

Designing and Implementing Intelligent Tutoring Instruction for Tactical Action Officers

Richard Stottler and Alex Davis
Stottler Henke Associates, Inc.
San Mateo, CA
stottler@shai.com; davis@shai.com

Sue Panichas
Northrop Grumman Corporation
Middletown, RI
susan.panichas@ngc.com

CDR Mark Treadwell
NAWC-TSD
Orlando, FL
mark.treadwell@navy.mil

ABSTRACT

The Tactical Action Officer on board a U.S. Navy Cruiser, Destroyer, or Frigate is responsible for the operation of the entire watch team manning the ship's command center. Responsibilities include tactical decision making, console operation, communications, and oversight of a variety of watchstander responsibilities in air, surface, and subsurface warfare areas. Stottler Henke, in concert with Northrop Grumman, has developed the PORTS TAO ITS, an Intelligent Tutoring System (ITS) for the instruction of Tactical Action Officers (TAOs) in training at the Surface Warfare Officers School (SWOS) using the PC-based Open-architecture Reconfigurable Training System (PORTS) as its basis.

This paper describes the instructional philosophy of the PORTS TAO ITS, resulting from close collaboration with SWOS instructors and Northrop Grumman's domain experts. The goal of the ITS is to train the student in 'command by negation,' in which watchstanders perform their duties autonomously, while the TAO supervises, intervening in order to correct mistakes and rectify omissions. The TAO must know what responsibilities belong to each watchstander, how and in what circumstances those duties are performed, and how to communicate with watchstanders to request information, acknowledge reports, and order appropriate actions.

The ITS is designed to instill and assess mastery of these TAO abilities over the course of a series of exercises which present increasingly difficult problems. These include intentional mistakes of omission or commission by automated watch team role players. When TAO actions are expected, such as when these intentional mistakes are made, the ITS provides hints, prompts, and feedback to the student, which are also summarized at the end of each exercise with a detailed debrief. These interventions are sensitive to real-time changes in the student's mastery of a wide variety of principles, which are continually assessed.

There were several challenges and lessons learned from the implementation of this ITS and the related government acquisition process. These are also detailed in this paper.

ABOUT THE AUTHORS

Richard Stottler co-founded Stottler Henke Associates, Inc., an artificial intelligence consulting firm in San Mateo, California, in 1988 and has been the president of the company since then. He has been the principal investigator on a large number of tactical decision-making intelligent tutoring system projects conducted by Stottler Henke including projects for the Navy, Army, Air Force and Marine Corps. Currently he is working on the PORTS TAO ITS for the US Navy, and a Littoral Combat Ship project. He has a Masters degree in Computer Science from Stanford University.

Alex Davis is an Artificial Intelligence Researcher at Stottler Henke Associates, Inc. He received his M.S. from the State University of New York at Buffalo in Computer Science. He has served as lead knowledge and software engineer for a variety of projects related to artificial intelligence, including simulations, behavior modeling for automated agents, and intelligent tutoring systems. Those projects include the PORTS intelligent tutoring system for Navy tactical action officers, a gaming testbed for the exploration of advanced C2 concepts, and a decision aid that incorporates behavior of adversarial forces in an unfamiliar cultural climate.

Susan Panichas, Program Manager at Northrop Grumman Corporation, has been the program manager for simulation based training systems for over ten years. She is the program manager for the PC-based Open-architecture Reconfigurable Training System (PORTS), which is the foundation for US Surface Navy operator and

tactical training programs. She has a Masters degree in Physical Oceanography from the University of Rhode Island.

CDR Mark Treadwell is the Deputy Program Director for Surface and Expeditionary Warfare Systems at the Naval Air Warfare Center, Training Systems Division (NAWCTSD) in Orlando. He is also the Government Program Manager for the Multi-Mission Tactical Trainer and all PORTS-related projects. He has a Masters degree in Engineering Management from Old Dominion University.

Designing and Implementing Intelligent Tutoring Instruction for Tactical Action Officers

Richard Stottler and Alex Davis
Stottler Henke Associates, Inc.
San Mateo, CA
stottler@shai.com; davis@shai.com

Sue Panichas
Northrop Grumman Corporation
Middletown, RI
susan.panichas@ngc.com

CDR Mark Treadwell
NAWC-TSD
Orlando, FL
mark.treadwell@navy.mil

INTRODUCTION

The mission of the Surface Warfare Officers School (SWOS) in Newport, Rhode Island is to provide professional education and training to prepare officers of the U.S. Surface Navy to serve at sea. As part of his training at SWOS, each Surface Warfare Officer learns how to "fight" his ship as a Tactical Action Officer. The TAO training consists of three months of classroom and simulator time wherein students are exposed to all elements of surface warfare; air, surface, subsurface, and amphibious operations as well as electronic and other support mechanisms. One major responsibility of a TAO is to be able to exercise command over the major systems of his ship (weapons, support platforms, radar and sonar, and navigation) during potentially hostile situations. The tactical decisions he makes during such situations can easily affect the outcome of the ship's mission, as well as have life or death consequences. In summary, to paraphrase the SWOS instructors, the TAO gathers information, analyzes it, and ensures the correct decisions are made and actions taken based on the tactical situation by issuing verbal orders and queries. Additionally TAOs "command by negation" in that much of what the Combat Information Center (CIC) team needs to accomplish in tactical situations is decided upon autonomously by individual watchstanders who also state their intentions before they execute their tasks. When these are correct, the TAO must merely acknowledge the watchstander's decision. However, if these decisions are incorrect or omitted, the TAO must negate the incorrect decision or proactively initiate the omitted actions.

PORTS DESCRIPTION

Previously, in order to address the need for watchstanders to practice tactical scenarios without requiring the use of expensive special purpose hardware, the Navy commissioned the development of the Generic Reconfigurable Training System (GRTS) now renamed the PC-based Open-architecture Reconfigurable Training System (PORTS), which replicates watchstation functionality with high fidelity, on low-cost generic PC hardware. One of the existing PORTS watchstations was the TAO's. It also included

a simulation of the naval tactical environment. The system was already used at SWOS to train TAO students in console operation, though an instructor was needed for every two students to play the role of other CIC team members and provide tutoring. SWOS had already set up an electronic classroom that included 42 student PCs for viewing electronic materials networked to a single instructor console.

Figure 1 shows the PORTS simulated TAO console. It includes panels for Variable Action Buttons (VABs), display selection (map control keys), radio control, tactical situation map (a scaled version of the large screen display), and Automatic Status Boards that, among other things, display information on the hooked track. The mouse is used to push buttons and select tracks. These displays are driven by a tactical simulation that simulates ownship's sensors and weapons, external platforms, and the environment. PORTS simulations are initialized from a PORTS scenario file created using a graphical scenario editor.

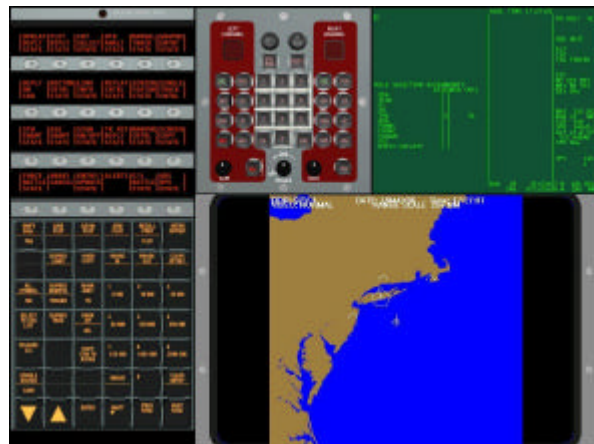


Figure 1. PORTS Simulated TAO Console

HIGH LEVEL INSTRUCTIONAL DECISIONS

The project began, as often the case with an ITS, without a clear understanding by the end-user (SWOS) of what ITS and AI technology was capable of and without a clear specification of functionality or requirements. The project kicked off with a

requirements workshop which included several SWOS instructors, several members of the ITS team, and several members of the acquisition organization, NAWC-TSD. It had already been decided to initially concentrate on the Air Warfare domain, specifically the Detect To Engage (DTE) sequence. One of the main watchstanders in the Air Warfare domain who reports directly to the TAO is the Air Warfare Coordinator (called simply "AIR")

We first presented a description of what ITSs could do. Initially it was not clear if an instructor would be partly involved or not and to what degree the instructional functions were to be fully automated. The SWOS instructors' initial concept for the ITS was to have it function very analogously to the way they were already using PORTS for scenario tactical training. One student played the role of TAO, one student played the role of AIR and the instructor played the roles of all other team members as well as evaluated the students' performance and provided coaching during the scenario and the debriefing at the end. Note that this setup required 1 instructor for every two students. Initially in the requirements meeting it was clear that the instructors wanted the ITS to fully eliminate the instructor from these scenario roles and assumed the ITS would just take over exactly the same functions currently being performed by the instructor.

The first major decision was to change the basic format of scenario training to eliminate the two student scheme in favor of individual TAO training. This decision was based primarily on efficiency. With the two student scheme, either a student wanting to practice on his own would need to find a partner or the ITS would have to support two different modes of operation, including tutoring two different roles - the AIR as well as the TAO. Second, while training two students with one human instructor is more efficient than training one student when the instructor is a limiting resource, this is not the case when the instructional functions are being performed by software which can be freely replicated.

In fact, greatly reducing the instructor workload during scenario training was the primary motivation for developing the PORTS TAO ITS. For training TAOs, it is most efficient for every student to assume the TAO role. Third, the better the AIR performed, the less the TAO had to do. So in the two student scheme, how much practice the TAO received in a scenario varied based on his partner's performance. It was unlikely that this performance was geared toward providing the TAO with the best training experience for him, based on his current state of development. With the change to a single student playing the role of the TAO, the ITS

had to provide agents, called Automated Role Players (ARPs), to play the role of all necessary CIC positions including AIR. Furthermore, the ITS had to assume all the other scenario-related instructor roles including automatically evaluating the student's performance, providing real-time coaching, and providing a debriefing.

It also became quickly clear that if the ARPs behave perfectly, the TAO would not be very challenged and it would be difficult for him to demonstrate that he had full mastery of the skills and knowledge required. As the TAO demonstrated, in scenarios, increased mastery, the ARPs would need to purposefully make more mistakes and mistakes which were more difficult to correct.

A second major decision involved the type of material that the ITS would teach. It was decided to concentrate on the material that was the most straight-forward and had the least ambiguity and controversy. These latter types of situations would be handled in the fully manned simulator where the ambiguous and controversial situations could be discussed with other students and instructors. This discussion was considered important for instruction. A second motivation for this decision was that it could be hard to get agreement by different instructors for the proper student choice in these situations which could stall the ITS development or cause the requirements to change after development had finished. The clearest, most agreed-upon material related the procedures and cadence of the detection, identification, warning, etc. process, when the results of each process step were clear.

It was also clear that the instructors wanted the ITS to take a spiral approach to instruction. When working with a new topic, the scenarios should be easy at first. As the student got more practice and demonstrated mastery, the principles and scenarios should increase in complexity and difficulty. When a topic was first introduced or when the difficulty first increased, it would be important to provide real-time coaching in a non-intrusive way. The coach should notice if the student was taking no action when required (or a wrong action) and provide hints, at first general but more specific if required, to try to get the student to perform the correct action. After exhausting the hinting process, the ITS should provide real-time feedback to student as to what he should have done. The ITS team originally assumed that the hints would take audio form such as tones or spoken words, but it became clear that for this domain, a simple line of text would be less intrusive and, since this was not a sensory

motor task, their eyes were free to quickly glance down at the few words required for a simple text hint.

The instructors also came to the meeting knowing that the system would have to support different modes of operations. Sometimes a single instructor would be present for an entire class of 30 to 40 students. In this mode all students would independently run the same scenario in parallel. This would be a scenario that practiced skills and knowledge (collectively called "principles") just covered by the instructor. It would be important for all students to be running the same scenario so if the instructor went over to help a specific student who was having problems that the ITS could not address, the instructor would already be familiar with what was going on in the scenario. But additionally, students needed to be able to come in and practice on their own, with no instructor present.

A third major set of decisions related to TAO verbal communications. It was decided early that the ITS would not tutor communications skills, including vocabulary and syntax. Communications would be tutored before the student encountered the ITS. It was also stated that SWOS was planning on teaching a more standard communication syntax than had been taught previously. It was later revealed that for Air Warfare, this communication syntax was rigid and all encompassing, that in a given situation what the TAO was supposed to say was rigorously defined and that every word could be exactly specified in advance, with the exception of track numbers and certain synonyms that would also be allowed. The instructors stated that the ITS could always assume the TAO students would use this correct syntax and that if an incorrect syntax was used, the system could interpret it any way that made sense from a software development perspective.

Many of the high level instructional decisions above became the basis for issues which needed to be addressed by the instructional design. To summarize, these issues were:

- More simulated scenario practice was needed for TAO trainees.
- More practice would require more instructors, while instructor billets were being reduced
- Instructional functions needed to be automated including:
 - Scenario coaching
 - Playing the role of other team members
- Automated role players need to make purposeful mistakes
- Spiral Instruction
- Instructor present and not present modes of operation

- TAOs primarily act by issuing verbal orders, but automatic processing of natural language is very difficult

INSTRUCTIONAL DESIGN

Levels

To address the spiral instruction issue, the PORTS TAO ITS is designed to present training content throughout a progression of difficulty, across a series of exercises. The basis of this progression is what TAOs and their trainers at SWOS call 'command by negation.' A TAO expects watchstanders to do their jobs correctly, concerning himself with maintaining awareness of those activities and of the tactical picture, stepping in for major decisions such as the engagement of an aircraft. When a watchstander takes an inappropriate action, either by making a mistake or for lack of information not normally available at that watchstation, the TAO negates it, and supplies a new order. This also applies when a watchstander fails to act, and the TAO 'negates' that inaction. Cases include a watchstander failing to make a routine report on the status of a friendly air asset, or incorrectly announcing intent to query an aircraft when it is within its territorial airspace.

From an instructional standpoint, the ITS is designed to require this 'negation' of the student more and more often, and in response to more nuanced mistakes, over a course of exercises. To this end, we partitioned instructional content into three levels, which were labeled simply by number (1-3) in order to avoid confusion with the SWOS terms 'basic', 'intermediate', and 'advanced,' which did not align with the usage here. In the DTE module, for instance, all three levels fall within what SWOS calls 'basic' DTE. More advanced DTE requires very nuanced, possibly ambiguous decision making, where the lessons lie less in right and wrong answers than they do in consideration of all the factors, though many of the factors do arise in the content of the ITS.

The three levels apply to principles, ARP behaviors, and exercises. A single principle may have three levels, in applying to three different kinds of situations. For instance, the principle 'query range' consists of knowing the correct range to issue a query to an aircraft. (Whether a query is appropriate for that aircraft in the first place is another principle.) To test this principle at level 1, an ARP announces intent to query at the correct range, and the TAO simply must acknowledge that announcement ("TAO, aye") in order to pass. No negation is necessary. At level 2, the ARP may remain silent when the aircraft reaches the correct

range, and the TAO must notice this omission, and request the query. At level 3, the ARP may intend to query at the wrong range, which the TAO must notice and correct. The ITS maintains an assessment of student performance at each level separately, so that the student must learn how to act correctly in each type of situation.

ARP Behaviors

The three levels of ARP behaviors exist in order to create situations where principles can be evaluated for all three levels. Level 1 is intended to help the TAO to learn to maintain awareness, and to habituate to the normal responsibilities of the watchstanders in a variety of situations. At this level, ARPs perform perfectly and the student witnesses no actions that need negating, so usually correct TAO response is simply to verbally acknowledge the action. Some level 1 principles don't involve ARP actions at all, but rather require the TAO maintain situational awareness in the tactical picture, performing console actions to select tracks and vary display ranges appropriately to the current tactical conversation being carried out between ARPs.

Level 2 requires the TAO to apply this awareness, by noticing when the usual procedures are not followed. At this level, ARPs tend to make mistakes of *omission*, usually by remaining silent when an utterance is appropriate. This requires the TAO to notice the omission and take action, most often by requesting the ARP to act.

Level 3 involves the ARPs tending to make mistakes of *commission*, where actions are taken incorrectly. Here, the TAO is expected to notice the mistake and request a correction of it. In the middle of an exercise, with many concurrent threads of activity with respect to a variety of simultaneous tracks, these ARP mistakes can be harder to notice and correct than when an ARP omitted an action altogether.

ARPs do not commit mistakes at every possible opportunity, but rather do their jobs correctly until a mistake is designated to happen. This is under control of the exercise author, who decides what mistakes will happen in an exercise, during what time range, and with what frequency. Thus the body of content on which the student is evaluated in a particular exercise is tightly controlled, but there is randomness to the mistakes from the student's perspective to prevent the student from 'learning' the exercise rather than the principles. Thus two exercises may have identical sets of events within the tactical picture, but depending on

the ARP behaviors, may evaluate entirely different skills.

Exercises are also divided into three levels, to suggest that they contain, for the most part, principle evaluations at that level and below. The course of exercises tends to progress from 1 to 2, although when new situations are introduced, such as the availability of friendly aircraft under ship control, they may start again at 1. Note that even for a level 3 exercise, there are plenty of level 1 (and possibly level 2) evaluations; in between mistakes, the ARPs go about their business correctly, and the TAO is still expected to acknowledge with "TAO aye."

The three levels are not always present for every principle, and do not always correspond to omission and commission mistakes by ARPs. When the TAO cannot verify the veracity of an ARP utterance—for instance, the ARP reports that a track has answered a query, on a communications net that only the ARP can hear—then only the omission is appropriate. Other principles, such as the TAO response to an equipment fault, involve no ARP mistake at all, but consist of more advanced knowledge that is not intended for instruction in early exercises. While the omission and commission mistakes provide a basic structure for progressive difficulty, difficulty may also be the result of inherent difficulty of a particular principle.

Hinting

Scenario coaching is handled primarily by hinting. Hinting is based on another progressive structure that mimics the training interventions of SWOS instructors. Hints begin to appear in the ITS when an action is expected of the TAO, for instance when the ARP has made an omission, or when an ARP has performed an action that the TAO is expected to acknowledge. (Hints pertaining to different principles that apply simultaneously appear together, and progress independently, possibly at different rates.)

There are four levels of hints: the red flag, the general hint, the specific hint, and the prompt. The *red flag*, the first to appear, is visually similar to other hints but has no content, consisting of a blue bar that appears on the screen. This is the instructional equivalent of a tap on the shoulder, indicating to the student that there is something that needs doing, without giving any further information.

The remaining hints contain text, gradually becoming more specific as they progress, each replacing the previous. The *general hint* tends to indicate some information about the situation that the TAO may not

have noticed, to which he is expected to react. The *specific hint* gives the gist of the type of action that is expected. Finally, the *prompt* provides the precise content of the action. The ITS assesses the student differently depending on how many hints had been displayed before the correct action is taken, with the lowest assessment resulting from no action at all. SWOS instructors, standing over the shoulder of a student performing in an exercise, often perform this sort of progressive intervention themselves.

As with principle levels, hinting levels are not required to follow this structure, and may consist of any sequence of texts that SWOS instructors desired the student to see.

Hints do not always appear, but are enabled by the ITS when the student's past performance (or lack of pertinent exercises) indicates a skill deficit for particular principles. This determination is made in real-time, so as the student demonstrates improvement by performing correct actions without an excess of hinting, the hints will disappear for just the principles for which the student has shown improvement. The ITS continues to assess performance on a principle by principle basis, however, and if the student subsequently fails a principle enough, the hints will again appear for that principle.

Communications

Most TAO actions consist of verbal communications with watchstanders, and most student actions in the ITS are student utterances into a microphone which creates the natural language processing (NLP) issue mentioned earlier. The microphone is activated with a physical push-to-talk button that also exists in the shipboard system. These utterances are expected to conform to a rigid syntax, which includes the variations and combinations that SWOS considers acceptable. This rigid syntax allows the system to use pre-defined grammars and standard commercial of the shelf software to achieve very high (over 97%) speech recognition accuracy and to have full understanding of the orders and queries issued by the student. Additionally, the ITS provides an intermediary utterance editor facility for the student to ensure the accuracy and syntactical correctness of utterances before they are passed on to the system and (and effectively 'said' to the ARPs).

Upon speaking, the ITS renders the text of the utterance in the editor. If the student is satisfied, it may be submitted with a mouse click, or the student may edit the text directly. This can be done with the type of typing operations available with a conventional word

processor, or through a hierarchical menu system that appears amid the text and allows expression of any allowable word. Red underlining indicates whether an utterance fails, so far, to be syntactic, and its positioning shows the location of the error. The student may also use the push-to-talk to try verbalizing again, rather than editing, which is the most common way students correct text. Since the ITS instructional content does not include communications syntax, the utterance is required to be correct when submitted.

Classroom and Homework Modes

The ITS has two primary modes: classroom and homework. The classroom mode is instructor-led with the entire group of 30 to 40 students running the same exercise as designated by the instructor. Here the instructor introduces the students to exercises that assess principles they have not seen before and, as the exercise goes on, may monitor the progress of the students (and their successes and failures) from a display on the instructor station.

Homework mode allows students to practice on their own, and does not allow the selection of every exercise. Rather, it follows on from classroom mode by retrieving an exercise containing principles that the student has already encountered in classroom mode, and on which the student needs to improve. This mode will not introduce new bodies of content, reserving that for the classroom environment, where the instructor is available to explain the essence of the principle and expected TAO behavior. Homework mode, in turn, gives the student a way to improve on those skills from the classroom, until they are shown to be proficient.

IMPLEMENTATION CHALLENGES/RESULTS

There were a number of challenges that needed to be overcome in the development of the PORTS TAO ITS which fell primarily into three categories - the government system acquisition process, instructor availability and turnover, and technical challenges relating to the development of the software. Each of these three categories is further broken down into specific challenges and described in the subsections below where the challenge and the means to overcome it are discussed.

Government Training System Acquisition Process

One of the main challenges related to the government training system acquisition process. On the one hand, it was geared toward primarily overseeing the development of training simulators or conventional Computer Based Training (CBT) systems. On the

other hand, the ITS development team was not particularly familiar with the formal training system acquisition process, including the involvement of government engineers, since their previous experience developing ITSSs generally involved working directly with the instructors in the target domain. These factors caused various forms of friction. The necessary acquisition requirements documents tended to reference items that were appropriate to CBT

development but not to ITS development. In particular, therefore, they were not conducive to communication of understood requirements to the instructors. This was solved by creating appendices to the requirements document, called the principle matrix and the ARP behaviors description that were more appropriate. Examples are shown below in Figure 2 and Figure 3.

Principle	Lev	Description	Behavior	Evaluation Cue	Exceptions	ARP Behavior	Learning Objective	Accepted Variations	Mistakes	Flag	Hint 1	Time from last hint	Hint 2	Time from last hint	Hint 3
Identification	1	Ensure proper identification of all tracks	Acknowledge AIR report of track ID	AIR reports track ID	None	Checkpoint result [any non-mistake]	"TAO aye."			3 sec	AIR has reported a track ID change.	5 sec	AIR is awaiting acknowledgment of a track ID change.	5 sec	Say "TAO aye."
	2		With AIR error in failing to report an evaluated track of interest, request such report	AIR fails to announce checkpoint result	None	Checkpoint result [any] omission	TAO says, "AIR TAO interrogative track ID 8\$track.TN\$?"; "AIR TAO interrogative"	"AIR TAO interrogative checkpoint track 8\$track.TN\$?"; "AIR TAO interrogative"		20 sec	A track ID has changed.	40 sec	All ID changes within CIEA and VA should be reported.	60 sec	Say "AIR TAO interrogative track ID 8\$track.TN\$?";
	3		With AIR error in providing incorrect identification of a track as a result of a CHECKPRINT procedure, request correct	AIR announces wrong track ID (symbology change always to correct ID, however)	None	Checkpoint result [any] commission	TAO says, "AIR TAO do not concur."	"AIR TAO interrogative track number?"; AIR TAO say again track number."; "AIR TAO negative ID, hold that"	TAO accepts report. OK for TAO to Say "TAO, aye" as long as he corrects himself within 5 secs	3 sec	Did you hook track last reported?	5 sec	Does track 8\$track.TN\$ correlate to ID given?	5 sec	"AIR TAO do not concur."
DCA VID	1		Acknowledge DCA VID order.	AIR orders DCA VID.		DCA VID					AIR has ordered DCA to VID.		AIR is awaiting acknowledgment of a DCA VID		Say "TAO aye."

Figure 2. Example Rows from the Principle Matrix

ARP Behavior	ARPs	Scenario Conditions	Actions
Checkpoint	AIR	New track of interest or TAO announces new track or requests checkpoint or requests correction of track number	"ALL STATIONS, AIR, CHECKPRINT track [num]"
Checkpoint omission	AIR	Same as above	none
Nonexistent track commission	AIR	Same as above	"ALL STATIONS, AIR, CHECKPRINT track [num, which does not exist]"
Checkpoint RSC	RSC	Checkpoint or TAO requests report of [ARP]	"RSC, aye, strength [strength]"

Figure 3. Example Behavior Descriptions

Test Procedures

A further issue related to Test Procedures mandated by the government. Because an ITS is based on an AI decision-making paradigm, and not based on pre-defined branching, as in most CBT systems, it was not possible to test every possible combination of student inputs. An initial attempt was made to develop procedures to test every meaningfully different student input on a per principle basis, but this created procedures that were overly redundant, especially

given the nature of the ITS software implementation. The ITS includes a single hinting behavior which looks up the hinting text and times from a file. The evaluation behaviors are essentially data that is processed by the behavior execution engine. Exhaustive test procedures effectively just tested the same code over and over.

Eventually a compromise was reached where enough testing was done to generate every possible hint and feedback message. Unfortunately, the time to create the test procedures had already been expended. One of the reasons it was so difficult is that test procedures are designed to be linear; the tester performs a series of actions in a certain order and expects the system to respond in a repeatable way. However, the scenario can play out slightly differently based on subtle differences in timing, so that even if the tester performs the same actions in the same order, it is practically impossible to perform them at exactly the same time in the scenario. This is especially true when multiple different things are occurring in the scenario at once, which sets up subtle race conditions.

There were many failed attempts to create test procedures based on realistic scenarios. Eventually, however, test-only scenarios had to be created, where, instead of many simultaneous tactical events occurring, the incoming enemy and commercial aircraft were spaced out enough so that, in spite of differences in the timing of tester actions (including significant delays in performing correct answers in order to generate all possible hints), the scenario would play out the same every time. This entire process took away from the more realistic testing that would normally be done with instructional scenarios. On the other hand, the development team learned valuable lessons on how to streamline the development of testing procedures and related test-only scenarios and to allocate enough time to it so as to provide more comprehensive, realistic testing with the actual instructional scenarios in the follow-on efforts.

Instructor Availability

One of the key elements in the development of a Navy training capability is having sufficient access to the subject matter experts – the schoolhouse Instructors who will have to incorporate the new training capability within their curriculum. In the current environment of the Iraq war, and resulting training budget cuts, there has been a decrease in the number of Instructors at SWOS. This results in a decrease in their availability to work with the contractor team in the definition of the requirements and in the design of what the Instructors need the ITS to be capable of. We have resolved this deficiency by using contractor subject matter experts (former Navy personnel) to initially define the detailed requirements, and then to have the Schoolhouse Instructors validate those requirements or modify as needed. This has worked fairly successfully, and has reduced the time spent waiting for Instructor availability.

SWOS Turnover

One other Implementation Challenge has been the turnover of Instructor Staff at SWOS. This creates issues with both the availability of Instructors and with the transfer of knowledge, history, and past decisions. Despite the best intentions, when an Instructor is called to other duty, the handoff to his replacement is never complete. Unfortunately, the Instructors who approved the requirements are not always the Instructors involved in the acceptance of the ITS software. There were several examples of this issue. For example, an initial group of instructors made the decision to require a rigid communication syntax and the specifics of what that syntax was. Later, during acceptance testing, an Instructor not privy to these decisions was brought in

for free-play testing, without a briefing on the syntax. Similarly this Instructor was not briefed on which principles were tutored. A related issue was foot pedal discipline. Almost without exception, students or testers playing the role of students tend to release the push-to-talk foot pedal slightly too soon, cutting off the last word or syllable, which, of course, causes difficulty with speech recognition. A certain amount of practice up front and vigilance throughout is required. Of course turnover magnifies this problem, because new groups of instructors have to learn and internalize this issue. The turnover challenge is mitigated by maintaining detailed notes of the decisions made in the requirements workshops and in reminding the government team of those decisions prior to the acceptance test process.

ITS/Simulation Integration

There is always the potential for complications when you integrate two existing technologies, especially if one or the other was not designed with this integration in mind. There have been several issues in the integration of the PORTS TAO Trainer with the TAO ITS which are worth mentioning as lessons learned.

The first issue was that although the PORTS TAO Trainer had been installed at SWOS for over a year, there were a very limited number of problem reports, most of which occurred within the first couple of months of delivery. We had expected this trainer to be used daily, and for any latent software issues to be uncovered, and brought to our attention. Once the problem reports stopped coming from the schoolhouse, we could only believe that the PORTS Trainer was fully meeting their requirements and was performing well. It turns out that after the first couple of months, the school had not incorporated the trainer into the day to day curriculum, and reserved its use for demonstration of the combat system capabilities to the students. This was an artifact of staff turnover. In the course of integrating the ITS with the PORTSPORTS TAO Trainer, a number of software issues in the PORTSPORTS software were uncovered, and required troubleshooting and fixing.

The second issue was in the interface defined between the PORTS TAO Trainer and the ITS. NAWC-TSD knew they wanted to add the ITS capability, and had provided funding to add an interface to an unnamed ITS system. The generic interface was built based on the CORBA technology with as much detail as was possible without knowing what it was going to be interfacing with. While most of this generic interface did work very well with the delivered ITS, there have been a number of issues arising from the way the

CORBA interface was designed, especially in periods where significant data is being passed across the interface. This has resulted in some redesign and redevelopment of this PORTS-ITS interface.

Classified Content

Some portions of instructional content are classified, such as the operating parameters of certain equipment, or the criteria for decisions such as the classification of an aircraft as hostile based on its behavior (PHID, or probable hostile identification). We wanted to allow for as much discussion, development, and testing as possible without restricting the work to a classified environment. The solution was to represent classified content as 'classified expressions,' which are all defined and stored in a single file and expressed in a simplified procedural computer language. This file is editable upon installation at SWOS, or by instructors any time thereafter, as part of the Exercise Authoring Tool which will propagate any changes to all other instructors' and students' computers. These expressions fulfill a black-box query role within the ITS. They are available both to the ARPs, when deciding when or how to perform an action, and to the evaluation machines, when assessing the correctness of a TAO action. An ARP or evaluation queries such an expression to determine, for example, whether a particular track should be considered hostile, and acts accordingly. The expressions used for the purposes of development and testing are unclassified, consisting of code that acts upon the same input as classified criteria, but making different decisions. This allowed SHAI to test the flow of data to the expression inputs, their causal linkage to ITS behavior in response to expression outputs, in a wide variety of situations, without requiring the use of classified code or documentation until the final stages of ITS deployment.

Incremental Development

One of the most important decisions made by the government acquisition team was to develop the complete PORTS TAO ITS in stages, one module at a time. For example, the first module was the Air Detect To Engage Sequence. It represented the first of 10 planned modules which eventually encompass all three warfare areas (Air, Surface, and Undersea). By phasing the development, all of the challenges associated with a complete development/acceptance cycle, most of which are described above, could be encountered early in the project and the resulting lessons learned applied to the later modules. Given that these challenges were encountered during the development of just the first 10% of the PORTS TAO

ITS, while 90% of the development was still to be accomplished, this provided a huge improvement the development efficiency and will result in a significantly enhanced final product.

FUTURE WORK

Applying the lessons learned as described above, we are currently implementing the next two modules of air warfare content, Inbound Anti-Ship Cruise Missile (ASCM) and Inbound Unknown Air Contact in the Vital Area (VA). These are elaborations on the basic DTE problem, introducing more novel and urgent circumstances than in the basic DTE.

In ASCM exercises, the student is briefed to expect imminent use of anti-ship cruise missiles. The DTE process is augmented with additional steps to heighten the ship's readiness, and new emergency procedures are introduced to defend the ship when a missile has been launched, either from an aircraft or from land. These tasks include additional DTE steps in prosecution of unknown tracks, maneuver of the ship to unmask both offensive (hardkill) and defensive (softkill) systems, and the deployment of these systems at the appropriate time. As opposed to the relatively routine and deliberate process of DTE, these steps must be performed at a rapid pace, because the potential threat is more imminent. Furthermore, actions such as softkill deployment are dependent on the type of missile, so the student must be able to deploy only the appropriate systems, and oversee the ARP's use of those systems, within this fast-paced environment. The applicability of each system to a particular missile type is classified, and is represented in a similar manner as are the PHID criteria in the DTE module.

ARP behaviors and ITS evaluation of the student's actions are more loosely coupled in this module, because the steps are less rigidly ordered. Unlike many of the steps of DTE, there is no single correct ordering of anti-ship missile defense procedures. A proactive TAO may choose to perform necessary actions before the time that the responsible ARP would choose to do so, but still be evaluated as correct. In this case, the ARP must avoid later duplication of that action. Also, because of the faster pace of this module, the ITS relies more heavily on the instructional technique of automatically pausing the scenario, rather than having ARPs compensate for TAO failure to keep the scenario moving forward.

The VA module consists of two general cases: Pops, where aircraft suddenly appear on RADAR at close range, and low-slow flyers (LSFs), which do not fit the ordinary profile of either an airliner or a tactical

aircraft. Popups are similar to missiles in that they precipitate a rapid series of events, but in this case those events conform more closely to basic DTE. ARPs and ITS evaluations incorporate modifications to the DTE module: some steps must be skipped in the interest of time, and other steps happen immediately in sequence, rather than awaiting the approach of the aircraft to a certain range. Beside those differences, of which the TAO must demonstrate knowledge, the Popup case is merely an accelerated DTE, requiring the TAO to demonstrate the same knowledge in a more rapid fashion. LSFs are a less urgent problem, but similarly require the omission of some steps from basic DTE, owing to their probable identity of a civilian aircraft without the usual communications ability, or of a threat that is dealt with by other crew, who achieve visual contact from atop the ship.

We are also beginning a set of modules that cover surface warfare (SUW), including SUW Recognized Maritime Picture Construction, and Short Notice Over-The-Horizon Targeting with requirements gathering and the initial workshop with instructors. As with the subsequent set of modules covering subsurface warfare, these modules cover instructional content that will provide focused exercises, but will support the creation of exercises that cover all warfare areas at once. The surface modules are likely to bear many similarities to the air modules, such as expectation of the TAO to maintain situational awareness by interacting with the tracks on the tactical display, and also pose new challenges, such as the fact that SUW is more concerned with the efficient use of scarce ID assets and that the communications protocol, while still being capable of being rigorously structured, has a greater number of variations.

REFERENCES

See <http://www.colorado.edu/NROTC/career/career-SWO.html> or <http://depts.washington.edu/uwnrotc/swo.htm> for a description of the role of a Surface Warfare Officer.