The Roles of Human Operator and Machine in Decision Aid Strategies for Target Detection

Nancy J.M. Allen and Ronald T. Kessel
Defence R&D Canada – Atlantic
P.O. Box 1012
Dartmouth, Nova Scotia B2Y 3Z7
CANADA

SUMMARY

Automated detection systems are more often disappointing than they are genuinely useful. The automated detection algorithms have borne too much of the blame thus far. Success and failure are re-examined here, at a system level, in light of a simple principle of balance: the responsibilities of each detector, human and machine alike, must be matched to its abilities. Giving the automation too much or too little responsibility makes it ineffective, either jeopardizing the mission on the one hand, or robbing it of purpose and making it an annoyance on the other hand. This notion of balance is substantiated using elementary detection theory. It is used then to diagnose what we call conventional detection aids within a larger framework of human-machine decision support. Three strategies for the development of detection systems are proposed, all based on the principle of balance, with many realistic examples given for illustration. The strategies are not new, but their derivation, now from a single principle, builds a coherent framework for purposeful improvement and design. This work should be of particular use to system developers, operators, and tactical planners.

1.0 INTRODUCTION

As technology advances, one expects to see ever-increasing automation, decreasing human intervention, and corresponding improvements in performance. This has not always been true of military detection systems, particularly when a haphazard, automation-for-automation's-sake approach has been followed for auto-detection, with no larger strategy of system design by which the benefits of automation can be assured. The results have been less than satisfactory. In our experience (submarine and mine hunting), operators generally consider auto-detection aids as an annoyance. Novice operators may experiment with them momentarily, whereas experienced operators are inclined to switch them off, almost as a matter of course. But the scope of system development has been expanding, from a conventional *close-to-the-sensor* technology (sensors and pattern recognition) to a more integrated view of human and machine working together as a single system with shared responsibilities.

In this work we argue that:

- top-down strategies for system design are required if the widespread dissatisfaction with auto-detection is to be overcome;
- following the lead of civilian pattern recognition falls short of military demands, especially in terms of the flexibility required for changing operational factors;

Paper presented at the RTO HFM Symposium on "The Role of Humans in Intelligent and Automated Systems", held in Warsaw, Poland, 7-9 October 2002, and published in RTO-MP-088.

- because it assists in the target/non-target decision, influencing mission progress and expenditure of resources, a detection system should properly be viewed as a decision support system rather than as a mere sensor alarm; and
- the system must adhere to a principle of balance between the responsibilities and abilities for each detector, human and machine alike.

The last is a simple, common-sense principle that has all too often been violated with disappointing results, as we demonstrate below by example. Three design strategies for detection systems are proposed, each aiming for a different balance between responsibilities and abilities.

To be clear, we define target detection to be the uncertain perception of new target-like contacts, where the uncertainties are due to:

- imperfect signal quality owing to limited sensor resolution and sensitivity, signal noise, and adverse environmental conditions;
- target-like and target-masking clutter in the environment;
- target camouflage and stealth; and
- the unpredictable nature of new targets (range, bearing, aspect, number, type, intent, and so forth).

Detection is often the first stage in a process working towards target classification, whose uncertainties likewise hinge on these factors. The progression from detection to target classification is a change in degree rather than kind; of increasingly decisive attention, attitude, and action toward a single contact, in response to progressively increasing certainty [Allen, 2002]. We focus here on the start of that chain.

Automatic detection for military systems entails pattern recognition, which is likewise central to many civilian applications such as medical diagnosis, process control, and security. Research into military systems can therefore draw on the significant body of research for civilian pattern recognition. The dramatic differences between civilian and military pattern recognition become most evident at the operational level, as opposed to the algorithmic level, the operational being in many ways the more significant determinant of overall system design. It is the operational side of pattern recognition that is of particular interest below.

2.0 CHALLENGES OF MILITARY TARGET DETECTION SYSTEMS

An operational contrast between civilian and military systems would be useful, to direct applied research across the military-civilian gap, and to highlight the motivation for the present work. The contrast has been discussed briefly in connection with decision support systems in [Van Der Wal, 1998] [Cox and Lloyd, 1984] and sea-mine detection in [Kessel, 2001c]. Here we summarize the three main elements that, for our purposes, distinguish military from civilian pattern recognition applications (with over-simplification for brevity):

(A) **Environments** are very diverse and continually changing in military applications (geography, time of day, seasons, weather, clutter density, stress level, etc.), whereas civilian pattern recognition systems tend to operate under fixed conditions (permanent indoor location, staged and repeatable scenes, low clutter, controlled illumination, low unchanging operator stress level, etc.).

(B) **Targets** are largely unpredictable in military applications (in range, bearing, aspect, type, technology, number, intent, etc.), whereas prior knowledge of "targets" is much greater in civilian applications.

(C) **Objectives, costs, and risks** of pattern recognition are continually changing in military operations between missions and during a single mission (time constraints, strategic importance of operation, available resources, assets at risk, nature of enemy threat, prior intelligence, politics, etc.), whereas civilian applications remain fairly constant and are often quantitatively defined. Military errors may also be irrevocable while incurring extreme costs, particularly when a high-threat target is missed, or civilian transport is mistakenly targeted. Costs may also be highly asymmetric, missed targets being in some cases significantly more costly than false alarms, as in mine hunting and anti-submarine warfare for instance. The perceived costs of errors in civilian pattern recognition applications, on the other hand, tend to be more proportionate between error types, and more acceptable – at least in the mind of the decision-maker.

Much that is taken for granted in civilian pattern recognition therefore no longer holds in operational military detection systems. The demands that changing costs and risks (C) place on operational pattern recognition are perhaps most often overlooked and bear reiteration. Given that optimal detection minimizes the costs and risks of operation (as shown in detection [Van Trees, 1968], in pattern recognition [Fukanaga, 1990], and in classical decision theory [Pearl, 1988]), and given that near optimal performance is required to prevail against challenges (A) and (B) above, it is clear that operational detection must either be extremely flexible and adaptive, or it must be limited to very narrowly defined circumstances. Both extremes are among the examples given below.

3.0 CAPABILITY AND RESPONSIBILITY: AN EFFECTIVE BALANCE

Detection is an uncertain operation insofar as some genuine targets may be missed and non-targets may be mistaken for targets. A cost and risk is associated with each type of error. These are defined in operational terms, as in point (C) of the previous section.

In detection theory, the susceptibility to error is quantified in terms of the probability of detection (ideally unity) and the probability of false alarm (ideally zero) [Van Trees, 1968]. Less than ideal probabilities are inevitable in practice, but they are adjustable, though not independently, with adjustments for increased probability of detection being made at the expense of increased probability of false alarm (see Figure 1). In simple detectors, the adjustments amount to changing a single detection threshold, raising it to decrease the detector's sensitivity to false alarms, when operating in regions of high clutter for instance, or lowering it to increase the detector's sensitivity to targets, when the slightest indication of a target is cause for alarm. A probabilistic approach can also be used to assess the performance of very complex detectors, even human operators, although their internal thresholds may be inaccessible to direct observation or adjustment. Their thresholds are inferred instead from performance tests, with adjustments made indirectly through training or heightened awareness [Green and Swets, 1988] [Wickens and Holland, 2000].

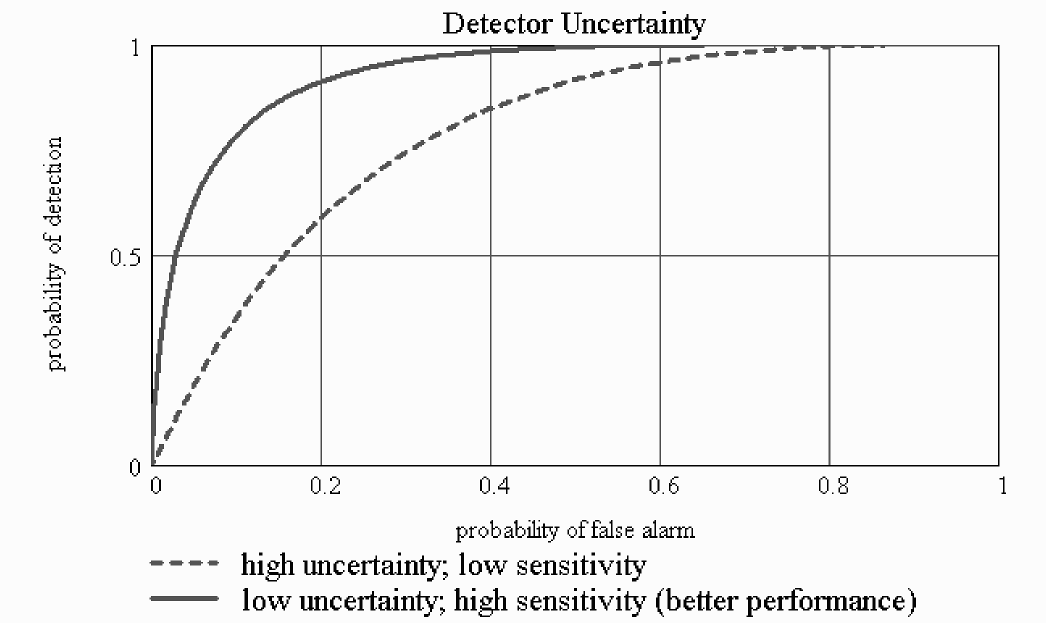

Figure 1: The full range of a detector's performance is given by its receiver-operator characteristics. High quality detectors achieve high probabilities of detection at low probabilities of false alarm. The adjustment of thresholds inside the detector determines the point at which a detector operates on its performance curve. This probabilistic view of detection uncertainties applies to a very broad range of detectors, from simple to complex, human or machine.

In any case, when faced with the prospect of error in detection, the system designer's quest for detection accuracy necessarily becomes one of cost and risk minimization. Cost and risk reduction is achieved in either of two ways, by:

(A) reducing the detector's uncertainty, sometimes called increasing the *sensitivity*, (i.e., increasing the probability of detection for a given probability of false alarm[1]), by improving sensors, improving pattern recognition, or resorting multi-sensor fusion perhaps; or

(B) reducing the costs and risks through operational measures or precautions (e.g., directing shipping around high-clutter areas where mine or submarine detection is difficult, excluding civilian and friendly traffic from areas where high-threat targets are expected to approach, using remote low-value sensor platforms for early warning, and so forth).

In practice, the system designer tries to do both in an effort to strike a working balance between a detector's abilities and responsibilities (See Figure 2).

---

[1] Maximizing the probability of detection for given probability of false alarm is known as the Neyman-Pearson condition for optimal detection [Van Trees, 1968]. It can be shown that this minimizes the expected total cost of operation.

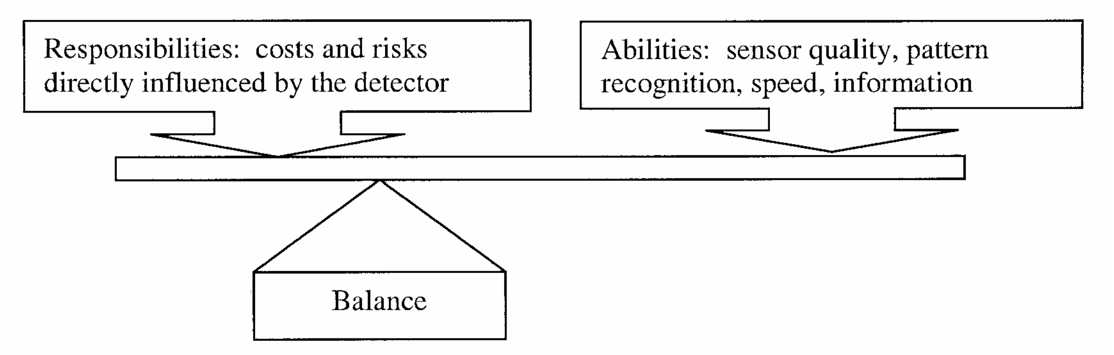

Figure 2: Effective detection requires that the responsibilities assigned to the detector (left) be balanced with the detector's abilities (right). A detector will be ineffective if it cannot be adjusted to keep the expected cost of operation reasonably low, or if it is given no responsibilities whatsoever. This principle of effective detection must be satisfied by every detector (human or machine) in a well-designed detection system.

Giving a detector too much responsibility means that the expected costs and risks directly influenced by the detector's alarms are unacceptably high. The detector jeopardizes the mission, by either failing to accomplish its assigned task, by incurring unacceptably high costs and risks, or both. Its use is ill advised. Giving a detector too little responsibility, on the other hand, means that the detector exerts no influence on mission costs and risks, which is to say that it serves little purpose. Effective detection therefore requires that the costs and risks borne by a detector be commensurate with the detector's ability, being neither so great as to jeopardize those who rely on the detector, nor so little as to make the detector inconsequential.

The same principle applies for every detector in a larger detection system. A detection system has at least as many conceptions of costs and risks, and as many detection thresholds, as there are detectors in the system, whether human or machine. Each detector is most effective when adjusted to minimize the costs and risks it directly influences. Thus a fundamental design strategy for a detection system is 1) to ensure that the costs and risks influenced by each detector are commensurate with its capabilities (balance principle), and 2) to devise methods for minimizing detection uncertainty and operational costs and risks. This is the approach followed in the more detailed strategies given in Section 5.

4.0 ACCOUNTING FOR THE DISSATISFACTION WITH *CONVENTIONAL* AUTO-DETECTION

Consider the common conception of an auto-detect feature which we call *conventional*. A computer and human operator simultaneously scan the same signals (synthetic aperture radar, sonar, or aerial photos, for example), each watching for target signatures, and each raising alarms in response to what is seen. The computer typically raises its alarms directly on the operator's display, for the operator to inspect, then confirm or dismiss. Such systems have been developed for use in underwater minehunting and anti-submarine warfare for instance [Kessel, 2001a] [Wort, 1996]. Conventional auto-detection is therefore characterized by its cautious stance of completely subordinating the automated alarms to the judgment of the operator – a justifiable caution if the performance of the automation is unknown, or is known to fail under realistic conditions.

In our experience, this approach to auto-detection systems has met with widespread dissatisfaction [Kessel, 2001b] [Wort, 1996]. They are perceived to overload and annoy rather than assist, with no evidence of overall performance gains above that of the operator working without the automation. All but novice operators are inclined to switch them off. Operators and system developers rightly point out that the auto-detect algorithms are inadequate. The algorithms make too many obvious errors and require continual expert adjustment to accommodate changing scenes, often overloading the operator as a result[2]. System developers have therefore spent considerable effort trying to improve detection algorithms. But success along that route seems far off. In Section 5, we consider alternate ways to realize more modest, immediate gains using the balance principle.

Shortcomings in the detection algorithms were in fact anticipated. This is why the automatic alarms were subordinated to the human operator from the very outset, as a matter of design. Note, however, that this solution by subordination does not address the problem of the automation's ineffectiveness. By removing *all* responsibility from the auto-detect, total subordination makes the automation ineffective because its alarms now serve little purpose. It is no wonder then that experienced operators do not see any use for activating conventional auto-detect features, or that they find them annoying. Being inadequate for the responsibilities originally envisioned for it, the auto-detect is left without any clear responsibilities, making its ineffectiveness complete, at least until other responsibilities more commensurate with the auto-detection's abilities have been clearly defined. This lead is taken up again in Section 5.3.

5.0 DECISION AID STRATEGIES FOR TARGET DETECTION

5.1 Advisory

In *conventional* auto-detection, the role envisioned for automation (but difficult to achieve) is that of an expert advisor. The automation captures the knowledge of a team of human experts, plus the accumulated experience of many missions, and reports its judgments (alarms) to the human operator, presumably in an effort to improve performance by reducing the operator's uncertainty. Needless to say, it is difficult to give constructive advice to expert operators on matters which to their mind are uncertain. The advice is necessarily viewed as doubtful unless the advisor is recognized to excel the operator, if only on occasion, in knowledge, speed, memory, training, clutter rejection, or combinations of these. Recognized superiority is especially important if the advice is given in a simple form, such as the on-screen alarms typically used in *conventional* auto-detection, with no further means for dialogue by which to rationally sway the opinion of the operator[3].

If auto-detection is to fill an effective advisory role, while keeping its alarms as simple and direct as possible (as they must be in high stress and chaotic situations), then it is necessary for the system designer to clearly identify deficiencies in the operator that the auto-detection can address without upsetting its balance between responsibilities and abilities.

To illustrate, consider one of the more sensational auto-detection systems currently operational, though used with extreme caution in a state of highest alert. Naval air-defence systems are capable of countering

---

[2] A review [Wort, 1996] of twelve existing auto-detection systems (mostly proprietary), all for sea minehunting, concluded that "it is reasonable to expect the operator may tune the selected algorithm on the order of, or less than, once per minute, for the type of ocean bottom or ambient conditions encountered". The operator typically requires thorough working knowledge of the image processing or pattern recognition schemes used by the auto-detection, and must continually apply that knowledge throughout a mission.

[3] The challenge of automating the advisory role lies at the heart of decision support systems where it has earned considerable attention [Woods and Roth, 1988] [Fischoff, 1985] [Dawes et al., 1989], as well as many warnings against unrealistically high expectations. System developers should proceed with caution, "starting small" with marginal but assured performance gains [Woods, 1986]. Much the same could be concluded from the dissatisfaction with the *conventional* auto-detection noted earlier.

coordinated missile and air attacks by automatically detecting, prosecuting, and evading targets. The primary sensor is radar, and automated alarms are signaled directly on the operator's viewing screen by highlighting newly acquired signatures on a geographic map of target positions and tracks. Under the advisory design strategy, the deficiency addressed here is the inability of the operator and crew to react quickly for effective defence against a missile attack, especially in times of high activity and stress. In the fully automatic mode, the automation directly influences tremendous costs and risks by taking quick action against a genuine attack or, in the case of error, by mistakenly targeting civilian or friendly aircraft that stray into the free-fire zone. To balance such risks, the uncertainties of the auto-detection must be very low. Every measure must be taken to ensure this. Uncertainty is lowered in practice by one basic rule understood by everyone who is aware of auto-engagement, namely, that any contact within the free-fire sector will be considered a target. The effectiveness of this rule in turn depends on the likelihood of non-targets straying into the sector, and on the likelihood of clutter generating false targets. If both can be safely ruled out by operational directives and disposition, then uncertainty is almost entirely eliminated by removing the false target class from among the allowable responses of the detector, since every new contact in the free-fire zone is considered a target when the fully-automatic option is exercised. Even relatively simple auto-detection algorithms, such as classical energy detection [Van Trees, 1968] for instance, may suffice when uncertainty is so thoroughly eliminated by operational measures in this way[4].

In this air-defence example, then, threats to the mission are reduced because the automation dramatically reduces the natural, but dangerously long reaction times of the crew for executing defensive action. In other words, automation's high speed reduces the weight of costs and risks acting on the left side of the balance in Figure 2, but it is still considerable. To balance the remaining weight, the ability of the auto-detection must be reinforced by the stringent operational directives imposed on its use. In other words, these directives increase the weight of the auto-detection acting on the right side of the balance in Figure 2. This illustrates the two key elements to the advisory strategy: 1) addressing a deficiency and 2) satisfying the balance principle.

More generally, the deficiency need not be the operator's. It may be a more mundane operational deficiency of the system, possibly hardware- or software-related. We call these *incidental* deficiencies. Consider the example of a sensor attached to a remotely controlled vehicle crossing hazardous zones while the operators remain a safe distance away. Real-time transmission of complete sensor signals may be prohibitive owing to bandwidth limitations on the radio link between the sensor and the control center. In such cases, an object-of-interest (or region-of-interest) detector might be inserted between sensor and radio link, to focus transmissions on signatures of interest. This might be done in remote minehunting, for instance, when a remotely controlled vehicle tows a high-resolution sidescan sonar to survey areas of the seafloor for mines. The vacant seafloor is of little interest, and therefore admits reduced resolution or outright suppression in the transmissions [Kessel, 2002]. The target class for the auto-detection is therefore defined broadly as any object or region of interest, and the automated "alarms" signify no more than signatures worthy of attention. Because it serves as a data-reduction filter between the sensor and the operator, the detector must obviously function with a very high probability of detection so far as true targets are concerned. At the same time, a high probability of false alarm can be tolerated because the transmitted data is scanned again for targets, this time by the human operator on receipt of transmission[5]. Focusing on interesting signals this way can reduce the transmission burden considerably, while the high false-alarm rates allowed for the auto-detection may permit the use of mediocre to poor detection algorithms with good results.

---

[4] Algorithm simplicity is often required to maintain the computational speed with which the auto-detection algorithms must work in such cases.

[5] By "high" false alarm rate we mean high relative to the false alarm rate to be achieved by the overall detection system, human and machine working together.

An auto-detection filter of this sort may also serve when the number of signals to be monitored requires too many operators, such as when dozens of beams of a towed-array sonar are to be monitored for submarine detection. The auto-detection would monitor all signals simultaneously, and alert one or more operators to "interesting" events, allowing the operators to focus more efficiently on signals of interest, while safely ignoring background noise. The workload reduction and improved performance may be dramatic.

As a final example, we can imagine auto-detection performing an advisory role during training missions, the automated alarms highlighting the kinds of signatures that operators should take into account during training. The deficiency addressed is operator skill, and the costs and risks borne by the automated alarms are the low-level risks associated with cultivating operator experience under realistic conditions. Only the experienced operator would find the alarms inconsequential, which is a positive outcome in this case.

Various deficiencies have been addressed in each of these examples. But what can be said of using automation to replace experienced operators in missions where operators have served well enough? There is evidence from civilian applications that straightforward actuarial rules can in fact outperform human experts [Dawes et al., 1989]. The reason is that actuarial rules[6] are optimized and unaffected by the biases of the human thought process. But these conclusions follow from experiments carried out in well-defined and limited conditions, not military operations. Further research into the relative performance of humans and machines, in detection tasks, is required for the reasons given in Section 2.

The move to replace human operators with automation must ultimately follow the balance principle of matching responsibilities to abilities. Even when they are less able than human operators, fully automated detection systems may be effective provided that the costs and risks of operation are sufficiently reduced by keeping operators clear of danger. In this way, a fleet of autonomous minehunting vehicles might be an operational improvement over a manned detection system.

5.2 Advanced Sensor

Automation can also improve mission performance without explicitly raising alarms. A suite of automated signature analysis and information management tools, for instance, may be used by the operator to reduce uncertainty in signals. Such tools may assume a purely passive role, performing computations and manipulations, and displaying the results at the operator's request. We call this the *advanced sensor* strategy inasmuch as the automation augments the signals, by assimilating and managing them into geospatial or chronological representations that convey more information than a real-time signal feed ordinarily provides. The mechanism for improved detection is that target/non-target discrimination improves with the addition of non-visual cues (derived by automated signature analysis) and with increased situational awareness (about the disposition of environmental clutter sources for instance, or the disposition of signatures gathered from different positions or times in the same region).

Signature analysis may typically include the computation of:

- target/clutter discriminating features that are too tedious for the operator to repeatedly undertake;
- non-visual cues derived by analysis such as measures of texture and features of signal processing (e.g., wavelet or Fourier analysis); and
- measures of absolute signature strength.

---

[6] Actuarial rules are simple mathematical rules for decision-making, usually based on linear decision variables.

The additional cues reduce the operator's uncertainty insofar as they provide additional information for detection that may not be immediately evident on inspection of a signature.

Information management, on the other hand, may typically:

- fuse signatures and geospatial information for improved situational awareness (i.e., topographic or bathymetric maps and charts for sense-making, clutter avoidance, orientation, hypothesis testing, and so forth);
- fuse historical and current signatures for discriminating difficult targets by change detection (i.e., the detection of a target because it is a new signature amid known clutter);
- fuse multiple looks of a region from different aspects;
- support geospatial imaging at various resolutions (from a wide-area mosaic of geo-referenced signatures, to single swath, and target-size windows of inspection);
- provide libraries of target and object images, for the operator to consult or for automatic matching of likely signatures; and
- provide management, ranking and sorting targets on the basis of their particular cues, class, or perceived threat level, in order to select the "most likely" targets from a list too numerous to prosecute.

In light of the balance principle, the advanced sensor strategy gives the automation no direct responsibility for operational costs and risks. The responsibility for raising alarms falls squarely on the operator. The absence of responsibility works in this case, where it failed with *conventional* auto-detection, because the automation now makes no pretence of raising alarms. Its effectiveness comes instead from the suite of services it provides as a tool set for the operator according to need and ingenuity as hypotheses about likely targets come to mind. The use of these tools would no doubt vary from one operator to another, and one mission to another, depending on environmental factors, mission objectives and constraints. Their utility therefore depends on the system developer's vision of the operator's typical workflow and cognitive processes, and on the tools' flexibility and convenience.

Virtually all of the information management tools mentioned here have been included in Canada's Remote Minehunting System (RMS), a Technology Demonstration Project entering its final development stages at the Defence R&D Canada – Atlantic. In this case, large-area surveys of the seafloor (rather like aerial photographs) are carried out by towing a sidescan sonar along straight tracks. Swath imagery is transmitted to operator consoles for real-time analysis if desired, where it is also stored for post-processing, in a database. Signature analysis works autonomously in the database, for generating clutter-density and seafloor-type maps, and populating fields with features associated with all contacts. The operator could ask for maps of clutter regions, for instance, where seafloor mines would be difficult to find, either to focus efforts on those regions, or, what is more likely, to chart a safe shipping route through low-clutter areas where clearance is more efficient. If the mission resources (e.g., the number of divers, total dive time available, etc.) are fixed and predetermined, then the target management tools can provide a map of targets ranked according to priority or level of certainty. The database also holds historic survey imagery for use in semi-automatic change detection, which compares present and historic (peacetime) surveys and which ranks among the most effective methods for clutter rejection.

5.3 Ergonomic

Operator vigilance and workload are among the foremost concerns during detection missions. Indeed, three of the top five “important issues” raised in a survey of sonar operators were 1) the ability to stay alert, 2) fatigue, and 3) work/rest cycle [Kobus and Lewandowski, 1991]. These were ranked higher than issues associated with the performance of the sonar equipment the operator’s control – the reverse, it would seem, of what many system designers might expect. An ergonomic design strategy focusing on just these human factors is therefore of particular interest.

In an ergonomic design strategy the costs and risks influenced by the auto-detection’s alarms are defined in terms of benefits to the operator’s physical and mental performance; e.g., vigilance, reduced decision bias, memory accuracy, reduced memory load, and workload reduction. There is no pretence now of the automation directly influencing the overall mission costs and risks, but its responsibility is clear and significant nonetheless: to increase the sensitivity of the most responsible and expert agent in the detection system – the human operator.

It is known, for instance, that lapses of attention during a watch can be countered by raising the occasional, conspicuous alarm (usually false) in order to increase the operator’s awareness, and as a reminder of what signatures are of interest. It has even been suggested that false alarms be purposely injected into the signal during periods of particularly low activity [Wickens and Holland, 2000]. The injected alarms can also be viewed as a training technique when genuine targets are rare [Woods and Roth, 1988]. The key, in any case, is to provide an occasional automated alarm without increasing the operator’s workload, without excessive, obvious false alarms to annoy the operator, and without prolonged periods with no alarms when attention lapses unchecked. In our experience with sea minehunting, these responsibilities could be filled by a number of auto-detection algorithms now available [Kessel et al., 2001]. Indeed, many algorithms previously thought ineffective might be turned to good use by a redefinition of responsibilities this way.

An operator-centered approach may represent a significant reversal of perspective for many developers. If low, regulated false alarm rates are crucial for vigilance gains, then the auto-detection must function with low probabilities of false alarm in high clutter regions, which for mediocre detectors necessarily results in a low probability of detection[7] – much lower than if the automation were entrusted with full mission costs and risks. But the auto-detection has very different responsibilities. It must keep the operator alert above all else, for which the automation’s probability of detection really does not matter. This redefinition of responsibility is detailed in [Kessel, 2001b], and with extension to workload in [Kessel, 2001a].

An ergonomic strategy may also aim to reduce the bias of operators, to ensure that they are not focusing on some target features to the exclusion of others. A checklist of key target features can counter bias, for instance, by requiring the operator to systematically inspect all features when making a target/non-target decision. Some navies have instigated a checklist for mine avoidance, for use by the operator who first detects an object in the view of a forward-looking sonar of a frigate, in order to reduce false alarms through a systematic process of elimination on the basis of key signature features. Bias and de-biasing are considered more generally in [Wickens and Hollands, 2000] [Kahneman et al., 1982].

---

[7] See lower-left corner of Figure 1.

5.4 Strategies in Practice

Each of the above strategies is suited to one or another application. No single strategy will satisfy all applications[8]. A combination of these strategies might also be followed, provided that the balance principle is satisfied. Given remotely controlled sensors, for instance, an advisory strategy might be used for data-reduction to facilitate radio links, while an ergonomic strategy might increase operator vigilance during real-time acquisition, and an advanced sensor strategy might decrease uncertainty through signal analysis and geospatial imaging. The key in any case is to ensure that a working balance of responsibilities and abilities is maintained, as much as possible, for all detectors in the system.

Any redefinition of responsibilities owing to a shift in design strategy may call for a radical readjustment of the auto-detect's operating characteristics. One should not be surprised to see it functioning best with probabilities of detection and false alarm that are significantly different from the probabilities expected from the detection system (human and machine) operating as a whole.

A good strategy would also recognize a dual role for human operators in development, as users and mission experts, to confirm the choice of a given strategy, guard against oversights, and ensure that their own responsibilities are commensurate with their abilities. A new automated detection aid may in fact meet with considerable distrust from an experienced operator owing to past experience with problematic conventional aids. To establish trust, it will be important to clearly explain the responsibilities and abilities of the new aid to operators [Muir, 1987], and of course give them experience working with the aid under realistic conditions.

Finally, well-defined strategies clarify the operational limitations of a detection system, which coincide with the marked imbalance of any detector in the system. The range of allowable clutter levels and environmental conditions, for instance, depends on the effect they have on the responsibility-ability balance of the auto-detection. The performance of automated air-defence systems designed for use on the open ocean where clutter is low, for instance, can suffer badly when brought into inland waters where land clutter may be high. The uncertainty regarding targets increases dramatically, possibly making the auto-detect-and-fire responsibilities insupportable. Reliance on the system may then jeopardize rather than help the mission. Much the same may be true of long-range submarine hunting sonars, where auto-detection designed for deep water becomes much more uncertain in shallow littoral waters owing to target-masking and target-like seafloor reverberation and geoclutter.

6.0 CONCLUSION

Much that is taken for granted in civilian pattern recognition solutions does not hold in operational target detection. Military applications routinely face complications – changing environments, unpredictable targets, and changing costs and risks – that are unprecedented in civilian applications. Hence it is more realistic to compare military detection systems with expert decision support systems for which the "rate, complexity, dimensionality, and uncertainty of events and information" [Wohl, 1981] likewise present unique challenges beyond the scope of civilian pattern recognition.

A common weakness in the design of detection systems is to envision the automation performing with the same probabilities of detection and false alarm as those expected of the overall detection system. This accounts for much of the dissatisfaction occasioned by *conventional* auto-detection. The situation is reviewed in light of the simple, common-sense principle of matching responsibilities to abilities for every

---

[8] The narrow-minded application of one strategy to all problems is a common pitfall in decision aiding, where it is ominously called the "Law of the Hammer": in a child's hand, a hammer is used indiscriminately, on everything [Hopple, 1986].

detector, human or machine, in a larger detection system. Three basic design strategies for target detection, all satisfying this balance principle, are outlined in this work: 1) an advisory strategy to fill operator or incidental deficiencies; 2) an advanced sensor strategy to reduce operator uncertainty, and 3) an ergonomic strategy to keep the operator in top form.

Military detection systems apparently demonstrate in microcosm many of the human-machine dynamics involved on a much larger, macroscopic scale in command and control. Detection systems may therefore be an ideal test bed for proving decision-aid principles on a smaller, more tractable scale, for use later at more complicated levels of command and control. The principle of matching responsibilities to abilities may generalize in this way.

7.0 REFERENCES

Allen, N. (2002). The Maritime Patrol Aircraft Acoustic Mission: A Phase Analysis. Defence R&D Canada – Atlantic, Canada, Technical Memorandum DRDC Atlantic TM 2002-088. Limited distribution.

Cox, I.J. and Lloyd, L.J. (1984). Artificial-Intelligence Systems in Antisubmarine Warfare: Results of a Pilot Study with Expert Systems. SACLANTCEN Memorandum SM-176.

Dawes, R., David, F., and Meehl, P. (1989). Clinical versus Actuarial Judgment. *Science*, 243, pp. 1668-1674.

Fischhoff, B. (1985). Decision Making in Complex Systems. In *Intelligent Decision Support in Process Environments*, E. Hollnagel, G. Mancini, and D. Woods Editors, Springer Verlag, Berlin, 1986; the published Proceedings of the NATO Advanced Studies Institute, Sept 16-27, 1985, San Miniato, Italy.

Fukunaga, K. (1990). Introduction to Statistical Pattern Recognition, 2nd Ed. New York: Academic Press.

Green, D.M., and Swets, J.A. (1988). Signal Detection Theory and Psychophysics. Los Altos, California: Peninsula Publishing.

Hopple, G.W. (1986). Decision Aiding Dangers: The Law of the Hammer and Other Maxims. *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. SMC-16, No. 6, pp. 948-964.

Kahneman, D., Slovic, P., and Tversky, A., Editors (1982). Judgment under Uncertainty: Heuristics and Biases. Cambridge: Cambridge University Press.

Kessel, R.T. (2002). Using Speckle to Identify Regions of Interest and for Mine Detection. In: *Proceedings of the International Society for Optical Engineering (SPIE)*, Volume 4742, Detection and Remediation Technologies for Mine and Mine-Like Targets VII, AeroSense 2002, Orlando, Florida.

Kessel, R.T. (2001a). On-Screen Alarms in Computer-Aided Detection Systems: Signal Processing, Human Factors, and System Design. Canadian National Defence Research and Development Branch Technical Memorandum DREA TM 2001-184, November 2001.

Kessel, R.T. (2001b), “On-Screen Alarms in Computer-Aided Detection Systems: A “Cost-Benefit” Analysis,” Proceedings of CAD/CAC 2001, Halifax Trade and Convention Center, Halifax, Nova Scotia, 12-14 November 2001.

Kessel, R.T. (2001c). Mission Adaptive Computer Aided Detection (CAD) for Mine Hunting. Defence Research Establishment Atlantic, Canada, Technical Memorandum DREA TM 2000-111.

Kessel, R.T., Myers, V., and Fawcett, J.A. (2001). Computer Aided Detection and Classification (CAD/CAC) for Sea Minehunting. Canadian National Defence Research and Development Branch Technical Note DREA TN 2001-195.

Kobus, D.A. and Lewandowski, L.J. (1991). Critical Factors in Sonar Operation: A Survey of Experienced Operators. Naval Health Research Center Report No. 91-19.

Muir, B.M. (1987). Trust between Humans and Machines, and the Design of Decision Aids. *Int. J. Man-Machine Studies*, Vol. 27, pp. 527-539.

Pearl, J. (1988). Probablistic Reasoning in Intelligent Systems. San Mateo: Morgan Kaufman Publishers.

Van Der Wal, A.J. (1998). The Importance of Softcomputing Methods for Military Observation Systems. *Proceedings of the 3rd International FLINS Workshop*, pp. 163-170.

Van Trees, H.L. (1968). Detection, Estimation, and Modulation Theory. New York: John Wiley and Sons.

Wickens, C.D. and Hollands, J.G. (2000). Engineering Psychology and Human Performance, Third Edition. New Jersey: Prentice-Hall.

Woods, D.D. and Roth, E.M. (1988). Aiding Human Performance II: From Cognitive Analysis to Support Systems. *Le Travail humain*, tome 51, n° 2, pp. 139-172.

Woods, D.D. (1986). The Design of Decision Aids in the Age of "Intelligence". In: *Proceedings of the 1986 IEEE International Conference on Systems, Man, and Cybernetics,* pp. 398-401.

Wohl, J.G. (1981). Force Management Decision Requirements for Air Force Tactical Command and Control. *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. SMC-11, No. 9, pp. 618-639.

Wort, P. (1996). CAD/CAC Study Report. MacDonald Dettwiler Report CL-RP-50-7394, Issue 1, Revision 0, June 27, 1996.

This page has been deliberately left blank

Page intentionnellement blanche