Context-enhanced maritime surveillance optimization

Steven Horn

Defence Research and Development Canada (DRDC)

ABSTRACT

A method for the optimization of wide area surveillance given contextual information in the form of intelligence about existing or potential maritime threats is presented. Given the complexity of the surveillance parameter space, exhaustive search is typically not possible. A Bayesian search using posterior sampling is proposed whereby a near optimal surveillance plan can be generated with reduced computational requirements. Given any new intelligence over time, the surveillance plan can then be quickly updated to integrate the new information. The method also allows the incorporation of past surveillance activities including nil observations.

Keywords: Surveillance, Pattern of Life, Modelling and Simulation, Bayesian Inference

1. INTRODUCTION

The problem of surveillance resource allocation and optimization given constraints and objectives is a widely studied area of research.[1,2] The goal of surveillance resource optimization is to maximize the available information and to minimize the operational risk. Typically, past research has focused on the management of a single sensor,[3] or a set of sensors[4] for performance optimization of the sensor or sensors. Additionally, attention has been given to the net objective of surveillance in relation to the reduction of operational risk.[5] Furthermore, the optimization process can be dynamic with repeated optimizations providing adjustments on short *tactical* minute-to-minute time-scales such as when a sensor is employed *in-situ*, or the optimization can be on a longer *operational* time-scale where the adjustments are days in advance.

The task of optimization of a wide-area surveillance plan which aims to conduct *operational* time-scale planning is considered in this paper. In contrast to some existing work, the optimization discussed here is focused on the singular detection (*i.e.* detected at least once) of one or more potential maritime threats and not necessarily the optimization of the tracking of a threat. Figure 1 presents an illustration of the operational wide area surveillance planning problem.

This paper is organized as follows: Section 2 presents the mathematical framework for describing the problem, Section 3 presents the proposed optimization technique, and Section 4 presents some simulated results for a notional scenario.

2. MODELLING OF SURVEILLANCE, THREATS, AND CONTEXT

The operational wide area surveillance problem is described mathematically as follows. A surveillance planner has a set of n surveillance resources $\mathcal{R} = \{R_i\}$, $i = 1 \ldots n$, which are each subject to a set of k constraints $\mathcal{C}_i = \{C_j\}$, $j = 1 \ldots k$. A surveillance plan is generated for the employment of a subset of $\mathcal{R}$, and each configured with parameters $\mathcal{P}_i$ which include the locations and times of surveillance subject to $\mathcal{C}_i$. The choice of which subset of $\mathcal{R}$ is employed in the surveillance plan is such that the total cost of surveillance should be minimized while maximizing the effectiveness of the plan for detecting threats.

To model the multiple types of concurrent threats for which to conduct surveillance against, we define a set of threats $\mathcal{T} = \{T_x\}$ each with a set of characteristics with uncertain values. The following characteristics are considered here for each threat: the origin (T_o), destination (T_d), speed (T_s), and size (T_z) of a threat such that the characteristics of T_x are $\{T_{x,o}, T_{x,d}, T_{x,s}, T_{x,z}\}$.

Each surveillance resource R_i is described by a sensor detection model d_i which is a function of the ability to detect a threat T_x with given characteristics such that $d_i = f(R_i, C_i, T_{x,z})$.[6,7] The simplest approach is if we

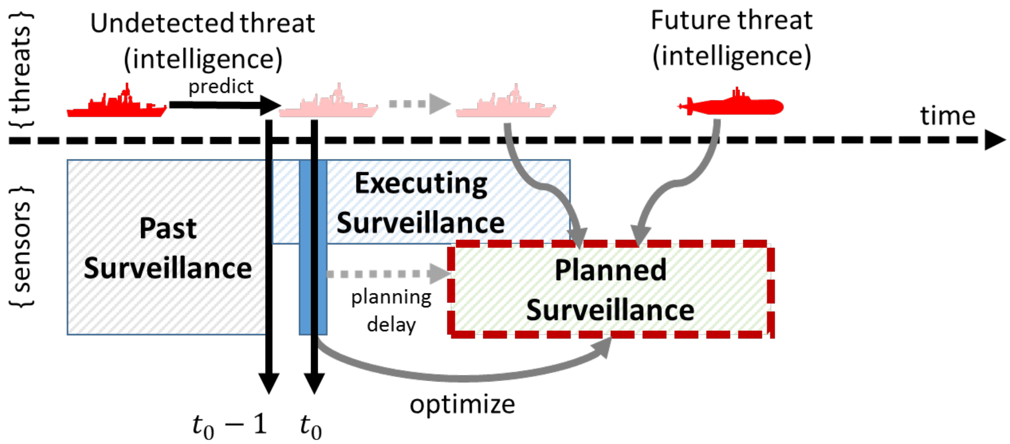

Figure 1. Illustration of the operational wide area surveillance planning problem. Surveillance optimization leverages past surveillance information, current information, and intelligence to plan future sensor allocations.

assume that each sensor is independent of one another. However, this assumption is not always valid in reality: for example, if one sensor is cross-cuing another sensor. Even so, it is not a bad assumption for many surveillance scenarios and sensors.[8] Consideration for sensor dependence can be included in a more advanced model but will not be discussed in this paper.

Elaborating on surveillance shown in Figure 1, there are three aspects to the surveillance model: 1) information about surveillance which has already occurred; 2) surveillance which is committed and will execute; and, 3) surveillance which is planned but not yet committed. The past surveillance is defined as the set of surveillance resources which were executed and the resulting combined information generated from that surveillance including nil detections. This is the *a priori* state of detection of threats. Executing surveillance is defined as the set of surveillance resources which are providing information at present. For example, this could consist of real-time radar information, satellite collections, or patrols (air and surface).

The route of the threat transiting the surveillance region can be contextually modelled using intelligence derived from the pattern of life[9] given T_o and T_d. Through simulation of the surveillance plan, the posterior probability distribution is estimated for the detection of threats ($f_{d|R,T}$) after the threats transit through the surveillance zone given the resources $\mathcal{R}$ and threat characteristics.

Figure 2 illustrates how the use of multiple independent sensors improves the effect of surveillance. For illustrative purposes, $\mathcal{R}$ consists of three fake sensors with parameters shown in Table 1. These were simulated against a threat transiting from $x_0 = 0 \pm 2$ to $x = 1000$ nautical miles with speed $v = U(10, 15)$ knots, starting at $t_0 = U(0, 100)$ hours, where U is the uniform distribution. The distribution is derived from the simulation using kernel density estimation with a Gaussian kernel with a bandwidth derived from Scott's rule.[10] The score for surveillance using one type of sensor independently shown in Figure 2 (a-c) versus when all of the sensors are used (d) illustrates the coordination effect of surveillance. Equation 1 shows how this score is calculated for threat x. In general, for all threats, the overall surveillance score is $s = \sum_x s_x$. Depending on when and where each sensor is in relation to the threats, the overall surveillance score will vary as the surveillance parameters are varied. This score function is a form of linearly weighted sum of the distribution which increases for improved expected detection probability.

$$s_x = \int_0^1 p \cdot f_{d(p)|R,T_x} \, dp \, . \tag{1}$$

For each sensor parameter free from constraints, the objective is to then find a surveillance plan which results in a maximum value of s given the surveillance parameters. The next section presents a method to search the parameter space to optimize the surveillance plan.

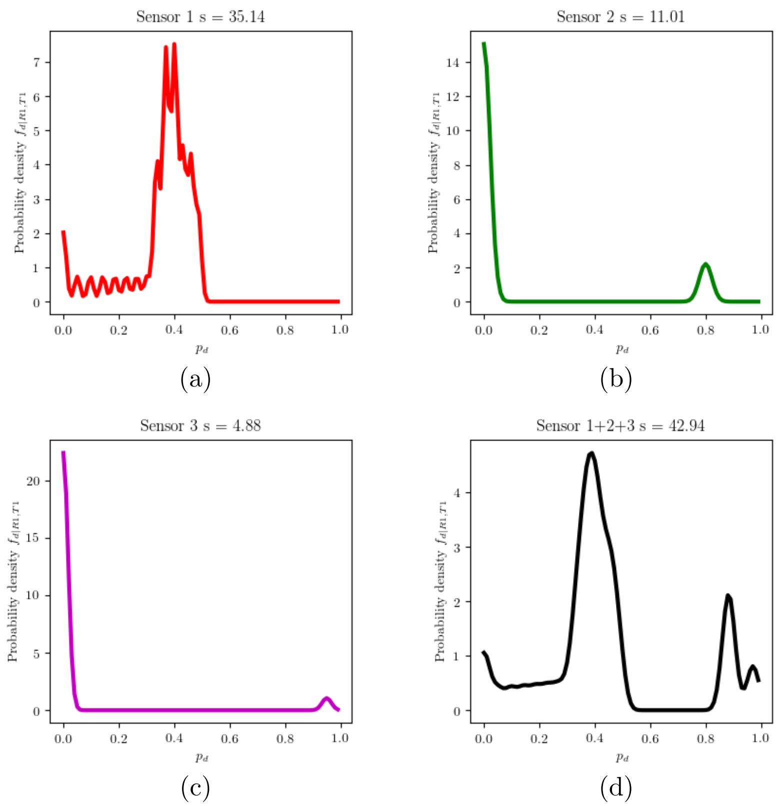

Figure 2. For a given threat T_1, the combined detection performance for three sensors each individually (a-c), and all sensors combined (d). The surveillance score is indicated by s.

Table 1. Surveillance parameters for three simulated sensors.

Sensor	**Description**	**Parameters**
1	Surveys a large area relatively often with low probability of detection. This is representative of a constellation of satellites.	$P_d = 0.04$, $dt = 2$ h, $x = [50, 300]$
2	Surveys a small area (50 nm wide) every 24 hours. This is representative of a long range patrol aircraft.	$P_d = 0.8$, $dt = 2$ h, $dx = 50$
3	Surveys a specific box every 20 hours. This is representative of a surface patrol.	$P_d = 0.95$, $dt = 2$ h, $x = [900, 1000]$

3. OPTIMIZATION THROUGH POSTERIOR SAMPLING

The objective is to select and configure the resources $\mathcal{R}$ by varying the surveillance parameters P_n subject to constraints $\mathcal{C}_R$ to maximize the detection of the set of threats, as represented by the surveillance score s. Given a surveillance configuration with a set of parameters, one can calculate through simulation the probability distribution for the detection of a threat $f_{d|R1,T1}$, as shown in Figure 2. The brute force approach would be to simulate and compute the posterior distribution of all possible surveillance options. However, the search space for possible surveillance configurations, even under constraints, is extremely large. Moreover, changes to the configuration are highly non-linear and finding an optimal (or better) solution is not straightforward. Depending on the complexity of the surveillance optimization, there can be numerous parameters (hundreds), and evaluating the combinations of each parameter through simulation is not tractable.

Proposed here is a Bayesian approach using posterior sampling[11] with Gibbs method[12] to *learn* and predict with uncertainty the effect of surveillance in order to explore more informative simulations. For each surveillance option parameters which satisfy the constraints, a distribution of the surveillance score s is defined conditioned on those parameters. A scaled score over the search space can be represented by a multivariate distribution of the surveillance parameters under the test that the score is maximized given the parameters

$$f(s = s_{max}|P_1, \ldots, P_n), \tag{2}$$

which is initialized as per standard Bayesian methods. Here, one can choose an uninformed (uniform) prior for each parameter P_n to be optimized. The objective of the optimization is to find the tuple of P which maximizes s. This optimization is equivalent to identifying the maximum after inference of the distribution in Equation 2

Step 0: Initialization

The distribution in Eq. 2 is initialized with an uninformative prior such that

$$\int_{P_1} \cdots \int_{P_n} f(s_{max}|P_1, \ldots, P_n)\ dP = 1. \tag{3}$$

Step 1: Sampling

The selection of the tuple of surveillance parameters P to simulate is made using collapsed Gibbs sampling.[13] The objective in making this choice is to evaluate the simulation with parameters which is expected to provide useful information for the inference of $f(s_{max})$, while still exploring the parameter space which has yet to be evaluated. To select a value for a parameter P_i, the search space is collapsed through integration respecting the constraints C.

$$f(s_{max}|P_i) = \int_{P|C} f(s_{max}|P_{1\ldots n})\ dp_{-i} \tag{4}$$

Where $f(s_{max}|P_i)$ is the collapsed marginal distribution, and the differential dp_{-i} represents integration over all parameters except i. The calculation of Eq. 4 is done numerically. Since the distribution of $f(s_{max})$ is normalized as in Eq. 3, and the values are all positive, then the maximum of the cumulative distribution is 1. The choice of value for P_i is then derived from a random number $0 < r < 1$, and the parameter is derived from the cumulative distribution:

$$r = \int_{min(P_i)}^{P_i} f(s_{max}|P_i = x)\ dx. \tag{5}$$

The calculations in Eqs. 4-5 are repeated for each parameter dimension to choose the tuple of surveillance parameters R_p.

Step 2: Simulation

The surveillance score is calculated as per Eq. 1 for the sampled R_p, resulting in a calculated score $\hat{s}$. Recall that this surveillance score now considers the uncertainties in the threat, and the plan with R_p is simulated against the threat intelligence (context).

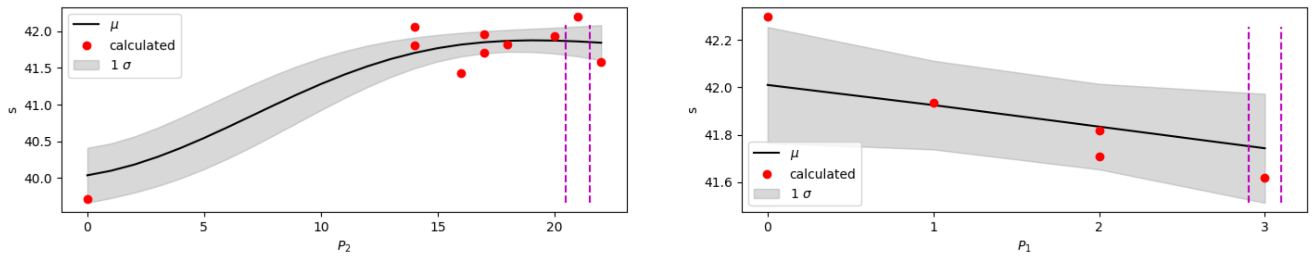

(a) (b)

Figure 3. Visualization of the local Gaussian Process Regression for the estimation of the likelihoods for two parameters. The latest iteration is shown between the vertical dashed lines.

Step 3: Calculate posterior

The joint posterior distribution of the form in Eq. 2 is estimated using Bayes method.

$$f(\hat{s} = s_{max}|P) = \frac{\prod_{i \in R_p} p(\hat{s}|P_i, R_{p-i})}{p(\boldsymbol{s})} \cdot f(s_{max}|P) \tag{6}$$

The fraction in Eq. 6 is the Bayes factor where the top term is the joint likelihood, and the bottom term is the marginal likelihood. It is difficult to calculate the exact value of these terms, and so here it will be approximately calculated using numerical non-parametric methods. For the parameters which were simulated in the tuple R_p, the likelihood of the simulation result $\hat{s}$, given a parameter $P_i = q$, is $p(\hat{s}|P_i, R_{p-i})$. This likelihood is estimated using local regression, which has been shown to be effective[14] when sampling from a computationally expansive process such as the simulations in the problem here.

The local regression method used here is a Gaussian Process Regression on each parameter. The benefit of this particular regression is that it itself is a Bayesian estimator, and the posterior of it's estimate provides a distribution with a variance, which when applied here models the level of ignorance from lacking evidence from yet to be computed simulations. Given the regression on parameter P_i, the likelihood is:

$$p(\hat{s}|P_i, R_{p-i}) = 1 - \int_0^{\hat{s}} N(s|\mu_i(q), \sigma_i(q)^2)\ ds, \tag{7}$$

where $\mu_i(q)$ and $\sigma_i(q)$ are the regression outputs for the Gaussian as a function of the parameter value. Figure 3 shows an example of a local regression for the optimization problem which will be presented in the next section. A radial-basis function kernel was used with a length parameter of 2 for the regression.

Given the set of simulation results $\{s\}$ for the simulations, the marginal likelihood $p(\boldsymbol{s})$ in Eq. 6 is estimated using kernel density estimation (KDE).[15] A standard Gaussian KDE was used with a bandwidth derived using Scott's method.[10] This integration is illustrated graphically in Figure 4. Finally, since the estimation methods used for the likelihoods are not-exact, the posterior is re-normalized after applying the Bayesian update.

Step 4: Iterate

Steps 1-3 are repeated until a stopping condition is reached. This stopping condition can be a function of the update such as when the information gain from each simulation is no longer significant. Once the information gained from additional simulations for the update of the multi-dimensional surface $f(s|P)$ is low (*e.g.* below a threshold), sampling can be stopped, and the best discovered parameters R_p are taken as the best available surveillance plan.

Figure 5 shows the Kullback-Leibler (KL) divergence mean and variance over multiple iterations for 10 runs of the same optimization problem. The quickest optimizations in this plot were completed after just 8 iterations. One can note that as the algorithm converges on the optimal parameter set, the KL divergence tends to zero.

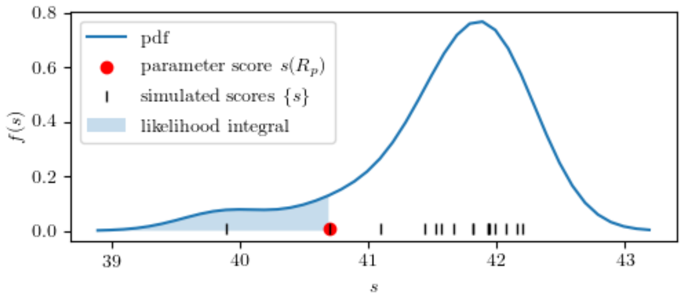

Figure 4. Example of the estimation of the likelihood that a given score for parameters is an optimum given a set of previous simulations.

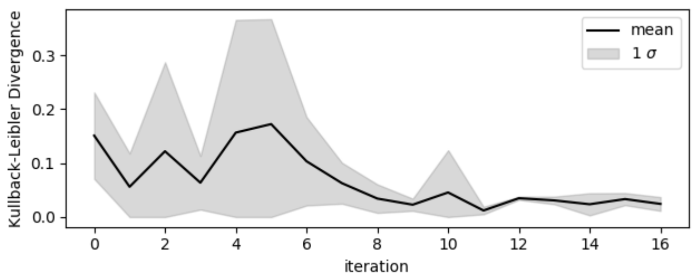

Figure 5. The Kullback-Leibler divergence of the posterior to the prior on updates from multiple simulation steps.

The steps with the higher divergence indicate locations where the information from new simulation runs was higher.

It must be noted that this algorithm does not guarantee a globally optimum result, but since the computational cost of evaluating each surveillance option is high, a good solution arrived by exploring as few as possible parameter variations.

4. RESULTS

To demonstrate the posterior sampling approach, the scenario shown in Section 2 was used while the parameters for sensor 2 were varied under arbitrary constraints for demonstration. These are not necessarily realistic constraints, but they serve to create a parameter space to demonstrate searching. The constraints used are:

1. $dx = 50$ (the surveillance box is 50 nm wide)
2. $x_0 = [0, 50, 100, 150]$ (the surveillance is constrained to one of these 4 boxes)

In Figure 6a, the posterior distribution is visualized for progressive solution search steps. One can note that the peak in the likelihood $f(s_max|P)$ moved and narrowed during the search procedure, as it arrived at a solution. The best solution found was with $P_1 = 1$ and $P_2 = 23$. Figure 6b shows the distribution function of this distribution for threat detection before and after optimization of sensor 2.

5. CONCLUSIONS

A method to quickly search a large parameter space for surveillance planning is proposed and was demonstrated on a simulated simple scenario. These types of searches are computationally costly, but the method here may make the optimization tractable. Future work includes exploring different Markov Chain Monte Carlo sampling techniques for sampling the prior, and applying the method to realistic high fidelity models of sensors and platforms.

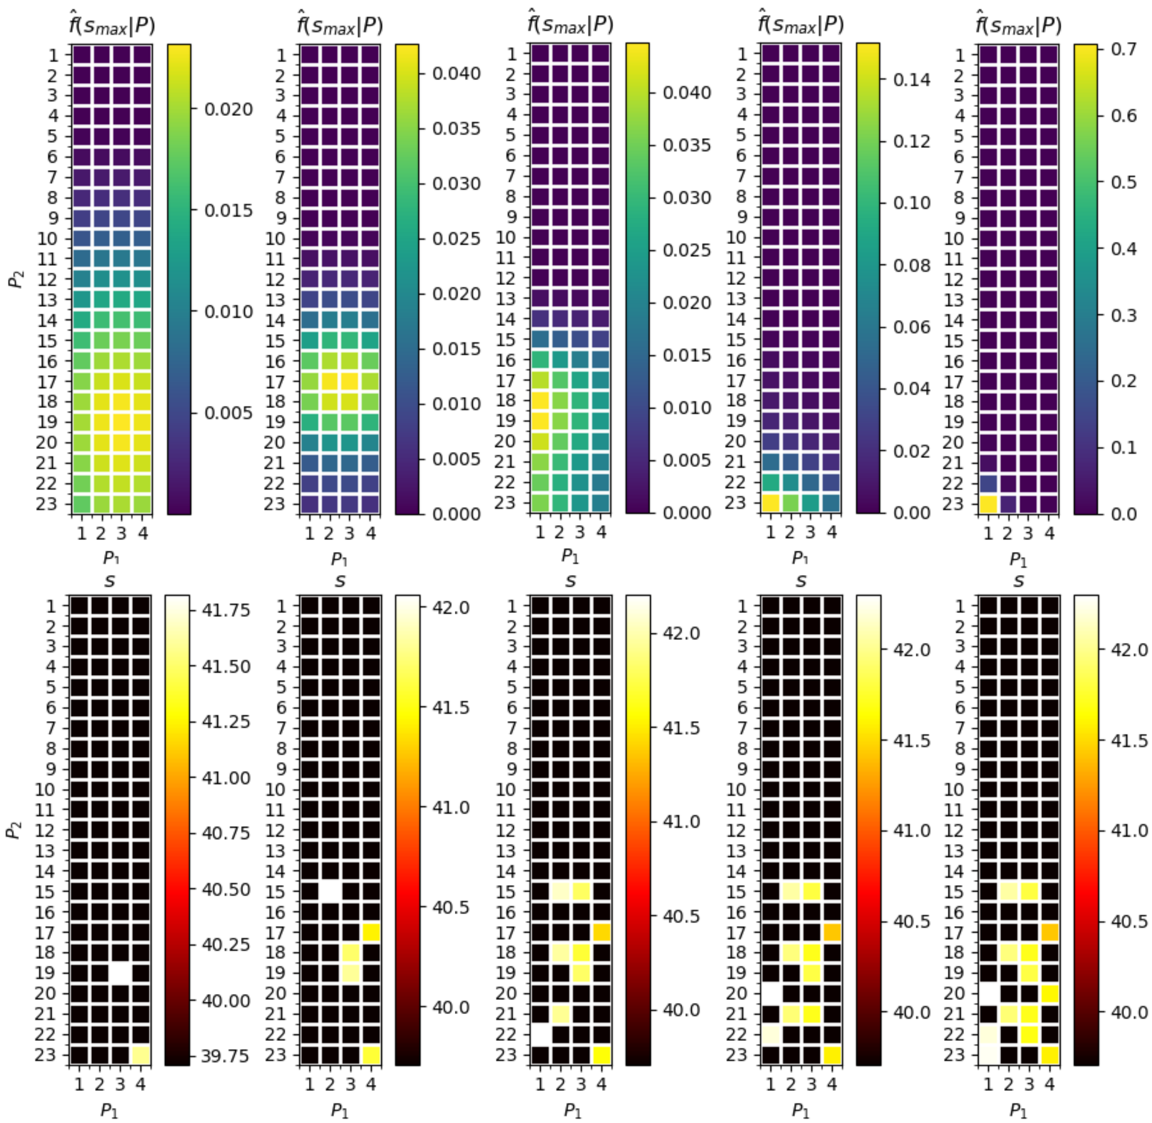

(a)

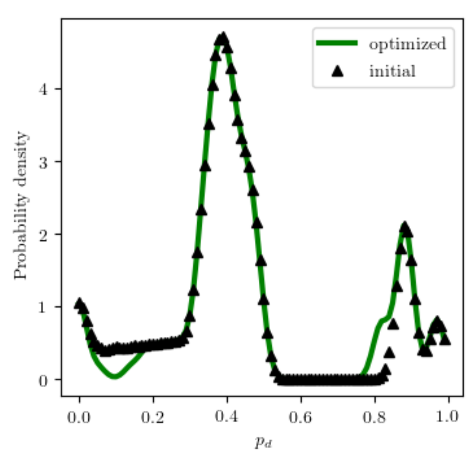

(b)

Figure 6. Progression of the posterior distribution during search of surveillance parameters (a), and distribution of detection probabilities for the optimized surveillance plan (b).

REFERENCES

[1] Dridi, O., Krichen, S., and Guitouni, A., "A multi-objective optimization approach for resource assignment and task scheduling problem: Application to maritime domain awareness," in [*2012 IEEE Congress on Evolutionary Computation*], 1–8 (June 2012).

[2] Berry, P., Pontecorvo, C., and Fogg, D., "Optimal employment of space surveillance resources for maritime target tracking and re-acquisition," in [*International Conference on Information Fusion*], 719–725 (2003).

[3] Hero, A. O. and Cochran, D., "Sensor management: Past, present, and future," *IEEE Sensors Journal* **11**, 3064–3075 (Dec 2011).

[4] Ponsford, A., "Multi-sensor resource management system for navy applications," (2019). DRDC-RDDC-2018-C240.

[5] Papageorgiou, D. and Raykin, M., "A risk-based approach to sensor resource management," in [*Advances in Cooperative Control and Optimization*], 129–144, Springer (2007).

[6] Horn, S., "Near real time estimation of surveillance gaps," in [*Proceedings of the 16th International Conference on Information Fusion*], 1871–1877 (July 2013).

[7] Horn, S. A., Mason, D., Zegers, A., and Edmond, E. J., "A bayesian approach to estimating detection performance in a multi-sensor environment," Tech. Rep. DRDC-RDDC-2014-R56, Defence Research and Development Canada CORA (2014).

[8] Papa, G., Horn, S., Braca, P., Bryan, K., and Romano, G., "Estimating sensor performance and target population size with multiple sensors," in [*2012 15th International Conference on Information Fusion*], 2102–2109 (July 2012).

[9] Millefiori, L. M., Braca, P., Bryan, K., and Willett, P., "Modeling vessel kinematics using a stochastic mean-reverting process for long-term prediction," *IEEE Transactions on Aerospace and Electronic Systems* **52**, 2313–2330 (October 2016).

[10] Scott, D. W., [*Multivariate density estimation: theory, practice, and visualization*], John Wiley & Sons (2015).

[11] Russo, D. J., Van Roy, B., Kazerouni, A., Osband, I., Wen, Z., et al., "A tutorial on thompson sampling," *Foundations and Trends® in Machine Learning* **11**(1), 1–96 (2018).

[12] Casella, G. and George, E. I., "Explaining the gibbs sampler," *The American Statistician* **46**(3), 167–174 (1992).

[13] Griffiths, T. L. and Steyvers, M., "Finding scientific topics," *Proceedings of the National academy of Sciences* **101**(suppl 1), 5228–5235 (2004).

[14] Conrad, P. R., Marzouk, Y. M., Pillai, N. S., and Smith, A., "Accelerating asymptotically exact mcmc for computationally intensive models via local approximations," *Journal of the American Statistical Association* **111**(516), 1591–1607 (2016).

[15] Thijssen, B. and Wessels, L. F., "Approximating multivariate posterior distribution functions from monte carlo samples for sequential bayesian inference," *arXiv preprint arXiv:1712.04200* (2017).