Damn the Torpedoes or *Let's Wait and See*?

Decisive Inference for Making the Right Choice

Robert J. (Bob) Jannarone and John T. (Tyler) Tatum
Brainlike Surveillance, Inc.

C. Michael Traweek
Office of Naval Research

Thomas Wettergren
Naval Undersea Warfare Center

ABSTRACT

When making critical combat decisions, commanders must often decide to either act at once or hold off and get more data. Immediate action may offer tactical advantages and improve success prospects, but it may also lead to heavy losses. Getting more data may improve situational awareness and avoid heavy losses, but resulting delays may cause other problems. Making the right choice depends strongly on knowing how much could be gained from gathering more data and how much could be lost by delaying action. Seeing what to expect will help decide if getting more information is worthwhile. An example is presented where a commander must decide to either launch an amphibious landing at once or hold off, pending more mine detection data. Results show that forecasting situational awareness may add substantial command and control value. This paper offers a data-based inference system that is intended to help the commander make the right choice. The system is presented in a minefield detection setting. The system allows the commander to see how the operational picture might change with the arrival of more data. The system produces "decisive inference" by effectively updating both situational awareness and choice tradeoffs as quickly as new information arrives. The system is based on an "auto-adaptive" data analysis process and an "optimal value" choice model. An analytical framework for the results is included, based on a decision model for optimal choice value, along with a statistical basis for auto-adaptive information processing.

INTRODUCTION

This paper offers an inference system for helping a commander make the right choice during time-critical operations. The paper focuses on situations where information that could alter the commander's awareness arrives very quickly. The paper focuses further on deciding whether or not to launch an invasion through a suspected minefield, based on anticipating what kind of mine location information to expect beforehand.

The results of this paper are made possible by powerful recent developments in littoral warfare. The emergence of unmannned underwater vehicles (UUVs) has inspired related analysis methods (Azimi-Sadjadi, Yao, Huang, & Dobeck, 2000; Harris, Avera, & Bibee, 2002; Yao, Azimi-Sadjadi, Jamshidi, & Dobeck, 2002; Wernli, 2001; Bachkosky, Brancati, Conley, 2000), some of which allow operational awareness to improve as quickly as UUV data are gathered (Jannarone & Tatum, 2005; Jannarone, 1997). Due to these developments, a commander can now anticipate how long it will take to obtain mine detection information and what kind of situational awareness improvements to expect throughout the mine detection process.

Beyond UUV-based minefield detection, the system offers added value in a wide variety of applications, military and otherwise, due to its auto-adaptive, decision-based nature. Closely related systems have recently been applied to dynamically optimizing distributed sensor fields (Wettergren, 2006), and a related process has also been validated for magnetometer-based submarine detection (Jannarone & Tatum, 2006). Efforts are also underway to validate and deploy the system for sonar-based submarine detection and camera-based image processing. In all such applications, the system adds value by clarifying situational awareness as conditions change, within cluttered field environments. The system produces "decisive inference" by effectively updating both situational awareness and choice tradeoffs as quickly as new information arrives.

The remainder of this paper describes two key system components within a minefield detection context and provides a detailed example. System components include an "auto-

adaptive" data analysis process and an optimal value (OpVal) choice model, as described in the next two sections.

AUTO-ADAPTIVE PROCESS

Suppose that a task force commander has been ordered to launch a landing within 48 hours, and has been further ordered to consider only two options. The first option (A) would require heading to a beachhead 48 hours away. The second option (B) would require heading to a potentially more hazardous beachhead in the opposite direction, only 24 hours away. Under option A, the commander would immediately launch an attack upon arrival. Under option B, the commander would take time to detect mines prior to launching the landing after arrival. The commander must either take the A path that is known to be less hazardous in the absence of additional information, or take the B path and hope that 24 hours worth of additional information will make it less hazardous.

Insights about information that might unfold under option B could come from prior results, such as a recently completed empirical analysis of minefield data (Jannarone & Tatum, 2005). Such results could imply that initial mine detection operations along the Option B path could produce a certain number of hits if no mines were present, but a much larger number of hits if a minefield was present. Hit expectations given a minefield would exceed hit expectations with no mines present to the point that minefield presence or absence could be evaluated after the first hour of mine detection operations (see the Detailed Example section below). Once an initial minefield presence determination was made, follow-up mine detection operations could clarify mine locations, as shown in Figure 1.

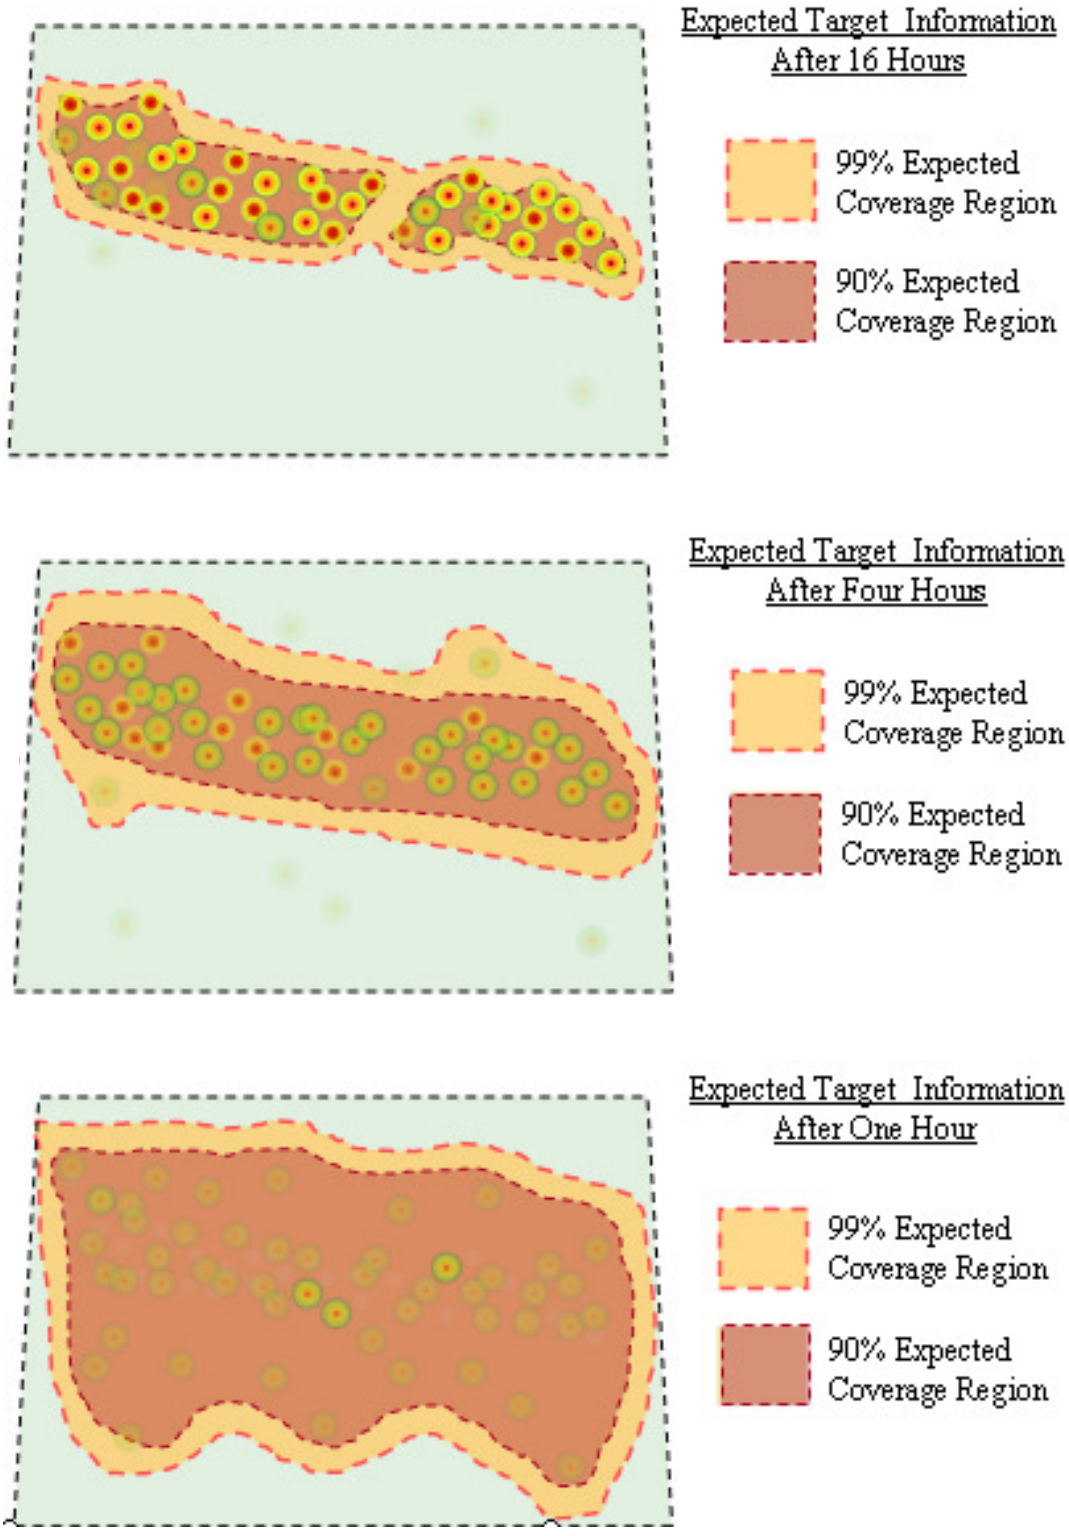

Figure 1. Mine detection likelihoods and Confidence Regions After One, Four, and 16 Hours

More specifically and as formulated below, each of the three graphs in Figure 1 represents cells that are organized into a grid on the suspected minefield region. Each cell represents a small volume of water below it. The circles in the graphs are based on estimates from an auto-adaptive process. The process first computes actual and expected values within each cell for several features, which are computed as functions of observed sonar echo strengths in and around the cell. The process then computes a global deviance value for each cell, based on the differences between expected and actual feature values for the cell. Actual feature values are computed as functions of echo strengths from sonar readings. Expected feature values are estimated from actual feature values in nearest neighbor cells. Global deviance values are standardized so that they have an expected value of zero and a standard deviation of 1.

The process achieves auto-adaptivity within UUV passes by continuously adjusting feature function estimation parameters, as well as global deviance standardizing values, as quickly as new cell data arrive. As a result, the process continuously adjusts for changing clutter and background conditions in the region. The process also achieves auto-adaptivity between passes by combining features from new passes with features from previous passes. For example, if four features are determined per cell in one pass, the process utilizes 16 features in four passes in order to compute deviance values with more precision.

Deviance values computed in this way are monotically related to likelihood ratios based on reasonable target versus no target data assumptions such as those given below. For example, if deviance values in each cell were independently and identically distributed with a mean of zero with no mine in the cell but with a higher mean with a mine in the cell, then high deviance values would correspond to a high likelihood ratio while low deviance values would not.

In the presence of historical data containing specific targets, functions for computing feature values are determined analytically. For the analysis that led to the Figure 1

results, four features were identified for which the auto-adaptive process produced relatively high deviance values when targets were present. In the absence of such historical data, the process computes deviance values directly from raw sensor data. Anomalies that are uncovered in that case simply reflect unexpected changes that are worthy of generating alerts.

Along with individual mine likelihood circles, each Figure 1 graph shows two expected mine coverage regions with shaded, dashed borders. The coverage regions bounded by long and short dashes are constructed so that the expected percentages of mines within them are 99% and 90%, respectively. While the circles in Figure 1 are based on analytical results from the auto-adaptive process, the confidence regions are based on work in progress. Elliptical confidence regions for each graph could easily be computed analytically, by first estimating deviance value mean and variance values along dimension, along with a deviance covariance value. Based on those estimates, an elliptical bivariate normal distribution profile could be computed that covered 90% or 99% of the detected mines. Since mines cannot be expected to fall within elliptical regions, however, more general procedures based on fitting higher-degree profiles to deviance values are being employed.

As Figure 1 shows, repeated UUV scans can be expected to produce sharper mine presence images. In the bottom graph, mine presence likelihoods are low relative to background levels, and mine coverage regions are large relative to the two top graphs. In the middle graph, information that has been added to the bottom graph has improved clarity. Many circles having low likelihoods from the bottom graph have disappeared, many others have become brighter, and the coverage regions have become smaller. In the top graph, nearly all circles are either very bright or very dull, substantially clarifying where mines are present and absent, to the point that a prospective channel through the minefield has emerged.

The bottom graph shows that mines are likely to be present, because target hits have significantly exceeded chance levels (see Detailed Example section below) and they

systematically distributed. However, the figure gives only a rough sense of where the mines are concentrated. After four hours, the task force could proceed much more safely through more water and concentrate mine countermeasures on a narrower region. After 16 hours, the minefield would be tightly identified, to the point that an open channel has emerged.

In order to be effective, the auto-adaptive process must produce event recognition signals such as those in Figure 1, in real time. The process does so by efficiently receiving information, producing event detection signals, updating learned information about background clutter, and correcting for background clutter. The process also refines target likelihoods and related displays immediately upon new data arrival. For the Figure 1 example, each new pass over the minefield region would result in a new set of feature values. The process would combine them with previously measured features values as described below to refine mine detection likelihoods and Figure 1 graphs based on them.

The process evaluates mine detection data as UUVs move along. It produces alerts like the circles shown in Figure 1 when within-cell intensity values are unexpected, relative to recently learned expectations, all in real time. In this way, the process can produce alert signals that are robustly accurate against changing background noise, which it continuously learns to expect and remove.

The auto-adaptive process is designed for a variety of applications, other than UUV-based mine detection. The process can operate with packets of information in spatial cells along one, two, or three dimensions, in distinct time slices. At the beginning of each time slice it receives a configurable number of feature values within each cell, and at the end of each time slice it produces a target likelihood value for each cell. Resulting likelihood values can be used to produce alerts, generate intensity plots such as those in Figure 1, or prioritize subregions for further examination.

A precursor to the auto-adaptive process was developed for broad use in monitoring, forecasting, and control (Jannarone, 1997). That process has since been used

commercially for electricity demand and price forecasting (Jannarone, 2003) as well as application performance monitoring (Netuitive, 2006). The new process described herein has been designed to produce anomaly signals very quickly and compactly on spatially registered grids such as Figure 1, (Jannarone, Tatum, & Gibson, 2006). By contrast with its precursor, the new process employs nearest neighbor estimation and staging toward that end, as described in the remainder of this section. New process effectiveness has been demonstrated by recent results (Smith & Jannarone, 2004, Jannarone & Tatum, 2005, Jannarone & Tatum, 2006).

During each time slice, the process is designed to operate in pipelined stages, which in turn are designed for implementation on fast, compact field programmable gating arrays (FPGAs). Each such stage operates by passing window functions over each spatial cell at each time slice. Window functions may be partitioned within each stage and processed in parallel if speed and memory requirements so dictate.

The following formulation covers the one-dimensional case, and values in brackets cover the Option B case in the example. Unless otherwise stated, all bold symbols below represent row vectors. Assume that a UUV is producing one set of echo intensity values above and below itself during each time slice as it moves from left to right in the minefield region (see Detailed Example section). Let $t = 1, \ldots, t^{[MAX]}$ represent time slices as the UUV moves from left to right. Let $c = 1, \ldots, c^{[MAX]}$ represent the cells above and below the UUV for which the UUV obtained an intensity value at that time slice. Let $f = 1, \ldots, f^{[MAX]}$ represent features that are being measured within each cell. Let x_{tcf} represent the measured UUV value on feature f within cell c at time slice t. The auto-adaptive process is partly based on the assumption that

$$x_{tcf} = \alpha_f + \beta_{tcf} + e_{tcf}, \quad (1)$$

where α_f is an anomaly value for feature f that would only be nonzero in the presence of a target, β_{tcf} is a continuously updated expected background value, and e_{tcf} is an

independent and identically distributed error value with a mean of zero. The process is also based on the assumption that

$$\beta_{tcf} = \mu_{t-1,cf} + \sum_{u=t-t^{[\mathrm{WINDOW}]}}^{t} \sum_{d=c-c^{[\mathrm{WINDOW}]}}^{c+c^{[\mathrm{WINDOW}]}} \sum_{g=1}^{f^{[\mathrm{MAX}]}} \rho_{t-1,udgcf} x_{udg} , \quad (2)$$

where $\mu_{t-1,cf}$ and $\rho_{t-1,udgcf}$ are background mean and regression weight estimates, respectively, which have been updated just prior to time point t for estimating x_{tcf} values from the recent and nearest neighbor x_{udg} values, and e_{tcf} is an error value. Thus, each feature estimate within each cell is formulated as a function of other current feature values in the cell, other current values in the $2\,c^{[\mathrm{WINDOW}]}$ nearest neighbor cells above and below it, and recent values in corresponding cells that were observed from one time slice to $t^{[\mathrm{WINDOW}]}$ time slices ago.

In order to produce background β_{tcf} estimates for cells that lie within $c^{[\mathrm{WINDOW}]}$ cells of boundary cells, the process sets $x_{tcf} = 0$ for $c' = c^{+}+1, \ldots, +c^{[\mathrm{WINDOW}]}$, and for $c' = -c^{[\mathrm{WINDOW}]}, \ldots, 0$, respectively. Also, in order to exclude estimating each x_{tcf} as a function of itself, the process sets $\rho_{t-1,tcfcf} = 0$. Proceeding in this way allows the process to produce reasonable estimates at grid boundaries, while avoiding the time consuming use of special logic to treat boundary cell estimation differently.

The auto-adaptive process continuously updates and adjusts estimates for background clutter, based on learning functions with the following properties (Jannarone, 1997). Process learning functions produce weighted least squares estimates based on prior observed values. As a result, their interpretation and computation are straightforward. Process learning functions also give more recent observed values more weight that less recently observed values. As a result, they can quickly reflect and compensate for changing background conditions. Process learning functions use recursive functions that quickly update estimates at each time point as a function of only observed values at each time point and values that were updated during the immediately preceding time point. As

a result, the process can update learned parameters very quickly. Process learning functions also require that only learned parameter values from the last time point and observed values from the last few time points be retained. As a result, Process modules can be sufficiently small to reside on off-board sensing platforms.

For $t = 1, \ldots,$ estimation formulas for background noise means have the recursive form,

$$\mu_{tcf} = (l_t \, x_{tcf} + \mu_{t-1,cf}) / (1 + l_t) \tag{4}$$

where l_t is a positive learning weight that is computed in a way that satisfies the above properties, and μ_{0cf} is either set to 0 or estimated from the first few time slices. Learning weights are formulated so that learning occurs in blocks, where the most recent block produces half of the overall prior impact in a weighted least squares sense, the next most recent block produces half of the remaining impact, and so on (Jannarone, 1997). More comprehensive formulations that deal with missing and deviant values are provided elsewhere (Smith & Jannarone, 2004, Jannarone, 1997).

Regression weights are updated as functions of recursively updated matrices, which may take on different forms that depend on different process models. For some models regression weights are updated by updating inverses of mean squares and cross-products (MSCPs) among current and recent values. Such matrices denoted by

$$\boldsymbol{\omega}_t^{[\mathrm{MSCP}]}, \tag{5}$$

are of order

$$m = t^{[\mathrm{WINDOW}]} \times (2\, c^{[\mathrm{WINDOW}]} + 1) \times f^{[\mathrm{MAX}]},$$

and are based on all nearest neighbor time points, features, and nearest neighbor cells that would be needed to estimate background effects for any given feature, cell, and time point value. In the $f^{[\mathrm{MAX}]} = 4$ case, for example, $t^{[\mathrm{WINDOW}]} = 2$, and $c^{[\mathrm{WINDOW}]} = 2$ would result in a fast and compact process based on $m = 40$.

Depending on the specific application, the process may update the $\boldsymbol{\omega}_t^{[\mathrm{MSCP}]}$ matrix each time the window function passes over each cell (e.g., 1,000 times per time slice if implemented on each UUV in the Option B case]. In that case, for $t = 1, \ldots,$ the updating formula would have the form (Jannarone, 1997)

$$\boldsymbol{\omega}_t^{[\mathrm{MSCP}]} = (1+l_t)[\,\boldsymbol{\omega}_t^{[\mathrm{MSCP}]} + l_t\,\boldsymbol{r}^{\mathbf{T}}\,\boldsymbol{r} \,/\, (1+l_t\,d)], \tag{6}$$

where

$$\boldsymbol{r} = \boldsymbol{x}_{tc}\,\boldsymbol{\omega}_{t\text{-}1}^{[\mathrm{MSCP}]}, \tag{7}$$

the $\mathbf{T}$ superscript denotes transposition, $\boldsymbol{x}_{tc}$ is a row vector containing all feature values necessary for estimating cell c values at time t,

$$d = \boldsymbol{x}_{tc}\,\boldsymbol{r}^{\mathbf{T}}, \tag{8}$$

and $\boldsymbol{\omega}_0^{[\mathrm{MSCP}]} = \mathbf{I}$.

For example, each $\boldsymbol{x}_{tc}$ would be made up of 40 values, including the 4 feature values in each of 10 cells: cell c, the two cells to its left, the two cells to the right, and the five corresponding cells that were observed one time slice ago.

Alternatively, the process may compile cumulative MSCP matrices of the form,

$$\boldsymbol{\delta}_t = \sum_{c=1}^{c^{[\mathrm{MAX}]}} \boldsymbol{x}_{tc}^{\mathbf{T}}\boldsymbol{x}_{tc} \tag{9}$$

as an updating window function passes over all cells during each time slice, and the process may update MSCP matrices by computing,

$$\boldsymbol{\sigma}_t^{[\mathrm{MSCP}]} = (\,l_t\,\boldsymbol{\delta}_t + \boldsymbol{\sigma}_{t-1}^{[\mathrm{MSCP}]}\,) \,/\, (1+l_t), \tag{10}$$

and then computing

(11)

$$\boldsymbol{\omega}_t^{[\mathrm{MSCP}]} = \boldsymbol{\sigma}_t^{[\mathrm{MSCP}]-1}.$$

In either case, once $\boldsymbol{\omega}_t^{[\mathrm{MSCP}]}$ has been computed regression weights may be obtained by computing

$$\boldsymbol{\rho}_t = \boldsymbol{\omega}_t\ [\mathbf{I}-\mathbf{D}^{-1}(\boldsymbol{\omega}_t)]. \tag{12}$$

where $\boldsymbol{\omega}_t$ is the inverse of the variance-covariance matrix, $\mathbf{D}$ is the diagonal matrix function,

$$\boldsymbol{\omega}_t = \boldsymbol{\omega}_t^{[\mathrm{MSCP}]} + \boldsymbol{s}^{\mathbf{T}}\boldsymbol{s}/\ (1+d). \tag{13}$$

Process models may be extended to include two and three spatial dimensions in a similar way.

Based on the above model, feature deviance values within each cell may be obtained by computing,

$$z_{tcf} = [(x_{tcf} - \beta_{tcf}) - \mu_{t-1,tcf}^{[\mathrm{Z}]}]/\sigma_{t-1,tcf}^{[\mathrm{Z}]}, \tag{14}$$

where $\mu_{t-1,tcf}^{[\mathrm{Z}]}$ and $\sigma_{t-1,tcf}^{[\mathrm{Z}]}$ are learned deviance value means and standard deviations, respectively, which are updated continuously. As a result, expected deviance values are zero if no mines are present but α_f if mines are present, resulting in a monotone relationship between z_{tcf} values and likelihood ratio values for any nonzero α_f value.

Based on the above model, feature deviance values within each cell may be obtained by computing,

Provisions have been included in the process for dealing with numerical problems such as linear redundancy, ill conditioning, and with highly deviant or missing input values. Linear redundancy and ill conditioning problems are precluded by initializing the MSCP matrix and its inverse to the identity matrix, continuously checking their main diagonal elements, and either removing offending features or readjusting the matrices as problems

arise. Whenever highly deviant or missing input values occur, they are replaced with estimated values, so that the learning process is not effected. As an added benefit, the process imputes missing values in ways that are useful for pattern completion.

OPTIMAL VALUE (OPVAL) CHOICE MODEL

Operationally, the decision-maker is continually faced with comparing the likelihood of successful gains from searching further against the utility cost of making that other search. In this context, utility cost includes items such as actual search cost along with time lost to searching, risk of missing mines if search isn't performed, and so on. Experienced operational commanders make these tradeoffs intuitively. However, situational awareness gained from an analytical representation of such expected gains and costs can provide useful information to facilitate the decision process. This insight is representative of the classical econometric approach of using utility costs to inform automated decision making. In the Figure 1 context, the likelihood of mine locations (represented by the deviance values from the auto-adaptive process) is typical of the information examined by an operator, whereas the analytical representation of risks and gains is typical of the information examined by a commander. The auto-adaptive process creates the operator view of mine likelihoods for each pass of the mine searcher (or simulated results from a planned pass). The OpVal model then takes this operator view and translates the mine likelihoods into measures of uncertainty in minefield existence, and creates an estimate of the improvement or degradation of that uncertainty that can be expected from future searches. This is the information that the commander may then weigh against the cost of performing successive searches.

The OpVal choice model first converts expected and measured detection, false alarm and missed target frequencies to a future search value function, in a way that allows field modifications of choice models on the spot as tactical situations change. The model could not conceivably automate command decisions in general, because such decisions are far too complex to model comprehensively. However, the model can point out relatively simple tradeoffs, so that commanders will be able to estimate the expected

utility of gathering more information. In the Figure 1 case comparable cost metrics and likelihoods could easily be assigned to false alarm rates, missed target rates, and their associated costs (see next section). Comparable UUV deployment cost metrics could be assigned as well, giving commanders a better understanding of choices they are about to make.

The focal decision within the Figure 1 context is whether it is worth spending the extra time to make another pass in the search for mines. Decision processes may operate at either the minefield level, which is the focus of the current OpVal model discussion, or the individual mine level, which would operate similarly. At the minefield level, the decision process is based on potentially employing further searches in order to reduce the uncertainty in the simple binary hypothesis of H_1: there is a minefield obstructing the path, and H_0: there is no minefield obstructing the path. The probability of the hypothesis decision $p = P(H)$ is the probability that there is a minefield obstruction, so that $p = 0$ and $p = 1$ represent the two hypothesis choices. Generally, values of p will fall within this range, and thus the decision to take another search is based on whether it is expected that another search will significantly push the value towards one of those limits. The decision process for the commander is to weigh the expected value (given as the reduction in uncertainty of the minefield/no-minefield hypothesis test) of the complete future sequence of potential searches against the cost of such a search strategy. The OpVal model represents binary decision (search again vs. no search) and the binary hypotheses as a two-state test T with values t_1 (looks like a minefield) and t_2 (does not look like a minefield) as shown in the following:

	$T = t_1$	$T = t_2$
H_1 (mines)	$P(H = H_1 \mid t_1)$	$P(H = H_1 \mid t_2)$
H_0 (no mines)	$P(H = H_0 \mid t_1)$	$P(H = H_0 \mid t_2)$

The values associated with each entry in this table represent the value of a given test result under each of the states of the system hypothesis. To the commander deciding whether or not there is a minefield, the important value is the uncertainty with which the

decision can be made given a test result. The OpVal model represents these payoffs as conditional probabilities, where $P(H = H_1 \mid t_1)$ represents the probability that there is a minefield given that the search showed what appears to be a minefield. As successive iterations are made, these probabilities will further depend on the prior assumption with regard to the minefield state. These conditional probabilities are data (and/or simulation) derived and thus provide a mechanism for adaptation in the decision process. This process presents a decision form that Jensen refers to as a *myopic hypothesis driven data request* (Jenson, 2001). The term myopic refers to the fact that, at each decision step, a choice is made based on maximizing the profit (where profit = benefit − cost) of performing that individual step, with no regard for the potential benefit of a future set of steps. In other words, a search is made if the expected benefit is greater than the cost.

The myopic hypothesis driven data request is solved by examining the expected profit of performing the next step in the search, $EP(S)$, which is given by

$$EP(S) = EV(S) - V(P(H)) - C_S \tag{15}$$

where $EV(S)$ is the expected value of the search, $V(P(H))$ is the value of the stated hypothesis, and C_S is the cost associated with the search. If this profit is positive, then a rational decision maker would perform the search. Rather than explicitly computing the expected profit, the OpVal model presents the decision maker with the expected benefit of a next search, $EB(S) = EP(S) - V(P(H))$, which may then be weighed against the known costs. In this manner, the model eliminates the need to express both costs and benefits in compatible units, which is a subjective matter that is better performed by skilled decision makers than by a computer.

To compute the expected benefit, the value of the added search must be quantified with respect to the simple binary hypothesis. Added value may be expressed as a reduction in the variance (or uncertainty) of the value of the hypothesis statistic. In other words, the commander is interested in performing the search to reduce uncertainty in regards to the existence of an obstructing minefield. For a hypothesis H with states {h}, the variance value function is given by

(16)

$$V(P(H)) = -\sum \left(h - \sum hP(h)\right)^2 P(h)$$

where both sums are taken over all of the potential hypothesis states. For the minefield decision problem with the simple binary hypothesis, the value function for variance reduces to

(17)

$$V(P(H)) = p(p-1)$$

where p is the probability of the H_1 hypothesis, which is assumed to be complementary to the H_0 hypothesis. This expression reflects a variance reduction rather than a variance value. Thus, in either case, a decrease in variance is sought. The choice of a variance value function is specific to the intended mine hunting application, however, the general approach is applicable to any other value functions.

The only other term to compute in the expected benefit of an additional search is the expected value of the search, $EV(S)$, which is found by summing the values of the conditional search outcomes over all of the potential search results. For the binary test case (test T with results t_1 or t_2) of the mine hunting example, we have

(18)

$$EV(S) = V(P(H \mid t_1)) \cdot P(t_1) + V(P(H \mid t_2)) \cdot P(t_2)$$

Using the variance value function of equation (17), the value of an added search under the condition of a given test (search) result t_j is given by

(19)

$$V(P(H \mid t_j)) = -p_{1j}^2 p_{2j} - p_{1j}(1-p_{1j})^2 ,$$

where $p_{ij} = P(H = H_i, t_j)$. Since the values of p_{1j} and p_{2j} are necessarily complementary, the following expression holds for the expected benefit of an added search:

(20)

$$EB(S) = p_0(1-p_0) - P(t_1)p_{11}(1-p_{11}) - P(t_2)p_{22}(1-p_{22})$$

For this simple search test, the likely outcome of the test is given by the assumed hypothesis value going into that decision step. Thus, given 90% prior certainty that there is a minefield, that is, $p_0 = 90\%$, the probability of a positive search outcome will be $P(t_1) = p_0$, and the probability of a negative search outcome will be $P(t_2) = 1\text{-}\, p_0$. This leads to the resulting value of

(21)

$$EB(S) = p_0(1-p_0) - p_0 p_{11}(1-p_{11}) - (1-p_0)p_{22}(1-p_{22})$$

If the search performance is not well-approximated by the prior decision p_0, then it is appropriate to assume that the test results are uniformly likely, leading to

(22)

$$EB(S) = p_0(1-p_0) - \tfrac{1}{2}p_{11}(1-p_{11}) - \tfrac{1}{2}p_{22}(1-p_{22})$$

Equation (21) represents the reduction in uncertainty that can be anticipated by taking a new search. For other value function choices (other than the variance value function of equation (17)), the expected benefit of equation (21) has a different interpretation, but is still a predicted benefit. If the expected benefit of equation (21) is viewed as cost-effective, then the next search is a rational decision choice. Given values of p_{11} and p_{22}, the benefits of added searches are readily evaluated.

The values of p_{11} may be obtained from preceding experience with the data. Recall that p_{11} represents the probability that there is an obstructing minefield given that we have a test result showing a minefield. Therefore, if the value of certainty that there is a minefield increases x-fold for each new search (*i.e.*, if the expected "resolution" increases by a factor of x), then a model of p_{11} is given by

(23)

$$p_{11} = 1 - \frac{(1-p_0)}{x}$$

Similarly, if the expected ability to remove false alarms improves by a factor of y for each pass, then

$$p_{22} = \frac{p_0}{y} \tag{24}$$

Finally, the new value of p_0 for the next step is given by weighting the result of the test with the appropriate prior probability,

$$p_0 = \sum hP(h \mid t_j) \tag{25}$$

for test result t_j, so that

$$p_0 = \begin{cases} p_{11} & T = t_1 \\ 1 - p_{22} & T = t_2 \end{cases}. \tag{26}$$

Assuming that the values of x and y are given by $x = y = 4$ (for example), then the decision process could proceed as follows (assuming further that the test results are based on prior data):

Step 0: p0 = 0.5 → EB(S) = 0.1406 → search (sees mines, t1) → p0 = 0.88
Step 1: p0 = 0.88→ EB(S) = 0.0615 → search (sees mines, t1) → p0 = 0.97
Step 2: p0 = 0.97→ EB(S) = 0.0173 → search (sees mines, t1) → p0 = 0.99
Step 3: p0 = 0.99→ EB(S) = 0.0044 → don't search (low benefit)

The sequence shows that the corresponding reduction in uncertainty in the estimate of p0 (probability of minefield) has lowered to a point where the expected benefit of another search (even if it refines the field-view) is not enough to merit another search. If same sequence of steps is performed, but assume the first search doesn't yield a positive field-level search result, then the resulting process is as follows:

Step 0: p0 = 0.5 → EB(S) = 0.1406 → search (misses mines, t2) → p0 = 0.13

Step 1: p0 = 0.13→ EB(S) = 0.0615 → search (sees mines, t1) → p0 = 0.78

Step 2: p0 = 0.78→ EB(S) = 0.0961 → search (sees mines, t1) → p0 = 0.95

Step 3: p0 = 0.95→ EB(S) = 0.0291 → search (sees mines, t1) → p0 = 0.99

Step 4: p0 = 0.99→ EB(S) = 0.0076 → don't search (low benefit)

In this case the search value after steps 0 and 1 is still large because of two conflicting search results. The expected benefit (in terms of reduction of uncertainty) drops significantly to a point where no more searches are desired, only after a sequence of multiple non-conflicting search results have occurred.

Seeing these expected benefits allows the commander to assess the relative merit of a new search, relative to prior benefits. In particular, the first search has a benefit of 0.1406. After scaling the benefit numbers by this value, then the sequence becomes {1.00, 0.44, 0.12, 0.03} for the first step and {1.00, 0.44, 0.68, 0.21, 0.05} for the second step. Searching would stop when the expected benefit from its corresponding step dropped below 10% of the initial benefit of a search, which is the assumed cost limit for this example.

The process becomes more complicated for changed values of x and y. Assume the ability to distinguish a minefield increases fourfold with a successful search when one exists ($x = 4$). Furthermore, assume the ability to distinguish the non-existence of a minefield (ie: clearing out false alarms) only increases two-fold after a successful test ($y = 2$). Then the expected search benefit sequences for cases with all successful searches (seeing minefield each time) and success after a single failed search are {1.00, 0.51, 0.14, 0.04} and {1.00, 0.66, 0.70, 0.22, 0.06}, respectively. These values are consistently higher than the previous cases, showing that searching is more useful when clearing false alarms conflicts with enhancing detections. The model is readily updated at each step by replacing p_0 by a better estimate obtained from the actual search operation.

A comprehensive OpVal choice model that covers a variety of other cost factors and has been implemented in spreadsheet form is available from the authors upon request.

DETAILED EXAMPLE

If the two UUVs were traveling from left to right and each UUV was pinging cells to its left and its right repeatedly as it went along, each such ping would represent a time slice. An auto-adaptive process residing on the UUV would display a circle like those shown in the figure if and only alert values computed by the UUV exceeded a criterion. The criterion would automatically compensate for changing background expectations along the way.

Suppose that two UUVs are moving from left to right within the region shown in Figure 1 at a rate of 6 km/hr. Suppose further that every minute each UUV is sensing 1,000 cell slices from left to right, along a swath that is 500 meters wide (250 meters above and 250 meters below the UUV). In that case, each cell would be 0.1 meters wide ([0.1 meters/ cell] = [6,000 meters/hour)]/[60 minutes/hour]/[1000 cells/minute]). If each ping on either side produced 500 equally spaced intensity values (i.e., $c^{[MAX]} = 1{,}000$) as reported in [1], then each corresponding cell would be 0.5 meters high (0.5 meters = [250 meters/ ping]/ [500 cells/ping]). Also, the minefield (assuming it to be rectangular) would be 30,000 cells wide (i.e., $t^{[MAX]} = 30{,}000$) and 4,000 cells high. Then two UUVs proceeding 500 meters apart would be able to cover the top half of the field in 30 minutes and the entire field one hour, as in the above example. Suppose further that the auto-adaptive process computes four features per cell (i.e., $f^{[MAX]} = 4$), each of which has been formulated to have a high value if the cell contains a certain type of target and a small value otherwise, while compensating for background (Jannarone & Tatum, 2005).

Regarding process stage operation in the Figure 1 case, the first stage would compute a number of feature values for each cell as simple functions of nearest neighbor window values. For example, four feature functions have been derived with the Figure 1 application in mind and reported previously (Jannarone & Tatum, 2005). The second stage would compute background noise estimation mean values for each stage; the third

stage would compute regression weights; and the fourth stage would compute alert values based on estimated anomaly values.

These results could have been obtained at the rate of a few second per thousand time slices on a conventional processor, easily keeping up with UUV data arrival rates. For on-board UUV implementation, having only limited capacity microprocessors available, the process would be implemented on FPGAs operating independently of and much faster than central UUV microprocessors.

All such formulations are based on reasonable error models, which can be used to produce straightforward expected values based on empirical data. For example, empirical results from the previous study were used to compute expected values that resulted in the bottom Figure 1 graph. Those results were then extended to generate expected values given more data, as shown in top two graphs, based on the following, well known sampling theory result: if results from independent samples are averaged and the process assumptions are satisfied, then four times as much sampling will double the accuracy, 16 times as much sampling will increase accuracy by a factor of 4, and so on.

Referring to Figure 2, the three two by two tables on the left side of the figure contain expected hit and false alarm frequencies after one hour, three hours later, and 12 hours after that, under Option B. Each corresponding table on the right side gives corresponding expected costs, obtained by multiplying each missed target and false alarm cost per entry that is provided in the top right figure by its corresponding cell size. The numbers in the top right figure reflect much larger costs for not marking a cell that contains a mine than for marking a cell that doesn't contain a mine, because hitting a mine unexpectedly could be disastrous while incorrectly concluding that a mine is present will only cause mine avoidance or mine clearing delays. These estimated costs per entry could be modified in the field to reflect different situations, allowing them to estimate just how much expected missed target and false alarm costs would vary under different data gathering scenarios. They could also be combined with data gathering costs, lost

time costs, and other costs as well, in order to determine which data collecting choice would be best under Option B or even whether to choose Option A over Option B.

The global frequency of 120,000 cells in the left tables corresponds to the number of cells in the example (30,000 time slices per pass times from left to right time 4,000 rows), while the total number of targets (48) is based on the density of actual targets in the dataset from the original study (Jannarone & Tatum, 2005). The hit and false alarm rates in the lower table were based on those found in the same study. Hit and false alarm rates in the top two tables were based on assuming missed targets and false alarms would reduce as the square root of the number of passes, in keeping with the same independent sampling rationale that was outlined earlier.

As mentioned earlier, the expected coverage regions shown in Figure 1 represent work in progress. The regions may be constructed by fitting a non-negative function of the two region coordinates to observed deviance values. The function would be one of a class of functions, including the bivariate normal distribution density, that are monotone with respect to a modal ridge along the trend curve that best fits the data. In the bivariate normal case, the curve would be a straight line. In the figure 1 case, the curve would go from the top left corner of the region the center right portion, in a slightly nonlinear way Expected coverage regions would then be constructed that covered all hit cells, along with areas outside the cells where hits have not been identified, based on prior information about target hit rates. For example, the regions within the bottom Figure 1 graph cover much larger areas than their counterparts within the top graph, because the hit rates for the one-hour case are much lower than the hit rates for the four-hour case, as shown in Figure 2.

The OpVal model takes these mine likelihood regions as shown in figure 1 and convolves them with the expected characteristics of a notional minefield. For instance, a uniformly distributed minefield of a certain spacing may be represented by a set of Gaussian "bumps" separated by such spacing. This convolved geometric map is then converted to a single number by simple geometric averaging and this number may provide the

posterior value of minefield likelihood that is used as the prior (p0 in the OpVal notation of equation (25)) for the next search decision. Similarly, the values of x and y in the OpVal equations (23) and (24) are updated by examining a prior history of searches with similar mine hunting technology or by examining recent results in the search operation under consideration. In this way the OpVal decision sequence allows the commander to weight expected future search benefit against future search cost.

EXPECTED COST BASIS

Mine Actuality:	UUV Results: Marked	Unmarked
Absent	1,000	0
Present	0	1,000,000

EXPECTED HITS AND FALSE ALARMS

	Mine Actuality:	UUV Results: Marked	Unmarked	Totals
	Absent	20	119,999,932	119,999,952
AFTER ONE HOUR	Present	32	16	48
	Totals:	52	119,999,948	120,000,000

	Mine Actuality:	UUV Results: Marked	Unmarked	Totals
	Absent	10	119,999,942	119,999,952
AFTER FOUR HOURS	Present	40	8	48
	Totals:	50	119,999,950	120,000,000

	Mine Actuality:	UUV Results: Marked	Unmarked	Totals
	Absent	5	119,999,947	119,999,952
AFTER 16 HOURS	Present	44	4	48
	Totals:	49	119,999,951	120,000,000

EXPECTED COSTS

Mine Actuality:	UUV Results: Marked	Unmarked	Totals
Absent	20,000	0	20,000
Present	0	16,000,000	16,000,000
Totals:	20,000	16,000,000	**16,020,000**

Mine Actuality:	UUV Results: Marked	Unmarked	Totals
Absent	10,000	0	10,000
Present	0	8,000,000	8,000,000
Totals:	10,000	8,000,000	**8,010,000**

Mine Actuality:	UUV Results: Marked	Unmarked	Totals
Absent	5,000	0	5,000
Present	0	4,000,000	4,000,000
Totals:	5,000	4,000,000	**4,005,000**

Figure 2. Option B Cost Analysis

PROGRESS TO DATE AND EFFORTS REMAINING

Many data-based results have already been obtained that could support the inference system described in this paper (Jannarone & Tatum, 2005; Smith & Jannarone, 2004; Jannarone, 1997). However, substantial further mine detection data could be processed in much the same way, in order to provide commanders and information officers more precise pictures with less guesswork. The authors intend to continue analyzing available mine detection data toward that end. The authors also hope that as war fighters begin using results like these and discover their value, commanders will support their further delivery, especially by working with Navy analysts to obtain more data.

CONCLUSIONS

This paper has introduced a system for guiding commanders in making optimal choices, in operations where situational awareness and information availability may change quickly. An example has been provided where a commander must choose between launching a landing at once or taking the time to improve situational awareness. Key to guiding the commander in making such decisions is (a) providing all available data in a form that will show how situational awareness would improve, and (b) being able to improve situational awareness as quickly as new information arrives. A process has been described that has been designed to meet these two aims. The system is based on an optimal value (OpVal) choice model and an auto-adaptive process for target detection. The OpVal model has been designed to weigh costs and benefits of acting at once versus gathering more information, and point toward optimal choices accordingly. The process has been designed to deliver the essential simplicity, speed, compactness, robustness, and affordability to make the process feasible. Process validation awaits processor testing and evaluation, planned to begin in the near future.

REFERENCES

Jannarone, R. J., and Tatum, J. T. 2005. "A Novel Process for Littoral Target Recognition: Preliminary Results." *Proceedings of the Oceans 2005 Conference,* MTS/IEEE.

Azimi-Sadjadi, M. R., Yao, D., Huang, Q., and Dobeck, G. J. 2000. "Underwater Target Classification using Wavelet Packets and Neural Networks. "*IEEE Trans. Neural Networks*, vol. 11, 784–794.

Harris, M., Avera, W. and Bibee, L. 2002. Environmental Data Acquisition from the AQS-20 Mine Hunting Sonar Requirements, Technical Feasibility and Cost. NRL/FR/7440—02-1002.

Yao, D., Azimi-Sadjadi, M. R., Jamshidi, A. R., and Dobeck G. J. 2002. A Study of Effects of Sonar Bandwidth for Underwater Target Classification. *IEEE Journal of Oceanic Engineering*, Vol. 27, No. 3, 619-627.

Wernli, R.L. 2001. "Recent U.S. Navy Underwater Vehicle Projects," 24th Joint Meeting of the United States/Japan Committee on Natural Resources, Marine Facilities Panel, Honolulu.

Bachkosky, J.M., Brancati, T., Conley, D.R. 2000. Unmanned Vehicles (UV) in Mine Countermeasures. Final Report. Naval Research Advisory Committee Report, Arlington, VA.

Smith, R .L., and Jannarone, R.J. 2004. "An Auto-Adaptive Statistical Procedure for Tracking Structural Health Monitoring Data," *Proc. SPIE*, **5391**, 166-176.

Jannarone, R. J., 1997. *Concurrent Learning and Information Processing: A Neurocomputing System that Learns* during *Monitoring, Forecasting, and Control.* Chapman & Hall.

Wettergren, T., "Dynamic Optimization of Distributed Sensor Fields," Office of Naval Research Maritime Sensor Program Review, Newport RI, August, 2006.

Jannarone, R. J., & Tatum, J. T., 2006. "Auto-Adaptive Clutter Removal for Improved Submarine Detection, Final Report under Office of Naval Research Phase I STTR Contract #N00014-05-M-0287, April, 2006.

Jannarone, R. J., 2003. "Electricity Monitoring Improvement," Brainlike Surveillance, Inc., Technical Report.

Netuitive, 2006. "Netuitive Real-Time Analysis," www.Netuitive.com.

Jannarone, R. J., Tatum, J.T., & Gibson, J.A. , 2006. "Efficient Processing in an Auto-Adaptive Network," U.S. and International patents pending.

Jensen, F. V. (2001). *Bayesian Networks and Decision Graphs*, Springer-Verlag.