3. *Forty-five years of man-machine systems: prospects for advanced robotics*

by *Thomas Sheridan*[*]

Abstract

The origins of development of intellectual discipline for interfacing human and machine are reviewed in terms of three fifteen-year phases from 1940 to 1985.
The first phase concentrated on empirical studies and design of the physical human-machine interface, i.e., displays and controls.
The second phase focused on the transformation of systems engineering models to characterize the entire closed loop communication, decision and control system containing the human operator.
The third and most recent phase emphasized the application of computers to aiding and implementing operator decision. The currently available knowledge combined with the technologies of advanced computers, sensors, robot effectors and the techniques of artificial intelligence and control is now producing a new phase of telerobotics which portends fundamental change in the way people work.
The emergence of telerobotics has newly accentuated four classical dilemmas: (1) determinism vs. free will; (2) reliability vs. creativity; (3) utilization of new technology vs. prevention of worker alienation; and (4) objectivity vs. advocacy. The relevance of each to telerobotics and human work is discussed. Society is now at the point of having to decide what mix of human and machine is best to produce desired goods and services and to satisfy the aspirations of workers and organizations.

1. Three phases in the development of man-machine systems as an intellectual discipline

Phase A (1940-1955): acuity, anthropometry and activity analysis: the emergence of human factors or ergonomics as a discipline

Assigning names and dates to intellectual developments is always precarious. Different reviewers are wont to offer different taxonomies and interpretations. Any one reviewer is unavoidably biased by his own

[*] Full professor of Engineering and Applied Psychology, Massachusetts Institute of Technology.

Studi organizzativi - Special Issue 2020 - Issn 0391-8769, Issn-e 1972-4969

DOI: 10.3280/SO2020-001-S1005

experience and cultural perspective. On that premise, I will claim that the Second World War, a time of terrible tragedy, was also a time of awakening for engineers and managers regarding the role of physiological and psychological factors in the design of machines and workplaces.

Of course, European university psychologists and physiologists had already begun a fine tradition of scientific experimentation and publication, though it was rather disconnected from technology. Inventors like Thomas Edison, Alexander Bell and the Wright brothers had already impacted society with their inventions. And Frederick Taylor (1911) had already imposed his “scientific management” on the production line, using his principles of man-machine analysis to justify workplace designs which increased efficiency (but often at the cost of worker morale).

Thus technology was already pervasive and most people were already aware that when technology was introduced into the workplace, not all effects of that technology were positive. Yet it took a world war to make the awareness truly widespread.

Machines of war were being produced in record numbers, and people were operating them in life-critical situations. No other circumstances could have made it more evident that the operation of machines depends critically on the human interface. Many instances occurred where aircraft cockpits, gun aiming systems, and radio communication systems were not designed to fit their operators. Design engineers often had ignored critical factors of whether displays could be seen and read (inadequate visual acuity). They often designed operating spaces that large operators could not fit into, or provided controls that small operators could not reach (inadequate anthropometry, statistical measurement of the human body). They often neglected to perform proper activity analysis, to understand exactly what the operator had to sense and decide and do under what contingencies and with what resources.

Government laboratories were put to work assembling as much relevant data as were available, and running experiments to make up for what were not available. University psychologists and physiologists and medical doctors were drawn into this new interdisciplinary effort, which was sometimes dubbed “knobs and dials engineering”. By the end of the war, many of the design errors had been corrected. But for us now what is of the most interest is that a new appreciation of human factors in the design of machines and large scale technological systems was established.

Many of the government laboratories were continued after the war. Many of the scientists involved went back to their universities and companies and set up laboratories there. In Europe, which had the immense task of rebuilding its industrial base, the factories and production machines were set up with a new appreciation for human factors. The results of many experiments were assembled into “handbooks” of human engineering data for design engineers to be used by industry (Van Cott and Kinkade, 1972). Professional societies such as the Human Factors Society in the United States and the Ergonomics Research Society (founded in the U.K.) were set up. Various professional journals emerged. Professional engineering societies set up technical committees and working groups. Some human factors leaders of this phase were Fitts, Grether and Chapanis in the U.S. and Bartlett in the U.K.

The commercial aircraft industry was probably the first to take human factors seriously. There was little tolerance for inadequate human factors analysis of piloting tasks and the pilot-aircraft interface, leading to pilot error. The automobile industry followed, but here the dictates of styling and marketing led to numerous compromises. What was best and safest from a human factors viewpoint was not necessarily cheapest and most appealing to the customer; chrome strips and useless paraphernalia sometimes won out. Other consumer product manufacturers also began to human factor their products, but faced the same compromise between what was a best performing human interface and what would make the most money. Military organizations continued to institutionalize human factors.

Though the intellectual innovation may have occurred during that first fifteen year period, the Phase A or empirical human factors kinds of efforts continue to this day, reforming new industries (e.g., following the Three Mile Island accident, the nuclear power industry suddenly “discovered” human factors), generating new and refined data for handbooks and the general scientific literature, and gradually making products and systems safer and more efficient.

Phase B (1955-1970): Models Borrowed from the Control, Communication and Decision Technologies

Since before the Second World War, exciting new systems engineering theories had been developing in communication and control. The war saw a rapid application of these theories to automation of the production line, to operations research and decision sciences for planning and management, and to radar, sonar, fire contro1, and other technologies of weapons systems. Tustin (1944) in the U.K. and James, Nichols and Phillips (1947) in the U.S. published the first works modeling the human operator in a control system (of tanks and aircraft) using the same kinds of equations as had been applied to machines. Wiener (1948) wrote his landmark

Cybernetics (from the Greek for “steersman”), which he defined as «communication and control in the animal and the machine». Soon afterward a great many others saw control theory as a nice way to describe from experiments and predict for design purposes what a human operator of a control system does. Applications went all the way from eye movements and postural reflexes in balancing one’s own body or external objects such as broomsticks, to flying aircraft and helicopters, driving cars and steering ships, to controlling large-scale chemical plants and other production systems.

During this period, mathematical models for this one class of human behavior became quite sophisticated, to the point of fourth-order differential difference equations which could predict human response with 95% accuracy. An important insight of McRuer and his colleagues (1965) was that the human operator generally adapts to whatever is being controlled so as to make the combination of human-plus-machine relatively invariant--and hence the entity to be modeled rather than the human operator only. During the latter part of this period, optimal control theory became popular for modeling the human operator.

Shannon (1949) published his mathematical theory of communication -- a way of putting all communication through a “channel” into a common denominator based on reducing uncertainty of the message receiver, and apart from any “meaning” of the message. Psychologists almost immediately saw this as a way to characterize human stimulus-response behavior, where the “channel” was now everything that happened between stimulus and response. For a wide variety of human-machine interactions, from industrial tasks to piano playing, it became fashionable to compute the number of bits of information in stimulus and response and the bits transmitted from one to the other. Miller (1956) published his well-known paper The magical number seven, plus or minus two, some limits on our capacity for processing information, in which he characterized many aspects of human memory by Shannon’s theory.

During this same period, Green and Swets (1966) showed how the theory of signal detection, developed during World War II to characterize radar and sonar could also be used to model human behavior in detecting signals in noise by hearing, vision or other senses. Edwards (1962) and others showed how normative Bayesian probability algorithms could be used as a yardstick for human decision-making. Game theory and other “borrowed systems engineering” ideas proved similarly useful.

In short, during this fifteen-year period a variety of systems engineering theories were shown to be useful for representing human capabilities in interacting with machines (Sheridan and Ferrell, 1974). Man-machine systems thus became quantitative, at least insofar as the tasks to be done were well-defined and therefore tractable to mathematical description. Though the intellectual ferment in converting engineering systems theory for modeling sensory-motor human performance is no longer so great (such modeling is now generally accepted) such efforts continue to this day.

Phase C (1970- 1985): computers and cognition

During the recent fifteen years, it is clear that the single factor that has changed the science and art of man-machine systems more than any other is the computer. Computers, of course, existed long before, first in analog form, subsequently in digital form.

Computer-based flight simulators had long been used for training pilots and "crunching numbers" in human-operated systems, but not until recently has the computer come to be regarded as a way of thinking about human behavior and human-machine systems. Namely, the recent surge of interest in artificial intelligence, with its emphasis on formalizing thinking, planning, pattern recognition and language understanding, has provided new emphasis for tackling the heretofore ill-defined aspects of operator behavior.

Cognitive science was born (see Aitkenhead and Slack, 1985, for a review).

Although the systems engineering models were useful for sensory-motor skills such as airplane piloting and narrowly-defined communication and decision tasks, they seemed relatively helpless in coping with a variety of human-computer interaction skills such as word processing or use of spread sheets, computer-aided design and maintenance, medical diagnosis or investment portfolio management. New computer-based question-answering, advice-giving expert systems are beginning to prove their worth in such applications.

Not very many years ago respectable psychologists would have nothing to do with "thinking" as a topic of science, for the reason that the sine qua non of the scientific method, direct and repeatable objective measures, were seen to be unavailable (Skinner, 1945).

Only observable stimuli and responses were admissible to scientific study of behavior (behaviorism). Now when verbalization by experimental subjects is correlated with the behavior of computational models that is accepted as sufficient model validation. Previously accepted criteria such as the number of free parameters somehow no longer seem as relevant.

Rich new heuristic computational procedures have replaced simple experimental paradigms. The line between metaphor and proof seems less clear. Stimulus- response behaviorism is asserted to be old-fashioned, rigid, and dead.

2. The development of telerobotics, a new man-machine relationship (1985-)

One might assert that Phase A provided the basis for designing a human-friendly interface between man and machine. Phase B allowed the design of a human-friendly system, one that is stable and matched to the user in terms of relatively low-level information considerations. Phase C opened the door to adding somewhat higher-level intelligence to the system which serves in a subordinate role to the human operator or user.

The cumulative knowledge from these three periods for designing machines to augment human sensory-cognitive-motor skills has offered new capabilities, which I shall generalize as telerobotics. Telerobotics is the science and art of building and programming devices which have prescribed sensing, mobility, manipulation and intelligence capability to perform rudimentary tasks autonomously (in short, robots) while remaining in continual communication with a human supervisor who is located elsewhere, providing useful information to, and accepting instructions in high-level language from that supervisor. That means the human supervisor specifies goals and constraints, typically in the form of "if (sensed pattern) then (action), else---" statements ("production rules"), and then the telerobot executes the task using its own sensors, memory, decision criteria provided, and actuators. The human supervisor plays the same role relative to the telerobot, as does the supervisor of a human subordinate in an organization. Another term for telerobotics is supervisory control (Sheridan, 1984).

The term telerobot originated with the U.S. National Aeronautics and Space Administration (NASA), though it is much more general than space robotics. For three decades, relatively unintelligent telerobots (usually called teleoperators) have been under development not only for space but also for undersea, nuclear power, construction, mining and other hazardous environments too dangerous for people but in which people nevertheless have to do work. In essence they provided video extensions for the operator's eyes and electromechanical extensions for the operator's legs and arms. Gradually, telerobots have embodied more sophisticated sensors, have become more mobile and dexterous, have incorporated more intelligence, and have provided the human supervisor more extensive aid in both situation assessment and command. New applications for telerobots include warehousing, aids for the elderly and handicapped, building cleaning and maintenance, fire fighting, police surveillance, and of course military operations of many kinds.

The industrial robot has been considered an entirely different problem from that of human operated remote control through teleoperation in hazardous environments. The industrial robot developed from the need for a more flexible numerically controlled transfer machine on the production line, a device that could be programmed to perform well-defined tasks but for a much smaller production run or batch size than would justify fixed-transfer machinery. The industrial robot has certainly proven itself for tasks such as paint spraying and spot welding where high positioning accuracy and force sensitivity are not required and where medium size runs are. However, in recent years it has become evident that for small runs teaching an industrial robot each new task is a costly bottleneck, a significant traction of the production cost. Furthermore, a human operator may have to stand by just in case something goes wrong. It has become evident that the needs for industrial robots are looking more and more like those for telerobots, and the human operator is becoming more and more recognized as an integral companion for both--though in both cases playing a supervisory role. The development of telerobotics follows six stages:

1. the human operator as a direct "hands on" controller of the physical task (the "controlled process");
2. the operator controls indirectly through intermediary displays and controls;
3. computer elements are imposed between the display interface and the controlled process to provide better information presentation, and between the controls and the controlled process to provide automatic open-loop execution of the task;
4. the computer has the capability to provide feedback to queries and simulation exercises posed by the operator, thus becoming a decision aid. It also has the capability to dose the loop through artificial sensors and actuators on the controlled process, thus making the latter a "robot" in the conventional sense, at least for short periods;
5. the two computer functions of step (4) are divided by an arbitrary distance and a communication channel. Tue first of these may be called a ''human interactive" computer (designed to assist the human to plan, teach the plan in the form of a program, monitor the program's

automatic execution, detect and diagnose failures, take over control when necessary, and learn from experience. The second may be called a “task interactive computer’’ designed to perform with speed and reliability at the remote location;
6. multiple processes, together with their task interactive computers (with appropriate sensors and effectors attached) are controlled in supervisory fashion by a single person.

Thus we see the telerobot as a generic new form of technology, that is, a versatile electromechanical “slave” to an individual person charged with being its supervisor, capable of performing one or multiple physical sensing or manipulation tasks in semi-autonomous fashion at locations arbitrarily distant from the human. Telerobots are useful for working on repetitive tasks as well as in remote and/or hazardous environments.

3. Four old dilemmas which are newly accentuated

Telerobotics is not only a new technology and a new stage of the industrialization of work. It is also a new paradigm for thinking about human-machine relationships. As such, it accentuates several fundamental dilemmas for people in relation to how they consider technology.

Dilemma 1 – Determinism versus free will, closed versus open criteria: how to judge what’s good

At an earlier time, human behavior was analyzed as a deterministic stimulus-response relation, with little role for free will in anything the Behavioristic Science of the time could accommodate. Now, following the impact of the computer and cognitive Science, thinking (and mind, and apparently free will) seem to be accommodated. Perhaps this is out of necessity, because the programmed stimulus-response activity bas been taken over by the telerobot sensors, effectors and computer.

Ultimately, the human supervisor is left with only the tasks of planning, setting goals for and teaching the telerobot.

At an earlier time, man-machine system analysts were happy to assume system goals as given, and performance was measured in terms of these goals. Analysis of the new man-machine system logically begs the question of what the human really does when he plans, sets goals and teaches.

The AI community is already trying to encroach on the human’s planning role. Perhaps the single final role is deciding on value – what is good and what is not. No one that I know asserts that a computer can do that!

Engineers and managers like determinism, and assumptions of determinism seem to have served them well in setting up experiments, modeling, predicting and designing systems, including those incorporating human operators. Engineers are uncomfortable with subjectivity and fuzziness.

Nevertheless, both subjectivity (in the form of Bayesian probability) and fuzziness (in the form of fuzzy set theory) are asserting themselves into engineering.

The idea of utility, as developed by Pareto, Von Neumann and other mathematical economists, has provided a way into valuation: deciding what is good and how to compromise among the multi-objectives, the various components of good (see Keeney and Raiffa, 1976, fora general exposition). Unhappily, recent research by decision psychologists has cast doubt on the credibility of classical utility theory as a way to do this. Yet providing high-tech decision aids and expert systems requires the computer to have some utility function, some basis for determining which alternative option (each of which may be represented mathematically by a point in multi-objective space) is best.

All of these factors attest to the "opening" of man-machine systems in terms of goals and criteria by which they may be evaluated. That is, the function of the human element is to provide the goals, the norms, the criteria of goodness. However, having open criteria makes it difficult to engineer expert systems and decision aids, and makes the new man-machine systems essentially impossible to analyze and evaluate, for it is difficult or impossible to elicit from the human supervisor at any level exactly what these are. It is always easier to assume determinism.

Therein lies the first dilemma. Philosophers, psychologists, and economists have long appreciated this dilemma. The refined role of the human component in the new telerobotic man-machine system has simply accentuated the problem.

Dilemma 2 – Reliability versus creativity: how to consider human error

The second dilemma is related, but has to do primarily with what position the engineer-manager takes with regard to human error.

Today there is much discussion about human error. It is blamed for automobile accidents, industrial accidents, and home accidents. It is blamed for nuclear power plant accidents at Three Mile Island and Chernobyl, fora chemical plant accident at Bhopal and for the shooting down of a civilian aircraft over the Persian Gulf. It is also blamed for poor decisions in engineering design, management, politics and personal relations, which cause waste, loose money and produce human suffering.

The public expectation is that the decisions of medical, engineering and management professionals be error-free. In the United States, liability litigation against professionals is at an all-time high.

Manufacturing and service companies, in addition to carrying large liability insurance coverage must go to great lengths to not to deviate from accepted standards. If they could, the lawmakers would legislate against human error. At the same time, there is a cry for greater creativity in both technical and people-related fields, for old solutions are obviously not serving many of our new problems. It is well known that creativity demands experimentation and exploration and variations in behavior--some would say “trial and error’’. It would be nice if such experimentation could all be conducted in the protected confines of the laboratory, but unfortunately that is seldom possible. Plant operators, managers and others repeatedly face new challenges in-situ, within and during their normal operating activities. They must be creative, do some “trial and error” in real-time, on-line.

Many who have analyzed industrial accidents have concluded that had the operators been allowed greater freedom (i.e., had training, procedures and technical systems provided some tolerance for small errors) the people involved would have discovered for themselves what was going wrong and what actions to take to prevent the large and ultimately costly errors. Creativity, and indeed what the human operator is best at compared to the machine, is often inhibited by tight and inflexible organization; it requires some tolerance for variation. Thus we have a second dilemma, not a new one but one spotlighted by the emergence of telerobotic systems. We have a clear difference between those who would search out and eradicate all human error by subjecting people in organizations to the same probabilistic risk analysis as they use for machine elements, and those who would accentuate the potential for human creativity by allowing some tolerance for variation from the conventional norms.

Dilemma 3 – Utilization of new technology versus prevention of worker alienation: how to retain human responsibility

Telerobotic developments clearly offer great advantages of safety for work in hazardous environments, extension of the dynamic range of the operator in spatial movement, exertion of and force and speed, and precision of movement. At the same time, their use in the workplace is likely to be alienating, where alienation takes on several meanings:

a. the operator is removed physically from hands-on contact with the product;
b. the operator performs different functions than he may have performed at an earlier time when he was in physical contact--which may lead to “deskilling’’ and corresponding loss of identity;
c. the operator may not understand what the computerized telerobot is really “thinking” (planning to do next, and therefore be mystified and anxious;
d. the operator may come to feel dependent upon and therefore inferior to the machine;
e. the operator may begin to feel that the machine is enslaving him rather than the other way around;
f. the operator may abandon responsibility for what the telerobot does.

One surely wishes to put modem telerobotics to use wherever it makes sense functionally and economically, but must at the same time be sensitive to the grim potential for alienation. Employing more skilled, better-educated workers may be one way of avoiding such alienation, but that may not be economically, legally or ethically feasible. Operators are best served by education and by participating in the decisions to introduce such machines in the workplace and in the ongoing decisions of how they will be used.

Dilemma 4 – Objectivity versus advocacy: how to serve the human interest

Within the International Federation of Automatic Control, there exist two technical subdivisions, Man-Machine Systems, and Social Effects of Automation.

The first seeks to develop and apply the mathematical theories of control and related systems engineering disciplines by modeling the human element as another objective, deterministic element, with quantitative specification of uncertainty. The second seeks to enhance worker satisfaction and the quality of work life in factory and other industrial settings. The first group promotes objectivity about human-machine relationships; the second promotes the human in such systems by advocacy.

The two groups have had rather interesting conversations over the past few years, agreeing that both viewpoints are essential, but acknowledging that the synthesis of the two approaches is very difficult, like the yin and yang of Chinese cosmology.

Man-machine technologists elsewhere are admitting to their own uncertainties about where the continuing evolution of telerobotics is likely to lead, and whether a perspective that limits itself to objectivity is sufficient. Improved productivity is a noble criterion--so long as the word means both the objective productivity of more product per unit of resources (plant equipment, manpower and money) and workers' subjective feelings of being productive in the sense of personal fulfillment.

4. How far to let the computer-telerobot go?

Telerobotics is coming on fast, including flexible robot machinery, computer control, and the capability for human supervisors to issue commands to and monitor the actions of one or multiple mechanical slaves from an arbitrary distance.

It is sometimes instructive to consider a scale of successive degrees of automation, starting at pure human control and ending at pure robotic control:

1. human worker does the whole task;
2. computer makes alternative suggestions for action, and proposes one action as the best, and turns over that alternative for execution by the telerobot if and when the human approves,
3. or the telerobot executes the action at a particular time if the human does not disapprove,
4. or the telerobot decides on and executes the action and necessarily reports it to the human,
5. or the telerobot decides, executes and reports to the human if it decides he should be told,
6. or the telerobot decides, executes and reports the action and ignores the human.

The questions is: how much to trust the telerobot, i.e., how far to let it go. We are now at the point of having to decide what mix of human and machine is best to produce desired goods and services and to satisfy the aspirations of workers and organizations.

References

Aitkenhead A.M., Slack J.M. (1985). *Issues in Cognitive Modeling.* London: The Open University.
Edwards W. (1962). "Behavioral Decision Theory". *Annual Review of Psychology,* 12: 473-498.
Green D.M., Swets J.A., (1966). *Signal Detection Theory and Psychophysics,* New York: Wiley.
James H.M., Nichols N.B., Phillips R.S. (1947). *Theory of Servomechanisms*. New York: McGraw Hill.
McRuer D.T., Graham D., Krendel E., Reisener W. (1965). *Human Pilot Dynamics in Compensatory Systems: Theory, Models and Experiments with Controlled Element and Forcing Function Variations.* AFFDL-TR-65-15, Dayton, OH, Wright Air Dev. Center.
Miller G.A. (1956). "The Magical Number Seven Plus or Minus Two: Some Limits on Our Capacity for Processing Information". *Psychological Review*, 63: 81-97.
Keeney R.L., Raiffa H. (1976). *Decisions with Multiple Objectives*. New York: Wiley.
Shannon C.E. (1949). "Communication in The Presence of Noise", Proc. IRE, 37: 10-22.
Sheridan T.B., Ferrell W.R. (1974). *Man-Machine Systems*. Cambridge (MA): MIT Press.
Sheridan T.B. (1984). "Supervisory Control of Remote Manipulators, Vehicles and Dynamic Processes: Experiments in Command and Display Aiding". In: Rouse W. (ed.). *Advances in Man-Machine Systems Research,* 1: 49-137. New York: JAI Press.
Skinner B.F. (1945). "The Operational Analysis of Psychological Terms", *Psychological Review*, 52(5), 270–277.
Taylor F.W. (1911). *The Principles of Scientific Management*. New York: Harper and Brothers Publishers.
Tustin A. (1944) "An Investigation of the Operator's Response in Manual Control of a Power Driven Gun". C.S. Memorandum, n. 169. Sheffield (U.K.): Metropolitan Vickers Co.
Van Cott H.P., Kinkade R.G. (eds.) (1972). *Human Engineering Guide to Equipment Design*. Washington: DC, U.S. Govt. Printing Office.
Wiener N. (1948). *Cybernetics*. New York: Wiley.