COGNITIVE ENGINEERING

Less is (sometimes) more in cognitive engineering: the role of automation technology in improving patient safety

K J Vicente

Qual Saf Health Care 2003;12:291-294

There is a tendency to assume that medical error can be stamped out by automation. Technology may improve patient safety, but cognitive engineering research findings in several complex safety critical systems, including both aviation and health care, show that more is not always better. Less sophisticated technological systems can sometimes lead to better performance than more sophisticated systems. This "less is more" effect arises because safety critical systems are open systems where unanticipated events are bound to occur. In these contexts, decision support provided by a technological aid will be less than perfect because there will always be situations that the technology cannot accommodate. Designing sophisticated automation that suggests an uncertain course of action seems to encourage people to accept the imperfect advice, even though information to decide independently on a better course of action is available. It may be preferable to create more modest designs that merely provide feedback about the current state of affairs or that critique human generated solutions than to rush to automate by creating sophisticated technological systems that recommend (fallible) courses of action.

A growing number of efforts to improve patient safety are underway, far more than even just a few years ago before the influential US Institute of Medicine report was released in December 1999.¹ As practitioners and researchers delve into the details of how to reduce medical errors, it may be worthwhile to take a reflective step back and examine the role that information technology can and should play in these efforts.

In many cases there is a natural tendency to assume that human error in complex sociotechnical systems such as health care can be reduced or even eliminated by more technology, especially automation. This viewpoint is intuitively appealing because the logic behind it seems bullet proof. If we have a problem (patient safety) we should do something to resolve it, and the more sophisticated the means we adopt (more technology), the more likely we will be to attain a definitive resolution to the problem (improve safety). After all, if we only make a half hearted attempt and do less than is within our means (use less technology), then the people faced with the problem (healthcare providers) have to take on more of the

burden than they otherwise would. As a result, we run the risk of not solving the problem at hand. I will refer to this potential role of technology in improving patient safety as the "more is better" view. As the patient safety agenda moves forward, many examples of this approach can be found as bar coding systems, computerized physician order entry systems, and automated diagnosis aids are introduced into clinical practice in attempts to reduce medical error.

The "more is better" perspective does have its place and has already led to improvements in patient safety, but sometimes it seems as if this view is pushed to an extreme, leading to an uncritical rush to automate whenever and wherever possible. For example, in November 2001 the Ontario Hospital Association held its annual convention under the theme "Touching Technology". To advertise the meeting a special eight page marketing supplement was published in the *Globe & Mail*, one of Canada's largest and most influential daily newspapers. The following five sentences appeared in that supplement, spread across two different articles^{2,3}:

- "We are using the new technology to allow hospital staff more time to focus on and provide that vital human touch."
- "Today there is so much high technology in almost any operation that there is no way even the best of nurses can understand all of it."
- "Technology places incredible burdens on every nursing departments [sic]."
- "The burden is certain to grow and perhaps never lessen."
- "Technology is the hope for the future of health care."

Perhaps there are limits to the "more is better" view.

Research from cognitive engineering—an interdisciplinary approach to the analysis, design, and evaluation of complex sociotechnical systems such as health care—has addressed this issue.⁴⁻⁶ Cognitive engineers have developed theories, models, and data to better understand the ways in which technology impacts safety—for better and for worse—in a number of different sectors such as aviation, nuclear power, and health care. In this article I will review some of this research to show the limitations of the "more is better" view that seems to prevail in health care, and to propose a more nuanced approach to leveraging technology (and people) to improve patient safety.

MORE IS NOT ALWAYS BETTER

Aviation is frequently cited as a role model for improving patient safety, so consider first the

Correspondence to:
Dr K J Vicente, Department
of Mechanical & Industrial
Engineering, University of
Toronto, 5 King's College
Road, Toronto, Ontario
M5S 3G8, Canada;
vicente@mie.utoronto.ca

Accepted for publication
4 June 2003

study by Layton *et al* which investigated computer support for en route flight planning.⁷ Three different designs were compared under four scenarios in an experiment with 30 airline pilots. The first design was a “sketching only system” that helped pilots sketch proposed flight plans on an electronic map while the low level details such as fuel remaining, recommended altitudes, and estimated time of arrival were taken care of by the computer. The second design was a “route constraints and sketching system” that had an additional feature—namely, the ability for pilots to specify the high level constraints on the path they desired and then have a computer optimization algorithm find the shortest route within those constraints. The pilots could decide whether or not to adopt the route suggested by the computer. The third design was an “automatic route constraints, route constraints, and sketching system” that had all of the features of the previous design as well as an additional capability—namely, that the computer would automatically generate a suggested flight plan deviation as soon as it detected a problem with the original plan.

How well did these three designs help the professional pilots in simulated en route flight planning tasks? The predictions made from the “more is better” view are straightforward: performance should increase as a function of the sophistication of the technology in the design, with the “sketching only system” being the worst and the “automatic route constraints, route constraints, and sketching system” being the best. Indeed, the sketching only group had difficulties identifying the most economic route in several scenarios, presumably because the large solution and data spaces in this complex domain exceeded the psychological information processing limits of the pilots.

However, the other two designs were not without their problems. In two scenarios several participants merely accepted the automatically generated plan suggested by the two more sophisticated designs, failing to critically evaluate the situation. Note, however, that the automatically generated plans are not perfect. The computer system treats the forecast as reality when in fact the forecast is merely an uncertain estimation of the future. If the actual weather deviates from the forecast in an operational setting, then pilots can be faced with severe difficulties that can threaten safety because the plan they selected did not anticipate this eventuality. In contrast, participants in the sketching only system explored more options, considered the uncertainty associated with the forecast more thoroughly, and chose more conservative routes that would have been more robust to unexpected changes in weather in an operational setting.

In this set of circumstances at least, less can therefore actually be more. Participants using the more sophisticated designs tended to eyeball and accept the computer suggested plans—even though they were not forced to adopt them—whereas participants using the less sophisticated design reasoned through the problem in a more critical fashion.

This “less is more” effect has been observed in other aviation studies. Sarter and Schroeder⁸ recently conducted a simulator study of computer support for in-flight icing decisions. Three different designs were compared under 20 scenarios in an experiment with 27 commercial pilots. The baseline group had to perform icing decisions without any specific computer support by relying on kinesthetic flight cues alone, the status display group also received information about the icing situation but had to determine which course of action to take, and the command display group received a computer generated recommendation about which course of action to take to cope with the icing situation. The “more is better” view would predict that performance should be worst in the baseline group and best in the command display group.

Again, the results were more subtle than this plainly monotonic view would suggest. When the information provided by the status and command displays was accurate,

both of these conditions led to better performance than the baseline group. However, when the decision support advice was not accurate (as would occasionally be the case in a real decision aid which would be less than perfect), performance was worse than that observed in the baseline group. This result may be surprising given that the same kinesthetic information was available in all three conditions. Apparently, participants in the status and command groups made less use of the kinesthetic flight cues than the baseline group. More importantly, the performance decrement caused by the inaccurate advice was significantly greater for the command display than for the status display. In conditions where the validity of the information provided by the decision support is uncertain, providing more sophisticated technology (for example, a command display rather than just a status display) can therefore actually make things worse rather than better.

The “less is more” effect is not restricted to aviation. Smith conducted an experimental study of management decision making that compared the performance of two groups of participants, one that was given a deliberately incomplete problem representation in the form of a decision tree diagram and another that was not given any representation aid at all.⁹ The incomplete representation impaired performance because participants tended to rely on it as a comprehensive and veridical representation of the problem, failing to consider the important factors that had been deliberately omitted. In contrast, because they did not have any aid, the control participants reasoned through the problem more thoroughly and took into account some of the relevant factors that had been left out of the diagram. In principle there is no reason why the other group could not have engaged in the same kind of critical evaluation, but apparently the mere presence of the (imperfect) representation discouraged them from doing so. Thus, being provided with an incomplete problem representation can lead to worse performance than having no representation—even nothing at all is sometimes more.

WHY LESS CAN BE MORE: UNANTICIPATED EVENTS LEAD TO FALLIBLE AUTOMATION

In a closed system where everything can be anticipated more is, indeed, better. If a decision aid can identify the optimal course of action, there is no reason not to accept the advice because it is impossible to do better—by definition. However, complex safety critical systems such as aviation and health care are open systems where unanticipated events are bound to occur eventually.⁵ In these types of systems the decision support provided by a technological aid will be less than perfect because there will always be situations that the technology does not take into account. The studies reviewed above, and others like them, show that less is sometimes more in open systems. Designing sophisticated technology that suggests an uncertain course of action seems to encourage people to accept the imperfect advice, even though information to decide independently on a better course of action is available. As a result, performance can be worse than with a decision aid that only provides more modest technological support, or even than with no decision aid at all.

Is the “less is more” effect relevant to health care? There are reasons to believe that it is. In an evaluation of a computer based physician order entry system, Bates *et al*¹⁰ observed a reduction in the overall number of non-intercepted serious medication errors compared with no such computer support. At the same time, however, the number of such errors for drug problems that were not addressed by the computer database doubled. Providing a sophisticated technological aid can therefore cause physicians to refrain from engaging in the cognitive processes that they would normally use in the absence of such an aid, increasing particular types of medical errors. Even in health care, less is sometimes more.

A MORE NUANCED VIEW: THE CONSTRAINT BASED APPROACH

Is there a way to make the most of the benefits that technology has to offer while minimizing the kind of insidious effects described above? One possibility is to adopt a constraint based approach which provides people with rich feedback about the current state of affairs but does not recommend particular courses of action, instead leaving it up to people to determine what to do, given their knowledge of the local contingencies, many of which cannot be anticipated offline during design but can be observed online during operations.⁵ This approach aims to help people to adapt to unanticipated events—a role for which people are uniquely suited—while eliminating errors that are caused by blind reliance on solutions generated by (imperfect) automation. At the same time, technology could be used in a more modest but constructive critiquing mode, pointing out potential deficiencies in the decision making processes or the course of action that people are considering. This approach would also take advantage of the benefits that technology has to offer in overcoming human information processing limitations.

Guerlain *et al.*¹¹ provide an example of how this critiquing approach could be applied to health care. They developed a decision support system to help blood bankers identify alloantibodies in patient's blood. Medical technologists were in charge of the decision making process, but the critiquing system provided a rich source of feedback and notified the technologists when (a) errors of omission or commission were committed; (b) a complete protocol was not followed; (c) the answers provided were inconsistent with data collected; and (d) the answers provided were inconsistent with prior probability information. An experiment with 32 professional blood bankers evaluated performance under four test scenarios on a critiquing system and a baseline information system that provided the same displays and controls but no critiquing advice.

The critiquing system led to significantly better decision making performance overall, and in three of four cases completely eliminated misdiagnoses. Furthermore, in a scenario that was not anticipated during the design of the critiquing system, half of the control participants misdiagnosed the case whereas only three of 16 critiquing participants did so. This result was only marginally significant ($p=0.072$), but it represents a trend opposite to that noted earlier where more sophisticated decision support systems led to worse performance than baseline groups for unanticipated scenarios.

These findings show that it is possible to reconcile the need to help healthcare practitioners make better decisions with the need to overcome the limitations of a computer aid in dealing with unanticipated events. Using a constraint based approach to decision support provides computer guidance that can reduce errors due to human information processing limitations while simultaneously providing people with the freedom and flexibility that can reduce errors due to imperfections in the computer aid, thereby making the most of both technology and people.

A MULTIDIMENSIONAL FRAMEWORK FOR HUMAN-AUTOMATION DESIGN

The "more is better" view fails as a generalizable explanation of the impact of technology on safety because it oversimplifies a multidimensional problem by treating it as if it could be captured along a single dimension—more or less technology. Parasuraman *et al.*¹² provide a framework of human-automation interaction that sheds light on this issue. As shown in fig 1, the core of their framework consists of two dimensions. The first distinguishes between four types of functions that can be automated: (a) data acquisition; (b) information analysis; (c) decision selection; and (d) action

Figure 1 Multidimensional framework of human-automation interaction from Parasuraman *et al.*¹² Two examples of systems with different types of automation profiles are shown. The solid line represents a design that uses technology to automate as much as possible, whereas the dotted line represents a design that uses technology primarily for information acquisition, giving people primary responsibility for the remaining functions. Reprinted with permission. © 2000, IEEE.

implementation. The second dimension distinguishes between various levels of automation from completely manual (no automation) to completely automatic (no human intervention). The key insight of this framework is that these two dimensions—types of automation and levels of automation—are conceptually orthogonal. It is possible to choose radically different levels of automation for different functions. For example, one could design a system that has fully automated data acquisition, information analysis, and action implementation yet fully manual decision and action selection. Such a design would relieve people of having to collect data, synthesize it into information, and physically implement an action, while still giving them complete autonomy to choose a course of action. Of course, many other combinations of types and levels of automation are possible, and no one combination is ideally suited for all circumstances. In addition to the two core dimensions of their framework, Parasuraman *et al.* discuss a number of additional criteria that need to be considered when making function allocation decisions, including: the degree of mental workload imposed by the design, the degree to which the design supports situation awareness, the impact of the design on operator complacency, the degree of skill degradation induced by the design, the reliability of the automation, and the potential cost of decision/action outcomes.¹²

This conceptual framework helps unravel the apparent contradiction in the findings reviewed above—that more technology is better in some cases whereas less is more in other cases. The crucial clarification is that level of automation alone cannot be used to predict performance; it is also essential to take into account which functions are being automated. For example, a command display may have a high degree of automation in data acquisition, information analysis, and decision and action selection because it delegates all of these functions to technology and presents the human with a course of action that can then be adopted or not. In contrast, a status display may only have a high degree of automation in the data acquisition and information analysis functions, leaving it up to the human to perform the decision and action selection function. Sarter and Schroeder⁸ showed that, under conditions of irreducible uncertainty, these two automation profiles can lead to significantly different performance outcomes that have critical implications for safety, even though both designs rely extensively on automation, albeit in different functions.

This finding is not anomalous. Indeed, Parasuraman *et al.* suggest that it is best to avoid high levels of automation in the

Key messages

- Adding more technology can worsen human performance and thus threaten safety.
- In an open system with irreducible uncertainty, advice provided by automation will sometimes be inappropriate.
- People tend to follow actions recommended by automation, even when information is available to decide independently on a better course of action.
- Allowing people to make decisions and using technology to provide feedback and critique them may sometimes be preferable to using automation technology to recommend (fallible) courses of action.

decision function for systems that require human intervention because such designs have been found to lower the performance of the combined human-machine system.¹² This recommendation is particularly important in open safety critical systems such as health care because, under these circumstances, “there will always be a set of conditions under which the automation will reach an incorrect decision” and the consequences of error can be fatal.¹² The bottom line is that, in complex safety critical systems, increasing the level of technological sophistication seems to have more of a performance decrement on the decision and action selection function than on other functions.

CONCLUSION

Technology has an important role to play in improving patient safety, but the cognitive engineering research literature clearly shows that more is not always better. Creating sophisticated technological systems that recommend courses of action can lead to worse performance than more modest designs that merely provide feedback about the current state of affairs or that critique human generated solutions. In some cases a technological support system can lead to worse performance than no technology at all.

In an open system like health care, unanticipated events can and do occur. This reality must be confronted. People are uniquely capable of adapting to change and novelty and constraint based technology can be designed to help healthcare providers play this essential but challenging role.

We can still take advantage of what technology has to offer but, tempting though it is, we must avoid the rush to automate and remember the counterintuitive but amply documented finding that less is sometimes more. As patient safety researchers we should keep this lesson in mind as we work towards reducing the number of people who are injured and die annually from preventable medical error.

ACKNOWLEDGEMENTS

The writing of this paper was sponsored in part by the Jerome Clarke Hunsaker Distinguished Visiting Professorship at MIT and by a research grant from the Natural Sciences and Engineering Research Council of Canada. The author would like to thank Paul Barach, Nancy Leveson, Joachim Meyer, and the reviewers for their helpful comments.

REFERENCES

- 1 Kohn LT, Corrigan JM, Donaldson MS. *To err is human: building a safer health system*. Washington, DC: National Academy Press, 1999.
- 2 Anonymous. Technology and the human touch. *The Globe & Mail* 2 November 2001, H1.
- 3 Anonymous. Health care's high-tech revolution. *The Globe & Mail* 2 November 2001, H5.
- 4 Rasmussen J, Pejtersen AM, Goodstein LP. *Cognitive systems engineering*. New York: Wiley, 1994.
- 5 Vicente KJ. *Cognitive work analysis: toward safe, productive, and healthy computer-based work*. Mahwah, NJ: Erlbaum, 1999.
- 6 Vicente, KJ. From patients to politicians: a cognitive engineering view of patient safety. *Qual Saf in Health Care* 2002;**11**:302–4.
- 7 Layton C, Smith PJ, McCoy CE. Design of a cooperative problem-solving system for en-route flight planning: an empirical evaluation. *Human Factors* 1994;**36**:94–119.
- 8 Sarter NB, Schroeder B. Supporting decision making and action selection under time pressure and uncertainty: the case of in-flight icing. *Human Factors* 2001;**43**:573–83.
- 9 Smith GF. Representational effects on the solving of an unstructured decision problem. *IEEE Transactions on Systems, Man, and Cybernetics* 1989;**SMC-19**:1083–90.
- 10 Bates DW, Leape LL, Cullen DJ, et al. Effect of computerized physician order entry and a team intervention on prevention of serious medication errors. *JAMA* 1998;**280**:1311–6.
- 11 Guerlain SA, Smith PJ, Obradovich JH, et al. Interactive critiquing as a form of decision support: an empirical evaluation. *Human Factors* 1999;**41**:72–89.
- 12 Parasuraman R, Sheridan TB, Wickens CD. A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans* 2000;**SMC-30**:286–97.