A HISTORY OF COGNITIVE ENGINEERING RESEARCH AT RISØ (1962-1979)

Kim J. Vicente
Cognitive Engineering Laboratory
Department of Mechanical & Industrial Engineering
University of Toronto
Toronto, Canada

Following the theme for this year's conference, this paper contributes to ongoing discussions defining the future of cognitive engineering research by examining a part of its past. The history of one particular line of research, that of the Electronics Department at Risø National Laboratory, is reviewed. A number of important studies, conducted between 1962 and 1979, are briefly described. Among these are operational experience acquired from the introduction of a prototype digital console in a nuclear research reactor, two field studies of professional operators conducting representative tasks in representative settings (electronic trouble-shooting and conventional power plant control), and analyses of over 645 human error reports in the nuclear and aviation industries. Some of the themes characterizing the Risø research program in cognitive engineering are briefly summarized. These themes help define what cognitive engineering is, and what it might be concerned with in the future.

INTRODUCTION

The Cognitive Engineering and Decision Making Technical Group is a relative newcomer to the Human Factors and Ergonomics Society. Consequently, it is useful to engage in efforts of self-examination to help define future research in this area. Just as it is important to reflect upon what we should be teaching our students (Woods et al., 1996) and the current state of our field (Endsley et al., 1995), it is also useful to study where we have come from. This, in turn, will help us make more deliberate decisions about how to proceed in the future. This paper tries to achieve this goal by presenting one history of cognitive engineering. I say "one" because there are several histories that can be told, each from a different intellectual lineage stemming from one of the many people who contributed towards defining the discipline in the 1960s and 1970s (e.g., Bainbridge, Duncan, Johannsen, Leplat, Moray, Rouse, Sheridan, Stassen). The history presented here is that of the cross-disciplinary research group in the Electronics Department of Risø National Laboratory. This research team included Jens Rasmussen, Len Goodstein, Morten Lind, and Erik Hollnagel. Rasmussen was the group's conceptual leader (Rouse, 1988), so I will focus mainly on his publications, although these were undoubtedly influenced by others, especially Goodstein and Lind who worked at Risø since the 1960s. Hollnagel did not join the group until April of 1978 (Erik Hollnagel, personal communication, 1997).

Motivation

There are several reasons why it is important to present the history of this particular lineage. First, the Risø group has been very influential in defining the field of cognitive engineering. For example, Reason (1990) observed that some of the ideas arising from the Risø group have become "market standards" in the cognitive engineering community (p. xiii). Moray (1988) stated that the Risø perspective is "nothing less than a paradigm shift" (p. 12) in the study of complex human-machine interaction. Second, despite this impact, the history of the Risø program is not very well known, except for the more focused analysis of Sanderson and Harwood (1988). Many of the insights from this line of work were only published as technical reports, which have not been widely available or read. As I will try to show, there is a great deal that can be learned from these old, obscure reports.

Because the more recent history of cognitive engineering is better known (e.g., Norman, 1981, 1986; Hollnagel & Woods, 1983; Rasmussen, 1986; Rasmussen, Pejtersen, & Goodstein, 1994), this paper will focus on the period 1962-79. Several key contributions published during this time will be reviewed. To add context, two additional historical data points are worth mentioning: (1) the first time, that I know of, that the term "cognitive engineering" appeared in print (Fischoff, Slovic, & Lichtenstein, 1978, p. 343); and (2) the first time that the term "cognitive engineering" was used in print to describe a discipline with a unique set of characteristics (Norman, 1981).

Caveats

Writing a five page history of almost two decades of research conducted by a productive and influential laboratory is bound to be an incomplete endeavor. To avoid any misinterpretations, it is important to make several caveats explicit. First, I have deliberately tried to present an intellectual history rather than a sociological one. Sociological histories of science are fascinating but they are notoriously subjective, presenting contradictory viewpoints caused by the biases introduced by the different participants. Consequently, I have based my account on scientific publications rather than on interviews with the researchers who participated in this history (although I have had contact with all of the participants listed earlier, starting when I first worked at Risø in 1987-88). The benefit of this approach is that my historical account is backed up in detail by citations that can be verified. The disadvantage is that I have ignored the interpersonal dynamics that characterize all scientific activities (Hull, 1988). Therefore, I have traded off comprehensiveness for demonstrability.

Second, because I am relying on scientific publications as documented evidence, the sequence of ideas I present should not be confused with the sequence of thought processes that the participants engaged in while they conducted the research before publication. As Feynman (1966) observed, "We have a habit in writing articles ... to make the work as finished as possible, to cover up all the tracks, to not worry about the blind alleys or describe how you had the wrong idea at first, and so on. So there isn't any place to publish, in a dignified manner, what you actually did in order to get to do the work". Thus, using the terms of philosophy of science, I am presenting a history of the context of justification, not a history of the context of discovery.

Third, I do not address the broader intellectual climate in which the Risø ideas evolved. An international community of researchers in human factors, psychology, systems engineering, biology, artificial intelligence, philosophy of science, and cognitive science influenced the Risø group (Morten Lind, personal communication, 1997). This broader influence is important and interesting, but there is not enough space to discuss it here.

Fourth, the fact that I have chosen to write this history paper does not mean that I believe that all of the important problems in cognitive engineering have been solved -- far from it. We have barely scratched the surface on many key issues, and much important and challenging work remains to be done. However, as I will try to show in the remainder of the paper, taking this reflective look at the past might help us in directing our future research activities.

RISØ NATIONAL LABORATORY

Risø National Laboratory (or Research Establishment Risø, as it was first known) was created in 1956 and Niels Bohr, the Danish Nobel laureate in physics, served as its first chairman of the board. Risø was given the charge of conducting research so that Denmark could effectively implement nuclear power within 5 years. Remarkably, this 5 year window was maintained for over a quarter of a century until it was decided that Denmark would not have any commercial nuclear power plants!

During this quarter century, Risø fostered an exceptionally unique environment for conducting research. Originally, the laboratory's funding came from the Danish ministry of Finance, providing a vast supply of financial support. There was no requirement at all to bring in large research contracts from external funding agencies. Furthermore, there was no requirement to publish research results in academic journals. Instead, much of the work described below was published in an internal series of green technical reports. Although Risø had collaborations with universities, its researchers were not required to teach classes, supervise graduate students, or take on extensive administrative responsibilities. What they were required to do was conduct research to address a practical problem of great social relevance -- how to effectively support Danish government and industry when the decision was made to introduce nuclear power, roughly within 5 years.

Anyone familiar with academic or industrial research will recognize that these were extraordinary circumstances. All of this has changed, particularly in the last 10 years. Like almost all other research institutions around the world, Risø is now under extensive external scrutiny and intense economic pressure. Researchers are required to spend substantial amounts of time writing proposals to bring in large amounts of external funding. They are also required to engage in extensive project management activities that are associated with large research contracts. Furthermore, researchers often cannot afford to adopt long-term research goals, but instead are required to produce demonstrable short-term products. They are also required to document their productivity by publishing extensively in academic circles -- the familiar doctrine of publish or perish. But during the period described in this paper (1962 - 79), none of these constraints were present. Although it is difficult to substantiate this claim with any certainty, I believe that it is not an accident that the research described in the remainder of this paper was conducted at an institution with this type of research environment.

PRE-HISTORY (62-63)

The first body of research to be reviewed is actually more accurately described as pre-history in that it was not directly concerned with issues that we readily recognize today as pertaining to cognitive engineering. Nevertheless, the activities of this period are relevant because they explain the origins of the Risø research program in cognitive engineering.

Three research reactors were installed on site to support Risø's mission. The Electronics Department was responsible for the commissioning and safety certification of the instrumentation of the research reactors. And since the head of the Electronics Department (Rasmussen) became chairman of the Risø Reactor Safety Committee, it should not be surprising to find that the Electronics Department's early efforts focused on analyzing the reliability of reactor equipment and instrumentation (Rasmussen & Timmerman, 1962; Jensen, Rasmussen, & Timmerman, 1963). Note that the focus was strictly on examining hardware reliability, not human reliability, using probabilistic mathematical models. Probability of equipment failure and the degree of redundancy required in backup safety systems to achieve the desired level of reactor safety were both investigated. Furthermore, considerable attention was paid to collecting failure data under representative conditions, a theme that would re-emerge in later work. Interestingly, there is one passage in these reports which foreshadows the next phase of research: "Low [probability] figures must be used with great care, because the reliability in this case may be governed by factors not dealt with in this report" (Jensen et al., 1963, p. 31).

MOTIVATING PROBLEMS (68-69)

The next phase of research in the Risø program was conducted during the mid-1960s. Based on the work described above, researchers found that they could design redundant reactor safety systems with extremely high technical reliability, yet accidents still occurred. To solve this mystery, a review of 29 cases with major consequences to either plant or personnel in the nuclear domain and of 100 accidents in air transportation was conducted (Rasmussen, 1969). The results of this review revealed that the reliability of complex systems, such as nuclear reactors, cannot be viewed from a strictly technical viewpoint without considering the human element in the system (Rasmussen, 1968b, 1968c). It became apparent that the human operator played a key role in overall system reliability and safety. The review also revealed that accident-causing errors arose because operators were confronted with unfamiliar situations that had not been, and could not have been, anticipated by system designers (Rasmussen, 1969). In contrast, under normal circumstances, a trained and experienced operator would often be able to compensate for deficiencies in the interface (Rasmussen, 1968a, 1969). Consequently, the single most important concern in improving system safety is to provide operators with the support required to adapt to unfamiliar and unanticipated abnormal situations. The research program subsequently undertaken at Risø's Electronics Department was directed at developing a design framework to deal with this challenging practical problem.

Around this time, a prototype console was installed in a room adjacent to one of the research reactors' control rooms. This console was instrumented with a set of displays that were well before their time in that they included: the use of digital computers, overview displays, emergent feature graphics, and display of higher-order functional information (Goodstein, 1968). Operator interaction with this experimental prototype design was observed over a number of years in the field. In today's language, we would say that this phase of the Risø research program consisted of prototype building and usability analysis in a naturalistic setting. These rich field observations led to a number of insights which had a significant impact on the direction of subsequent research.

EMPIRICAL & CONCEPTUAL DEVELOPMENTS (72-79)

One of the important contributions to stem from Risø is a field study of electronic trouble-shooting strategies (Rasmussen & Jensen, 1973). From a methodological point of view, the study is interesting because most of the data were based on verbal protocols, despite the fact that such reports were considered unreliable by many psychologists at that time. Also, the study was conducted under highly representative conditions, with professional technicians diagnosing complex faults in commercially available electronic equipment. In this sense, this study was a precursor of subsequent research which also shared these methodological features. From an empirical point of view, this study was interesting because it showed that there can be several, very different strategies that can be used to perform the same task, and that people would switch between those strategies during their problem solving activities. The study also showed that the way in which a person formulates the task (e.g., find the faulty component as quickly as possible vs. understand why the equipment is faulty in an elegant way with minimal observations) has a very strong effect on their strategies. Thus, to understand why people tend to use certain strategies, one has to identify the performance criteria that people choose to adopt. Otherwise, adaptive behavior can actually seem irrational from an analytical perspective. Thus, the electronic trouble-shooting field study was important because it showed that expert strategic behavior in a representative setting could be described systematically, despite its apparent complexity.

Another important contribution was the development of the decision ladder which represents information processing activities (activation, observation, identification, interpretation, evaluation, task definition, procedure formulation, execution) in the shape of a ladder with shortcuts in between the two legs of the ladder (Rasmussen, 1974). This conceptual development originated from a field study of the cognitive activities of professional operators starting up a conventional power plant over a one week period, again using verbal protocols. The decision ladder was based on two insights. First, it is possible to parse a descriptive or normative timeline of human cognitive activities into a basic number of recurring decision tasks. Second, and more importantly, experts do not follow all of the information processing steps that a novice would perform. Instead, experts rely on their knowledge and experience and thereby exhibit direct shortcuts (or associations) which allow them to bypass several cognitive activities. These shortcuts account for the increased speed and reduced effort that are hallmarks of expert performance.

The verbal protocols from the two field studies just described also gave some insights into how people represent complex systems during problem solving. These insights eventually coalesced in the development of the abstraction hierarchy (Rasmussen, 1979b), a framework that describes complex systems at various levels of abstraction in a psychologically relevant fashion. The abstraction hierarchy is based on several important observations. First, when experienced operators are solving problems in the context of complex human-machine systems, they spontaneously adopt, and switch between, different models of the system in order to match the immediate task demands. Some of these models provide physical information whereas others provide functional information. Second, it is possible to represent engineering systems in a way that makes reference to purpose, thereby spanning material form and functional meaning. This is accomplished through the means-ends links between levels of the abstraction hierarchy. The result is a representation framework that bridges the gap between the technical and the psychological, thereby supporting goal-directed human problem solving.

Another important activity in the Risø program was the analysis of 516 human error reports from the nuclear industry (Rasmussen, 1978, 1979a). Once again, the focus was on analyzing data on human behavior in representative conditions. These analyses shed light on the task characteristics and psychological mechanisms that were responsible for human error in complex systems. Implications were also derived for quantitative risk analysis.

More importantly, these analyses eventually led to the skills, rules, knowledge (SRK) framework (Rasmussen, 1979b), which categorizes three qualitatively different ways in which people can interact with the environment. As such, it provides a taxonomy for models of human performance, and it can be used to derive useful implications for design as well (e.g., Vicente & Rasmussen, 1992). The utility of the taxonomy is evident by the influence it has had on the cognitive engineering community (Sanderson & Harwood, 1988; Reason, 1990).

CONCLUSIONS

Contributions

The concepts and findings generated by the Risø program have had an exceptionally significant impact on both applied practice and basic research. For example, these concepts served as the conceptual basis for the design of advanced control rooms for nuclear power plants by Westinghouse in the US (Easter, 1987) and by Toshiba in Japan (Itoh, Sakuma, & Monta, 1995). The Toshiba implementation is particularly noteworthy since it has been constructed on a very large scale, a prototype advanced control room connected to a full-scope nuclear power plant simulator. More recently, some of these concepts have also led to a new theoretical account of the relationship between expertise and memory recall (Vicente & Wang, in press), making a significant contribution to basic theory in psychology. The Risø concepts have also had an acknowledged influence on prominent researchers in cognitive science (e.g., Norman, 1993, p. 257). It is very difficult to think of another line of research which has had such a demonstrable influence on practical problems in industry, basic theory in psychology, and basic research in cognitive science.

Limitations

Despite these notable contributions, there are still many gaps in the Risø line of research. First, very few of the concepts have been rigorously tested experimentally. Most of the emphasis to date has been on inductively developing concepts from field studies. The predictions that arise from these concepts should be directly tested through more controlled experimentation. Second, the generalizability of these ideas to different applications domains needs to be evaluated further. Originally, the ideas were developed in the context of process control, and some attempts have been made to apply the same concepts to other domains (e.g., aviation, information retrieval, engineering design). However, more of this type of work should be done to evaluate the breadth of applicability of the Risø concepts. Third, more effort needs to be devoted to formalizing these concepts so that they can be used by researchers and designers who were not involved in their development. Currently, it is difficult for anyone to just pick up these ideas and apply them to the problems with which they are concerned. It is important to make these conceptual tools readily available to a broader community.

Implications

Because multifaceted contributions of the type discussed above are rare in human factors research, it is worthwhile speculating as to what characteristics of the Risø program may have led to this success. There are a number of themes that can be identified:

- the research started with a practical problem of social and economic relevance, not particular theories, methods, or generic academic curiosity
- the concepts developed were continually informed by, and focused by, very intensive analyses of data collected under representative conditions (field studies, operating experience, human error reports)
- early empirical results changed the original focus of the research (hardware reliability) to a new set of issues (supporting operator adaptation to novelty) that were found to be of greater relevance to the problem of interest.

This research strategy is markedly different from that adopted in most North American human factors research, which has focused primarily on well-controlled laboratory experiments (Meister, 1995). Perhaps, then, the methodological example offered by the Risø program can complement traditional human factors practices (cf. Vicente, 1997). Seeking out practical problems, analyzing them in the field, being sensitive to the results obtained, and changing the focus of the research as need be can help cognitive engineering make its mark, not only on applied practice, but on psychology and cognitive science as well.

Adopting this research strategy will not be an easy task, however, because it requires a particular type of infrastructure and institutional support. As mentioned earlier, the research environment at Risø during the period covered in this paper was quite unique. Researchers were explicitly problem-driven rather than paradigm-driven; they had the luxury of being able to tackle research issues that were relevant to the problem at hand rather than research issues that were fashionable with funding agencies or journal editors and reviewers; they could afford to adopt a long-term rather than a short-term approach to their work; they could choose to adopt meaningful research methods that were very time-consuming and laborious rather than having to resort to methods that generated any kind of results efficiently; they could let the research findings dictate their next step rather than having to stick to the deliverables defined in a research contract; they had the time to focus and think uninterruptedly about their research rather than having to time-share most of their attention among a number of other activities such as administration, teaching, and proposal writing. A research setting with these characteristics is almost unheard of in contemporary society. Yet one could argue that these are ideal conditions for cognitive engineering research. If so, then perhaps the largest challenge facing our discipline may not be to do the appropriate research, but rather to create the conditions so that such research can be conducted.

ACKNOWLEDGEMENTS

The writing of this paper was sponsored by a research contract with the Japan Atomic Energy Research Institute (Dr. Fumiya Tanabe, Technical Monitor) and by grants from the Natural Sciences and Engineering Research Council of Canada. I would like to thank Dr. Tanabe and Dr. Kazuo Monta for their encouragement, and Cathy Burns, Jeff Caird, Klaus Christoffersen, Erik Hollnagel, Greg Jamieson, Morten Lind, Neville Moray, Jens Rasmussen, and David Woods for their comments on earlier versions.

REFERENCES

Easter, J. R. (1987). Engineering human factors into the Westinghouse advanced control room. Nuclear Engineering International, 32(May), 35-38.

Endsley, M. R., Klein, G., Woods, D. D., Smith, P. J., & Selcon, S. J. (1995). Future directions in cognitive engineering and naturalistic decision making. In Proceedings of the Human Factors and Ergonomics Society 40th Annual Meeting (pp. 450-453). Santa Monica, CA: HFES.

Feynman, R. P. (1966). The development of the space-time view of quantum electrodynamics. Science, 153, 699-708.

Fischoff, B., Slovic, P., & Lichtenstein, S. (1978). Fault trees: Sensitivity of estimated failure probabilities to problem representation. Journal of Experimental Psychology: Human Perception and Performance, 4, 330-344.

Goodstein, L. P. (1968). An experimental computer-controlled instrumentation system for the research reactor DR-2. In Application of on-line computers to nuclear reactors (pp. 549-566). Halden, Norway: OECD Halden Reactor Project.

Hollnagel, E., & Woods, D. D. (1983). Cognitive systems engineering: New wine in new bottles. International Journal of Man-Machine Systems, 18, 583-600.

Hull, D. L. (1988). Science as a process: An evolutionary account of the social and conceptual developments of science. Chicago: University of Chicago Press.

Itoh, J., Sakuma, A., & Monta, K. (1995). An ecological interface for supervisory control of BWR nuclear power plants. Control Engineering Practice, 3, 231-239.

Jensen, A., Rasmussen, J., & Timmerman, P. (1963). Analysis of failure data for electronic equipment at Risø (Risø Report No. 38). Roskilde, Denmark: Danish Atomic Energy Commission, Research Establishment Risø.

Meister, D. (1995). Divergent viewpoints: Essays on human factors questions. Unpublished manuscript.

Moray, N. (1988). Ex Risø semper aliquid antiquum: Sources of a new paradigm for engineering psychology. In L. P. Goodstein, H. B. Andersen, and S. E. Olsen (Eds.), Tasks, errors,and mental models: A festschrift to celebrate the 60th birthday of Professor Jens Rasmussen (pp. 12-17). London: Taylor and Francis.

Norman, D. A. (1981). Steps toward a cognitive engineering: System images, system friendliness, mental models (Tech. Rep.). La Jolla, CA: UCSD.

Norman, D. A. (1986). Cognitive engineering. In D. A. Norman and S. W. Draper (Eds.), User centered system design: New perspectives on human-computer interaction (pp. 31-61). Hillsdale, NJ: Erlbaum.

Norman, D. A. (1993). Things that make us smart: Defending human attributes in the age of the machine. Reading, MA: Addison-Wesley.

Rasmussen, J. (1968a). On the communication between operators and instrumentation in automatic process control plants (Risø-M-686). Roskilde, Denmark: Danish Atomic Energy Commission, Research Establishment Risø.

Rasmussen, J. (1968b). On the reliability of process plants and instrumentation systems (Risø-M-706). Roskilde, Denmark: Danish Atomic Energy Commission, Research Establishment Risø.

Rasmussen, J. (1968c). Characteristics of operator, automatic equipment and designer in plant automation (Risø-M-808). Roskilde, Denmark: Danish Atomic Energy Commission, Research Establishment Risø.

Rasmussen, J. (1969). Man-machine communication in the light of accident records (S-1-69). Roskilde, Denmark: Danish Atomic Energy Commission, Research Establishment Risø.

Rasmussen, J. (1974). The human data processor as a system component: Bits and pieces of a model (Risø-M-1722). Roskilde, Denmark: Danish Atomic Energy Commission.

Rasmussen, J. (1978). Operator/technician errors in calibration, setting, and testing nuclear power plant equipment (N-17-78). Roskilde, Denmark: Risø National Laboratory, Electronics Department.

Rasmussen, J. (1979a). Preliminary analysis of human error cases in U.S. licencee event reports (N-8-79). Roskilde, Denmark: Risø National Laboratory, Electronics Department.

Rasmussen, J. (1979b). On the structure of knowledge - A morphology of mental models in a man-machine system context (Risø-M-2192). Roskilde, Denmark: Risø National Laboratory, Electronics Department.

Rasmussen, J. (1986). Information processing and human-machine interaction: An approach to cognitive engineering. Amsterdam: North-Holland.

Rasmussen, J., & Jensen, A. (1973). A study of mental procedures in electronic trouble shooting (Risø-M-1582). Roskilde, Denmark: Danish Atomic Energy Commission, Research Establishment Risø.

Rasmussen, J., & Timmerman, P. (1962). Safety and reliability of reactor instrumentation with redundant instrument channels (Risø Report No. 34). Roskilde, Denmark: Danish Atomic Energy Commission, Research Establishment Risø.

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

Reason, J. (1990). Human error. Cambridge, England: Cambridge University Press.

Rouse, W. B. (1988). Ladders, levels and lapses; and other distinctions in the contributions of Jens Rasmussen. In L. P. Goodstein, H. B. Andersen, and S. E. Olsen (Eds.), Tasks, errors and mental models: A festschrift to celebrate the 60th birthday of Professor Jens Rasmussen (pp. 315-323). London: Taylor and Francis.

Sanderson, P. M., & Harwood, K. (1988). The skills, rules, and knowledge classification: A discussion of its emergence and nature. In L. P. Goodstein, H. B. Andersen, and S. E. Olsen (Eds.), Tasks, errors,and mental models: A festschrift to celebrate the 60th birthday of Professor Jens Rasmussen (pp. 21-34). London: Taylor and Francis.

Vicente, K. J. (1997). Heeding the legacy of Meister, Brunswik, and Gibson: Toward a broader view of human factors research. Human Factors, 39.

Vicente, K. J., & Rasmussen, J. (1992). Ecological interface design: Theoretical foundations. IEEE Transactions on Systems, Man, and Cybernetics, SMC-22, 589-606.

Vicente, K. J., & Wang, J. (in press). An ecological theory of expertise effects in memory recall. Psychological Review.

Woods, D. D., Watts, J. C., Graham, J. M., Kidwell, D. L., & Smith, P. J. (1996). Teaching cognitive systems engineering. In Proceedings of the Human Factors and Ergonomics Society 40th Annual Meeting (pp. 259-263). Santa Monica, CA: HFES.