PROCEEDINGS OF SPIE

SPIEDigitalLibrary.org/conference-proceedings-of-spie

Multi-phase performance evaluation for modern minehunting systems

Øivind Midtgaard, Narada Warakagoda, Gary L. Davies, Warren Connors, Marc Geilhufe

Øivind Midtgaard, Narada Warakagoda, Gary L. Davies, Warren Connors, Marc Geilhufe, "Multi-phase performance evaluation for modern minehunting systems," Proc. SPIE 11012, Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XXIV, 110120P (10 May 2019); doi: 10.1117/12.2518837

SPIE.

Event: SPIE Defense + Commercial Sensing, 2019, Baltimore, Maryland, United States

Multi-phase performance evaluation for modern minehunting systems

Øivind Midtgaard*[a], Narada Warakagoda[a], Gary Davies[b], Warren Connors[c], Marc Geilhufe[a]
[a]Norwegian Defence Research Establishment (FFI), P. O. Box 25, NO 2027 Kjeller, Norway;
[b]Defence Science and Technology Laboratory (Dstl), Porton Down, Salisbury, Wiltshire SP4 0JQ, United Kingdom; [c]Defence Research and Development Canada (DRDC), Atlantic Research Centre, P. O. Box 1012, Dartmouth, Nova Scotia, B2Y 3Z7, Canada

ABSTRACT

Many NATO navies are in the process of replacing their dedicated minehunting vessels with systems of heterogeneous, unmanned modules. While traditional ship-based assets prosecute sonar contacts in sequence through to neutralisation, modern systems employ unmanned vehicles equipped with side-looking sonar to detect and classify minelike contacts in a full area segment before proceeding with contact identification and mine neutralisation. This shift in technology and procedure brings important operational advantages, but also introduces a need to modify the traditional minehunting performance evaluation based on the percentage clearance metric. Previous works have demonstrated that the achieved detection and classification performance of modern minehunting systems can be estimated from the collected sonar data (through-the-sensor) and reported as detailed geographical maps. This paper extends the map-based evaluation approach to the identification and neutralisation phases, and also includes the case where some of the contacts or mines intentionally are left unprosecuted, e.g. disposal of only the specific mines required for establishing a safe sailing route. Each map cell is assumed to be sufficiently small to contain at most one sonar contact and can thus be assigned a status based on the hunting results for that cell: minelike contact, identified mine, etc. To this end we derive Bayesian formulations of a new performance metric: the probability of a remaining mine in a given cell. Furthermore, we show that this metric provides consistent multi-phase performance evaluation and estimates of the mine impact risk for a follow-on ship transiting a specified route.

Keywords: Underwater Minehunting, Performance Evaluation, Unmanned Maritime Systems, Residual Risk

1. INTRODUCTION

Many navies are transitioning from the traditional approach to underwater minehunting, where a multi-beam forward-looking sonar is either mounted on or deployed from a specialized ship, to an approach using unmanned, modular systems. The traditional approach requires that the sonar is able to detect and classify mines at a safe distance ahead of the mine counter-measures vessel (MCMV), in order to minimize risk to the vessel and its crew (typical detection range is many hundred meters). This restriction limits the resolution of the sonar imagery that can be obtained which, in turn, limits the mine counter-measures (MCM) effectiveness. Another restriction is that potential mines usually have to be dealt with (identified and, if necessary, disposed of) before the ship can safely continue to search for other mines. Robotic vehicles can be fitted with shorter-range side-looking sonars that provide higher resolution and therefore have the potential to be far more effective. In general, side-looking sonar provides better image quality if the carrier platform travels in straight lines, minimizing the number of turns. This means that minehunting missions using unmanned systems are more efficient if an area is first searched for mines using a series of straight legs and then each potential mine is prosecuted in turn, possibly with the use of other platform types. This leads to a serial phase approach to minehunting that is described later in this paper.

The long-range, forward-looking sonars used in traditional minehunting generate large quantities of data. This has historically made it impractical to store and analyse the sonar data following a mission. The inputs relating to sonar performance that are required by the mission evaluation algorithms therefore have to be based on average values over a large area (and these values often require an element of operator judgement). This limits both the fidelity and resolution of mission evaluation that can be performed using the traditional approach. The lower data rate of side-looking sonars allows all the sonar data to be stored in the vehicle, enabling an objective assessment of performance based on high

*Oivind.Midtgaard@ffi.no; phone +47 6380 7653; ffi.no

Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XXIV, edited by
Steven S. Bishop, Jason C. Isaacs, Proc. of SPIE Vol. 11012, 110120P · © 2019 SPIE
CCC code: 0277-786X/19/$18 · doi: 10.1117/12.2518837

resolution sensor data. This should lead to increased fidelity evaluation carried out at a much higher resolution. In addition, the approach used by most NATO navies to evaluate minehunting missions contains several assumptions that tend to break down when applied to missions performed by unmanned vehicles with side-looking sonars [1]. These factors provide a strong motivation for re-thinking the approach to evaluation of modern minehunting operations. The use of detailed performance maps fits well into such a paradigm shift.

Previous works on map-based planning and evaluation of the search phase of minehunting missions using side-looking sonar have led to the computer models AUVPET [2] and MCM Insite [3]. This paper develops these map-based evaluation approaches to include the identification and neutralisation phases in addition to the search phase. The approach exploits the sensor data recorded during each phase to produce a high resolution geographical map of the probability of remaining mines in the area. The probability map can be used to estimate the residual risk of mine impact for a transiting ship. This approach is related to the use of probabilistic occupancy grids for robotic path planning [4,5].

2. MINEHUNTING PHASES

Minehunting is the process of finding and countering naval mines. This process is structured around a set of well-defined phases, as visualized in Figure 1. First, sonar data of the seafloor and water column is collected for the detection of objects as minelike echoes (MILECs). More detailed features extracted from both sonar echo and shadow are then used to classify each object as either a minelike contact (MILCO) or a non-minelike contact (Non-MILCO). Next, each MILCO is identified to determine whether it is a mine, including its specific model, or a non-mine minelike bottom object (NOMBO). Identification typically involves the use of visual means, such as a camera on an unmanned vehicle or the employment of a diver. The final neutralisation phase concerns the actual countering of the identified mines (MINEs), where an explosive charge is placed close to the mine through human or robotic intervention. Mine disposal may be verified by re-inspecting the object location after charge detonation.

Traditional minehunting has been conducted by low-signature MCMVs equipped with advanced forward looking sonars for long-range detection and classification, as well as divers or remotely operated vehicles (ROVs) for close-up identification and neutralisation. The process starts with the execution of a pre-planned set of survey tracks designed to provide the required sonar coverage of the area. Upon discovery of a MILEC, the survey will be halted. The MCMV remains at expected safe stand-off range while the object of interest is prosecuted through classification, identification and neutralisation, unless being deemed a Non-MILCO or NOMBO (Figure 1, top). During these phases, the classification sonar is fixed onto the detected echo to monitor its relative position, thus ensuring that the same object is considered in all phases. After completed prosecution of the object, the MCMV resumes the detection phase survey.

Modern minehunting employs multiple robotic platforms which have been designed for a specific phase (Figure 1, bottom). These scalable and modular systems offer important advantages including increased efficiency, improved sensor data, relative low cost and reduced risk to personnel. For the detection and classification in the search-classify-map phase, autonomous underwater vehicles (AUVs) or unmanned surface vehicles (USVs) equipped with high-resolution, side-looking sonar are used to search the full area and locate all MILCOs (without first calling MILECs). Then, identification and neutralisation can be conducted with divers, ROVs or dedicated mine disposal weapons, all deployed from either a traditional MCMV or a vessel of opportunity. Alternatively, identification can be performed by an AUV equipped with an optical camera, thereby significantly reducing the number of contacts for which a separate disposal asset must be launched. This capability is particularly relevant in areas with high clutter density or complex seafloor. If the purpose of the MCM operation is to establish a safe route through a wide area, neutralisation can then also be limited to the minimum number of mines that cannot be circumnavigated.

One of the primary differences between the two approaches is the sequential nature of the traditional process based on MCMVs. These platforms will complete all phases on a detected object *in situ*, before continuing the search phase. On the contrary, modern techniques apply specialized modules to detect and classify all objects of interest during the search-classify-map phase before proceeding with the subsequent phases. The separation of the phases requires consideration of reacquisition, i.e. the process of locating a previously detected target based on a global or relative position. This can pose a challenge as the navigation error within each phase must be accounted for to ensure successful reacquisition of the object. As shown in Figure 1 (bottom), each mine object needs to be reacquired twice when two different platforms are employed for identification and neutralisation, respectively.

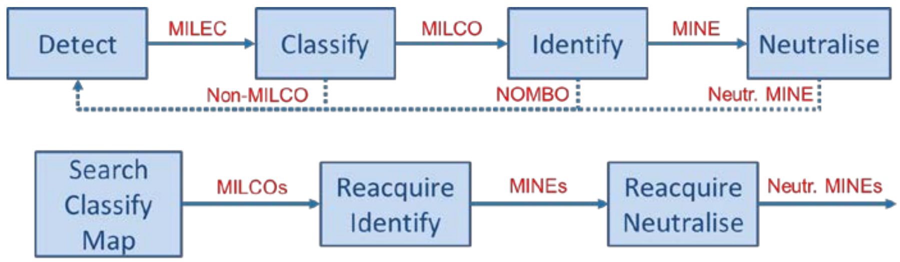

Figure 1. Minehunting phases for traditional MCMV (top) and modern, modular systems (bottom).

3. PERFORMANCE AND RISK EVALUATION

The aim of MCM performance evaluation is to quantify the achieved effect of the invested efforts. This is essential in order to verify to what extent the demanded results have been obtained and, if necessary, to plan the best use of further MCM resources in the area. Estimates of performance and number of countered mines are also needed to quantify the residual risk, which is the probability of mine impact for subsequent ship traffic. Figure 2 shows the processing cycle of *planning-execution-evaluation* in modern minehunting. Maximum efficiency requires automated, in-vehicle processing.

Figure 3 illustrates the prosecution of a single object in a modern minehunting system using two separate platforms for identification and neutralisation. An object that is detected and classified as minelike is labelled MILCO, otherwise it is labelled Non-MILCO (this covers both Non-MILEC and Non-MILCO in traditional minehunting shown in Figure 1, top). Reacquisition failures have different implications for identification and neutralisation. A failed MILCO reacquisition is assumed to result in the mistaken identification of a nearby, background object, which may lead to a

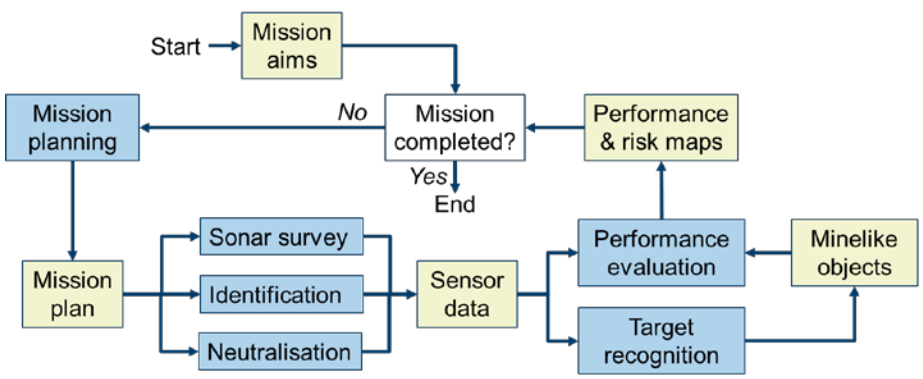

Figure 2. Diagram of processing cycle for modern minehunting operations. Blue boxes represent processes and yellow boxes represent data/results. The evaluation of performance and risk is based on *in situ* sensor data and recognition of minelike objects by a human operator or an automated system. The mission plan is regularly adapted based on the estimated performance and risk maps to ensure both mission effectiveness and efficiency. The mission is completed when its specified aims have been reached or are considered unfeasible based on the recorded sensor data (e.g. unhuntable seafloor).

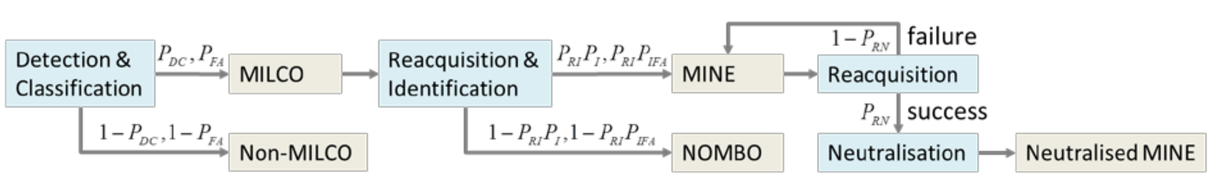

Figure 3. Object prosecution in modern mine hunting system. Blue boxes represent stochastic processes and gray boxes represent the possible results for a given object or grid cell. The processing probabilities for mines and non-mines (false alarms) are shown for each result, except the probability for successful neutralisation, P_N, because the output label is Neutralised MINE whether the process succeeds or fails.

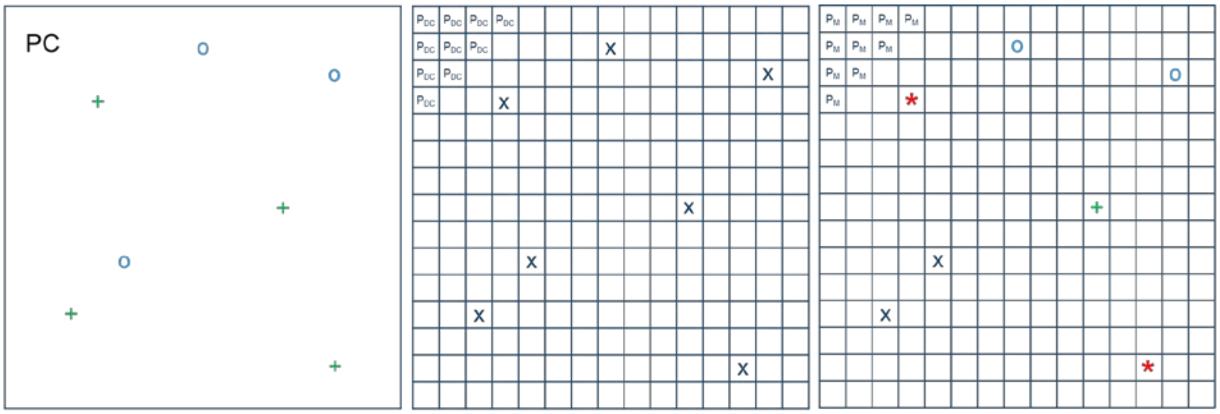

Figure 4. **Left:** Traditional, global minehunting performance evaluation outputs a percentage clearance (PC) value averaged over an area. No geographical information is reported, except positions of identified NOMBOs and Neutralised MINEs. **Centre:** Map-based performance evaluation of the sonar search phase reports the detection and classification probabilities per cell for mines (P_{DC}) and non-mines (P_{FA}), as well as MILCO positions. **Right:** Map-based, multi-phase minehunting evaluation outputs the posterior probability of a remaining mine per cell P_M given achieved results, as well as positions of contacts at various prosecution levels. Symbols represent: MILCO x, NOMBO o, MINE * and Neutralised MINE +.

mine MILCO being identified as a NOMBO. These mistakes may reduce the mine identification rate and typically occur in environments with dense clutter and low visibility. On the other hand, a failed MINE reacquisition should be evidently noticed and prohibit any attempt of neutralisation, thus leaving an identified MINE in that position.

Approaches for performance and risk evaluation can be categorized as global or map-based as displayed in Figure 4 and described in the following sections.

3.1 Global evaluation

Global evaluation methods report performance as values averaged over an area. A large operation area is typically divided into smaller segments, where constant parameter values are assumed within each segment. It can be challenging to establish representative values, because the input parameters, in particular detection probabilities, depend heavily on the underwater environment.

3.1.1 Percentage clearance

The prevailing MCM performance metric for both traditional and modern systems is percentage clearance *PC*, which is defined as the percentage probability to detect and dispose a given mine. When a single mine type is considered, *PC* represents the expected fraction of mines in the area that have been detected and neutralised. Assuming conditional independence, *PC* can be estimated as the product of the probabilities of a successful outcome from each of the minehunting phases. For modern, heterogeneous systems (Figure 3) this can be written as:

$$PC = \underline{P}_{DC}\,\underline{P}_{RI}\,\underline{P}_{I}\,\underline{P}_{RN}\,\underline{P}_{N} \quad (1)$$

where any underlined parameter denotes the area averaged value. Further, P_{DC} is the mine detection and classification probability, P_{RI} is the probability of reacquiring a given contact for identification, P_I is the mine identification probability, P_{RN} is the probability of reacquiring a given identified mine for neutralisation and $\underline{P}_N$ is the probability of successful mine neutralisation. If a single platform is used for concurrent identification and neutralisation of a mine, P_{RN} is typically set to 1. The probabilities are based on the total efforts in each minehunting phase. If for example the seafloor has been imaged multiple times by side-looking sonar (e.g. using two perpendicular lawn-mower survey patterns), P_{DC} constitutes the combined probability for mine detection and classification based on all sensor views. Upon fusion of the individual views, the sensing geometries and sonar settings must be compared to determine the degree of statistical dependence between the data sets and this dependence must be taken into account. Similarly, fusion is needed when several attempts of reacquisition are performed for the same object during a single phase or multiple camera views are available for identification. As Equation (1) includes the product of two reacquisition probabilities, maximum achievable *PC* may be severely limited by the absolute or relative navigation accuracies of the platforms.

3.1.2 Percentage prosecuted

An important restriction with the *PC* metric is the assumption that all MILCOs have gone through identification and all MINEs have been attempted neutralised by the end of the minehunting operation. This metric thus cannot be used for e.g. the case of deliberately neutralising only the minimum subset of the identified mines needed to establish a safe sailing route. However, corresponding performance metrics for the detection/classification and identification phases can be defined as the expected fractions of mines correctly called as MILCOs and MINEs, respectively. These metrics are given as: percentage classified = $\underline{P}_{DC}$ and percentage identified = $\underline{P}_{DC}\,\underline{P}_{RI}\,\underline{P}_{I}$.

We define a new performance metric, percentage prosecuted *PP*, as the expected fraction of successfully prosecuted mines, given the selected prosecution level for each object. This fraction comprises all mines still labelled either MILCO or MINE after completion of the MCM operation, in addition to truly neutralised mines. Importantly, the positions of all remaining MILCOs and MINEs must be reported together with the *PP* value, as these uncountered objects must be treated as potential mines and circumvented at safe range by subsequent traffic. *PP* can be written as a weighted sum of percentage classified, percentage identified and percentage clearance:

$$\begin{aligned} PP &= \underline{P}_{DC}\left(1-f_I\right)+\underline{P}_{DC}\,\underline{P}_{RI}\,\underline{P}_{I}\,f_I\left(1-f_N\,\underline{P}_{RN}\right)+f_I\,f_N\,PC \\ &= \underline{P}_{DC}\left(1-f_I\left(1-\underline{P}_{RI}\,\underline{P}_{I}\left(1-f_N\,\underline{P}_{RN}(1-\underline{P}_{N})\right)\right)\right) \end{aligned} \tag{2}$$

where f_I is the fraction of MILCOs selected for identification and f_N is the fraction of MINEs selected for neutralisation. It is assumed that the true mines are proportionally distributed between the identified/non-identified MILCOs and between the neutralised/non-neutralised MINEs (i.e. fraction f_I should not consist of only the most minelike MILCOs). The two reacquisition probabilities appear differently in Equation (2), because P_{RI} reduces the mine identification rate while P_{RN} reduces the fraction of neutralised MINEs, as described for Figure 3. *PP* collapses into the traditional *PC* for the special case where all objects are fully prosecuted (i.e. f_I, f_N and P_{RN} all equal 1). One intriguing aspect of Equation (2) is that *PP* decreases when contacts go through identification and neutralisation, because more probabilities are factored in. However, these additional hunting efforts eliminate MILCOs and MINEs which otherwise would require further consideration.

The *PC* and *PP* performance metrics do not consider false alarms (non-mines). The number of non-mine MILCOs must be sufficiently small to allow necessary processing within the available time window. Otherwise the area is simply regarded as unhuntable with the given minehunting systems. A zero error rate is usually assumed for identification with traditional MCMVs, as the operator/diver can thoroughly investigate the contact from different aspects until he or she is confident whether the contact is a MINE or NOMBO. Challenges may arise for densely cluttered seafloors, poor visibility and degraded object surfaces (e.g. corrosion, marine growth). Still, the product of the probabilities for contact reacquisition, identification and neutralisation is assumed to be close to 1 for MCMVs. On the contrary, performance of modern minehunting systems can be limited by failure rates in automated reacquisition systems and significant false alarm rates from automated target recognition (ATR), particularly in challenging environments [6].

3.1.3 Number of deployed mines

Evaluation of the risk of damage to follow-on ship traffic requires an estimate of the remaining number of mines in the area. This is the difference between the number of deployed mines *n* and the number of truly neutralised mines. We can apply Bayes' rule [7] to estimate the posterior probability distribution for *n* given the minehunting results:

$$p(n\,|\,Results)=\frac{p(Results\,|\,n)P(n)}{p(Results)}=\frac{p(Results\,|\,n)P(n)}{\sum_{i=0}^{\infty}p(Results\,|\,i)P(i)} \tag{3}$$

where *Results* denote detected MILCOs and identified MINEs. *P(n)* is the prior mine number probability which can be estimated from intelligence information regarding the adversary's capacity and opportunity for laying mines. Alternatively, a flat prior probability distribution can be selected, up to a maximum number of mines determined from the size of the operation area and the absolute minimum distance between mines.

The binomial distribution [7] $P_b(k|n,p)$ specifies the probability of achieving exactly *k* successful outcomes in a sequence of *n* independent, binary valued experiments, given the probability *p* of success for each individual experiment:

$$\mathrm{P}_b(k\,|\,n,p) = \binom{n}{k} p^k (1-p)^{n-k} \tag{4}$$

The detection and classification of a given mine can be modelled as a binary, stochastic process with probability P_{DC} of a positive outcome. Thus, defining m_{DC} as the number of detected MILCOs, Equation (4) gives an expression for $p(m_{DC}|n)$ which can be inserted in Equation (3) to estimate $p(n|m_{DC})$. However, this presumes that all MILCOs correspond to true mines. To account for false alarms, we need to consider the joint probabilities of detecting exactly k mines and m_{DC}–k non-mines, for all possible values of k:

$$p(m_{DC}\,|\,n) = \sum_{k=0}^{m_{DC}} P_b(k\,|\,n,\underline{P}_{DC})\; P_b(m_{DC}-k\,|\,N-n,\underline{P}_{FA}) \tag{5}$$

where N is the number of non-overlapping area cells surveyed by sonar. As each cell is sufficiently small to contain at most one sonar contact, N–n is the total number of false alarm opportunities. P_{FA} is the probability of falsely calling a MILCO in a given cell. Assuming that all MILCOs have been prosecuted for identification, the probability $p(m_I|n)$ of observing the number of identified MINEs m_I given the total number of deployed mines n can similarly be estimated from Equation (5) by replacing m_{DC} with m_I and also replacing P_{DC} and P_{FA} with the products $P_{DC}P_{RI}P_I$ and $P_{FA}P_{RI}P_{IFA}$, respectively. Alternatively, we can modify Equation (5) to estimate the joint probabilities of observing both m_{DC} and m_I given n, by including the probabilities of identifying exactly i out of k mines and m_I–i out of m_{DC}–k non-mines, for all possible values of i:

$$\begin{aligned} p(m_I, m_{DC}\,|\,n) = \sum_{i=0}^{m_I} \sum_{k=i}^{m_{DC}-(m_I-i)} & P_b(k\,|\,n,\underline{P}_{DC})\; P_b(i\,|\,k,\underline{P}_{RI}\,\underline{P}_I)\; P_b(m_{DC}-k\,|\,N-n,\underline{P}_{FA})\ldots \\ & P_b(m_I - i\,|\,m_{DC}-k,\underline{P}_{RI}\,\underline{P}_{IFA}) \end{aligned} \tag{6}$$

where P_{IFA} is the probability of falsely identifying a non-mine MILCO as a MINE instead of a NOMBO. After insertion in Equation (3) the peak of the resulting distribution $p(n|m_{DC},m_I)$ will be more distinct (higher) than that of p$(n|m_{DC})$, because the identification phase discriminates far better between mines and non-mines than the sonar search phase (see example in Figure 6).

3.2 Map-based evaluation

A significant restriction of global evaluation is the requirement of estimating input values that are representative for a complete area segment. Traditionally, this input has been generated with a certain degree of operator judgement. Also, any information on performance variations within the area will be lost due to averaging. A better approach for modern systems, in which all high-resolution sensor data is stored and georeferenced, is to evaluate local performance from this data. The outputs are then detailed geographical maps with a large number grid cells, so that a homogenous environment can be assumed within each cell. The cell size is indeed set sufficiently small (typically around 2m x 2m) to contain at most a single sonar contact. Each cell can thus be assigned a status based on the hunting results for that cell, as illustrated in Figure 3. If no MILCO is declared for a cell, its status is Non-MILCO, which will be the case for the vast majority of the cells. If a contact is declared, the cell status will be MILCO, NOMBO, MINE or Neutralised MINE depending on the level of prosecution and recognition results. The product is a heterogeneous phase performance map (Figure 4, right).

3.2.1 Through-the-sensor estimation

The mine recognition performance depends on three general parameters [6]: mine threat, sensor system and environment. Each of these parameters can be decomposed further into a range of sub-parameters. The mine threat is determined by the mine type, configuration, surface condition (corrosion, marine growth, deposits) and degree of burial. The sensor system consists not only of its hardware and settings, but also the platform motion, sensing geometry, signal processing and data analysis system. The environment affects performance through many factors, some of which vary mainly in space (e.g. sediment type, bathymetry, seafloor roughness and clutter density), while others vary in both space and time (e.g. sea state, sound velocity profile, marine life, currents/tides, waves and turbidity). Variability over relevant spatial and temporal scales can be particularly large in the littoral waters typical for MCM operations. This implies that complete and accurate environmental information is rarely available before a minehunting operation and predictions will be unreliable without assured inputs. *In situ* measurements are thus required to establish the influence on performance from the experienced physical conditions during the mission. The use of the actual payload sensor data for this is referred to as through-the-sensor estimation.

As shown in Figure 3, the main performance metrics for the sonar search phase are the probabilities of detection and classification for mines (P_{DC}) versus non-mines (P_{FA}). Complex dependencies make it difficult to infer these two probabilities directly from the sensing and environmental parameters. A more viable approach [3] uses a performance model combining through-the-sensor estimation with a priori system specifications to calculate two latent parameters for each grid cell. The latent parameters image quality and image complexity estimate how well a potential mine would be represented in the sonar image and how difficult it would be to recognize a mine against the local image background, respectively. These parameter values are mapped into performance values using prior learnt look-up tables for different mine threats. The resulting performance maps of P_{DC} and P_{FA} (Figure 4, centre) can be used for planning and evaluation of the search phase (Figure 2). Global estimates ($\underline{P}_{DC}$ and $\underline{P}_{FA}$) are easily found by averaging over all grid cells.

The performance of the identification phase should also be evaluated using a through-the-sensor approach. Latent parameters can be calculated from camera images to estimate *in situ* the probabilities of MINE identification for mines (P_I) and non-mines (P_{IFA}). Identification performance varies significantly with e.g. water turbidity, sensor range and object surface conditions. The estimate for the probability of reacquisition should take into account the reported navigation system accuracy at the times of both the contact detection and the attempted identification.

3.2.2 Probability of remaining mine

The global, multi-phase performance metrics, *PC* and *PP*, basically estimate the area-averaged probability that a certain minehunting result (e.g. identified MINE) has been achieved for a given mine. These metrics do not comply with map-based evaluation where a more relevant metric would be the probability that a mine is present in a cell given its minehunting results. Conceptually, this is a risk metric, rather than a performance metric.

We define M as the event that a mine remains in the cell with indices (i,j) and $Results(i,j)$ as a variable-length sequence of single minehunting events for that cell: DC denotes a detected MILCO, DC^C denotes a Non-MILCO, I denotes an identified MINE, I^C denotes an identified NOMBO and N denotes a Neutralised MINE. We can apply Bayes' rule to estimate the posterior probability of M as an update of the prior probability $P(M)$ based on the observed results sequence for the cell (cell indices i,j are omitted for simplicity):

$$P(M \mid Results) = \frac{P(Results \mid M)P(M)}{P(Results)} = \prod_{\forall k} \frac{P(Results_k \mid M)}{P(Results_k)} P(M) \tag{7}$$

where $Results_k$ is the k'th element in the *Results* sequence for cell (i,j). The iterative product formulation to the right demands conditional independence between the *Results* event elements. The prior mine probability $P(M)$ depends on the total number of deployed mines which can be estimated from Equation (3) using Equation (5) or (6) with the observed numbers of MILCOs and MINEs as inputs.

The two mine probability updates for the search-classify-map phase (MILCO and Non-MILCO) are given as:

$$P(M \mid DC) = \frac{P(DC \mid M)P(M)}{P(DC)} = \frac{P_{DC}P(M)}{P_{DC}P(M) + P_{FA}\left(1 - P(M)\right)} \tag{8}$$

$$P(M \mid DC^C) = \frac{P(DC^C \mid M)P(M)}{P(DC^C)} = \frac{(1 - P_{DC})P(M)}{(1 - P_{DC})P(M) + (1 - P_{FA})\left(1 - P(M)\right)} \tag{9}$$

where P_{DC} and P_{FA} are the detection and classification probabilities for mines and non-mines, respectively, fused for all sonar views of the map cell at indices (i,j). Alternatively, Equations (8) and (9) can be applied iteratively with the recognition outcome for each single view of the cell, on the premise of conditional independence between the views. If this premise is invalid, the dependence model between the views needs to be incorporated in the updates. If the recognition system provides confidence values with each positive outcome (declared MILCO) and local Receiver Operator Characteristics (ROC) curves are available for each grid cell, P_{DC} and P_{FA} values corresponding to the confidence value's operating point on the ROC curve can be used in Equation (8). High confidence MILCOs will then lead to higher mine presence probability values than low confidence MILCOs.

Similarly, the mine probability can be updated based on identification phase results (MINE and NOMBO):

$$P(M \mid DC, I) = \frac{P(I, DC \mid M)P(M)}{P(I, DC)} = \frac{P_I P(M \mid DC)}{P_I P(M \mid DC) + P_{IFA}\left(1 - P(M \mid DC)\right)} \tag{10}$$

$$P(M \mid DC, I^C) = \frac{P(DC, I^C \mid M)P(M)}{P(DC, I^C)} = \frac{(1 - P_{RI} P_I)P(M \mid DC)}{(1 - P_{RI} P_I)P(M \mid DC) + (1 - P_{RI} P_{IFA})\left(1 - P(M \mid DC)\right)} \tag{11}$$

Finally, the probability update after the neutralisation phase is given by:

$$P(M \mid DC, I, N) = (1 - P_N)P(M \mid DC, I) \tag{12}$$

3.2.3 Mine impact risk

After an MCM operation, it is of fundamental importance to reliably estimate the residual risk for damage to follow-on traffic from remaining mines. One commonly used risk metric is the probability of mine impact for the first vessel transiting the area. This metric is typically estimated while ignoring ship counters (the mine may have implemented logic to allow a pre-set number of ships to pass it before detonating) and assuming the mine has 100% detection probability for ships passing within its sensors' range. Mine impact will then occur when the transitor comes closer to a mine than the minimum of the mine's damage radius (W_d) and activation distance (W_a). Assuming conditional independence between grid cells, the residual risk can be calculated from the probability that none of the cells within an effective channel around the ship's path contains a mine (Figure 5):

$$Risk = 1 - \prod_{\forall i,j} (1 - p_{i,j}(r < R)\, P(M \mid Results(i, j)) \tag{13}$$

where $R=min(W_d, W_a)$ and $p_{i,j}(r<R)$ is the probability that cell (i,j) lies within distance R from the ship's path given uncertainties in ship navigation, grid cell position and R. The mine damage radius can vary within the operation area, because it depends not only on the mine type, but also on bathymetry and seafloor type.

The map-based evaluation approach makes it possible not only to calculate the risk for an arbitrary route based on achieved minehunting results, but also to find the safest route through the area using automated search for the minimum risk route (assuming the required channel width is smaller than the area width). This search is closely related to how robotic planning algorithms search probabilistic occupancy grids for collision-free paths in a static environment [4,8].

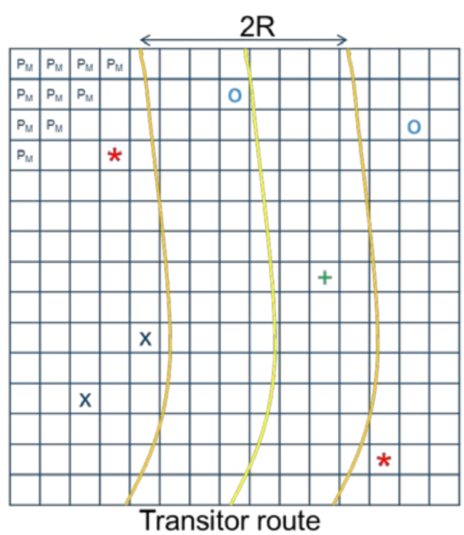

Figure 5. The probability map for remaining mines can be used to estimate the mine impact risk to the first transitor following a defined route (yellow line). Impact occurs if at least one mine lies within the damage/activation distance, *R*, of the route. Symbols represent: MILCO x, NOMBO o, MINE * and Neutralised MINE +.

4. EXAMPLE NUMERICAL RESULTS

This chapter presents example numerical results for the proposed evaluation approaches described in Chapter 3 using the default input parameter values listed in Table 1 unless other values are explicitly specified. Note that the probabilities in Table 1 and Table 2 are given as percentage values.

Figure 6 presents posterior probabilities for the number of deployed mines, estimated from Equation (3) with a flat prior distribution up to n=100. The dotted curves have been created with Equation (5) for the indicated input number of MILCOs m_{DC}, while the solid curves have been created with Equation (6) for the indicated input numbers of both MILCOs and identified MINEs m_I. The results show that identification significantly reduces the variance in the probability distribution for the number of mines. Comparison of the blue and red curves reveals that lower P_{DC} increases both the distribution variance and the expected number of mines (for fixed numbers of MILCOs and MINEs), due to an expected larger number of undetected (Non-MILCO) mines. The green curves have been calculated with a lower m_{DC} value of 50, which corresponds to the expected value for the number of non-mine MILCOs with the used values for P_{FA} and number of map cells. The dotted, green curve thus has maximum at n=0, as all the MILCOs can be explained as false alarms. However, updating the estimated probabilities for n with the observation m_I=10 (which provides little support to the hypothesis that all 50 MILCOs are false alarms) gives the solid, green curve which is almost similar to the solid, blue curve obtained for m_{DC}=60. This shows that for the given default parameters the value m_{DC} has only a minor influence on the output estimate, when m_I is available.

Table 1. Default parameters values used in Figure 6 and Table 2.

Area size	Cell size	P_{FA}	P_{DC}	P_{RI}	P_I	P_{IFA}	P_N
2.0 km^2	4.0 m^2	0.01%	90%	95%	99%	2%	99.9%

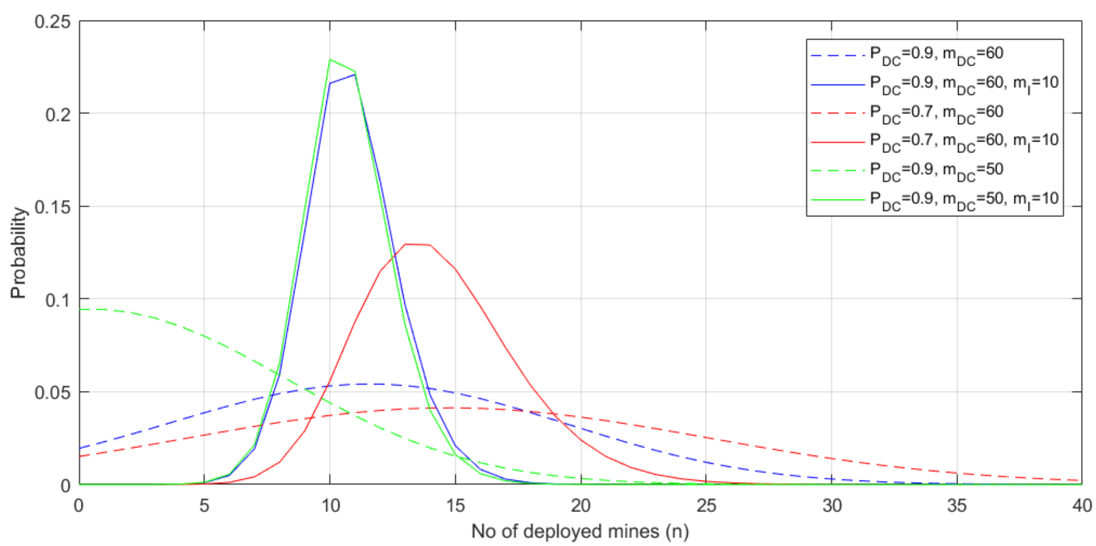

Figure 6. Posterior probabilities for number of deployed mines estimated from observed numbers of MILCOs m_{DC} and MINEs m_I, as well as parameters listed in Table 1 and a flat prior probability $P(n)$. Dotted curves are produced with Equation (5) using only m_{DC} and solid curves with Equation (6) using both m_{DC} and m_I.

Table 2 lists cell-wise mine presence probabilities for the different cell statuses calculated from Equations (8)-(12) for 10 mines distributed over the area. For each table row, the indicated parameter value is changed, while the other parameters are set to the default values in Table 1. For the row with increased cell size, the P_{FA} is increased accordingly to maintain the same expected number of non-mine MILCOs over the full area. The table shows to what extent the mine presence probability is shifted onto the MILCO cells during the sonar survey. With the selected input parameter values, the mine probability in a MILCO cell is around 7500 times larger than the prior probability, while the mine probability in a Non-MILCO cell is 10 times smaller than the prior probability. This asymmetry is of course due to the significantly smaller number of MILCO cells relative Non-MILCO cells. Also, these probability values are sensitive to the ratio of the numbers of actual mines and false alarms (non-mine contacts).

The mine probability values in identified MINE column in Table 2 are high, but still well below 100%, which would be the result with the assumption of zero false alarms often used for traditional identification systems (i.e. $P_{IFA}=0$). If false alarms (non-mine MINEs) were not considered, all three solid curves in Figure 6 would be zero for $n < 10$ (m_I). The mine probability values in NOMBO cells are low, but still considerably higher than in both Non-MILCO and No-Coverage cells. The reason is that the MILCO cells contain a large fraction of the mines in the area (in fact, the expected fraction is P_{DC}), thereby producing a finite posterior mine probability for NOMBO cells even though identification yields high mine/non-mine discrimination. Similarly, the mine probabilities in identified MINE cells are high, because a large fraction of mines is located in these relatively few cells. The mine probabilities of Neutralised MINE cells are thus approximately $1–P_N$, which in our example actually exceeds the mine probability in No-Coverage cells, even for P_N=0.9999. This illustrates differences between performance and risk evaluation, as the minehunting performance is zero in cells without sensor coverage.

Table 2. Probabilities of remaining mine given the cell's minehunting results. The default parameter values are specified in Table 1. For each row, one parameter value is changed as indicated. The number of deployed mines *n* is set to 10.

Cell status	No sensor coverage	Non-MILCO	MILCO	NOMBO	Identified MINE	Neutralised MINE
Mine probability	$P(M)$ [%]	$P(M\|DC^C)$ [%]	$P(M\|DC)$ [%]	$P(M\|DC,I^C)$ [%]	$P(M\|DC,I)$ [%]	$P(M\|DC,I,N)$ [%]
Default parameters	0.0020	0.00020	15.25	1.080	89.91	0.0899
P_{FA}=0.005%	0.0020	0.00020	26.47	2.137	94.69	0.0947
No of mines=5	0.0010	0.00010	8.257	0.5429	81.67	0.0817
Cell size=16.0m^2 P_{FA}=0.04%	0.0080	0.00080	15.26	1.080	89.91	0.0899
P_{DC}=70%	0.0020	0.00060	12.28	0.8420	87.39	0.0874
P_{RI}=75%	0.0020	0.00020	15.25	4.494	89.91	0.0899
P_I=95%	0.0020	0.00020	15.25	1.758	89.53	0.0895
P_{IFA}=1%	0.0020	0.00020	15.25	1.070	94.69	0.0947
P_N=99.99%	0.0020	0.00020	15.25	1.080	89.91	0.0090

5. SUMMARY

In this paper we have discussed evaluation of minehunting performance and pointed out some of the deficiencies connected with applying traditional, global techniques to missions executed by modern, unmanned, modular systems. We propose to replace the prevailing percentage clearance metric with a novel percentage prosecuted metric, which supports heterogeneous phase minehunting, i.e. missions where not all minelike contacts have been fully prosecuted. Positions of remaining sonar contacts and identified mines must then be reported together with the percentage prosecuted value, and these positions need to be circumvented at safe distance by subsequent traffic.

However, full utilization of the advantages of modern minehunting systems requires an evaluation approach using high-resolution, geographical maps. For the sonar search phase, the probability of mine detection and classification, together with the false alarm probability, can be used for iterative mission planning and performance evaluation. Due to large spatial and temporal variability of the underwater environment, estimation should be based on *in situ* sensor data for each map cell.

As new metric for map-based, multi-phase evaluation, we propose the probability of a remaining mine in a cell given the hunting results obtained for the cell. This metric is applicable to all minehunting phases, even before the sonar search starts. Calculation of this metric requires a prior estimate of the mine density distribution in the area, which can be estimated from mine laying intelligence, if available. The mine probability estimate in each cell is then refined as results from the sonar search and the optical contact identification become available.

The mine probability map facilitates estimation of the damage risk from remaining mines to a vessel sailing a specified route. The map-approach also enables automated search to find the minimum risk (i.e. safest) route through a wide area.

ACKNOWLEDGEMENTS

This work has been supported by the Norwegian Defence Material Agency.

REFERENCES

[1] Percival, A. M., Couillard, M., Midtgaard, Ø., Fox, W., "Unmanned systems, autonomy and side-looking sonar: a framework for integrating contemporary systems into the operational architecture", NATO STO CMRE report CMRE-FR-2013-013, La Spezia, Italy (2013).
[2] Myers, V., Davies, G., Petillot, Y., Reed, S., "Planning and evaluation of AUV missions using data-driven approaches," Proc. MINWARA Conference, Monterey, CA, U.S.A. (2006).
[3] Midtgaard, Ø., Alm, I., Sæbø, T. O., Geilhufe, M., Hansen, R. E., "Performance assessment tool for AUV based mine hunting," Proc. of the Institute of Acoustics, vol. 36, pt. 1, 120-129, Lerici, Italy (2014).
[4] Moravec, H. P. and Elfes, A., "High resolution maps from wide angle sonar," Proc. IEEE Robotics and Automation Conf., 116-120, St. Louis, MO, U.S.A. (1985).
[5] Baylog, J. G. and Wettergren, T., "A ROC-based approach for developing optimal search strategies in UUV path planning," IEEE J. Ocean. Eng., 43(4), 843-855 (2018).
[6] Stack, J., "Automation for underwater mine recognition: current trends and future strategy," Proc. SPIE 8017, 1-21 (2011).
[7] Duda, R. O., Hart, P. E. and Stork, D. G., [Pattern Classification], 2nd ed., Wiley, New York, U.S.A. (2001).
[8] LaValle, S. M., [Planning algorithms], Cambridge University Press, U.K. (2006). http://planning.cs.uiuc.edu/