Operational Design Review of an Envisioned World: Empirical Lessons Learned on reviewing a future CSE-based design for a large scale DSS Environment

Szymczak, Samantha S., McKenna, Brian P., Potter, Scott S., & Elm, William C.
ManTech – Cognitive Systems Engineering Center

Introduction

Accidents frequently occur in systems because the evaluation plan for the design of that system was not well thought out (Chiles, 2002). In many cases, the evaluation is a one-time event at the end of the design process. The subject matter experts (SMEs) are essentially asked to validate the first version prototype of the display. This, unfortunately, leaves very little time and even less money to fix the problems that are discovered during the evaluation cycle. Due to this, one common solution is to ensure the known major problems are fixed and assume that those gaps that are identified later can either be solved on the fly through patches (e.g., Microsoft routinely releases patches to shore up security flaws that have been found in their code) or that the worker will develop stop-gap measures to work around the problem (e.g., an operator may use a notepad to write down important information which is not readily available in the interface). The latter is often the preferred method, as people will generally make any system workable, given that they have spent enough time with it (Woods and Hollnagel, 2006). These solutions, at best, delay eventual system failure.

When evaluations are performed early in the design process, they often focus on superficial aspects of the display. As there is no working system with which a user can interact, it can be difficult to truly gain an understanding for what the system is attempting to accomplish, something that is easy to see when observing an operational system during use. A static picture makes it difficult to convey the pace and flow in which a system normally operates. Therefore, the early evaluation is used to ensure that each item is in its desired place and that the controls will be obvious to the user; at best, SMEs might be used to provide subjective feedback about what they think about design. What these early evaluations lack is the ability to see work in context – it is during these times where much insight can be gained about the effectiveness of a system. Seeing an operator use a system provides more useful information than can be gleaned from a talking about a static picture. However, with no functioning system in place, it is a challenge to get these data.

Evaluations of new systems may be complicated by other factors as well. Often times these new systems are coupled with a new operational concept. The evaluation becomes not only a review of the system, but also a review of the operational concept behind the system. The challenge then comes in identifying differences in an SME's review of the concept versus their review or understanding of the design, and to identify when this distinction needs to be made. Observational data from SMEs will vary from items related to the fundamental concept of the system to those related to the instantiated design choice, where the concept may be valid, but the system design may need to be improved.

The scale of the new design scope can cause problems as well. It can be difficult to get a comprehensive evaluation of the interactions of one person with one display. New systems are likely to have multiple displays being operated on by multiple people. These multi-user, multi-display systems make up a Decision Support Environment (DSE), which also must be validated within the evaluation process. Because of this, it becomes difficult to narrow the focus of an evaluation well enough to get feedback on each item while still evaluating the DSE as a whole.

Evaluations need to occur frequently throughout the design process, but getting valuable feedback on the depth of decision support provided by the system early in the development process requires an evaluation style that covers more than the superficial aspects of the display. The key to running a successful evaluation at this critical juncture lays in portraying a situation

similar to what would be encountered in the real world, even when using a lo-fidelity prototype; to do so, this situation must be mocked up just well enough to allow the SMEs to get into their role

By crafting a storyboard, the design team is able to determine exactly what the SMEs see. In this way, they can ensure that events that stress the ability of the system occur. This is often lost when observers watch events in the real world – where what happens is detected by external sources. This style of evaluation, therefore, allows for Staged World observations (Woods & Hollnagel, 2006), creating a balance between control of the situation and realistic operations. The use of scripted scenarios also creates the effect of dynamism despite the static prototypes that are being used. This enables operators to shift their focus from the physical components of the display to the functional merits of the work the operator is required to accomplish. ManTech's Cognitive Systems Engineering Center (CSEC) explored this concept for use with a lo-fidelity prototype during the review of a command center for a naval frigate (N.B., for current purposes, the system is considered lo-fidelity since it had not advanced beyond the design phase).

Levels of Evaluation

For a true indication of decision support, an evaluation needs to go further than examining the usability of the system. The focus on usability is too narrow and does not allow for adequate insights into the effectiveness of the system. This is generally done during an evaluation early in the design process, because this is the easiest place to focus. As there is no physical system, the Usability of a system to complete a task cannot be tested. It is possible, however, to determine if a button looks like it can be pushed, or if a control looks like it can be used. These issues are important, but generally do not require an SME to be there. Even during evaluations at the end of the design process, with fully functional systems, the evaluation is generally expanded only to see the operator accomplish some routine tasks or confront a problem that has been anticipated by the designer.

Using a well-developed storyboard, however, allows other insights to be gleaned from the evaluation. It is important to see how well the system supports the user at work. A system that does not support the user in accomplishing his goals will ultimately result in a system failure, regardless of the usability of that system. This was the motivation for using the U^{3}SE framework, which espouses three levels of evaluation: Usability, Usefulness, and Understanding. The second and third levels take evaluation beyond Usability and are concerned more with in depth issues related to how well the system supports the user in accomplishing his goals.

Usefulness is concerned with the decision-making effectiveness of an operator during standard operational conditions. A Useful system would support the higher-level mission objectives, rather than just the specific mission tasks. Successful system operation is more than random tasks that need to be accomplished. Systems are designed to assist the user in accomplishing goals. In order to accomplish goals, the user must be able to identify what tasks are important for that goal and the best way to go about it. Asking a user to navigate from point A to point B in a system or to push the correct button does not show the system to be Useful.

The highest level at which a system should be evaluated is Understanding. At the Understanding level, the decision-making effectiveness of an operator during abnormal conditions is tested. It is possible to see how the operator would react to unexpected events in the environment. The better insights the operator has into the work domain, the better these levels will be supported by the display. Success for the operators will be determined by how well they are able to act effectively in the world (Woods & Sarter, 1993).

A comprehensive human-system interface evaluation program must address issues at all three of these levels. The Usability/Usefulness/Understanding System Evaluation (U^{3}SE) approach – consists of a complementary set of evaluations that address three levels that, taken together, address the net decision-making effectiveness of the joint cognitive system. With U^{3}SE,

not only are human-system interface problems identified (based on the specific JCS under evaluation), but also decision support deficiencies are uncovered (based on the support required vs. specific JCS under evaluation).

The CSEC's focus during an evaluation is aimed towards the higher two levels: Usefulness and Understanding. By placing the evaluator into an envisioned world, it is possible to move beyond Usability towards the Usefulness and Understanding levels. Even without a physical system to work with, it can be possible to evaluate how well a system will do under operational pressures.

Lessons Learned

It is possible, and effective, to simulate a command-room environment using every-day resources

The CSEC team had to determine how to present the multi-sub-system envisioned world operational concept to the reviewers. Traditional design reviews might sit a reviewer in front of a single computer screen, and walk them through the system one display at a time. Each screen would be critiqued in isolation on its own merits. The storyboard-based review, on the other hand, seeks the exact opposite; the reviewers are surrounded by the information as it is designed, and immersed in a scenario that they are faced with daily, so that they can see the consequences of the new system design in their every day decision-making.

In order to complete this effect during their evaluation, it was important that CSEC's design review room simulate the atmosphere and layout of the designed Envisioned World Command Center. There was a wide gap between the atmosphere of the room where the review was going to take place and the atmosphere of the Envisioned World. The room needed to facilitate the suspension of disbelief for the SMEs. Ideally, the evaluation would have been conducted in a space that was similar to the proposed layout of the Command Center. However, in this case, the room shape and size could not be recreated (the actual review took place in a conference room at the customer facility, which is significantly smaller than the space of the actual Command Center).

Creating the illusion of being in the actual Command Center was vital for a successful evaluation. It was critical, therefore, that the room contained the features of the proposed system, namely the four shared displays designed for the room as well as individual workstations, which the evaluators would sit. In addition to the all of the system features being simulated for the review, the atmosphere of the room was also important. The lighting of the conference room was set at levels that would imitate a Command Center. Only in this manner could the evaluators truly get a feel for what a decision maker within the envisioned Command Center would go through.

The room size dictated how much space was available to place all the components as well as where the evaluators could be situated. Space considerations and the hardware available placed limitations on the manner in which CSEC was able to set up for the design evaluation. From the room specifications of the conference room where the actual design review would take place, however, it was possible to determine the best arrangement of people and hardware within the room. In order to do so, CSEC created an in house replica of the customer's conference room prior to the scheduled evaluation. Not only did this allow for continual adjustments to the layout, wherein the optimal layout could be configured, but it also provided a suitable arena for practicing the design evaluation. The replica room would ensure that all components necessary for the evaluation were in proper working order and that the evaluation could commence as anticipated. Using a set of four projectors, each connected to a laptop computer, the CSEC was able to simulate the designed command wall for the four group displays. The four independent

projectors/laptop systems allowed each display to show the information in each group display that corresponded to the desired point in the storyboard, just as it was designed in the Envisioned World operational concept. Flat-panel computer screens were placed in the room as the individual displays (in a place where neither the monitor nor the person sitting at that location cast a shadow onto the screens) serving as a workstations for the reviewers to sit at. Every key aspect of the Envisioned World was considered for the review and practiced in-house before the actual design review dates. The result was a "full-functioning" (lo-fidelity) Envisioned World Command Center prototype.

It is Critical for Reviewers to Understand the Envisioned World Operational Concept Prior to Beginning the Design Review

One of the major challenges in conducting the design review was bringing the members of the reviewing team "up-to-speed" on the Envisioned World. Each subject matter expert attending the design review session was given a booklet of supporting material weeks before the design review was scheduled. This booklet contained:

- Objectives and approach of the project
- *Summary of the concept of operations*
- Overview of the design review
- *Concept-of-operations detailed report*
- Sample screen shots of each interface including
  - Operational description of interface (what it means to the practitioner)
  - Explanation of interface (how to read the interface) for the operational case shown
- *A CD containing an animation of the envisioned world*

The italicized items proved to be key artifacts to help train the reviewers on the new operational concept, which was much different than the Navy's current practice. The key differences between the two concepts were summarized at the beginning of the report to help familiarize the practitioners with the new envisioned world. For example, the crew structure and room layout were completely re-designed based on a functional modeling of the work domain (using the Applied Cognitive Work Analysis methodology of Elm et al. 2003). The concept called for operators to use new types of interface displays, namely the group displays visible by all crew members at once. The layout itself was non-traditional as the seats were arranged for optimal viewing of the group displays and communication between crew members, rather than in traditional rows. These kinds of major differences in the new design, which would otherwise have stayed the same or similar, were outlined at the beginning of the booklet, both for emphasis and to ensure that any practitioner that may have paged through the booklet rather than read it cover-to-cover would have seen them. Details of these features were contained in the Concept of Operations report, also included in the supporting material booklet.

The most beneficial artifact given to the reviewers prior to the review dates was an animated model of the control room, provided on a CD, to illustrate the dynamism of the newly proposed Envisioned World. The CD contained two animated tracks:

1) The Scenario – actors for each of the newly-proposed crew roles played out the scenario while a 3rd person narrator explained the events.
2) Concept of Operations – a narrator describes the functional layout of the room, operator role-allocation, and purpose of the information displayed at strategic places throughout the room.

The animated scenario provided an operational view point for the reviewers by illustrating the work flow and decision making flow from beginning to end of a realistic situation. Operator roles were made apparent as the animation would focus on one character or another for varying decision making tasks. The animation illustrated the system functions by featuring display components and particular displays as the operators would use them for decision making in the scenario context. One of the largest differences between current state of practice and the envisioned world was that an operator's work space consisted of his individual work station in coordination with multiple types of group displays. The scenario track helped bring together the operational concept of the coordination of the different display devices as one would use all of them in a normal state of practice. The concept of operations track explained the system in more detail. It complimented the scenario's operational view point by providing information about the systems functions, including the roles of the operators and purpose of the information space displayed by the computer system. Reviewers commented that the video helped them envision how the new operational concept would work, proving to be critical in suspending any dis-belief of the envisioned world that the operators may have had. One reviewer commented that his understanding of the operational concept from the CD was greatly different than how he envisioned the operations after reading the provided literature. Seeing an animation of the control room helped them visualize the concept the CSEC was proposing.

Once an understanding of the operational concept is achieved, it is necessary to train the reviewers on the use of the new system prior to requesting feedback about the system.

One of the most useful lessons learned for an Envisioned World design review was the necessity of training the reviewers on the new system prior to presenting the scenario. The review technique focuses on an immersive scenario-driven presentation. This is effective for illustrating an operational concept, but not sufficient for receiving feedback on the new Envisioned World system. The following technique evolved as a lesson learned during the CSEC's presentation:

- When showing a new interface, begin with a simple and nominal use-case
- Explain the interface in an operational way – describe what the interface is communicating to an operator and how the interface is related to interfaces already discussed
- After the review team is comfortable with the meaning of the interface, then begin to explain the specific coding mechanisms – how to read the screen
- Explain each presentation concept on the interface one by one (i.e. each graph, information display, etc.)
- Begin with the frame of reference of the presentation concept (explain what the axes mean, how information is laid out the screen)
- Then explain in detail each marking within the frame of reference. A separate presentation was sometimes needed to explain complex markings.
- Once the reviewers understood how to read the interface, then the more complex use case can be shown.

It was critical for the CSEC to first present a simple use case so that the reviewers had a context for where the interface and information it presented fit within the new operational concept. Once this was understood the details of the interface could be discussed. The training for each interface took, on average, only thirty minutes before the reviewers were comfortable with reading the

information presented and understanding what it meant about the world. This may be attributed to the expertise of the reviewers chosen to be on the review committee. The best in the field were chosen to review the new innovative concept. One would naturally expect them to understand things quickly. It may also be a property of the representational design technique of Cognitive Systems Engineering and that the designs themselves may have afforded for quick understanding by the experts of the domain. Regardless, training on the system design was critical in order to receive feedback about the usefulness and understanding levels of the interface for supporting domain specific decision making. The CSEC found the training to be minimally time consuming and non-burdensome to running the design review they had planned for.

The review proved to be effective for receiving operational feedback on the usefulness and understanding levels of an envisioned-world system design

One of the most critical steps in the Envisioned World design review was the interpretation of the reviewers' feedback. The Usability, Usefulness, and Understanding framework was used as a reference for placing the feedback into one of the three categories:

Usability

Some of the questions from the reviewers were of the style "What would happen if I did…" This type of statement or question was considered a Usability issue and most often referred to the dynamics within a display, or the possible user interactions with a display. Because the focus of the review was at a conceptual level, most of the questions were about details that had been speculated by the design team, but not necessarily defined at the time of the design review. These issues are important to consider when the system implementation specifications are being written. They do not, however, change the direction of the project or seed any re-design efforts. The technique did inspire some usability questions, as people are naturally thinking about how to interact with the screen. It would not be considered, in this experience, to be sufficient for a complete usability review, as it was not intended to be.

Usefulness

A useful display provides the necessary information to support a practitioner in nominal or non-stressed situations. The system interface must provide enough information to be useful to a practitioner in normal routine. If a user cannot carry out his or her normal state of business, the interface is considered non-useful, and must be re-designed. For example, the CSEC received much feedback about the interfaces designed to support ship navigation. The representation in the design presented to the reviewers was contained in a frame of reference of absolute coordinates, where the ship was shown to move and propagate on the screen through the world. This information was technically "correct," but as it turns out, not the frame of reference used conceptually by the practitioners. A more useful frame of reference for this information was a relative representation, where the ship was the center and maintained a "forward" orientation, while the world moved and rotated around the ship. In this case, the interface shown to the reviewers did not support them in nominal situations.

Another example that illustrates the reviewers' operational statements related to usefulness was when they would ask clarifying questions of how and when to take action in the world. For example, a reviewer clarified that when the display looked a certain way it meant that the threat was within weapon range and clarified how he would interact with the system to engage the target. The details about the controls relate to the usability of the system. The fundamental premise behind the interface, however, was to support decision making about target engagement. When the reviewer displayed an understanding of how the interface related to his world, and expressed his thought process of when to engage a target, in the nominal case presented, he was

in essence validating the Usefulness of that particular visualization. The information was complete and sufficient for the operator to deal with situations within nominal range.

Understanding

The extent to which a display affords true Understanding of a domain will be apparent when the system is stretched to its limits. A system that provides an Understanding level of decision making support will help an operator in extreme conditions. For the naval domain, examples of extreme conditions would be when there are an extensive amount of objects in the ship's environment, if the ship's resources are low, if a weapon system or sensor system is down, etc.

The most prominent example from the design review of the reviewers providing feedback which validated that the Understanding of the domain was portrayed through the interface representation was an intense discussion between the practitioners about a display meant to support decision-making related to scheduling engagements. The practitioners were debating over whether a particular combat power would be effective against a particular target. The discussion was intense, and at the surface level the reviewers seemed to be confused by the system interface. But, when listening to the discussion between the reviewers, one realized that they were not arguing about the meaning of the display at all. They were arguing about the capabilities of their weapon system in a non-nominal situation. They were discussing the tactics of how they would deal with the non-nominal case that the display was portraying. What seemed to be a mis-understanding of the display, turned out to be exactly the opposite: a very in-depth understanding through the interface into the limits and constraints of their weapon systems, and the fact that a target was pushing those limits to a point where the practitioners disagreed on the best practice for handling the edge which they were approaching.

Another Understanding-level statement the reviewers had was about displaying the history of some values over time. During the review, the practitioners felt that they did not need this history. They did not see the value of having that information as dedicated real estate on the screen. The CSEC team felt that this history was necessary to see the trend of the particular piece of information so that decisions could be made. The CSEC team believed that the practitioners did in fact use historical information as context for decision making within their heads, but that the practitioners were simply not used to seeing the history explicitly indicated on an interface. In hindsight, this emphasized the need for an understanding level of evaluation for this part of the screen. In the nominal case, the operators may not need this information, if in fact they did retain this information in their head, it may be trivial in nominal circumstances. It would be beneficial to stress decision making tasks of the operators by providing this screen interface with and without the history in a non-nominal situation, to test of this information in indeed part of the Understanding of the domain. As part of this lesson learned, it would have been beneficial in general to have more explicit "system limit" points within the scenario to stress the understanding level of the system. The reviewers' statements did not cause the history to be removed from the design, as might be the result of standard reviews. Instead, it was identified as a point for further review as it had been identified as a necessary context for information in previous knowledge elicitation opportunities. Users are often unaware of all the actions that they are performing to complete a goal (Woods & Hollnagel, 2006). If the history information was presented in a way which stretched the operators to the Understanding of the domain, it might be found that this information was critical for decision making in a non-nominal situation. The important aspect to note, however, is that the reviewers challenged the need for this information, and because of the Usability, Usefulness, and Understanding framework, the designers are able to identify this as a place for further testing on the Understanding level of the interface, to ensure the proper design is created before system implementation occurs.

An in-house "full dress rehearsal" prior to the actual presentation was key to the success of the design review

Once the scenario production and the plan for the design review room was complete, CSEC conducted a practice of the presentation in-house. The following proved to be critical to the success of the actual review:

- *Ensuring that the proposed set-up was conducive to a design review*, that the equipment and participants would "fit" in the room in such a way as to create a simulated immersive environment, without disruption due to physical constraints of the space.
- *Ensuring the accuracy of the created storyboard* - improper information on the screen can disrupt the mindset of the reviewer – losing the effectiveness and credibility of the simulated world. The focus would shift from the concepts the displays were presenting to the feasibility of the data. This "check" step was important because of the lo-fi prototype being used. All screens shown to the SME were actually manually populated screen shots of the proposed system. Therefore, there was room for population errors, which, as explained, could have been detrimental to communicating the intent of the displays.
- *Designating hardware (laptops and projectors) to specific displays* – With the set-up details already figured out the day of set-up was streamlined and routinized , to ensure there were no unnecessary obstacles. Tradeoffs between available projectors and required resolution for each of the four group displays were also worked out prior to the evaluation dates.
- *Practicing the control of the screen coordination* – The system designed for the naval control room consisted of 4 large screen displays, in addition to individual workstations. Each reviewer was shown the same workstation screen, so this resulted in 5 displays to be coordinated and in sync at all times. It was extremely critical for the presenter in charge of the screen coordination to practice and develop a rhythm to the screen synchronization.

Concentrate production effort on only those screens that are exercised at that point in the storyboard

The scenario used was a distillation of an actual scenario used by the Navy. Unlike in an operational command center, where interesting events are likely to occur at any point in the work domain, the design team was able to choose when and where events were to occur. While this allowed for flexibility in what could be shown, the team also needed to focus the SMEs to particular system areas for the evaluation. The parts of the scenario that the CSEC team used were the operationally interesting parts, where the operators' cognitive workload was thought to be high, and quick and informed decision-making was critical.

The most challenging part of developing the scenario was choosing the snapshots of time that would be shown to the reviewers. The designed system contained a display architecture allowing the users to navigate between interfaces. At any one time, the operator might be looking at one of 13 displays. For the sake of the scenario, the design team chose one interface for each snapshot of time to display to the design team during the review. The display chosen held the crux of action for that state of the world. The design team attempted to pinpoint where decision that was important at that point in time was being made, so that the usefulness of the display could be made evident. At some states of the world, the operational concept was shown best if time was frozen, and the navigation through the system was shown, as an operator would be expected to navigate through the system in order to solve the given issue.

Because there was no working prototype, the CSEC team had to manually populate each display. In order for the scenario to be believable, there had to be consistency both within the same display across time and across multiple displays. However, each new display state was also required to show change over time, since the state of the world is constantly changing during operations. This change must be greater than the one or two items of interest within the display. Each item on the screen must act as though it was in an operational system, meaning every change in the scenario had to be meticulously detailed in the screens in order to convey a sense of realism to the reviewers. The evaluation also required that each screen interface be equally represented throughout the scenario so that each screen received the same amount of operational feedback.

These limitations, however, guided the design of the scenario. The storyboard focused only on a selected portion of the system. As the decisions were focused on certain displays, the decision-making focus would be on those displays. Therefore, the effort spent on production of the scenario was dedicated on these displays only; there was no need to populate displays that were uninteresting at each point in time. This approach saved time during production because of the tediousness of populating the displays given the complexity of the scenario. The SMEs were still able to evaluate each display, but would do so only at the appropriate times.

One possible downfall of this approach is that, during the review, the review team is limited to only the information on the screen chosen to be shown by the design team (During this review, the SMEs rarely asked navigate to a different screen). However, by crafting a scenario with a narrow enough focus, the design team can ensure that these requests for more information are unlikely.

Conclusion

By immersing evaluators into a lo-fidelity prototype of an envisioned world and conducting a review at an operational level, the CSEC overcame a major challenge of early evaluations - evaluating a system that has yet to be built. They were able to create an immersive Envisioned World for a room sized system at a relatively low cost during the design stage of the project, prior to any system implementation. With the hardware configured to simulate the room, the seating arranged in an immersive manner, and the ambiance of the room adjusted to be authentic, the environment can be realistic enough to suspend the disbelief of the reviewers. An in-situ observation using a series of static-screen images provided enough dynamism and context to result in operational feedback about the Usefulness and Understanding levels of the system interface.

Prior to beginning the evaluation, the reviewers must be caught-up with the Envisioned World operational concept. This can be done extremely effectively through an animated storyboard, so that the reviewers can see the dynamism of the concept prior to diving into system specifics. Once the operational concept is understood the reviewers should be trained on the new system so that they are able to suspend disbelief and immerse themselves as operators, in-role, into the lo-fidelity prototyped system being used for the evaluation.

The production of the prototype can be tedious; however, time can be saved without detriment to the evaluation by only producing those displays which are operationally exercised at given points in the scenario. It is extremely beneficial for the presenting team to perform a full-dress rehearsal prior to the evaluation dates. Most importantly this rehearsal ensures that the presentation runs smoothly so that the reviewers remain in-role and perceive the presented lo-fidelity prototype realistically enough to give operational feedback about the system.

The review resulted in operational feedback at a conceptual level of the system design, namely the Usefulness and Understanding levels. The feedback received occurred during the

design stage of system development, leaving time and money for radical changes in the system design.

CSEC believes the lessons learned from this design evaluation can be generally applied to evaluations across a broad range of systems. The challenges encountered form a pattern which designers normally face during early evaluations (Woods & Hollnagel, 2006). Using these lessons learned to conduct an early evaluation will result in a more tightly coupled Joint Cognitive System (Hollnagel & Woods, 2005) once implementation occurs.

References

Chiles, J. (2002). *Inviting Disaster: Lessons From The Edge Of Technology*. New York: HarperCollins Publishing, Inc.

Elm, W.C., Potter, S.S., Gualtieri, J.W., Roth, E. M., & Easter, J.R. (2003). Applied cognitive work analysis: A pragmatic methodology for designing revolutionary cognitive affordances. In E. Hollnagel (Ed) Handbook for Cognitive Task Design. London: Lawrence Erlbaum Associates, Inc.

Hollnagel, E., & Woods, D D. (2005). *Joint Cognitive Systems: Foundations of Cognitive Systems Engineering.* Boca Raton, FL. CRC Press / Taylor & Francis.

Woods, D.D. & Hollnagel, E. (2006). *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering.* Boca Raton, FL: CRC Press / Taylor & Francis.

Woods, D.D. & Sarter, N. (1993). Evaluating the Impact of New Technology on Human-Machine Cooperation. In J. Wise, V. D. Hopkin, and P. Stager, editors, *Verification and Validation of Complex Systems: Human Factors Issues*, Berlin: Springer-Verlag.