Applying Cognitive Theories & Methods to the Design of Computerised Medical Decision Support

Eugenio Alberdi (e.alberdi@abdn.ac.uk)
Department of Psychology
University of Aberdeen
Aberdeen AB24 2UB, Scotland (UK)

Robert Logie (r.logie@abdn.ac.uk)
Department of Psychology
University of Aberdeen
Aberdeen AB24 2UB, Scotland (UK)

Abstract

This paper describes an approach to cognitive engineering which promotes a symbiosis between the theories and methodologies of cognitive psychology and the practices of human-computer interaction design. We ground the description of our approach in a particular design problem: the development of computerised decision support in medical intensive care. We review the psychological literature of medical reasoning and decision making, highlighting its potential to inform the design of medical computerised aids. We also discuss how addressing this design problem may in turn benefit cognitive theory. This is followed by a brief description of our proposed methodology.

Introduction

Intensive care monitoring is an area of health care in which computerised aids may play a significant role. The automatic collection, interpretation and display of monitoring data offer considerable potential for improving the quality of clinical care. However, it has been shown repeatedly that computerised monitoring does not always lead to the sought-for improvements (see e.g. Green, Gilhooly, Logie, & Ross, 1991; McIntosh, Cunningham, & Elton, 1994). It can be argued that the main reason for these difficulties is the failure to develop an understanding of the ways in which the eventual system users make clinical decisions at the intensive care unit (ICU).

In this paper we discuss a cognitive engineering approach to the problem. The goal is to understand the cognitions and working practices of clinical staff, with a view to developing design principles for computerised decision support in intensive care. We characterise cognitive engineering as a discipline which integrates, in a symbiotic way, psychological research and the practices of human-computer interaction design.

Symbiotic Approach to Cognitive Engineering

In the last two decades, a great deal of effort has gone into applying theories and models developed by cognitive psychologists to inform the design of human-computer applications. This enterprise has often been referred to as "cognitive engineering" (Norman, 1986). However there are strong suggestions that psychological knowledge has not had a significant impact on system design (see e.g. Barnard & Harrison, 1988); Carroll, 1991; Landauer, 1987). This has led to new characterisations of cognitive engineering which essentially exclude psychological practices from human-computer studies (e.g. Long & Dowell, 1989, 1996).

In contrast, we argue that cognitive psychology can play an important role in engineering design, and that system design and psychological theories and methods can support each other by maintaining a symbiotic relationship.

We would like to emphasise that cognitive science can play (and, in fact, has played) important roles in the development of usable knowledge for human-computer interaction. If rightly applied, a great deal of what is known about human cognition can have important implications for design. In particular, cognitive science research can play a number of important roles in the development of interactive computer systems:

1. To suggest potential biases, limitations, and cognitive abilities which could play a central role in specific human-computer interactions.
2. To investigate whether and how those cognitive constraints actually influence the interactive task being studied. Using psychological methodology one can study, for example, whether predictions derived from a given cognitive model apply to a specific interactive task.
3. To inform the development of prescriptions or guidelines for the design of specific computational aids; that is, a specification of the characteristics a system

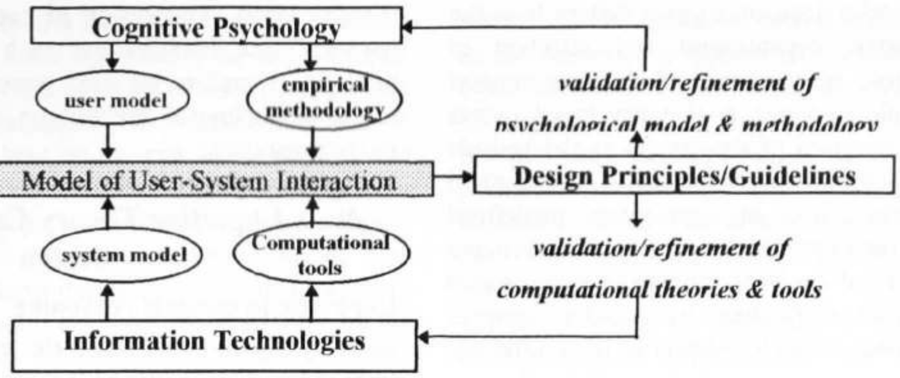

Figure 1: Symbiotic approach to cognitive engineering.

must possess to adapt to user's constraints, so that user performance is effective and pleasurable.

4. To provide tools for the evaluation of a particular computer design. Again empirical psychological methodologies can contribute to the assessment of the utility of the designs.

Most of the above roles are illustrated in Figure 1, which summarises our symbiotic approach to cognitive engineering. This paper concentrates on the top half of the figure (i.e. the one concerned with the roles of psychology in cognitive engineering). But, as the figure suggests, the symbiotic relationship we propose could possibly be mirrored in the roles of computational sciences in cognitive engineering.

In fact, cognitive psychology has been successfully applied to the design of technology in a variety of domains. For example, the design of auditory alarms in aircraft and in intensive care applications (Edworthy & Stanton, 1995); the design of bibliographic software (Egan, Remde, Gomez, Landauer, Eberhardt, & Lochbaum, 1990); the evaluation of workstation design for telephone operators (Gray, John, & Atwood, 1993); the evaluation and design of computer displays in intensive care monitoring (Green, Logie, Gilhooly, Ross, & Ronald, 1996); etc. In all the above examples, the application of a cognitive model (e.g. Card, Moran, & Newell, 1983; Wickens, 1992), coupled with a sound analysis of the application domain and extensive empirical psychological investigations, resulted in successful contributions to the design process.

Additionally the results of design-oriented task-specific psychological investigations can feed back into the cognitive theory from which they were generated. The drive to develop a particular piece of technology has often forced questions on the psychological theories which informed the development. And, as a result, these theories have been refined and enhanced. An improvement of the cognitive theory can have useful repercussions for the solution of future related design problems as the theory (as well as the methodology) has been adapted to meet specific constraints of real world domains. In other words, a potential benefit of a symbiosis between psychological and human factors practices is actually the development of more powerful tools for addressing engineering design.

The Design Problem: Computerised Decision Support in Intensive Care

The clinical monitoring of patients in the ICU has three objectives:

1. to allow confirmation that the patient is stable;
2. the early detection of physiological events, with a view to rectifying problems before they become too established;
3. to detect situations in which the patient fails to respond to treatment, thereby requiring alternative actions.

Intensive care has seen a rapid increase in the number of different parameters which can be continuously monitored: heart rate, blood pressure, blood gases, etc. Conventionally, each physiological parameter is displayed on a separate monitor and in a different format. However, physiological conditions can be indicated by changes in several of these parameters; therefore such an arrangement results in significant complications for scanning and assimilating the data displayed.

Computers and ICU Monitoring

An important development in the last decade has been the use of computerised systems to facilitate intensive care monitoring (e.g. Cunningham, Deere, Elton, & McIntosh, 1992; Green et al., 1996). Computers can collect data from the different monitors and display them in a more uniform way. Furthermore, computers have the ability to store data and keep a record of physiological trends over long periods of time, whereas conventional monitors only give a measure of a particular moment in time.

Additionally, a potential contribution of computer technologies is the development of decision support systems to assist in the interpretation of monitored data (Coiera, 1993). The major difference between a computerised monitoring system and a decision support system is in the level of interpretation, organisation and selection of available data. A role of computerised decision support would be, for example, to detect, as they developed, events which merited the attention of the person at the bedside (Salatian & Hunter, 1996). A decision support system could also suggest the selection of an appropriate, predefined display to facilitate the interpretation of data. Research and development in clinical decision support systems has a history going back for more than two decades; however only a limited number of such systems is in routine use (Shortliffe, 1993).

Problems with the Implementation of Computerised Aids at the ICU

Sophisticated computer systems often require the employment of computer support staff, and time-consuming training of staff on the ward. In addition, legal considerations and the perceived or actual reliability of a system may still require paper records to be kept. This parallel recording, coupled with training requirements results in the computer adding to, rather than easing, the workload of ward staff.

In a wide variety of health care domains, it has been shown repeatedly that computerised aids in medicine may not be readily accepted or widely used by medical or nursing staff (Green et al, 1991). Furthermore, even where staff are positive about the utility of the computerized system, there may be no impact on clinical outcome (McIntosh, Cunningham, & Elton, 1994). However, in spite of all of these difficulties, there is a widespread belief within the clinical professions that computers have significant potential in patient care provided that the system is designed and implemented appropriately (Imhoff, 1992).

Computerised monitoring systems may be designed through consultation with a panel of medical experts. However typically it is the more junior physicians and the nursing staff who are the main users of the system. Hence, a major issue is whether junior staff can effectively use the data displayed to spot trends and to take appropriate action: a) the nurses are continually by the bedside but may not see it as their role to use computerised data displays; b) junior physicians visit the bedside in turn and may not be present when a trend starts to develop or may not have the clinical experience to detect a trend inherent in the presented data pattern; c) experts, on the other hand, may be able to detect trends but they are the least frequent professional visitor to each patient.

The above analysis suggests that understanding differences in clinical expertise may be an important requirement for the design and successful application of computerised monitoring systems in the ICU. Other psychological factors which may influence the successful use of such systems are the so-called cognitive biases (Evans, 1989). The design of computerised aids do not normally take into account such limitations of human thinking. Therefore we next present a brief summary of recent psychological findings on expertise and biases in medical reasoning.

How Cognitive Theory Can Contribute to Design

Expertise in medical reasoning

Gilhooly (1990) summarised the major findings emerging from research in expertise in various complex cognitive activities and compared them with what was known at the time about expertise in medical reasoning. In general, expertise research had suggested that experts perform better than novices because they possess superior domain knowledge accumulated after many years of extensive practice, and not because of superior basic capacities. It is generally assumed that experts' superior domain knowledge is in the form of schemata which include recommendations for action. The assumption is that, because they have a richer repertoire of schemata, experts can remember better than novices new information in their field, and have better problem representations, that is, experts focus on the deep structure of the problem (i.e. they are able to abstract relevant principles) whereas novices are led by the surface features of the problem. Additionally, experts work forwards (i.e. from the starting state to the goal state), whereas novices work backwards from the unknown to the givens.

Most of the above findings have been replicated by research on medical reasoning. However, Gilhooly noted two significant discrepancies between research on medical problem solving and general research in expertise. One such discrepancy is that medical experts do not seem to show a strong tendency to work forward to a goal state. In pioneering studies of diagnostic thinking (Elstein, Shulman, & Sprafka, 1978), it was found that expert physicians generated hypotheses very early in the process, after seeing just a few signs or symptoms; these hypotheses were then tested, checking for the presence or absence of symptoms deduced from the hypotheses. This approach, which can be characterised as one of hypothetico-deductive reasoning, involves reasoning backwards from the goal (the hypothesis) to the given (the symptoms). A reasonable explanation for this type of processing is that in diagnostic thinking not all the necessary information is presented initially; hence, the task requires information search, and this search is usefully guided by hypotheses (see discussions in Gilhholy, 1990; Elstein, Shulman, & Sprafka, 1990).

A second discrepancy between research on medical diagnosis and general expertise research is that medical experts do not remember information about new cases any better than less experienced subjects (Claessen & Boshuizen, 1985). However, even if experts have poor memory for the specifics of a case, the diagnoses produced by the experts for that same case tend to be more accurate than those produced by less experienced subjects. This data pattern may be explained by the different forms of knowledge brought to bear on the task by expert and less expert subjects (Schmidt & Norman, 1988). Experts seem to use compiled knowledge in the form of "illness scripts" which contain prototypical information about diseases. This usually allows a reasonable diagnosis with little processing of the textual information. Less experienced subjects, on the other hand, may rely on general pathophysiological knowledge (not grounded on personal experience) which leads them to a slower and often less accurate processing of the information.

The findings just outlined can be of particular use if they are incorporated in some way in a computerised decision support system. One can view the function of such a system as that of making available to less experienced staff the expertise of senior clinicians. This may involve, for example, highlighting or making explicit those correlations or causal relationships (known by the experts) which may not be obvious to more junior staff; thus supporting the detection and prevention of risky pathological conditions. Additionally, by understanding the differences in domain knowledge and cognitive processes of experts and novices, one can decide, for example, on the sort of information which should be made available to junior staff as well as the sort of representation (e.g. level of abstraction) in which that information should be provided.

Heuristics and biases in medical decision making

It is well established that human decision making is prone to a range of biases (Evans, 1989). A well-known bias is the so-called confirmation bias; that is, once an individual forms an hypothesis there is a tendency to seek evidence which confirms that hypothesis rather than to gather information to refute it. Given the open-ended nature of medical diagnosis which, as noted above, is at least partly driven by hypothetico-deductive reasoning, the potential implications of confirmation bias are particularly relevant. In fact, there is evidence that suggests that physicians (either experienced or novice) do not use scientific principles based on systematic evaluation of evidence (see e.g. Patel, Evans, & Groen, 1989).

Another weakness of human cognition is the tendency to select inappropriate base rates when making probability judgements or assessing risk. For example, if a test to detect a disease whose prevalence is 1/1000 has a false-positive rate of 5%, then what is the chance that a person found to have a positive result actually has the disease? The majority of a sample of Harvard Medical School students reported that the answer to this question was 95%. But this answer fails to take into account the very low base rate of the disease, and the correct answer is 2% (Cascells, Schoenberg, & Grayboys, 1978). It is possible to increase somewhat the number of correct respondents if the figures in the scenarios are given as frequencies (e.g. 50 out of every 1000) rather than as probabilities (e.g. 5%) (cited in Gilhooly, 1996, pp. 188-189). Qualified medical practitioners are apparently no less prone to faulty probabilistic reasoning and cognitive biases than are medical students or other members of the adult population (see e.g. Christensen-Salanski & Bushyhead, 1981; Schwartz, 1994).

These findings (if confirmed in real medical practice) have important implications for the design of a computerised medical decision support system. Such a system would have severely limited utility if the system designer assumed (implicitly or explicitly) that decisions would be made by expert clinicians who systematically evaluate evidence, or if the system failed to take account of possible biases in probabilistic reasoning or in hypothesis testing. More specifically, the design of appropriate data displays may compensate for user limitations by, for example, showing rates and other numerical data in an easily interpretable manner or by highlighting patterns, or other information, which are likely to be ignored by the user.

How Solving A Design Problem Can Benefit Cognitive Theory

Research in medical reasoning has typically involved simulated patient scenarios, either using actors in the role of patients, or vignettes comprising a medical history and sets of symptoms (for a review see Gilhooly, 1990). Rarely have studies of medical decision making taken place with genuine patients or genuine patient data (Koehler, 1996). An advantage of working on an applied medical domain such as intensive care is that the physiological data are recorded and archived. This offers a rich source of stimulus material for use in experimental studies. In other words, it facilitates controlled studies of the cognitions of the staff concerned in realistic settings involving genuine patient data (i.e. computerised records of monitored physiological data); hence increasing the ecological validity of the research. Such investigations can therefore provide insights as to whether medical decision making is affected in practice by the biases and expertise effects encountered in laboratory studies. As noted earlier, a potential outcome is the refinement of the cognitive theories which inspired the investigations.

Methodologies

In order to build a system which will actually help its intended users, we need to achieve an understanding of its role by studying the decision-making environment in some detail. More specifically, we need to understand what factors (i.e. biases, expertise, availability of information, etc.) affect the decision making of medical and nursing practitioners at the intensive care ward. In addressing these issues, our approach has been to use a range of methodologies which we outline in the rest of this section (for a more detailed discussion see Logie, Hunter, McIntosh, Gilhooly, Alberdi, & Reiss, 1997).

- *Questionnaires and semi-structured interviews with staff.* These will yield a subjective view of working practices, staff attitudes and perceived expertise. The interviews also can focus on the clinicians' reports as to how they interpret data and the sources on which they rely.
- *Systematic observations and recording of daily activities in the working environment.* This provides a more objective picture of staff's working habits and performance.
- *Experimental approaches to examine the cognitions of physicians and nurses.* Our current approach is to present staff with data patterns recorded from previous real patients, and then ask them to "think aloud" during their decision making. The "think-aloud" protocols are then transcribed and analysed (Ericsson & Simon, 1984) to examine: on the one hand, clinicians' hypothesis generation and testing procedures; and, on the other hand, their ability to identify relevant physiological patterns on the computer display. This form of analysis then allows an assessment of the decision making in terms of the incidence of decision making biases, the role of expertise, and so on.

Once the factors affecting decision making have been identified, this insight can be used to inform the design of computational algorithms in computerised decision support.

Conclusions

In this paper, we have highlighted the important roles that cognitive psychology can play (and in fact has played) in the design of human-computer interactions. In contrast with pessimistic views which question the applicability of psychological theories and methods to the design of interactive computer systems (Long & Dowell, 1996), we have shown the potential mutual benefits of a symbiosis between psychological practices and the practices oriented to solving the practical problems of computer design. In particular, we have discussed in some detail how such an approach can have its application in a particular design problem which we are currently addressing, that is, the development of decision support systems in medical intensive care.

Acknowledgments

We would like to acknowledge the support of the UK Economic and Social Research Council, grant number L127251019.

References

Barnard, P & Harrison M. (1988). Integrating cognitive and system models in human-computer interaction. In A. Sutcliffe & Macaulay, L. (Eds.), *People and Computers V*. Cambridge, UK: Cambridge University Press.

Card, S. K., Moran, T. P., & Newell, A. (1983). *The psychology of human-computer interaction.* Hillsdale NJ: Lawrence Erlbaum.

Carroll, J. M. (1991). The Kittle House manifesto (Introduction). In J. M. Carroll (Ed.), *Designing Interaction*. Cambridge, MA: Cambridge University Press.

Cascells, W., Schoenberger, A. and Grayboys, T. (1978). Interpretation by physicians of clinical laboratory results. *New England Journal of Medicine, 299*, 999-1000.

Christensen-Szalanski, J.J. and Bushyhead, J.B. (1981). Physicians use of probabilistic information in a real clinical setting. *Journal of Experimental Psychology: Human Perception and Performance, 7*, 928-935.

Claessen, H. F. and Boshuizen, H. P. (1985). A. Recall of medical information by students and doctors. *Medical Education, 19*, 61-67.

Coiera, E. (1993). Intelligent monitoring and control of dynamic physiological systems (Editorial). *AI in Medicine, 5*, 1-8.

Cunningham, S., Deere, S., Elton, R. A., & McIntosh, N. (1992). Neonatal physiological trend monitoring by computer. *International Journal of Clinical Monitoring and Computing, 9*, 221-227.

Edworthy, J. & Stanton, N. (1995). A user-centred approach to the design and evaluation of auditory warning signals: 1. Methodology. *Ergonomics, 38*, 2262-2280.

Egan, D. E., Remde, J. R., Gomez, L. M., Landauer, T. K., Eberhardt, J., & Lochbaum, C. D. (1990). Formative design-evaluation of SuperBook. *ACM transactions on Information Systems, 7*, 30-57.

Elstein, A. S., Shulman, L. S., & Sprafka, S. A. (1978). *Medical Problem Solving: An analysis of Clinical Reasoning*. Cambridge, MA: Harvard University Press.

Elstein, A. S., Shulman, L. S., & Sprafka, S. A. (1990). Medical problem solving: A ten-year retrospective. *Evaluation and the Health Professions, 13*, 5-36.

Ericsson K. A. and Simon H. (1984). *Protocol Analysis: Verbal Reports as Data*. Cambridge MA: MIT Press.

Evans, J. St. B. T. (1989). *Bias in human reasoning: causes and consequences*. Hove, UK: Lawrence Erlbaum.

Gilhooly, K. J. (1990). Cognitive psychology and medical diagnosis. *Applied Cognitive Psychology, 4*, 261-272.

Gilhooly, K.J. (1996). *Thinking: directed, undirected and creative*. London, UK: Academic Press.

Gray, W. D., John, B. E., & Atwood, M. E. (1993). Project Ernestine: Validating a GOMS analysis for predicting and explaining real-world task performance. *Human-Computer Interaction, 8*, 237-309.

Green, C. A., Gilhooly, K. J., Logie, R., & Ross, D. G. (1991). Human factors and computerisation in Intensive Care Units: A review. *International Journal of Clinical Monitoring and Computing, 8*, 95-100.

Green, C. A., Logie, R. H., Gilhooly, K. J., Ross, D. G., & Ronald, A. (1996). Aberdeen polygons: computer displays of physiological profiles for intensive care. *Ergonomics, 39*, 412-428.

Imhoff, M. (1992). Acquisition of ICU data: concepts and demands. *International Journal of Clinical Monitoring and Computing, 9*, 229-237.

Koehler, J. J. (1996). The base rate fallacy reconsidered: Descriptive, normative, and methodological challenges. *Behavioral and brain sciences, 19*, 1-53.

Landauer, T. K. (1987). Relations between cognitive psychology and computer design. In J. M. Carroll (Ed.), *Interfacing Thought: Cognitive Aspects of Human-Computer Interaction*. Cambridge, MA: MIT Press.

Logie, R.H., Hunter, J., McIntosh, N., Gilhooly, K., Alberdi, E. & Reiss, J. (1997). Medical cognition and computer support in the intensive care unit: A cognitive engineering approach. In D. Harris (Ed.), Engineering Psychology and Cognitive Ergonomics. Aldershot, UK: Ashgate.

Long, J. & Dowell, J. (1989). Conceptions of the discipline of human-computer interaction: Craft, Applied Science, and Engineering. In A. Sutcliffe & Macaulay, L. (Eds.), *People and Computers V*. Cambridge, UK: Cambridge University Press.

Long, J. & Dowell, J. (1996). Cognitive engineering or 'Getting users interacting with computers to perform effective work'. *The Psychologist, 9*, 313-317.

McIntosh, N., Cunningham, S., & Elton, R. (1994). A randomised trial of computerised trend monitoring. *Pediatric Research, 36*, 28A (155).

Norman, D. A. (1986). Cognitive engineering. In D. A. Norman & S. W. Draper (Eds.), *User Centered System Design. New Perspectives on Human-Computer Interaction*. Hillsdale, NJ: Lawrence Erlbaum.

Patel, V. L., Evans, D. A., & Groen, G. J. (1989). Biomedical knowledge and clinical reasoning. In D.A. Evans & V.L. Patel (Eds.), *Cognitive Science in Medicine*. Cambridge, MA: MIT Press.

Salatian, A. & Hunter J. (1996). ASSOCIATE: An approach to the interpretation of ICU data. *Working Notes of IDAMAP-96 (Workshop on Intelligent Data Analysis in Medicine and Pharmacology), ECAI-96 (Budapest 1996)*, 73-78.

Schmidt, H. G. & Norman, G. R. (1988). *On the development of expertise in medicine.: Evidence from case-representation studies. (Technical Report No. 123).* Maastricht, The Netherlads: University of Limburg.

Schwartz, S. (1994). Heuristics and biases in medical judgment and decision making. In L. Heath (Ed.), *Applications of heuristics and biases to social issues*. New York, NY: Plenum Press.

Shortliffe, E. H. (1993). The adolescence of AI in medicine: Will the field come of age in the '90s? *AI in Medicine, 5*, 93-106.

Wickens, C. D. (1992). Engineering Psychology and Human Performance. New York, NY: Harper Collins.