

User Acceptance of an Intelligent User Interface: A Rotorcraft Pilot's Associate Example

Christopher A. Miller
Honeywell Technology Center
3660 Technology Dr.
Minneapolis, MN 55418
cmiller@htc.honeywell.com

Matthew D. Hannen
The Boeing Company
5000 E. McDowell Rd.
Mesa, AZ 85215
matt.hannen@boeing.com

ABSTRACT

The U.S. Army's Rotorcraft Pilot's Associate (RPA) program is developing an advanced, intelligent "associate" system for flight demonstration in a future attack/scout helicopter. A significant RPA component is the intelligent user interface known as the Cockpit Information Manager (CIM). This paper describes the high level architecture of the CIM, with emphasis on its pilot-perceptible behaviors: Crew Intent Estimation, Page Selection, Symbol Selection/Declutter, Intelligent Window Location, Automated Pan and Zoom, and Task Allocation. We then present the subjective results of recent full mission simulation studies using the CIM to illustrate pilots' attitudes toward these behaviors and their perceived effectiveness.

Keywords

Cockpit Information Management, Rotorcraft Pilot's Associate, Associate Systems, Page Selection, Symbol Selection/Declutter, Automated Task Allocation, Pan & Zoom, Window Location, Intent Estimation.

THE ROTORCRAFT PILOT'S ASSOCIATE PROGRAM

The US Air Force's Pilot's Associate programs were among the first efforts to implement large, adaptive, intelligent user interfaces (IUIs) [2]. The US Army's Rotorcraft Pilot's Associate (RPA) program is extending and implementing this work [4]. In this section, we will briefly describe the approach to IUI we have developed for the RPA with emphasis on the intelligent information management cockpit behaviors which result in task-sensitive, dynamically generated cockpit configurations.

The RPA program is a five year, \$80 million research contract managed by the U.S. Army's Aviation Applied Technology Directorate at Ft. Eustis. It currently represents the U.S. Army's largest research and development commitment and is one of the largest ongoing IUI application efforts in the world. The goal of RPA is to develop and

demonstrate in flight an "associate" system in a next-generation attack/scout helicopter. Associate systems are collections of intelligent aiding systems that, collectively, exhibit the behavior of a capable human [13,14]. They can (a) perform roughly the same breadth of activities as a human expert in the domain, (b) take initiative when necessary, but generally follow a human colleague's lead, and (c) integrate over ongoing activities to exhibit robust, coordinated, appropriate behavior.

A critical goal of RPA is to manage the information available in future helicopter operations so that human crews can attend to all and only relevant portions at any given time. Further, RPA must accomplish this without increasing pilot workload or decreasing situation awareness.

Nature of the RPA Task Domain

While other IUI applications [e.g., 5,8,11,12] are characterized by truly vast quantities of information, highly configurable interface and automation technologies, highly variable and unstructured task needs, and comparatively mild time constraints, RPA differs on all these fronts. The information available to a military helicopter pilot is extensive and growing, but it is constrained by available sensors and pre-formatted datalink communications. Thus it is less broad and variable than that available to e.g., a military operations planner—much less a university student with a Web query. While military glass cockpits are highly flexible, capable of presenting information in multiple modalities and formats, they maintain strict constraints and conventions on information formatting to facilitate transfer of training and ease of uptake. Thus, a truly novel interface, generated dynamically for each situation, is not acceptable. While the Army attack/scout helicopter mission is one of the most flexible in the military, and future battle scenarios call for even more dynamic mission planning and re-planning, procedures are exhaustively thought out at all levels from broad force strategies to specific mission plans and pilots are extensively trained in these procedures. Thus, user tasks are reasonably well scripted and shared, and most situations will have appropriate doctrine created to enable users to 'fall back on their training'. Finally and importantly, response times are critical. The huge number of potential threats (passive and active) in the operational

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IUI 99 Redondo Beach CA USA

Copyright ACM 1999 1-58113-098-8/99/01...\$5.00

environment demands constant and complete situation awareness. Flexibility of use, and even operator autonomy, may sometimes be sacrificed if the payoff is faster reaction time.

These aspects of the RPA domain have colored our development of an IUI for it. As is described below, we have emphasized decreasing pilot workload to access and comprehend available information and adhering to the standards and expectations for interface construction, over dynamically accessing novel information or generating interfaces. Above all, we have emphasized gaining pilot acceptance. This paper will report both the IUI design approach we took, and the insights gained from recent full mission simulation evaluations as to those aspects of the RPA IUI which pilots are willing and unwilling to accept.

Overall RPA Architecture

Figure 1 illustrates the whole RPA architecture. There are two major parts. First, an Advanced Mission Equipment Package (AMEP) provides a suite of sophisticated automation including advanced sensors, communications and targeting systems. While highly capable, these are all "traditional" automation systems in that they serve a single function without explicit reference to operator goals and have little autonomous capability. While relatively intelligent, they are what Riley [13] refers to as assistant or slave systems rather than true "associates". The second major component of the RPA system is the Cognitive Decision Aiding System (CDAS). CDAS integrates the functionality of the Mission Equipment Package with explicit models of crew tasks to sew the traditional automation systems into a semi-autonomous "associate."

CDAS itself consists of six modules. Data Fusion is responsible for correlating the stream of incoming data from multiple external sources into a single, accurate and unified view. External Situation Assessment consists of battlefield and target assessment subsystems that reason about the significance of external conditions to known mission goals. Internal Situation Assessment performs similar functions on internal aircraft health and status monitoring equipment. A series of six real-time Planners use this assessment data to offer suggestions for maximizing successful accomplishment of known goals. Each Planner is responsible for a single functional area: route planning, survivability, communications, sensor management, attack and reconnaissance. The Cockpit Information Manager (CIM) is the IUI component for the RPA CDAS as a whole.

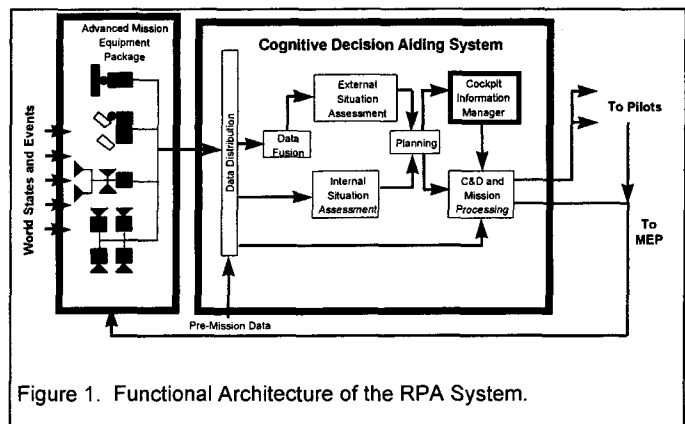


Figure 1. Functional Architecture of the RPA System.

THE RPA COCKPIT INFORMATION MANAGER

Cockpit Information Manager Architecture

The architecture and knowledge representation of the Cockpit Information Manager (CIM) is described briefly in this section. For more detailed information, see [9,10].

CIM is primarily a task-based, rather than state-based, artifact-based, or user-based IUI (cf. [12] and Jones' discussion of 'boundary objects' as organizational structures for coordinating interactions in [6]). This task-centered focus works well in a highly proceduralized and trained domain such as helicopter operations. CIM is responsible for determining the current and near-future tasks of the crew, and then adjusting the cockpit configuration to meet task needs. A cockpit configuration consists of an allocation of all active tasks to a mixture of cockpit actors (two pilots and automation), and an allocation of interface functions required for those tasks to some mix of cockpit display and control devices. The ultimate goal of interface management in CIM is enhanced task performance, thus CIM prioritizes and filters presentations so that the most important tasks have their information needs met first, and crew workload and display device capacity thresholds are not exceeded. The Controls and Displays (C&D) and Mission Processing logic is then responsible for issuing the low-

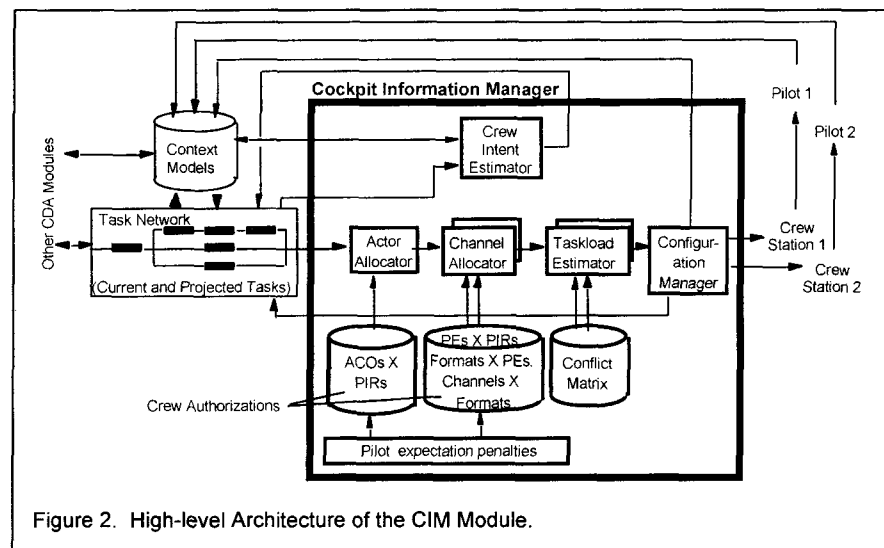


Figure 2. High-level Architecture of the CIM Module.

level avionics commands to achieve automation tasks and cockpit configurations commanded by CIM.

Figure 2 illustrates the CIM we have developed. Task and context information are provided to the CIM by two shared-memory resources of the CDAS as a whole: the Task Network and the Context Model. The Context Model represents the CDAS' current beliefs about the state of the aircraft and the world. The Task Network represents the CDAS' current beliefs about the tasks that are now being performed and upcoming. In essence, all tasks that can be performed in the RPA aircraft are modeled in the Task Network (cf. [7]), along with alternate methods for performing them. Before takeoff, a mission-specific task network is crafted to represent the mission plan. During a mission, activation and completion conditions on tasks are triggered by pilot actions or world states (recorded in the Context Model), thus enabling the Task Network to maintain, in real time, a model of the active, expected, and completed tasks at any point.

CIM is not just a consumer of task information. The Crew Intent Estimator component (see Figure 2) interprets pilot actions and world events against mission plans in the Task Network and, using knowledge of goals and side effects, ascertains whether the pilots are following the mission model or are attempting alternate plans or goals. The Crew Intent Estimator can revise the Task Network model to reflect new crew intent to perform a different set of tasks. More details on the Intent Estimator's function and mechanism can be found in [1].

The set of active tasks, constantly updated and prioritized, is CIM's basis for interface management. CIM's reasoning is conducted in four stages as illustrated in Figure 2 and, in more detail, 3. Prior to mission execution, a designer will use a data capture tool to record constraints and pilot preferences for task allocation, symbol use, channel use, pop-up window position, etc. S/He also captures the information needs for each task, and the display alternatives for satisfying those needs. This information feeds a flight da-

tabase that is accessed by the CIM algorithms.

Whenever the set of active tasks changes, CIM computes a new cockpit configuration (though this only occasionally produces a change in what the crew sees). For each task, CIM decides how the task should be allocated across the available, legal combinations of actors (human and automation). Each legal task allocation is called an Actor Configuration Option (ACO). ACO selection is performed by CIM's Actor Allocator component.

For each ACO, CIM's knowledge base contains a list of prioritized pilot information requirements. CIM's Channel Allocator component attempts to satisfy 'PIRs' by associating presentation elements (PEs--concrete means of conveying information: graphical symbols, acoustic tones, etc.) with them and allocating cockpit input and output channels to those PEs.

Throughout the selection of actors and presentation methods, CIM maintains a record of its violations of pilot expectations, information needs, display stability, or any of a number of other desirable (but often conflicting) interface characteristics. Once an actor's tasks cockpit configuration is determined, CIM's Taskload Estimator component computes a human workload estimate for doing those tasks in that way. This estimate, along with the violation scores described above, is input to a scoring function resident in CIM's Configuration Manager component. This evaluation metric is described in more detail in [9].

Finally, if time permits, CIM will iterate through its design decisions to develop cockpit configurations with better scores on its evaluation metric. When time is up, CIM reports its best configuration to the cockpit C&D and Mission Processing logic for generation in the crew stations.

CIM Interface Management Behaviors

For communicating CIM's capabilities to pilots, we found it helpful to describe CIM behaviors from an operator's viewpoint instead of, or before, presenting architectural descriptions like those above. CIM performs six primary activities observable by the pilot. These are:

1. Intent Estimation
2. Task Allocation
3. Page (or Format) Selection
4. Symbol Selection/Declutter
5. Window Placement
6. Pan and Zoom

We will describe each of these behaviors below in the context of a representative mission scenario (depicted in Figure 4). This scenario covers roughly two minutes. It begins with an "Ingress" task as a part of a larger reconnaissance plan. During this task, the crew encounters an unexpected threat which triggers a previously unplanned (and inactive) "Perform Actions on Contact" (AOC) task. Subtasks under AOC in this scenario (others are possible) include performing a "Low Probability of Detection Ma-

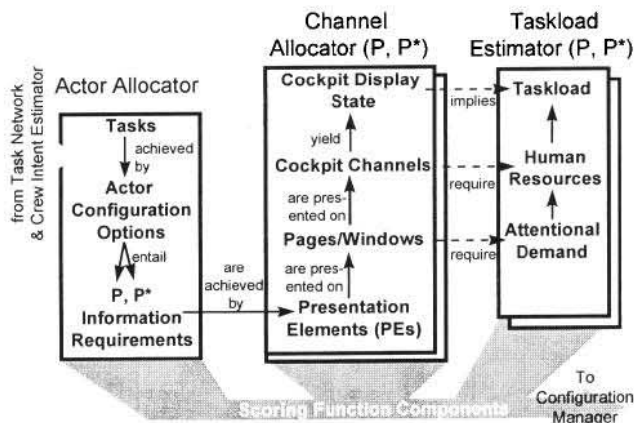


Figure 3. Knowledge Representation and Reasoning in CIM.

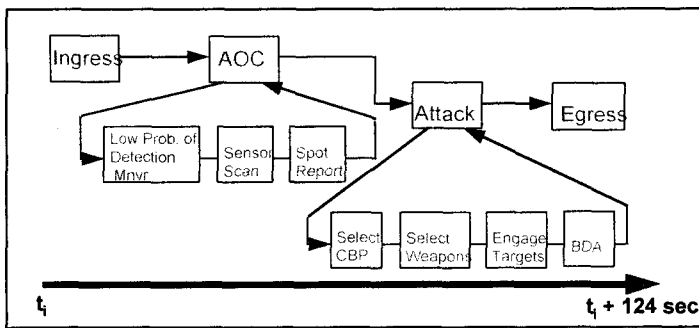


Figure 4. Representative mission scenario illustrating CIM behaviors. neuver,” a “Sensor Scan” to determine the nature and extent of the threat, and then submitting a “Spot Report”. The Crew Intent Estimator then determines that the crew intends to engage the threats and therefore asserts a newly active “Attack” task, which in turn contains the subtasks “Select Combat Position,” “Select Weapons,” “Engage Targets” and “Battle Damage Assessment”. These tasks are followed, in this scenario, by an “Egress” task.

Intent Estimation

The intent estimation behavior of the CIM has little visible output to the RPA crew, but it can dramatically affect the other CIM behaviors through the set of active tasks it reports. At the beginning of our scenario, the AOC task is neither active nor a scheduled part of the mission. Detection of an unexpected enemy threat (e.g., through Data Fusion’s interpretation of sensor data) is an automatic Task Network trigger for the AOC task, but sensors can be wrong and there are battlefield threats that are not detectable with current technology (e.g., small arms fire). Since it is important for CDAS to remain ‘in the loop’ with the pilot even in these instances, the Crew Intent Estimator constantly tracks crew behavior to infer their intent. Thus, while sensors might be unable to detect a small arms threat, the pilot’s pattern of evasive maneuvering, weapon slewing, communications, etc. and will be evaluated using template-based plan recognition techniques [1] to infer the presence of an otherwise undetectable instance of AOC.

While intent estimation is mostly invisible, it is not entirely so. The RPA cockpit includes an ‘intent interface’ to provide the crew with both insight into, and some control over, CIM’s understanding of their intent. This “Crew Coordination and Task Awareness” display consists of four small LED buttons located in the upper right portion of each pilot’s main instrument panel. It reports, in text, the current inferred (1) high-level mission context, (2) highest priority pilot task, (3) highest priority copilot task, and (4) highest priority CDAS task. Pressing these buttons permits the pilot to override CIM’s current inferred tasks and assert new ones via push button input. Since the inclusion of a direct method for viewing and interacting with intent estimation was a new development in the RPA cockpit (over prior as-

sociate system work), we were especially interested in how pilots would regard it.

Task Allocation

Under this behavior, for not-yet-active tasks, CIM determines how best to allocate the task among a ‘legal’ (i.e., pre-approved) mixture of human and automation actors—an ACO. For example, the pilots may have authorized CIM to consider allocating the “spot report” task either (a) to the copilot alone, (b) to automation with the final approval by the copilot, or (c) to automation alone with no copilot approval. Each of these would then be a ‘legal’ ACO which CIM will consider, though the pilots may also indicate that one method is heavily preferred (a factor included in CIM’s evaluation metric and traded off there against potential benefits such as workload reduction.

Page Selection

Given each actors’ tasks, CIM determines the best set of pages (i.e., formats) to present on the three available multi-function displays (MFDs) for each crew member. In our

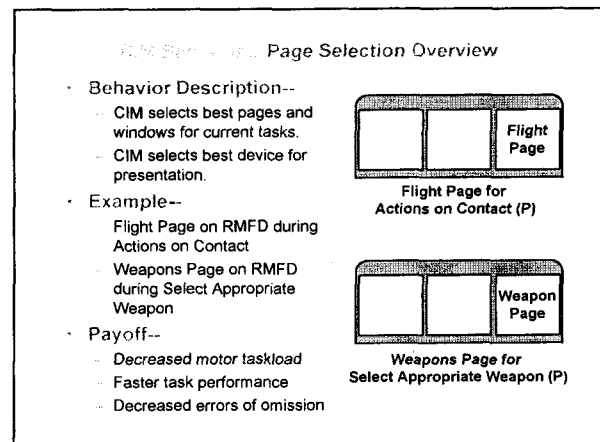


Figure 5. Page Selection Example.

scenario, when AOC becomes active, the crew will have high need for information about the threat and about the maneuvering capabilities of the aircraft. Thus, in this context, the three MFDs will generally be configured to show a sensor page, a tactical situation map display and a flight page with primary flight symbology—though the concurrent presence of other, high priority tasks might result in a different configuration. When the decision is made to engage the threat (when AOC ends and the “Attack” task begins), weapons configuration and control information are temporarily more important for the copilot, and thus his cockpit may be re-configured to present a weapons page instead of the flight page (see Figure 5).

Symbol Selection/Declutter

For each of the selected pages, CIM determines the best set of symbols for meeting the current information needs, removing unnecessary symbology when it taxes either the MFD’s or the human’s processing capacity. For example, during the “Ingress” and “Egress” tasks, the crew needs

navigational information and their map displays will generally show routes, phase lines, passage points, etc. When an unexpected threat appears, however, this navigational information is less critical than information about the threat and its relative position. Thus, CIM will generally suppress navigation symbology and replace it with threat symbology such as threat icons, lethality and intervisibility envelopes, etc. (see Figure 6).

Window Placement

The RPA crew station can present many types of information in pop-up windows overlaid on a portion of a larger page. When pop-up windows are used, CIM determines their best placement to minimize obscuring other needed information and symbology, yet adhering as closely as possible to expected locations for each window type. For example, when the pilots decide begin the "Attack" task, they need to select a combat position. In support of this "Select CBP" task, the AMEP provides them with a recommended position and an explanation of the recommendation in terms of the Army's standard position evaluation criteria. This explanation is presented in a pop up window as 9 criteria scores. While useful, placing this large window in any default position on the map risks obscuring threat symbology critical to the ongoing attack task. CIM dynamically selects a position for the window that minimizes obscuration and, when impossible, ensures that only lower importance symbology is obscured (see Figure 7).

Pan and Zoom

CIM also controls the pan and zoom settings of the tactical situation (i.e., map) display to ensure presentation of im-

CIM Behaviors: Window Placement Overview

- Behavior Description--


- CIM positions windows to minimize obscuration of other needed information.

- Example--


- Default position for accept route window is lower right on Tactical Situation Display.
- Important engagement tracks are located there, so window is relocated to alternate position.

- Payoff--

- Decreased Motor taskload
- Improved Situation Awareness
- Faster task performance



Without CIM: Window obscures important symbols



With CIM: Intelligent selection of secondary locations

Figure 7. Window Placement example.

portant symbols, yet avoid clutter. For example, during Ingress and Egress, pilots need a large area presented on the map, though high resolution for terrain is less important. When a pop up threat appears, needs shift: high resolution for the area around the threat and possible maneuver paths is important, though the total area shown may be reduced. CIM reasons through these various needs and asserts pan and zoom commands that adjust the map to appropriate pan and zoom settings (see Figure 8).

SIMULATION TEST RESULTS

While a primary focus of the RPA program is to carry the associate architecture into a flight demonstration phase be-

ginning in July, 1998, extensive full mission simulations were carried out to help evaluate CDAS behaviors and implementation, and to assist in prioritizing issues for flight demonstration. These full-mission simulations were conducted at Boeing Mesa in February and March, 1998. Below, we report the experimental design and primary findings with regards to pilot acceptance of the CIM IUI design. Quantitative performance data are still being analyzed and will be the topic of a future paper.

The RPA Simulator Evaluation emphasized military mission realism. The 4 Crews (8 pilots) trained and flew together, as they do in field operations. Crews were given realistic objectives and permitted to make their own tactical decisions about how to achieve them. Tests were flown in Boeing's full mission simulator and included full fidelity RPA cockpits,

Symbol Level Declutter Overview

- Behavior Description--


- CIM selects best symbols for current tasks
- Removes unnecessary or lower priority symbols

- Example--

- Navigation symbols during Perform Ingress and Egress,
- Threat symbology during Actions on Contact and Attack.

- Payoff

- Reduced visual taskload
- Improved situation awareness
- Improved tactical decision making under stress



vs.




Figure 6. Symbol Selection example.

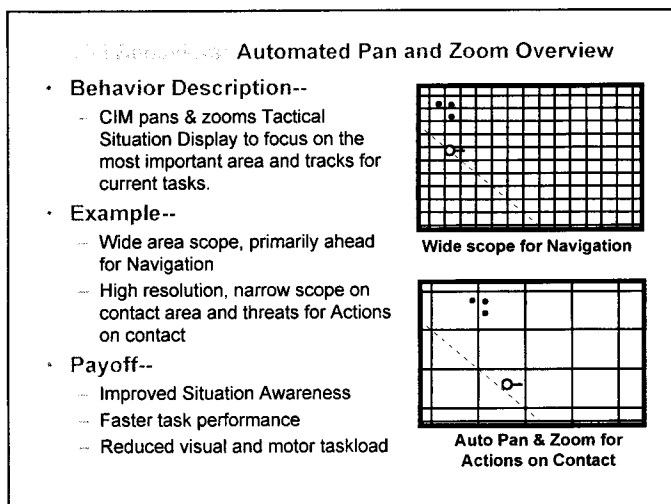


Figure 8. Pan and Zoom example.

dome visuals, extensive passive and active threats, and human control of the Tactical Operations Center, friendly artillery, and 1 to 3 wingmen. Realistic communications, including change of mission Fragmentary Operations Orders, were maintained between these players. Each pilot received an average of 10.8 hours of training in the simulator and 13.9 hours of classroom training over a two-week period.

Each crew flew 14 part-mission test scenarios of 20-50 minutes duration—7 with the RPA CDAS, and 7 with the AMEP alone. The focus of these part-mission scenarios was on exercising a particular CDAS or AMEP behavior and thus context, objectives and task flow were permitted to be somewhat fragmentary or unrealistic. Each crew also flew four 1-1.5 hour full-mission scenarios—two with the AMEP alone and two with CDAS. Full-mission scenarios were designed to be highly realistic and crews were given free reign to pursue their commander's objectives via whatever methods they thought appropriate.

The AMEP vs. CDAS conditions were chosen to balance the evaluation over a common baseline of advanced automation technologies. The only difference between the two conditions was the addition of the integrative, associate and IUI technologies of the CDAS (cf. Figure 1). All missions were balanced for complexity. Crews flew the two AMEP or CDAS full missions in sequence and then switched technology conditions and flew the remaining two missions with the other set of technologies. The sequence was counterbalanced to minimize training effects.

The simulation tests were constructed to include numerous examples of the CIM page selection, window location, pan & zoom, and symbol selection behaviors in a variety of tactical mission contexts. Crew Intent Estimation was implemented for the Actions on Contact task alone, and CIM task allocation behaviors were not implemented in the simulation due to time and budget constraints.

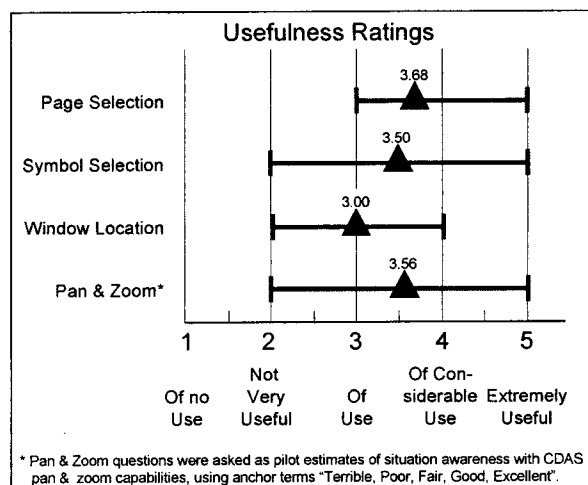


Figure 9. Pilots' average ratings and ranges for the usefulness of the four CIM behaviors.

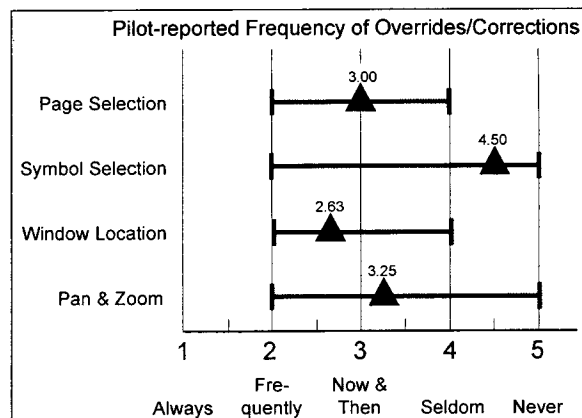


Figure 10. Pilots' average reported perceived frequency of overrides and corrections to the four CIM behaviors.

To obtain crew acceptance data, a questionnaire was administered after each of the final AMEP and CDAS full-mission test trials. All of the questionnaire responses utilized complete verbal anchoring and a linear response scale with five equal intervals, in accordance with [3]. The pilots were instructed to mark one point anywhere along the linear scale from one to five. The criteria value (for satisfactory CIM behavior) established before the simulation test was an average score of 3.5.

This criterion was met for three of the four CIM behaviors tested directly. The average and range of pilots' ratings of the behaviors, is presented in Figures 9. In general, pilots found the CIM behaviors to be 'Of Use' or 'Of Considerable Use.' Figure 10 presents pilots' perceptions of the frequency with which they overrode or corrected CIM's actions. The average over the CIM behaviors fell between 'Seldom' and 'Now and Then' with symbol selection capabilities performing notably better.

With regards to the less observable Crew Intent Estimation behavior, pilots believed it was fairly accurate at recognizing crew intent to perform AOC (average rating 4.15 \approx 'Frequently' triggered when crew intent or mission context made it appropriate). But this came at the cost of false positives (average rating 2.40 \approx CDAS 'Seldom' or 'Now and Then' triggered AOC when pilot intended to continue past threats). In spite of these perceived occasional inaccuracies, and in spite of some pilot complaints about inadequate training in their use, most pilots found the inclusion of the LED Task Awareness display 'Of Considerable Use,' as shown in Table 1.

Table 1. Perceived usefulness of the LED Task Awareness Display (where 4.0='Of Considerable Use' and 5.0='Extremely Useful'.

LED Button for:	Score
Mission Task	4.4
Pilot Task	4.3
Copilot Task	4.3
Associate Task	4.0

benefits provided by CDAS, it is difficult for both pilots and experimenters to parcel out some of these benefits. Table 2 presents a comparison of pilot ratings of their effectiveness over four high-level mission tasks with CDAS versus with the AMEP alone. On the average, pilots found themselves to be more than half a point more effective (12.5% of the scale length) with CDAS than without.

CDAS also produced overall benefits relative to AMEP in

Table 2. Perceived effectiveness in different mission tasks with CDAS and AMEP alone (where 3.0= 'Fair', 4.0='Good' and 5.0='Excellent'.

Average Rating	AMEP	CDAS
Zone Reconnaissance	3.75	3.88
Area Reconnaissance	3.75	4.25
Deliberate Attack	4.13	4.75
Change to Attack	3.63	4.63

one other critical area. Using TLX measures of subjective workload collected at the end of each part- and full-mission trial, perceived workload levels were consistently higher for AMEP

conditions than for CDAS conditions (57 points versus 46 points). This difference was significant in an Analysis of Variance [$F(1,6)=11.524, p<.05$]. There were no significant differences between pilots and copilots and the interaction effects were not significant.

Furthermore, separate ANOVAs were conducted for each of the six TLX subscale ratings to determine CDAS' contributions to overall workload reduction. As can be seen in Table 3, the reduced workload for the CDAS configuration is apparent in the mental, physical, and temporal demand and effort subscales. There is also a marginal finding for the frustration subscale ($p>.07$). Means in all cases indicate

Figure 11 shows pilot ratings of CIM as a whole. CIM was seen as 'Frequently' providing the right information at the right time and, of critical importance to our subjects, was seen as almost always predictable in its behaviors.

Finally, while CIM cannot claim credit for all of the




Figure 11. Pilots' average overall ratings of the CIM.

that CDAS provides a workload benefit to the pilot. Examination of the perceived performance ratings, however, shows no effect of technology. This may indicate that pilots use a different subjective criteria in rating their own performance, possibly judging it based on how well they felt they should have done in the given context, which would include the cockpit configuration.

CONCLUSIONS

Although we will have a more complete picture when the objective performance data have been evaluated, the subjective pilot responses described above suggest that the CIM behaviors we identified and implemented are generally meeting mission expectations, contributing to perceived pilot effectiveness, reducing workload and are gaining pilot acceptance. It is worth noting that perfection in aiding and tracking pilot intent is not a prerequisite to the levels of acceptance we have gained. Pilots rated the CIM behaviors 'of considerable use' and said that CIM 'frequently' provided the right information at the right time in spite of perceived false positives and 'now and then' having to override or correct CIM's behaviors. High degrees of predictability and the addition of a highly regarded (if simple) Task Awareness and Crew Coordination Display may have contributed to pilots' willingness to tolerate these occasional 'mistakes' on CIM's part.

Table 3. Analysis of the TLX subjective workload subscale ratings.

TLX subscale	AMEP mean	CDAS mean	F-Value (df: 1,6)
Mental Demand	61.77	46.25	10.487*
Physical Demand	54.48	40.31	12.042*
Temporal Demand	62.08	45.73	14.061*
Perceived Performance	35.00	42.08	2.429
Effort	62.60	48.54	20.470*
Frustration	52.81	45.63	4.961

* $p<.05$

** $p<.01$

Areas for CIM improvement identified during this evaluation, include: (1) more predictable and accurate CIM window locations, (2) finer CIM control of the digital map pan and zoom selections, and (3) improved inhibits and/or improved intent discriminability on the AOC CIM behaviors when performing a deliberate attack. These modifications are being made for the RPA Flight Demonstration.

As a final indication of pilot acceptance of the CIM behaviors, when conducting the full-mission simulations using CDAS, pilots were given the option of turning off any or all of the CIM behaviors via a TAILOR page available both before and throughout the trials. Nevertheless, all eight pilots chose to leave all CIM behaviors turned on throughout their full mission trials—a sign of trust in, and perceived benefit from, CIM's management of the displays in response to the changing mission context.

The RPA CIM is producing a range of reliable, predictable, useful cockpit interface management behaviors. The fact that pilots rank CIM behaviors highly and choose to use them is an indicator, albeit preliminary, that CIM's IUI functions will be useful and will contribute to mission performance. As the RPA program moves through its flight demonstration in 1998, CIM behaviors will continue to be refined and evaluated, but these results give us reason to believe that they will be one of the core benefits provided by the RPA Cognitive Decision Aiding System as a whole.

ACKNOWLEDGEMENTS

This work was supported by the US Army AATD, contract number DAAJ02-93-C-0008. We would like to thank Harry Funk, Mark Ringer, Denny Rother, Erik Steinmetz, Ruben Yanas, Bernie McBryan and Ray Higgins for their help in performing the work reported here.

REFERENCES

- Andes, R. Assuring Human-Centeredness in Intelligent Rotorcraft Cockpits: Using Crew Intent Estimation to Coordinate RPA Functions. Available at www.asinc.com/papers/ahs97fnl.html. 1997.
- Banks, S. and Lizza, C. Pilot's associate; A cooperative knowledge-based system application. *IEEE Expert*, (June, 1991), 18-29.
- Charlton, S. G. (Ed.). *Operational Effectiveness Questionnaire Handbook*, Subjective Assessment Working Group, US Air Force Operational Test and Evaluation Center, Technical Paper 7.1, July, 1989.
- Colucci, F. Rotorcraft pilot's associate update: The army's largest science and technology program. *Vertiflite*, (March/April 1995), 16-20.
- Frank, M. and Szekely, P. Adaptive Forms: An interaction paradigm for entering structured data. In *Proceedings of the ACM International Conference on Intelligent User Interfaces* (San Francisco, California, January 6-9, 1998) 153-160.
- Jones, P. and Chin, G. Varieties of representational forms: Towards designing an information infrastructure for an organization. To appear in *Proceedings of the 42nd Annual Meeting of the Human Factors and Ergonomics Society*, (October 5-8: Chicago, Ill), 1998.
- Keller, K. and Stanley, K. Domain specific software design for decision aiding. *Proceedings of the AAAI Workshop on Automating Software Design*. (NASA Ames Research Center, July, 1992), 86-92.
- Maybury, M. *Intelligent Multimedia Interfaces*. AAAI Press/MIT Press, Menlo Park, CA, 1993.
- Miller, C. Bridging the information transfer gap: Measuring goodness of information "fit". Submitted for inclusion in a special issue of *Journal of Visual Languages and Computing*, 10(3).
- Miller C. & Funk, H. Knowledge requirements for information management; A Rotorcraft Pilot's Associate example. In Mouloua, M. and Koonce, J. (Eds.), *Human-Automation Interaction: Research and Practice*. Lawrence Erlbaum: Hillsdale, NJ, 1997.
- Mitchell, C.M. and Saisi, D.S. Use of model-based qualitative icons and adaptive windows in workstations for supervisory control systems. *IEEE Transactions on Systems, Man and Cybernetics*, 17(4), 594-607.
- Oppermann, R. *Adaptive User Support*. Lawrence Erlbaum: Hillsdale, NJ, 1994.
- Riley, V. A general model of mixed-initiative human-machine systems. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Denver CO, October 1989) 124-128.
- Sheridan, T. Supervisory Control. In G. Salvendy (Ed.), *Handbook of Human Factors*. John Wiley & Sons, New York, 1987. 1244-1268.