Enabling self-configuration of fusion networks via scalable opportunistic sensor calibration

Murat Üney*, Keith Copsey†, Scott Page†, Bernard Mulgrew*, and Paul Thomas‡

*University of Edinburgh, School of Engineering, EH9 3FB Edinburgh, UK
†Cubica Technology Ltd, Unit 4, Woking 8, Forsyth Road, GU21 5SB Woking, UK
‡Defence Science and Technology Laboratory (DSTL) , SP4 0JQ Porton Down, UK

ABSTRACT

The range of applications in which sensor networks can be deployed depends heavily on the ease with which sensor locations/orientations can be registered and the accuracy of this process. We present a scalable strategy for algorithmic network calibration using sensor measurements from non-cooperative objects. Specifically, we use recently developed separable likelihoods in order to scale with the number of sensors whilst capturing the overall uncertainties. We demonstrate the efficacy of our self-configuration solution using a real network of radar and lidar sensors for perimeter protection and compare the accuracy achieved to manual calibration.

Keywords: sensor registration, sensor networks, network self-configuration, pseudo-likelihood, belief propagation

1. INTRODUCTION

Multi-sensor exploitation enhances situational awareness significantly in a number of aspects including coverage, sensitivity and accuracy. One of the biggest obstacles in the way of fusion systems with large number of sensors is the difficulty in network calibration following sensor deployment. In particular, sensor data is collected in the local coordinate frame of the sensor, and, data from multiple sensors can be combined only after they are registered in a common coordinate system.[1] In order to do so, sensor locations and orientations need to be found out. This process also needs to be sufficiently accurate in order not to cause degredation in the target detection/tracking performance. Inaccuracies in these parameters might result in ghost tracks and erroneous estimates for the number of objects in the surveillance region.

In order to illustrate the complications involved, let us restrict ourselves to only locating the sensors. Sensor networks for underwater surveillance (see Fig. 1) cannot exploit global navigation space systems (GNSS) due to physical constraints of the environment on signal propagation.[2] Terresterial applications can benefit from the Global Positioning System (GPS) using on-board GPS receivers. This solution might fail to perform reliably given the vulnerability of GPS to, for example, jamming.[3] Another solution is to use trained personnel with special equipments such as laser location finders for surveying in sensor locations. This process, however, needs to be repeated if a new sensor is to be added or an existing sensor is knocked.

Other alternatives include the use of reference (or, cooperative) vehicles,[4] radio frequency front-end and/or network statistics such as received signal strength (RSS) and time of arrival (TOA).[5] Localisation based on RSS and TOA type noisy distance measurements, however, is often not sufficiently accurate for networks with the degree of connectivity typical in fusion applications [6, Chp.6], and, pose the same vulnerabilities as GPS receivers.

In all of the above approaches, network configuration is tedious, costly to repeat, and, prone to errors in the external equipment/sensors used. On the other hand, sensors can be algorithmically located with respect to a

Murat Üney and Bernard Mulgrew are with the Institute for Digital Communications (IDCOM), School of Engineering, University of Edinburgh (E-mail: {M.Uney, B.Mulgrew}@ed.ac.uk).
Keith Copsey and Scott Page are with Cubica Technology Ltd (E mail: {keith.copsey, scott.page }@cubicatechnology.co.uk, Homepage: https://www.cubicatechnology.co.uk).
Paul Thomas is with DSTL, Cyber and Information Systems Division (E-mail: pathomas@dstl.gov.uk)

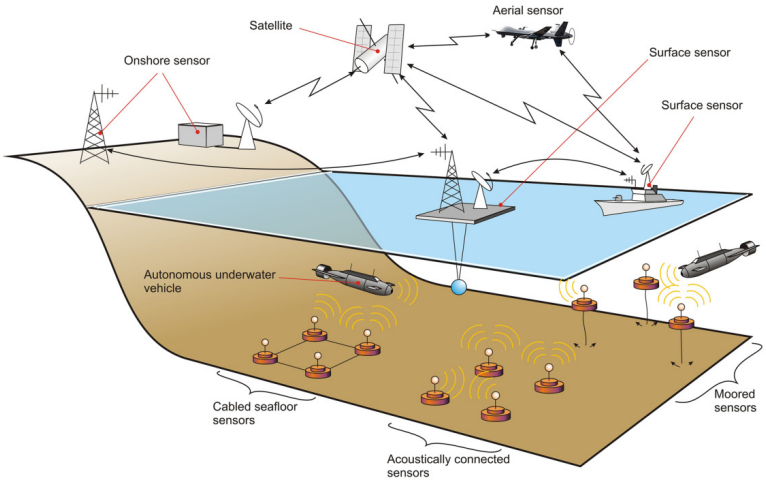

Fig. 1. An example underwater surveillance network (illustration inspired from[7] and[2]).

selected one as the origin using measurements already collected from the objects in the surveillance region.[8,9] The underlying mathematical problem is often referred to as latent parameter estimation in state space models (see, e.g.,[10] for a review) and the solutions thereof have been used to estimate, for example, sensor biases.[11,12] In this perspective, the parameters to be estimated might be "intrinsic" to the sensors and/or respective among the sensors. Network calibration problems exclusively consider the latter and often involve estimation of the sensor orientations together with their locations. The parameter likelihood of the problem, however, is intractable when typical uncertainties in the case of fusion scenarios are involved, e.g., unknown number of objects, unknown data association and false positives. In the best case, scalability with the number of sensors which specifies the dimensionality of the unknown should be addressed.

In order to address these challenges, pseudo-likelihood approaches have been developed, recently, which approximate the intractable exact model of the network calibration problem by a combination of local models.[13,14] These likelihood like functions can be evaluated using local Bayesian filtering of individual sensor streams. There is a high degree of flexibility in accommodating different sensing modalities (e.g. radar, lidar and/or visual) and filtering algorithms. The dual-term separable likelihood family[13] can be evaluated using random finite sets[15] tracking algorithms. The quad-term separable likelihoods are more accurate alternatives[14] which lead to feasible computational schemes only with hypothesis based trackers. Calibration parameters are estimated using a Bayesian approach in which separable likelihoods for pairs of sensors are combined with (non-informative) priors yielding a Markov random field posterior model. Sensor parameters are found using belief propagation (BP)[16] on this model.

In this article, we demonstrate network self-configuration capability underpinned by separable likelihoods using data from radar and lidar sensors in a SAPIENT (Sensing for Asset Protection with Integrated Electronic Networked Technology) autonomous sensor system for base and perimeter protection,[17] under UK Ministry of Defence funding. This capability is extremely useful for "movable" or "rapid deployment" scenarios where pre-made site measurements are not available. We outline sensor data processing procedures in order to exploit separable likelihoods. Real scenarios pose the additional challenge of handling asynchronous data which is handled in the pre-processing step. The accuracy achieved using scalable likelihoods within particle BP is demonstrated in comparison to manual calibration, which is the de-facto method used in this application. This approach can be used for the estimation of any such parameters that is needed to configure such networks. For scenarios with mobile sensors, the proposed framework can be used together or within simultaneous localisation and mapping (SLAM).

The structure of the article is as follows: An overview of separable likelihoods is given in Section 2 in which we introduce both dual-term and quad-term families. Section 3 explains computational methods for parameter estimation using separable likelihoods in BP. In Section 4 we present results with real data as aforementioned.

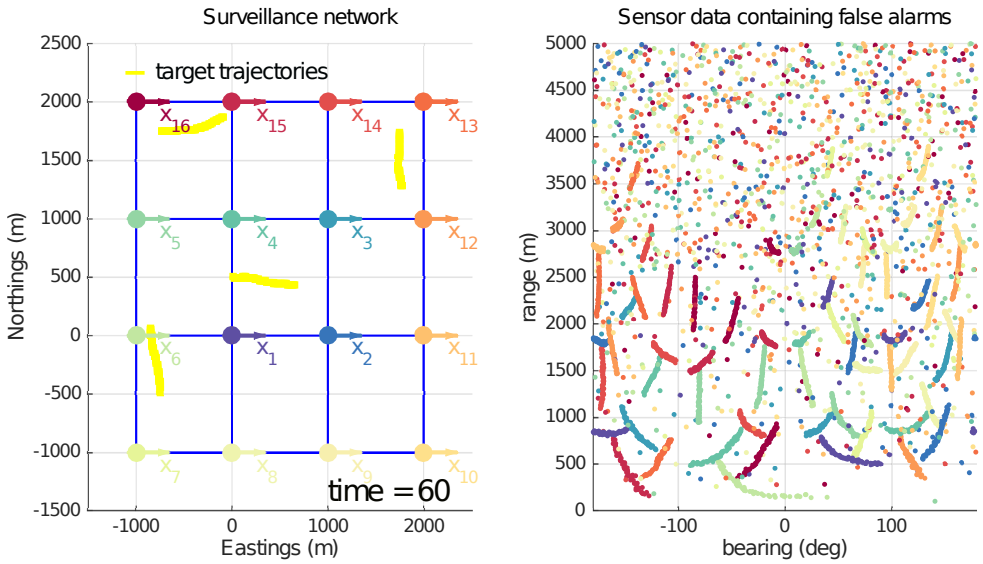

Fig. 2. A typical multi-object multi-sensor scenario: (left) 16 sensors (colour coded circles on the grid) collect data consisting of false alarms and noisy range-bearing measurements from 5 objects (yellow lines) taken with less than one probability and data association uncertainty, i.e., sensor measurements do not have labels that identify targets. (right) Network wide collected measurements overlaid for 60 time steps with colour codes associated with the sensor. The data is collected in the local coordinate frame of the associated sensor. The bearing measurement references (i.e., local x-axis) of sensors are depicted with arrowed lines on the left hand side figure – which are all aligned with the global East direction. The elongated tracks are object measurements and tend to appear at the bottom of the bearing-range plot as the objects are close to most of the sensors compared to their maximum range of 5000m, in this example.

First, the system structure is illustrated and then the pre-processing is outlined. Several figures illustrate the comparitive performance of the algorithm. Finally, we conclude in Section 5.

2. SEPARABLE LIKELIHOODS FOR NETWORK SELF-CONFIGURATION

Network configuration involves finding the values and registration of parameters that model respective quantities among the sensors as well those that are "intrinsic" to the sensors that specify the mapping of measurements onto a desired coordinate frame (see, e.g. [18, pg. 70]). For the sake of simplicity, let us consider the case in which individual sensors are calibrated so that intrinsic parameters are known [*], and, sensor locations and orientations need to be found in order to configure the network. Let us denote these variables by

$$\theta = [\theta_1, \ldots, \theta_N], \tag{1}$$

where $\theta_i = [\alpha_i, s_i]$ is the concatenation of the location s_i and the azimuth orientation angle α_i of sensor i with respect to an arbitrary (global) reference coordinate system.

The sensors collect measurements with a high level of uncertainty in typical fusion scenarios: There is an unknown and time varying number of objects in the surveillance region. The sensor measurements contain false positives. Measurements from objects are taken with less than one probability and contain noise. The object-measurement association –or, data association– is unknown, i.e., the measurements do not contain any indication of which object they are induced by, or, whether they are false alarms. Each sensor collects these measurements in its own coordinate system. An example scenario is illustrated in Fig. 2. On the left pane, sensor locations and local coordinate reference are depicted with circles and arrowed lines using colour codes. On the right pane, the same colour code is used to illustrate the network-wide measurements. In this example the global reference

[*]Unknown intrinsic parameters can be estimated in the same framework described in this section. The explicit relation of such parameters to the measurements vary for different modalities and sensors. This manuscript restricts the discussion to parameters with a specified role in network configuration.

frame is selected as that of the sensor 1 and the other sensors' orientations are perfectly aligned with the global frame. Let us denote the network-wide collected data in the time window of $1, \ldots, t$ by

$$\mathbf{Z} \triangleq [\mathbf{Z}^1_{1:t}, \ldots, \mathbf{Z}^N_{1:t}]. \tag{2}$$

The mathematical problem for network calibration using measurements from cooperative measurements is hence the estimation of θ (e.g., the colour coded circles and orientation angles of the arrowed lines on the left hand pane in Fig. 2) based on only $\mathbf{Z}$ (e.g., the colour coded scatter plot on the right hand pane in Fig. 2). Such inverse problems with uncertainties are solved by using the likelihood of the problem. One can proceed, for example with the maximum likelihood (ML) approach, or, the minimum mean squared error (MMSE) approach. The first finds the parameter configuration that maximises the likelihood whereas the latter combines this likelihood with prior parameter distributions and finds the expectation of the posterior. In both cases, the likelihood of θ, i.e. $l(\mathbf{Z}|\theta)$ needs to be evaluated. This function, however, is intractable in the sense that its exact evaluation does not scale with the number of sensors N,[13] the number of objects, and, otherwise does not lead to a closed form expression due to the non-linearities involved in the measurement process.[19]

The main reason that the parameter likelihood in multi-sensor state space models does not scale with the number of sensors is the multi-sensor filtering involved in the evaluation of the likelihood. The details of this point is left out of the scope of this manuscript and can be found in [13, Section II], [14, Section II]. Separable likelihoods overcome this problem by building upon filtering of only sensor streams individually for which there are several well established algorithms in the literature.[20] The network calibration likelihood takes the form

$$\begin{aligned} l(\mathbf{Z}|\theta) &\approx \prod_{(i,j)\in\mathcal{E}} \psi_{i,j}(\theta_i, \theta_j) \\ \psi_{i,j}(\theta_i, \theta_j) &= \tilde{l}(\mathbf{Z}^i, \mathbf{Z}^j|\theta_i, \theta_j) \end{aligned} \tag{3}$$

where $\mathcal{E}$ is a set of pairs of sensors, and, $\tilde{l}$ is a "separable" approximation to the exact likelihood for the pair (i, j) in that its evaluation requires only local filtering of the measurements collected by sensors i and j.

2.1 Dual-term separable likelihoods

Dual-term separable likelihoods replaces the calibration likelihood of the sensor pair (i, j) in (3) with[21]

$$\tilde{l}(\mathbf{Z}^i, \mathbf{Z}^j|\theta_i, \theta_j) = \prod_{k=1}^{t} p(\mathbf{Z}^i_k|\mathbf{Z}^j_{1:k-1}, \theta_i, \theta_j) p(\mathbf{Z}^j_k|\mathbf{Z}^i_{1:k-1}, \theta_i, \theta_j) \tag{4}$$

where the first factor on the right hand side is given by

$$\begin{aligned} p(\mathbf{Z}^i_k|\mathbf{Z}^j_{1:k-1}, \theta_i, \theta_j) &= \int l_i(\mathbf{Z}^i_k|\tilde{X}) p(X|\mathbf{Z}^j_{1:k-1}) \delta X, \\ \tilde{X} &= \{\ \tilde{x}\ |\ \forall x \in X,\ \ \tilde{x} = R(\alpha_i)\big((R(\alpha_j)^T x + s_j\big) - s_i\}, \end{aligned} \tag{5}$$

with X denoting a set valued multi-object state variable and the integral denoting a set integral,[15] both without the loss of generality[†]. The first term inside the integral is the likelihood of sensor i. The state variable argument of the sensor likelihood $\tilde{X}$ is a set with the elements of X transformed from sensor j's coordinate frame to sensor i's coordinate frame with the proposed values of the calibration variables θ_i, θ_j. Here, $R(\alpha_j)$ is the Euler rotation matrix for the set of azimuth, elevation and pitch angle vector α_j which represents the orientation of sensor j with respect to the (selected) global coordinate system. The other term inside the integral is the prediction distribution for the multi-object state from a random finite set filtering algorithm that processes the sensor stream of sensor j [‡].

[†]Similar expressions can be found for vector valued multi-object states estimated by hypothesis based tracking algorithms.

[‡]The case that θ_j contains intrinsic parameters can be handled by using the version of this density that is also conditioned on θ_j.

The second factor in (4) is found by switching i and j in (5). It has been proved that the accuracy of this approximation is related to the accuracy of the local trackers.[13] The information contained in the measurements are combined over a time window of length t which renders the global approximation (3) fairly accurate for active range-bearing sensing modalities. The term in (4) leads to a closed form expression when the prediction density is propagated by a probability hypothesis density (PHD) filter[13] also when the sensor field-of-view only partially overlap.[22]

2.2 Quad-term separable likelihoods

An alternative separable term to use in (3) is a combination of four density terms[23] given by

$$\tilde{l}(\mathbf{Z}^i, \mathbf{Z}^j|\theta_i, \theta_j) = \prod_{k=1}^{t} \frac{1}{\kappa_k} \left(p(\mathbf{Z}_k^i|\mathbf{Z}_{1:k}^j, \theta_i, \theta_j) p(\mathbf{Z}_k^i|\mathbf{Z}_{1:k-1}^i) p(\mathbf{Z}_k^j|\mathbf{Z}_{1:k}^i, \theta_i, \theta_j) p(\mathbf{Z}_k^j|\mathbf{Z}_{1:k-1}^j) \right)^{1/2}, \tag{6}$$

where the factors are similar to (5), and, κ_k is a scaling factor to ensure that the square rooted term behaves as a regular probability density function for the measurements $\mathbf{Z}_k^i$ and $\mathbf{Z}_k^j$.

The quad-term approximation is provably more accurate than the dual-term approximation, under typical sensing conditions.[14] The computation of κ_k, however, does not lead to closed form expressions for set valued multi-object state variables. Instead, vector valued multi-object state variables and associated prediction/posterior distributions propagated by hypothesis based tracking algorithms[20] can be used with (6).

3. BELIEF PROPAGATION WITH SEPARABLE LIKELIHOODS AS EDGE POTENTIALS

The pairwise structure of the separable pseudo-likelihood in (3) is beneficial when it is used with prior distributions for the unknowns. The parameter posterior corresponding to a selection of priors and the pseudo-likelihood is a pairwise Markov random field over $\mathcal{G} = (\mathcal{V}, \mathcal{E})$:[16]

$$p(\theta|\mathbf{Z}) \propto \prod_{i \in \mathcal{V}} \psi_i(\theta_i) \prod_{(i,j) \in \mathcal{E}} \psi_{ij}(\theta_i, \theta_j), \tag{7}$$

$$\psi_i(\theta_i) = p_{0,i}(\theta_i), \tag{8}$$

$$\psi_{i,j}(\theta_i, \theta_j) = \tilde{l}(\mathbf{Z}^i, \mathbf{Z}^j|\theta_i, \theta_j), \tag{9}$$

where $\mathcal{V} = \{1, \ldots, N\}$ is the set of sensor nodes, and, $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is a set of pairs (selected arbitrarily, in this work). Here, ψ_i is a node potential function which is the selected prior for sensor i (e.g., uniform distributions over bounded sets θ_i take values from) in this model, and, ψ_{ij} is an edge potential function which equals to the dual/quad-term separable likelihood for the pairs (i, j).

The MMSE estimate in this model is the concatenation of the expected values of posterior marginals, i.e., $p(\theta_i|\mathbf{Z})$ for $i = 1, \ldots, N$. These marginal distributions are output by the belief propagation (BP) message passing algorithm exactly when $\mathcal{G}$ is a tree. BP iterations are given by

$$m_{ji}^{(s)}(\theta_i) = \int \psi_{ij}(\theta_i, \theta_j)\, \psi_j(\theta_j) \prod_{i' \in ne(j) \setminus i} m_{i'j}^{(s-1)}(\theta_j)\, \mathrm{d}\theta_j, \tag{10}$$

$$\tilde{p}_i^{(s)}(\theta_i) \propto \psi_i(\theta_i) \prod_{j \in ne(i)} m_{ji}^{(s)}(\theta_i), \tag{11}$$

for all $i \in \mathcal{V}$, where s is the iteration index and $\tilde{p}_i$ is the local belief over θ_i that is updated by messages from the neighbours $j \in ne(i)$. For a tree graph, after a certain number of iterations related to the diameter of the tree, $\tilde{p}_i^{(s)}$ converges to the marginal posterior of concern. The expectations of $\tilde{p}_i^{(s)}$s constitute the MMSE estimate of the network configuration, for the case.

Alternatively, $\mathcal{G}$ can be selected to have more sensor pairs contribute to the end result thereby containing cycles. For cyclic graphs, BP message and update equations in (10) and (11) can still be used by the nodes

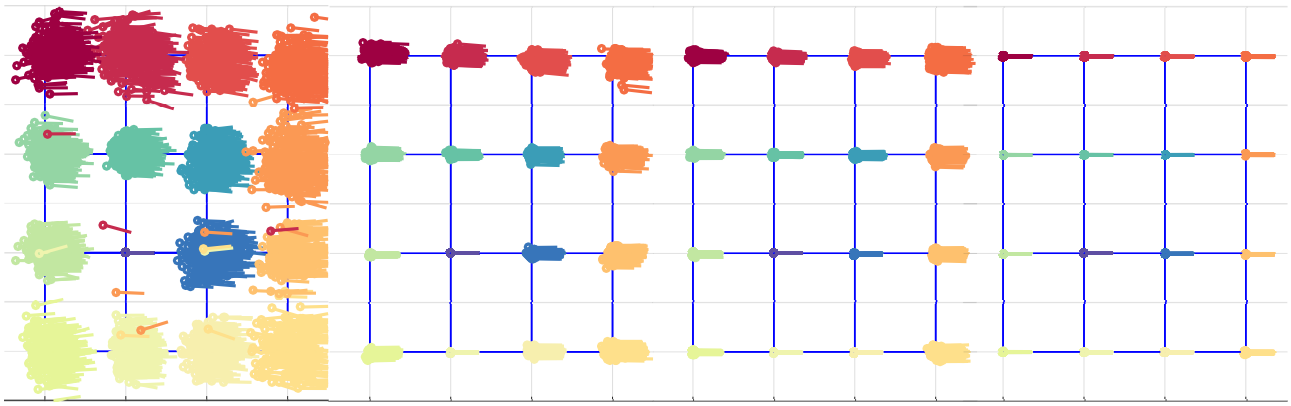

Fig. 3. Particle BP iterations self-calibrating the network in the left pane of Fig. 2 using the data in the right pane: For the iteration steps $4, 8, 10$ and 15, scatter plots of the particles generated from node beliefs are depicted. Each particle represets a location and an orientation depicted with the colour code of the sensor used in Fig. 2.

simultaneously. This message passing scheme is known as loopy BP (LBP), and, is not guaranteed to converge.[16] Nevertheless, in this problem setting it is often the case that models over spanning trees of a loopy $\mathcal{G}$ are consistent in that they lead to similar marginal parameter distributions. This suggests the existence of loopy BP fixed points[24] that will be converged when priors (equivalently, node potentials) are selected reasonably.[25]

In summary, separable pseudo-likelihoods enable us to benefit from the scalability properties of pairwise MRF models and parallel message passing algorithms over them in order to compute the MMSE estimation of network calibration parameters. Next, we overview computational procedures that unleash these benefits.

3.1 Particle BP in self-calibration

The pairwise MRF model in (7) and the BP iterations in (11) do not lead to tractable expressions. In order to realise the proposed approach, we have developed Monte Carlo computational procedures and particle representations for the densities/likelihoods involved. Key to these developments is a kernel mixture[26] representation for the edge potentials given their evaluations at particular values of the configuration parameters selected by particle sampling.

Suppose we are given $\tilde{L}$ pairs of points $\{(\theta_i^{(l)}, \theta_j^{(l)})\}$ and (dual-term/quad-term) separable likelihood evaluations $\tilde{l}(\mathbf{Z}^i, \mathbf{Z}^j | \theta_i^{(l)}, \theta_j^{(l)})$ for $l = 1, \ldots, \tilde{L}$ (see, (4) and (6)). The proposed approximation is given by[14]

$$\psi_{i,j}(\theta_i, \theta_j) \approx \frac{1}{\tilde{L}} \sum \omega_{i,j}^{(l)} \mathcal{K}(\theta_i, \theta_j; \theta_i^{(l)}, \theta_j^{(l)}) \tag{12}$$

$$\omega_{i,j}^{(l)} = \frac{\tilde{l}(\mathbf{Z}^i, \mathbf{Z}^j | \theta_i^{(l)}, \theta_j^{(l)})}{\sum_{l'=1}^{L} \tilde{l}(\mathbf{Z}^i, \mathbf{Z}^j | \theta_i^{(l)}, \theta_j^{(l)})}, \tag{13}$$

where $\mathcal{K}$ is a Gaussian radial basis function.[26] This representation leads to a kernel mixture representation for the BP messages (10) as in the non-parameteric BP algorithm.[27] The details of the resulting kernel mixture representations for the BP messages in (10) can be found in [14, Section VI.C], together with procedures for sampling from the beliefs in (11). As a result, independent identically distributed (i.i.d.) samples from the (scaled) product of the ith local belief are output at each step s of the BP iterations, i.e.,

$$\theta_i^{(s),(l)} \sim \tilde{p}_{0,i}(\theta_i) \prod_{j \in ne(i)} m_{ji}^{(s)}(\theta_i) \text{ for } l = 1, ..., L. \tag{14}$$

where $p_{0,i}(\theta_i)$ is the prior selected for i (see (8)). The empirical average of these samples constitute the MMSE estimate sought for.

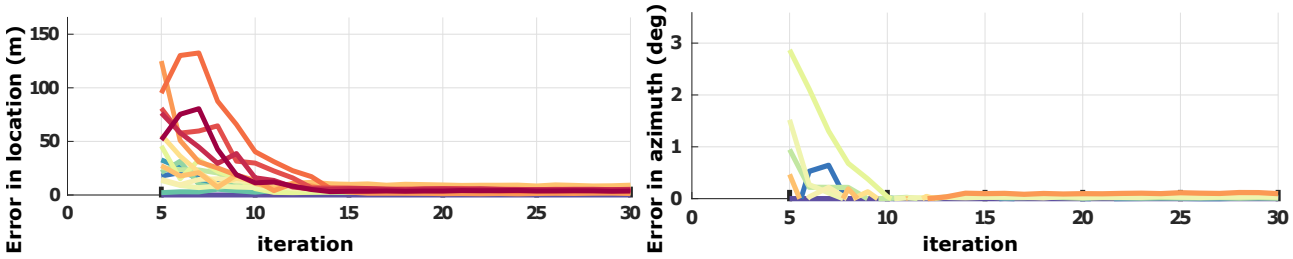

Fig. 4. Location and orientation estimation error during particle BP iterations with the dual-term separable pseudo-likelihood.

This particle BP approach is used for self-calibration of the sensor network in Fig. 2. Each sensor is associated with a 3-dimensional vector representing its location on the plane and azimuth orientation. Therefore, the dimensionality of the network configuration vector θ is 48. The global coordinate system is selected as the local coordinate frame of sensor 1, i.e., sensor 1 is the network anchor. This is reflected to the BP by selecting a Dirac's delta prior for θ_1 in (8). As a result, the dimensionality of the space over which the posterior is distributed is 45. Samples from the prior calibration distributions are selected as equally spaced grid points over $[\ -2000,\quad 2500\ ]\times[\ -2000,\quad 2500\ ]$ (the entire planar zone in Fig. 2), and, between $+10^\circ$ and -10° which mimics uniform distributions over the aforementioned space. The number of grid points used is 125. The number of particles generated from the node beliefs at each iteration (see (14)) is 100. Therefore, particle BP explores this space with a total of 1500 particles, which is a fairly small number [§]. The separable likelihood edge potentials used are dual-term likelihoods computed using the Gaussian Mixture PHD filter,[28] in this example.

The convergence speed and the accuracy achieved with so few particles is remarkable. Fig. 3 illustrates loopy BP iterations through scatter plots of particles representing the location and orientation of sensors. For the first four iterations, a tree graph is used to propagate information from sensor 1 towards its one-hop, two-hop, three-hop, and, four-hop neighbours (at iterations $1, 2, 3$ and 4, respectively). The beliefs at the end of these four steps is given in the first plot in Fig. 3. These beliefs are used as node potentials (equivalently, samples from the prior distributions) for the remaining iterations for which the graph $\mathcal{G}$ is also updated to the cyclic pairwise graph in Fig. 2 depicted with blue lines. In Fig. 3, the other plots are hence the node beliefs after loopy BP iterations $8, 10$ and 15 (left-to-right).

The resulting error in location and orientation estimation are given in Fig. 4. Note that convergence is reached before 15 iterations with a remarkable accuracy in both location and orientation configuration for all sensors. The particle efficiency and convergence speed are underpinned by the tree BP initialisation which effectively bridges non-informative calibration priors to distributions selected based on measurements.

4. REAL DATA EXPERIMENTS

In this section we demonstrate the proposed scalable network self-configuration approach for the calibration of a SAPIENT autonomous sensor network for perimeter protection.[17] The experimental setup is illustrated in Fig. 5 and consists of three lidar and three radar sensors feeding into the Cubica Technology Higher Level Decision Making Module (HLDMM) which is an integrated data fusion and sensor management system as described in.[29] Measurements from the sensors are recorded in a database which drives the network self-configuration software. The parameters found are then used to run the data through the HLDMM and situational awareness performance is assessed.

Fig. 6 illustrates the network configuration details of which are given in Table 1. The lidar sensors cover a 240° field-of-view (FoV) with a range of 25m (blue, red and green sectors). The radar FoVs are sectors with 40° angle and 25m range (magenta, black and cyan sectors). The global coordinate frame's origin is the location of sensor 1. The orientation angles in Table 1 are between the sensor boresights (solid lines splitting the FoV into equal sectors in Fig. 6) and the East direction.

[§]The Kernel representation in (14) is used with all the sample points evaluated thus far.

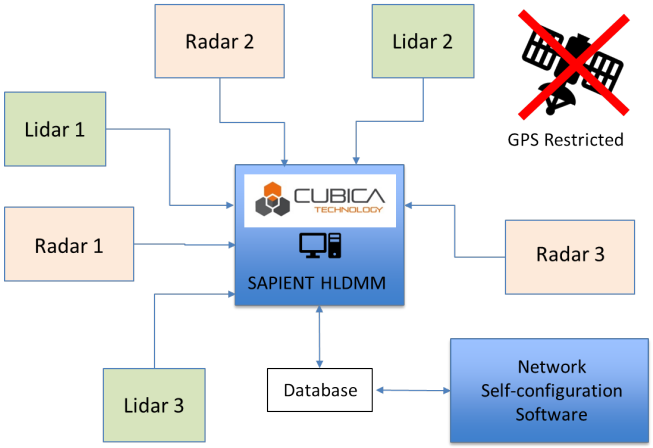

Fig. 5. Block diagram of the experimental system: Three lidar and three radar sensors form a SAPIENT network. Detections from the sensors are simultaneously recorded in a database which feeds into the network self-configuration software implementing particle belief propagation with separable likelihood edge potentials.

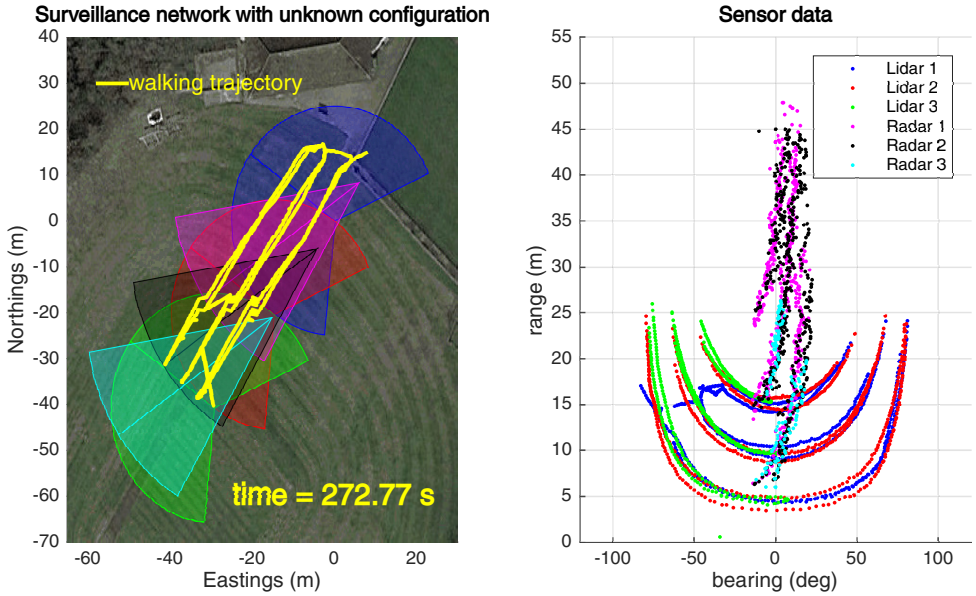

Fig. 6. SAPIENT network and data collection: (left) Configuration of the network with the parameters given in Table 1 and the pedestrian trajectory inducing sensor measurements. (right) Range-bearing measurements overlaid for the experiment period for all six sensors (colour coded).

The sensors collect range-bearing measurements from pedestrians walking in their FoVs. The data points are depicted in the right pane in colour code associated with sensors in the left pane. The goal is to find the locations and orientations of sensors $2, 3, 4, 5$ with respect to sensor 1 solely using the range-bearing data depicted on the right pane in Fig. 6. The self-configuration software in Fig. 5 implements the separable likelihood edge potentials and the particle BP algorithm given in Sections 2 and 3 together with a pre-processing stage for handling the asynchronous nature of the data.

4.1 Pre-processing for network self-calibration

The evaluation of any calibration likelihood including the separable approximations requires the identification of time periods in which pairs of sensors collect measurements from the same object(s). In order to fully automatise calibration, we developed a time stamp matching algorithm which exploits the fact that the sensors transmit

Table 1. SAPIENT network configuration (manually measured ground truth)

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

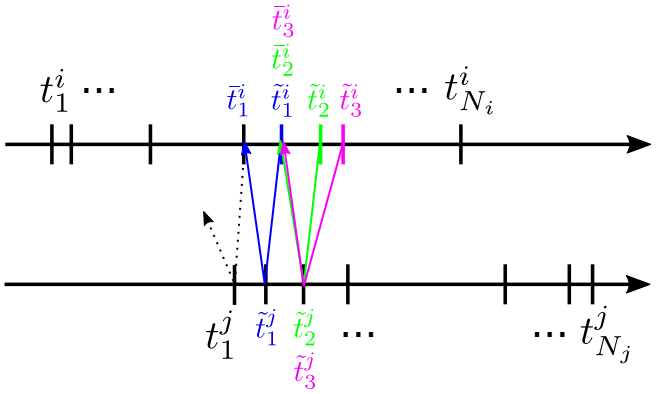

Fig. 7. Handling asynchronous data: Illustration of the time stamp matching based selection of the sensor i measurements that contribute to the separable likelihood. The blue, green and magenta coloured and tilde marked measurements are selected to contribute as their matching time stamps have reverse assignments.

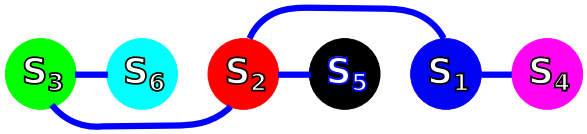

Fig. 8. The Markov graph $\mathcal{G}$ used for the real data experiment.

reports upon only detection of targets. The operation of the algorithm is illustrated in Fig. 7. Time stamps of sensor i are paired with those of sensor j which are collected in the last two seconds period and can also be paired back with one of sensor i's measurements on this basis. This approach produces fast and accurate results when the targets in the scene move together in a way that guarantees that sensors detect (at least one) common object. The separable likelihood in (4) is adopted for these asynchronous measurements as

$$\tilde{l}(\mathcal{Z}^i, \mathcal{Z}^j|\theta_i, \theta_j) = \prod_{(\tilde{t}^i,\tilde{t}^j)\in S^i} p(z^i_{\tilde{t}_i}|z^j_{t^j_1:\tilde{t}^j}, \theta_i, \theta_j) \prod_{(\tilde{t}^j,\tilde{t}^i)\in S^j} p(z^j_{\tilde{t}_j}|z^i_{t^i_1:\tilde{t}^i}\theta_j, \theta_i), \tag{15}$$

where,

$$S^i = \{(\tilde{t}^i, \tilde{t}^j)|\tilde{t}^i - \tilde{t}^j \in (0, 2] \; and, for some \; \bar{t}^i \in \{t^i_1, \ldots, t^i_{N_i}\}, \tilde{t}^j - \bar{t}^i \in (0, 2]\}. \tag{16}$$

4.2 Scalable network self-calibration

The pre-processing stage outputs data set pairs to use when evaluating the pseudo-likelihood edge potentials in particle BP (Section 3). A useful feature of the proposed framework is that the edge potentials for different pairs do not have to use data collected during the same time period. Thus, the proposed framework is capable of handling the cases in which not all sensor FoVs have an overlap by selecting different measurement time windows for different sensor pairs. This provides great flexibility for handling different network configurations.

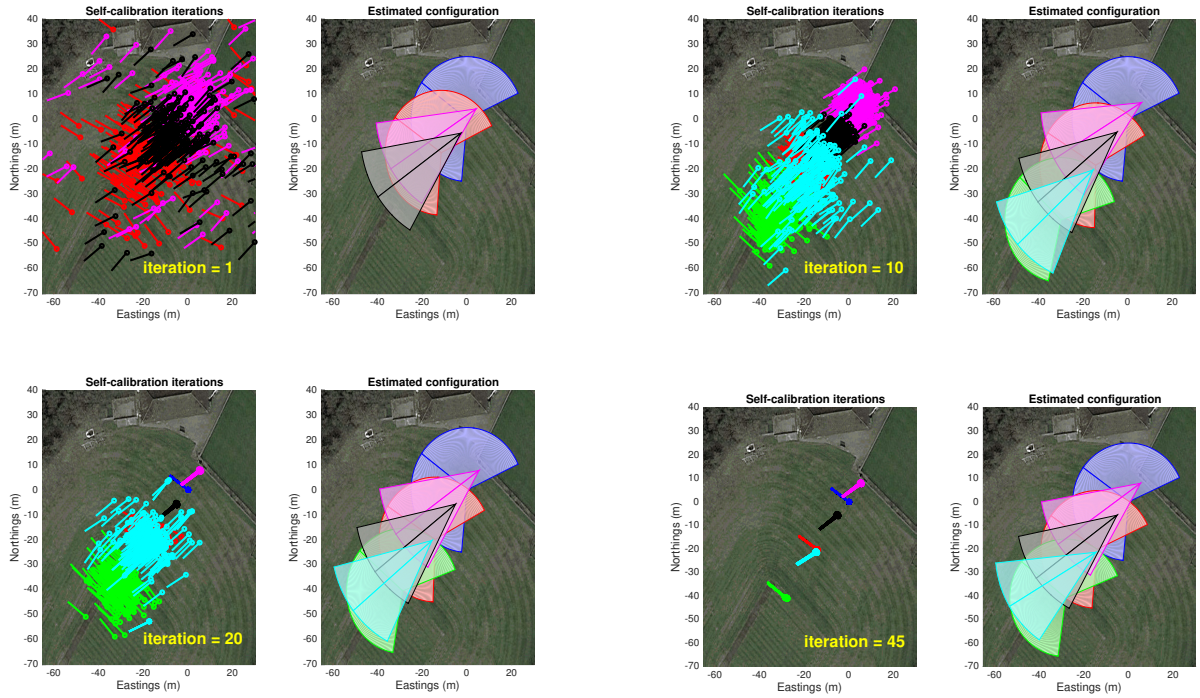

Fig. 9. Network self-calibration iterations with the SAPIENT data.

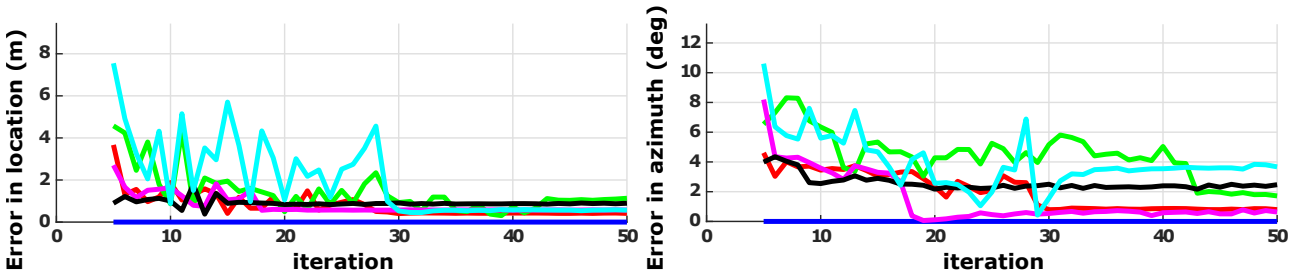

Fig. 10. Location and orientation estimation errors using the SAPIENT data (with respect to the manually measured ground truth given in Table 1).

The Markov graph used for this experiment is given in Fig. 8. Scatter plots of network configuration particles are given in Fig. 9. The speed of convergence is slower compared to the simulation example in Section 2 which is expected as the probabilistic models used to capture the sensor characteristics – for example, the noise variance in range-bearing measurements– do not necessarily match the actual characteristics. Nevertheless, the proposed framework estimates the network configuration with high accuracy in less than 50 iterations.

The estimation error with respect to the manually measured ground truth (Table 1) is depicted in Fig. 10. All sensor locations are found within ± 1.5m of the ground truth. The orientation estimation errors are within $\pm 2^\circ$ for the lidar sensors and the magenta radar sensor which is a one hop neighbour of sensor 1. The least accurate result is obtained for the cyan radar which is a three-hop neighbour of sensor 1. These error figures are calculated using manual measurements which are prone to errors, on the other hand, and should be interpreted accordingly. We elaborate this point, next.

4.3 Comparison of manual and algorithmic calibration

In this section, we provide a visual comparison of the manually measured ground truth and the algorithmically obtained results by transforming the lidar measurements to the global coordinate frame using these quantities,

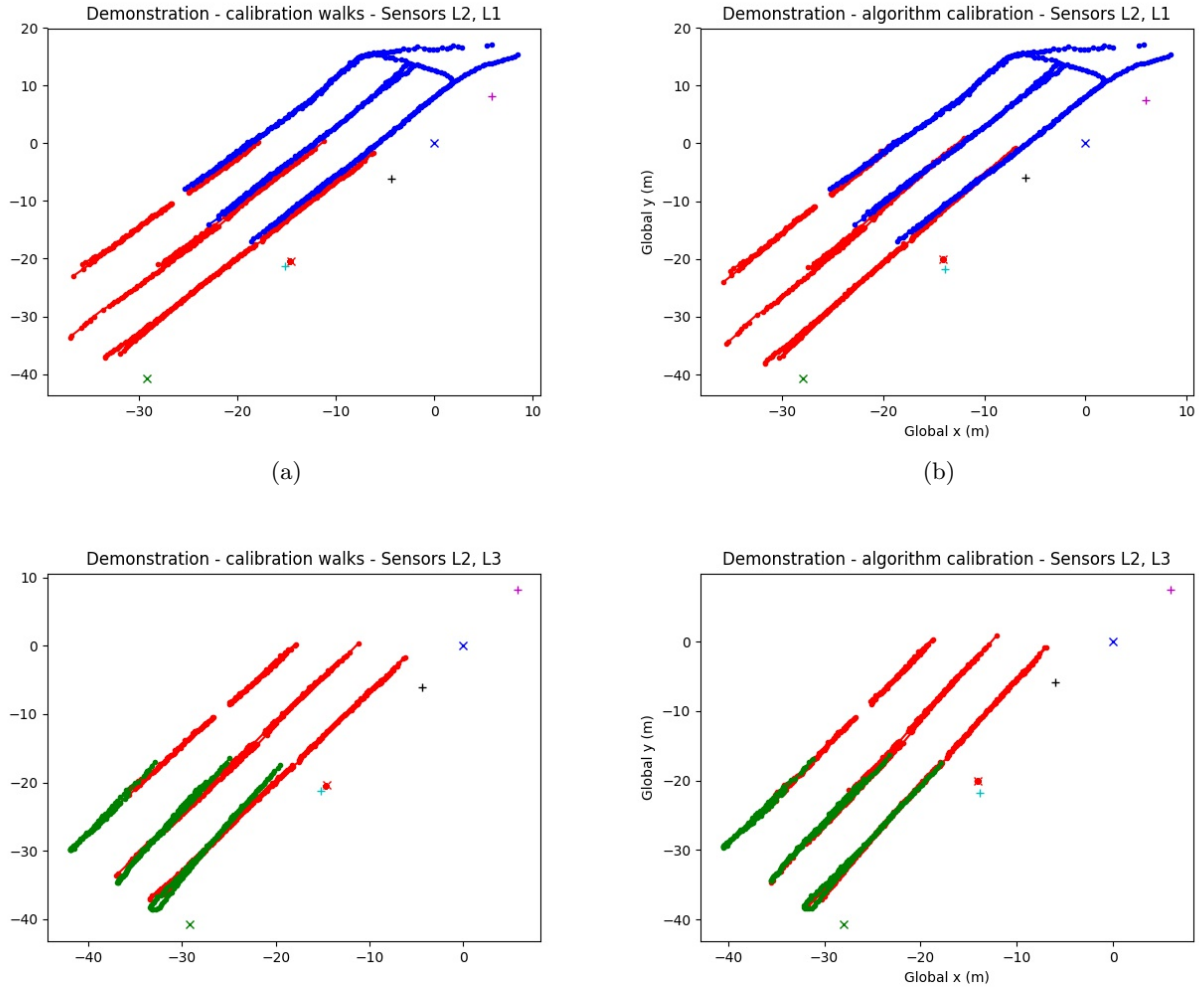

Fig. 11. Comparison of manual and algorithmic calibration via lidar data alignment: (a) Data from lidar 1 and 2: Manual calibration used to transform lidar 2 measurements. (b) Algorithmic calibration results used to transform lidar 2 measurements. (c) Data from lidar 2 and 3: Manual calibration used to transform both of the lidar measurements. (d) Algorithmic calibration results used to transform both of the lidar measurements.

and, examining the alignment of walking trajectories. As lidar range-bearing measurements are reasonable noise-free and the walking trajectories are straight, this reconstruction gives us an opportunity for this comparison.

Fig. 11 provides this comparison. The data alignments of lidar 1 and 2 for manual and algorithmic calibration are given in Fig. 11(a) and (b) respectively. The offset in the manual calibration alignment indicates that the ground truth is not very precise. The algorithmic calibration leads to a better alignment of the tracks as seen in Fig. 11(b). The comparison with the data from lidars 2 and 3 is also in favour of algorithmic calibration and can be seen in Fig. 11(c) and Fig. 11(d). These results indicate that the error figures in Fig. 10 might appear larger than the actual due to the imperfections in the manual calibration process.

5. CONCLUSION

This article has presented a scalable approach for self-calibration in fusion networks using sensor data collected from non-cooperative objects in the surveillance region. The proposed method is underpinned by separable

likelihoods which are computationally feasible approximations to the intractable likelihood of the problem. These approximations are then used to specify a Markov random field model for the calibration posterior. Parameters such as sensor locations and orientations are then estimated using particle BP. We demonstrated that this approach is capable of facilitating network self-calibration in real fusion networks and achieves better data alignment compared to manual surveying. The results motivate further research in fusion networks for exploiting the rich information content of measurements collected for situation awareness in order to also develop network self-awareness.

ACKNOWLEDGMENTS

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) Grant number EP/K014277/1 within the MOD University Defence Research Collaboration (UDRC) in Signal Processing, and, the MOD Defence and Security Accelerator grant ACC102185 "Rapid multi-sensor deployment using automatic calibration."

REFERENCES

[1] Dana, M. P., [*Registration: A prerequisite for multiple sensor tracking*], ch. 5, 155–185, Artech House (1990).

[2] Akyildiz, I. F., Pompili, D., and Melodia, T., "Underwater acoustic sensor networks: research challenges," *Ad Hoc Networks* **3**(3), 257 – 279 (2005).

[3] The Royal Academy of Engineering, "Global navigation space systems: Reliance and vulnerabilities," (2011).

[4] Erol-Kantarci, M., Mouftah, H., and Oktug, S., "A survey of architectures and localization techniques for underwater acoustic sensor networks," *Communications Surveys Tutorials, IEEE* **13**, 487–502 (Third Quarter 2011).

[5] Patwari, N., Ash, J., Kyperountas, S., Hero, A., Moses, R., and Correal, N., "Locating the nodes: cooperative localization in wireless sensor networks," *Signal Processing Magazine, IEEE* **22**(4), 54–69 (2005).

[6] Ihler, A. T., *Inference in Sensor Networks: Graphical Models and Particle Methods*, PhD thesis, Massachusetts Institute of Technology (2005).

[7] Heidemann, J., Stojanovic, M., and Zorzi, M., "Underwater sensor networks: applications, advances and challenges," *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **370**(1958), 158–175 (2012).

[8] Kantas, N., Singh, S., and Doucet, A., "Distributed maximum likelihood for simultaneous self-localization and tracking in sensor networks," *IEEE Transactions on Signal Processing* **60**(10), 5038–5047 (2012).

[9] Uney, M., Mulgrew, B., and Clark, D., "Target aided online sensor localisation in bearing only clusters," in [*Sensor Signal Processing for Defence (SSPD) 2014*], 1–5 (Sept 2014).

[10] Kantas, N., Doucet, A., Singh, S. S., Maciejowski, J., and Chopin, N., "On particle methods for parameter estimation in state-space models," *Statistical Science* **30**, 328–351 (08 2015).

[11] Lin, X., Bar-Shalom, Y., and Kirubarajan, T., "Multisensor multitarget bias estimation for general asynchronous sensors," *Aerospace and Electronic Systems, IEEE Transactions on* **41**, 899–921 (July 2005).

[12] Ristic, B., Clark, D., and Gordon, N., "Calibration of multi-target tracking algorithms using non-cooperative targets," *Selected Topics in Signal Processing, IEEE Journal of* **7**(3), 390–398 (2013).

[13] Uney, M., Mulgrew, B., and Clark, D., "A cooperative approach to sensor localisation in distributed fusion networks," *IEEE Transactions on Signal Processing* **64**, 1187–1199 (March 2016).

[14] Uney, M., Mulgrew, B., and Clark, D., "Latent parameter estimation in distributed fusion networks using," *IEEE Transactions on Signal and Information Processing Over Networks* (2018). accepted.

[15] Mahler, R. P. S., [*Statistical Multisource Multitarget Information Fusion*], Springer (2007).

[16] Wainwright, M. J. and Jordan, M. I., "Graphical models, exponential families, and variational inference," *Found. Trends Mach. Learn.* **1**, 1–305 (Jan. 2008).

[17] Thomas, P. A., Marshall, G., Faulkner, D., Kent, P., Page, S., Islip, S., Oldfield, J., Breckon, T. P., Kundegorski, M. E., Clark, D. J., and Styles, T., "Toward sensor modular autonomy for persistent land intelligence surveillance and reconnaissance (isr)," *Proc.SPIE* **9831**, 9831 – 9831 – 18 (2016).

[18] SanMiguel, J. C., Micheloni, C., Shoop, K., Foresti, G. L., and Cavallaro, A., "Self-reconfigurable smart camera networks," *Computer* **47**, 67–73 (May 2014).

[19] Jiang, L., Singh, S. S., and S.Yildirim, "Bayesian tracking and parameter learning for non-linear multiple target tracking models," *IEEE Transactions on Signal Processing* **63**, 5733–5745 (Nov 2015).

[20] Vo, B.-n., Mallick, M., Bar-shalom, Y., Coraluppi, S., Osborne, R., Mahler, R., Vo, B.-t., and Webster, J. G., [*Wiley Encyclopedia of Electrical and Electronics Engineering*], ch. Multitarget Tracking, John Wiley & Sons, Inc. (2015).

[21] Uney, M., Mulgrew, B., and Clark, D., "Cooperative sensor localisation in distributed fusion networks by exploiting non-cooperative targets," in [*IEEE Workshop on Statistical Signal Processing (SSP) 2014*], 516–519 (June 2014).

[22] Uney, M., Mulgrew, B., and Clark, D., "Distributed localisation of sensors with partially overlapping field-of-views in fusion networks," in [*2016 19th International Conference on Information Fusion (FUSION)*], 1340–1347 (July 2016).

[23] Uney, M., Mulgrew, B., and Clark, D., "Distributed estimation of latent parameters in state space models using separable likelihoods," in [*Proceedings of the ICASSP 2016*], (2016).

[24] Wainwright, M., Jaakkola, T., and Willsky, A., "Tree-based reparameterization framework for analysis of sum-product and related algorithms," *Information Theory, IEEE Transactions on* **49**, 1120–1146 (May 2003).

[25] Yedidia, J. S., Freeman, W., and Weiss, Y., "Constructing free-energy approximations and generalized belief propagation algorithms," *Information Theory, IEEE Transactions on* **51**, 2282–2312 (July 2005).

[26] Murphy, K. P., [*Machine Learning: A Probabilistic Perspective*], The MIT Press (2012).

[27] Sudderth, E. B., Ihler, A. T., Isard, M., Freeman, W. T., and Willsky, A. S., "Nonparametric belief propagation," *Commun. ACM* **53**, 95–103 (Oct. 2010).

[28] Vo, B. N. and Ma, W. K., "The gaussian mixture probability hypothesis density filter," *IEEE Transactions on Signal Processing* **54**, 4091–4104 (Nov 2006).

[29] Page, S. F., Oldfield, J. P., Islip, S., Benfold, B., Brandon, R., Thomas, P. A., and Stubbins, D. J., "Threat assessment and sensor management in a modular architecture," in [*Proc. SPIE*], *Proc.SPIE* **9995** (2016).