TECHNICAL REPORT 1915
July 2004

An Intelligent Threat Assessment Tool for Decluttering Naval Air Defense Displays

M. St. John
D. I. Manes
H. S. Smallman
Pacific Science & Engineering Group, Inc.

B. Feher
J. G. Morrison
SSC San Diego

SSC San Diego

TECHNICAL REPORT 1915
July 2004

An Intelligent Threat Assessment Tool for Decluttering Naval Air Defense Displays

M. St. John
D. I. Manes
H. S. Smallman
Pacific Science & Engineering Group, Inc.

B. Feher
J. G. Morrison
SSC San Diego

SPAWAR
Systems Center
San Diego

SSC San Diego
San Diego, CA 92152-5001

EXECUTIVE SUMMARY

Operational displays can quickly become congested with large numbers of symbols. This report discusses our study of a naval air defense task in which users monitored a cluttered airspace, evaluated aircraft for their levels of threat, and executed defensive responses against significant threats. A heuristic threat assessment algorithm continuously evaluated aircraft for their levels of threat, and it "decluttered" the less threatening ones by decreasing the salience of their symbols on the geographical display. As expected, 27 expert U.S. Navy users appropriately distrusted the automation and continuously checked its assessments. Nonetheless, decluttering improved response timeliness to threatening aircraft by 25% compared with a baseline display with no decluttering, and it was especially beneficial for detecting and monitoring threats in more peripheral locations on the display. Decluttering did not affect which aircraft were deemed threatening, and 25 out of 27 users preferred using a declutter display over the baseline display.

CONTENTS

Figures

Tables

INTRODUCTION

Clutter can become a serious problem for users monitoring situation displays. For example, in naval air defense warfare, users monitor airspaces for threatening aircraft. These airspaces are frequently in busy environments near land that contain multiple commercial airlanes and other air traffic. Clutter increases search times by increasing the number of objects that must be sifted through or searched to find objects of interest (Treisman & Gelade, 1980). Clutter also increases the chance for “change blindness,” the chronic human inability to detect any changes in a scene when attention is focused in one location while critical changes occur elsewhere (Rensink, 2002). These problems can lead to reduced situation awareness (SA) and delayed response times to fast-changing events.

A common method for reducing clutter and facilitating SA is to identify important objects and somehow highlight them. Highlighting, when the identification process is reliable, allows users to focus on a subset of objects and thereby effectively reduces the number of objects that must be sifted through or monitored (e.g., Fisher, Coury, Tengs, & Duffy, 1989). However, one downside of highlighting and cueing is that it can impede the detection of important objects that are mistakenly left unhighlighted when the automation is imperfect or the situation is uncertain (e.g., Baddeley, 1972; Posner, 1980; Yeh & Wickens, 2001).

A related method for reducing clutter is to identify *less* important objects and declutter them from the display by making them somehow less visually salient. This method also reduces the effective search space by eliminating some objects from the focus of attention. Several studies have shown that users appreciate and benefit from decluttering tactical displays for search tasks (Johnson, Liao, & Granada, 2002; Nugent, 1996; Osga & Keating, 1994; Schultz, Nichols, & Curran, 1985).

Many methods have been used to declutter objects by reducing their visual salience, including size reduction, dimming, turning symbols into dots, and complete removal. Ideally, a good declutter method should visually segregate important from less important objects, but with minimal disruption to the information content of the symbols. St. John, Feher, & Morrison (2002) found that dimming symbols to one-third of their initial luminance supported easy segregation, but without removing any information.

A separate issue is how the highlighted or decluttered objects are identified in the first place. In most experimental studies, the identification function is simply assumed to exist, but is left unspecified. In applied tactical domains such as air warfare, the identification functions are typically simple classification rules such as all friendly aircraft or all aircraft with altitudes over 25,000 feet (standard U.S. Navy practice). Although attractive because of their simplicity, these rules often fail to meet the needs of sophisticated users because they do not align with the categories of most interest to tactical users.

A more sophisticated approach is to define meaningful categories of objects and then use these categories as the basis for decluttering. For example, in air defense, rules can be defined to identify commercial versus tactical aircraft, and then the commercial aircraft can be decluttered (standard U.S. Navy practice). Of course, such rules are necessarily heuristic and miscategorize aircraft on occasion. Moreover, the identification function of most interest to tactical users is the threat level of aircraft. The U.S. Navy users monitor tactical situations to assess threats and then execute responses to minimize them. Threat, however, is an ill-defined and complex function of many aircraft attributes that requires years of experience.

Development of reliable automated threat assessment algorithms have long been a "holy grail" for aiding SA generally, and air defense in particular. Unfortunately, there are several challenges to producing reliable threat evaluation automation. First, the problem can grow extremely complex in attempting to account for all possible variables, including aircraft kinematics, coordinated aircraft behaviors (the big picture), intelligence information, and situational factors such as the geopolitical context. Second, the problem can suffer from tremendous ambiguity because data may be unknown or unknowable. For example, aircraft identity is often based on electronic emissions that may not be detectable or available, and ultimately, the intent of an aircraft can never be established with certainty.

Third, expert decision-makers frequently disagree about the threat of individual aircraft (Marshall, Christensen, & McAllister, 1996). Consequently, an automated algorithm can never match the threat ratings of every expert user. Fourth, well-known problems of automation trust, complacency, and confirmation bias (Parasuraman & Riley, 1997) can undermine the effective use of automation and lead to disastrous consequences. For example, a user might monitor only those aircraft that the automation indicates as threats. If the automation missed a threat, the user might be significantly delayed in noticing it. Or if the automation mistakenly overrated the threat of an aircraft, a user might treat it more aggressively than necessary. On the other hand, distrust of automation might actually increase workload by driving users to increase their monitoring of lower threat aircraft.

Our philosophy is to treat the automation and the user as a "mixed initiative" system that combines heuristic automation that is known to be imperfect with engaged, knowledgeable users. According to the "Trust but Verify" design strategy (St. John & Manes, 2002; St. John, Manes, & Osga, 2002; St. John, Oonk, & Osga, 2000), users understand how and where the automation is likely to be trustworthy or to make errors, and they verify the automation accordingly. The Trust but Verify design strategy fits well with what Parasuraman, Sheridan, and Wickens (2000) term "lower level" automation, which might involve merely identifying alternative solutions rather than recommending a single best solution or executing a solution unless countermanded by the user. For example, in a visual search task, St. John and Manes (2002) used heuristic automation to make a rough first cut at identifying the locations of hidden targets. Users then exploited this information to guide their own searches. This approach led to a 23% improvement in search times, even when the automation was far from perfect and only 70% reliable.

In a situation monitoring paradigm such as air defense, heuristic automation could be used to identify and highlight threatening aircraft and declutter less threatening aircraft. Importantly, the less threatening aircraft would continue to be displayed, but with reduced salience. Therefore, the decluttered aircraft would not distract from the higher threat aircraft, yet would still remain available for inspection. Users could exploit the information provided by the automation by focusing most of their attention toward the highlighted aircraft while periodically scanning the entire display and verifying the automation's assessments of the decluttered aircraft. This decluttering method should enhance SA and facilitate a timely response to significant threats because the significant threats would be clearly visible on the display. This enhanced visibility might be especially useful for facilitating the early detection of significant threats at longer ranges from own ship. Yet, because the less threatening aircraft remain visible, although at a reduced level, users should be able to maintain awareness of the entire situation. The improved efficiency for monitoring the significant threats should allow ample time to verify the automation's evaluations.

The current experiment tests these predictions in a scenario-based, quasi-realistic air defense task with expert naval users. Figure 1 shows a snapshot of the display used in the experiment. The blue circle near the center of the display represents own ship. Unknown, potentially threatening aircraft appear as yellow clover shapes. Less threatening aircraft appear faded, and the significantly threatening aircraft stand out as bright yellow amid the clutter.

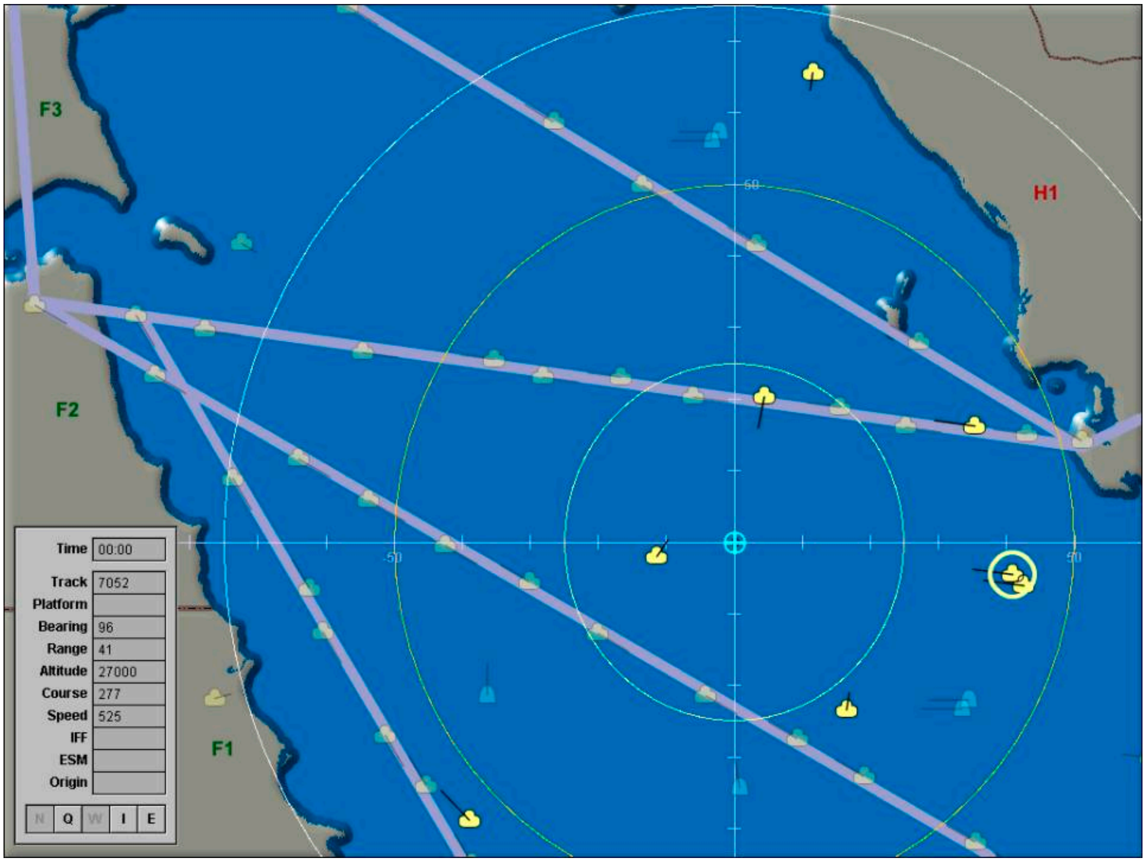

Figure 1. Task Display (decluttered aircraft appear faded).

Participants performed the normal tasks involved in air defense, monitoring an airspace, evaluating aircraft, and responding to the “significantly threatening” ones by issuing queries and warnings. Significantly threatening aircraft were defined as aircraft scoring an 8 or higher on a 10-point scale of threat. While the actual air defense task involves a team of naval personnel, the experiment was designed to be performed by a single individual by removing many subsidiary technical tasks such as correlating raw radar data and operating radio circuits. The scenarios were designed to be cluttered cluttered and reasonable challenging by making a number of aircraft ambiguously threatening.

A heuristic threat assessment algorithm evaluated the aircraft every second as they moved about the display, and decluttered the less threatening ones. The algorithm did not have to be very precise to perform this categorization, and participants were warned that the algorithm was likely to make occasional mistakes. Hence, a heuristic model of aircraft threat assessment was adequate (see below).

To declutter the less threatening aircraft, we followed the method used by St. John, Feher, and Morrison (2002) for reducing the luminance of the aircraft symbols to one-third of their initial values. This method allows good segregation between fully visible and decluttered symbols while continuing to represent information about the decluttered aircraft so that overall SA can be maintained.

The natural place to set the declutter threshold was to declutter all but the significantly threatening aircraft. However, given the heuristic nature of the automated threat algorithm, it was likely that the algorithm would occasionally declutter an aircraft that one or more participants might determine to constitute a significant threat. Lowering the threshold to keep more "borderline" threatening aircraft fully visible might reduce this problem, but at the cost of leaving more aircraft fully visible and increasing clutter on the display. Figure 2 shows hypothetical distributions of threat scores for threatening and nonthreatening aircraft, given a *reasonably* reliable but not perfectly reliable algorithm. Most aircraft are nonthreatening, and their threat scores tend toward low values. A minority of aircraft are threatening, and their scores tend toward high values.

Unfortunately, the two distributions are likely to overlap. As Figure 2 shows, moving the declutter threshold involves trading off risk and clutter. A high threshold reduces clutter and should aid users in focusing on significant threats, but at the risk of the automation inappropriately decluttering a significant threat that must still be detected and monitored despite its reduced visibility. On the other hand, a lower threshold keeps more aircraft fully visible and therefore provides less reduction of clutter. In turn, decluttering provides less aid to users who must spend more time searching among and evaluating a large set of fully visible aircraft, only some of which are actually significantly threatening. However, the lower threshold should reduce the risk of the automation inappropriately decluttering a significant threat.

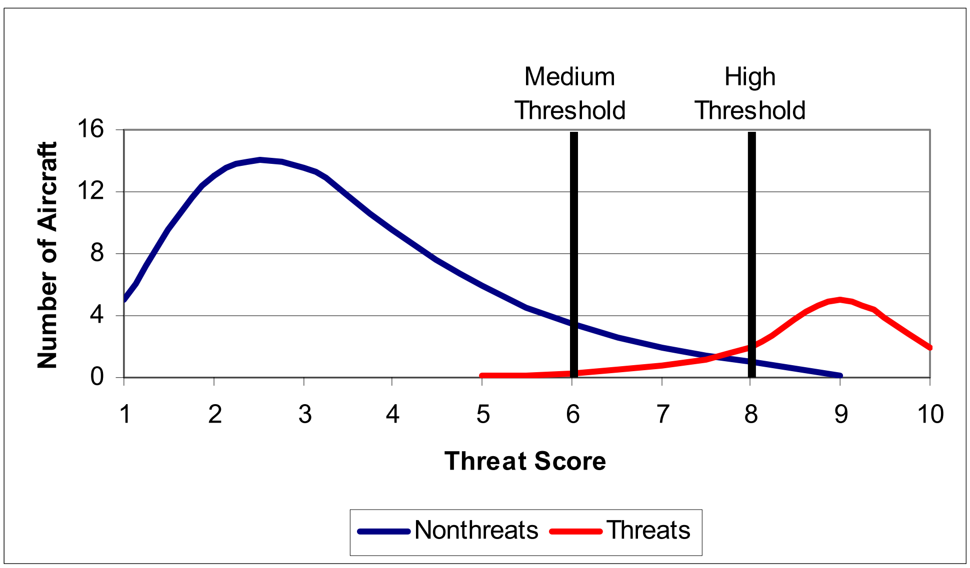

Figure 2. Distributions of threat scores for threatening and nonthreatening aircraft, and effects of high and medium declutter thresholds. Red regions to left of the thresholds are "misses," and blue regions to right of thresholds are "false alarms."

To investigate this trade-off empirically, the declutter threshold was manipulated as an independent variable in the study. In the high threshold condition, only the significantly threatening aircraft remained fully visible (8 or higher on a 10-point scale). In the medium threshold condition, significantly threatening and borderline threatening aircraft remained fully visible (6 or higher on a 10-point scale).

A final consideration was the difficulty of measuring performance in experiments on tasks such as air defense that involve substantial expert user judgment. This difficulty arises for two reasons. First, the assessment of which aircraft are threatening varies among experts (e.g., Marshall, Christensen, &

McAllister, 1996). Second, the timing of responses to threatening aircraft is also known to vary among experts (e.g., Morrison, Kelly, & Hutchins, 1996). In the experiment presented here, the assessment variability problem was addressed by allowing participants to exercise their own judgment in identifying threatening aircraft. Only aircraft that individual participants determined as significantly threatening were included in the analyses of response timeliness.

The response variability problem was addressed by explicitly defining when and how participants were required to respond to significantly threatening aircraft. Participants were required to respond at specific ranges to any aircraft that they determined to be significantly threatening. For example, participants were required to query (i.e., hail over a radio channel) every significantly threatening aircraft immediately if they crossed within 50 miles of own ship. Participants were also required to respond immediately to any aircraft that became significantly threatening due to it performing some suspicious or menacing behavior. These explicit Rules of Engagement (ROE) meant that any delays in responding could be attributed to poor SA rather than differences in judgment.

METHOD

PARTICIPANTS

The participants were 27 U.S. Navy personnel, 26 male and 1 female. Ages ranged from 24 to 54 years, with a mean of 35 years. Eight participants were chiefs or senior chiefs (E-7 to E-8) from the Aegis Training and Readiness Center Detachment San Diego; three were senior officers (O-5 to O-6) from the Tactical Training Group, Pacific; and 16 were junior officers (O-2 to O-4) from the Airborne Early Warning Wing, Pacific. The participants had from 3 to 30 years of service in the U.S. Navy, with an average of 13 years. A subject matter expert rated each participant's air warfare expertise and experience on a three-point scale. Fourteen participants received a very high rating, two received a high rating, and 11 received a moderate rating.

TASK

The task was a quasi-realistic naval air warfare task in which users monitored an airspace filled with more and less threatening aircraft (called "tracks"). The geographical display showed a 170- x 120-nautical mile area reminiscent of the Persian Gulf (Figure 1). Three relatively friendly countries, F1, F2, and F3, appeared on the left, and a relative hostile country, H1, appeared on the right. Own ship appeared near the center of the area and was designated by a blue circle. Commercial airlanes appeared as faded violet lines. The experiment was run on a laptop with a 15-inch screen that displayed 1024 x 768 pixels per inch.

Standard military symbols (MIL-STD-2525B) represented aircraft. Unknown tracks, including commercial airliners, oil platform helicopters, and tactical aircraft appeared as yellow amorphous cloverleaf shapes, with black speed leaders indicating their heading and speed (long leaders indicated fast speeds). Friendly military aircraft appeared as blue bullet shapes. No ships other than own ship and no submarines appeared on the display.

In all conditions, users could access a variety of information about a track by clicking on a track and then viewing a set of track data that appeared in a window in the lower left corner of the display. The track data included a track number for identification; the platform or type of aircraft; the bearing and range of the track from own ship; the altitude, course, and speed of the track; two types of electronic/radar information (identification friend or foe [IFF] and electronic support measures [ESM]); and its country of origin. For realism, not all information was available for every track. For example, track 7052 in Figure 1 is emitting no identifying electronic or navigational radar information; therefore, its IFF and ESM are unknown, and consequently, the platform is also unknown. Additionally, the track flew in from the East over water, so its country of origin is also unknown.

There were three equivalent scenarios, each lasting 15 minutes. During each scenario, tracks moved about the display at realistic rates from 95 to 560 nautical miles per hour, which is equivalent to 10 to 55 pixels per minute. Approximately 50 tracks were always on the display, with tracks occasionally entering or exiting the displayed area. Most tracks appeared benign, behaving like normal commercial airliners, oil platform helicopters, or other light commercial aircraft. At each moment, however, approximately seven tracks appeared significantly threatening (8 or higher on a 10-point scale), behaving, for example, like tactical fighter aircraft moving at a high speed from hostile origins toward own ship. Approximately 12 additional tracks appeared potentially threatening or "borderline" (6 or 7 on a 10-point scale of threat). These tracks presented a mix of benign and threatening attributes.

As tracks moved about the display, their threat levels changed. For example, as tracks approached own ship, their threat levels rose, and then dropped once they passed. Occasionally, aircraft would start out by appearing as a commercial airliner following an airlane, and then would abruptly change course and head inbound at high speed. This action would raise their threat score abruptly. Other tracks appeared suddenly from islands or oil platforms. In general, the scenario presented a range of aircraft behaviors and kept the participants busy.

There were three conditions: no declutter, medium threshold declutter, and high threshold declutter. Assignment of scenarios to conditions was counterbalanced across participants. In the no declutter condition, all track symbols appeared equally bright, and the user received no aid in evaluating the tracks for their levels of threat to own ship. In the two declutter conditions, less threatening tracks were decluttered by reducing the luminance of their symbols to one-third of their initial value.

Declutter was implemented in two parts: (1) each track on the display was evaluated every second and assigned a threat score, and (2) the less threatening tracks were decluttered. The threat assessments were accomplished by using a "declutter algorithm" based on research into how navy experts evaluate threat (Liebhaber & Feher, 2001; Liebhaber, Kobus, & Feher, 2002; Marshall, Christensen, & McAllister, 1996). The algorithm consisted of two steps. First, a raw threat score was computed by summing the threat values for each track's attributes (Table 1).

Second, the final threat score was produced by transforming the raw score to accentuate the differences among intermediate values, and then rescaling the result between 1 and 10. Accentuating the mid-range of the threat scale was useful because few tracks ever received extreme scores. In more detail, the raw threat score was first rescaled within the range from –0.5 to 0.5. Then, the score was transformed using the logistic function, and then it was rescaled again within the range from 1 to 10. Equation 1 and Figure 3 show how the rescaled raw scores (R) are transformed into final scores.

The participants monitored the tracks and responded to the significantly threatening ones. Participants were instructed that the evaluation part of the task was their own expert judgment. They were also told that the threat algorithm and declutter operation was only an imperfect aid designed to provide a reasonable "first cut" at evaluating threat. These instructions allowed and encouraged users to judge for themselves which tracks were significantly threatening.

Once a track was judged as a significant threat, however, the ROE determined how participants were required to respond. Two types of "significant events" required responses: (1) ring crossings, and (2) threat-level increases. For ring crossings, participants were required to "notify alpha bravo" (i.e., notify a superior command element) if a significantly threatening track crossed a ring at 75 nautical miles from own ship, "query" the track if it crossed a ring at 50 nautical miles from own ship, and "warn" the track if it crossed a ring at 25 nautical miles from own ship. Participants were required to perform these responses immediately at the ring crossings. For threat-level increases, if a previously less threatening track became a significant threat by performing some threatening action such as turning inbound and increasing speed, then participants were asked to respond immediately with the response appropriate for that distance from own ship. The declutter algorithm identified 25 significant events during each scenario. It also identified 29 "borderline events" (when a borderline track crossed a ring or a track increased its threat level to become a borderline track) and 40 "low-threat events." Of course, participants were only required to respond to those events that they personally judged as significant. Finally, at the beginning of each scenario, participants were required to "come up to speed" on the situation by immediately responding to each significantly threatening track on the display.

Table 1. Threat values for track attributes.

Attribute	Value	Score	Attribute	Value	Score
Affiliation	Neutral	2.5	Platform	Unspecified	4.3
	Unknown	3.5		737	2.0
	Hostile	4.5		E-2	4.5
	Friendly	NA		F-14	4.5
				S-3	4.5
Origin	Unspecified	0.0	ESM	Unspecified	0.2
	Own ship	-0.8		APQ-65	-0.2
	F1	-0.8		APS-137	-0.2
	F2	-0.8		AWG-9	0.2
	F3	-0.8		RDR-1	0.0
	H1	1.8			
IFF	Unspecified	1.6	Airlane	No	1.2
	Mode 3a	0.2		Yes	-1.0
	Mode 4	-2.0			
Feet Wet	No	-0.4	Group	No	0.0
	Yes	1.6		Yes	1.2
Approach	≤ 90 degrees	1.8	Range	≤ 5 nm	2.0
	≤ 180	-0.4		≤ 25	1.8
				≤ 50	0.8
				> 50	-0.4
Altitude	≤ 500 feet	1.4	Speed	≤ 150 nmi per hour	0.2
	≤ 1000	1.2		≤ 250	0.3
	≤ 5000	1.0		≤ 350	0.4
	≤ 10,000	0.8		≤ 450	0.6
	≤ 20,000	0.4		≤ 550	1.8
	> 20,000	0.0		> 550	2.0

Note: Affiliation refers to whether the track is known friendly, known neutral/commercial, unknown, or known hostile. All tracks in the scenario were either known friendly or unknown. Platform refers to the type of aircraft. Origin refers to the country of origin. ESM and IFF refer to electronic/radar emissions from the track. Airlane refers to whether the track is flying on a known airplane. Feet wet refers to whether the track is over water (or land). Group refers to whether the track is flying in a group. Approach refers to the angle of approach of the track toward own ship. Range, altitude, and speed refer to the kinematics of the track in units of nautical miles, feet, and nautical miles per hour.

$$\text{final score} = 10 * 1/(1 + \exp^{-7*R}) \tag{1}$$

Figure 3. Transformation of raw scores to final scores.

These rules provided a firm measure for the timeliness of responses. Because participants were required to respond at specific ranges and immediately following specific changes in track behaviors, any delays could be measured in time and range from own ship.

All responses were executed by first hooking the track (by clicking on it), and then pushing the appropriate button underneath the track data display (either N, Q, or W for notify, query, warn, respectively). Two additional responses, “illuminate” and “request to engage” were also available to participants if they felt tracks represented an especially elevated level of threat. Unlike notify, query, and warn, no ROE or other guidance was provided for when such actions should be taken. The method for executing these responses was the same as for notify, query, and warn responses, except the buttons were labeled "I" for illuminate and "E" for request to engage. These extra response options were included to provide added realism and to keep users occupied and engaged with the most threatening tracks, as they would be in the real task.

PROCEDURE

Participants were given a basic description of the task and then asked to sign informed consent forms. They were then given a detailed orientation to the display, the task, the ROE, and the tactical situation using a static view of the no declutter condition. Participants were then briefly exposed to all three conditions and told that we were interested in how the different displays might influence their performance. They then ran through a practice scenario with assistance from the experimenter. The practice scenario used the no declutter condition and lasted 5 minutes. Following the practice, participants rated their expected difficulty in performing the task when using each of the three interfaces.

Each participant performed in all three declutter conditions, one with each scenario in a counter-balanced order. Twenty-four participants were administered the NASA Task Load Index (TLX) (Hart & Staveland, 1988; NASA Ames Research Center Human Performance Group, no date) following each scenario to assess their subjective workload levels. The three other participants wore a head-mounted eye-tracking system during each scenario to assess their eye movements. Those results are not reported here. Because administering the TLX to these participants would have required removing and then re-attaching the headgear, we elected to exclude these participants from the TLX task.

Finally, following all three scenarios, participants filled out a questionnaire. The questionnaire again asked participants to rate the difficulty of the task when using the three interfaces. It also asked a number of questions concerning the participants' strategies and the usability, strengths, and weaknesses of the interfaces.

RESULTS

BEHAVIORAL MEASURES

The declutter algorithm categorized tracks into significant threats (scores of 8 to 10), borderline threats (scores of 6 to 7), and low threats (scores of 1 to 5). To assess how the declutter interfaces influenced the number and types of tracks that elicited responses, the number of responses (notifies, queries, and warnings) taken, overall and for each category of threat, was tabulated for each scenario for each participant. Occasionally, participants made multiple notifies, queries, or warnings toward the same track. Because participants were not allowed to maintain written notes about what responses they made to each track, the standard operating procedure or memory lapses may have initiated these responses. In either case, only the first response was used in the analyses.

To assess how participants monitored the display, the number of tracks hooked, overall and for each category of threat, was also tabulated for each scenario for each participant. Finally, to assess how the declutter interfaces influence response timeliness, response times were computed by taking the difference between the time a response occurred (i.e., the N, Q, or W button was clicked) and the time of the most recently significant event. Mean response times, overall and for each level of threat, were then computed for each scenario for each participant.

First, we asked which tracks elicited notify, query, and warn responses from participants. Participants responded an average of 21.4 times during each scenario. On average, each participant responded to less than one low-threat track, and 81% of the participants responded to no low-threat tracks. On average, each participant responded to 3.5 borderline threat tracks and 17.2 significantly threatening tracks (Figure 4). Recall, for comparison, that the declutter algorithm identified 25 significant events during each scenario.

In summary, 80% of participants' responses were made to tracks that the declutter algorithm identified as significant threats, and 17% of their responses were made to tracks that the declutter algorithm identified as borderline threats. Only 3% of participants' responses were made to low-threat tracks. These results indicate that the declutter algorithm and the participants corresponded quite well with one another in evaluating threat at this basic, yet critical, level of categorization.

Second, we asked how the declutter operation influenced responding. The overall number of responses in each condition was submitted to a one-way repeated measures analysis of variance (ANOVA). There were no differences in the number of responses among declutter conditions, $F(2, 52) = 0.9$. Looking at each level of threat in separate one-way repeated measures ANOVAs, there were no differences in the number of response to low threat tracks, $F(2, 52) = 1.1$, $p = 0.33$, or to significant threat tracks, $F(2, 52) = 1.2$, $p = 0.30$.

However, there was a difference in the number of responses to borderline threats, $F(2, 52) = 5.3$, $p = 0.008$. The medium declutter condition had more responses than in the high declutter condition ($p < 0.05$ by Tukey–Kramer post hoc test). This difference is understandable because borderline threat tracks were fully visible in the medium declutter condition, but decluttered in the high declutter condition. However, the difference was very small in absolute terms: 4.4 responses in the medium declutter condition versus 2.7 responses in the high declutter condition. Keeping the borderline threat tracks fully visible appears to have slightly elevated the likelihood that participants would respond to them, and decluttering the borderline threat tracks appears to have slightly lowered the likelihood that participants would respond to them.

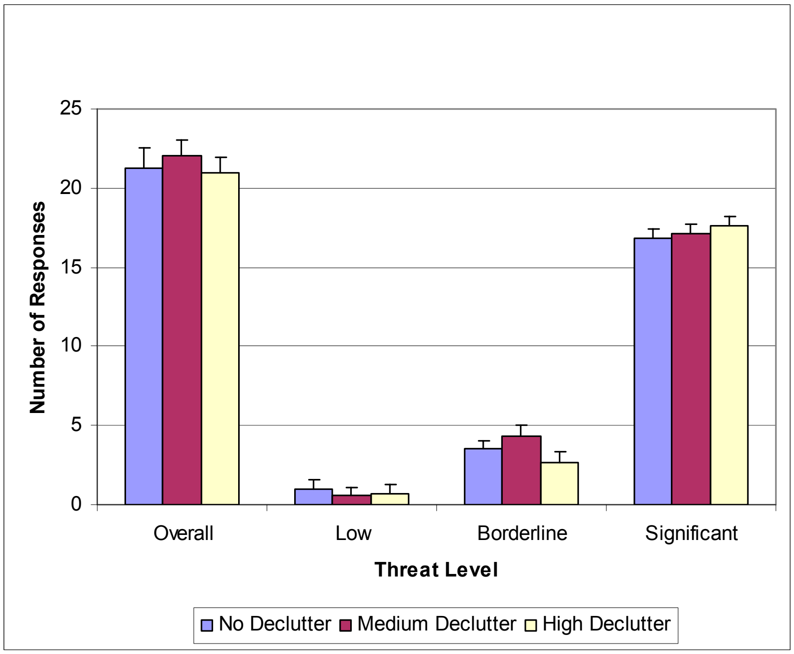

Figure 4. Number of responses (overall and within each threat category).

The small difference, however, and the lack of any differences overall, indicate that decluttering did not appreciably influence the threat assessment process. This finding of no confirmation bias, except for a slight bias among borderline cases, is important. The participants continued to apply their own judgment in deciding which tracks constituted significant threats.

The third question we asked was how the declutter operation might have influenced the process of monitoring the display and maintaining SA. Situational awareness was measured by tabulating which tracks participants selected, or "hooked," to view and evaluate their detailed attribute values. The assumption was that participants would tend to hook tracks repeatedly that were threatening or otherwise worth close examination. Our hypothesis was that decluttering would help participants focus on monitoring high-threat tracks. The number of hooks for each declutter condition and level of threat were submitted to a two-way repeated measures ANOVA. Confirming the assumption, participants primarily hooked the significantly threatening tracks, $F(2, 52) = 145.5$, $p < 0.0001$ (Figure 5).

The declutter operation had no effect on the overall amount of hooking, $F (2, 52) = 0.5$. This finding is important because it indicates that decluttering did not reduce participants' attention to and close monitoring of the situation, nor did it create extra work for participants by influencing them to increase their hooking. Rather, participants continued to hook and evaluate tracks at their normal levels.

Decluttering did influence which tracks were hooked, however, as indicated by a significant interaction between declutter condition and threat level, $F(4,104) = 13.9$, $p < 0.0001$. To examine this interaction, we looked separately at each threat level in one-way repeated measures ANOVAs. For significantly threatening tracks, high threshold declutter increased the amount of hooking, $F (2, 52) = 9.4$, $p = 0.0003$. This finding indicates that participants watched and evaluated the significantly threatening tracks more closely when the declutter operation kept these tracks fully visible and decluttered the rest. Interestingly, this increase did not occur in the medium declutter condition, even though the medium declutter condition also kept these tracks fully visible. Instead, the medium threshold declutter condition increased the number of *borderline* threats that were hooked, $F(2, 52) = 19.7$, $p < 0.0001$.

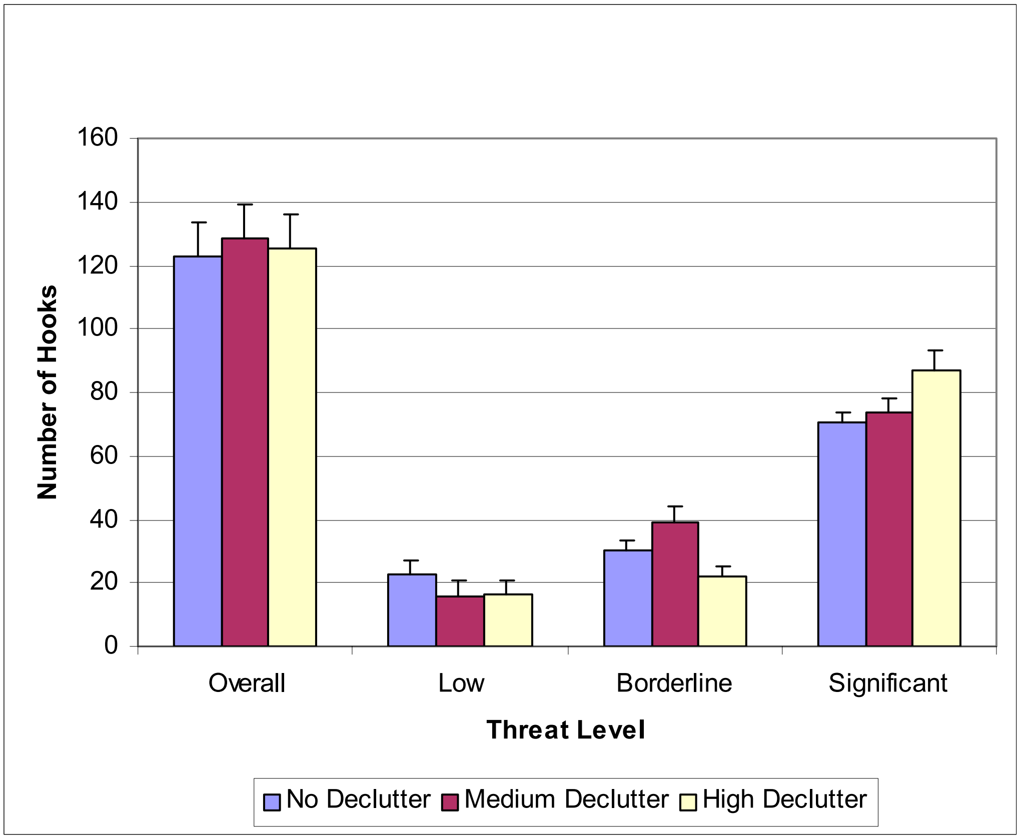

Figure 5. Number of hooks (overall and within each threat category).

In summary, making only the significantly threatening tracks fully visible (high threshold declutter) led participants to hook them more frequently. Making the significant and borderline threat tracks fully visible (medium threshold declutter) led to an increase in hooking only the borderline threats. Perhaps participants hooked at these borderline tracks more frequently than otherwise to understand why they were fully visible. This increase in hooking borderline threats is not necessarily counterproductive, though maintaining close surveillance over the significant threats is arguably more useful.

The most important question, of course, is how decluttering influenced actual air warfare performance. Decluttering did not influence which tracks received responses, but how did decluttering influence the timeliness of those responses? Our hypothesis was that decluttering the low-threat tracks would facilitate noticing and responding to the ring crossings and threat changes of significantly threatening tracks. To test this hypothesis, overall response times for each declutter condition were submitted to a one-way repeated measures ANOVA. Decluttering significantly reduced response times, $F(2, 52) = 3.5$, $p = 0.037$ (Figure 6). Response times were 25% faster in the high declutter condition than in the no declutter condition. In a separate one-way repeated measures ANOVA of response times to only the significantly threatening tracks, response times were 28% faster in the high declutter condition than in the no declutter condition, $F(2, 52) = 3.6$, $p = 0.035$.

Response times to only the borderline threat tracks were not significantly different between declutter conditions, $F(2, 40) = 1.2$, $p = 0.31$. Note that the reduced degrees of freedom in this analysis was because six participants responded to no borderline threatening tracks in one or more declutter conditions. Response times to low-threat tracks could not be analyzed because so few participants ever responded to these tracks. The infrequency of responses to borderline and low-threat tracks limited their impact on the overall results. Decluttering substantially improved response times in most cases.

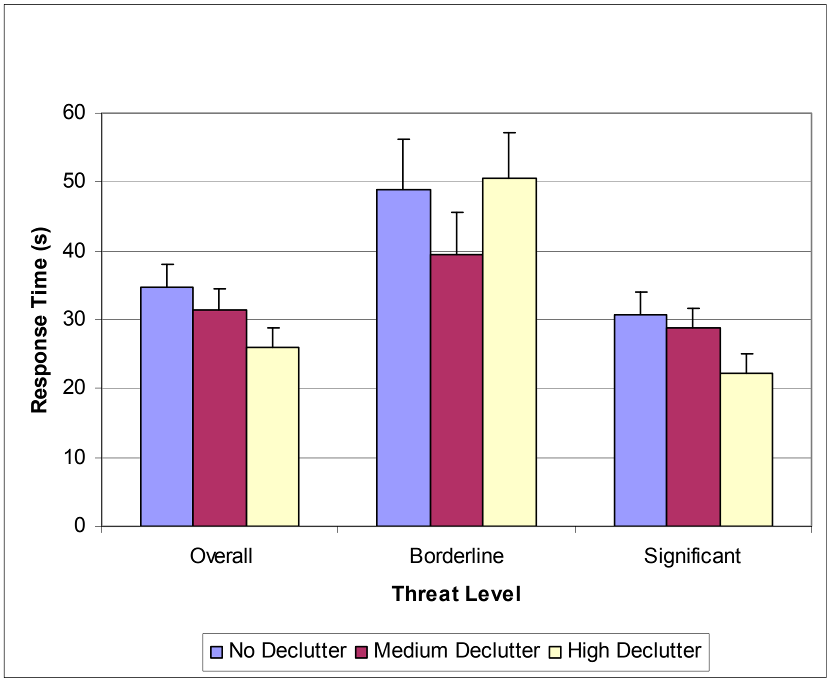

Figure 6. Mean response times (overall and within each threat category).

It is interesting that the response times were as long as they were—the mean response time was 31 seconds. These long times suggest that participants did not, or could not, continuously and rapidly sweep around the display. Instead, monitoring for significant threats and critical events must have required careful evaluation and close observation of individual tracks, which sometimes delayed the detection of other critical events.

To investigate the effect of decluttering more closely, we split the response times based on the type of significant event that prompted them: ring crossings or threat-level increases. The overall response times for each declutter condition and significant event type were submitted to a two-way repeated measures ANOVA. Response times to ring crossing events (27 seconds) were on average faster than threat-level increase events (40 seconds), $F(1, 26) = 57.6$, $p < 0.0001$. However, the same pattern of response times was found for ring crossing events and threat-level increase events, indicating that decluttering facilitated the detection of relatively salient and predictable ring crossing events and relatively less obvious threat change events by approximately the same amount. The main effect of declutter condition was significant, $F(2, 52) = 4.6$, $p = 0.015$, but the interaction between declutter and event type was not significant, $F(2, 52) = 0.9$.

Next, we split the response times based on the type of response: notify, query, or warn. Because these responses were designated to occur at different ranges from own ship, the three responses provided a convenient way to examine the effects of decluttering at different ranges from own ship and the center of the display. As Figure 7 shows, response times were fast and similar across declutter conditions for warnings, which occurred within 25 nautical miles of own ship. However, for the queries at 50 nautical miles and the notifies at 75 nautical miles, response times were slower and strongly influenced by decluttering. The response times were submitted to a two-way repeated measures ANOVA of response type and declutter condition. The main effect of response type was significant, $F(2, 50) = 21.1$, $p < 0.0001$, and the main effect of declutter condition was significant, $F(2, 50) = 3.9$, $p = 0.028$. The interaction of response type and declutter condition was also significant, $F(4, 100) = 2.9$, $p = 0.025$. These results indicate that even the baseline display was sufficient for monitoring tracks close to own ship, and that the real benefits of decluttering lie in facilitating the rapid detection and response to threats further away from own ship. For the peripherally located notify responses, high threshold decluttering improved response times by an impressive 44%.

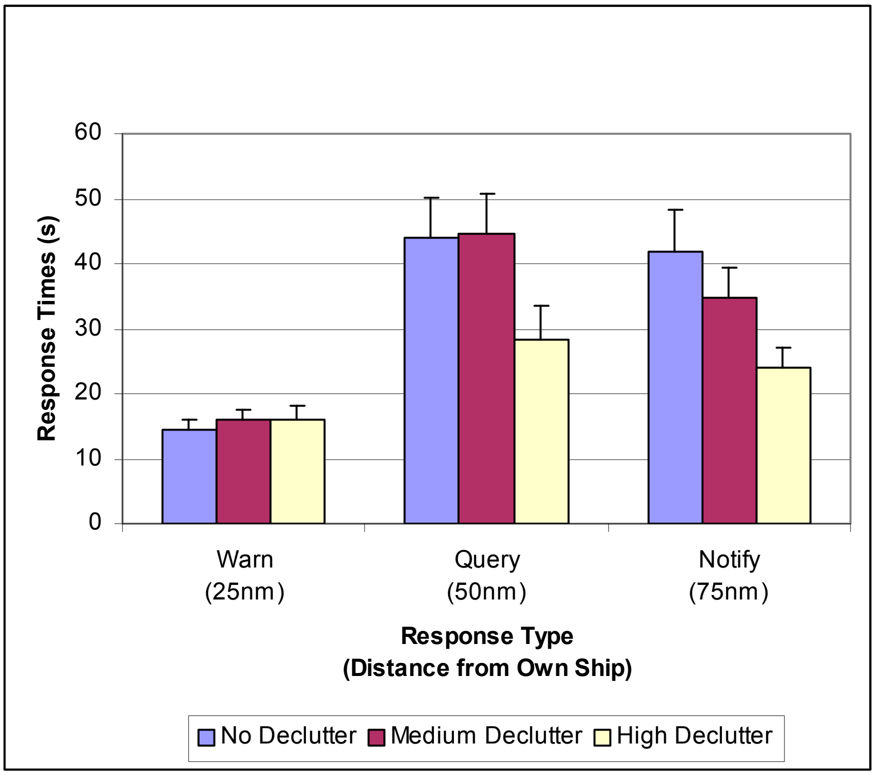

Figure 7. Effect of decluttering for different response type and distances from own ship.

Finally, we split the response times by participants' level of experience at air warfare. To perform this analysis, the two highly experienced participants were dropped due to the small sample size. This reduction left 14 very highly experienced participants and 11 moderately experienced participants. Experience level led to several differences among participants, though no differences in the effects of decluttering. First, the overall number of responses in each declutter condition and experience level were submitted to a two-way, mixed-effects ANOVA. Moderately experienced participants responded to more significant events (24) than very highly experienced participants (19), $F(1, 23) = 8.5$, $p = 0.008$. Looking separately at each level of threat, the moderately experienced participants primarily responded to more borderline events, $F(1, 23) = 4.4$, $p = 0.048$. In a similar analysis of the number of hooks, the moderately experienced participants also hooked more borderline threat tracks, $F(1, 23) = 4.2$, $p = 0.051$, and more low-threat tracks, $F(1, 23) = 5.8$,

$p = 0.024$, than the very highly experienced participants. This increase in responding was similar for all three declutter conditions. In contrast with the number of responses and number of hooks, experience level did not influence response time. In a similar analysis of response times, experience level had no main effect, $F(1, 23) = 0.08$.

The most likely explanation for these results is that the moderately experienced participants played the task more conservatively by judging more tracks to warrant responses. In the high threshold declutter condition, this higher rate of responding meant that moderately experienced participants were actually more likely than the very highly experienced participants to disregard the automation's threat assessments. Contrary to conventional wisdom (including that of the participants themselves), the less experienced participants did not doggedly follow the automation. If we assume that experience leads to greater self-confidence at the task, then the very highly experienced participants should have been the most confident, and therefore, the least likely to trust the automation (Lee & Moray, 1994). Instead, the moderately experienced participants appeared *more* skeptical of the automation than the very highly experienced participants. Furthermore, if we take the very highly experienced participants' responses as the standard, then the moderately experienced participants' conservatism and skepticism was actually somewhat counterproductive.

SUBJECTIVE MEASURES

Immediately following each scenario, 24 participants rated their subjective workload using the NASA TLX (Hart & Staveland, 1988; NASA Ames Research Center Human Performance Group, no date). The overall indices for each declutter condition were submitted to a one-way repeated measures ANOVA. The effect of declutter was not significant, $F(2, 46) = 1.1$, $p = 0.35$. We then examined only mental demand, which was the workload subscale that participants judged as most relevant to the task. In a similar analysis of only mental demand, the effect of declutter was significant, $F(2, 46) = 6.1$, $p = 0.004$. The subjective mental demand in the no declutter condition was given an average rating of 49 out of 100, while the medium and high declutter conditions were given average ratings of 40 out of 100. In terms of mental demand, decluttering reduced subjective workload by an average of 18%.

Following all three scenarios, all participants completed a questionnaire in which they rated numerous aspects of the experiment. The full report of questionnaire results appears in Feher (2003). On a scale of one to five, with five the highest rating, participants rated the task as reasonably realistic in terms of the scenarios (3.5) and their tasking (3.6), although simplified. They also rated the task to be moderately difficult (2.9 for the baseline no declutter condition), which was lower than expected, given the high degree of clutter during the scenarios. The rating may be due to several factors, including the expertise of the participants, the slow rate of change of the display, and the fact that most low-threat tracks were clustered along well-defined airlanes, which helped to declutter the display in some respects.

In a rank ordering of preference for the interfaces, participants overwhelmingly preferred the decluttered interfaces (Table 2). Twenty-five of the 27 participants preferred either one or both decluttered interfaces over the no declutter interface. Additionally, participants rated, on five-point scales, the task as less difficult with the high threshold declutter interface (1.9) and with the medium threshold declutter interface (2.2) than with the baseline interface (2.9). They also rated both declutter interfaces as more useful and better for overall situation awareness than the no declutter interface. They rated the medium threshold declutter interface as better for detecting threats and better for detecting changes in threats than the no declutter interface (all results significant by t-test, $p < 0.05$).

Table 2. Rank order preferences for interfaces.

Experience Level	No Declutter	Medium Declutter	High Declutter	Either Declutter
Overall	2	18	7	25
Moderate	0	11	0	11
High and very high	2	7	7	14

In interviews, participants claimed that these benefits reduced their workload, relieved the pressure to act and decide quickly, allowed time to concentrate on suspects, and aided SA. Comments included the following:

"I didn't have to waste time on low-threat tracks."

"I actually had more time to spend scanning the display because I could see where the high threats were."

"With no declutter, it is possible to get behind the power curve since there is a lot of mental math to keep track of (while conducting air defense warfare)."

"With decluttering, I had more time to loiter on a track of interest and put the puzzle pieces together."

"Decluttering allowed me to get ahead in my ROE—instead of behind it when mistakes are more likely to happen."

Among the two declutter interfaces, highly and very highly experienced participants split their preferences between the high declutter and the medium declutter interfaces. Moderately experienced participants overwhelmingly preferred the medium declutter interface. Similarly, highly and very highly experienced participants rated both declutter interfaces as more useful and better than the baseline interface while the moderately experience participants rated only the medium threshold declutter interface as better than the baseline interface. In effect, the moderately experienced participants appeared to take a more conservative stance toward decluttering. A common opinion was that "medium threshold declutter helped narrow down the tracks that were better candidates to recheck," while the "high threshold left me more suspicious of the decluttered tracks (leading to) greater workload." This more conservative stance matches the behavioral data on number of responses and number of hooks, but contrasts with the data on response times. Participants at all experience levels benefited similarly and solely from the high declutter interface. The medium declutter interface may have felt "safer," but it was the high declutter interface that improved response times.

Participants reported using the interfaces in the manner that we expected. Even though the participants rated the threat assessment automation as reasonably accurate (4.0 out of 5.0), and they concentrated most of their attention on the fully visible, significantly threatening tracks, they continued to intermittently sample the decluttered tracks. The result was more efficient monitoring because significant events were responded to more quickly. But this efficiency was not accompanied by any increase in automation complacency because decluttered tracks continued to be checked and verified.

CONCLUSIONS

Decluttering a naval air defense display using a heuristic threat assessment algorithm was successful in the following ways:

1. U.S. Navy experts (25 out of 27) preferred one or the other of the two declutter interfaces over a baseline no declutter interface. They rated the declutter interfaces as easier to use and better for detecting threats and maintaining SA.
2. Participants rated the overall task as easier and its mental demands as lower when using the declutter interfaces.
3. The high threshold declutter interface significantly improved the timeliness of responding to significantly threatening tracks. Responses were 28% faster to tracks that the declutter algorithm identified as significant threats, and they were 25% faster overall.
4. Despite these benefits, the declutter algorithm had little influence on which tracks received responses. In other words, there is little evidence of confirmation bias.
5. Decluttering increased SA. Participants spent significantly more time looking at the tracks that the declutter algorithm identified as significantly threatening and spent significantly less time looking at the less threatening tracks, as measured by which tracks were hooked during the scenarios.

In important respects, the declutter algorithm performed quite well, even though it used relatively simple heuristics to assess threat. Rather than attempting to strictly rank tracks from most threatening to least threatening, it merely attempted to categorize tracks as significant, borderline, or low threat. At this less ambitious task, the algorithm was reasonably successful in that it reasonably closely matched the judgments of participants. In the no declutter condition in which the algorithm rated tracks but did not influence the display, 5% of participants' responses, on average, were to low-threat tracks, 17% of participants' responses were to borderline threat tracks, and fully 79% were to significantly threatening tracks. Most importantly, this good, but not perfect, categorization performance by the declutter algorithm enabled the task performance benefits described above. These benefits, we believe, derive from the way in which the automation was designed into the interface and used by the participants. It suggested where users should focus their attention while still allowing them to scan the entire situation and respond as they saw fit.

The response time benefits for the high threshold declutter interface are easy to understand. For the tracks that the algorithm assessed as significant threats, ring-crossing events were clearly visible because these events were the only fully visible tracks on the display. Threat-level increase events were also easy to observe because these events typically caused a decluttered track to turn fully visible. Even if a participant did not see the actual change in status, once a track became fully visible, it was easy to notice quickly. On the rare occasion when participants determined that a decluttered track was a significant threat, response times were substantially longer. However, these longer times were about the same length as those in the baseline condition. Therefore, the high threshold declutter interface led to substantial response time benefits when the participants and automation agreed, and led to no delays when they disagreed.

In contrast, for the medium threshold declutter interface, detecting ring crossings was more difficult because of substantially more fully visible tracks to monitor, only some of which were significantly threatening. Similarly, threat-level increases that turned a borderline track into a significant threat would have been difficult to detect because the borderline tracks were already fully visible. Consequently, this interface required close monitoring of the borderline tracks. These extra burdens on participants in the medium threshold declutter condition may explain the relative lack of response time benefits.

Participants were split, however, in their preference for the medium and high threshold declutter interfaces. The medium declutter interface was viewed as safer, and it fit with a more conservative stance toward decluttering. The less experienced participants overwhelming preferred this interface. Like these participants, our hypothesis going into the experiment had been that the medium threshold declutter interface represented a sensible compromise between the “aggressive” decluttering of the high threshold declutter interface and the baseline no declutter interface. By leaving borderline threats fully visible, participants would never miss a threat, and yet would still realize benefits from monitoring a reduced set of fully visible tracks.

However, behavioral evidence to support this conservative stance is minimal. Instead, it appears that the actual benefits lay with the high threshold declutter interface. As it turned out, participants could focus easily on the unambiguous threats of the fully visible tracks and still maintain a broader awareness of additional potential threats.

Of course, an important limitation of these findings is the short-term nature of the scenarios, which limits our ability to generalize the findings to operational settings. In practice, users stand watch for hours at a time, over periods of weeks and months, and significant threats are typically few and far between. Whether these differences in task duration and threat frequency would change the results of the study are unknown. Users who guard against automation complacency in the short term might be lulled into complacency over the long term. Future research and development of the declutter concept must consider this issue. For instance, it may be possible to implement design features to help guard against this potential hazard. One such possibility, which resonated with participants during their interviews, would be to allow users to modify the declutter threshold to suit the situation. Users could set the threshold low during relatively benign situations to see any potential threats, and set the threshold high during more tense situations to focus on the more significant threats.

Finally, while the current experiment demonstrated the basic benefits of decluttering, the declutter interface could be improved in numerous ways. Appendix A lists many suggestions from the participants.

REFERENCES

Baddeley, A. D. 1972. "Selective Attention and Performance in Dangerous Environments," *British Journal of Psychology, 63,* 537–546.

Feher, B. A. 2003. "Geoplot Declutter Interface Usability Evaluation." SSC San Diego TR 1917. Space and Naval Warfare Systems Center, San Diego, CA.

Fisher, D. L., B. G. Coury, T. O. Tengs, and S. A. Duffy. 1989. "Minimizing the Time to Search Visual Displays: The Role of Highlighting," *Human Factors, 31,* 167–182.

Hart, S. G. and L. E. Staveland. 1988. "Development of a Multi-Dimensional Workload Rating Scale: Results of Empirical and Theoretical Research." In *Human Mental Workload*, P. A. Hancock & N. Meshkati, Eds. Elsevir Science Publishers B. V., Amsterdam, North-Holland.

Johnson, W. W., M. Liao, and S. Granada. 2002. "Effects of Symbol Brightness Cueing on Attention during a Visual Search of a Cockpit Display of Traffic Information," *Proceedings of the Human Factors and Ergonomics Society 46th Annual Meeting* (pp. 1599–1603). Human Factors and Ergonomics Society, Santa Monica, CA.

Lee, J. D. and N. Moray. 1994. "Trust, Self-Confidence, and Operators' Adaptation to Automation," *International Journal of Human–Computer Studies, 40,* 153–184.

Liebhaber, M. J. and B. Feher. 2001. "Description and Evaluation of an Air Defense Threat Assessment Algorithm." Pacific Scientific and Engineering Group, Inc. San Diego, CA..

Liebhaber, M. J., D. A. Kobus, and B. A. Feher. 2002. "Studies of U.S. Navy Air Defense Threat Assessment: Cues, Information Order, and Impact of Conflicting Data." SSC San Diego TR 1888. Space and Naval Warfare Systems Center, San Diego, CA.

Marshall, S. P., S. E. Christensen, and J. A. McAllister. 1996. "Cognitive Differences in Tactical Decision-Making," *Proceedings of the 1996 Command and Control Research and Technology Symposium* (pp. 122–132). Department of Defense, Command and Control Research Program, Washington, DC.

Morrison, J. G., R. T. Kelly, and S. G. Hutchins. 1996. "Impact of Naturalistic Decision Support on Tactical Situation Awareness," *Proceedings of the Human Factors and Ergonomics Society 40th Annual Meeting* (pp. 199–203). Human Factors and Ergonomics Society, Santa Monica, CA.

NASA Ames Research Center Human Performance Group. No date. Task Load Index (TLX) V 1.0 Users Manual. Available at http://iac.dtic.mil/hsiac/docs/TLX-UserManual.pdf

Nugent, W. A. 1996. "Comparison of Variable Coded Symbology to a Conventional Tactical Situation Display Method," *Proceedings of the Human Factors and Ergonomics Society 40th Annual Meeting* (pp. 1174–1178). Human Factors and Ergonomics Society, Santa Monica, CA.

Osga, G. and R. Keating. 1994. "Usability Study of Variable Coding Methods for Tactical Information Display Visual Filtering." NRaD Technical Report 2628. *Naval Command, Control and Ocean Surveillance Center RDT&E Division, San Diego, CA.

Parasuraman, R. and V. Riley. 1997. "Humans and Automation: Use, Misuse, Disuse, Abuse," *Human Factors, 39,* 230–253.

Parasuraman, R., T. B. Sheridan, and C. D. Wickens. 2000. "A Model for Types and Levels of Human Interaction with Automation," *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans, 30,* 286–297.

Posner, M. I. 1980. "Orienting of Attention," *Quarterly Journal of Experimental Psychology, 32*, 3–25.

Rensink, R. A. 2002. "Change Detection," *Annual Review of Psychology, 53,* 245–277.

Schultz, E. E., D. A. Nichols, and P. S. Curran. 1985. "Decluttering Methods for High Density Computer-Generated Graphic Displays," *Proceedings of the Human Factors and Ergonomics Society 29th Annual Meeting* (pp. 300–303). Human Factors and Ergonomics Society, Santa Monica, CA.

St. John, M., B. A. Feher, and J. G. Morrison. 2002. "Evaluating Alternative Symbologies for Decluttering Geographical Displays." SSC San Diego TR 1890. Space and Naval Warfare Systems Center, San Diego, CA.

St. John, M. and D. I. Manes. 2002. "Making Unreliable Automation Useful," *Proceedings of the Human Factors and Ergonomics Society 46th Annual Meeting* (pp. 332–336). Human Factors and Ergonomics Society, Santa Monica, CA.

St. John, M., D. I. Manes, and G. A. Osga. 2002. "A 'Trust but Verify' Design for Course of Action Displays," *Proceedings of the 2002 Command and Control Research and Technology Symposium*. Department of Defense, Command and Control Research Program, Washington, DC.

St. John, M., H. M. Oonk, and G. A. Osga. 2000. "Designing Displays for Command and Control Supervision: Contextualizing Alerts and 'Trust but Verify' Automation," *Proceedings of the Human Factors and Ergonomics Society 44th Annual Meeting* (pp. 646–649). Human Factors and Ergonomics Society, Santa Monica, CA.

Treisman, A. M. and G. Gelade. 1980. "A Feature-Integration Theory of Attention," *Cognitive Psychology, 12,* 97–136.

Yeh, M. and C. D. Wickens. 2001. "Attentional Filtering in the Design of Electronic Map Displays: A Comparison of Color Coding, Intensity Coding, and Decluttering Techniques," *Human Factors, 43,* 543–562.

---

*now SSC San Diego

APPENDIX A

SUGGESTIONS FOR TOOL/INTERFACE IMPROVEMENTS

A.1 DECLUTTER

- Manual ability to declutter individual tracks
- User-modifiable declutter threshold level
- Modifiable weights in the declutter algorithm
- Better change of declutter status indicators
- Procedures for sharing declutter information across the Battle Group

A.2 GENERAL TO AIR WARFARE

- Extra factors in the algorithm for change in altitude or speed with notification of user
- A user-assignable suspected hostile symbol
- Mission commander can designate tracks-of-interest that are highlighted on all displays
- Track information provided by pre-hook (roll-over)
- Means of keeping track of contacts acted on
- Territorial airspaces and air lanes

The suggestion to better indicate changes in threat and declutter status is especially interesting. During the experiment, threat-level increases that changed a track from nonsignificantly threatening to significantly threatening produced a relatively salient change in visibility—from a faded decluttered symbol to a fully visible symbol. However, these relatively large visibility changes still led to fairly long response times. It seems likely that in many cases, participants did not actually observe the status and symbol changes, but found the already changed tracks during their normal scanning around the display. Research in change blindness (Rensink, 2002) supports the idea that small changes in the display are difficult to observe unless the participants happen to be directly attending at the moment of change. The participants call for more salient and longer lasting indicators of threat changes in the display matches recent empirical findings that "change history" tools that preserve a record of important changes can be very useful for maintaining and re-acquiring situation awareness (Smallman & St. John, 2003).

Another important suggestion was to add a "response manager" to the interface. The concept of a response manager is to maintain a visible record of responses taken toward each track and to recommend appropriate responses. The benefits of well-designed response management tools are well-documented (Morrison, Kelly, & Hutchins, 1996; St. John, 1998; St. John, Manes, Moore, & Smith, 1999).

APPENDIX A REFERENCES

Morrison, J. G., R. T. Kelly, and S. G. Hutchins. 1996. "Impact of Naturalistic Decision Support on Tactical Situation Awareness," *Proceedings of the Human Factors and Ergonomics Society 40th Annual Meeting* (pp. 199–203). Human Factors and Ergonomics Society, Santa Monica, CA.

Rensink, R. A. 2002. "Change Detection," *Annual Review of Psychology, 53,* 245–277.

Smallman, H. S. and M. St. John. 2003. "CHEX (Change History EXplicit): New HCI Concepts for Change Awareness," *Proceedings of the 46th Annual Meeting of the Human Factors and Ergonomics Society* (pp. 528–532). Human Factors and Ergonomics Society, Santa Monica, CA.

St. John, M. 1998. "The Impact of a Response Management Tool on Air Warfare Tactical Decision Making," *Proceedings of the 20th Annual Conference of the Cognitive Science Society* (pp. 1013–1018). Lawrence Erlbaum Associates, Mahway, NJ.

St. John, M., D. I. Manes, R. A. Moore, and C. A. P. Smith. 1999. "Development of a Naval Air Warfare Decision Support Interface Using Rapid Prototyping Techniques," *Proceedings of the 1999 Command and Control Research and Technology Symposium*. Department of Defense, Command and Control Research Program, Washington, DC.

INITIAL DISTRIBUTION

20012	Patent Counsel	(1)
202753	Archive/Stock	(6)
202752	Library	(2)
2027	M. E. Cathcart	(1)
20275	F. F. Roessler	(1)
202753	D. Richter	(1)
244210	B. Feher	(5)
24410	J. G. Morrison	(5)

Defense Technical Information Center
Fort Belvoir, VA 22060-6218 (4)

SSC San Diego Liaison Office
C/O PEO-SCS
Arlington, VA 22202-4804 (1)

Center for Naval Analyses
Alexandria, VA 22311-1850 (1)

Office of Naval Research
ATTN: NARDIC
Philadelphia, PA 19111-5078 (1)

Government-Industry Data Exchange
Program Operations Center
Corona, CA 91718-8000 (1)

Pacific Science and Engineering Group (15)
San Diego, CA 92122