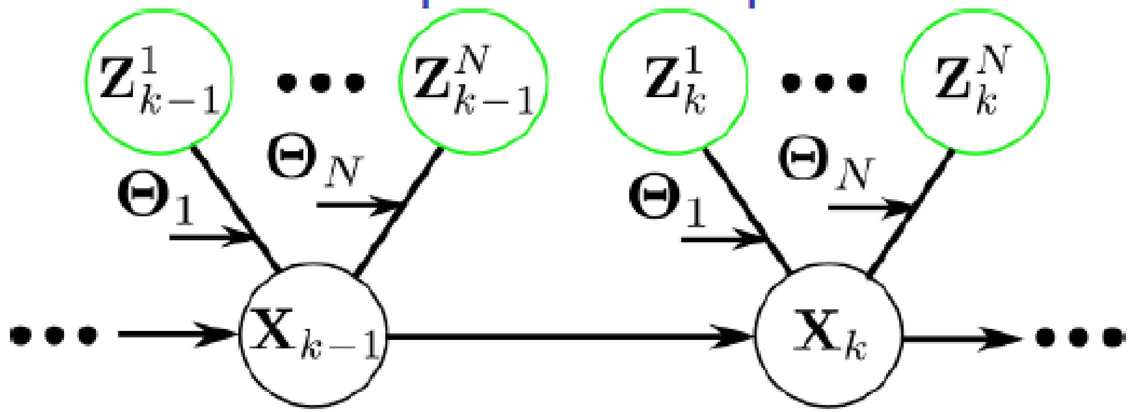

$p(X_{k-1}|Z^1_{1:k-1},\cdots,Z^N_{1:k-1},\theta)$ $\quad$ $p(X_k|Z^1_{1:k},\cdots,Z^N_{1:k},\theta) = \dfrac{\prod_{i=1}^N l_i(Z^i_k|X_k,\theta_i)\,p(X_k|Z^1_{1:k-1},\ldots,Z^N_{1:k-1},\theta)}{p(Z^1_k,\ldots,Z^N_k|Z^1_{1:k-1},\ldots,Z^N_{1:k-1},\theta)}$

Prediction $\quad$ Update

$p(X_k|Z^1_{1:k-1},\cdots,Z^N_{1:k-1},\theta)$

- Example: $N = 1$ sensor, $M_k = 1$ object, and, $L_k = 1$ measurement in a linear Gaussian model is Kalman filtering.
- Good solutions for $N = 1$, when $M_k, L_k > 1$ and association uncertainties present: Multi-object filtering algorithms.
- **For $N > 1$, combinatorial complexity with N.**

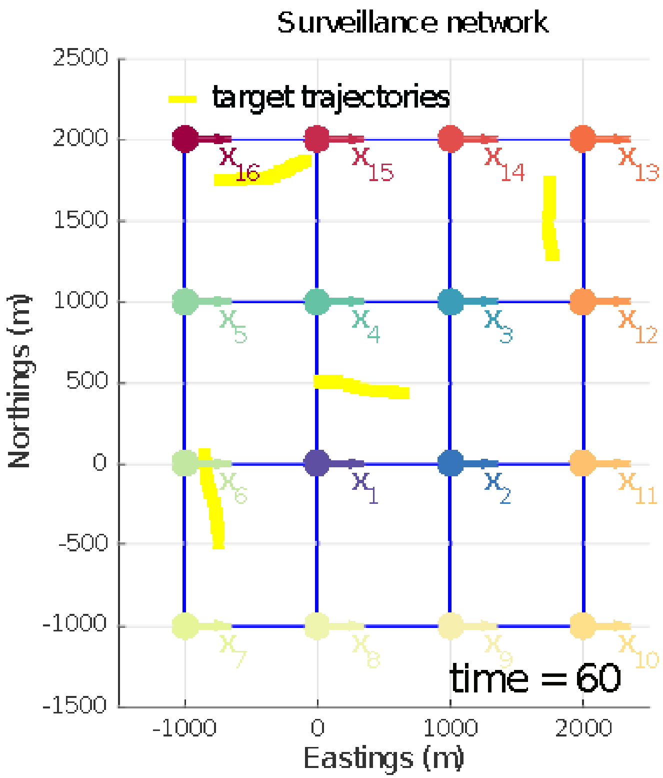

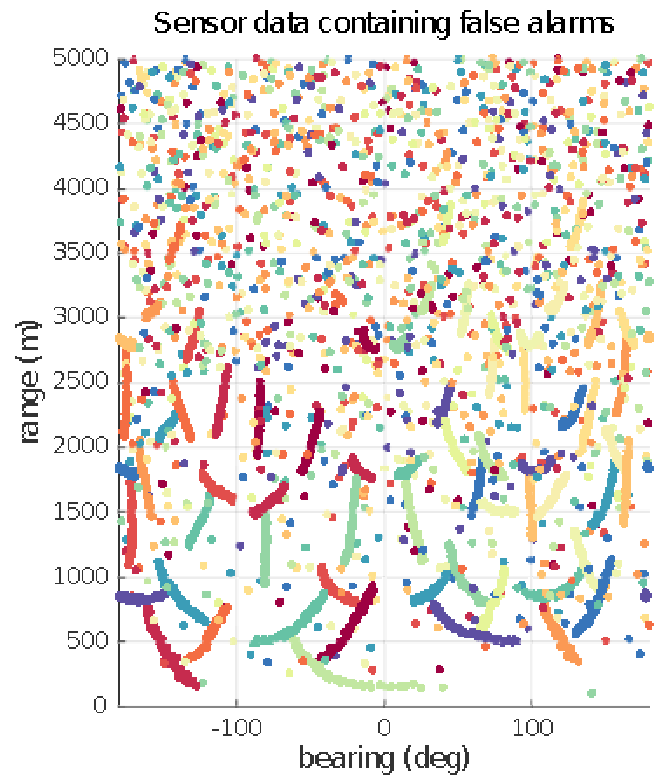

16 sensors collecting measurements from 4 targets with the following uncertainties:

i) unknown measurement-target association

ii) false alarms,

iii) less than one probability of detection.

Unknowns sensor locations and orientations (left pane), the network-wide collected data (right pane)

Problem Definition

Challenge 1: Intractability of the likelihood of θ

The likelihood for "parameter estimation in state space models"

$$l\left(Z^1_{1:t}, \ldots, Z^N_{1:t} | \theta = [\theta_1, \ldots, \theta_N]\right) = \prod_{k=1}^{t} p(Z^1_k, \ldots, Z^N_k | Z^1_{1:k-1}, \ldots, Z^N_{1:k-1}, \theta)$$

- Evaluation of likelihood update involves joint multi-sensor filtering:

$$p(Z^1_k, \ldots, Z^N_k | Z^1_{1:k-1}, \ldots, Z^N_{1:k-1}, \theta) =$$
$$\int \underbrace{\left(\prod_{i=1}^{N} l_i(Z^i_k | X_k, \theta)\right)}_{\text{Multi-sensor likelihood}} \times \underbrace{p(X_k | Z^1_{1:k-1}, \ldots, Z^N_{1:k-1}, \theta)}_{\text{Prediction distribution of a (centralised) Bayesian multi-sensor filter.}} \delta X_k$$

- **Combinatorially complex with the number of sensors N (dimensionality of θ)**

Problem Definition

Challenge 2: The likelihood is non-negligible only over a very small subset of the possible configurations (and potentially multiple local maxima).

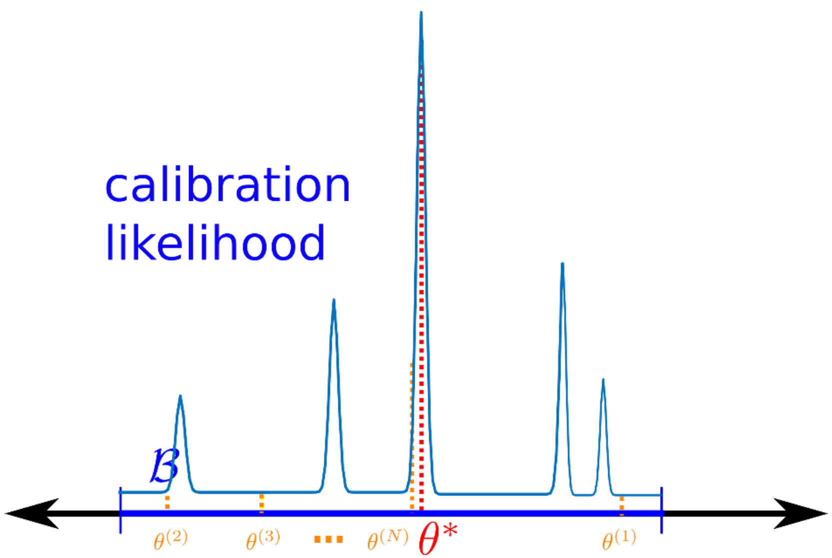

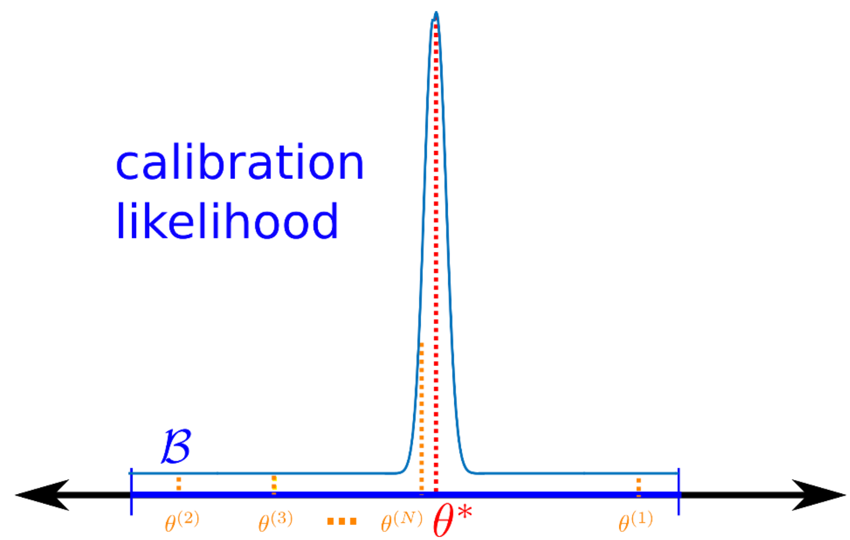

Hint: Potentially smoother if the time window length t is high.

Separable pseudolikelihoods

A general pseudo-likelihood form for problems with intractable likelihoods

$$Z \triangleq [Z^1_{1:t}, \ldots, Z^N_{1:t}].$$

$$\tilde{l}(Z|\theta) = \prod_{s \in S} \tilde{l}(Z_{d_s}|Z_{c_s}, \theta)^{\omega_s}$$

Varin, Reid, Firth “An overview of composite likelihood methods,” Statistica Sinica, no. 21, 2011.

A pairwise pseudo-likelihood on $\mathcal{G} = (\mathcal{V} = \{1, ..., N\}, \mathcal{E} \subset \mathcal{V} \times \mathcal{V})$

$$\tilde{l}(Z|\theta) = \prod_{(i,j) \in \mathcal{E}} l(Z^i_{1:k}, Z^j_{1:k}|\theta_{i,j})$$

- $l(Z^i_{1:k}, Z^j_{1:k}|\theta_{i,j})$ is intractable.

Separable pseudolikelihoods

Separable likelihoods

- Computationally feasible surrogates of $l(Z^i_{1:k}, Z^j_{1:k}|\theta_{i,j})$
- For example, approximations based on single sensor filtering, i.e.,
  1. local prediction $p(X_k|Z^j_{1:k-1})$
  2. local filtering distribution $p(X_k|Z^j_{1:k})$

  for $j \in \mathcal{V}$:
  1. provide scalability with the number of sensors
  2. align well with distributed data fusion (DDF) architectures
  3. exploit recent linear complexity updates for the single sensor multi-object model (e.g., PHD filtering).

     Mahler, Statistical multi-source multi-target information fusion, 2007.

Separable pseudo-likelihoods

Dual-term separable likelihood $\tilde{l}\left(Z^i_{1:t}, Z^j_{1:t}|\theta\right)$

approximates $l\left(Z^i_{1:t}, Z^j_{1:t}|\theta_{i,j}\right)$ with

$$\tilde{l}\left(Z^i_{1:t}, Z^j_{1:t}|\theta_{i,j}\right) = \prod_{k=1}^{t} s(Z^i_k, Z^j_k|Z^i_{1:k-1}, Z^j_{1:k-1}, \theta_{i,j})$$

$$s(Z^i_k, Z^j_k|Z^i_{1:k-1}, Z^j_{1:k-1}, \theta_{i,j}) \triangleq \underbrace{p(Z^i_k|Z^j_{1:k-1}, \theta_{i,j})}_{\int l(Z^i_k|X_k;\theta_{i,j})\,\underbrace{p(X_k|Z^j_{1:k-1}, \theta_{i,j})}_{\text{info from sensor } j \text{ to } i}\,\delta X_k} p(Z^j_k|Z^j_{1:k-1}, \theta_{i,j})$$

- Linear complexity evaluation when $\mathbf{X}_{k|k-1}$ is a Poisson random finite set

Proposition

Kullback-Leibler divergence between the exact and the dual-term update

$$
\begin{aligned}
&D(p(Z^i_k, Z^j_k | Z^i_{1:k-1}, Z^j_{1:k-1}, \theta_{i,j}) || s(.|.)) \\
&\quad \leqslant H(\boldsymbol{X}_k | \boldsymbol{Z}^i_{1:k-1}, \theta_{i,j}) - H(\boldsymbol{X}_k | \boldsymbol{Z}^i_{1:k-1}, \boldsymbol{Z}^j_{1:k-1}, \theta_{i,j}) \\
&\qquad + H(\boldsymbol{X}_k | \boldsymbol{Z}^j_{1:k-1}, \theta_{i,j}) - H(\boldsymbol{X}_k | \boldsymbol{Z}^i_{1:k-1}, \boldsymbol{Z}^j_{1:k-1}, \text{worst}\{\boldsymbol{Z}^i_k, \boldsymbol{Z}^j_k\}, \theta_{i,j}),
\end{aligned}
$$

where H denotes the Shannon Entropy

Corollary

Kullback-Leibler divergence between the exact and the dual-term separable likelihood is linearly bounded

$$
D(l(Z^i_{1:t}, Z^j_{1:t} | \theta_{i,j}) || \tilde{l}(Z^i_{1:t}, Z^j_{1:t} | \theta_{i,j})) \leqslant \max_{k=1,\ldots,t} \Delta_k \times t
$$

where Δ_k is the the right hand side of previous proposition.

Uney, Mulgrew, Clark "A cooperative approach to sensor localisation in distributed fusion networks," *IEEE Trans. Signal Proc., 2016*.

Separable pseudo-likelihoods

Separable likelihood asymptotics

$$p(X_{1:t}, Z^1_{1:t}, \dots, Z^N_{1:t} | \theta = [\theta_1, \dots, \theta_N]) = \mathcal{N}(.; \boldsymbol{\mu}(\theta), \boldsymbol{\Sigma}(\theta))$$

- $O(N^2 t)$ vs. $O(Nt)$ per likelihood evaluation
- Asymptotics become relevant as t grows, e.g., $\lim_{t\to\infty} D(p(\theta|Z_{1:t})||q(\theta|Z_{1:t}))$

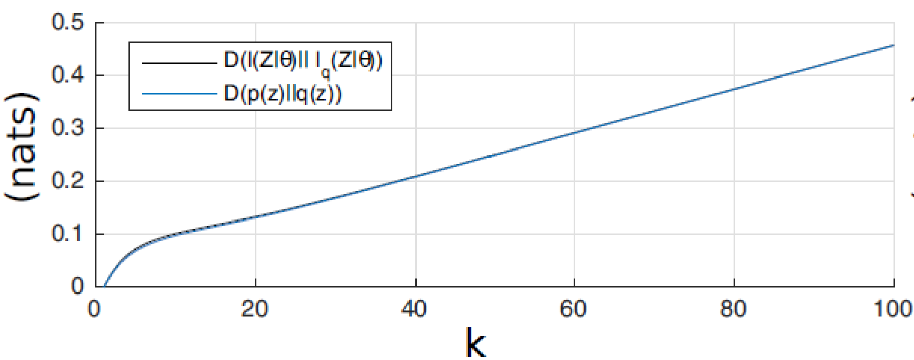

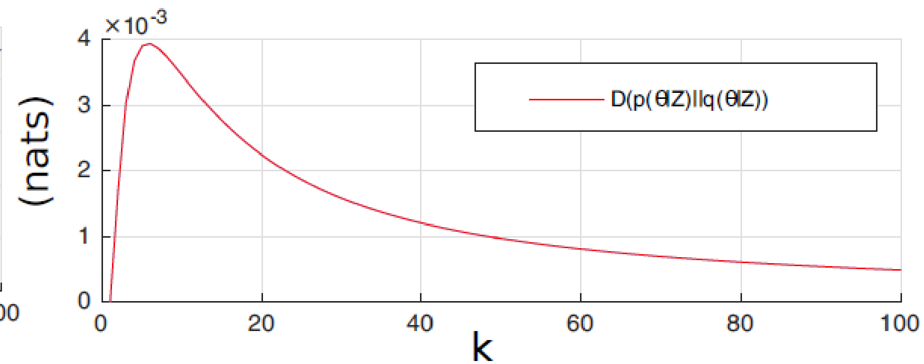

- $D(p(\theta|Z_{1:t})||q(\theta|Z_{1:t})) = D(l(Z_{1:t}|\theta)||\tilde{l}(Z_{1:t}|\theta)) - D(p(Z_{1:t})||q(Z_{1:t}))$

Example: Separable likelihoods in sensor localisation

$\theta_i = [\alpha_i = 0, s_i = [0,0]]$ (blue sensor) and $\theta_j = [\alpha_j = 0, s_j = [\delta x, \delta y]]$ (black sensor)

$$p(Z_k^i | Z_{1:k-1}^j, \theta_i, \theta_j) \quad = \quad \int l_i(Z_k^i | \tilde{X}) p(X | Z_{1:k-1}^j) \delta X$$

$$\tilde{X} \quad = \quad \{ \ \tilde{x} \ | \ \forall x \in X, \ \ \tilde{x} = R(\alpha_i)\big((R(\alpha_j)^T x + s_j\big) - s_i \}$$

$s\left(Z_k^i, Z_k^j \middle| ., ., \delta x, \delta y\right)$ $\quad$ $\tilde{l}(Z_{1:k}^i, Z_{1:k}^j | \delta x, \delta y)$

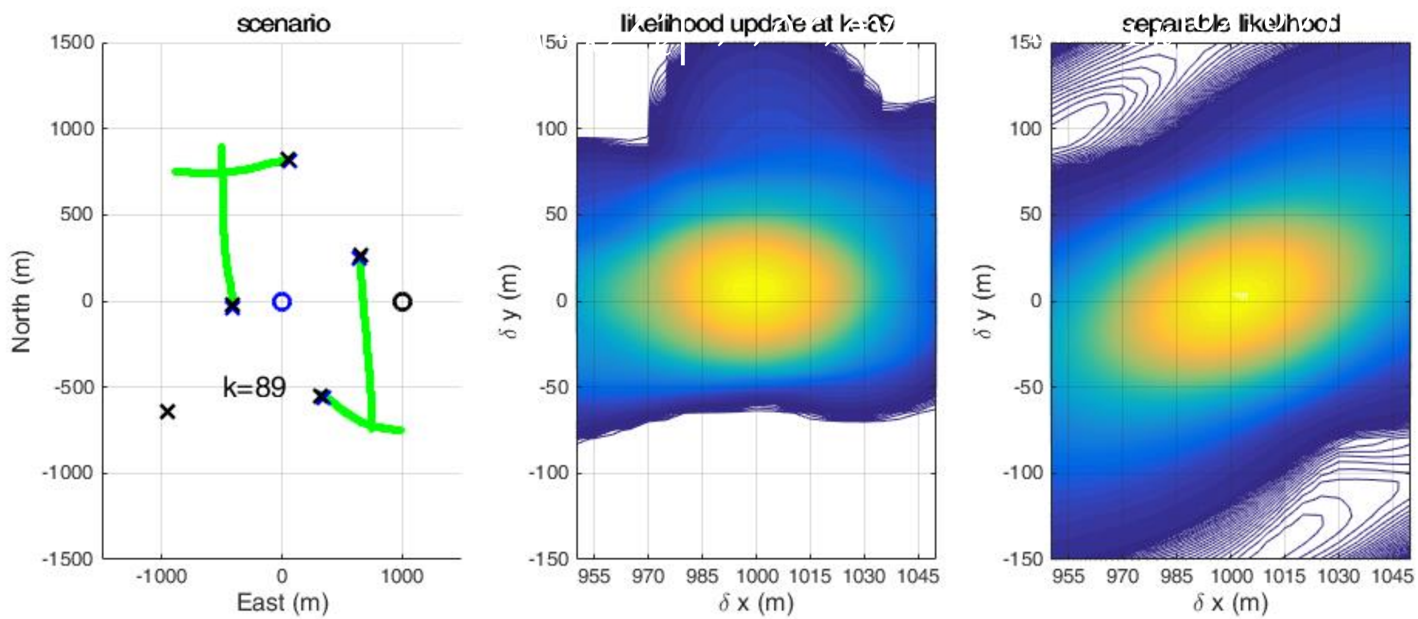

Separable pseudo-likelihoods

- Different approximations to $p(Z^i_k, Z^j_k | Z^i_{1:k-1}, Z^j_{1:k-1}, \theta_{i,j})$ are possible:

Quad-term separable likelihood $\tilde{l}\left(Z^i_{1:t}, Z^j_{1:t} | \theta_{i,j}\right)$

approximates $l\left(Z^i_{1:t}, Z^j_{1:t} | \theta_{i,j}\right)$ with $\prod_{k=1}^{t} q(Z^i_k, Z^j_k | Z^i_{1:k-1}, Z^j_{1:k-1}, \theta_{i,j})$

$$
\begin{aligned}
q(Z^i_k, Z^j_k | Z^i_{1:k-1}, Z^j_{1:k-1}, \theta_{i,j}) \triangleq \frac{1}{\kappa_k(\theta_{i,j})} & \left(p(Z^i_k | Z^j_{1:k}, \theta_{i,j}) p(Z^j_k | Z^j_{1:k-1}, \theta_{i,j})\right)^{1/2} \\
& \times \left(p(Z^j_k | Z^i_{1:k}, \theta_{i,j}) p(Z^i_k | Z^i_{1:k-1}, \theta_{i,j})\right)^{1/2}
\end{aligned}
$$

where $\kappa_k(\theta_{i,j})$ is the normalisation constant.

- It can be shown that $D(p(.|.)||q(.|.)) \leqslant D(p(.|.)||s(.|.))$ under reasonable conditions*.

Uney, Mulgrew, Clark "Latent parameter estimation in fusion networks using separable likelihoods," *IEEE Trans. Signal and Information Proc. Over networks 2018*.

Markov random fields with separable likelihoods

Latent parameter posterior

- $\tilde{l}(.|.)$ together with independent priors on θ_is lead to a pairwise MRF posterior $\mathcal{G} = (\mathcal{V}, \mathcal{E})$

$$p(\theta = [\theta_1, \ldots, \theta_N] | Z^1_{1:t}, \ldots, Z^N_{1:t}) \propto \prod_{i \in \mathcal{V}} \psi_i(\theta_i) \prod_{(i,j) \in \mathcal{E}} \psi^t_{ij}(\theta_i, \theta_j),$$

$$\psi_i(\theta_i) = p_{0,i}(\theta_i), \quad \psi^t_{ij}(\theta_i, \theta_j) = \tilde{l}(Z^i_{1:t}, Z^j_{1:t} | \theta_i, \theta_j)$$

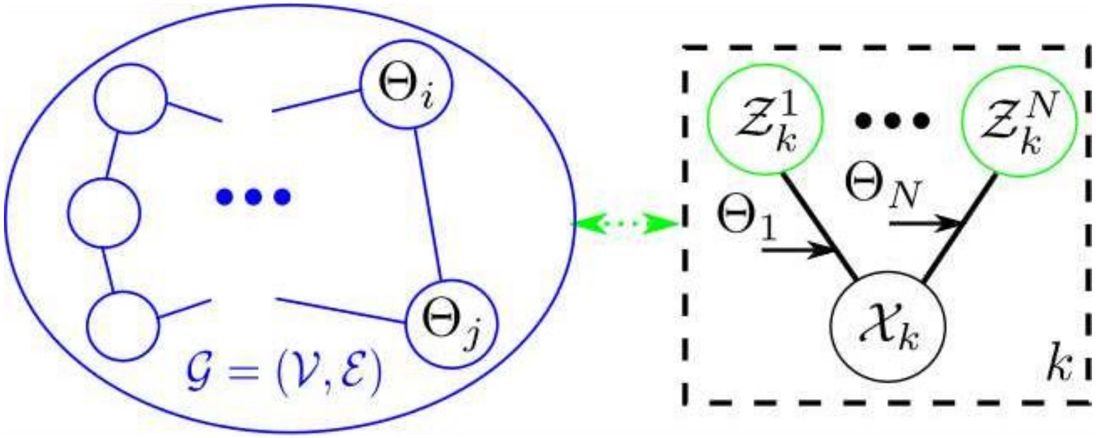

Markov random fields with separable likelihoods

$$p(\theta_1, \dots, \theta_N) \propto \prod_{(i,j)\in\mathcal{E}} l(.,.|\theta_i, \theta_j) \prod_{i\in\mathcal{V}} p_{0,i}(\theta_i)$$

The graph $\boldsymbol{\mathcal{G}}$ represents conditional independence relations by graph separation (Corollary to the Hammersley-Clifford theorem)

$$p(\theta_i | \theta_{\setminus i}) = p(\theta_i | \theta_{ne(i)})$$

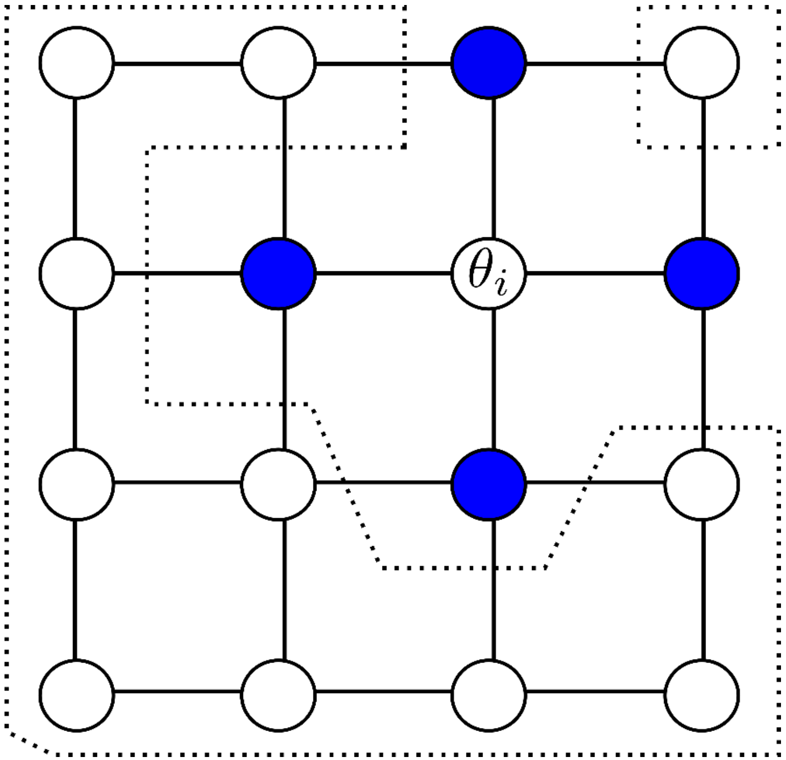

Self calibration via belief propagation

- Minimum mean squared error (MMSE) estimation $E\{\theta|Z\} = [E\{\theta_1|Z\}, \dots, E\{\theta_N|Z\}]$

- Belief propagation over $\mathcal{G}$ finds the marginals underlying the expectation in the MMSE estimate (approximately unless $\mathcal{G}$ is a tree) iteratively with the following messaging passing equations at step s

$$m_{ji}^{(s)}(\theta_i) = \int l(.,.|\theta_i, \theta_j) p_{0,j}(\theta_j) \prod_{i' \in ne(j) \backslash i} m_{i'j}^{(s-1)}(\theta_j)\, d\theta_j$$

$$\tilde{p}_i^{(s)}(\theta_i) \propto p_{0,i}(\theta_i) \prod_{j \in ne(i)} m_{ji}^{(s)}(\theta_i)$$

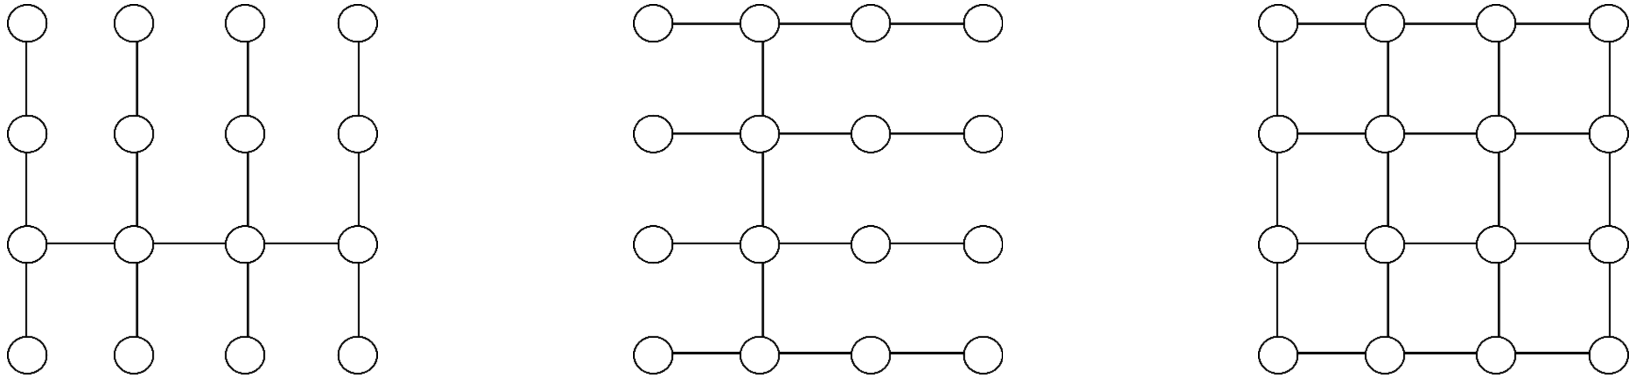

Wainwright, Jordan, “Graphical models, exponential families, and, variational inference,” *Foundations and Trends in Machine Learning 2008*.

Self calibration via belief propagation

Suppose we are given $\tilde{L}$ many configuration samples $\{\theta_i^{(l)}, \theta_j^{(l)}\}$. Approximate the edge potential by the Kernel sum

$$\psi_{i,j}(\theta_i, \theta_j) \approx \frac{1}{\tilde{L}} \sum \omega_{i,j}^{(l)} \mathcal{K}(\theta_i, \theta_j; \theta_i^{(l)}, \theta_j^{(l)})$$

$$\omega_{i,j}^{(l)} = \frac{\tilde{l}(Z^i, Z^j | \theta_i^{(l)}, \theta_j^{(l)})}{\sum_{l'=1}^{L} \tilde{l}(Z^i, Z^j | \theta_i^{(l)}, \theta_j^{(l)})}$$

For the case, BP messages become Kernel mixtures, as well. Samples from the node marginals are generated by

$$\theta_i^{(s),(l)} \sim \tilde{p}_{0,i}(\theta_i) \prod_{j \in ne(i)} m_{ji}^{(s)}(\theta_i)$$

$$E\{\theta_i | Z\} \approx \frac{1}{L} \sum_{l=1}^{L} \theta_i^{(s),(l)}$$

Uney, Mulgrew, Clark “Latent parameter estimation in fusion networks using separable likelihoods,” *IEEE Trans. Signal and Information Proc. Over networks 2018*.

Self calibration via belief propagation

- Efficiency in sampling and fast convergence via initialisation of node potentials with BP over a (spanning) tree:

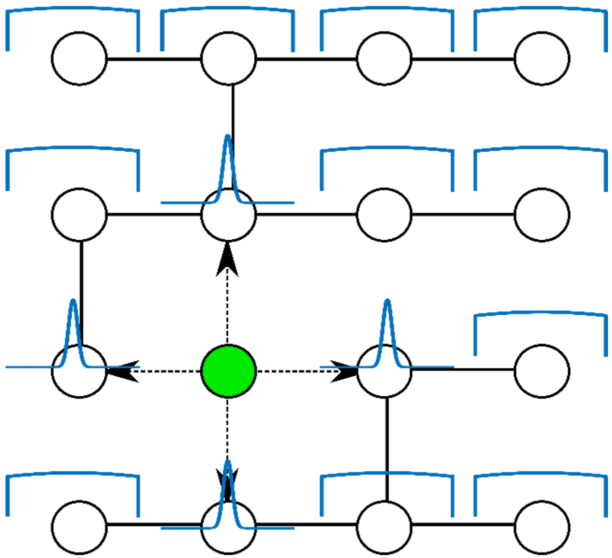

$$p_{\mathcal{T}_1}(\theta_1, \dots, \theta_N)$$

$$\propto \prod_{(i,j)\in \mathcal{T}_1} l(.,.|\theta_i, \theta_j) \prod_{i\in\mathcal{V}} p_{0,i}\,(\theta_i)$$

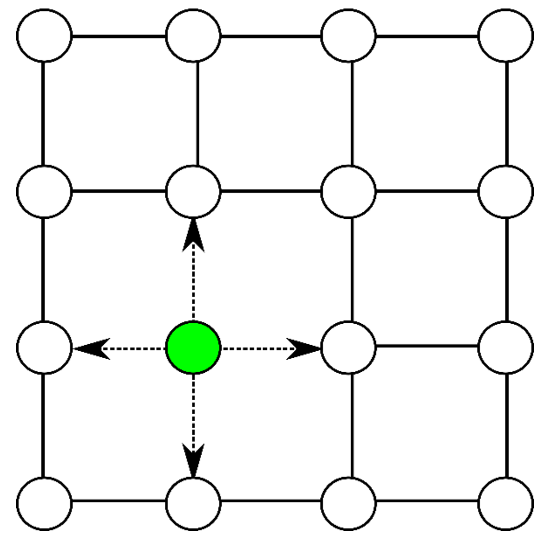

$$p_{\mathcal{G}}(\theta_1, \dots, \theta_N)$$

$$\propto \prod_{(i,j)\in \mathcal{T}_1} l(.,.|\theta_i, \theta_j) \prod_{i\in\mathcal{V}} p_{\mathcal{T}_1,i}\,(\theta_i)$$

Simulation Example

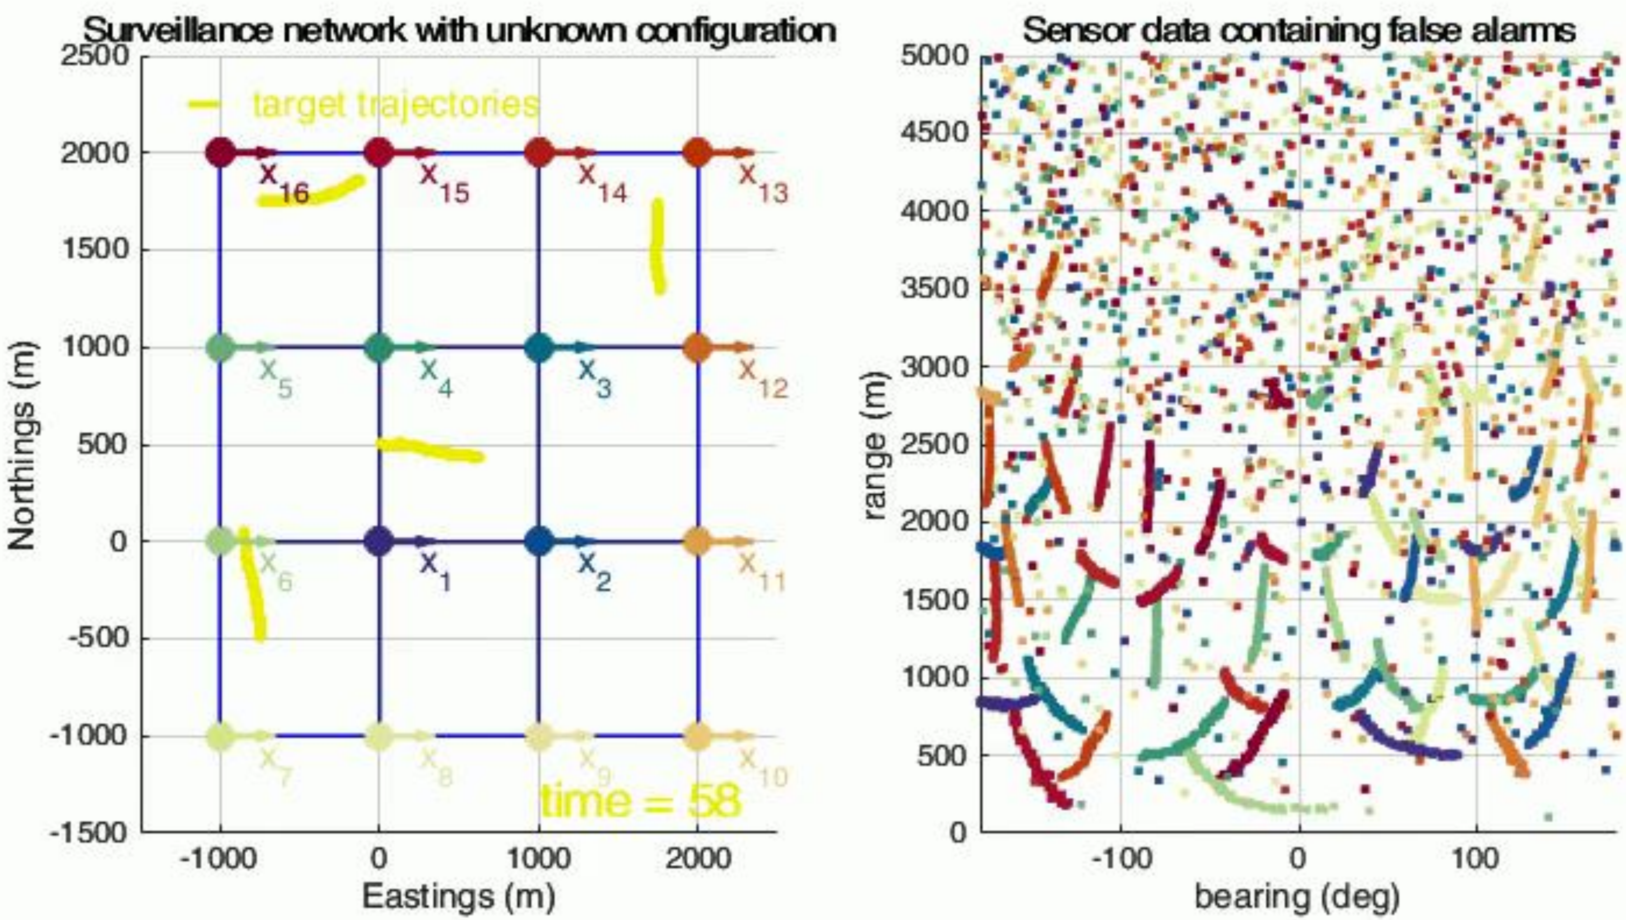

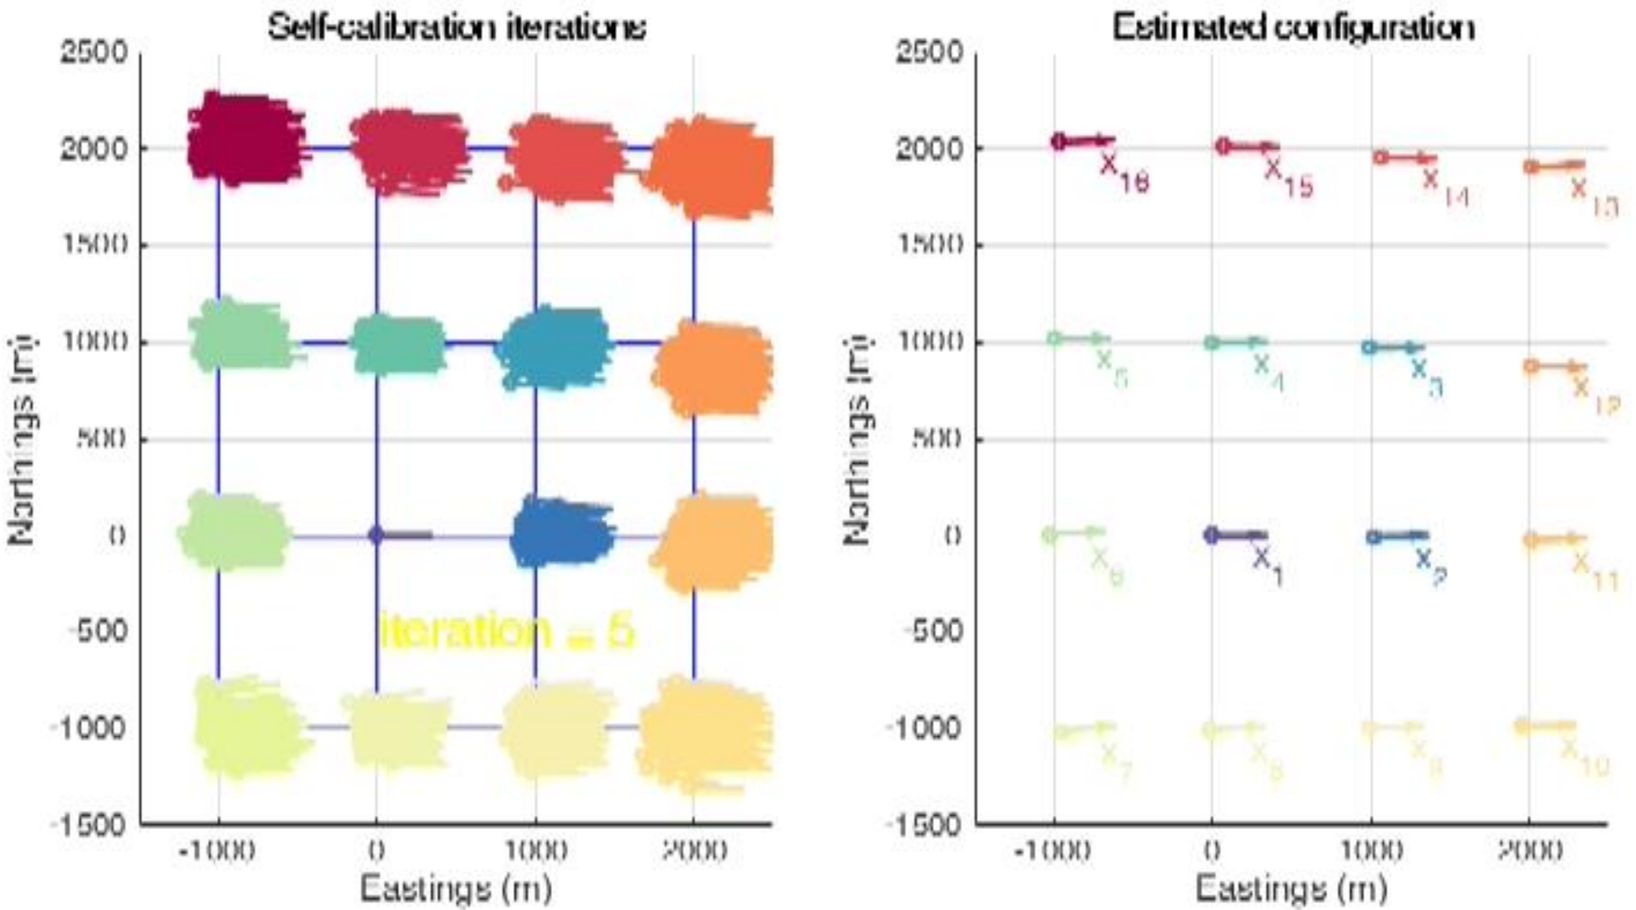

Particle BP over the pairwise MRF with separable likelihood edge potentials.

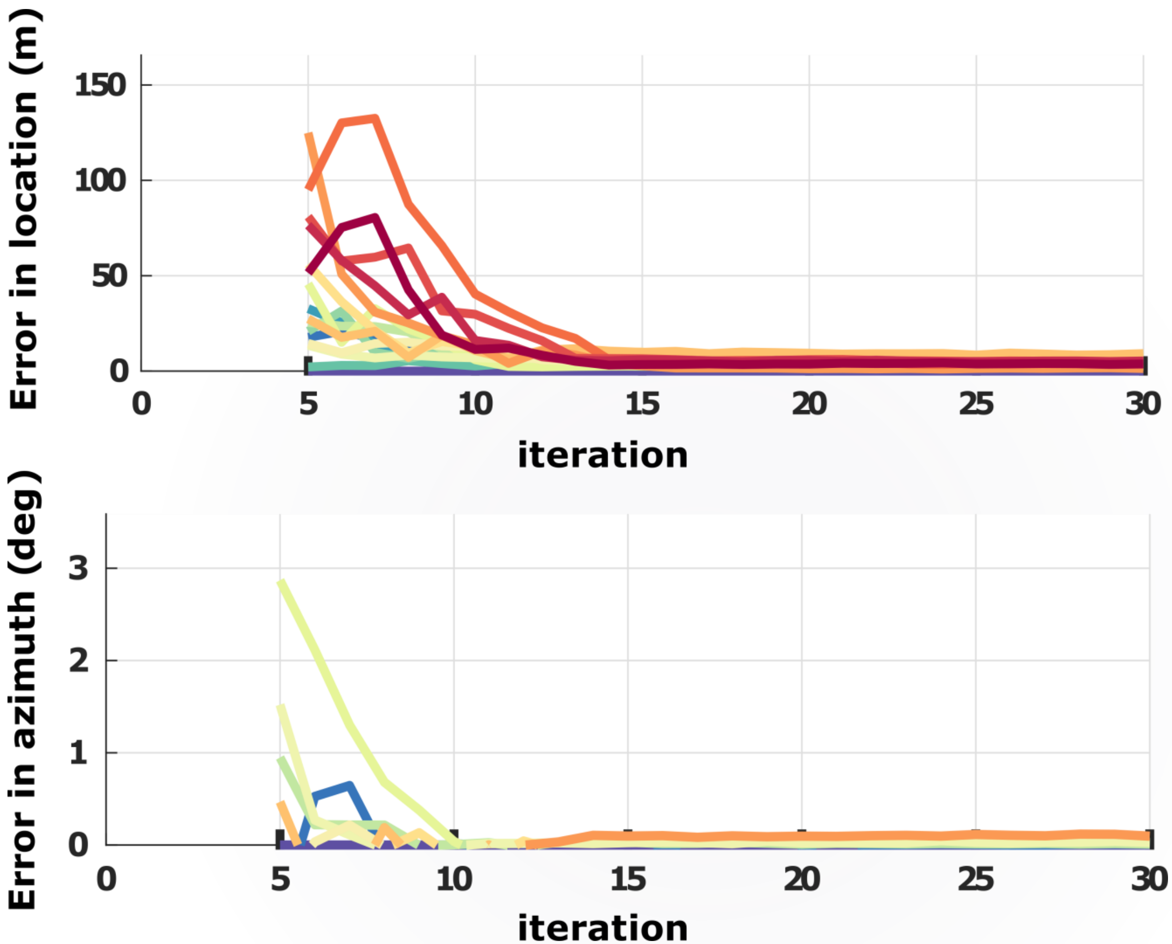

Location and orientation estimation error during particle BP iterations with the dual-term separable pseudo-likelihood.

Demonstration on a SAPIENT network

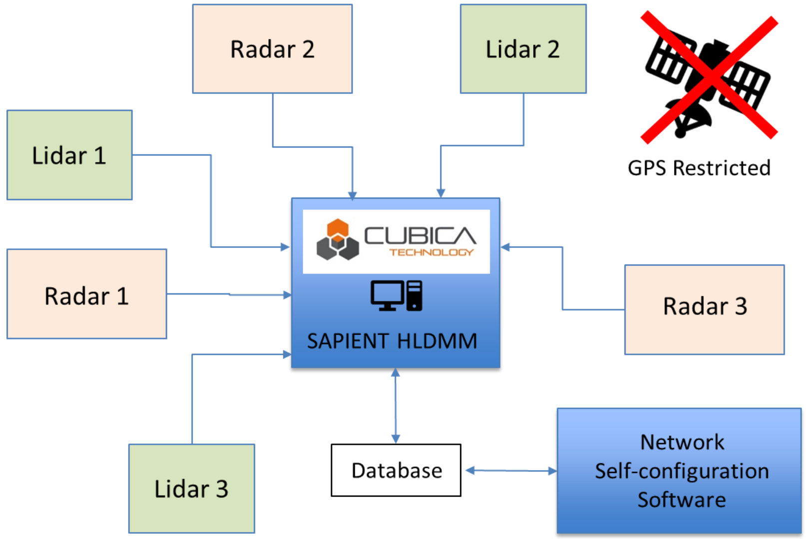

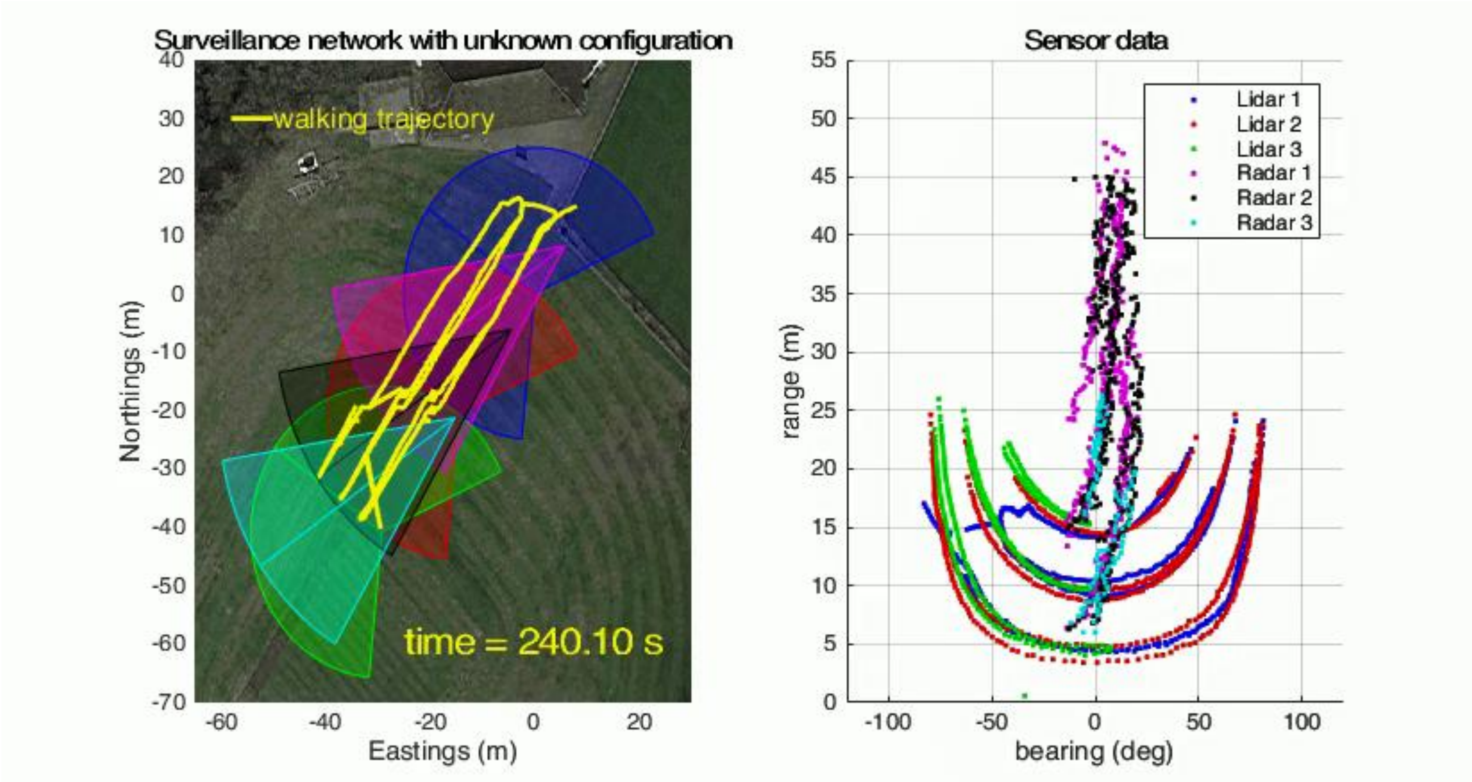

Data collection: (left) Sensor field-of-views and pedestrian trajectories, (right) Range-bearing measurements overlaid for the experiment period for all six sensors (colour coded)

Manually measured ground truth

Sensor #	Type	Location	Orientation
1 (blue)	Lidar	$[0 \quad 0]$ m	144.4°
2 (red)	Lidar	$[-14.6 \quad -20.4]$ m	144.4°
3 (green)	Lidar	$[-29.2 \quad -40.8]$ m	144.4°
4 (magenta)	Radar	$[5.8 \quad 8.2]$ m	209.4°
5 (black)	Radar	$[-4.4 \quad -6.1]$ m	209.4°
6 (cyan)	Radar	$[-15.2 \quad -21.2]$ m	209.4°

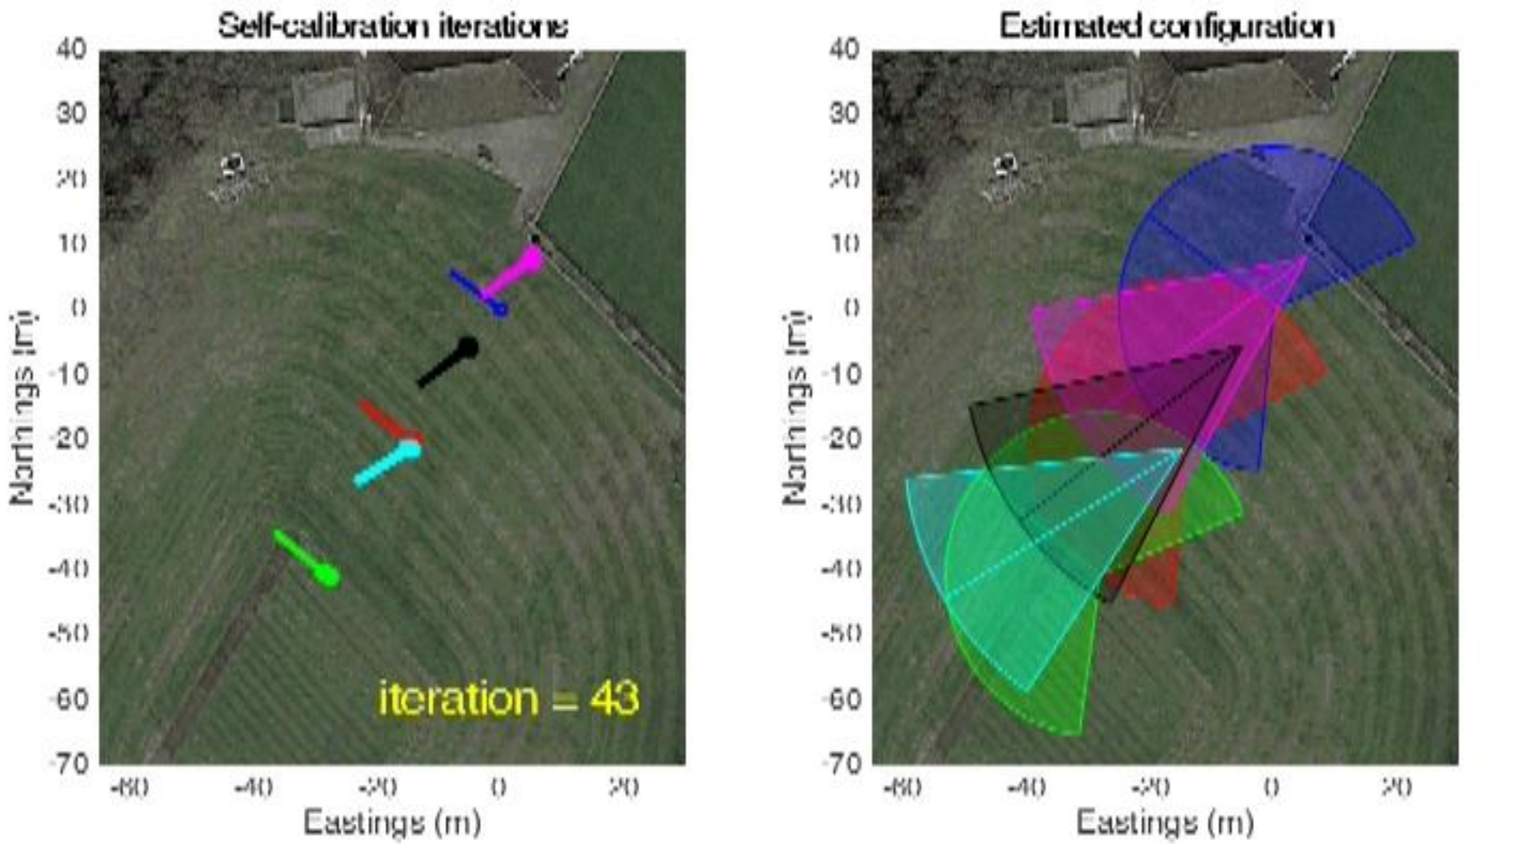

Real data with particle BP
over the graph on the left and
separable likelihood edge potentials.

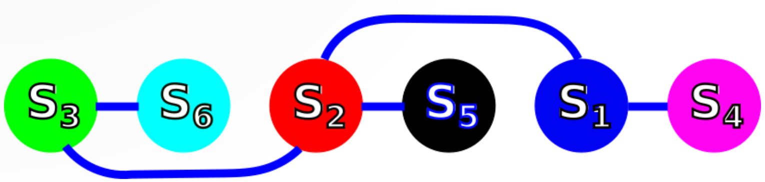

Demonstration on a SAPIENT network

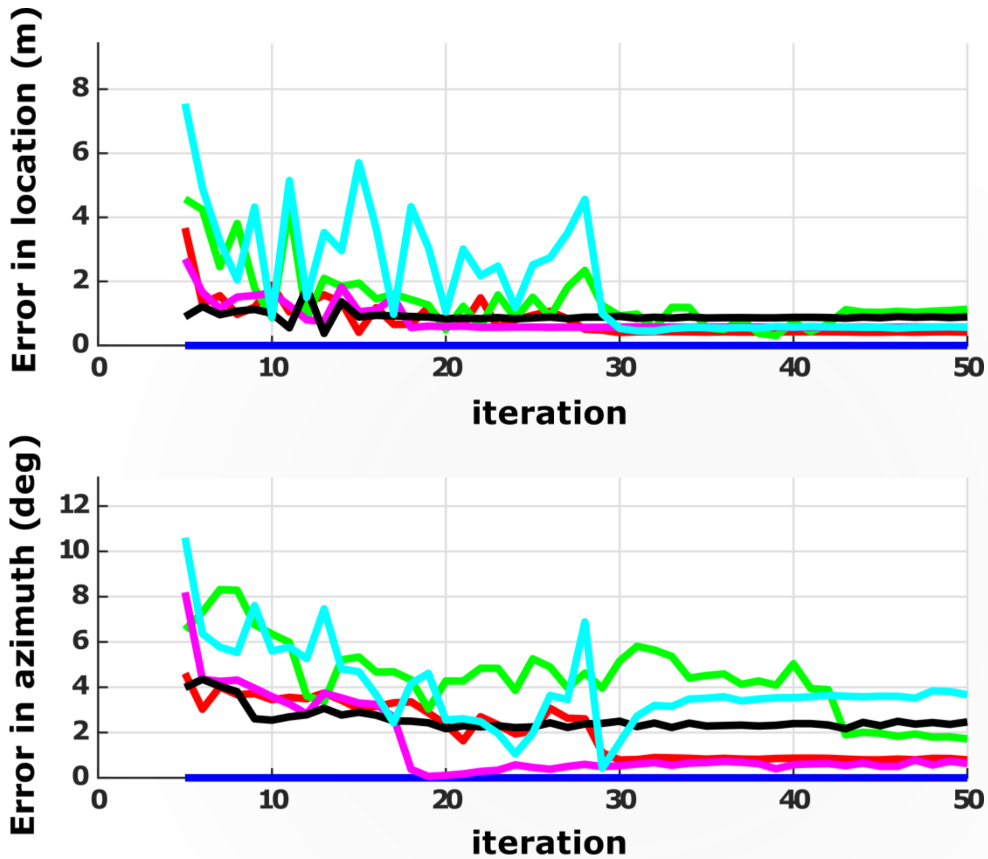

Error in location and orientation estimation with respect to the manually measured ground truth

Uney, Copsey, Page, Mulgrew, Mugrew, Thomas "Enabling self-configuration of fusion networks via scalable opportunistic sensor calibration," *SPIE Defence + Security 2018*.

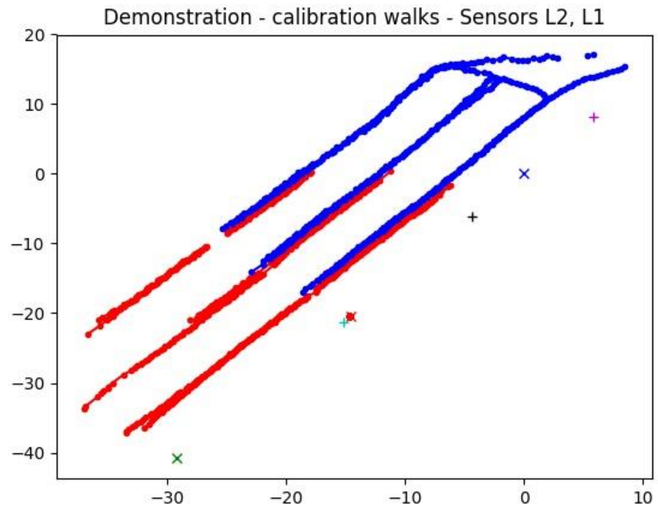

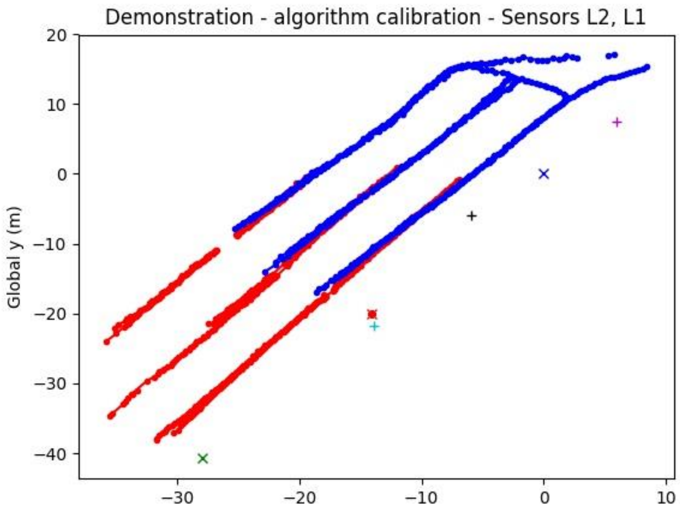

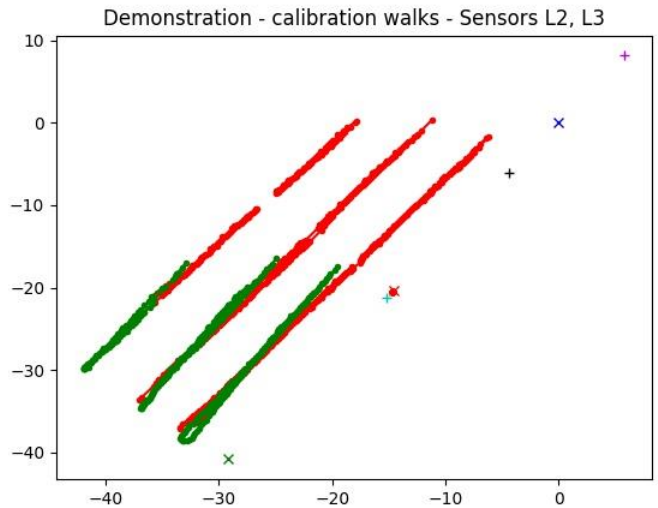

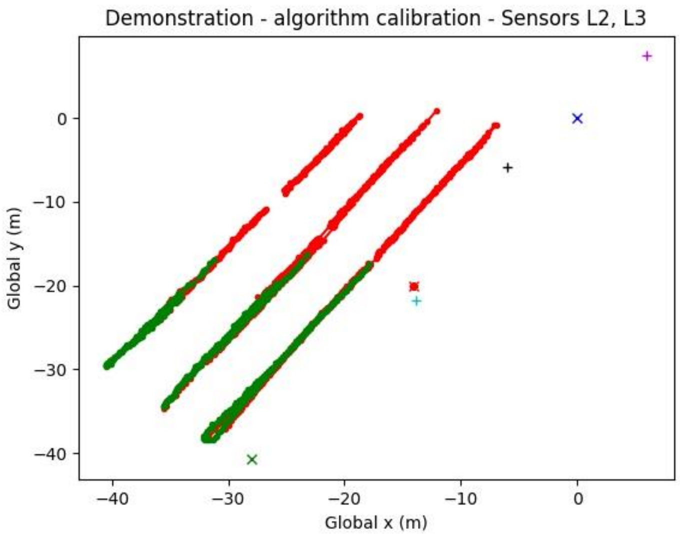

Lidar data projected to the global coordinate system using manual calibration parameters..

Lidar data projected to the global coordinate system using the automatic calibration parameters..

Conclusions

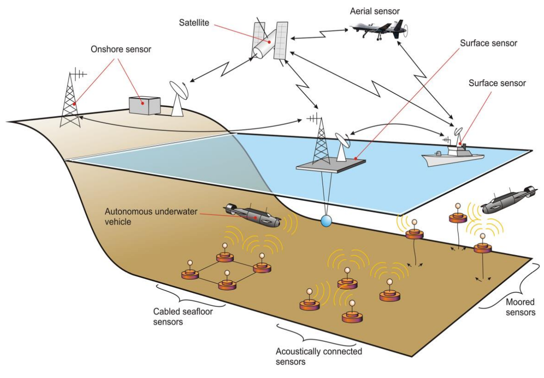

- Separable likelihoods are powerful approximations for scalable latent parameter estimation in multi-sensor fusion problems.
- MRF parameter posterior with separable likelihood edge potentials
- Particle BP for MMSE parameter estimation
- Demonstrated real world self-calibration capability on a SAPIENT network.
- MATLAB code available: https://github.com/muratuney
- Work in progress: Sampling in continuous valued MRFs
  - Tree reparameterised loopy BP for annealing the edge potentials

Thank you for your attention...
Questions?

Multi-object likelihood

$$l\left(\boldsymbol{Z}_{\boldsymbol{k}}^{\boldsymbol{i}} \middle| \boldsymbol{X}_{\boldsymbol{k}}; \theta_i\right)$$

$$= \mathcal{P}(Z; \lambda_{FA}) \prod_{x \in \boldsymbol{X}} (1 - P_D(x)) \sum_{\tau:\{1,\ldots,n\} \rightarrow \{0,1,\ldots,m\}} \prod_{i:\tau(i) \neq 0} \frac{P_D(x_i) l(z_{\tau(i)} | x_i ; \theta_i)}{\left(1 - P_D(x_i)\right) e^{\lambda_{FA}} \mathcal{P}(\{z_{\tau(i)}\}; \lambda_{FA})}$$