J. Webster (ed.), *Wiley Encyclopedia of Electrical and Electronics Engineering*

HUMAN MACHINE SYSTEMS

The phrase *human-machine system* refers to a technological system operated by or otherwise interacting with one or more human beings. Interpreted broadly, this includes practically all human interactions with the physical environment. Interpreted narrowly, it usually is used to refer to humans operating vehicles such as aircraft, automobiles, trains, or ships; humans operating computers in home or business settings, and humans operating industrial machinery or other complex equipment in factories, stores, hospitals, homes, etc.

More than anything, the term refers to a system point of view regarding human interaction with the physical environment, wherein specific variables of information or energy can be identified in a two-way, cause-effect relation between human and machine, as illustrated in Fig. 1. Specifically, a set of signals from the machine are observed by human senses, and another set of signals from the human muscles affect the machine in defined ways.

Problems of Human-Machine Systems

Usually, the problems of interest to human-machine systems engineers have to do with meeting goals of system performance and safety. Performance includes communication, decision-making, feedback control, reliability, and the negatives of error and cost. Safety means minimization of death and injury to personnel and minimization of property damage, and includes consideration of both probability and consequences of accidents.

This, of course, includes the design and operation of systems to achieve desired performance and safety. The human-machine systems engineer, therefore, assumes responsibility for design of displays and controls for the human operator, design of procedures for normal and emergency operation and for maintenance and repair, and design of instructions and training programs.

These days, the computer figures in some form into most sophisticated human-machine systems, even though the purpose of the system is not operation of the computer per se but rather the operation of some vehicle or machine or industrial process in real time. Computers acquire information from sensors and data bases when needed, generate displays for the human senses, render advice for human consideration and decision, transform human commands into machine actions, and store data. If requested, computers also close low level automatic decision and control loops under human supervision. For these reasons, the human-machine system engineer must be familiar with the use and abuse of computer hardware and software.

Human Control: Emerging New Roles

Control Theory Applied to Direct Manual Control. During the 1940s, aircraft designers needed to characterize the transfer-function of the human pilot mathematically. This is necessary for any vehicle or controlled physical process for which the human is the controller, see Fig. 2. Here, both the human operator **H** and the physical process **P** lie in the closed loop (where **H** and **P** are LaPlace transforms of the component

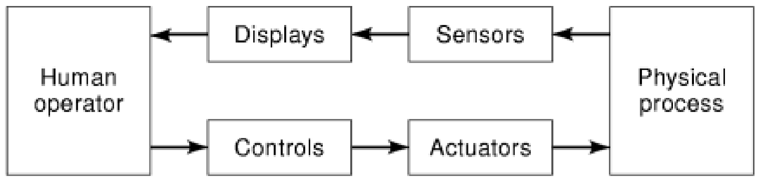

Fig. 1. Human-machine system.

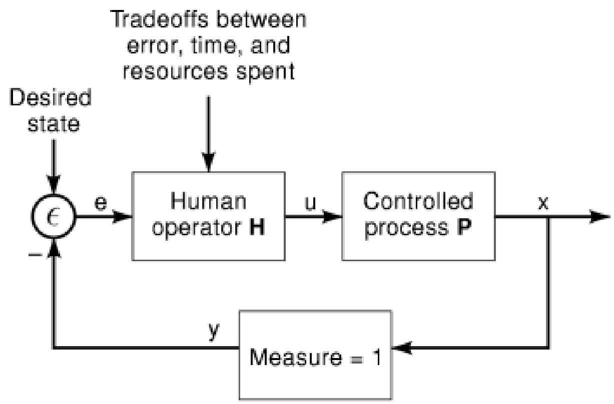

Fig. 2. Direct manual control.

transfer functions), and the **HP** combination determines whether the closed loop is inherently stable (i.e., the closed loop characteristic equation $1 + \mathbf{HP} = 0$ has only negative real roots). **H** must be chosen carefully by the human for any given **P** not only to ensure stability but also to provide good transient and steady-state response.

Efforts to characterize the pilot in these terms resulted in the discovery that the human adapts to a wide variety of physical processes so as to make $\mathbf{HP} = K(1/s)(e^{-sT})$. In other words, the human adjusts **H** to make **HP** constant. The term K is an overall amplitude or gain, $(1/s)$ is the LaPlace transform of an integrator, and (e^{-sT}) is a delay T long (the latter time delay being an unavoidable property of the nervous system). Parameters K and T vary modestly in a predictable way as a function of the physical process and the input to the control system. This so-called "simple crossover" model (1) is now widely accepted and used, not only in engineering aircraft control systems but also in designing automobiles, ships, nuclear and chemical plants, and a host of other dynamic systems.

The New Form: Supervisory Control. Supervisory control may be understood in terms of the analogy between a supervisor of subordinate staff in an organization of people and the human overseer of a modern, computer-mediated, semi-automatic control system. The human supervisor gives human subordinates general instructions which they, in turn, may translate into action. The human supervisor of a computer-controlled system does the same.

Strictly, supervisory control means that one or more human operators are setting initial conditions for, intermittently adjusting, and receiving high-level information from a computer that itself closes a control loop in a well-defined process through artificial sensors and effectors. For some time period, the computer controls the process automatically.

More generally, supervisory control is used when a computer transforms human operator commands to generate detailed control actions or makes significant transformations of measured data to produce integrated summary displays. In this latter case, the computer need not have the capability to commit actions based

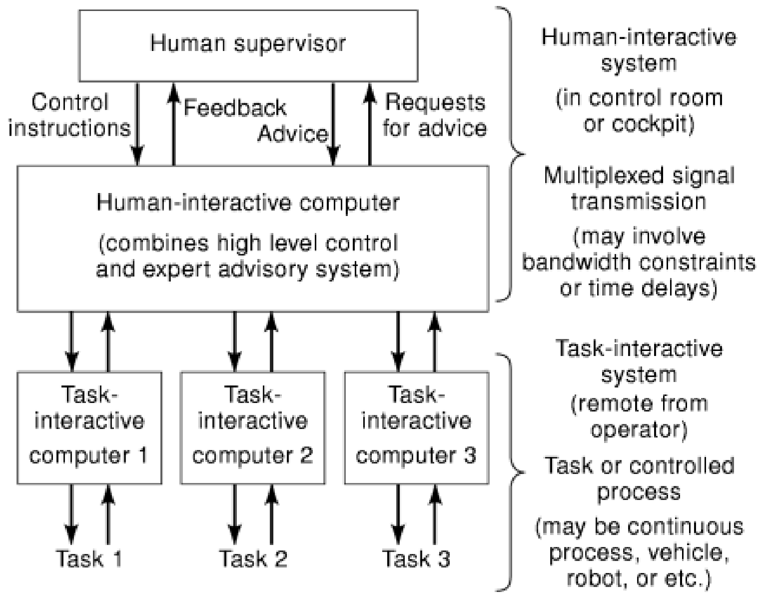

Fig. 3. Supervisory control.

upon new information from the environment, whereas in the first, it necessarily must. The two situations may appear similar to the human supervisor since the computer mediates both human outputs and human inputs, and the supervisor is thus removed from detailed events at the low level.

Figure 3 shows a supervisory control system. Here, the human operator issues commands to a human-interactive computer which can understand high level language and provide integrated summary displays of process state information back to the operator. This computer, typically located in a control room or cockpit or office near to the supervisor, in turn communicates with at least one, and probably many (hence the dotted lines), task-interactive computers, located with the equipment they are controlling. The task-interactive computers thus receive subgoal and conditional branching information from the human-interactive computer. Such information serves as reference inputs, and then, the task-interactive computers close low-level control loops between artificial sensors and mechanical actuators, thereby accomplishing the low level automatic control.

The process being controlled is usually at some physical distance from the human operator and his human-friendly, display-control computer. Therefore, the communication channels between computers may be constrained by multiplexing, time delay, or limited bandwidth. The task-interactive computer, in this case, sends analog control signals to and receives analog feedback signals from the controlled process. The controlled process does the same with the environment as it operates (vehicles moving relative to air, sea, or earth; robots manipulating objects; process plants modifying products; etc.).

Command and feedback channels for process state information are shown in Fig. 3 to pass through the left side of the human-interactive computer. Represented on the right side are decision-aiding functions with requests of the computer for advice and displayed output of advice (from a data base, expert system, or simulation) to the operator. Many new developments in computer-based decision aids for planning, editing, monitoring, and failure detection are forming an auxiliary part of operating dynamic systems.

The nervous system of higher animals reveals a similar kind of supervisory control wherein commands are sent from the brain to local ganglia, and peripheral motor control loops are then closed locally through receptors in the muscles, tendons, or skin. The brain, presumably, does higher level planning based on its own stored data and "mental models," an internalized expert system available to provide advice and permit trial responses before commitment to actual response.

Supervisory control ideas grew from the realization that many systems were becoming automated, and that the human was being replaced by the computer for direct control responsibility and was moving to a new role of monitor and goal-constraint setter. The space program posed another form of supervisory control problem: how a human operator on earth could control a manipulator arm or vehicle on the moon through a three-second communication round-trip time delay. The solution which most easily avoided instability was to make the operator a supervisor communicating intermittently with a computer on the moon which, in turn, closed the control loop there. The rapid development of microcomputers has forced a transition from manual control to supervisory control in a variety of applications (2,3).

The next section provides some examples of human-machine interaction which constitute supervisory control. First, we consider three forms of vehicle control, namely, control of modern aircraft, *intelligent* highway vehicles, and high speed trains. All three have both human operators in the vehicles as well as humans in centralized traffic control centers. Then, we consider remotely operated manipulators and vehicles for distant and/or hazardous environments.

Examples of Salient Human Machine Systems That are Undergoing Active Change

Advanced Control of Commercial Aircraft: Flight Management Systems. The aircraft industry has understood the importance of human-machine interaction from its beginning and today, exemplifies the most sophisticated forms of human-machine interaction. The flight management system is an excellent example of supervisory control, where the pilot flies the aircraft by communicating in high-level language through a computer intermediary. The flight management system is a centralized computer which interacts with a great variety of sensors, communication from the ground as well as many displays and controls within the aircraft. It includes many functions and mediates most of the pilot information requirements. In former days, each sensor had its own display, operating independently of all other sensor-display circuits. The flight management system brings together all of the various autopilot modes, from long-standing, low-level control modes wherein the aircraft is commanded to go to and hold a commanded altitude, heading, and speed, to more sophisticated modes where the aircraft is instructed to fly a given course, consisting of a sequence of waypoints (latitudes and longitudes) at various altitudes, and even land automatically at a given airport on a given runway.

One type of display mediated by the flight management system, shown in Fig. 4, integrates many formerly separate components of information. It is a multicolor, plan-view map showing position and orientation of important objects relative to one's own aircraft (the triangle at the bottom). It shows heading (compass arc at top, present heading 175°), groundspeed plus windspeed and wind direction (upper left), actual altitude relative to desired altitude (vertical scale on right side), programmed course connecting various waypoints (OPH and FLT), salient radio navigation beacons to the right and left of present position/direction with their codes and frequencies (lower left and right corners), the location of key navigation beacons along the course (three cornered symbols), the location of weather to be avoided (two gray blobs), and a predicted trajectory based on present turn rate, showing that the right turn is appropriately getting back on course.

The flight management system is programmed through a specialized keyboard and text display unit (Fig. 5) including all the alphanumeric keys plus a number of special function keys. The displays in this case are specialized to the different phases of a flight (Taxi, takeoff, departure, enroute approach, land, etc.), each phase having up to three levels of pages.

Designing displays and controls is no longer a matter of what can be built; the computer and the flight management system allows essentially any conceivable display/control to be realized. The computer can also provide a great deal of real-time advice, especially in emergencies, based on its many sensors and stored knowledge about how the aircraft operates. But pilots are not sure they need all the information which aircraft

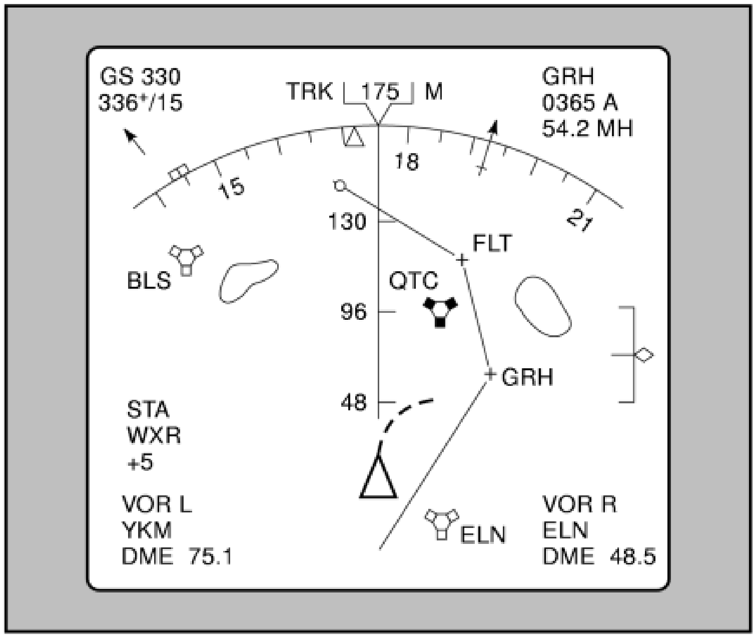

Fig. 4. Integrated aircraft map display [from Billings (4)].

designers would like to give them and have an expression, *killing us with kindness* to refer to this plethora of available information. The question is what should be designed based on the needs and capabilities of the pilot.

The aircraft companies Boeing, McDonnell Douglas and Airbus have different philosophies for designing the flight management system. Airbus has been the most aggressive in automating, intending to make piloting easier and safer for pilots from countries with less well established pilot training. Unfortunately, it is these most automated aircraft which have had the most accidents of the modern commercial jets—a fact which has precipitated vigorous debate about how far to automate.

The Evolution of Air Traffic Control. Demands for air travel continue to increase and so do demands on air traffic control. Based on what is currently regarded as safe separation criteria, air space over major urban areas is already saturated. Simply adding more airports is not acceptable (in addition, residents do not want more airports with their noise and surface traffic). The hope is to reduce separations in the air without compromising safety and to land aircraft closer together or on parallel runways simultaneously. The result is greater demands on air traffic controllers, particularly at the terminal area radar control centers (*TRACON*), where trained operators stare at blips on radar screens and verbally guide pilots entering the terminal airspace from various directions and altitudes into orderly descent and landing patterns with proper separation between aircraft.

Many changes are now being introduced into air traffic control which have profound implications for human-machine interaction (5). Previously, communication between pilots and air traffic controllers was entirely by voice, but now, digital communication between aircraft and ground (a system called *datalink*) allows both more frequent and more reliable two-way communication. Through datalink, weather, runway, and wind information, clearances, etc. can be displayed to pilots visually. However, pilots are not sure they want this additional technology. The demise of the partyline of voice communications with which they are so familiar, and which permits all pilots in an area to listen in on each other's conversations, is a threatening prospect.

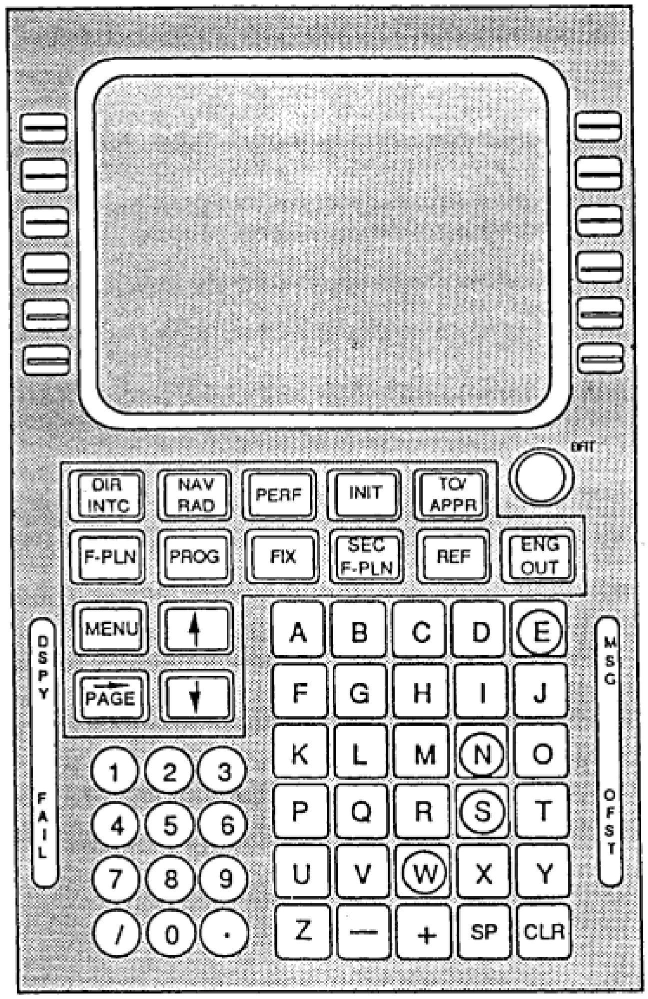

Fig. 5. Flight management system control and display unit [from Billings (4)].

Now, there are aircraft-borne radars which allow pilots to detect air traffic in their own vicinity. There are improved ground-based radars which detect microbursts or windshear which can easily put an aircraft out of control. But with both types of radar, there are questions as to how best to warn the pilot and provide guidance as to how to respond. These new systems also pose a cultural change in air traffic control, since until now, pilots have been dependent upon air traffic controllers to advise them of weather conditions and other air traffic. In addition, because of the new weather and collision avoidance technology, there are current plans for radically altering the rules whereby high altitude commercial aircraft must stick to well defined traffic lanes. Airlines want pilots to have more flexibility as to altitude (to find the most favorable winds and therefore save fuel) and to be able to take great-circle routes straight to their destinations (also saving fuel). However, air traffic controllers are not sure they want to give up the power they have had, becoming passive observers and monitors and functioning only in emergencies.

Guidance and Navigation Systems for Highway Vehicles. Having *GPS* (global positioning system) satellites, high density computer storage of map data, electronic compass, synthetic speech synthesis, and computer-graphic displays permits cars and trucks to know where they are located on the earth to within 100 meters or less. New navigation systems can guide a driver to a programmed destination by a combination of a map display and speech. There are human factor challenges in deciding how to configure the maps (how much detail to present, whether to make the map north-up with a moving dot representing one's own vehicle position, or current heading up and rapidly changing with every turn). Computer graphics can also be used to show what turns to anticipate and when to get in which lane. Computer-generated speech can reinforce these turn anticipations, can caution the driver if he is perceived to be headed in the wrong direction or off-course, and can even guide him or her how to get back on course. One experimental question is what the computer should say in each situation to get the driver's attention, to be understood quickly and unambiguously, but without being an annoyance. Another question is whether such systems will distract the driver's attention from the primary tasks, thereby reducing safety. The various vehicle manufacturers have developed and evaluated such systems for reliability and human use, and they are beginning to market them in the United States, Europe, and Japan.

Intelligent Cruise Control. Conventional cruise control has a major shortcoming. It knows nothing about vehicles ahead, and one car can easily collide with the rear end of another car if the driver is not careful. Intelligent cruise control means that a microwave or optical radar detects the presence of a vehicle ahead and measures that distance. But then there is the question of what to do with this information: just warn the driver with some visual or auditory alarm or automatically brake? A warning sound is better than a visual display because the driver does not have to be looking in the right place. But can a warning be too late to elicit braking or surprise the driver so that he brakes too suddenly and causes a rear-end accident to his own vehicle? Should the computer automatically apply the brakes by some function of distance to the obstacle ahead, speed, and closing deceleration? If the computer did all the braking, would the driver become complacent and not pay attention, to the point where a serious accident would occur if the radar failed to detect an obstacle, say a pedestrian or bicycle, or the computer failed to brake? Should braking be some combination of human and computer braking, and if so, by what algorithm? These are human factors questions which are currently being researched.

It is worth noting that current developmental systems only decelerate and down-shift mostly because if the vehicle manufacturers sell vehicles which claim to perform braking, they would be open to a new and worrisome area of tort litigation.

It should also be mentioned that the same radar technology that can warn the driver or help control the vehicle can also be applied to cars overtaking from one side or the other. Another set of questions then arises as to how and what to communicate to the driver and whether to trigger some automatic control maneuver in certain cases.

Control of High Speed Passenger Trains. Railroad technology has lagged behind that of aircraft and highway vehicles with respect to new electronic technology for information sensing, storage, and processing, but currently it is catching up. The proper role of the human operator in future rail systems is being debated since for some limited right-of-way trains (e.g., in airports), one can argue that fully automatic control systems now perform safely and efficiently.

The principal job of the train driver is speed control (although there are many other monitoring duties he must perform). In a train, this task is much more difficult than in an automobile because of the huge inertia of the train, since it takes 2 to 3 km to stop a high-speed train. Speed limits are fixed at reduced levels for curves, bridges, grade-crossings, and densely populated areas, while wayside signals temporarily command lower speeds if there is maintenance being performed on the track, there are poor environmental conditions such as rock slides or deep snow, or especially if there is another train ahead. It is mandatory that the driver obey all speed limits and get to the next station on time. It can take months to learn to maneuver the train with its long time constants, given that for the speed control task, the driver's only input currently is an indication of current speed.

A new computer-based display which helps the driver anticipate the future effects of current throttle and brake actions has been developed in the Human-Machine Systems Laboratory at MIT. It is based on a dynamic model of the train and gives an instantaneous prediction of future train position and speed based on current acceleration. It allows speed to be plotted on the display assuming the operator holds to current brake-throttle settings and also plots trajectories for maximum emergency braking and maximum service braking. The computer generates an optimal speed trajectory which adheres at all (known) future speed limits, gets to the next station on time, and minimizes fuel/energy.

Remote Manipulators for Hazardous Radiation, Space, and Other Environments. The development of master-slave remote manipulators started when nuclear power was first adopted in the late 1940s. Using such devices, a human operator at one location could position and orient a device attached to his hand, and a servomechanism-controlled gripper would move in correspondence and handle objects at another location which was too hazardous for a human. Such manipulators remotely controlled by humans are called teleoperators.

Teleoperator technology got a big boost from the industrial robot technology, which came in a decade or so later, and provided improved vision, force and touch sensors, actuators, and control software. Large teleoperators were developed for rugged mining and undersea tasks, and small teleoperators were developed for delicate tasks such as eye surgery. Eventually, teleoperators have come to be equipped with sensitive force feedback, so that the human operator cannot only see the objects in the remote environment but also can feel them in his grasp.

The space program naturally stimulated the desire to control lunar manipulators from earth. However, the unavoidable round trip time delay of three seconds (speed of light from earth to moon and back) would not permit simple closed loop control. Supervisory control provided an answer. The human on earth could communicate to the moon a subgoal to be reached and a procedure for getting there, and the teleoperator would be turned loose for some short period to perform automatically. Such a teleoperator is called a telerobot.

The Flight Telerobotic Servicer (*FTS*) developed by Martin Marietta for the US Space Station Freedom is pictured in Fig. 6. It has two seven-degree of freedom (*DOF*) arms (including gripper) and one five DOF leg for stabilizing itself while the arms work. It has two video eyes to present a stereo image to its human operator. It can be configured either as a master-slave teleoperator (under direct human control) or as a telerobot (able to execute small programmed tasks using its own eyes and force sensors).

Remotely Operated Submarines and Planetary Rovers. The same principles of supervisory control as can be applied to telemanipulators can also be applied to submarines and planetary roving vehicles. In Fig. 7 is shown the remotely operated submersible Jason developed by Woods Hole Oceanographic Institution. It is the big brother of Jason Junior which swam into the interior of the ship Titanic and made a widely viewed video record when the latter was first discovered. Jason has a single manipulator arm, sonar and photo sensors, and four thrusters which can be oriented within limited range and which enable it to move in any direction. It was designed for the severe pressures at depths up to 6000 m. It can be operated either in direct teleoperator mode or as a telerobot.

The Mars Pathfinder roving vehicle is yet another example of supervisory control, in this case operating over one-way time delays of more than 10 min.

Methods of Human-Machine Systems Engineering

The following four steps are typically carried out in designing new human-machine systems or making improvements in existing such systems.

(1) Function and Task Analysis This is a detailed paper and pencil exercise wherein system operation is characterized as a sequence of steps, and for each step, requirements are specified for

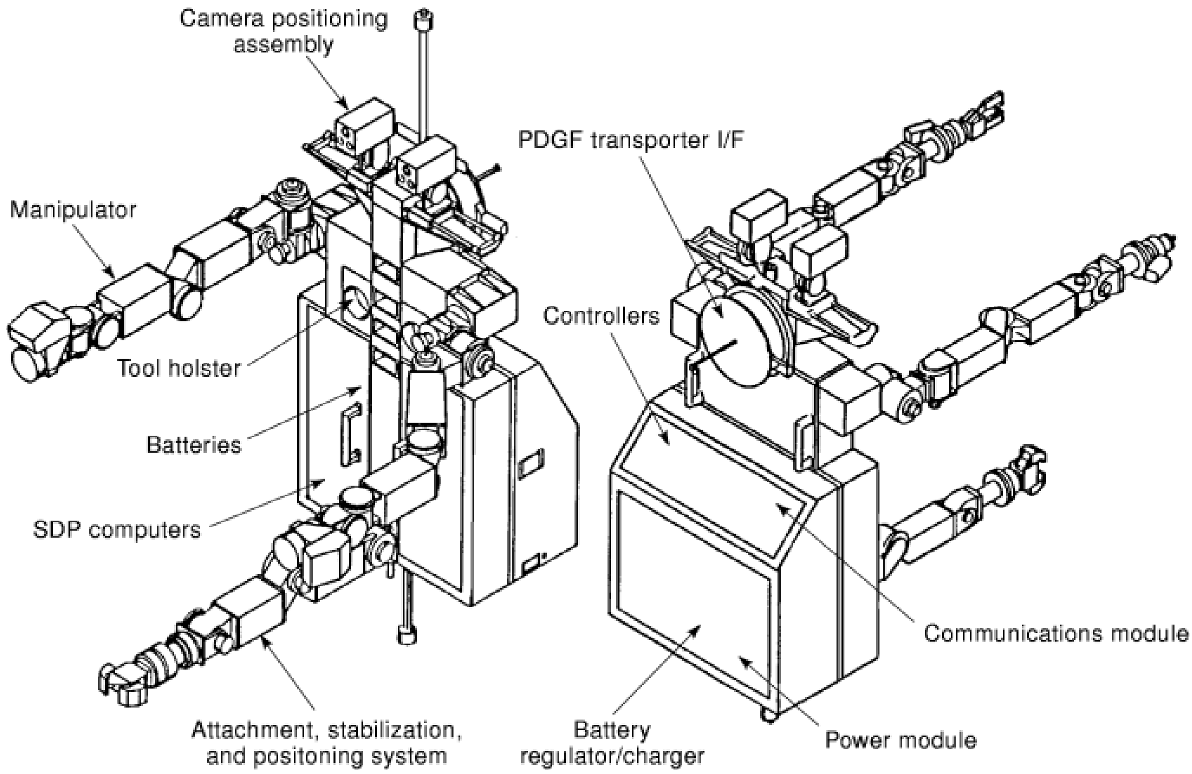

Fig. 6. Flight telerobotic servicer prototype design (courtesy NASA).

(1) what variables must be controlled to what criterion (what constitutes satisfactory completion of that step). Each step can be accomplished by a machine or a human or a combination of these. A notation is made of what decisions must be made at that step and what are the tradeoffs among resources (such as time, energy, raw materials used) for that decision/control. It is noted what control devices must or could be actuated.

(2) what information is needed by the human or machine entity performing that step. The current source of that information, whether direct observation of the process, viewing of a display, reading of text or printed illustration, or verbal communication, is noted. The probabilistic aspects of inputs and disturbances are also noted, as well as how the operator or machine can determine whether the system is failing in some mode.

(3) Tentative allocations are made at this point as to whether human or machine ought to perform each task or function in that step and what improvements are suggested in displays, controls, or procedures.

(4) Design Based on Common Criteria for Human-Machine Interfacing Design of human-system interactions poses the same types of problems no matter what the context. The displays must show the important variables unambiguously to whatever accuracy is required, but more than that must show the variables in relation to one another so as to clearly portray the current situation (situation awareness is currently a popular test of the human operator in complex systems). Alarms must get the operator's attention, indicate by text, symbol, or location on a graphic display what is abnormal, where in the system the failure occurred, what is the urgency, and if response is urgent, even suggest what action to take. (For example, the ground proximity warning in an aircraft gives a loud, *Whoop, whoop!* followed by a distinct spoken command, *Pull up, pull up!*).

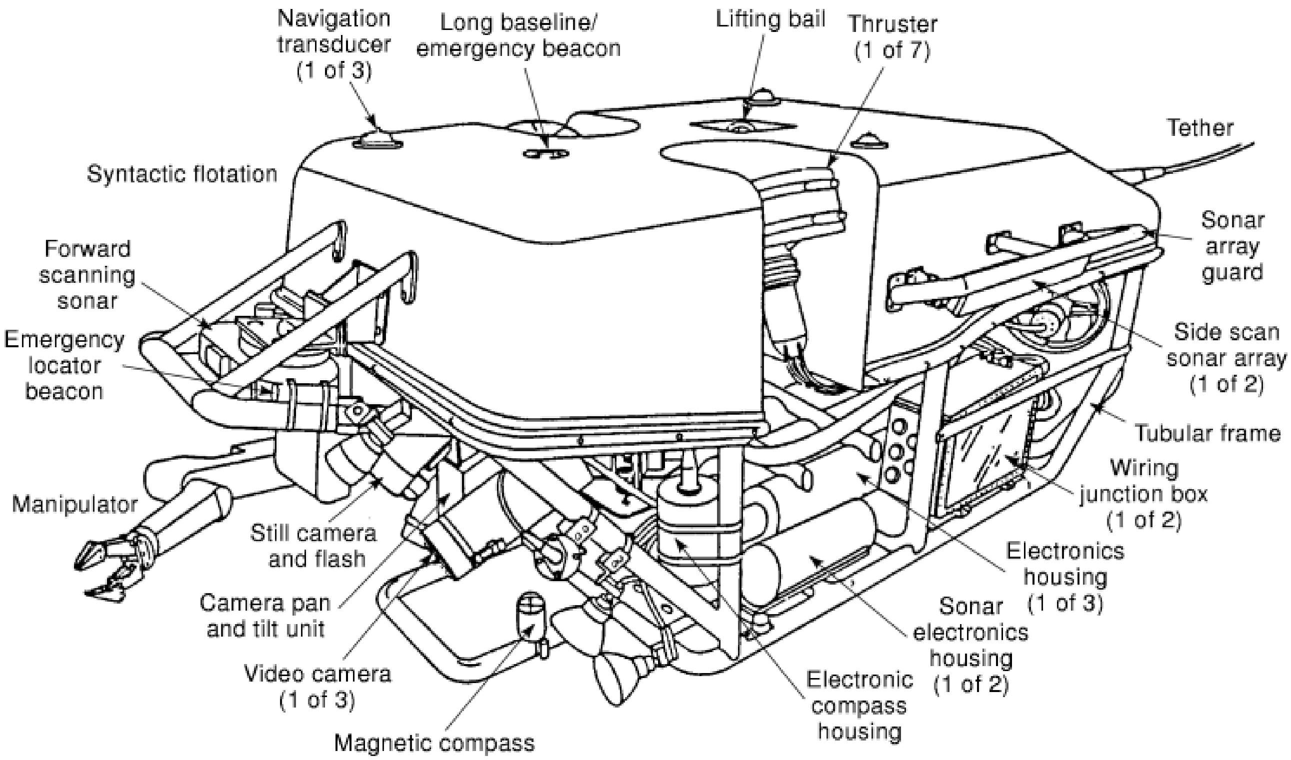

Fig. 7. Deep ocean submersible Jason (courtesy Woods Hole Oceanographic Institution).

(5) Controls—whether analogic joysticks, master-arms, or knobs, or symbolic special-purpose buttons or general purpose keyboards—must be natural and easy to use and require little memory of special procedures (computer icons and windows do well here). The placement of controls and instruments and their mode and direction of operation must correspond to the desired direction and magnitude of system response.

(6) Human engineering design handbooks are available but require experience to interpret for the same reason that medicine cannot be practiced from a handbook. Humans are not as simple as machines, and we do not have a tool kit of well codified engineering principles.

(7) Mathematical Modeling With the task analysis in hand to rough out the topology of the process and a preliminary design, appropriate mathematical tools are often brought into account. The most useful ones for human-automation interactions are listed below: Control models are normally essential to characterize the automation itself (robot or vehicle movement, etc.); however, we are most concerned with characterizing control at a higher level, where computer or human decision-making adjusts parameters or make changes. Fuzzy control techniques are particularly useful to characterize those aspects of control which are not amenable to linearization and precise quantification, particularly where adaptation to environmental variables or variations in raw material require adjustment. Control decisions by the human operator may also be characterized as fuzzy control in terms of several key variables (which need not be continuous or even numerical).

Information models are useful to characterize the complexity (entropy) of a process and the effort required to narrow down to decision and action. Any input-output array can be characterized by information transmission measures to show how consistent the process is. Information measures are always averages over a set of events.

Bayesian decision-making models are appropriate to provide a normative model of inferring truth about the state of the system from various forms of evidence.

Conventional decision models specify what action should be taken to maximize the subjectively expected utility (value or relative worth) of consequences, given the prior probabilities of events which copule to those consequences.

Signal detection models are appropriate to human/machine detection of signals in noise, determination of failures, etc. where there is some continuous variable which correlates with the degree of decision confidence, and probability density functions on that variable can be derived conditional on true positive (actual failure) vs. false positive (false alarm).

(8) Human-in-the-Loop Simulation In parallel with mathematical analysis, it is generally appropriate to develop a real-time simulation of the process or some aspect of the process that needs to be changed. Simulating the physics of the process are not necessary but instead, only the dynamics of the interactions of key variables, to the extent that the displays and controls can be simulated and appear to allow for realistic human interaction. The simulation can perhaps be implemented on a PC, but a graphics workstation is often more appropriate. Experiments are then carried out on the simulation with trained human subjects, while parametric changes are made which simulate design changes in the process, the automation, the displays and controls, or the procedures. Sufficient trial runs are made to allow for simple statistical comparisons of system performance under the various parameter treatments. System performance includes objective measures of product quality, time, cost, human error, and mental workload as well as subjective appraisals. This gives evidence of what design alternatives are best and by how much.

(9) Refinement, Optimization, and Pilot Testing All the evidence gleaned from the above three steps is codified in the form of specific design improvement recommendations. Insofar as economically feasible, these are implemented on a pilot automation system and evaluated in conventional ways.

Workload and Error

New technology allows combination, integration, and simplification of displays compared to the intolerable plethora of separate instruments in older aircraft cockpits, plant control rooms, and instrument display/control panels. The computer has taken over more and more functions from the human operator. Potentially, these changes make the operator's task easier. However, it also allows for much more information to be presented, more extensive advice to be given, and more complexity in operation, particularly when the automation fails. These changes add many cognitive functions that were not present at an earlier time. They make the operator into a monitor of the automation who is supposed to step in when required to set things straight. Unfortunately, people are not always reliable monitors.

Mental Workload. It is imperative to know whether the mental workload of the operator is too great for safety (physical workload is of less and less concern). Human–machine systems engineers have sought to develop measures of mental workload, the idea being that as mental load increases, the risk of error in performance increases, but presumably, measurable mental load comes before actual lapse into error. Thus mental workload is a more sensitive predictor of performance deterioration than performance itself.

Mental workload may be measured by three techniques:

(1) The subjective rating scale, typically a ten-level category scale with descriptors for each category from no load to unbearable load. This is the most reliable measure, though it is necessarily subjective.

(2) Physiological indices which correlate with subjective scales, including heart rate and the variability of heart rate, certain changes in the frequency spectrum of the voice, electrical resistance of the skin, diameter of the pupil of the eye, and certain changes in the evoked brain wave response to sudden sound or light stimuli.

These techniques would be attractive were it not for the wide variability between human subjects and the *noise* in measurements.

(3) The so-called secondary task, an easily measurable additional task which consumes all of the operator's attention remaining after the requirements of the primary task are satisfied. This technique has been used successfully in the laboratory but has shortcomings in practice in that operators may refuse to cooperate.

These techniques are now routinely applied to critical tasks such as aircraft landing, air traffic control, certain planned tasks for astronauts, and emergency procedures in nuclear power plants. The evidence suggests that supervisory control relieves mental load when things are going normally, but when automation fails, the human operator is subjected to rapidly increased mental load.

Human Error. While human error has long been of interest to psychologists and industrial engineers, only in recent decades has there been serious effort to understand human error in terms of categories, causation, and remedy (7).

Human error may be classified in several ways. One is according to whether an error is one of omission (something not done which was supposed to have been done) or commission (something done which was not supposed to have been done). Another is slip (a correct intention for some reason not fulfilled) versus mistake (an incorrect intention which was fulfilled). Errors may also be classified according to whether they are in sensing, perceiving, remembering, deciding, or acting. There are some special categories of error worth noting which are associated with following procedures in operation of systems. One, for example, is called a capture error, wherein the operator, being very accustomed to a series of steps, say A, B, C, and D, intends at another time to perform E, B, C, F. But he is "captured" by the familiar sequence B, C and does E, B, C, D.

With regard to effective therapies for human error, proper design to make operation easy, natural, and unambiguous is surely the most important. It is always best if system design allows for error correction before the consequences become serious. Active warnings and alarms are necessary when the system can detect incipient failures in time to take such corrective action. Training is probably next most important after design, but any amount of training cannot compensate for an error-prone design. Preventing exposure to error by guards, locks, or an additional "execute" step can help make sure that the most critical actions are not taken without sufficient forethought. Least effective are written warnings such as posted decals or warning statements in instruction manuals, although many tort lawyers would like us to believe the opposite.

Social Implications of Modern Human-Machine Interaction

Trust. Trust is a term not often taken an engineering variable, but it is rapidly taking on such a connotation. When an operator does not trust his sensors and displays, expert advisory system, or automatic control system, he will not use it or will avoid using it if possible. On the other hand, if an operator comes to place too much trust in such systems, he will let down his guard, become complacent, and when it fails, not be prepared. The question of operator trust in the automation is an important current issue in human-machine interface design. It is desirable that operators trust their systems, but it is also desirable that they maintain alertness, situation awareness, and readiness to take over, so there can be too much trust.

Alienation. There is a set of broader social concerns that the new human-machine interaction can have, which can be discussed under the rubric of alienation.

(1) People worry that computers can do some tasks much better than they themselves can, such as memory and calculation. Surely, people should not try to compete in this arena.

(2) Supervisory control tends to make people remote from the ultimate operations they are supposed to be overseeing—remote in space, desynchronized in time, and interacting with a computer instead of the end product or service itself.

Table 1. Scale of Degrees of Automation

1.	The computer offers no assistance: the human must do it all.
2.	The computer offers a complete set of action alternatives, and
3.	narrows the selection down to a few, or
4.	suggests one alternative, and
5.	executes that suggestion if the human approves, or
6.	allows the human a restricted time to veto before automatic execution, or
7.	executes automatically, then necessarily informs the human, or
8.	informs the human only if asked, or
9.	informs the human only if it, the computer, decides to
10.	The computer decides everything and acts autonomously, ignoring the human.

From Sheridan, 1987

(3) People lose the perceptual-motor skills which in many cases gave them their identity. They become deskilled, and if ever called upon to use their previous well-honed skills, they could not.

(4) Increasingly, people who use computers in supervisory control or in other ways, whether intentionally or not, are denied access to the knowledge to understand what is going on inside the computer.

(5) Partly as a result of factor four, the computer becomes mysterious, and the untutored user comes to attribute to the computer more capability, wisdom, or blame than is appropriate.

(6) Because computer-based systems are growing more complex, and people are being elevated to roles of supervising larger and larger aggregates of hardware and software, the stakes naturally become higher. Where a human error before might have gone unnoticed and been easily corrected, now such an error could precipitate a disaster.

(7) The last factor in alienation is similar to the first, but more all-encompassing: namely the fear that a "race" of machines is becoming more powerful than the human race. These seven factors and the fears they engender whether justified or not, must be managed. Computers must be made to be not only human friendly but also not alienating with respect to these broader factors. Operators and users must become computer-literate at whatever level of sophistication they can.

How Far to Go With Automation. The trend toward supervisory control is surely changing the role of the human operator, posing fewer requirements on continuous sensory-motor skill and more on planning, monitoring, and supervising the computer. As computers take over more and more of the sensory-motor skill functions, new questions are being raised regarding how the interface should be designed to provide the best cooperation between human and machine. Among these questions are: To what degree should the system be automated? How much help from the computer is desirable? What are there points of diminishing returns?

Table 1 lists ten levels of automation, from 0 to 100% computer control. Obviously, there are few tasks which have achieved 100% computer control, but new technology pushes relentlessly in that direction. It is instructive to consider the various intermediate levels of Table 1 in terms not only of how capable and reliable is the technology but also what is desirable in terms of safety and satisfaction of the human operators and the general public.

BIBLIOGRAPHY

1. D. T. McRuer H. R. Jex A review of quasi-linear pilot models, *IEEE Trans. Hum. Factors Electron.*, **HFE-4**: 231–249, 1967.
2. T. B. Sheridan *Telerobotics, Automation, and Human Supervisory Control*, Cambridge, MA: MIT Press, 1992.
3. T. B. Sheridan Supervisory control. In G. Salvendy (Ed.), *Handbook of Human Factors/Ergonomics*, New York: Wiley, 1987.
4. C. E. Billings *Human-Centered Aircraft Automation: A Concept and Guidelines*, Moffet Field, CA: NASA Ames Research Center, 1991.
5. C. Wichens National Academy of Science, *The Future of Air Traffic Control*, Washington, D.C.: National Academy Press, 1998.
6. S. Y. Askey *Design and Evaluation of Decision Aids for Control of High Speed Trains: Experiments and a Model*, PhD Thesis, Cambridge, MA: Massachusetts Institute of Technology, June, 1995.
7. J. Reeson *Human Error*, Cambridge Univ. Press, 1990.

THOMAS B. SHERIDAN
Massachusetts Institute of Technology