Reactive Data Gathering for Underwater Mine Countermeasures†

By Jordi Barr‡, Mark Williams, David Nicholson
BAE Systems Advanced Technology Centre

Abstract

We demonstrate how mathematical techniques for Reactive Data Gathering (RDG) can be used to increase the probability of correctly classifying seabed threats. The problem is one of separating underwater objects into mine or non-minelike bottom object (NoMBO) classes using noisy and incomplete sequences of Sonar data. The incorporation of new individual sensor observations is undertaken at two levels. We operate on sensor output directly using Sequential Importance Sampling (otherwise known as the Particle Filter). We also undertake inference at the decision level using confusion matrix output from an Automatic Target Recognition (ATR) algorithm. Predicted observations are used as the basis for our RDG scheme, which seeks to maximise the one-step-ahead information gain. RDG is evaluated against current operating procedure and standard benchmarks, e.g. rosette patterns or random views of the targets.

We use a Sonar simulator developed by the Ocean Systems Laboratory at Heriot-Watt University and, where appropriate, a BAE Systems-developed ATR to conduct a series of experiments to test the benefits of RDG in realistic mine countermeasure (MCM) scenarios. Various types of mines and NoMBOs are modelled. The nature of our ATR is likely to be sensor (and in our case simulator) specific and so we ensure that our RDG techniques are ATR-neutral. The results demonstrate that using RDG algorithms increases the probability of correct classification within a specified number of observations. Alternatively, RDG reduces the number of observations required to attain a target classification accuracy. We speculate that RDG could deliver a threefold increase in operational tempo for MCM operations.

1. Introduction

Autonomous Underwater Vehicles (AUVs) are used increasingly in mine countermeasure (MCM) operations as a method by which the man can be removed from the minefield. AUVs are typically equipped with forward- or side-looking active Sonar sensors, which enable them to construct (sound) images of their environment as they progress through it. Currently, these images are downloaded at the end of a complete mission and processed by an operator who identifies objects by eye. In most cases, the nature of targets is ambiguous and follow-up missions are necessary to ensure the required probability of a correct classification. Our work addresses the need for a shorter duty cycle between the observation of the environment and the identification of the objects within it. There is a stated desire to bring these activities 'in-step' - that is to classify objects while the AUV is in the water. Reactive Data Gathering (RDG) algorithms build toward this goal by allowing an AUV to decide on the best route

† This work was funded under MoD Competition of Ideas project contract no. RT/COM/5/013
‡ jordi.barr@baesystems.com

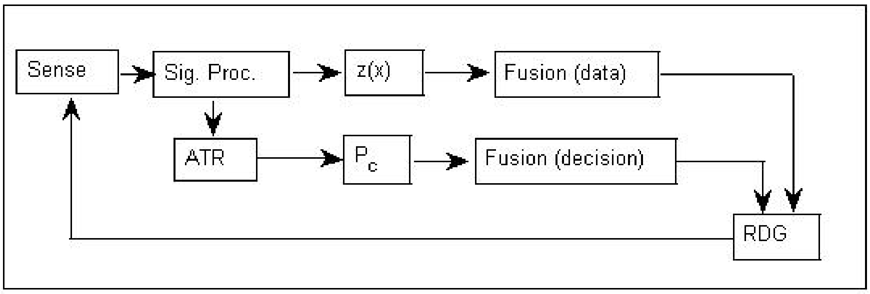

FIGURE 1. Two routes from sensing to RDG. The uppermost path conducts inference following signal processing to extract an emission, z, indicative of the underlying state, x, of an object. In this work we assume that x is a [type, aspect] pair and that z is the level of Sonar return (in dB). The alternative route in the lower path assumes that an ATR is available and delivers P_c, the probability that the target corresponds to a model type. In this case the inference relies on a statistical characterisation of the ATR through a (learned) confusion matrix. The RDG element operates as described in § 2.4.

to classify seabed objects. This work builds on previously developed RDG methods suited to complex target classification problems in various domains [1,2]. We detail the method from simulation to RDG in §2. An experimental plan and results are the subject of §3 with a summary and brief details of ongoing and future work presented in §4.

2. Method

This section details the process by which data are used by the RDG algorithm to generate classification outputs. Figure 1 shows our abstract scheme for the MCM AUV decision cycle. Note that when we speak of data fusion we mean the act of fusing new sensor data, or a new probability distribution over states, with our current state estimate. This is more properly called inference, though we will use the terms 'data fusion' (DF) and' inference' interchangeably. We undertake inference at the data level using Sequential Importance Sampling (otherwise known as the Particle Filter, PF). An alternative route is to employ an Automatic Target Recognition (ATR) algorithm on a single datum - thereby generating a probability distribution over the target set and fusing this with the currently held probability distribution at the decision level. RDG is then accomplished for either method by maximising the expected entropy reduction over a predicted set of future observations.

2.1. *Simulation*

The project relies on a simulation tool supplied by Heriot-Watt University Ocean Systems Lab. The simulation in question embodies a high degree of sophistication, having been produced by a team with long experience in this area [3]. It is currently impractical to test RDG methods with real data as the RDG algorithm would have to direct the real-time data gathering of an AUV, or have access to a vast database of real imagery so that its measurement choices could be serviced. Two flavours of simulator are used, the simpler generates an 'A-scan' solving an energetic equation over a modelled beam path. Spatial correlation cannot be achieved and geometrical phenomena such as sand ripples are not modelled. Only direct paths are considered. PF methods are employed directly on the output of this simulation. The second simulator generates the full sidescan images which are typical of MCM operations. This requires that an ATR is run on the image and consequently the fusion is done at the decision level. In all cases, a target and background type are modelled together with environmental parameters such as the depth, sea state, temperature, salinity, and pH.

2.2. *ATR*

The purpose of this study is not to develop an ATR, rather to investigate whether RDG techniques have potential for improving the performance in MCM. To that end the ATR we use has only to be representative of other more sophisticated ATR algorithms developed for MCM. As no ATR is supplied with our simulator, we developed one which relies on a set of templates of the expected shape of the shadow cast by each target of interest at a comprehensive set of evenly spaced observation angles. These templates are correlated with a thresholded version of the side-scan image and the adjusted scores are used to decide the probability of target class and orientation.

2.3. *Data Fusion: Sequential Importance Sampling (The Particle Filter)*

The goal here is to infer the underlying (hidden) states $x_{0:t}$ based on a sequence of observations $z_{1:t}$, i.e. $P(\mathbf{x}|\mathbf{z})$. In our analysis $x_i \equiv [T, \theta_i]$ where T is the model type and θ is its orientation with respect to the world. For the A-scan simulations we use a PF method. Data fusion at the decision level is trivial, requiring no more than a multiplication and normalisation of confusion matrices.

Previous classification work carried out by BAE Systems ATC under the MOD-funded RCSC18b project [4] fused data using Hidden Markov Models (HMMs). There is a considerable body of academic work devoted to multi-aspect classification using HMMs [5,6]. However, the necessary assumption of our targets as piecewise stationary with respect to observing angle breaks down for the simulated targets. For this reason we use a method which is less sensitive to the piecewise stationary assumption. Particle filters work by tracking samples (particles) and assigning weights to these particles according to their probability of corresponding with observations. By splitting particles with large weights and removing those with low weights, particles converge to high-probability solutions. See [6] for background on PFs.

We follow the method of sequential importance sampling as described by [7]. Suppose we have a set of support points, $\mathbf{x}_{0:t} = \{x_j, j = 0, ..., t\}$, producing observations, $\mathbf{z}_{1:t}$, with weights $\mathbf{w}_{1:t}$. We seek to approximate the posterior probability density,

$$p(\mathbf{x}_{0:t}|\mathbf{z}_{1:t}) \approx \sum_{i=1}^{N} w_t^i \delta(\mathbf{x}_{0:t} - \mathbf{x}_{0:t}^i), \tag{2.1}$$

for N particles. The recipe in [7] gives the recursive update of this estimate. The 'vanilla' version of the PF is a recursive propagation of weights and support points as each measurement is received. The canonical problem, and one we would encounter here, is that of degeneracy. After a small number of iterations, all the particles save one have negligible weight. In order to avoid this problem we resample our particles at each timestep. The idea is to discard particles with small weights, and split those with larger weights into two or more points in about the same location. The new points are sampled at equal intervals from the posterior cumulative density function. The new weights are reset to $1/N$.

There are two general observations about the particle filter which may inform our investigation. Firstly, unlike the HMM, there is no assumption regarding state transitions. This is both an advantage and a disadvantage depending on the nature of the problem. It may take more timesteps to converge to a solution. However, tailoring model parameters, as is needed in the HMM, is not necessary. Secondly, convergence depends on the choice of support points; if too few are chosen then the solution will take longer to converge and its accuracy will be diminished. This is mitigated by using more support points (in the limit $N \to \infty$ the approximation in Equation 2.1 becomes an equality). Obviously the key is to sample the appropriate number of particles to maintain an acceptable computing time.

2.4. *Reactive Data Gathering*

Reactive Data Gathering seeks to reach some level of confidence in detection and/or identification in as few measurements as possible. This is accomplished by estimating the expected benefit of future observations. The objective of our RDG method is to predict where next to view the target to maximise expected information gain in the posterior probability distribution over the target class. In practice, a number of possible moves are considered and for each move the relative entropy is calculated. The move that minimises the relative entropy is executed by the vehicle. To do this we calculate current entropy as,

$$H(T|\Delta\theta_{1:t}, \mathbf{z}_{1:t}) = -\sum_{l=1}^{L} P(T_l|\Delta\theta_{1:t}, \mathbf{z}_{1:t}) \log P(T_l|\Delta\theta_{1:t}, \mathbf{z}_{1:t}), \tag{2.2}$$

for the $\Delta\theta_{1:t}$ set of moves undertaken so far. The expected entropy is,

$$g_{t+1}(\Delta\theta_{t+1}) = E_{z_{t+1}|\boldsymbol{\Delta}\theta_{1:t+1}, \mathbf{z}_{1:t}} H(T|\Delta\theta_{1:t+1}, \mathbf{z}_{1:t+1}) \tag{2.3}$$

$$= \textstyle\sum_{l=1}^{L} H(T_l|\Delta\theta_{1:t+1}, \mathbf{z}_{1:t+1}) \times \frac{\sum_{l=1}^{L} P(\mathbf{z}_{1:t+1}|\Delta\theta_{1:t+1}, T_l)}{\sum_{l=1}^{L} P(\mathbf{z}_{1:t}|\Delta\theta_{1:t}, T_l)}. \tag{2.4}$$

The optimum angular displacement maximizes the reduction in the expected entropy,

$$\Delta\theta_{t+1} = \arg\max_{\Delta\theta}(H(T|\Delta\theta_{1:t}, \mathbf{z}_{1:t}) - g_{t+1}). \tag{2.5}$$

So given a sequence of angular displacements, observations, and the expected return of z_{t+1} we can select the $\Delta\theta_{t+1}$ that maximizes the reduction in the expected entropy. In practice this is achieved by sampling uniformly at a number of angles for each potential target.

3. Application to MCM

This section sets out a number of experiments that are run to test the RDG algorithm against benchmarks. We test the RDG methods for both simulated A-scan and sidescan returns using a number of targets against a variety of backgrounds.

3.1. *Targets and Backgrounds*

We include 7 targets in the experiments. These constitute just about all the targets we are able to model with the simulator. The targets are 3 mines: a circularly-symmetric Manta; an irregularly-shaped Rockan; and a round-ended cylindrical ‘encapsulated torpedo’ type. The non-minelike bottom object (NoMBOs) set comprises a vertical and horizontal cylindrical oil drum and a sphere. The absence of a target is also included as a ‘target’ to simulate the effect of a false detection. More information on the various mine types is publicly accessible through various online resources (e.g. [8]). Three background seabeds are modelled: fine sand, sandy gravel and rough rock. In the sidescan simulation we also include sand ripples at various angles. Other parameters of interest which are varied are sea state (0-9), seabed roughness (0-1) and a noise parameter. Salinity is fixed at 35ppt, Temperature at 281K and pH at 8. The variables have the effect of increasing the (Rayleigh distributed) noise relative

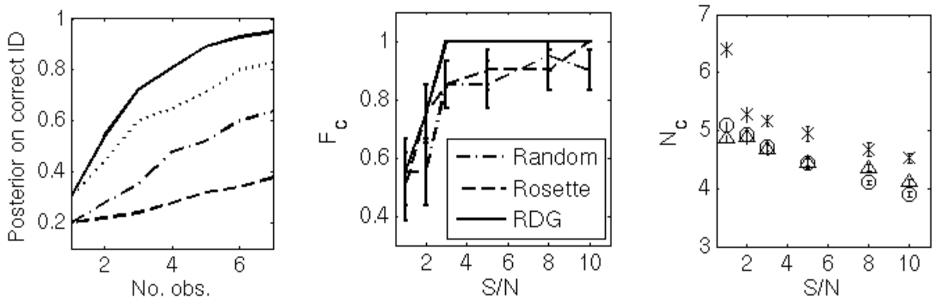

FIGURE 2. *Left:* Posterior probability of classification for the Manta type in sidescan simulation using the confusion-matrix fusion method (solid and dotted lines). The dashed and dot-dashed lines show the results of using a voting scheme for fusing. The solid and dot-dashed lines employ RDG whereas the other methods use a random-look scheme. *Centre:* Fraction of correct classifications of the Rockan type using the A-scan from among the seven types versus Signal-to-Noise ratio. The dashed line indicates results from the rosette observation schedule, the dot-dashed line from random observation and the solid line the RDG results. In all cases the target was observed 7 times. *Right:* Number of observations required to reach $P_c = 0.99$ for those correctly classified as Rockan. Crosses denote random observations, triangles the rosette scan and circles are the RDG method.

to the target. In our results, we quantify S/N rather than list the full gamut of parameters used in determining the noise.

3.2. *Observing schedules*

Targets are observed at a range of 30m. For the A-scan the source level is 200dB and the observation frequency 4.5kHz. The sidescan source observation frequency is set at 300 kHz. We test target classification against three sequential observing schedules:

- **Random** The next observation angle is chosen at random from the set of angles $\theta = \{\theta + n\pi/4\}$ where $n = \{1, ..., 8\}$.
- **Rosette** The sequence of observations is designed to simulate a rosette pattern – a standard AUV observation pattern which seeks to compile as many observation angles as practical for target identification. The angles are $\theta = \{0, \pi, \pi/4, 5\pi/4, \pi/2, 3\pi/2, 3\pi/4, 7\pi/4\}$. These may not correspond exactly to a rosette as it is usual to undertake several $\theta, \theta + \pi$ pairs at different ranges before switching to a different angle.
- **RDG** The next angle is chosen according to the method in §2.4. In order to retain a fair comparison the set of future observing angles is restricted to $\theta = \{\theta + n\pi/4\}, n = \{1, ..., 8\}$.

In all cases where the A-scan is used for simulation, DF is undertaken using the PF with 300 particles. We test the efficacy of DF against a simple voting method. In the voting scheme the ATR decisions are counted as votes and after N measurements the identification with the highest number of votes is selected. If the vote is tied then one of the identifications that ties for the most votes is selected at random. To test RDG we examine how many views are required to classify a target to a required level of confidence. Equivalently, we quantify the level of confidence reached after a given number of views of the target. We vary simulation parameters discussed in 3.1 in order to vary the S/N ratio.

3.3. *Results*

The utility of data fusion can be quantified by examining its performance against the voting scheme. Figure 2 (left) shows that, for a given number of views of a Manta mine simulated in sidescan, data fusion methods give a higher posterior probability of classifying the correct target than the voting scheme. This panel also shows that the RDG method (either with the confusion-matrix-based data fusion, or the voting scheme) always produces a higher posterior probability of correct classification than random views of the target.

Figure 2 (centre) shows the results of a series of tests to classify the Rockan type using the A-scan simulator. The graph shows the fraction of the tests which result in the highest P_c corresponding to the correct class within seven observations. The results are plotted for varying signal-to-noise conditions. The RDG method gives a higher fraction than either the random or rosette patterns. This indicates that RDG will deliver a higher P_c for a given number of views of the target. Therefore we should expect that RDG allows one to classify to a given P_c in fewer views than either of the other methods. Figure 2 (right) shows this; for those instances where the Rockan type is correctly classified, the number of views taken to reach $P_c = 0.99$ is lower for the RDG method. This is true when tested against a set of random views (where replications are allowed), but ceases to be true when tested against the rosette pattern at lower S/N. However in these regimes the RDG method outperforms the rosette pattern in correctly classifying targets. These results hold when other mine and NoMBO types are tested.

4. Summary and future work

We have shown how principled techniques for Reactive Data Gathering (RDG) can be used to improve the classification rate for AUVs undertaking MCM missions. A Sonar simulator was used to simulate various mine and NoMBO seabed targets in both A-scan and sidescan images. A series of experiments was then run to test the RDG method against an observation schedule composed of random views and one designed to simulate a rosette pattern. Data fusion was shown to improve classification accuracy over a simple voting scheme. The RDG method, which is ATR-neutral, was found to require fewer observations to reach a particular confidence level, or equivalently to reach a higher confidence level for a given number of views. Recent and future work (e.g. SEAS-DTC Exemplar 1 [10]) seeks to raise the technology readiness of this work by testing these algorithms on AUVs in the water. By applying RDG in-water, this work seeks to demonstrate a three-fold increase in operational tempo.

REFERENCES

1. Bladon P., Day P. 2006, *Reactive Data Gathering,* Proceedings of the 1st SEAS-DTC Technical Conference.
2. Bladon P., Hester G., Lloyd C., Nicholson D., Williams M.L. 2007, *Proactive Data Gathering,* Proceedings of the 2nd SEAS-DTC Technical Conference.
3. Pailhas Y., Petillot Y., Capus C., Brown K. 2009, *Real-time Sidescan Simulator and Applications,* IEEE Oceans 2009 Europe Conferences, Bremen
4. Litchfield M., Nicholson D., Anderson J. 2006, *RCSC18: Project B, UAVs for Future Target Detection,* Contract No. RT/COM/2/015, Project No. 103433
5. Runkle P. et al. 1999, *Hidden Markov Models for Multiaspect Target Classification,* IEEE Transactions on Signal Processing, 47(7), 2035
6. Ji S., Liao X., Carin L. 2005, *Adaptive Multi-Aspect Target Classification and Detection with Hidden Markov Models,* IEEE Sensors Journal, 5(5), 1035
7. Ristic B., Sanjeev Arumpalam M., Gordon N. 2004, *Beyond the Kalman Filter: Particle Filters for Tracking Applications,* Artech House, ISBN 9781580536318
8. Sanjeev Arulampalam M. et al. 2002, *A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking,* IEEE Transactions on Signal Processing, 50(2), 174
9. IHS Jane's Defense & Security Intelligence & Analysis. www.janes.com
10. Doggart T., Mew M., Andrews G. 2010, *The Role of Technology Integration, Evaluation and Demonstration in the SEAS DTC,* Proceedings of the 5th SEAS-DTC Technical Conference.

