Sheridan, T., Hancock, P.A., Pew, R., Van Cott, H., & Woods, D. (1998).
Allocating functions rationally between human and machines. Ergonomics in Design, 6 (3), 20-25.

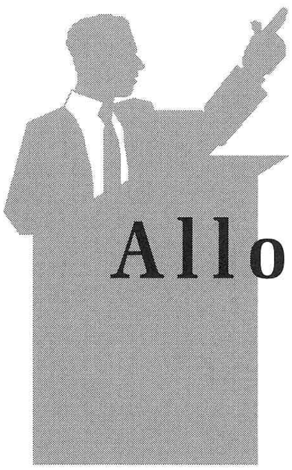

Allocating Functions Rationally

An ongoing debate about how to divide tasks results in (surprise!) lack of agreement.

THIS IS A SUMMARY OF A DEBATE that took place at the 1997 HFES Annual Meeting. The proposition was "Resolved: allocating functions between humans and machines can never be done on a rational basis." The debaters were Harold P. Van Cott and Richard W. Pew for the proposition, and David D. Woods and Peter A. Hancock against the proposition. I was organizer and chair. Harold "Smoke" Price was cochair and made invited comments at the end, as did Alphonse Chapanis, who was a coauthor of the original Fitts list report.

Shortly after the 1977 Annual Meeting, a three-day symposium on the theme "Revisiting the Allocation of Functions," sponsored by the IEA, was held in Galway, Ireland. I attended and gave a plenary paper there. The second part of this article briefly comments on that meeting and mentions a few points from my paper.

Introduction

Function allocation is generally taken to mean the assignment of required functions or tasks to resources, instruments, or agents (either people or machines). It can also mean assignment of the people or machines to the functions, which is equivalent. (*Function* and *task* are used here to mean the same thing, though some authors make subtle distinctions.). For years, human factors engineers have been seeking better ways of doing function allocation as an inherent part of system design.

Fitts and his contributors (1951) were perhaps the first to make well-publicized assertions about the proper allocation of functions between humans and machines. The report attributed to him as senior author asserted that humans are better at detecting small amounts of visual, auditory, or chemical energy; perceiving patterns of light or sound; improvising and using flexible procedures; storing information for long periods and recalling appropriate parts; reasoning inductively; and exercising judgment. Machines, on the other hand, were said to be better at responding quickly to control signals, applying great force smoothly and precisely, storing information briefly and erasing it completely, and reasoning deductively.

The so-called Fitts list has sparked much controversy over the years. Jordan (1963) criticized use of the Fitts list by people who assume that the goal is to *compare* people and machines, then decide which is best for each function or task element. He quoted Craik (1947), who had earlier pointed out that to the extent that the human is understood as a machine, we know how to replace him/her with a machine. Early studies by Birmingham and Taylor (1954) revealed that in simple manual control loops, performance can be improved by quickening, wherein visual feedback signals are biased by derivatives of those signals, thereby adding artificial anticipation and saving the human operator the trouble of performing this computation cognitively.

Birmingham and Taylor concluded that "man is best when doing least" in this case. Jordan suggested that this is a perfect example of Craik's tenet. He also quoted Einstein and Infeld (1942), who discussed the development and then the demise of the concept of *ether* in physics, and how when empirical facts do not agree with accepted concepts, it is time to throw out the concepts (but retain the empirical facts). Jordan's point was that we should throw out the idea of comparing human and machine but keep the facts about what people do best and what machines do best, and that the main point of retaining the Fitts list is that people and machines are *complimentary*.

Meister (1971) suggested a straightforward procedure for doing task allocation: Write down all the salient mixes of allocation, then write down all the applicable criteria. Following this, one should rank order all combinations of allocation mix and criteria, thus determining a rank-order score. Some authors, however, have pointed to serious difficulties with any such direct method: hidden assumptions, interdependence of tasks, interdependence of criteria, unanticipated criteria, nonlinearities that invalidate simple multiplication of weight by rating and addition of products, and, most of all, the fact that a very large

Sheridan, T., Hancock, P.A., Pew, R., Van Cott, H., & Woods, D. (1998).
Allocating functions rationally between human and machines. Ergonomics in Design, 6 (3), 20-25.

number of possible interactions between human and computer compete for consideration, not simply "human vs. computer."

Price (1985) asserted that in order to make use of the Fitts list, one needs data that are context-dependent, but these data are rarely available. Acquisition of these data is exacerbated by the fact that technology is not static, that the capabilities of machines to perform "intelligent" acts such as automation and decision support are ever improving. But, claimed Price, automation can "starve cognition" if the human is not

between Humans and Machines

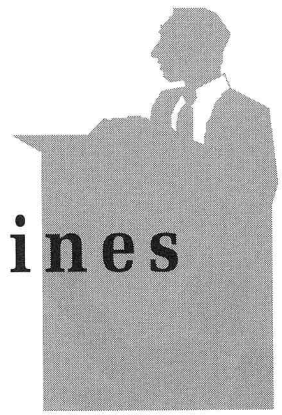

BY THOMAS B. SHERIDAN

kept in sufficient communication with what the automation is doing or intending. He seems to agree with Jordan when he points out that human performance and machine performance are not a zero-sum game, implying that the combination can be much better than either by itself. Kantowitz and Sorkin (1987) and Price (1990) provided reviews of the literature in task allocation.

The public (and, unfortunately, too many political and industrial decision makers) have been slow to realize that task (function) allocation does not necessarily mean allocation of a whole task to either human or machine, exclusive of the other. For example, in the space program, it has been common to consider that a task must be done by either an astronaut or a "robot"; that if a spacecraft is manned, then astronauts must do almost everything; and that if a spacecraft is unmanned, every task must be automated. In fact, on manned spacecraft, many functions are automatic, and on unmanned spacecraft, many functions are performed by remote manual control from the ground.

Obviously, various combinations of human and machine can function interactively and cooperatively, as when a human programs a computer and the computer then performs a task automatically. Which does the task in this case – human or computer? It is a cooperation. This is what has come to be called *supervisory control.* I have suggested qualitative scales of degrees of human and computer control.

If system goodness could be stated as a function of two variables – human and machine – and the variables were themselves scalars, then the greatest goodness could easily be found by any of a number of search techniques. When this so-called objective function is much more complex than just "human" and "machine" being two scalar variables, optimization is theoretically possible but is far more complex. When no quantitative function is available, a quantitative optimum human-machine mix is not defined. In reality, goodness itself is generally judged to be a function of many factors, including performance, efficiency of operation, safety, aesthetics, justice, and so on – factors that may not even seem commensurable.

On the other hand, if realizable system allocation alternatives can be rendered in a clear but other-than-mathematical way (say, verbal descriptions plus computer simulations) and people could interact with them, there might be ways to sample responses of these observers and make statistical predictions of what would be best for the general population. But such a method could be very expensive, particularly as the number of alternatives gets large.

Is there anything in between? Is rationality in this process somehow impossible? Are we always left to cope with an infinity of variables? Must we just "muddle through" the allocation problem in the process of doing system design?

The proposition for this debate was worded with the full appreciation that debaters on both sides would have a go at the term *rational.* The fuzz was intentional. The debaters were selected because they were known to have made significant contributions to the field and were provocative and entertaining. The audience was given the opportunity to make comments and at the end of the session to vote on the resolution based on the arguments presented.

Again, the debaters took positions on the proposition, "Resolved: Allocating functions between humans and machines can never be done on a rational basis." Summaries of the main points made by each debater appear on the following pages, using his own words as much as possible.

Sheridan, T., Hancock, P.A., Pew, R., Van Cott, H., & Woods, D. (1998).
Allocating functions rationally between human and machines. Ergonomics in Design, 6 (3), 20-25.

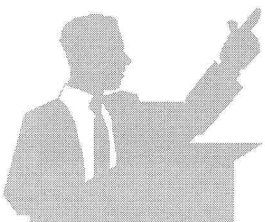

Jordan's point was that we should throw out the idea of comparing human and machine but keep the facts about what people do best and what machines do best.

For: Harold Van Cott

The *Oxford English Dictionary* defines *rational* as the use of reason and calculation for analysis and action. Reasoned argument is supported by observation, hypothesis, measurement, and calculation. It's the way science eventually arrives at the truth. But full rationality is not achievable in a complex, not fully understood world: Operating environments change, humans err, and machines fail. Therefore, function allocation in a real world can never be rational.

The original Fitts list stating what men are better at and what machines are better at has been revised from time to time as technology got better. But such lists retain the drawback that they are static and emphasize fitting the person to the machine. They ignore the nonrational characteristics that make us human and delightfully unpredictable. They promote automating what we understand and forcing humans to do the leftover tasks. And, as our colleague Hal Hendrick has noted, insufficient attention has been given to social, organizational, and other macro factors. Price (1990) emphasized that a systematic, conscious approach to function allocation has been an elusive goal of human factors specialists for years.

Probably, even if we came up with an allocation procedure that worked well, design engineers would not use it, and it would not fit into their tool kit. Human factors engineering is still too foreign, especially cognitive engineering. Design engineers are good at automating systems. Human factors engineers get what is left.

Henry Petroski (1992) has said about design, "It is quite impossible for any design to be the logical outcome of requirements. Simply because of the requirements being in conflict, their logical outcome is impossible. Design decisions are driven politically, economically, culturally as well as technologically, mathematically, rationally and logically."

Eugene Furgeson (1993) said "There is no good algorithm for the design process and there never will be. No matter how vigorously a science of design is pushed, the successful design of real things in a contingent world will always be based more on art than on science. Unquantifiable judgments and choices are the way good design comes together."

So let's keep on muddling. It's the way of the world, and it's not such a bad way. Even if rationality in function allocation is impossible, we can continue to influence the design, operation, and evaluation of systems in many useful ways.

Against: David Woods

Design has much rationality and order, but this order is not simply a matter of dividing up independent tasks between person and machine. It is a matter of understanding that design of new technology is always an intervention into an ongoing world of activity, itself a world of transformation and adaptation (Flores, Graves, Hartfield, & Winograd, 1988). New tools alter the tasks for which they were designed; indeed, they alter the situations in which the tasks occur and even the conditions that cause people to want to engage in the tasks.

Allegedly much of the equipment deployed in the Gulf War was designed to ease the burden on the operator, reduce fatigue, and simplify the tasks involved in combat. Instead these advances were used to demand more from the operator. Virtually every advance in ergonomics was exploited to ask military personnel to do more in more complex ways (Cordesman & Wagner, 1996).

Technology change transforms operational and cognitive systems by adding new roles, new needs for coordination, new kinds of errors, new paths to failure. It changes what is canonical and what is exceptional (Woods, Johanneson, Cook, & Sarter, 1994).

Prototypes and products express hypotheses about how artifacts shape cognition and collaboration. For example, new computational technologies are justified on their presumed impact on human thinking and performance. In so doing, the designers commit the error of substituting the designer's vision for empirically based and generalizable findings about the actual effects on users in real practice.

What designing is not so much about is allocation of fixed tasks to people or machines. What it is more about is predicting the dynamic process of transformation and adaptation, given proposed changes in technology or organization. It is about predicting new roles, new paths to failure, and new investments to support the new roles or to mitigate the new forms of error. If, after the fact, conditions are too difficult, too expensive, too risky or brittle, then any prototype design must be rejected.

Sheridan, T., Hancock, P.A., Pew, R., Van Cott, H., & Woods, D. (1998).
Allocating functions rationally between human and machines. Ergonomics in Design, 6 (3), 20-25.

It is worth noting the importance of understanding coordination, a "teamwork" theme common to early human factors, the supervisory control paradigm of the 1970s, cognitive systems of the 1980s, and computer-coordinated cooperative work of the 1990s.

To provide design rationality, we need experimental data and models to predict how organization and technology change cognition and collaboration. Such design will never be done by formula but, rather, will be an open-ended process of innovation informed by science.

For: Richard Pew

Design involves (a) an operational context, (b) scenarios of use, (c) a set of physical and human performance constraints, (d) a set of design goals and requirements, (e) a set of performance metrics against which the design can be evaluated to decide if the goals are met, (f) a budget, and (g) a schedule. Nothing was included about function allocation. It is wrongheaded to assert that what one does when one designs is to allocate functions. I would even go so far as to say that to make function allocation an explicit part of the design process can degrade the result.

Human-machine system design must be attacked from the perspective that humans and machines must work in concert, in harmony. Do symphonic composers start by trying to assign parts to instruments? I believe they start with efforts to create a unitary experience of sound.

Just as composing a symphony is a creative process, so too is design of a human-machine system. Constraints, both hard and soft (e.g., the system should be error forgiving), are added to those that exist until the remaining degrees of freedom are things the designer wants the user to control. If no constraints remain, the system is transparent automation. The control panel for human-imposed constraints is the user interface. The constraint imposition is an iterative exercise that is intuitive and not algorithmic. The final allocation of functions is not an explicit intention but, rather, the result of a process that focuses on achieving goals.

All designs are compromises between what is desired and what is achievable. Further, beyond all the human, physical, economic, and other constraints (even if they were all quantifiable – many are not), there are intangibles such as aesthetic factors and the need to grow or to be adaptable. In a Boston Zoo, the monkey colony is placed next to a children's park, an arrangement that creates a nice symbiosis. This arrangement did not result from some mechanical allocation algorithm.

Of course, there are some systematic and methodical aspects to design, but some very significant part of it is not rational but, rather, creative and aesthetic. Allocating functions is not a typical step the designer travels, and it is a distraction to parse the problem that way.

Referring again to composing the symphony, the detractor can readily point out that some music is computer-generated. However, the computer musician would surely maintain that the creative process comes first. Most human-machine system design poses an even more difficult problem: that of meeting a complex set of goals given an even more complex set of constraints.

Against: Peter Hancock

Coming back to the language of the proposition, according to the dictionary, rationality implies reason, and reason implies the exercise of intelligence. Intelligence is the capacity to acquire and apply knowledge. People are already engaged in acquiring and applying knowledge to function allocation and so perforce are doing it rationally. Since they are already doing it, hopefully by proper use of such practices as adaptive allocation (Hancock & Scallen, 1996), reference to the future through the term *never* is rendered illogical and without disputational merit.

Design is about handling constraint or delimitation. Specification of constraint must be specific enough to allow physical construction of a thing, a process, or a system. Engineers are more comfortable closer to certainty, and they naturally seek quantification of constraints. Nevertheless, we must admit to the inevitability of underspecification of constraint.

The most valuable lesson I learned in preparing for this debate was that indeterminacy is not necessarily a bad thing. Underspecification actually provides the degrees of freedom to allow for progress and the evolution of systems. No one wishes to live in a totally predetermined world, but neither

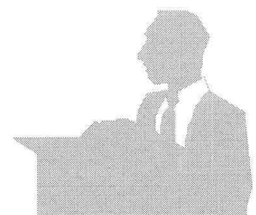

Probably, even if we came up with an allocation procedure that worked well, design engineers would not use it, and it would not fit into their tool kit.

Sheridan, T., Hancock, P.A., Pew, R., Van Cott, H., & Woods, D. (1998).
Allocating functions rationally between human and machines. Ergonomics in Design, 6 (3), 20-25.

do we wish to live in chaos. Life and, by extension, our technically supported existence survive at a balance of determinacy and indeterminacy. Our human-machine systems evolve with emerging properties beyond the conception of their designers.

Ultimately the allocation question comes down to our view of ourselves. At one end of the spectrum is the unrepentant materialist, for whom the human is nothing more than a biological machine, and for whom Licklider's dream of human-machine symbiosis (Licklider, 1960) is a rational goal. At the other is the unapologetic deist, for whom living is partaking of the divine spark and understanding of human symbiosis of any kind is beyond the pale. The more accepted view is that of the materialist, where human-machine and human-environment links are not all unique (Flach, Hancock, Caird, & Vicente, 1995). Some generalities do emerge, which we can rationalize.

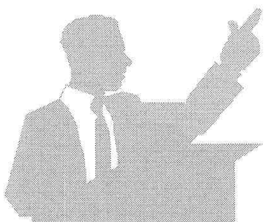

Even if rationality in function allocation is impossible, we can continue to influence the design, operation, and evaluation of systems in many useful ways.

It certainly would be a useful exercise for each of us in the ivory tower to take on some practical allocation problem, and for those of us who survive as practitioners to seek the theoretical foundations that underlie that practice. The real victory will come when function allocation is developed to the point that it forms a major weapon in our collective armory.

Audience Vote and Conclusion

Cochair Smoke Price congratulated the debaters and reiterated the importance of context and of dynamic changes, ideas for which he is well known (see the introduction). Alphonse Chapanis gently chided the chair for organizing a debate on such a loaded proposition. The audience was ready with a number of good comments, both pro and con. Space allows mention of only two.

Barrett Caldwell urged acceptance of Herbert Simon's observations that human rationality is bounded, and the audience seemed to agree. Raja Parasuraman brought up the "never say never" paradox and the troubles it leads to. He also suggested that we at least acknowledge the important human-machine system design tenet that is at the basis of the Hippocratic oath: *primum non nocere* (first do no harm).

At the end the audience was asked to vote on whether they agreed or opposed the resolution (or to choose which side was most convincing). I recall that most people voted to oppose the resolution, signifying there is some hope for rationality.

Did the outcome of the vote matter? Everyone was made to think, and they all seemed to have a good time. Was there consensus on where we are in function allocation and what needs to be done to enhance the rational process? Surely not!

Reflections on a Conference

Quite by chance, only a few days after the HFES meeting ended, a symposium called ALLFN '97 on the theme "Revisiting Function Allocation" was held October 1–3, 1997, in beautiful Galway, western Ireland. It was sponsored by the International Ergonomics Association and organized by Professor Enda Fallon of the University of Galway. Emboldened by the wisdom I had just gleaned at HFES in Albuquerque, I attended the conference and gave a plenary paper.

The conference offered a variety of both practical and philosophical perspectives to which no brief summary could do justice. As with the HFES debate, there were arguments on all sides: (a) You cannot do function allocation explicitly and objectively. (b) Of course you can, and do, every time you design a human-machine system. There were some prescriptions on "how to do it" in specialized areas. There were many presentations on "here's how we did it in our application." The poor individual commissioned to sum up at the end could only attest to the number of great ideas that had been offered but could find no consensus. Surprised?

My plenary paper, "Function Allocation: Algorithm, Alchemy, or Apostasy?" offered thoughts: Automation has moved progressively from open-loop mechanization of the Industrial Revolution, to closed loop control, to crisp or fuzzy rule-based decision, to neural nets and genetic algorithms and other mechanisms that truly learn. We warm to the idea of "human-centered automation," but when the alternative meanings of the phrase are examined, we find the substance thin and the real potential questionable. As humans become supervisors and as computers become mediators or intelligent agents, we realize that teaching some tasks to the machine isn't worth the trouble, and other tasks, those we can't define easily, we cannot program machines to do. We also see that goal-setting remains a secure function for

Sheridan, T., Hancock, P.A., Pew, R., Van Cott, H., & Woods, D. (1998).
Allocating functions rationally between human and machines. Ergonomics in Design, 6 (3), 20-25.

humans (there is some prospect for the human role!). Furthermore, satisfying our needs as humans is rarely a machine-doable function at all. Human productivity and machine productivity may not even be positively correlated. We don't want machines in the most precious parts of our lives!

However, throughout history, *homo faber*, man the tool maker, the designer of technology, has always used the human as the model and the challenge. As we understand what can be done with technology to mimic human function, we naturally embody it in machines, for the machine is simply the way *homo faber* renders his understanding. This is our history, and designing intelligent systems for today's complex needs is no exception. The challenge and therefore the natural tendency is to embody in technology whatever can be understood. Physics and engineering analysis seek to specify the constraints, and designers synthesize around those constraints, motivated by some goal or objective function. Setting the objective function remains the sine qua non role for humans. Ultimately all the rest, if the problem is clearly stated, can be mechanized.

We have always celebrated our technology as a spiritual endeavor. Earlier it was cathedral building; now it is space probes, artificial intelligence, and genetic engineering. We will continue to make grievous errors in applying those machines carelessly and inappropriately, failing, and then revising, according to our momentary criteria of success. That is the Darwinian dictate, the evident reality. We may try to control Darwin – that is, rationalize our human-machine allocation by our qualitative axioms and scientific algorithms – but in the end we will have to go with the flow. We can try to moderate evolution, and in fact evolution seems to prefer a pace that sometimes is too slow for our liking. The human-machine mix will continue to evolve, partly by chance and partly to conform to our fancies of the moment. The fancies themselves will come from us, as long as we have the upper hand, and that is where we should focus.

References

Birmingham, H. P., & Taylor, F. V. (1954). A design philosophy for man-machine control systems. *Proc. I.R.E.*, *42*, 1748–1758.

Cordesman, A. H., & Wagner, A. R. (1996). *The lessons of modern war, Vol. 4: The Gulf War*. Boulder, CO: Westview Press.

Craik, K. J. W. (1947). Theory of the human operator in control systems: 1. The operator as an engineering system. *British Journal of Psychology*, *38*, 56–61.

Einstein, A., & Infeld, L. (1942). *The evolution of physics*. New York: Simon and Schuster.

Fitts, P. M. (Ed.). (1951). *Human engineering for an effective air navigation and traffic control system*. Washington, DC: National Research Council.

Flach, J., Hancock, P. A., Caird, J. K., & Vicente, K. (Eds.). (1995). *Global perspectives in the ecology of human-machine systems*. Mahwah, NJ: Erlbaum.

Flores, F., Graves, M., Hartfield, B., & Winograd, T. (1988). Computer systems and the design of organizational interaction. *ACM Transactions on Office Information Systems*, *6*, 153–172.

Furgeson, E. S. (1993). *Engineering and the mind's eye*. Cambridge: MIT Press.

Hancock, P. A., & Scallen, S. F. (1996). The future of function allocation. *Ergonomics in Design*, *4*(4), 24–29.

Jordan, N. (1963). Allocation of functions between man and machines in automated systems. *Journal of Applied Psychology*, *47*, 161–165.

Kantowitz, B., & Sorkin, R. (1987). Allocation of functions. In G. Salvendy (Ed.), *Handbook of human factors* (pp. 365–369). New York: Wiley.

Licklider, J. C. R. (1960). Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics*, *HFE-1*(1), 4–11.

Meister, D. (1971). *Human factors: Theory and practice*. New York: Wiley.

Petroski, H. (1992). *The evolution of useful things*. New York: Knopf.

Price, H. E. (1985). The allocation of functions in systems. *Human Factors*, *27*, 33–45.

Price, H. E. (1990). Conceptual system design and the human role. In H. Booher (Ed.), *Manprint*. New York: Van Nostrand.

Woods, D. D., Johanneson, L. J., Cook, R. I., & Sarter, N. B. (1994). *Behind human error: Cognitive systems, computers and hindsight*. Dayton, OH: Crew Systems Ergonomics Information Analysis Center.

Thomas B. Sheridan is director of the Human-Machine System Laboratory, Massachusetts Institute of Technology, 3-346, Cambridge, MA 02139. Harold P. Van Cott is with Van Cott and Associates in Bethesda, Maryland. David D. Woods is at Ohio State University. Richard W. Pew is with BBN Technologies in Cambridge, Massachusetts. Peter A. Hancock is at Liberty Mutual Insurance, Hopkinton, MA.