PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

Effective Decisions and their Verbal Justification

D. E. Broadbent

Phil. Trans. R. Soc. Lond. B 1990 **327**, 493-502
doi: 10.1098/rstb.1990.0092

Phil. Trans. R. Soc. Lond. B **327**, 493–502 (1990)
Printed in Great Britain

Effective decisions and their verbal justification

By D. E. Broadbent, F.R.S.

Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford OX1 3UD, U.K.

Recent experimental work on human control of complex systems has drawn attention to the discrepancy that may exist between the person's reported knowledge of the system and their ability to control it. Sometimes people act correctly but cannot answer questions about what they are doing; sometimes they can say verbally what they should do (perhaps having had verbal instruction in the right answers), but still do not do it. This discrepancy is of major practical importance, for example in designing training programmes or in eliciting expert knowledge for incorporation in a mechanical 'expert system'. It is also puzzling for psychological theory, as it rules out certain plausible models of the functioning of the brain. This paper considers what mechanisms are still possible.

1. Introduction

The practical problem

Suppose we consider some practical task that involves human beings controlling a complex dynamic process. Why should we be concerned about the links between their performance and their ability to tell us what they are doing? As some might argue, all that matters is that their handling of the system should be safe and efficient. From that point of view, worrying about their conscious experience (whatever that means) is a problem for the philosopher or perhaps the psychiatrist, not something that need occupy the safety engineer or the assessor of human reliability. That however would be a mistake; there are at least three reasons why we should think hard about the connection between words and actions.

First, there is the question of training. A safe and competent operator of a nuclear power station, or a chemical plant, or an air traffic control system, must clearly possess a substantial database of knowledge about physics, or chemistry, or the rules governing allocation of heights and headings to aircraft. The most familiar and probably still the most common method of passing that information is to assemble the potential operators in a lecture hall and talk to them; or to assign them printed readings. Is that in fact the best way of guaranteeing successful action?

Secondly, there is the assessment of the operator once trained. We need to be sure that the training, however given, has truly imparted the necessary knowledge. The student may have been inattentive, or of too low ability, or a crucial part of the training may have been missed. Again, the most common form of assessment is verbal; in traditional British universities students write essays on broad topics to illustrate their mastery of the symbolic knowledge about those topics. In other teaching systems, they may answer large numbers of short questions, which can be scored more objectively than essays. A third possibility is a spoken interview, in which the interviewer can modify the questions to probe particular areas of knowledge about which the individual student seems uncertain. But all these are verbal methods of assessment. If it were true that a person can answer questions about a system, and yet fail to act correctly on the same system, none of these methods would be adequate.

Thirdly, most of the systems that are potentially hazardous seem clear cases for providing the operator with automatic assistance for carrying out portions of the task, for advising what actions should be taken in an unfamiliar situation, or for reminding the operator of some action that may have been forgotten. Such assistance can reduce the number of decisions that the operator has to take each minute, and compensate for failures of memory, training, or attention. But the knowledge that is embodied in an 'expert system' has to come from somewhere. Can the knowledge engineer extract it from a suitable human expert by talking to that person?

The intellectual problem

More generally, we can only guard against real-life errors by a theoretical understanding of the manner in which the human controller functions. It would not be safe to accept a chemical process merely because it seems to produce the correct product in one set of background conditions; to be sure that it will avoid instability in some other circumstances, we need to understand the underlying equations. In the same way the human component of the control loop needs to be analysed, if we are to be sure that a person who can carry out some emergency procedure on leaving the training school can also do so a year later, in the middle of the night, and in danger of death.

Examples of the kind of analysis we need have been provided by Rasmussen (1983) or by Anderson (1984). As the latter argues, the fact that humans can perform certain functions implies that they have within their heads some representation of general facts about the world, which we can think of as a set of propositions. This 'declarative' knowledge can be accessed in a variety of different situations when the task requires it, and can thus act as a kind of model of the world on which the decision-making system can try out the consequences of various actions. The system can also predict what will happen from seeing some event that it has not caused itself, and this will be true even though the event has not previously been experienced. Probably none of us have ever seen a reactor from which all the control rods have been removed, but most of us could make a rough prediction of the future!

It has been shown by Johnson-Laird (1984) that such a 'mental model' explains the results of a number of experiments on human handling of language and of reasoning, far better than various alternative theories that have been suggested. The notion of a mental model also fits very well with a distinction made by Rasmussen (1983). He discusses a 'symbolic' level of function that can be seen in task analysis of the performance of real operators. However, Rasmussen also distinguishes a 'rule-based' level of performance, in which the operator acts according to a classification of the present problem into some broad category. In this level of performance there may not be an exhaustive 'look-ahead' calculation of all future possibilities, and therefore, bizarre human errors may result.

In the same way, Anderson (1984) incorporates in his system a second kind of knowledge. In addition to facts about the world, a decision system needs to know the actions to take under certain conditions; without such knowledge the system would remain lost in thought. This 'procedural knowledge' can very naturally be represented as a set of condition–action rules, a 'production system' (Newell & Simon 1972), and in Anderson's formulation the two kinds of knowledge are represented in these two different ways. As Anderson points out, the logical

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

distinction between the two kinds of knowledge does not mean that they are represented differently in the brain; computationally it would be quite possible for knowledge about appropriate actions to be held as a set of propositions rather than of rules. Equally, the propositional knowledge about the world could be held as a set of rules, and most expert systems intended for practical use do use such a representation. (The propositions are then implicit in the rules, but could be computed from them if necessary). It does however help to remind us of the logical distinction if, like Anderson, we think of different representations for declarative and for procedural knowledge.

At this stage we come back to the question of verbal knowledge about one's own actions. If all knowledge is represented in the nervous system in the same way, then we might expect it to be equally easy or difficult to elicit by questioning and in actual action. From a set of propositions that describe the functioning of a nuclear reactor, the operator could either derive the verbal answer to a question, or compute the correct action. If all knowledge were held in the same form, the lecture hall and the written examination would be good methods of teaching and assessing operators. If on the other hand there is some distinction of declarative and procedural knowledge, as Anderson supposes, there might also be greater difficulty in testing procedural knowledge by questions, or declarative knowledge by actions. The practical issues from which we started have come back to a theoretical issue about the form of representations in the brain.

2. The evidence

Signs of dissociation

Experiments on control

In a number of laboratory experiments, people have been practiced on some artificial control task in which the relations between each possible action and its consequences are known. In other experiments, people have been given verbal instruction about the relations and then asked to control with relatively limited practice. In either case, they can then be tested both for successful performance and for ability to answer objective-type questions about the relations. Cases have been found by Broadbent *et al.* (1986) and by Berry & Broadbent (1984), in which practice improves performance without improving success at question answering; and also in which verbal instruction improves question answering without improving performance. Funke & Muller (1988) also found that controlling a system (as opposed to merely watching events) improved performance without improving the ability to predict future events; and conversely, practice at predicting the future improved later predictions, but with no improvement in actual success at control. Both forms of training made the person less good at constructing a causal model of events within the system. Similarly, Berry & Broadbent (1984) found a negative correlation between the ability to control a system and the score obtained on questions about it; the people who could do one were to some extent less good at doing the other.

A criticism that may occur to some readers is that the particular questions chosen by the experimenter may not have been appropriate to the particular ideas of the person learning to control the system. They might have knowledge that is relevant to controlling the task, but which has not been tested by the questions. Stanley *et al.* (1989) asked people first to practice the task, and then to explain verbally to somebody else how to control it. Although in this case people could choose their own form of words, their own performance improved before they could tell somebody else how to succeed.

At a simple factual level, then, it is clear that people can sometimes be able to control a system without being able to talk about it, or to talk about it without being able to control it. The two ways of assessing knowledge are not totally separate, however. At a moderate level of practice they are dissociated, but more highly practiced 'expert' operators were found by Stanley *et al.* (1989) to be able to give verbal statements that would indeed help novices to perform more successfully. Even though a simple statement of the relations within the system may not help, and some kinds of true instruction may even be harmful (Berry & Broadbent 1988), yet it is possible to find certain kinds of verbal instructions that will improve performance (Berry & Broadbent 1988; Stanley *et al.* 1989).

It is particularly important that the presence of the dissociation may depend on the nature of the task. Berry & Broadbent (1988) found two tasks that were identical except for the presence of a lag in one of them, between each action and the appearance of the effect. The unlagged task was easier to learn, but with enough practice both tasks could be brought to the same level of performance. At that point, however, the operator could answer questions about the relations within the unlagged task, but not in the lagged one. There were also other differences (Hayes & Broadbent 1988); the lagged task could be disturbed by an unexpected reversal of the relations within the task, whereas the unlagged task was not. Again, the unlagged task was disturbed by asking the person to perform another task at the same time; but the lagged task was not. These facts suggest that the two, apparently similar, tasks are involving partly different mechanisms within the operator, as Anderson would suggest. That would be consistent also with another difference between the two versions of the task. The lagged task, but not the unlagged one, is disturbed by a verbal hint that might be expected to encourage the operator to try and learn the task in an 'explicit' and hypothesis-testing fashion rather than a more passive exploratory one (Berry & Broadbent 1988).

In the tasks that show the dissociation between verbal behaviour and correct decision, the key relations might be described as 'non-salient'. In those for which the operator's ability to perform agrees with their verbal knowledge, the variables are 'salient', obvious when the task is first approached. This can be illustrated by an experiment that required operators to control a simulated company, while at the same time interacting with a person representing the workers in the company (Berry & Broadbent 1987). Some of the relations in this system might well be expected; changes in the work-force affected the output of the company, and changes in the intimacy with which the person was treated had an effect on the intimacy shown in return. But there were also less obvious, hidden, relations; the degree of intimacy with the worker's representative also affected the output, and the size of the work-force affected the attitude of the worker's representative. Verbal questions about the salient and obvious relations revealed quite adequate knowledge, but questions about the non-salient relations were answered less well by practiced than by unpracticed people. Yet a test based on performance in the actual situation showed that the knowledge had in fact been acquired. Similarly, Porter (1986, 1988) asked people to control a task with two aspects; one was a straightforward sensori-motor reaction, with few and obvious variables. The other was a complex handling of strategic decisions. In the first, the verbal statements of the people agreed well with objective performance; in the second, they did not. Correspondingly, the first task was badly disturbed by a simultaneous other task, while the second was not.

Experiments on tasks other than control

While these developments have been taking place in the study of control, similar observations have come to light by using different tasks. Reber (1967) (Reber *et al.* 1980) examined the learning of artificial grammars; the person is required to learn sequences of letters, which are in fact all governed by abstract rules that allow only some kinds of sequence. After this experience, the person is able to discriminate new test sequences into grammatical and ungrammatical, but cannot give a coherent verbal account of the rules. It is important that the person does not know that there are rules when experiencing the sequences, but merely tries to learn the individual items; deliberate instructions to try and learn rules may actually slow down learning (Reber *et al.* 1980). This is suggestively parallel to some of the results of Berry & Broadbent (1988).

A second line of attack comes from memory experiments; in such experiments people may show effects of past experience by tests such as ease of identifying words in a brief flash, or judging how much they like one word rather than another. Yet they are unable to recall or even to recognize that the words have been seen previously (for a review, see Schacter (1987)). The different kinds of task may be affected differently by brain injury.

A third group of experiments involves 'concept learning', for example, Lewicki (1986). People are asked, for example, to look at photographs of young women and to judge their intelligence. Before doing so, the judges have previously seen a sample of similar pictures together with verbal descriptions of the women; for some judges, women with short hair were described as more intelligent, while for other judges, women with long hair were described as intelligent. The judgements made on a fresh set of pictures reflected the particular experience the individual judge had been given, yet the judges were unable to identify that they were being influenced by hair length.

There are therefore several lines of evidence showing that the verbal accounts people give of their mental processes may be out of agreement with their actual performance.

3. Theoretical implications

The evidence that has been reviewed is reasonably consistent with Anderson's distinction between declarative and procedural knowledge, as involving different mechanisms or representations in the human decision system. Although consistent with it, however, some of the evidence does not exclude alternative interpretations.

Accessibility of memory by various routes

It has to be recognized that this area is controversial; in everyday life we know very well that somebody who understands a situation can tell us about it and also take appropriate actions. If they are wrong in their actions, we are not likely to trust anything they may say to us about the situation. It seems to run counter to this commonsense view if we suppose that people can show knowledge in one way and yet not in another. Furthermore, admitting that such a thing can happen seems to require two decision mechanisms, or two kinds of knowledge, within the brain. Perhaps it also needs some kind of barrier or insulation between them. That is a complication that most scientists would prefer to avoid if they possibly can.

Admittedly, we have seen plenty of experiments in which the things people say do not correspond to the things they do; but the weak point is that correct action may be based on a belief that is not the one the experimenter has asked about. It may indeed be based on a false belief. If I believe my neighbours are spraying my lettuces with arsenic, that can be false, and yet the fact that it makes me wash each salad before I eat it may be good for my health. Perhaps some similar fantasy underlies action in each of the experiments I have mentioned.

Logically, one cannot convince a sceptic about this, because it is impossible to ask questions about every conceivable relation that might exist. The best one can do is to make such a view implausible, by asking questions directly about the same measures that are taken from performance (Broadbent *et al.* 1986), by trying a variety of different kinds of questions (Berry & Broadbent 1984), or by leaving the person free to express their own beliefs whatever they may be (Stanley *et al.* 1989). It is always possible that some other measure of verbal knowledge would agree with action; one cannot prove a negative.

In addition, it is clear from the evidence that different measures of verbal knowledge are not equivalent. From general knowledge about human memory, we could not expect that they would be. If we show a person a list of words and then test memory, different tests will reveal different amounts of knowledge. Typically, a simple request for the person to recall all the words will give only a limited degree of success. Many of the words that have been 'forgotten' will however be recognized correctly if they are shown mixed with others (for example, Shepard 1967). They may even be recalled if some hint associated with the correct word is given (Tulving & Pearlstone 1966). Although information is in the brain, it is not accessible to one kind of test even though it can be shown through another kind of test.

To some extent this is a matter of one test being more sensitive than another, but that is not the whole story. If the context is deliberately varied, it is possible to make a recognition test fail for items that the person can recall (Tulving & Thomson 1973; Tulving 1983). Under special circumstances, people will recognize items they cannot recall just as well as items they can; the two tests are independent, rather than one being more sensitive than the other (Broadbent & Broadbent 1975, 1977; Rabinowitz *et al.* 1977). The key factor is that the situation should provide the correct context when memory is interrogated, so that the information about the original event can be retrieved. Just as a computer database may find it much easier to access a record when probed with a descriptor assigned to that record when originally filed, so people may also remember something in one context and forget it in another.

Such mechanisms are undoubtedly active in some of the situations usually described as studies of implicit and explicit knowledge. Thus, for example, Dulany *et al.* (1984) used Reber's technique of learning artificial grammars, but tested for explicit knowledge in a different way. Instead of asking merely for statements of the rules, they presented sequences of letters and asked people to show which part of the sequence was guiding their decision. This method revealed some knowledge of the regularities in the strings, that may have been enough to control the decisions about grammaticality. Similarly, in the control task of Berry & Broadbent (1987), the method used to reveal knowledge of non-salient relationships between personal relationships in the factory, and output of the factory, was to put the person under test back in a sample of the task rather than to give questions on paper.

Perhaps particularly striking is a result of Marescaux *et al.* (1989) using the same control task as that employed by Berry & Broadbent (1984). After people had practiced the task and could control it, they were asked 'questions' about the task in the form of a series of problems shown to them at the computer, so that they saw just what they might have seen in the full task. In that case, some knowledge of the task was revealed, but only on some of the questions. The key factor appears to be that the person was tested on specific situations which they had themselves experienced while practicing. They did not seem to have learned something that could be used in other novel situations.

It is therefore clear that knowledge may be elicitable by tests that recreate the original context, even though it may fail to appear on tests that differ too much from the situation during learning. Some of the studies mentioned in the last section do not explore this possibility. For example, the work of Reber or of Lewicki concentrates on the discrepancy between performance and the ability to state principles, without going into the degree of knowledge that might appear in situations resembling those of original learning. For this reason, one can only draw limited conclusions from the experiments that show discrepancies between verbal behaviour and action. They do not make it necessary to accept Anderson's distinction of declarative and procedural knowledge. Action may be based on some untested form of knowledge, and the original context may fit better with the measure of action than it does with the measure of verbal knowledge. Again, in other situations verbal knowledge may be more accessible than correct action is. It is also possible that different tests of verbal knowledge may disagree with each other, and that some of them will agree with action while others do not.

Modes of learning and forms of representation

Merely demonstrating a dissociation then will not convince anybody that there are fundamentally different mechanisms for implicit and for explicit knowledge. If indeed the mechanisms are different, however, they should be affected by different experimental variables. It should in that case be possible to find tasks in which action and verbal report appear to gain access to the same knowledge, and closely similar tasks in which they do not. By looking at the exact differences between such tasks we may shed light on the difference between mechanisms. This is a much more hopeful strategy for the future than merely demonstrating or denying dissociation.

So far, perhaps the clearest case of such a comparison is the pair of tasks devised by Dianne Berry and described above (Berry & Broadbent 1988). In both tasks, each control action by the person produces an output from the system that depends only on that input and simply adds a constant to it. In the 'salient' case, the output appears immediately; and when people have learned to control this version they can answer questions about it. In the other, 'non-salient', version, the output appears only after the next input, so that each output is undetermined by the last input. This version shows the dissociation of action and performance, and as noted earlier, the two tasks show a number of other differences. Thus the two tasks do seem to be learned in some rather different way.

Although these two tasks are the most nearly identical ones in the literature, the differences between them are consistent with a number of other findings. First, as has been noted above, Porter (1986, 1988) observed two, fairly different, tasks within a single general situation. The task on which verbal knowledge was better was also the one more disrupted by a secondary task. Secondly, the concept of salience was derived from a study by Reber *et al.* (1980), in which strings of letters were shown to the person either grouped according to the grammatical rule they illustrated, or mixed together. The first form of presentation showed more beneficial effects of explicit instructions to learn the rules, just as the salient task of Berry & Broadbent did. Thirdly, in the tasks of Broadbent *et al.* (1986), practice of a single relations on its own improved verbal knowledge as well as performance, whereas practice of multiple relations together did not; and experience of a complex crossed relation caused some deterioration of verbal knowledge when a simpler version of the same relation did not.

It is worth mentioning also the suggestive parallel between the concept of 'salience' in the Anglo-Saxon literature, and that of 'transparency' in German; the interest in that case has not been so much concerned with the relation of performance to verbal report, as with the relation to tests of intelligence. Studies of system control have been found several times to show little relation to intelligence, as long as there are many variables and the relations between these variables are concealed. When however, key relations are available to the operator, then test intelligence becomes more relevant. (For reviews in English, see Putz-Osterloh & Lemme (1987); Funke (1988); Funke (1990).)

The tasks that show dissociation seem therefore to be those in which the person has to learn to act in a situation, and yet the key features of that situation are unclear because they are accompanied by many other events. Even if the key features have been explained previously to the person, they may not be recalled at the correct moment in the task unless the person is deliberately asked to say them at that time (Berry & Broadbent 1984). Hence the learning takes the form of building links between situations and actions, 'procedures' in Anderson's sense. Such learning is best built up unselectively by exposure to the task itself.

Tasks that show good agreement between action and speech, on the other hand, are those with few variables and with clear relations between them, which can reasonably be learned by selective concentration on those few variables and by testing hypotheses about possible relationships. The resulting knowledge can most simply be described, not as links between situations and actions, but as a set of propositions about the underlying variables; if the proposition is known, it may produce an entirely new correct action if the situation is new. Propositional knowledge however corresponds more to Anderson's declarative than to his procedural knowledge.

The strategy or mode of learning that is going to be most successful is likely to be different for the two kinds of task, so that for a salient or transparent task, learning will be better with verbal instruction and will transfer easily to new situations; for a non-salient task, learning will be better through experience, badly tested by questioning, and transfer only weakly to new situations. If experiment shows that the best conditions for learning reverse when the task is different, that is much stronger evidence for the existence of two mechanisms than is given by the mere demonstration that people can perform without being able to talk about it. Although there are few reversals of relations in the literature, we have seen some.

Remember for example, that Berry & Broadbent (1984) found, across their whole body of experiments, a significant negative correlation between the ability to perform well and the ability to answer questions about the situation. The people who were best at one were actually worse at the other. Similarly, Funke & Muller (1988) found that factors improving scores by one measure gave significantly negative results on another measure.

It should not be surprising to find negative relations because in machine intelligence it is found that the order of merit of various computational strategies may reverse from one situation to another. From declarative knowledge of the possible moves in chess, a chess-playing program can compute possible future consequences for any action and therefore do much better than a human amateur who fails to notice some possible trap. But in certain situations, such as the end-game, such a program may be handicapped by the very large number of alternatives. It may then lose to another program (or even a human) that operates on procedural knowledge, choosing an action by consulting a previously calculated table that lists actions appropriate to situations. On that strategy, the results of all the other possible moves are never calculated. The kind of tasks that are best for calculation from propositional knowledge may well be worst for the application of condition–action rules, and vice versa. Any realistic chess program uses both techniques as appropriate; and it is reasonable that human beings should do the same (Michie & Johnston 1984).

Conclusions

At the very least, we have to recognize that the conditions of testing knowledge are vitally important in deciding whether it will be revealed. The evidence for that is overwhelming. Even if all knowledge is ultimately represented in the same form, therefore, access to it is by different paths.

Although most of the evidence merely shows this more restricted point, there are a few studies that go further than that. The comparisons of different tasks do establish that a distinction of procedural, condition–action, knowledge from declarative, event–event knowledge, is functionally useful. Some kinds of representation are useful only if the conditions of performance are very similar to those of learning; others are more portable from one situation to another. The ultimate basis in the nervous system may well be the same for both types, much as neurones operate on the same basis in many different parts of the brain. Just as it is useful to distinguish the auditory from the visual pathways, however, without implying a different kind of neurone, so it is useful to keep declarative and procedural knowledge separate in our thinking.

4. Practical implications

If we return now to the practical questions raised in the introduction, we have seen evidence that supports the views of many non-academics. It is not safe to rely on classroom instruction for the training of those who take decisions in complex systems. Equally, one cannot always assess the knowledge of such people by, for example, a written test or a promotion interview. If in despair one decides to build into a non-human system the knowledge possessed by the human expert, it will not necessarily be possible for the expert to reveal that knowledge to you.

However, the evidence also points to a real role for knowledge that is verbal, propositional, or declarative. That type of knowledge is superior when the situation is transparent; it is useful in turning apparently opaque tasks into transparent ones by isolating the correct key variables. It can be used to train operators to call up relevant knowledge at the right moment in the task. Above all it allows the person who possesses it to react in novel situations when those who are basing action on intuitive procedures may be adrift.

The true lesson therefore is not that intuition should always triumph, and that in risky situations you should automatically let 'the force be with you'. Rather, we need to design tasks and training situations so that the computational advantages of each way of operating are exploited. To do this will need explicit statements about the scope and limits of implicit knowledge.

The author is employed by the Medical Research Council. The experiments of Dr Dianne Berry were supported by the Economic and Social Research Council.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

References

Anderson, J. 1984 *The architecture of cognition*. Cambridge, Massachusetts: Harvard University Press.

Berry, D. C. & Broadbent, D. E. 1984 On the relationship between task performance and associated verbalizable knowledge. *Q. J. exp. Psychol.* **36**A, 209–231.

Berry, D. & Broadbent, D. E. 1987 The combination of explicit and implicit learning processes in task control. *Psychol. Res.* **49**, 7–15.

Berry, D. & Broadbent, D. E. 1988 Interactive tasks and the implicit–explicit distinction. *Q. Jl exp. Psychol.* **36**A, 209–231.

Broadbent, D. E. & Broadbent, M. H. P. 1975 The recognition of words that cannot be recalled. In *Attention & performance V* (ed. P. M. A. Rabbitt & S. Dornic), pp. 575–590. New York: Academic Press.

Broadbent, D. E. & Broadbent, M. H. P. 1977 Effects of recognition on subsequent recall. *J. exp. Psychol.: general* **106**, 330–335.

Broadbent, D. E., FitzGerald, P. & Broadbent, M. H. P. 1986 Implicit and explicit knowledge in the control of complex systems. *Br. J. Psychol.* **77**, 33–50.

Dulany, D. E., Carlson, R. A. & Dewey, G. I. 1984 A case of syntactical learning and judgment: how conscious and how abstract? *J. exp. Psychol.: general* **113**, 541–555.

Funke, J. 1988 Computer-simulated scenarios: a review of studies in the F.R.G. *Simul. Games* **19**, 277–303.

Funke, J. 1990 Solving complex problems: human identification and control of complex systems. In *Complex problem solving: principles and mechanisms* (ed. R. J. Sternberg & P. French). Hillsdale, New Jersey: Erlbaum.

Funke, J. & Muller, H. 1988 Eingreifen und prognostizieren als determinanten von systemidentifikation und systemsteuerung. *Sprache kog.* **7**, 176–186.

Hayes, N. & Broadbent, D. E. 1988 Two modes of learning for interactive tasks. *Cognition* **28**, 249–276.

Johnson-Laird, P. N. 1983 *Mental models*. Cambridge University Press.

Lewicki, P. 1986 *Nonconscious social information processing*. Orlando, Florida: Academic Press.

Marescaux, P.-J., Luc, F. & Karnas, G. 1989 Modes d'apprentissage selectif et non-selectif et connaissances acquises au controle d'un processus: evaluation d'un modele simule. *Cah. Psychol. Cog.* **9**, 239–264.

Michie, D. & Johnston, R. 1984 *The creative computer: machine intelligence and human knowledge*. Harmondsworth, Middlesex: Penguin.

Newell, A. & Simon, H. A. 1972 *Human problem solving*. Englewood Cliffs, New Jersey: Prentice-Hall.

Porter, D. B. 1986 *A functional examination of intermediate cognitive processes*. D.Phil thesis, University of Oxford.

Porter, D. B. 1988 Computer games and human performance. In *Proceedings, 11th Symposium on Psychology in the Department of Defense*, pp. 251–255. Colorado Springs: United States Airforce Academy.

Putz-Osterloh, W. & Lemme, M. 1987 Knowledge and its intelligent application to problem solving. *Ger. J. Psychol.* **11**, 286–303.

Rabinowitz, J. C., Mandler, G. & Patterson, K. E. 1977 Determinants of recognition and recall: accessibility and generation. *J. exp. Psychol.: general* **106**, 302–329.

Rasmussen, J. R. 1983 Skills, rules, and knowledge; signals, signs, and symbols, and other distinctions in human performance models. *IEEE Trans. Syst. Man Cybern.* **13**, 257–266.

Reber, A. S. 1967 Implicit learning of artificial grammars. *J. Verbal Learn. & Verbal Behav.* **5**, 855–863.

Reber, A. S., Kassin, S. M., Lewis, S. & Cantor, G. 1980 On the relationship between implicit and explicit modes of learning a complex rule structure. *J. exp. Psychol.* **6**, 492–502.

Schacter, D. L. 1987 Implicit memory: history and current status. *J. exp. Psychol.* **13**, 501–518.

Shepard, R. N. 1967 Recognition memory for words, sentences, and pictures. *J. Verbal Learn. Verbal Behav.* **6**, 136–143.

Stanley, W. B., Mathews, R. C., Buss, R. R., & Kotler-Cope, S. 1989 Insight without awareness: on the interaction of verbalization, instruction and practice in a simulated process control task. *Q. Jl exp. Psychol.* **41**A, 553–577.

Tulving, E. 1983 *Elements of episodic memory*. Oxford University Press.

Tulving, E. & Pearlstone, Z. 1966 Availability versus accessibility of information in memory for words. *J. verbal Learn. verbal Behav.* **5**, 381–391.

Tulving, E. & Thomson, D. M. 1973 Encoding specificity and retrieval processes in episodic memory. *Psychol. Rev.* **80**, 352–373.