Elliott, G., Crawford, J., Watson, M., Sanderson, P., & Naikar, N. (2000). Knowledge elicitation techniques for modelling intentional systems with Cognitive Work Analysis. *Proceedings of the Fifth Australian Aviation Psychology Symposium.* Manly, November 20-24. [CD-ROM]

Knowledge elicitation techniques for modelling intentional systems with Cognitive Work Analysis

GLENN ELLIOTT, JENNIFER CRAWFORD, MARCUS WATSON, PENELOPE SANDERSON
Swinburne Computer Human Interaction Laboratory
Swinburne University of Technology

NEELAM NAIKAR
Defence Science Technology Organisation

Introduction

The design of new human machine systems, and the analysis of current human machine systems pose interesting dilemmas for the human factors specialist. Cognitive Work Analysis (CWA), (Vicente, 1999; Rasmussen, Pejtersen, & Goodstein, 1994) is an analytical framework for modeling both work environments and safe, flexible strategies workers can use to act upon this environment. CWA has proven to be a particularly effective modeling tool not only for design, but also throughout the entire system design life cycle (Sanderson, Naikar, Lintern & Goss, 2000). An important consideration for CWA is the human requirements of a new system, or the evaluation of how humans interact with existing systems.

In this paper we describe four knowledge elicitation (KE) techniques developed over the last year to assist in the rapid development of human-centered CWA based analyses. These KE techniques were developed for work in two different aviation domains, an existing civilian Search and Rescue (SAR) service, and a new airborne defence platform called Airborne Early Warning and Control (AEW&C) which is yet to exist physically.

A common link between the SAR domain and the AEW&C domain is that they are both systems run largely through rules of practice that contain a ‘domain of potential risk’ for the workers involved. In the case of SAR the domain of risk can be thought of as ‘passive’, in that there are no malicious intentions (i.e. simply a patient lying in the water waiting to be rescued). However in the case of AEW&C, the domain of risk can be thought of as “active” in that there is a intelligent adversary with their own intentions and beliefs about the state of the environment. The appropriateness of the KE techniques will be discussed; a) with respect to the two aviation domains and b) with respect to the use of CWA products across different stages of the system design life cycle.

Use of Cognitive Work Analysis to answer human-system integration questions

Cognitive Work Analysis is an interdisciplinary, systems-oriented approach to the analysis, modelling, design, and evaluation of human-machine systems that provides a framework for modeling factors that shape human activity. It contains five phases of analysis beginning with Work Domain Analysis (WDA), which focuses on the environmental constraints and affordances presented to workers. For example, an environmental constraint for an airline pilot is to fly the plane within its operational limitations, whereas an example of an affordance would be the ability to use a winch to haul a patient onto a helicopter. The rationale behind this sequence of modelling is that only once after we have highlighted constraints placed on workers by their environment that we can begin to explain such things as the activities workers need to engage in to gain control over the environment, the possible safe strategies for accomplishing these activities, and the organisational structures and the worker competencies required to engage in these activities. In this sense, CWA begins from an ecological perspective and slowly moves towards more cognitively based analyses. CWA is considered a “formative” modelling technique, as opposed to a normative or descriptive technique (c.f. Sanderson, 2000), because it supports worker flexibility and safety, especially in previously unexperienced situations (Vicente, 1999; Rasmussen et al, 1994).

1. **Work Domain Analysis (WDA)**
2. **Control Task Analysis (CTA)**
3. **Strategies Analysis (SA)**
4. **Social Organisational Analysis (SOA)**
5. **Worker Competencies Analysis (WCA)**

Ecological

Cognitive

Figure 1: The five phases of Cognitive Work Analysis, with abbreviations.

Although a Control Task Analysis was also conducted for the two aviation domains in this paper, for the sake of brevity we will be focusing only on the first phase of CWA, Work Domain Analysis, with respect to the products developed from the knowledge elicitation techniques. WDA employs a structural means-ends abstraction hierarchy (AH), first used by Jens Rasmussen (1979, 1985) to describe a work domain from the user’s point of view.

WDA	Label	What it means
	Functional purpose	Why work domain exists
Purposive	**Priorities and Values**	Things to balance, conserve, etc
	Purpose-related functions	Work domain properties and functions
Physical	**Object-related functions**	Object properties and functions
	Physical objects	Object location and appearance

Figure 2: The five levels of abstraction in the Work Domain Analysis Abstraction Hierarchy.

As Figure 2 shows, the AH typically employs five levels of abstraction with the objective being to link the ‘purposive’ aspects of the work environment (i.e. why the system exists, elements of the environment that need to be preserved, and the general functions which can be acted upon in order achieve the domains purpose) with the physical elements of the world and the functionality that these elements afford. One major advantage of using WDA as a modeling framework is that unlike existing task analysis, the products of WDA are event independent. They can therefore help to design interfaces and communications structures that help workers behave flexibly during unanticipated events.

Intentional systems

Until recently, CWA has typically been applied to physically coupled process control systems such as feed-water process control microworlds (Vicente, 1999), pasteurising plants (Riesing, 1999), and nuclear power plants (Gualtieri, Roth & Eggleston, 2000). Unlike these systems, “intentional” systems are primarily driven by the rules of practice that guide the considerations, intentions and decisions of human operators rather than by physical constraints. Intentional systems can be highly dynamic and uncertain, since they are not engineered, raising the issue of whether CWA can model such systems. To date, no study has systematically investigated how well CWA can model intentional systems. An illustrative paper by Wong, Sallis, & O’Hare (1998) attempted to use WDA to model an emergency dispatch system. However, the authors argued that CWA was not an adequate framework for modeling intentional systems. In response, Hajdukiewicz, Burns, Vicente

& Eggleston (1999) presented a WDA of emergency dispatch and military command and control, and concluded that WDA can be performed for intentional systems. It is clear that further research is needed.

Domain of Risk

A further consideration for the work domains we are investigating is that they both contain an element of risk for the workers involved. The idea of juxtaposing a domain of risk and a domain of response to, or "mitigation' of, that risk, in a CWA analysis is not new (Moray, Sanderson, & Vicente, 1992; Hajdukiewicz, et al 1999). It is possible to divide a WDA into the 'domain of mitigation resources', or the resources available to workers to moderate or negate the element of risk, and the 'domain of risk', which is the source and nature of risk in the work environment that requires a mitigating response. Furthermore, we wish to make the distinction between domains of risk that are 'passive' in nature as opposed to those which are 'active'. As already noted, passive domains of risk could be thought of as a set of priorities and functions that pose a risk to an "opposing" set of interests, but without having malicious intent to do so. Active domains of risk could be thought of as an a set of priorities and functions in the work domain, that has a degree of intelligence, and that is actively working to maximise or somehow alter the risk presented to workers (e.g. a wartime enemy).

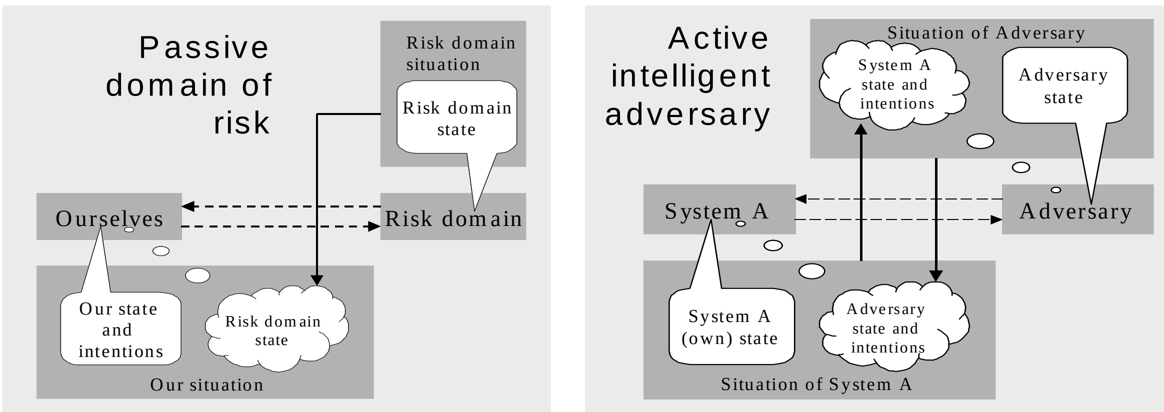

Figure 3: Schematic representations of Passive and Active Domains of Risk.

Need for suitable Knowledge Elicitation techniques

Advocates of human-centered design argue that the needs of users should be addressed as early as possible in the design process. A practical consideration for human factors specialists is the development of suitable KE techniques for analysing how workers interact with their work domain. We are trying to develop models of the work domain and control tasks required to operate in the domain that are objective, invariant and event independent, and at the same time, trying to understand the information workers use in order to extract subjectively meaningful dimensions of the work domain. How does the worker view and measure the work environment, and what information do they use when forming intentions and carrying out decisions? For example, when playing an air defence game, subjects were found to use such information as current fighter fuel load and missile status, as well as the immediate perceived enemy threat when formulating a subjective dimension of the "effective deterrent lifetime" of the fighter aircraft they had in the air (Elliott, Crawford, Solodilova, Sanderson & Naikar, 2000). Klein (1989) states; "By studying in detail the general knowledge, specific information, and reasoning processes an expert uses, a model of the task can be constructed that exhibits some of the properties of the expert being modeled." (pp, 462). In contrast, we are eliciting the SME's understanding of how the domain of work is constructed. Needless to say, the knowledge elicitation techniques used for CWA should be developed according to the purpose and scope of the analysis, and how the information is to be used with respect to the human-machine system being analysed.

In the present paper, three of the KE techniques; a) Crew Cognitive Walkthrough (CCW), b) Adversarial Crew Cognitive Walkthrough (ACCW), and c) Adversarial Crew Cued Recall Debrief (ACCRD), were developed to provide advice on human engineering for a system yet to exist in its physical form, namely Australia's new Airborne Early Warning and control (AEW&C) platform. The fourth technique, Adapted Critical Decision Method (AdCDM), was developed with the intention of assessing its suitability for analysing a system that currently exists, namely an airborne Search and Rescue (SAR) service.

KE technique for systems that currently exist – Adapted Critical Decision Method

The goal of the KE techniue was to develop a WDA of an airborne SAR service. Airborne SAR was chosen as a particularly relevant, as it is a domain in civil aviation that has a large intentional component. Furthermore, we believe SAR to be an excellent test bed for the adequacy of WDA for modeling a domain of risk. We were able to identify two potential passive domains of risk in this environment; a) physical elements, and b) people in distress.

The KE technique used was an adaptation of the Critical Decision Method (CDM), (Klein, 1989). The CDM is a semi-structured interview designed to elicit information about how expert performers make decisions in naturalistic settings, under time pressure. As with the CDM our adapted method focused on non-routine cases, took a case-based approach rather than a general approach, and used semi-structured cognitive probing. However, unlike the CDM, which uses probe questions to enable participants to reflect on their thought processes when making important decisions, the AdCDM used an informal set of questions to probe participants about, (1) the boundary constraints shaping their actions, (2) elements of the world that they were trying to conserve or preserve, (3) the physical objects they were acting upon, and (4) the functionality afforded by these objects. The goal was to develop a set of probe questions that would elicit information about each level of abstraction in the WDA abstraction hierarchy. Examples of these questions and their relationship to each level of the AH can be found in Figure 4. Three members of the SAR crew were interviewed (two winch crew and one pilot). The interviews lasted for an average of one hour each.

WDA level of abstraction	**Sample Probe Question**
Functional Purpose	If you were to sum up the overall purpose of your role in one sentence, what would that be?
Priorities/Values	What aspects of your environment are you trying to maximize and minimize. What are your priorities in order to achieve the mission goal?
Purpose Related Function	What is the goal of doing that?
Object Related Function	What does that piece of equipment actually do? What function does it actually carry out?
Physical Object	What physical objects are you exploiting at this stage?

Figure 4: Sample probe questions form the Adapted Critical Decision Method

From the three interviews we were able to develop a WDA of sea-based SAR consisting of 94 separate elements or nodes. We identified one functional purpose, five priorities and values, 11 purpose related functions, 35 object related functions, and 50 physical objects. A sample of this WDA can be seen in Figure 5.A series of nodes in this diagram are in bold. We can use this analysis to demonstrate how WDA could be used for a performance evaluation of the system. Let us assume that the SAR company has decided that their response time (from the time they get the call until the time they get the patient to medical assistance) needs improvement (c.f. Figure 5, “minimise time to care needed”). The company can use the WDA to guide their thinking towards areas of their operations that may be improved in order to achieve this priority. Let us assume that the company has top-of-the-range aircraft and that their crews are trained to a high degree of proficiency at evacuating people in distress.

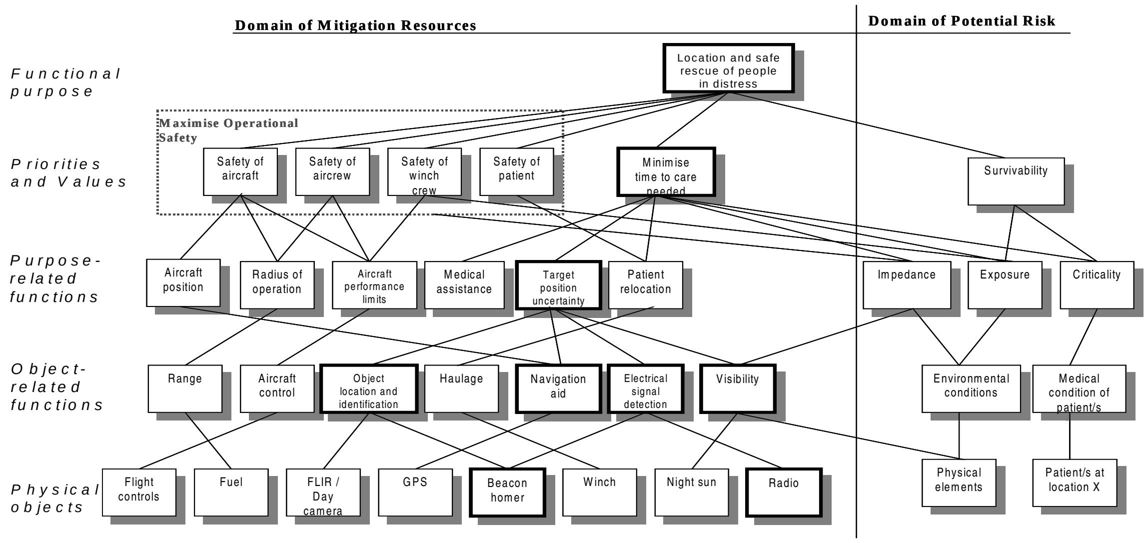

Figure 5: Sample of the SAR work domain, showing the domain of mitigation resources and the domain of risk.

At the Purpose-Related Functions level, we can see that the remaining means for shortening response times is to locate the patient’s position more rapidly and accurately (c.f. Figure 5, “target position certainty”). If we move down another level of abstraction to the Object-Related Functions level, the company is able to discover that the means of locating a patient, which they could possibly improve, is through electronic signal detection. Again moving down to the Physical Object level the company realises that they need to upgrade their beacon homers, or explore the market for new equipment designed to detect and locate emergency beacons.

This example demonstrates that through the use of effective KE techniques, CWA and in particular WDA can be usefully applied to a stage of the system design life cycle such as performance evaluation for an existing system. The following section will investigate the same issues for a system that is yet to exist in its physical form (AEW&C), and is therefore at the requirements/specifications stage of the system design life cycle.

Three KE techniques for analyzing a system that is yet to exist in its physical form – CCW, ACCW, ACCRD for AEW&C

The Crew Cognitive Walkthrough (CCW) procedure, (Elliott et al, 2000); (Naikar, Drumm, Pearce & Sanderson, 2000) was developed as an extension of the Cognitive Walkthrough procedure (Lewis, Polson, Wharton & Rieman, 1990) to perform walkthroughs of group interactions in the AEW&C domain. In this procedure, air defence Subject Matter Experts (SME’s) were required to design two typical air defence scenarios. They were then taken through a pencil and paper walkthrough of these scenarios, and after important phases, they were asked questions about how they were allocating various tasks to their aircrew. Two major limitations with the CCW technique were identified. First, the SME’s designed the scenarios that they played out, as is quite usual for the CW technique. However, because the aim of the simulations was to reflect an adversarial, unpredictable, intentional system, then having the SME’s know what was to happen was quite artificial. Second, the SME’s lapsed into a “comfort zone” where they played out the scenario as if all of their actions had the desired consequences and nothing unpredicted was assumed to have occurred. The SME's never placed themselves in a situation that they could not handle with the resources they had.

In order to address these concerns we developed the Adversarial Crew Cognitive Walkthrough (ACCW) technique. The rationale behind the ACCW was to develop an environment more representative of a dynamic, uncertain, system, so that a wider range of issues (due to the addition of an adversary) would emerge to support modeling. We believed that the resulting WDA would cover the AEW&C work domain in more representative situations (e.g. combat, uncertainty, greater exercise of priorities, criteria, functions, resources, etc…). A generic air-defence simulation was generated using a commercial of the shelf war game (Janes Fleet Command) as an environment generator. Participants were four war-gaming experts. Once the participants had been briefed on the nature of the simulation and familiarised with the Janes Fleet Command environment, the simulation began, with a research assistant playing the side of the adversary in real time. As with the CCW, at important phases of the simulation the game was paused and participants were asked a set of probe questions.

A potential concern with any walkthrough technique is its seemingly reactive nature. Will asking participants questions "on the fly" change the way that they are naturally thinking about the task, and will they alter their behaviour accordingly? We needed a technique that was non-intrusive and that did not disrupt the dynamic nature of the simulation. Until recently, the lack of suitable investigative methodologies for capturing data in field settings has compromised the empirical investigation of command and control work environments. A recent technique addressing this problem is the Cued Recall Debrief (Omodei, Wearing & McLennan 1997). First, the subject wears a lightweight head-mounted video camera, as a non-intrusive method for the recording of sequences of events in the simulated task. A head-mounted camera is used because it captures the events as the subject sees them from their "own point of view". When the video is replayed to the subject afterwards, the subject becomes fully re-immersed, cueing the comprehensive recall of psychologically relevant experiences. Second, the subject uses the video replay as a stimulus to recall all mental events (e.g. thoughts, uncertainties, constraints upon action, etc…), which occurred, and to verbalize these onto a second videotape. The main advantage of this technique is that it does not interrupt the natural flow of events of the simulation in order to elicit information from the subjects.

Comments made by participants were divided into two main categories of interest; comments made about their side, and comments made about their adversary. These categories were further divided into three sub-categories relating to; a) the situation, b) intentions, and c) constraints upon action. The ACCW process resulted in an approximately equal number of statements in each category of interest (comments about self and comments about adversary), due mostly to the structured questioning at regular intervals. The ACCRD leads to more statements overall about one's own activities and around the same number of statements about the adversary as the ACCW. Also there are markedly more statements about one's own situation and intentions for the ACCRD than for the ACCW, and if anything slightly fewer statements about the adversary's intentions and constraints. For a more comprehensive account of this study see Elliott et al (2000).

With respect to the development of CWA products, it appears that the ACCW technique was effective for generating an adversarial-based WDA incorporating a domain of risk. The ACCRD technique was particularly effective for allowing participants to talk about what they were doing and why, and thus seems more effective for measuring the subjective dimensions by which workers measure their work environment. A sample of the WDA resulting from this study can be seen in figure 6. By eliciting both objective information about the work domain, and subjective information about the type of information workers are thinking about, we can begin to specify design requirements for a human-machine system that is yet to exist in its physical form.

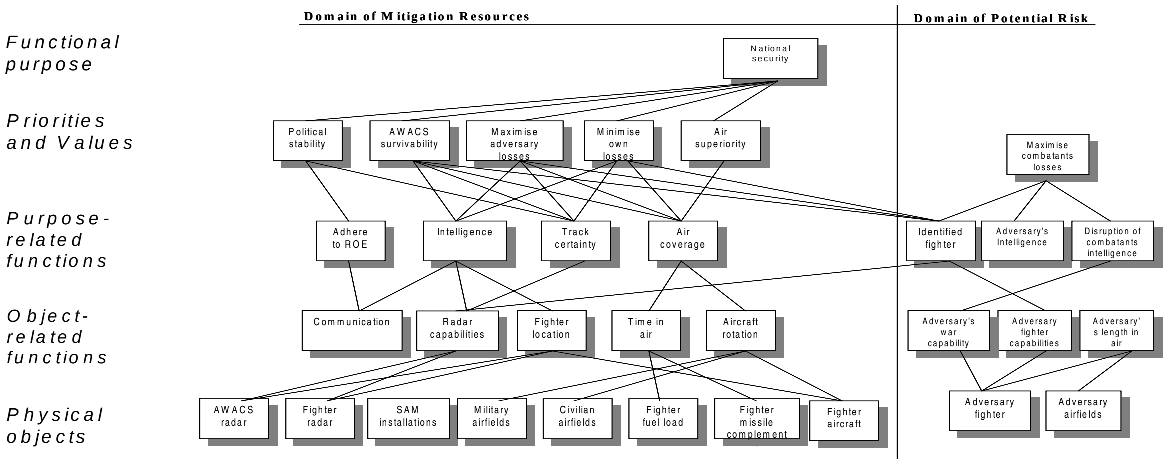

Figure 6: Sample of the AEW&C work domain, showing the domain of mitigation resources and the domain of risk.

Conclusions

We have highlighted what we believe to be four particularly useful knowledge elicitation techniques for developing CWA products. For a system in civil aviation that already exists (airborne SAR) we adapted Klein's critical decision method as the knowledge elicitation technique. From only three interviews, and four days data reduction and analysis we were able to create a WDA consisting of approximately 100 nodes. By observing merely a small slice of this WDA we are able to show the practical use of such a product for stages of the system design life cycle such as system performance evaluation. With respect to AEW&C, which is yet to exist in its physical form, we developed two KE techniques (ACCW & ACCRD) that were designed to elicit information for producing a WDA that was adversarial in nature. Although the analysis of these results is still incomplete, early findings reveal that the ACCW results in a more even distribution of comments made between participant's own forces (domain of mitigating resources) and their adversary's forces (domain of risk). On the other hand, the ACCRD results in participants talking proportionately more about their own forces activities. These initial results are yet to be validated by a comparison of the WDA products generated from the ACCW and ACCRD. To date a WDA illustrating the relationship between the domain of mitigating resources and the domain of risk has been developed only for the ACCW (see Figure 6).

When developing an appropriate methodology for eliciting knowledge from domain experts, one should be mindful of the type of information required with relation to the type of modelling products to be developed, and/or the design questions to be answered. Furthermore, researchers should be aware of the stage of the system design life cycle for which the information elicited will be used. For a system that is yet to exist in its physical form, we took the approach of building a generic air defence simulation, which was designed to closely approximate the demands of such environments. By using 'borrowed' subject matter experts we effectively allowed the workers to "finish the design requirements", much the same way that Rasmussen and Goodstein (1987) talked about allowing workers to "finish the design" when designing flexible work environments.

References

Elliott, G., Crawford, J., Solodilova, I., Sanderson, P., & Naikar, N. (2000). Comparison of Adversarial Crew Cognitive Walkthrough and Cued Recall debrief techniques for eliciting information in a D^{3}M environment. Paper proposed for the *Proceedings of the 2000 Annual Conference of the Computer-Human Interaction Group (CHISIG) of the Ergonomics Society of Australia (OzCHI2000).* Sydney, Australia. December 4-8.

Gualtieri, J., Roth, E. M., & Eggleston, R. G. (2000). Utilizing the Abstraction Hierarchy for Role Allocation and Team Structure Design. *Proceedings of the 5th International Conference on Human Interaction with Complex Systems.* Urbana, Illinois. April 30-May 2.

Hajdukiewicz, J. R., Burns, C., Vicente, K. J., and Eggleston, R. G. (1999). Work Domain Analysis For Intentional Systems. *Proceedings of the 43rd Annual Meeting of the Human Factors and Ergonomics Society*. Santa Monica, CA: Human Factors and Ergonomics Society, 333-337

Klein, G. A., Calderwood, R., & MacGregor, D. (1989). Critical decision method for eliciting knowledge. *IEEE Transactions on Systems, Man, and Cybernetics* 19(3), 462-472.

Leveson, N. G. (2000). Intent specifications: An approach to building human-centred specifications. *IEEE Transactions on Software Engineering* 26(1), 15-35.

Lewis, C., Polson, P., Wharton, C., & Rieman, J. (1990). Testing a walkthrough methodology for theory-based design of walk-up-and-use interfaces. *Proceedings ACM CHI '90 Conference Seattle,* WA, pp 235-242.

Moray, N., Sanderson, P. M. & Vicente, K. (1992). Cognitive Task analysis of a complex work domain: A case study. *Reliability Engineering and System Safety*, 36, 207-216.

Naikar, N., Drumm, D., Pearce, B., & Sanderson, P. (2000). Designing new teams with cognitive work analysis. Paper presented at the Fifth Australian Aviation Psychology Symposium. Manly, November 20-24.

Omodei, M., Wearing, A., & McLennan, J. (1997). *Head-mounted video recording: A methodology for studying naturalistic decision making.* In Flin, R., Salas, E., Strub, M., & Martin, L. (Eds), Decision Making Under Stress: Emerging Themes and Applications. Aldershot : Ashgate.

Rasmussen, J. (1979). *On the structure of knowledge – A morphology of mental models in a man-machine system context* (Report No. Riso-M-2192). Roskilde, Denmark: Riso National Laboratory, Electronics Department

Rasmussen, J. (1985). The role of hierarchical knowledge representation in decision making and system management. *IEEE Transactions on Systems, Man, and Cybernetics, SMC-15*, 234-243.

Rasmussen, J. & Goodstein, L. P. (1987). Decision support in supervisory control of high-risk industrial systems. *Automatica*, 23, 663-671.

Rasmussen, J., Pejtersen, A. M., & Goodstein, L. P. (1994). *Cognitive systems engineering*. New York: John Wiley & Sons.

Reising, D. V. (1999). The impact of instrumentation location and reliability on the performance of operators using an ecological interface for process control. Unpublished PhD thesis. Department of Mechanical and Industrial Engineering, University of Illinois, Urbana, IL.

Sanderson, P. (2000). Cognitive Work Analysis across the system life-cycle: Achievements, challenges and prospects. Paper presented at the Fifth Australian Aviation Psychology Symposium. Manly, November 20-24.

Sanderson, P., Naikar, N., Lintern, G., & Goss, S., (1999). Use of Cognitive Work Analysis across the system life cycle: Requirements to decommissioning. *Proceedings of the 43rd Annual Meeting of the Human Factors and Ergonomics Society*. Santa Monica, CA: Human Factors and Ergonomics Society. Pp 318-322.

Vicente, K. (1999). *Cognitive work analysis: Toward safe, productive, and healthy computer-based work.* Mahweh, NJ: LEA.

Wong, W. B. L., Sallis, P. J., & O'Hare, D. (1998). The ecological approach to interface design: Applying the abstraction hierarchy to intentional domains. *Proceedings of the Eighth Australian Conference on Computer-Human Interaction OzCHI'98.* Pp 144-150. Australia: IEEE Computer Society Press.