Systems Practice, Vol. 3, No. 1, 1990

Automation Design: Its Human Problems

Sasan Rahmatian[1]

Received July 14, 1987; revised July 18, 1988

This paper examines some of the general features of automation technology and explores certain human issues pertaining to the design of automated systems. After preliminary conceptual clarification of the components of an automated system and the level of knowledge that would make automation feasible, the paper explores the human problems facing designers of automated systems.

KEY WORDS: design; automation; mental activity; system; knowledge.

1. THE RISE OF AUTOMATION: A LOGICAL BEGINNING

The emergence of technology as well as the entire rationale behind it can probably be traced back to the realization on the part of primitive humans that they had a relatively weak body (as compared with many other creatures) but a strong mind. The strength of the mind made up for the weakness of the body by enabling human beings to invent tools that could enhance their physical capabilities. The lever, the pulley, and other "simple machines" are examples that immediately come to mind. Thus conceived, technology flourished for a few thousand years as a muscle-saving phenomenon.

At some point in the recent past, the question must have been raised, "If the human mind can amplify the muscle, why can't it amplify itself?" With this question, a new generation of technology (referred to as cybernetic, automated, etc.) was born. Mind-enhancing technologists aim at understanding a piece of the human mind (as manifested in a particular mental activity) so well as to discover the exact logic (algorithmic procedure) behind it, hence enabling humans to implant that logic into a machine which would then

---

[1] School of Business and Administrative Sciences, Department of Information Systems and Decision Sciences, California State University, Fresno, California, 39740-0007.

perform that activity for them much as they would have done it themselves, or perhaps even more efficiently (Ackoff, 1974, p. 17).

2. MENTAL ACTIVITY AS AN AUTOMATABLE FUNCTION

How are we to define "a piece of the human mind?" From a pragmatic viewpoint, we must define what the human mind is in terms of what the human mind does, namely, mental activities. Mental activities (or processes) are directed, integrated successions of mental states. Since mental states are definable in terms of functional properties of purposeful behavior (Ackoff, 1972, p. 72), it seems reasonable to conceptualize mental activities in terms of their observable manifestations, i.e., result-producing purposeful behavior. Like any purposeful activity, a mental activity is aimed at producing a result. But it must be distinguished from purely physical activities, for they produce results too. The distinction has to be found in the unique features of the mind itself. We are not going to delve into the intricacies of the human mind and its various functions (cognition, volition, emotion, etc.), which have kept philosophers arguing for thousands of years (Peters and Mace, 1967). Fortunately, we do not have to. All we need to do is to capture those mental elements which allow us to understand how a machine mimics the human mind. In order to understand the nature of automatic design, it is useful to clarify all mental activities as falling into one of the following four categories.

First, there is the *goal-setting* type of activity. This has to do with imaging a situation, usually different from what is actually present, which is either desirable (and therefore ought to be reached) or undesirable (and therefore ought to be avoided). The two key elements are

- imagining (which goes by other names such as envisioning, conceiving, visualizing, picturing in the mind, fantasizing, etc.) and
- values (which is related to such dichotomies as desirable/undesirable, good/bad, acceptable/unacceptable, satisfactory/unsatisfactory, etc.)

Goal-setting is perhaps the most human characteristic of the mind.

The second type of mental activity, which is not so characteristic, is *measuring*. This has to do with sensing an external entity and describing (one or several of) its attributes in terms of a predefined scale. Measuring is more than mere sensing, for it involves translation of what is sensed into predetermined units of measurement. The design of those units is part of what makes measurement a mental activity. The other part is the translation (description) effort, which is surely more than physical. The point is well made in psychology and also in philosophy of science (Hanson, 1958), where a distinction is drawn between "seeing" and "seeing as." Seeing an object does not become a mental activity unless it is seen as. Seeing the full moon is not measurement per se. But seeing the full moon *as* a "luminous disk" does involve a mental act of measurement, however primitive. Likewise, seeing a door is not a mental activity per se, but seeing it as a "closed door" (by an automatic system which is programmed to open the door anytime an object reaches within 6 ft of it) is a mental activity. This is a significant distinction because the monitoring mechanism in an automated system always measures a particular aspect of the object being monitored and hence sees it as.

Comparing two things is the third type of mental activity. In actuality, what we compare (in automatic design) are not two objects, but the actual state of an object vs its desired state. More specifically, we compare what should be (as produced during the imagining phase) with what is (as provided by the measurement phase). The comparison takes place in such a way that the outcome is either "yes" or "no." For such a concreteness to be achieved, it is obviously necessary that we should have already defined "what should be" and "what is" in very specific terms. Otherwise, the comparison would not yield an unambiguous response.

The last type of mental activity has to do with *deciding* (choosing, selecting, etc.). If the outcome of the above comparison is "yes" (i.e., what should be coincides with what is), then the decision will be to maintain the status quo. Otherwise, if the desirable and the actual do not coincide, then the decision will be to change the status quo. But what should be the nature of the change? In automatic systems, the nature of the change is predefined. For example, in very simple systems with only two states, the nature of the change would be to switch to the other state (from the one the system is currently in). In more complicated systems, more comparisons may be required before the specific nature of the change is determined.

The concept of automation expounded above is not new. Others (for example, Litterer, 1965, p. 237) have put forth essentially the same idea. Still, there appears to be an overlap between automation and mechanization (Ginzberg, 1982). The reason for confusing automation with mechanization may be that automation involves the use of machines as substitutes for human labor. However, the equating of the two is unfortunate in that mechanization (the use of machines) can take place without true automation. An "automatic" garage door opener is really not an automated system because it is a labor-saving device rather than a mind-saving device in the sense of the four-component model described above.

Perhaps an "automatic" garage door opener is so called for marketing reasons. Or maybe because it is electronic and high-tech. There is indeed a popular tendency to equate automation with electronic high-technology, and vice versa. This is quite unfortunate in that such an equation overshadows the ingenuity of some of the earlier automated systems which were completely mechanical (such as Watt's steam engine), while unduly glorifying those electronic technologies which are really an extension of the human body and not of the mind.

Mechanization (i.e., the utilization of mechanical tools in performing work) may be the first step in automation, but it is surely not the essence of automation. If it were, then the distinction between (say) a typewriter and a word processor would disappear. Yet one is a case of mechanization, while the other is a case of automation. Of course automation itself has various levels, depending on the number of mental functions in performing a job that are automated. Word processors which are automated with respect to (say) only return-at-the-end-of-the-line are obviously not as advanced as those which are automated with respect to other functions as well (such as right-margin justification, spelling correction, page feed, etc.). One of the early exponents of automation commented on the nature of automation as follows:

> (Automation) differs from mechanization in the very way it regards the problem of production. Automation requires us to view the production processes as an INTEGRATED SYSTEM and not as a series of individual steps divided according to the most economic distribution of human skills—or even of individual machines. (Diebold, 1955)

Perhaps it would be misleading to suggest that there is a simple dualistic dichotomy between mechanization and automation. Rather, various researchers have suggested that there is an entire spectrum of levels between the two. Of the taxonomies addressing this issue, one of the earliest is that developed by Bright (1958) in which a 17-point scale spans the gamut from manual work at level 1, to the capability to anticipate action required and adjust it at level 17. Similar taxonomies have been provided by Amber and Amber (1962), in which there are 10 levels, and also by Khandwalla (1974), Tracy and Azumi (1976), and Marsh and Mannari (1981).

3. KNOWLEDGE FOR AUTOMATED DESIGN

In constructing the above model of automated design, we adopted a pragmatic approach toward understanding the human mind. We continue this pragmatic attitude, trying to understand the variety of types of knowledge which produce automated design. From a pragmatic standpoint (see Ackoff, 1972, pp. 46, 47, 144), knowledge is to be distinguished from information and understanding. Information has to do with factual statements which are descriptive of entities or events and which produce a change in a person's probabilities of choice. Understanding has to do with the reasons why things are related the way they are. Knowledge, in contrast, is taken in the sense of how to make things happen according to predefined criteria. In more detail, we want to distinguish among the following three:

- *Know that* (know who/know what/know when/etc.), which is information. More formally, let us define "information" as the symbolic representation of an attribute of an entity or event, useful only when communicated to a purposeful system for reducing uncertainty about reaching its desirable states or avoiding undesirable states. Knowledge of a factual nature, although it may be false, is information.
- *Know why*, which is understanding. This type of knowledge, unlike the above type, is theoretical and hypothetical (i.e., hypothesis-based) because its aim is to explain why things happen the way they do. One explanation may be more powerful or comprehensive than another but not more true in a factual sense. Explanations are statements of reasons why relationships among certain events exist. The events they try to explain may be very real, but the relationships connecting them do not belong to the reality of the same order.
- *Know how*, which is knowledge in our sense of the word, namely, ability to pursue a course of action that produces a desirable outcome. Since this paper deals with the know-how sense of knowledge, it becomes necessary to explore the various levels of know-how in order to gain a deeper understanding of the role of knowledge in automation.

As far as automation is concerned, there are three distinct levels of knowledge (i.e., know-how) corresponding to three different ways of doing things. Their correspondence is explored in what follows.

One way of performing any task is doing it *randomly*. With this method, the outcome of action is completely determined by chance, rather than by any reasoning or strategy. As an example, most people tend to buy lottery tickets following this method. The random method represents the lowest level of knowledge, namely, *ignorance*. At this point in time, our knowledge regarding how to predict earthquakes, how to cure cancer, how to control the climate, and how to do a host of other things is at this level. Random selection is the only "method" that could be available under these circumstances.

A second way of performing a task is doing it *intuitively*. This means doing it based on some method evolved from many trial/error experiences but which (method) is hard to articulate explicitly in a step-by-step manner. Hence, the method, although there is one, has remained tacit and intuitive (Polanyi, 1966). This corresponds to a level of knowledge we may call the "*know-how*" level, where we are dealing with things we know how to do, but we do not quite intellectually comprehend exactly how we do them. Examples are talking, riding a bicycle, and recognizing a familiar face. We can talk without needing to know the exact procedure for vibrating our vocal chords, activating our muscles, and moving our jaws. The same with a personnel manager who may be a recruiter of capable employees but who may not be able to articulate clearly the exact procedure he follows in making such effective decisions.

A third possible way of performing a task is doing it *systematically*. This means doing it based on a method that can be stated explicitly and unambiguously, so clearly in fact that any person or any machine performing it would produce the same result. This corresponds to a level of know-how that may be termed "*self-knowledge*." The reason for this peculiar terminology has to be found in the fact that with this type of knowledge, not only do we know how to do things, but also we are aware of the exact procedure we follow in carrying out those tasks, so much so that we can objectively express and communicate them to others. The emphasis on "self" derives from the focus of attention being not merely on performing the task, but on catching oneself (as it were) doing it. The self-awareness thus achieved is the real key to making possible the transition from intuitive understanding to systematic method. Examples of systematic procedures are payroll processing, balancing a checkbook, and well-structured games such as chess.

While the practical implication of ignorance (lack of knowledge) is suffering (being at the mercy of uncontrollable events), the practical implication of know-how is that with experience and repetition, we become better practitioners and get closer and closer to the self-knowledge level. The practical implication of self-knowledge is automation, as explained in the section on the rise of automation. If we discover the exact procedure behind performing a task, then we can code those procedures and communicate them to a machine which would perform the task for us.

It is worthwhile pointing out that there is also a fourth level of knowledge, which may be termed *self-automation*. Whereas the third level dealt with understanding the logic behind this or that process so that automation could be accomplished, the fourth level deals with understanding the logic behind automation itself as a process. In other words, on the fourth level, we try to comprehend the process and nature of automation so well that we can automate the automation process itself. The variety of generations of programming languages is indicative of this phenomenon. Specifically, a third-generation language (such as COBOL) allows us to write code for performing a particular task. Some fourth-generation languages (Martin, 1986) are capable of automatically generating that third-generation code. Thus the object of automation becomes automation itself rather than a specific mental task. In the limiting extreme, we see the case of a completely automated robot which may be indistinguishable from a human and which, through many layers of software, can translate human-like messages into machine languages. Some (i.e., mechanists—the advocates of the philosophical doctrine of mechanism) tend to go as far as saying that that robot is really us!

4. HUMAN PROBLEMS OF AUTOMATION DESIGN

"Automation design" and "automated design" are different. The former has to do with the process of automating a mental function, while the latter refers to its product. The area of the human problems of automated designs (i.e., products) is well researched. However, little work has been done on understanding the human problems facing those who design such products. In this vein, we will then understand "automation design" to refer to the process of developing an automated system, rather than of using its product.

So far in this paper, we have defined automated design to consist of four elements: goal-setting, measuring, comparing, and deciding. After discussing the general form of the logic behind automated design, we examined the epistemological conditions under which this logic can be discovered. Specifically, we argued that the logic is discovered only if and when we go beyond mere know-how and analyze the nature of the know-how itself. If we can know how the know-how works, then this would mark the beginning of automation for the particular function under consideration. We now reexamine these issues in the case of automated design taking place within its real-world context, namely, a social system context.

By a social system context, we mean situations where purposeful humans are working together to accomplish certain goals. Right off, it must be admitted that ALL automation works within a social system context. Automatic systems are designed *by* human beings *for* human beings. Moreover, they are maintained, serviced, and operated by human beings. A human environment surrounds all automated design. Hence, the following issues are not special cases of more general situations but are universally applicable in themselves. Speaking of automation in a social system context, we are not implying that there is any other context within which it can be examined meaningfully. Rather, we are acknowledging that the context of all automation is the human context and that obstacles (to effective automation) due to specifically human factors must be identified and analyzed. These obstacles exist in both the design phase (where the four elements are determined and integrated) and the operation phase (where the system is activated and, depending on the input provided, produces certain outputs). The following discussion focuses more on the human problems typically faced *during* the design phase rather than *after* the system has become fully operational.

It may be more clear to make the following points in reference to an illustrative example. There are relatively simple automated systems, and then there are somewhat more complex automated systems. How complex an example do we need? What, in the first place, do we mean by a complex automated system? Before providing an illustrative example, we need to delve deeper into the question of complexity.

In a recent conceptual analysis, Flood (1987) disassembles complexity in terms of systems and people. Systems are complex to the extent of the number of their parts and the relationships among these parts. Further support for this view comes from Klir (1985), who states the above as the first general principle of systems complexity. Flood goes on to explain people complexity in terms of interests, capabilities, and perceptions.

Flood's distinction of systems vs people complexity is highly relevant to this paper. Systems complexity is of a *technical* nature. Engineers deal with highly complex systems which, from a people perspective (say, in terms of project management), may not be that complex. On the other hand, there are examples of people complexity (say, in the area of job satisfaction or employee morale) which are not necessarily accompanied by technical complexity. This is because people complexity is of a *human* nature. The relevance of Flood's ditinction to this paper is that our main concern here is with human (people) rather than technical (systems) complexity. In what follows, we explore this relevance in more detail.

Of the four elements of automated design discussed earlier, three correspond to the three ingredients of people complexity listed by Flood. Specifically, goal-setting relates to human interests, measurement corresponds to human perceptions, and deciding incorporates capabilities. Following these relationships, we define the complexity of an automated design in terms of

- the extent to which the various users have different interests in the system's function,
- the degree to which different perceptions result in divergent measurement methodologies, and
- the extent to which there is lack of unanimity about how to enable the system to act under various conditions.

An example of a very simple automatic system is a word processing package which provides for right-margin justification. A slightly more complex system would be a word processing package with automatic spelling correction. The complexity lies in whether, for example, "fianc" must be corrected to "fiance" or "fiancee." A more complex example would be an automatic teller machine (ATM) which dispenses at most $200 per day to each customer. The goal may be questioned and challenged (i.e., why not a $300 maximum per day?) and the measurement may be erroneous (it may dispense nine $20 bills, while registering the transaction as involving 10 such bills). A still more complex example would be that of an automatic system that approves or rejects new applications for credit cards. A case in point would be the one used by the Hertz Corporation, which has generated over 7 million credit cards for people throughout the world (Parker, 1984). When a potential customer fills out an application for a credit card, the information is then fed into the computer, which scores the application on factors such as credit references, income, and current employment. If the score is high enough, a credit card is issued; otherwise, the computer system automatically types up a letter of denial or indicates that more information is required. This is a relatively complex level of automation in that all four elements can easily become subject to controversy.

Historically, automation started with highly mechanical systems, then entered the human organization via mostly clerical tasks. With the emergence of decision support systems, expert systems, and artificial intelligence, automation is penetrating the social system profoundly. More and more human activities are being automated, with farther and farther impacts on human beings. As the complexity of automation is thus increasing, it would be appropriate to select an example, for analysis purposes, that touched human lives in an intimate manner. Let us take the case of an automated system which receives as input data about the performance of employees in an organization and provides, as output, a letter addressed to each employee informing that person of his or her employment status (contract renewed or terminated) for the following year.

I am personally not familiar with any such system being operational at the moment. But this does not detract from the appropriateness of the example for several reasons. First, it has many features common with similar systems, such as the Hertz credit card system discussed earlier. They both receive multidimensional data about certain people and produce a binary decision of importance to those people.

Second, the proposed example takes our attention away from the mechanical (hardware) details and directs it toward the logic of automated design, where it belongs. Automation is less a matter of mechanized decision-making than of structured production. Therefore, automation can be viewed as a philosophy of production (Diebold, 1955). In the limiting case, where every behavior in an organization becomes governed by rigid structures, that organization turns into a bureaucracy, with every human member of it having become a cog in the wheel. As such, automation is also a philosophy of organization.

Third, the selected example is appropriate because it foreshadows the direction in which automation is taking us. We can already witness many vital decisions being made in highly structured, automated ways. For instance, many universities, faced with thousands of faceless applicants, follow a strict "formula" in figuring out who to admit. If we realize that decisions of such importance (to candidates) are being made automatically, then our hypothetical example may appear less farfetched. We now proceed to the task of deriving, from the framework established, the human problems of automation design.

Goal-setting, the first mechanism of automation, becomes problematic to the extent that the human element enters the picture. The specific human element being referred to here is that, as has been argued elsewhere (Rahmatian, 1985), goals in social systems exhibit a distinctly hierarchical nature. A goal is, at one and the same time, an end with respect to lower goals, and a means with respect to higher ones. Hence every goal (with the exception of the lowest and the highest) has a dual nature—it is both an end and a means. This chain of ends and means makes it difficult to speak of "the goal," for "the goal" typically turns out to be sandwiched between a higher and a lower goal. Automation does not begin until these goals find a measurable expression reflecting their totality. Let us refer to our illustrative example to see what this means.

In our hypothetical case, we have a system which makes a binary decision (keep vs dismiss the employee) based on certain characteristics of the employee. What is the goal here? It appears to be something like, "to retain only competent employees." This goal is sandwiched between two other goals.

- The higher goal is "to have a competent work force," itself a means to higher ends which probably have to do with the mission of the organization. A subgoal of the stated higher goal would be "to hire competent employees."
- The lower goal is "to retain only those employees whose Overall Performance Score exceeds X," where X is some critical cutoff point. It is clear that automation cannot be achieved in the absence of such concrete operationalization.

So far, so good! But, two critical questions remained unanswered: How is the Overall Performance Score (OPS) to be operationalized? and To what value should the cutoff point be set?

The academic response would be that these two issues are part of what operationalizing the qualifier "competent" involves. This is correct but, in reality, not as trivial as it sounds. The first issue, the choice of criteria which become part of the measurement of OPS is, in practice, fraught with controversy and conflict. Typical criteria would be quality of output, quantity of output, interpersonal skills, growth potential, communication skills, etc. Is this list complete? Maybe "perceived loyalty" should also be added, for how useful can a competent employee be if she views her experience with a firm only as a stepping stone to a better position in another organization? It can be argued that many other factors must be considered too. Even if the given list is deemed complete, then the question of relative weight remains. How significantly should each criterion figure in the operational definition of OPS?

The second issue is, How should the critical cutoff point be determined? Elaborate schemes may be designed for answering these questions, although in practice answers seem to be dictated more by convenience and expediency. In the case of the ATM mentioned earlier (in connection with dispensing at most $200 per day to each customer), it is hard to imagine that "scientific research" yielded the perfectly round number of 200! We know it did not for another reason: that this number keeps changing every so often.

Difficulties in establishing a consensual measure of performance are rooted in the diversity of human goals. What constitutes success or failure is often a matter of controversy. Different actors in a social unit, looking at the same reality, may see it as desirable or undesirable, depending on their perspective. Being rooted in many complex factors (personality, culture, past experience, social pressure, individual motives, etc.), these diverse perspectives are not easily reconcilable. The point has been beautifully illustrated in a cartoon that shows two derelicts, one happy, the other sad. The happy one is saying to his friend, "You're a failure because you wanted to be a tycoon, I'm a success because I wanted to be a bum!"

Measurement is the second element in automated design. The human problems of measurement are to be found in its subjectivity. On the one hand, automation is inconceivable apart from purely objective measures. The reason for this lies in the fact that, in principle, automated functions must be performable by machines, and machines work on objectively expressible principles. They cannot handle human subjectivity. On the other hand, achieving complete objectivity is not human; it is superhuman. Take our automated contract renewal system again. How is the OPS to be measured for each employee? It is to be measured based, ultimately, on its ingredient criteria (quality and quantity of output, interpersonal skills, etc.). But how are these to be measured? Some criteria (such as the quality of output) may be measured objectively. But many other criteria remain for which it would be hard to imagine completely objective measures, such as the quality of output (where relevant) and perceived loyalty to the organization. The inevitable presence of a subjective element in measurements involving human judgment seriously weakens the prospects of their becoming automated.

The third element of automated design discussed in this paper is comparing. This may sound very straightforward. After all, it may be argued, if there is agreement on what the goal is and how to measure it, does it not follow that comparing the "is" and the "ought" must be a trivial task? Again, from an academic standpoint, the answer would be yes. But when the human element enters the picture, complications arise. Recall that comparisons between the actual and the desired, in order to become part of automated design, must yield an unequivocal yes/no result. Fuzzy outcomes such as "maybe," "to some extent," "mostly," and "not quite" are to be ruled out since they are too ambiguous. Moreover, as our model indicates, the yes/no outcome of the comparison process must be conveyed to the deciding mechanism which, based on preprogrammed decision rules, would determine the appropriate course of action.

To understand how the human element enters the picture, let us refer to our illustrative case. Suppose the threshold point is set at 80, i.e., only those workers whose OPS is equal to 80 or greater will be retained. Now suppose a certain employee's OPS equals 79.99. Strictly speaking, this person must not be retained. But shouldn't the figure be rounded to 80? After all, given the subjectivities inherent in measurement, it is quite likely that an error of 0.01 was generated by one of our measuring instruments. No measurement process is 100% reliable. Admitting this possible error and subsequently rounding off the OPS to 80 would entirely reverse the decision, of course.

As a solution to this problem, it may be proposed that a margin of error (say 2) be established. This will not work. First, it runs counter to the fundamental nature of automation, which requires the comparison process to yield a yes/no answer. Second, with the proposed approach we are only pushing the problem further down. What about the person whose OPS now equals 77.99 (again, 0.01 point short of the new threshold point of 78)?!

In the fourth element of automated design (i.e., deciding) we find the greatest challenge to designing automation within social systems contexts. The two fundamental issues are as follows:

- What is the totality of possible decisions (system states) that the deciding mechanism can choose from when changing the status quo? And how do we know if the repertory is really complete?
- Even with the most complete repertory of possible states to switch to, how do we know which state to switch to from the one the system is currently in?

The first issue takes us to the very heart of design as we understand the term in systems theory (Weinberg, 1982), namely, the generation of alternatives. This is the most imaginative step in automated design. Imagining every possible contingency that may arise and devising counteractive responses for each is the kind of task that calls for mental discipline in addition to creativity. As applied to our illustrative case, the question becomes, Why a binary decision (retain vs dismiss employee)? There are other possibilities such as extending a probationary contract. This itself is really a set of alternatives, for there can be probationary contracts of various lengths (1 month, 2 months, etc.).

The second issue is, What decision rules govern the assignment of actions to situations, and what assures us that these rules will really work? The derivation of decision rules is guided by the statement of the goal. For instance, going back to our example, if OPS has been operationalized, and if the cutoff point has been set to 80, then the decision rule would be

```
IF OPS ≥ 80 THEN
    Retain Employee
ELSE
    Fire Employee
```

In automated systems with binary outputs it is not difficult to derive the decision rules from the goal statements. However, as we argued, binary systems do not quite capture the complexities of reality. With more than two possible states for the system to have, the multiplicity of cutoff points parallels the multiplicity of goals in complexity. Therefore, the same points made about the human problems of goal-setting apply here again.

5. CONCLUSION

If and when automated systems produce results that are undesirable, those results can be overriden by humam beings. In overriding automated systems, people use their intuition and sense of judgment. They do not override automated systems systematically (i.e., based on preestablished criteria), for if they did, then the overriding function itself could become automated, and the automatic overriding system may itself be overriden, ad infinitum! In this endless regression, at some point, the human judgment which operates outside the automation sphere will come into play. It is precisely this situation that we have characterized as the human element.

REFERENCES

Ackoff, R. L. (1974). *Redesigning the Future*, Wiley, New York.

Ackoff, R. L., and Emery, F. (1972). *On Purposeful Systems*, Aldine Atherton, Chicago.

Amber, G. H., and Amber, P. S. (1962). *Anatomy of Automation*, Prentice-Hall, Englewood Cliffs, N.J.

Bright, J. R. (1958). *Automation and Management*, Harvard University Press, Boston.

Diebold, J. (1955). Applied automation: a practical approach. In Heyel, C. (ed.), *John Diebold on Management*, Prentice-Hall, Englewood Cliffs, N.J., p. 52.

Diebold, J. (1959). *Automation*, National Planning Association, Washington, D.C.

Dreher, C. (1957). *Automation*, W. W. Norton, New York.

Flood, R. L. (1987). Complexity: A definition by construction of a conceptual framework. *Syst. Res.* **4**, 177–185.

Foster, D. (1963). *Modern Automation*, Pitman, London.

Ginzberg, E. (1982). The mechanization of work. In *The Mechanization of Work*, W. H. Freeman, San Francisco.

Hanson, N. R. (1958). *Patterns of Discovery*, University Press, Cambridge.

Khandwalla, P. (1974). Mass output orientation of operations technology and organizational structure. *Admin. Sci. Q.* **19**, 74–97.

Klir, G. J. (1985). Complexity: Some general observations. *Syst. Res.* **2**, 131–140.

Litterer, J. A. (1965). *The Analysis of Organization*, Wiley, New York.

Marsh, R. M., and Mannari, H. (1981). Technology and size as determinants of the organization structure of Japanese factories. *Admin. Sci. Q.* **26**, 33–57.

Martin, J. (1986). *Fourth-Generation Languages*, Prentice-Hall, Englewood Cliffs, N.J.

Parker, C. S. (1984). *Understanding Computers and Data Processing*, Holt, Rinehart and Winston, New York.

Peters, R. S., and Mace, C. A. (1967). Psychology. In Edwards, P. (ed.), *The Encyclopedia of Philosophy*, Macmillan, New York.

Polanyi, M. (1966). *The Tacit Dimension*, Doubleday, New York.

Rahmatian, S. (1985). The hierarchy of objectives: Toward an integrating construct in systems science. *Syst. Res.* **2**, 237–245.

Tracy, P., and Azumi, K. (1976). Determinants of administrative control: A test of a theory with Japanese factories. *Amer. Sociol. Rev.* **41**, 80–94.

Weinberg, G. (1982). *Rethinking Systems Analysis and Design*, Little, Brown, Boston.