The Précis of Project Nemo, Phase 1: Subgoaling and Subschemas for Submariners

Wayne D. Gray*, Susan S. Kirschenbaum§, and Brian D. Ehret*

*Human Factors & Applied Cognition
George Mason University; m/s 3f5
Fairfax, VA 22030
{gray, behret}@gmu.edu

§Naval Undersea Warfare Center Division
Code 2214, Building 1171/1
Newport, RI 02841
kirsch@c223.npt.nuwc.navy.mil

Abstract

Project Nemo examines the cognitive processes and representational structures used by submarine Commanders while attempting to locate an enemy submarine hiding in deep water. This report provides a précis of the first phase of this effort. Protocol data, collected from commanders with 20 years of submarine experience, have been transcribed and analyzed. The data suggest a shallow goal structure with a basic level of subgoals that are used by all Commanders throughout the task. Relatively few operators are required for each subgoal. The results are congruent with a schema theory interpretation in which the process of schema instantiation provides the control of cognition.

Introduction/Significance

The submarine Approach Officer (AO) performs the role of senior decision maker during an encounter with a hostile target. His[1] job in identifying and locating enemy submarines is difficult, interesting, and important. It is difficult because it requires locating an enemy who is hidden in the vast and acoustically uncertain environment of the ocean. It is interesting to cognitive scientists as the expertise of AOs is similar to, but different from other, better studied, types of expertise. Finally, whatever the changing nature of warfare, locating hostile targets in an uncertain ocean environment is an important job that is important to do well.

This short report focuses on our interests as cognitive scientists in the AOs' task. It is the précis of a much longer report (Gray, Kirschenbaum & Ehret, 1997) that summarizes the protocol analysis phase of an attempt to understand the AOs' task by building computational cognitive models. These analyses have led us to postulate that a cognitively plausible simulation of AO behavior must incorporate the following:

1) A relatively flat goal hierarchy. Our analyses suggests 2-3 levels, certainly whether it is two, three, or even four levels it cannot be 7, 8, or 9 levels.
2) A small set of basic level operators. Our analyses suggested 9 operators. Some of these may need to be combined, whereas others may need to be subdivided. Whatever the exact number, the functionality of these 9 operators must be maintained.
3) A subgoal level that, together with the operators, may be regarded as the basic unit of AO expertise. These subgoals are used with little modification in the two main higher goals LOCALIZE-MERC and LOCALIZE - SUB.
4) Relatively few operators per basic level goal. Our mean was around four.

These conclusions have lead us to postulate a large, domain-specific schema structure (VanLehn, 1989) that resides in long term memory. Assuming an ACT-R (Anderson, 1993) type production rule architecture, each goal cycle must conclude by returning a change in value for one or more schema slots. The schema is reevaluated and the results of this reevaluation are used to select the next production rule that fires. This requirement implies that the reevaluation is subsymbolic and is not, typically, a conscious process. In modeling terms, it directs the selection of production rules but is not itself modeled by production rules.

Having given our conclusions above, in the next section we introduce the task. This is followed by a brief discussion of the empirical and analytic methodology used to collect and analyze the data. We then summarize the data that supports points 1-4. The schema representation and its use in the control of cognition are hypotheses based upon interpreting our conclusions in the light of existing cognitive theory. We conclude this paper with a short discussion of our plans for testing and verifying these hypotheses.

The Mission, Scenario, and Situation assessment

The part of the battle that we have studied is the situation assessment phase. Situation assessment is a part of a scenario which, in turn, is a part of the AO's mission. The mission defines why the submarine, and its crew, is in the water in the first place. For example;

---

[1] All crew and officers onboard submarines are men.

> A war has broken out between Russia and Ukraine over control of the former Soviet Navy. Each nation is trying to get the U.S. involved by attacking U.S. shipping with submarines and then blaming the other. You are patrolling in the eastern Caribbean sea. Your mission is to protect U.S. shipping lanes. Your orders are to search and destroy any enemy submarines. The month is April and the sea state is 2.

A patrol defined by such a mission could take weeks or even months to complete. Within this larger mission, we have defined scenarios that begin when the AO assumes control and receives a status report. For example;

> STATUS REPORT. You hold a merchant bearing 300° but no subsurface contacts. The broad-band range of the day for a subsurface contact is 3-6 thousand yards. Intelligence has indicated an Alfa-type formerly Soviet Submarine is patrolling in the area. There are no friendly submarines in your patrol area.

This scenario involves searching until the enemy is detected, obtaining a solution on the enemy, and moving to firing point procedures (i.e., ready to engage the enemy). A scenario could take hours or even days to complete.

Situation assessment is the part of the scenario we have studied. It begins after the enemy submarine is detected by ownship (OS). It consists of *determining a solution* that identifies how far the enemy is from OS (range in thousands of yards), its course, bearing from OS, and speed. When the AO is confident of his solution, he enters it in his computer and goes on to the next phase of the exercise (firing point procedures). The situation assessment phase lasts from tens of minutes to several hours.

Tools for Situation Assessment

The AO has a variety of tools to detect and locate the enemy. First, there are a variety of passive sonar detectors that can listen for different sound frequencies. Knowing when to use what detector to search which frequency is an important component of the AO's expertise. Second are changes to OS speed, course, or depth. The action chosen depends upon what problem the AO is trying to solve. For example, if he believes that OS noise is masking the signal from the enemy SUB, he may slow down. If he believes that the sound transmission properties of the water are better at a different depth, he may raise or lower the depth of OS. If he believes that the data he has gotten is as good as it is going to get, he may change course to gain additional data by *viewing* the situation from a different angle. The third set of tools involves statistical processing of the data. Processed data is provided by various algorithms. Some of these are analogous to multivariate analysis. They generate parameters that optimize the current best fit of the data to variables such as range, speed, bearing, etc. As with any such analysis, the algorithms will happily generate best-estimates to any set of data; however, the estimate for any given parameter may be absurdly unrealistic or subtly wrong. A large part of the AO's art involves comparing his estimates to those generated by these algorithms.

Situation assessment requires the AO to intermix and interplay these three sets of tools (sonar, changing OS, and statistically processed data) to determine a solution. The interplay among these three changes over time as noisy and uncertain data is transformed by the AO into familiar and reasonable patterns.

Summary

Although the AO works in a unique environment, his situation assessment task is a prime example of information processing. The rationality principle applies and much of his behavior can be characterized as search in a problem space (Card, Moran & Newell, 1983; Newell & Simon, 1972). Like most studied forms of expertise, the AO knows how to apply the tools available to transform the current state of the problem into a new state, until the desired, or goal state is reached. However, unlike many types of studied expertise, a large part of what the AO does well, resides in his ability to recognize a solution when he has found one.

This characterization of the AO's task implies an approach to analyzing the process of situation assessment that looks at (1) the elementary information operators the AO uses, (2) the constraints (goals and subgoals) under which they are applied, and (3) the nature of the schemas that guide the selection of goals and enable the AO to recognize a well-formed solution.

Methodology and Empirical Data

The empirical data was collected from current or former submarine AOs as they attempted to find a solution to one or more scenarios presented on the Combat Systems Engineering and Analysis Laboratory (CSEAL), a simulation residing at the Naval Underwater Weapons Center Division Newport (NUWCDIVNPT) in Newport, RI.

Participants

All subjects used in the study were Naval Commanders with the exception of one subject who was a Captain (select). The average number of years in the Navy was 20.2 and the average number of years spent at sea was 11.4. The average number of months since subject's last tour of duty prior to the study was 10.3 months.

Scenarios

All AOs were given the mission statement discussed above. This was followed by a status report for each scenario. The AO was then provided information regarding the course, speed, and depth of OS at the start of each scenario; e.g., course 190°, speed 9 knots, and depth 400 feet. Thus, at the beginning of a scenario the AOs had confirmed contact and bearing information on a merchant, knew the course, speed, and depth of OS, and had not detected, but had intelligence information regarding the nearby presence of a hostile, "alpha" class submarine.

Four scenarios were developed and used throughout the two stages of the main study. In developing the scenarios, the researchers solicited the advice and assistance of various NUWCDIVNPT personnel on issues relating to submarines, oceanography, and CSEAL. The resulting scenarios were tested and refined over the course of a pilot study.

CSEAL Simulation

CSEAL is a high-fidelity simulation of the ocean environment that was developed to facilitate the creation of various algorithms and instrumentation relevant to the detection of acoustic signals in the ocean (an acoustically noisy environment).

CSEAL supports a dynamic OS whose speed, depth, and course can be changed at any time. CSEAL keeps track of these changes and dynamically recomputes the range, bearing, etc of the targets. In addition, the signal-to-noise ratio (SNR) and path of the signal from the target is continually updated to reflect the relative position of the two vehicles in the simulated ocean. Hence, although we used a limited set of scenarios, after the first OS maneuver, each AO had a functionally different scenario.

Data Collection

All AO interactions with CSEAL were mediated through a computer operator who we refer to as *own ship operator* or OS-op. This arrangement had three advantages. First, the AOs did not have to learn to use the CSEAL interface. Second, on board real submarines, AOs do not have a workstation and all information they receive comes from various crew members. The crew members interact with the instruments and feed data to the AO. Having the OS-op use the computer mimicked these on board arrangements. No AO objected to this arrangement and no AO requested to use CSEAL himself. Third, since we were interested in what information the AOs wanted, when they wanted it, and when they received it, going through the OS-op meant that much information seeking and receiving was verbalized and therefore recorded. Videotapes recorded the screen and all AO and OS-op dialogue.

Procedures

Our study proceeded through two stages with 3 AOs at each stage. The major difference between stages was the use of four 30 min scenarios in stage 1 versus two 60 min scenarios in stage 2. Once the scenario started, the AO continued in free play mode until the time expired or until he ended the situation assessment phase by beginning the next phase, firing point procedures. All sessions lasted between 2 and 3 hours.

Protocol Encodings

For the 3 AOs in stage 1 (with four, 30 min scenarios), we have encoded and analyzed their last 2 scenarios. For the 3 AOs in stage 2 (with two, 60 min scenarios), we selected the one scenario in which we had the least technical problems (i.e., problems with CSEAL) and/or in which the AO achieved the best solution (closest to truth).

Operators

Each utterance and action made by the AO and OS-op was encoded using three main categories of operators: information, ship operations, and other. The *information* category includes one operator involved in seeking and two concerned with acquiring information about OS, the merchant vessel, or the hostile submarine. *Ship operations* includes four operators used by the AO to manage target information or change the position of OS. The *otherOps* category includes one operator related to the usability of the displays, as well as one that encodes all non-task related utterances. Neither of these otherOps operators are related to the situation assessment task, per se.

Goals

The AO's overall goal, as provided by the scenario statement, is to destroy the hostile submarine. To accomplish this goal the AO must first detect the submarine on sonar and then determine its bearing, range and course. Although the AO was primarily focused on locating the hostile submarine, an important component of situation assessment entails considering the hostile submarine in relation to the merchant. (Is the hostile submarine approaching the merchant? Is the merchant in the line of fire?) Also of importance is keeping track of the status of OS instruments and OS position. We have found a three-level goal hierarchy to be sufficient to characterize the basic situation assessment process. Our five level-1 goals are: DETECT-SUB, LOCALIZE-MERC, LOCALIZE-SUB, ENGAGE-SUB, and STATUS-OS. Of these, LOCALIZE-MERC and LOCALIZE-SUB form the core of the situation assessment phase. DETECT-SUB occurs prior to the start of situation assessment; ENGAGE-SUB marks the end of situation assessment; and STATUS-OS can be performed at any time. There are thirteen level-2 subgoals that fall into six major types (detailed definitions, examples, and discussions of all goals and operators are provided in Gray et al., 1997). The third level in our hierarchy includes one subgoal related to software usability issues in addition to the relatively rare instances in which one of our 13 level-2 subgoals are encoded as a subsubgoal for another.

Pushing, Popping, and *Successful* Goal Accomplishment. Our analysis presumes a simple goal stack mechanism in which the currently active goals are organized in a last-in, first-out basis. When the AO is working on a level-2 subgoal, the goal stack consists of two goals, with the level-1 goal on the bottom and the subgoal on the top. Before another level-2 subgoal can be pushed onto the goal stack, the first subgoal must be popped off: only one level-2 subgoal can be on the goal stack at any one time. Sequences of level-2 subgoals are not obtained by pushing a predetermined series of subgoals onto the goal stack. Rather, sequences emerge from reevaluating the schema between the popping of one subgoal and the selection (pushing) of the next. (We have more to say on the topic of *reevaluating the schema* in the discussion section.)

In our encodings, a goal may be popped in one of two circumstances. First, if it is interrupted by another goal (e.g., while trying to get a more exact bearing rate on the merchant, the AO notices another target being drawn on the waterfall display). This is an event-driven intrusion into the AO's goal-stack. Second, it may be popped if it completes. However, a *completed* goal (or subgoal) is not synonymous with the colloquial definition of a *successful* goal.

A completed goal is one that returns some information to the schema regarding the purposes of the goal. For example, the level-1 goal, `DETECT-SUB`, may be pushed to mark the AO's search through sonar displays for any trace of the hostile submarine. If no trace is found, then `DETECT-SUB` may be popped and, e.g., `LOCALIZE-MERC` pushed. Although, `DETECT-SUB` did not find a hostile submarine, the search yielded the information that at the present time, to the limits of the instrumentation (various sonar detectors) there was no hostile submarine within range. Although in the colloquial sense, the goal, `DETECT-SUB`, was not successful; in our sense it was successfully completed.

The mission is successful if the enemy submarine is detected and localized with a solution good enough for the AO to move into the firing point procedures phase. Any goal pushed during this mission is either completed or interrupted. As in the `DETECT-SUB` example above, any completed goal may be pushed again later, to obtain updated information regarding the purposes of the goal.

Summary

Three types of data were discussed in this section; operators, goals and subgoals, and the information returned to the AO's schema when a goal is completed. These data are *process data* and as such differ from more typical *outcome data* or *accuracy data*; that is, rather than being content to provide some overall summary measure of AO performance, these analyses attempt to provide an explanation of the AO's observed actions and overt behavior in terms of more elementary information processes.

Data Analysis and Results

Validation of encodings

After being transcribed and segmented, each protocol was independently encoded by two individuals or teams, using the operators described above. Cohen's Kappa corrects for chance agreement and yields a z-score that estimates the probability of getting such agreement by chance. The Kappa scores for our matches ranged from 0.58 to 0.76 with z-scores that ranged from 16.5 to 30.0. These are highly significant; e.g., a z of 16.5 would be expected by chance, less than 0.1% of the time ($p < .001$). After the IRR was calculated, the two encoders met and resolved all discrepancies.

Encoding the goals required greater domain knowledge as well as greater complexity of encoding. Hence, six of the nine trials were encoded during sessions in which three encoders participated. Each encoder had worked on the project for a minimum of 18 months and each had been actively engaged in the transcription, segmentation, and encoding of the operators for each trial. Before convening these *consensus* sessions (sometimes referred to as *delphi* sessions) for encoding the goals, we estimate that each of the three encoders had spent a minimum of 25 hours on each of the 9 trials. During the consensus sessions, the encoders had access to the videotape, truth files, the transcribed and segmented protocols, as well as the encoded operators.

Following the consensus encoding of the six trials, the remaining three trials were independently encoded by two encoders. The inter-rater reliabilities of those encodings were significantly greater than chance ($p < .001$ or less). However, as expected, the mean agreement for goals (44%) was less than that for operators (75%).

Cleaning the encodings

After differences had been reconciled among encoders we had 421 goals and 2,882 operators to encode the 9 scenarios from our six AOs. For operators the largest category of encodings was *OtherOps*. These 1,445 operators represented half of the 2,882 total operators. As these operators do not contribute to the AO's situation assessment process, we have dropped them from further analysis. Likewise, non-task subgoals and their operators have also been dropped (refer to Gray et al., 1997 for more details). We refer to the product of this effort as our *clean set*. This clean set consists of 397 goals and 1,269 operators. It is the set of encodings used in all subsequent analyses.

Goals and Subgoals[2]

Our focus here is in making sense of the differences within and between AO-trials in the two main situation assessment goals, `LOCALIZE-MERC` and `LOCALIZE-SUB`. Our conclusions are that the data seem inconsistent and random at the goal level, but consistent and meaningful at the subgoal level. (In the Summary section, we will try to make sense of the entire pattern of data, goal as well as subgoal, by appealing to a schema-based explanation of the control of cognition.)

There is nothing consistent about the data when viewed in terms of `LOCALIZE-MERC vs. LOCALIZE-SUB` goals. For example, across all AO-trials there are 33 `LOCALIZE-MERC` goals vs. 44 `LOCALIZE-SUB` goals. Likewise, the total duration per AO-trial (3:41 min vs 23:40 min), and mean number of subgoals (see left-side of Figure 1) per `LOCALIZE-MERC vs. LOCALIZE-SUB` goals differ, as does the number of operators per goal (208 versus 906).

`LOCALIZE-MERC` and `LOCALIZE-SUB` also differ in how they are used by individual AOs. For example, subject 5 in each of his two 30 min trials (s05tr3 and s05tr4), pushed more `LOCALIZE-SUB` and `LOCALIZE-MERC` goals (13 in each trial) than did subjects 8 (s08tr2), 9 (s09tr2), and 10 (s10tr1) in their one 60 min trial (12, 5, and 8

---

[2] For ease of exposition, level-1 goals will be referred to simply as *goals*; level-2 subgoals as *subgoals*.

respectively). Number was not correlated with time as, e.g., s05 spent a total of 24:12 and 23:30 min working on these goals; whereas s08, s09, and s10 spent 44:04, 28:04, and 19:48 min respectively.

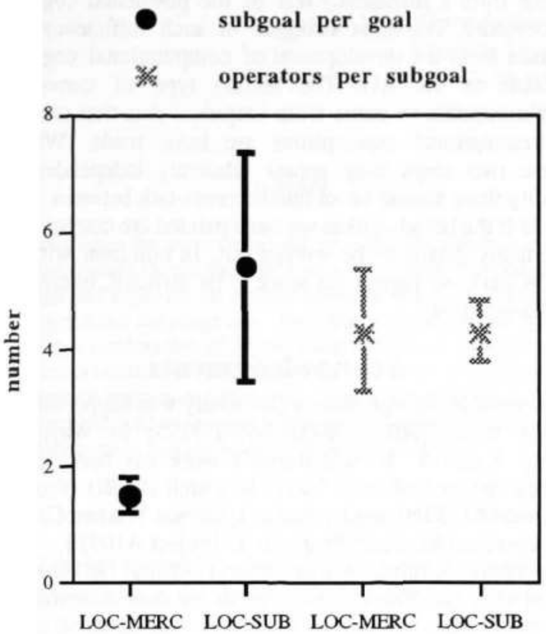

Figure 1: Mean number of subgoals per goal (left two) and mean number of operators per subgoal (right two) for LOCALIZE-MERC and LOCALIZE-SUB. Also shown are the 95% confidence intervals for the standard error of the mean.

Fortunately, the apparent instablity and randomness of the data at the goal level disappears when viewed at the subgoal level. For example, the number of operators per subgoal does not differ for subgoals of LOCALIZE-SUB versus LOCALIZE-MERC (see right-side of Figure 1). This similarity at the subgoal level holds for individual subgoals, as well as the overall duration of subgoals per goal. Even more striking, as shown in Figure 2, within AO-trials, the number of operators per subgoal does not vary with goal (LOCALIZE-MERC vs LOCALIZE-SUB). (The plot of duration of subgoals by AO-trial, not shown, yields a similar overlap between goals.) Hence, the number of operators per subgoal as well as the duration of subgoals appears to be constant between LOCALIZE-MERC vs LOCALIZE-SUB, both over all AO-trials (e.g., see right-side of Figure 1), as well as within AO-trials (e.g., see Figure 2). As would be expected by these analyses, the correlation between time per LOCALIZE-MERC and number of subgoals is high (r = 0.79); as is that between time per LOCALIZE-SUB and number of subgoals (r = 0.82).

These data are consistent in suggesting that the variability observed in the time to perform level 1 goals can be attributed to systematic differences in the number of subgoals used by each AO. Given that (1) these data were collected via talk-aloud procedures while the AOs used a free ranging simulation, (2) the task is a knowledge intensive attempt to direct search in an uncertain and noisy environment, and (3) over half of the encoded utterances had nothing to do with the task (OtherOps), the fact that so much of the time to do a goal can be explained by the number of subgoals involved seems astounding. This suggests that expert performance in a complex environment does yield to cognitive task analysis techniques and that the subgoal level captures both the consistency and differences between goals and experts.

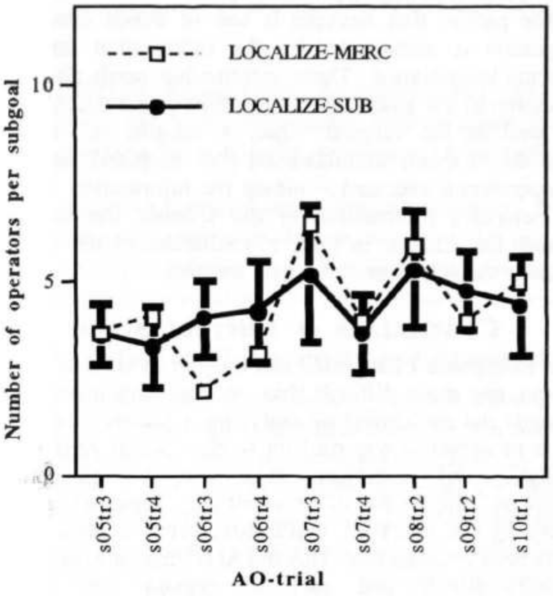

Figure 2: For each goal, the figure shows the mean number of operators per subgoal for each AO trial. The 95% confidence intervals for the standard error of the mean are shown for LOCALIZE-SUB only. The confidence intervals for LOCALIZE-MERC are not shown as they tend to be larger and, for three AO-trials, would extend beyond the graph (primarily due to the low number of subgoals, 2-3, per LOCALIZE-MERC in those trials).

Summary

The protocol analysis has provided a rich source of data that raises a number of interesting issues for the control of cognition. First, the data show a mean of 1.5 subgoals per each LOCALIZE-MERC versus 5.4 per each LOCALIZE-SUB goal (see left-hand side of Figure 1). We do not interpret these sequences as plans of 1.5 or 5.4 subgoals; rather we interpret them as an emergent property of reevaluating the schema after popping one subgoal and before pushing the next. The fact that such sequences are longer for LOCALIZE-SUB than for LOCALIZE-MERC is seen as a consequence of the greater importance of LOCALIZE-SUB. Hence, deciding what subgoal to push to obtain information about a goal can almost be considered a situated response to the current state of the schema. Each subsequent subgoal is not pushed to execute the next step in some master plan but, more simply, to obtain the *now* most important piece of missing information regarding the goal. Congruent with this position, we find no evidence of sequential dependencies between subgoals (within goals).

Second, the data does show striking consistency at the subgoal level. Within and between AO-trials, duration of subgoals, and the number of operators per subgoal, is constant. The relatively low number of operators per subgoal (mean of 4.28), suggest the execution of well-known methods (in the sense of Card, et al., 1983) rather than a search through a problem-space.

The picture that emerges is one in which control of cognition is orchestrated by the information needs of schema instantiation. These information needs direct the selection of the goal (`LOCALIZE-MERC` vs `LOCALIZE-SUB`) as well as the subgoal. Once a subgoal is launched (pushed), a short, well-rehearsed (but adaptable) sequence of operators is executed to obtain the information. Hence, by returning information to the schema, the subgoals change the schema; in turn, re-evaluation of the schema leads to the selection of the next subgoal.

Conclusions & Current Status

Our long phase 1 has ended and it is fair to say that it was longer and more difficult than we had anticipated. The rewards and excitement of analyzing a hitherto unstudied form of expertise was paid for in the coin of hard work, analysis, and reanalysis. Even with a well-explored expertise, the costs of transcribing, segmenting, and encoding six hours of multimedia protocol data would have been considerable. That the AOs' task in many ways appears simple and easy to explain was simply misleading. Like much expertise, the apparent simplicity was bought at the cost of 20+ years of hard work and training. Just as the Olympic diver makes the perfect dive appear effortless and graceful, so too did our AOs make the mental calculus required for situation assessment. In both cases, the sweat and years of hard work is concealed from the spectator.

This first phase of Project Nemo (Gray et al., 1997) tackled the task of studying an esoteric form of expertise using an approach that treats expert performance as a type of problem solving involving search in a problem space (Ericsson & Smith, 1991). At present, we believe that we have approached the limits of our current data set; that, although our AO-trial protocols have been very fruitful, they have yielded about as many insights as they will.

Key issues for the next phase revolve around our use of schemas as an explanatory construct. VanLehn (1989) discusses three phases of schema use; selection, instantiation, and execution. For our AOs, selection and execution are straight forward; both are determined by military doctrine and years of practice and experience. Situation assessment seems to involve schema instantiation. Equating schema instantiation with situation assessment requires a variety of claims regarding the processing and representation of knowledge.

The next two steps in this research effort are clear. Two types of converging evidence are required to support our protocol-based conclusions regarding the AO's possession of schema and the role that schema instantiation plays in the control of cognition. One type of evidence needs to come from a sufficiency test of the postulated cognitive processing. The most stringent of such sufficiency tests comes from the development of computational cognitive models of the AO. The second type of converging evidence needs to come from empirical data that tests the representational assumptions we have made. Whereas these two steps may appear relatively independent, in reality there should be an intense cross-talk between them. Even if the broad strokes we have painted are correct, there are many details to be worked out. In common with the AOs' task, we expect our work to be difficult, interesting, and important.

Acknowledgments

The work at George Mason University was supported by a grant from ONR (#N00014-95-1-0175) to Wayne D. Gray. Susan S. Kirschenbaum's work has been jointly sponsored by Office of Naval Research (ONR) (Program element 61153N) and by Naval Undersea Warfare Center's Independent Research Program, as Project A10328.

We thank Jennifer Bohan, Sheryl Miller, Michael R. Benjamin, and Tracy D. Hutcheson for their contributions to this project.

References

Anderson, J. R. (1993). *Rules of the mind.* Hillsdale, NJ: Erlbaum.

Card, S. K., Moran, T. P., & Newell, A. (1983). *The psychology of human-computer interaction.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Ericsson, K. A., & Smith, J. (Eds.). (1991). *Toward a general theory of expertise: Prospects and limits.* New York: Cambridge.

Gray, W. D., Kirschenbaum, S. S., & Ehret, B. D. (1997). *Subgoaling and subschemas for submariners: Cognitive models of situation assessment* (GMU-ARCH 97-01-16): George Mason University.

Kirschenbaum, S. S., Gray, W. D., Ehret, B. D., & Miller, S. L. (1996). When using the tool interferes with doing the task. In M. J. Tauber (Ed.), *Conference companion of the ACM CHI'96 Conference Human Factors in Computing Systems,* (pp. 203-204). New York: ACM Press.

Newell, A., & Simon, H. A. (1972). *Human problem solving.* Englewood Cliffs, NJ: Prentice-Hall, Inc.

VanLehn, K. (1989). Problem solving and cognitive skill acquisition. In M. I. Posner (Ed.), *Foundations of Cognitive Science,* (pp. 527-579). Cambridge, MA: The MIT Press.