28 May 1971, Volume 172, Number 3986 SCIENCE

Computer Graphics as an Aid to Learning

Computer graphics can facilitate the rapid learning of an important cognitive skill.

Kenneth R. Hammond

The psychology of learning has traditionally concerned itself with the problem of the *acquisition* of knowledge and skill. The focus here is on the problem of learning to *apply* knowledge already acquired. The role of computer graphics as an aid to learning this important cognitive skill is shown to be critical. The context chosen to illustrate the empirical significance of the argument is that of medical knowledge and diagnostic skill.

Knowledge versus Skill

In line with tradition, virtually all efforts to improve medical education attempt to provide more efficient transmission of medical knowledge. Thus, for example, the latest technological development, computer-aided instruction (CAI), is rapidly being developed by medical educators (and others) for that purpose. The principal reason for the rapid development of CAI programs for teaching knowledge is that a paradigm for such programs is already available—CAI programs simulate the traditional teacher-student model. Developmental, not basic, research is all that is required to make CAI programs useful instruments for conveying knowledge (*1*).

It is quite another matter as far as the application of knowledge is concerned. At this time, there is no established method (or model) for teaching the student to apply his knowledge. There is, for example, no reliable, efficient method for teaching a medical student to develop competent judgment in diagnosing a patient's illness. Such a cognitive skill is difficult, if not impossible, to learn from books. It is also difficult to learn from teachers. Moreover, it requires considerable practice and, therefore, is extremely time-consuming and expensive for all concerned. Worse still, these time-consuming efforts often prove frustrating, for teacher and student seldom know why a student's progress is slow; the source of the student's difficulty is hard to locate. This problem has recently begun to attract substantial research interest (*2*).

Model of the Judgment Process

Diagnosis is a specific instance of the more general process of judgment, which involves the integration of information conveyed by several cues (items of information about some state of affairs not immediately apparent to the judge, or diagnostician (*3*). This process is represented schematically in Fig. 1.

Three premises are included in our conception of the diagnostic task.

Premise 1. There is irreducible uncertainty in diagnostic tasks. *Implication*: The student cannot learn—cannot discover—an infallible rule for organizing and evaluating information. Because the task includes irreducible uncertainty, it requires the exercise of judgment, as well as knowledge.

Premise 2. Diagnostic tasks require integrating cues of various degrees of uncertainty. *Implication*: In exercising his judgment, the student must learn to assign differential weights to various cues, depending upon the degree of uncertainty of the cues.

Premise 3. Cues will differ in the form of their functional relation to the state to be inferred. *Implication*: The student must learn to integrate data from nonlinear relations as well as linear ones.

All of these premises are included in a type of learning task that psychologists have used to study human judgment (*4*)—the multiple-cue probability learning (MCPL) task. Because such tasks require judgmental learning analogous to the learning needed to make diagnostic judgments, it will be useful to relate what is known about MCPL to diagnostic judgments.

Multiple-Cue Probability Learning

Typically, studies of MCPL require the subject to learn more or less blindly by means of outcome feedback, sometimes called reinforcement. That is: (i) the subject is told little or nothing about the structure of the task; (ii) the materials are neutral in that they rarely carry meaning, and, therefore, past experience is of no use; and (iii) learning proceeds as a consequence of the subject's being "reinforced," that is, being told (in one form or another) the correct answer on each trial. There are, of course, variations in this theme, but the essential elements hardly vary.

There are two points to be noted about the above paradigm. First, no form of feedback other than outcome feedback has ever been employed within the traditional learning paradigm—involving MCPL or other types of learn-

The author is professor of psychology and director of the Research Program on Human Judgment and Social Interaction, Institute of Behavioral Science, University of Colorado, Boulder 80302.

ing. Second, under these conditions, learning is ordinarily slow and inefficient (*4*). The learning is also "stupid," in the sense that learners, when they do learn, are frequently unable to give a clear explanation of what it is they have learned (*5, 6*). They would, therefore, be poor teachers.

Slow and inefficient learning of such tasks is due to the nature of the feedback provided. Outcome feedback can be of little help to the learner, for two reasons: first, providing the correct answer (outcome) after having made a judgment is virtually useless, since outcomes are related to the cues in a complex, multidetermined, and uncertain way. There is no simple, rule-bound connection to be discovered, for the same answer can be produced by various combinations of cues. Conversely, identical combinations of cues can provide different answers. Second, the irreducible uncertainty in the task requires a long series of trials in order to distinguish between that which is nearly regular and that which is wholly accidental. There is no alternative to a long series of trials, if the learner is limited to outcomes as feedback.

In principle, then, two things will be required of any attempt to improve MCPL. First, "stupid" learning that is based only on outcome feedback and repetition will have to be avoided; the learner must be able to avoid the frustration engendered by not knowing what it is that he is doing wrong—or right. Second, information will have to be provided for the learner so that he can quickly grasp the essential characteristics of the task, despite its irregularities. He will need easily perceived information about the task and about himself, information that goes far beyond the comparison of, for example, the correct judgment and his own judgment. He needs to know why these are different, or the same.

Improving Diagnostic Skill

On the argument that the MCPL paradigm provides a reasonable simulation of diagnostic tasks, the above conclusions may be applied to the process of learning to make diagnostic judgments. Just as "stupid" learning that is based on outcome feedback and repetition will have to be avoided in MCPL, it will have to be avoided in learning diagnostic tasks. And just as the learner in the MCPL task will need to be pro-

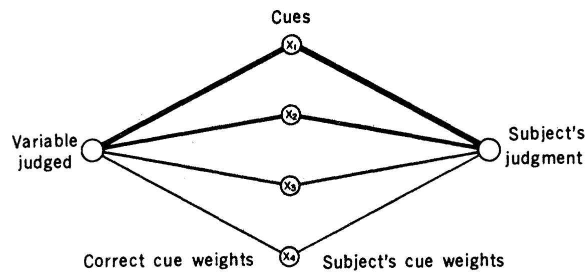

Fig. 1. Schematic representation of judgment process, indicating that the subject (judge) must integrate differentially weighted cues of various degrees of dependability. (Thickness of line indicates varying degrees of dependability.) The illustration indicates that the different cue weights in the task have been matched by the cue weights in the judgmental system of the subject.

vided with information as to why his judgment and the correct answer are not the same, so also will the medical student need this information.

But it is precisely this information which the teacher finds difficult to convey to a student. The teacher's skill in integrating information based on uncertain data has been accrued over a long period of practice, and he is apt to be vague about the basis of his own judgments. He cannot be certain about the nature of his judgments, for only subjective introspection, notoriously faulty, provides information about one's own judgmental processes. Further, the teacher will be unsure about the student's cognitive processes in arriving at a judgment, and he may attribute the student's error to the wrong cause. Moreover, the student, relying on his own introspective processes, may (correctly or incorrectly) disagree with his teacher's observations of him—covertly if not overtly. Communication under these circumstances is more likely to produce cooperative delusion than accurate understanding.

Hoffman, Slovic, and Rorer (*7*) have provided excellent documentation not only of the fact of wide disagreement among nine medical diagnosticians, but of the several ways in which their judgmental processes differed. The diagnosticians were, of course, quite unaware of the nature of these differences.

It cannot be assumed that experience will increase one's awareness of his cognitive processes. For in quite a different context, Slovic, Fleissner, and Bauman (*8*) found that the more experienced stockbroker was less able to describe accurately how he arrived at his judgments than was the less experienced broker. In a study of undergraduates' judgments about the future socioeconomic growth of underdeveloped nations, Summers, Taliaferro, and Fletcher (*9*) observed that their subjects (i) reported using more cues than they, in fact, did, and (ii) inaccurately described the weight they attached to various cues. Moreover, these authors were doubtful about the ability of their subjects to describe accurately the form (linear, nonlinear) in which they related data to their judgments. Significantly, they point out that "the consequences of this failure of self-report could well include misunderstanding, mistrust, and even conflict."

These results have highly significant implications. If expert judges, whether they be clinicians or stockbrokers, unwittingly mislead the person whom they are trying to teach to achieve similar judgments in the same way they do, the student would be taught one thing while his teachers practiced quite another—and both would be unaware of the disparity. Cooperative delusion may be an empirical fact, and it may extend beyond the confines of formal educational settings. It may be a general human characteristic (*10*).

Requirements for a Solution

If ignorance and uncertainty about one's own cognitive, judgmental processes are indeed at the root of the problem of ineffective application of knowledge, then one solution would be to devise procedures to (i) make explicit the characteristics of a person's judgmental system and (ii) relate these to the characteristics of the judgmental task. The student should be able to compare his functional judgmental system with that required by the diagnostic task. More specifically, he should be provided with an opportunity to compare (i) the differential weights he actually assigns to the cues with the weights required by the task, and (ii) the form of the functional relations between the cues and his judgments with the form of the functional relations required by the task. In brief, the student should be able to compare what should be done with what he is doing. Ideally, he should be provided with a picture of the properties of the task and a picture of his own (cognitive) judgmental system, in terms that will allow him to compare the two.

The critical research question is whether such comparisons, if they could be provided, would enhance learning. Can people make effective use of such information? Can they exercise sufficient control of their cognitive processes to modify them readily and thus learn to exercise better judgment? And can this be done without long periods of practice?

A Study of Cognitive Control

A study was designed to provide three different types of information to learners (*11*). One type consisted of traditional outcome feedback, merely as a baseline control. A second type provided two different amounts of information about the task in two different ways—verbally and pictorially. A third type made use of computer graphics techniques in order to provide the kinds of pictorial comparisons indicated above.

A learning task was constructed to simulate a diagnostic problem. The subject was required to arrive at a judgment that integrated the information provided by three cues. Each cue was differentially related to the criterion (the correct judgment): cue A was correlated 0.8, cue B 0.4, and cue C 0.2 with the criterion. The form of the relation between cue and criterion was not linear; rather, these relations were of an inverted-U shape. Finally, irreducible uncertainty was built into the task; no rule could be formulated which would permit the learner to achieve the right answer on every trial. The learner could not, in other words, become an infallible diagnostician.

On any one trial, the three cue values were presented on a 5-inch by 8-inch card in the form of bar graphs, the height of each bar indicating the value (from 1 to 10) of that cue (see Fig. 2). The subject was requested to interpret

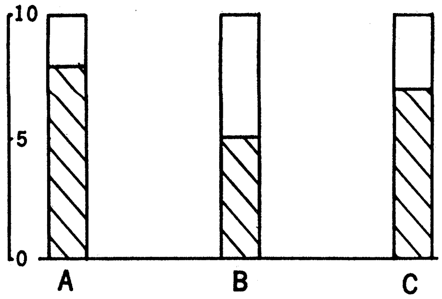

Fig. 2. Example of one display with cue values of 8, 5, and 7.

these three cue values and arrive at his answer (on a scale from 1 to 20) for this particular display. After recording his answer, he was shown the correct answer for that display. Two hundred trials, or displays, were used.

Five groups, each composed of ten college sophomores, were run. Group I was a control group receiving only outcome feedback. The remaining groups received task information after 20 outcome feedback trials as follows: Group II was informed about correct weights pictorially; group III was informed about correct functional relations pictorially; group IV was told both correct weights and functional relations verbally; and group V was provided with both correct weights and functional relations pictorially. All groups received outcome feedback on each trial throughout all 200 trials.

The results are quite clear: information about task properties aids learning (Fig. 3).

As expected, this task was too difficult to be learned by means of outcome feedback alone. The fact that group II did no better than group I indicates that information limited to differential weights and outcomes is not useful if the functional relations between cue and criterion are not simple linear relations. If these relations had been linear, the task would have been far simpler and information about differential weights would have produced learning (*4*).

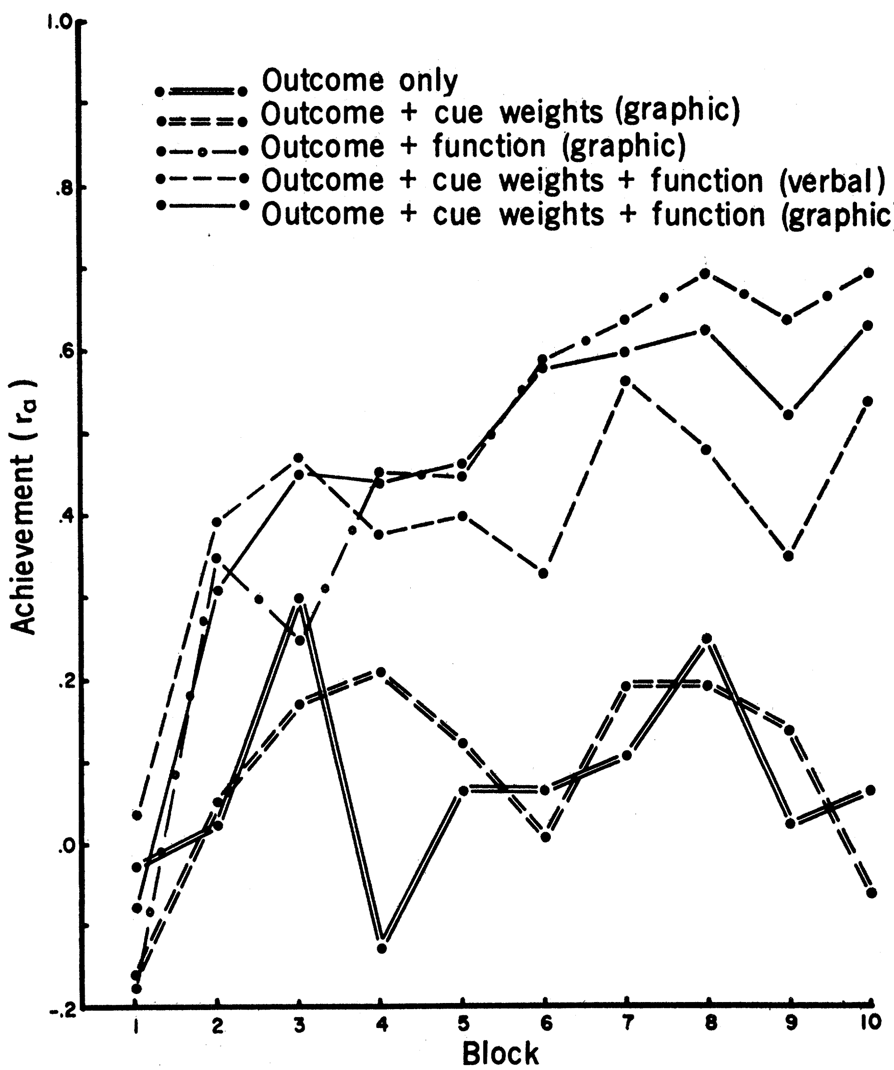

Fig. 3. Comparison of achievement, measured in terms of the mean correlation (r_a), between subject's judgments and the correct answers over ten blocks of 20 learning trials for five groups with different information about the task.

Groups III, IV, and V did learn to apply their knowledge, demonstrating that information concerning functional relations is critical when nonlinearity is involved (*12*). Learning is still slow and inefficient in these groups, however. Subjects approach asymptote gradually (indicating that they do not fully grasp the nature of their errors), and they never reach the statistical limits of achievement (that is, they never reach the potential "ceiling," r_a = .92, of the task).

Computer Graphics

As indicated above, learning should be enhanced if learners can compare pictorial, easy-to-grasp representations of the properties of their judgmental systems with the properties of the task. In an effort to provide such comparisons, a CDC 282 visual interactive display console, interfaced with a CDC 6400 (at the University of Colorado Computing Center), was used to display the appropriate information (*13*).

The subjects are seated before the console, and the cue values are presented on the face (a cathode ray tube) of the console, in the same form illustrated in Fig. 2. Subjects enter their judgments directly into the central (CDC 6400) computer by means of the keyboard on the console. After making a series of judgments, the subject may then, for example, call for information about his judgmental system, his teacher's judgmental system, or information about the task itself. And, of course, he may ask for comparisons. The function of the computer program is to make the necessary analyses of these systems and to provide the appropriate graphic displays in a form intelligible to the learner.

As illustrated in Fig. 4, such displays allow the learner to compare the weights he has assigned to cues with the weights he has been instructed to employ (Fig. 4A). Also, he may compare the form of the relation between his judgment and the cue values with the relation he has been instructed to employ (Fig. 4B). These representations are, of course, imprecise; their aim is explication rather than exactness. (Several other displays not shown here have been developed. For example, a subject may call for a "history" of his judgments in relation to the correct answers. He requests the presentation on the console of specific trial displays, in order that he may consider why he was in error on these particular trials.)

In an initial effort to evaluate the utility of such visual comparison, five subjects were presented with the judgmental task described above. Each subject was first given ten outcome feed-

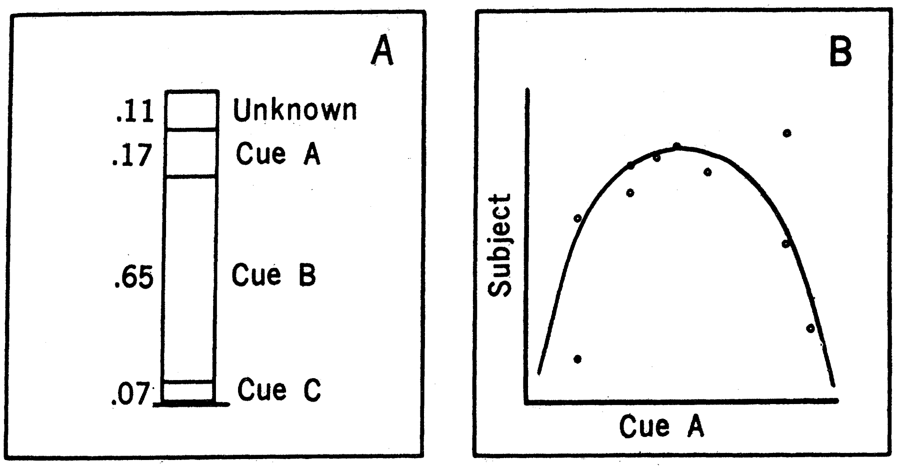

Fig. 4 (top left). Example of two displays presented to subjects by means of computer graphics. (A) Indicates the weight a subject is placing on each cue in terms of a percent of the variance in his judgmental system. The term "unknown" refers to the variance not accounted for by a specific mathematical representation of a learner's judgmental system. (B) Illustrates the nonlinear relation between a cue and a subject's judgment. Each point represents a judgment, and the curve is a least-squares, best-fit line.

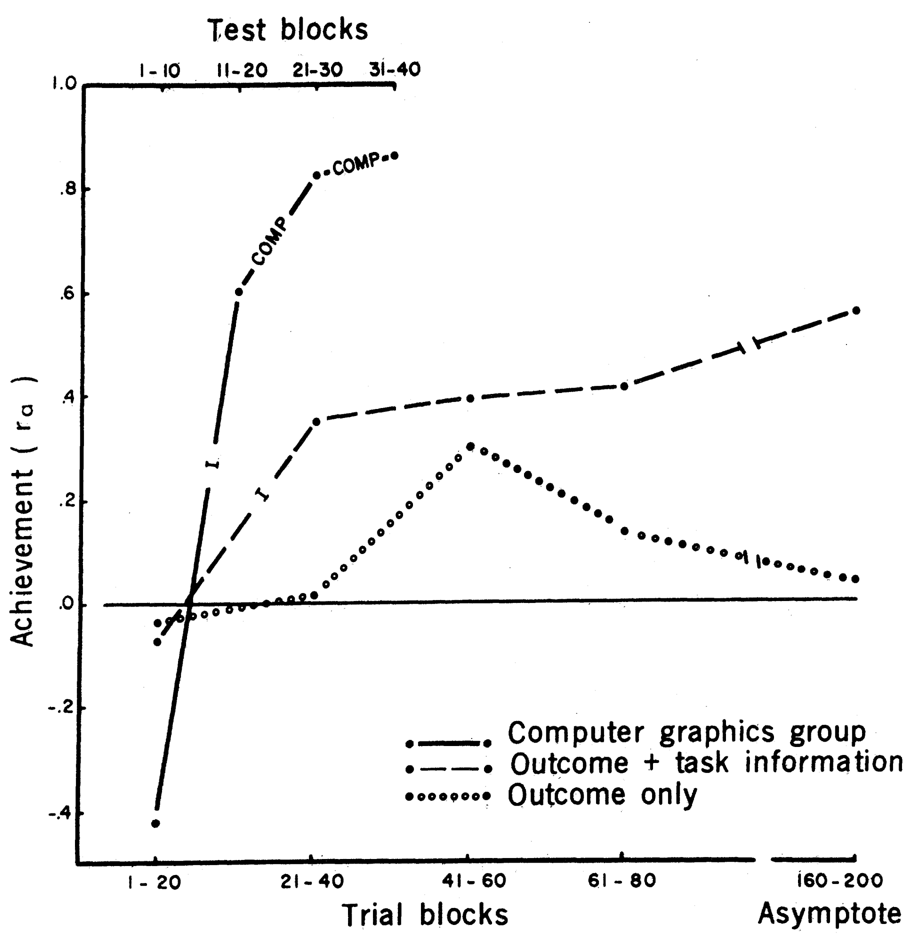

Fig. 5 (bottom left). Learning curve for the computer graphics group, compared with groups III, IV, and V (who learned the task slowly) and group I (who did not learn). The point at which task information was provided is indicated by *I*; *COMP* indicates where comparisons were provided for the computer graphics group. (Data from group II are not included in this graph.)

back trials on the console; as expected, no improvement in performance occurred. He was then given information about the properties of the task, as in group V above. After a set of test trials (with no outcome feedback), the learner was shown, by means of the displays indicated in Fig. 4, (i) the cue-weighting system he had employed in the test trials, and (ii) the functional relations between his judgment and each cue, thus enabling him to compare his actual performance with the performance he intended. He was then given a second set of test trials—a second opportunity (without outcome feedback) to make judgments and, in so doing, to bring his functional judgment in line with the system he intended to develop. After these test trials, he was provided with a comparison of his judgmental system and the task system, and was tested again with another series of trials. No more than three such comparisons were needed before the subjects approximated the statistical limit of achievement.

The results are presented in Fig. 5. They clearly indicate that comparisons of this kind are useful and can lead to rapid learning (*14*). These results are supported by results obtained with a larger number of subjects ($N = 24$) in an earlier study, carried out before computer graphics techniques were available, which provided similar comparisons in verbal terms; similar rapidity in learning was observed (*15*).

Encouraging results were also obtained with four medical students who rapidly learned simulations of two different medical diagnostic tasks (biochemical cues to the diagnosis of jaundice, and biochemical cues to a respiratory disease). The students clearly liked the use of the console. These results should be considered as merely implying that it is feasible to use the hardware described above for research with medical tasks and medical students.

Three conclusions can be drawn: (i) evidently the comparisons provided for the learner were appropriate ones, for judgmental accuracy improved rapidly; (ii) computer graphics successfully provided those comparisons; (iii) learning under these conditions was not "stupid" learning, since the learner was able to see clearly what he was expected to do and what he actually did.

Finally, the fundamental questions to which the study was directed were given clear, if tentative, answers. Human beings can exercise sufficient control over their cognitive processes to rapidly modify them in the direction intended. In other words, good judgment can be learned rapidly.

Future Research

The results above point toward the basic research and development that will be needed in order to improve the learning of an important cognitive skill —the exercise of good judgment in the application of knowledge already acquired. Of course, considerable basic research needs to be carried out in a wide variety of MCPL tasks, in order to gain more knowledge about human judgment and to discover how general the initial results are.

Moreover, considerable developmental research needs to be carried out before computer graphics techniques can be applied to medical diagnostic judgments. Uncovering the statistical characteristics of various diagnostic tasks, learning to scale various kinds of medical data, and discovering the range of tasks that can be dealt with by these procedures are major research problems. The development of informational displays will provide a considerable challenge to the ingenuity of those interested in computer graphics and the improvement of clinical competence.

It should not be assumed, however, that computer graphics of the sort we have described here must be restricted to use by medical *students*. Its potential use in refresher courses or specialized courses is obvious. Not so obvious is its potential use in training "medical associates," a new kind of personnel who will assist the physician in screening patients. Computer graphics should make it possible to teach such associates to recognize various common diseases in a fraction of the time it now takes. The learner can be presented with many more cases in much less time than it takes to see live patients; in addition, he can be presented with a greater variety of diseases. And, of course, skill in recognizing illness, rather than full knowledge about the disease and its treatment, is the primary function of medical associates.

Finally, computer graphics should be of considerable value in teaching clinical competence in those medical schools that have a very large number of students (for example, medical schools in Latin American countries). In such cases, the students have access to knowledge, but they have very little opportunity to develop skill in the application of their knowledge because of the limited opportunities for seeing patients.

But the sharp challenge to the research ingenuity of psychologists, computer scientists, engineers, and mathematicians will lie in using computer graphics to widen our knowledge about human thought processes—clearly a desperate need of the human race. Of prime significance is the fact that, although crude, these representations of a person's cognitive processes mean something to the person. It may well be that such representations will make it possible for man to enhance his understanding not only of his own judgmental processes, but those of others as well. Computer graphics may well mark a new era of research in the study of man.

Summary

Departing from the traditional model for teaching and learning, this article deals with the problem of teaching and learning the effective application of knowledge already acquired. To this end, a model for the process of exercising judgment was outlined, and the results of an empirical study of judgmental learning were employed to show the inadequacy of the traditional outcome feedback procedures. Computer graphics techniques were used to provide new forms of information to the learner; the results are promising for the rapid learning of a task that would otherwise be difficult to learn.

References and Notes

1. L. Stolurow, L. Peterson, A. Cunningham, Eds., *Computer-Assisted Instruction in the Health Sciences* (Entelek, Newburyport, Mass., 1970).
2. See L. Goldberg [*Psychol. Bull.* **73**, 422 (1970)] for a recent review of research.
3. K. R. Hammond [*Psychol. Rev.* **62**, 255 (1955)] provides the theoretical framework for this apprcach. See also K. R. Hammond, C. Hursch, F. J. Todd, *Psychol. Rev.* **71**, 438 (1964).
4. D. A. Summers and K. R. Hammond, *J. Exp. Psychol.* **71**, 751 (1966).
5. P. Hoffman, *Psychol. Bull.* **47**, 116 (1960).
6. S. Oskamp, *J. Consult. Psychol.* **29**, 261 (1965).
7. P. J. Hoffman, P. Slovic, L. G. Rorer, *Psychol. Bull.* **69**, 338 (1968).
8. P. Slovic, D. Fleissner, W. S. Bauman, *J. Bus.*, in press.
9. D. A. Summers, J. D. Taliaferro, D. Fletcher, *Psychonom. Sci.* **18**, 249 (1970).
10. K. R. Hammond and B. Brehmer, in *Human*

Judgment and Social Interaction, L. Rappoport and D. A. Summers, Eds. (Holt, Rinehart & Winston, New York, 1971).
11. D. Deane, K. R. Hammond, D. A. Summers, *Human Judgment and Social Interaction Program Report No. 131* (Institute of Behavioral Science, Univ. of Colorado, Boulder, 1970).
12. For similar results, see T. Earle, *Human Judgment and Social Interaction Program Report No. 129* (Institute of Behavioral Science, Univ. of Colorado, Boulder, 1970).
13. The computer graphics programs were written by P. J. R. Boyle, J. Little, M. Marshall, and J. Wilson.
14. The study was limited to five subjects for two reasons: (i) the results were clear; and (ii) finances were limited—the cost was approximately $100 per subject run on the interactive graphics terminal. Full details of this study are reported in K. R. Hammond and P. J. R. Boyle (*Bull. Brit. Psychol. Soc.*, in press).
15. F. J. Todd and K. R. Hammond, *Behav. Sci.* **10**, 429 (1965).
16. A somewhat different version of this paper was presented at the Conference on Information Systems for the Health Sciences Center, 3 April 1970, at the State University of New York, Stony Brook. The research was supported by National Institute of Mental Health grant MH-16437 and by the Commonwealth Fund. I thank F. Kern, D. Deane, and K. Will for their assistance.

Whither United States Universities?

George E. Pake

Someone has suggested dropping the first letter "h" from the title of this article. Let us hope that is only a bad joke—because some of us consider the university to be the most significant human institution for the future of free men. Yet there are knowledgeable, responsible people who have asked apprehensively, "What is the future of the university?" The question admits of speculation. But I believe that it is not susceptible to real prediction: The university's future hangs in precarious balance, and the direction in which the scale finally tips will be determined by as yet unresolved matters of institutional and, particularly, faculty government.

Where universities are concerned these days, few dare to claim expertise: I claim only involvement. I have spent 39 of my 46 years on a university campus. This is because my elementary and high school education took place in the university laboratory school of an institution now known to everyone—Kent State University. I received the B.S. and M.S. from Carnegie Tech (now Carnegie-Mellon University) and then spent a year at Westinghouse Research Laboratories. This was followed by a Ph.D. at Harvard and faculty positions at Washington University and Stanford University. For the next 7 years I was provost and executive vice chancellor of Washington University. Having just entered industry in 1970, I retain involvement as a new trustee of Washington University. Thus, I have seen the university from the vantage point of student, alumnus, faculty member, administrator, and now trustee. I have seen universities hold to high principle during the McCarthyist assault from without, and I have seen them stagger under recent assaults from within and, occasionally, from without. So much for my perspective.

The author is vice president, Xerox Corporation, Palo Alto Research Center, Palo Alto, California 94304.

Power Elements in the University

Whether it is possible at all to understand the present state of U.S. universities is debatable. But anyone who even hopes to understand universities must recognize that the faculty holds the de facto power in the university. Trustees, presidents with their administrative colleagues, and students each, as a group, has a modicum of power. But they can scarcely wield that power without the backing of the faculty, or at least a substantial portion of the faculty. In an ultimate, hypothetical showdown, the trustees probably could assert control by intervening in the firing and hiring of faculty and in the expulsion and admission of students. Even so, there seem to be serious doubts as to whether the courts would permit this exercise of absolute trustee power over administration, faculty, and students if there were the slightest indication that no form of due process was involved. Practically speaking, the trustees who went on such a rampage would likely find it impossible to recruit qualified new faculty and administration.

The philosophical basis for faculty power rests in the professional expertise of teacher and scholar and in the much-heralded principle of academic freedom. The practical basis rests in the implementation of the philosophy through faculty tenure and, since the 1950's, through what has been called the "star system."

Basically, the star system is the quest for well-known, internationally prestigious scholars and scientists, who help attract (i) institutional prestige, (ii) bright young faculty members, (iii) bright young students, and (iv) financial support. The Nobel laureate is the prototype, but of course there are stars in every field of academic endeavor, whether or not the Nobel bequest stipulates the field as one in which prizes are awarded. The nature of the power of these individual faculty stars emerges more clearly as one understands the diffuseness of faculty power generally and the weakness of trustees and administrators in the academic power structure. Such individual stars can exact from the administration commitments to better salaries, to an increased number of student assistants, to new office and laboratory space, and so on, simply by threatening to accept one of the standing offers they have from other star-seeking universities. How often the administrator wants to say, "Go ahead and take it!" That might soothe his frustrations. In fact, it would probably measurably weaken the university less often than faculty, trustees, and alumni may think.

Stated inelegantly but simply, academic freedom is the freedom of the scholar to search for truth, to reach his conclusions with intellectual honesty, and to retain his rights and privileges as scholar and teacher however unpopular his professional conclusions may be. The traditional example is that a Galileo or a Darwin should be allowed to hold his professorship even though the conclusions of his experimental or theoretical research, or both, are held by prevailing view to be heresy.