Int. J. Man–Machine Studies (1978) **10**, 225–232

Man–computer communication—what next?†

B. R. GAINES
Man–Machine Systems Laboratory, University of Essex, Colchester, Essex, U.K.

(*Received 16 May 1977*)

The future as she is told

As a topic, "What next?" has the obvious virtue of suffering no (immediate!) contradiction—one needs only a vivid imagination and a plausible manner. However, "What next?" also has the defect that the great ones of science fiction have done it so much better—in the realms of man–computer communication who can compete with the "feelies" (and the "smellies"), Simak's "dimensino" which gives "total participation" in the events portrayed, or even conversations with HAL.

Yet the level-zero problem of steering between plausible extrapolation and science fiction is compounded at level-one by the fact that much of the most outrageous science fantasy is remarkably close to the state-of-the-art today. We do have tactile input "gloves" that enable us to feel non-existent "objects", and automated psychological experiments on multivariate nasal perception get remarkably close to the technology of the "smellies". HAL can now speak to us clearly and incisively. His hearing is somewhat erratic and his conversation somewhat moronic—it is a geriatric HAL we have built, but speech and knowledge science hold the promise of youthful vigour yet to come.

The "dimensino" is still in pieces, with echoes of the future in Pask's scheme for audience participation in the theatre (the players receive on-line voting from the audience on their course of action!), Hash's suggestion for an "international ARPA" in which questions on ancient Chinese wisdom at one terminal are relayed to an ancient Chinese sage at another, and so on. Total participation" is still a dream—and yet that may be just where the solution lies. We *do* have perceptual processes in dreams that are not invoked by the corresponding sensory stimulation. Aids to the blind that operate by direct cortical stimulation are already with us. Perhaps our hallucinations may also be subject to such direct control. Those familiar with the multi-level hallucinatory hierarchy of Stanislav Lem's *Futurological Congress* will not need persuading of the possibilities. What company would spend millions of dollars and years of effort on developing a new technology if it could persuade its customers to hallucinate that such already existed?—in that world we shall all be driving mind-controlled Rolls Royces as we cycle to work!

Normative and predictive forecasting

Enough of the fantasy—I hope it makes my excuses for not looking too far ahead. Essentially, we must rate progress by the rate at which technology hits the market place, and I would like to concentrate on what is happening in the near future—in many respects that is fantastic enough!

†This paper formed the keynote address at the banquet of the 5th Man–Computer Communications Conference, Calgary, May 1977.

0020–7373/78/030225+08 $02.00/0 © 1978 Academic Press Inc. (London) Limited

Forecasts may be of two distinct types: *predictive*, in which one attempts to project current trends and develop required technologies; and *normative*, in which one attempts to impose future trends and develop required technologies.These two aspects are complementary and both are important in practical applications. Purely predictive forecasting is extrapolation from essentially noisy data and hence subject to comulative error. It also has the danger that when things are going badly because a technology is lacking, then the only prediction is that they will get worse! Purely normative forecasting is interpolation between the present situation and future goals and thus less subject to random variation (provided the goals are actually accepted!). It has the dangers of setting unrealistic targets or of routing effort to less rewarding objectives.

I mention both aspects of forecasting because I feel we should always read "What next?" not only as predictive but also as normative; what *should* come next, as well as what may. As system designers there is a strong normative component to our actions—we are setting standards in each new system delivered. For some time yet, most users will be having their *first experience* of man–computer communication through the systems we are now developing. Even the few "experienced" users may still be sufficiently resistant to the brain-washing of past systems to react favourably to improved man–computer communication (even Asimov, who should know better, makes his computer technicians more like priests than engineers—communication with Multivac is through rather bizarre specialists only—we have a lot to live down!).

Thus there is a real need to take continual cognizance of the normative role of our work. Some of the technology we need is subject to technical limitations of display rates, costs, computational power, etc., but much of it is in the knowledge-limited areas. *Knowledge only* is the limitation in so many aspects of man–computer interaction—the current hardware technologies are adequate for many tasks in which our current lack of success is due solely to our lack of knowledge of how to use them. Speech recognition is clearly in this category—electronic signal processing capabilities go way beyond what is needed at the lower level, computer data storage capacities are much greater than those needed at higher levels, and, more speculatively, I believe that current computational power is adequate to cope with co-ordinating the necessary "knowledge-base". However, we do not yet know enough about either the problem or the techniques required in its solution.

At a far more mundane level I will make the same comment about interactive dialogue with the non-professional user at a computer terminal. What we most often see is unimaginative, sloppily designed, often appalling in human factors terms—why? Main answers: the users do not know enough to complain; and we *do not know enough about dialogue in the general case* to have the background to deal with the man–computer one. For many linguists the listener seems an embarrassment rather than a necessary part of dynamics of language. That he may actually interrupt and reply—that a conversation may itself have organic unity and purpose—these aspects seem relegated to the study of rhetoric and are anathema to much of the present scholarship of language.

Thus, each user of an interactive system to whom we demonstrate what *can* be done and send back dissatisfied and complaining about the systems *he* is forced to use; each linguist we persuade to take an interest in man–computer dialogue; and so on—all of these are part of our normative role in "What next?". It is easy to adopt the passive predictive role and gear one's work to the natural inertia of the world, technological or otherwise, but a major component of progress must always be a quest for what is better, what is appropriate, what is right. We have to engineer systems within our limitations

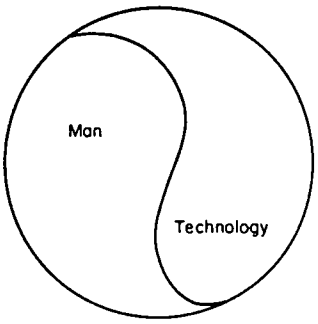

FIG. 1. Man and technology—the basic situation.

and would be foolish to attempt otherwise, but those limitations must be subject to continual questions, testing and open criticism.

Perhaps these comments seem so obvious that an example is necessary. Computer-based instruction was over-sold a decade ago. The computer was going to bring new freedom into education by individualized instruction, tailoring it to the personal needs of the student. And where was the knowledge of these needs to come from? Through a few characters per second link to a student laboriously typing capital letters at a teleprinter terminal. Would we have promised the educational advance of the millenium if it had been offered to link *our* students to *us* in the same way? Are we surprised not to have achieved it? Do some of the statements about helping students to understand their errors now not seem subject to rather obvious information-theoretic limitations?

The same thing is happening today—put all your information into a database and clearly all your problems can be solved; forecasts, resource allocation, management information and automated control of all kinds must surely follow. In another decade we may look back wryly and note that if you put all your information in a database then you have just that—all your information in a database. Information is not knowledge, and even knowledge is not necessarily power—we have a long way to go.

The man–computer situation

Figure 1 shows the basic man–technology situation. I have deliberately chosen this ancient symbology rather than the usual inter-connected boxes because it is rather more suggestive of our real objectives. The tools of man are essentially part of man himself—they have long extended our muscles and our senses—with the computer we may extend our information storage and processing capabilities—ultimately we may extend our *minds*.

The central division in Fig. 1 is initially a *barrier* dividing us from our technologies. Progress consists of making that barrier more and more permeable at appropriate points —flow of control from man to technology—flow of information from technology to man (those who claim that the flows are tending the other way point to a very real danger—what permeation *you* feel is appropriate varies, for example, with whether your social status is the governor or the governed!).

It is perhaps more realistic to draw Fig. 1 as Fig. 2, at least pre-computer. What seems plausible is that Fig. 1 is now becoming a more realistic picture as we add the computer to

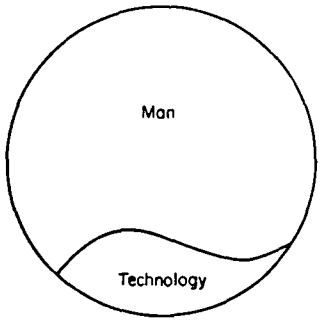

FIG. 2. Man and technology—pre-computer.

the technologies available. It is the first technology that presents the possibility of a *partnership* where *both* parties can comprehend and respect the full capabilities of each other. However, *neither* side of the division has yet achieved that level of understanding of the other.

As shown in Fig. 3, we can think of the man–technology partnership facing its common environment together. Such a system has the potential for greater comprehension and control of the environment than either of its parts alone. The term "*intelligence amplifier*" has been coined for the role of technology in this situation. However, the potential and its realization are quite distinct—we can provide many examples already of the computer as an "*ignorance amplifier*"—the potential is equally high!

Figure 4 illustrates the points made previously about the current state of man–computer communication. What it does demonstrate clearly is the current asymmetry in information-flow—man to computer is weak, slow in data rate, inaccurate and unnatural—computer to man is (or can be!) strong, fast, precise and well-matched to our data-input systems. Text preparation, colour graphics and speech synthesis are low-cost technologies —we still have plenty of ground to cover in the "knowledge-systems" to control these technologies but that is Carnapian in its incremental steady progress. Handwriting, pictures and speech as computer inputs are still research projects rather than technologies, and so far behind our requirements that a Kuhnian "paradigm-shift" seems necessary before progress can be made.

The direction of hope here is illustrated in Fig. 5. The processes of man–man communication that we wish to extend to the computer seem dependent on some assumption

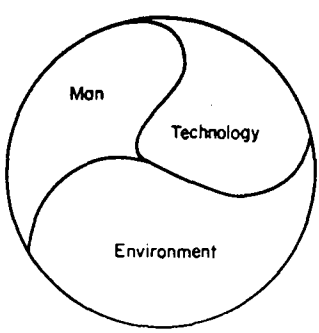

FIG. 3. Man, technology and environment.

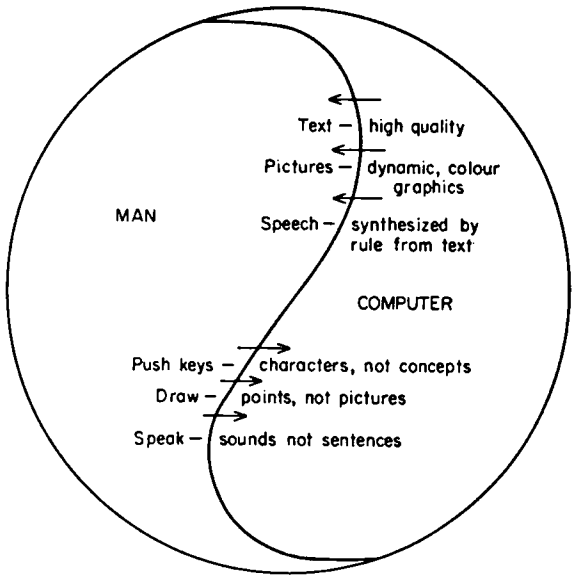

FIG. 4. The current state of man–computer communication.

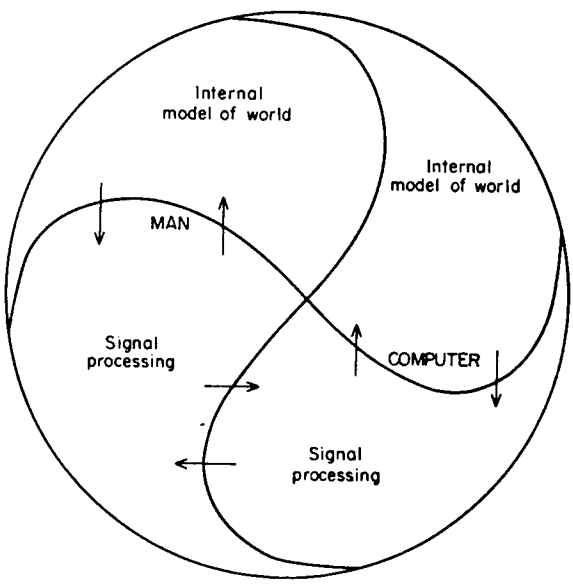

FIG. 5. The role of world-models in man–computer communication.

that those communicating share much of their internal world models—not in a Jungian sense (although a metaphor of a "racial memory" is a good one), but at least as related domains of knowledge. The source of this communicatability is difficult to analyse rigorously. Post-Humean philosophy seems to have arrived at an acceptance of his definitive argument that we cannot justify our acquisition of knowledge. We may derive what consolation (or inspiration) we like from Kant's transcendental arguments (we do it

because we do it—probably programmable in LISP but not in FORTRAN!) and agree with Wittgenstein that, "If anyone said that information about the past could not convince him that something would happen in the future, I should not understand him".

Equally, the debates in linguistics as to *why* we share a common language (given that many of us do not!), and *how* we can communicate with one another at all (and what that means anyway), are still hot-beds of argument and speculation. We come from a common genetic stock in a common, relatively stable, physical environment—perhaps that is sufficient constraint to pull our minds into some degree of synchronization. As a science the study of ourselves is severely restricted by our limited range of examples of intelligent system evolving in a common environment.

I did not digress into this slough of philosophic and linguistic despond for its own sake. It serves to make one important point about Fig. 5 that goes back to an elementary aspect of computer programming—*we cannot program that which we do not understand.* A clear illustration of this has been in the syntactic aspect of linguistics. We fumbled with natural language parsers for many years—now post-Thorne and Dewar we set such parsers as undergraduate projects. The moral is very much shown in the motivation for their work—it was to *demonstrate the consistency of a linguistic parsing schema.* Where it was weak or failed (e.g. in the conjunctions) was where the linguistic definitions were themselves vague or contradictory. Similarly, today when we argue as to whether a program has produced the correct parsing, it is generally what *is* a "correct parsing" that is the source of debate —not the correctness of the program.

It is our lack of understanding of both sides of the upper half of Fig. 5 which limits the man–computer communication at the lower level. We do not understand what constitutes the "world-model" used by a person. We have only begun to realize the difficulties of constructing adequate models of the world within computers—not miniature, circumscribed, artificial worlds, but the real world that is large, complex and unbounded, consisting of interacting *dynamic* processes, some of which are themselves *intentional.*

Note that Fig. 5 does not show the "environment" of Fig. 3. The "world models" of man and computer can be considered separately from *the* "world". It is our common "hallucination" that is significant, not its source or accuracy. This is significant because it emphasizes that much of the "world" relevant to discourse is not the hard empirical world of the physicist, but rather the constructed, imagined world of artist, novelist, and poet—the world of concepts and possibilities that we ourselves fabricate. We can speak of "unicorns", "beauty" and "god", as readily as we speak of "length", "atoms" and "physical dynamics"—perhaps more readily because myth, emotion and faith are buried deeper in human nature than is our veneer of numeracy and measurement.

Part of this conceptual structure is of course, Karl Popper's "third world" of knowledge and civilization. A world that exists in books, minds and society, but need have no physical counterparts—very much a linguistic world, but only one part of the entire linguistic world since there must be many aspects of our conceptual frameworks which Sir Karl would not wish to allow as part of his "third world".

"Artificial intelligence" research has during the past decade become primarily concerned with the questions that arise in considering Fig. 5—how do we structure the "world model" within the computer—how can man and computer communicate in "natural language"? Hence the close links of this research area to that of cognitive psychology—the one remaining component is the internal world of the mind of man himself.

The direction of effort means that most work in AI is very much narrower in scope than the title might imply. Intelligent interaction with the world does *not* imply that either the "intelligence" or the "world" should be that of man. Surely one dream must be that we can make machines *more* intelligent than ourselves, or with *differing* intelligences from ourselves. But here we are back to the realm of science fiction speculation, for it is not at all clear what this means or, if achieved, that we could either comprehend the achievement or communicate with the resultant machine. Thus, the objectives of research on "man–machine communication" and that on "artificial intelligence" are bound to be closely related for both historic and fundamental reasons.

A new paradigm

Reading back over this essay I am conscious that it does much to whet the appetite, but little to satisfy it. Perhaps that is as it should be, for it is always the questions that are dynamic and fashion the future. The answers once we have them are soon assimilated into the architecture of the "third world" and, by their familiarity, become unnoticed. Let me finish, however, by suggesting one "paradigm-shift" that I believe could occur in the near future and be very significant for many of the problems discussed.

"Programming" is a new term, entering our vocabulary since the mid-forties. We have erected and refined a career structure around computers that involves the "system analyst", "programmer" and "user". I suggest that the distinction is without a difference, or, at least, that any differences are themselves artefacts that must go because they are increasingly counter-productive. A computer is a universal machine, capable of actualizing any defineable algorithm. To "program" a computer is to communicate to it the algorithm in such a way that it can implement it.

By force of habit we have come to regard "programming" as a human cognitive skill whereby we translate the specification of the systems analyst into a computer program that will satisfy the end-users. The concept of the *virtual machine* has already relativized this distinction—for me as a micro-programmer, you as assembler writer are the end-user, yet you regard yourself as providing an "intermediate language" for the writers of several compilers who themselves are providing virtual machines for FORTRAN, COBOL, ALGOL, etc., programmers and so on, through a chain of "library programs", "application modules", ultimately getting to the end-user who may well be a "virtual machine" to someone else, e.g. a stock-exchange dealer providing a "service" (a historic term for a virtual machine!) to clients who are themselves institutional investors providing a service to further clients.

At every level of the virtual machine we tend to provide "options", means whereby an upward flow of control can vary the behaviour of the machine above. We may call them "status-words", "instructions", "parameters", etc., but fundamentally they all have the same function—to vary the operation of the virtual machine from the user level. I suggest that a shift to this broader concept of "programming" will occur—is already occurring. Psychological studies of programming and debugging are becoming increasingly concerned with the broader context of these activities as human cognitive, problem-solving skills. Database systems that offer "query-by-example" are giving the end-user substantial program control. Much of artificial intelligence research may be seen as essentially developing new means for programming computers. A "self-organizing" system is one which programs itself by interpolating a current state of affairs and a

desired state. A "learning" system is one which programs itself by storing the results of exploratory experience.

It is this more general concept of "programming" that may break down some of the barriers we have erected in the past that stifle progress now. In particular, for man–computer communication we should be prompted continually to ask, "what is the communication for—what purpose does it serve?", and attempt to answer that question in these broader programmatic terms. For example, the key problem with a visual display is no longer the technical mechanism for obtaining a picture at low cost, but rather that of *controlling* that mechanism to get not just the *right* picture but one whose construction and change matches our conceptual representation of the structure in the picture. Picture manipulation involves choice and, wherever there is choice, there is a program which has to be induced by interaction. Recent progress in graphic languages, particularly those suited to both picture analysis and picture generation, suggests that the programmatic aspects of interaction must come first. Until we have a language that can enable man and computer to communicate naturally about pictures, we do not have graphics, only the capability to reproduce pictures.

Conclusions

Well, "What next?"—clearly the steady incremental but rapid progress that has borne computer technology so far so fast will continue. Next is smaller, faster, cheaper, more reliable and so on. But also there are qualitative changes apparent which must come to fruition before the promises of the technology itself may be realized.

If I speak of "mutual respect" between man and computer it may seem strange, for what is the computer as seen by the user except a program written by another man. It is indeed understanding of, and mutual respect amongst ourselves, that is needed. The computer can be a vehicle to vent the worst extremes of a "technical-cognitive" authoritarian view upon others. Equally, it can be the means for the "emancipation of cognition" through a participatory partnership.

These may seem grandeloquent terms fit for sociologists but not system engineers. However, this dimension of control and emancipation is apparent in even the simplest interface between man and computer. Man must control the computer—the computer must emancipate man (and emancipation is *not* to give unnecessary and undesired choice, quite the contrary). To allow adequate control we must broaden the concept of programming to that of the flow of control up to a hierarchy of virtual machines. Our need for a world-model inside the computer can then be seen as necessary to allow communication in terms and concepts comprehensible to man.

However, all this involves far greater understanding of ourselves, and it is perhaps to the applied psychology of man–computer studies that we must turn to most in asking, "What next?" . . . Until I know what it is for *you* to understand me, how can I hope to program a machine to do likewise. As in all dialectial questions the resolution is a synthesis—a combination of philosophical, psychological and linguistic scholarship with electronic and computational technologies, under the pressure of commercial requirements for improved man–computer communication. It is in *that* mixture in *that* crucible that the answer to the question, "What next?", is currently being formed.

Bibliography

International Journal of Man–Machine Studies, 1969–1978. Volumes 1–10.