

EXTRAIT DE

LA RECHERCHE
EN ENSEIGNEMENT
PROGRAMMÉ
Tendances actuelles

PROGRAMMED
LEARNING RESEARCH
Major trends

ACTES D'UN COLLOQUE OTAN

NICE - Mai 1968

DUNOD
PARIS
1969

ADAPTIVELY CONTROLLED INSTRUCTION FOR A TRACKING SKILL

B. R. GAINES

Department of Electrical Engineering Science, University of Essex, Colchester, Essex, UK.

Le problème de contrôle de la régulation de l'apprentissage humain est étudié dans la recherche présentée au moyen de l'entraînement vu comme un comportement adaptatif se modifiant sous l'effet de la performance aux tâches ou sous-tâches qui le composent.

L'auteur propose un système de contrôle adaptatif (entraîneur automatique avec feedback) faisant varier la difficulté de la tâche en accord avec la situation du sujet concerné, et destiné à résoudre automatiquement ce problème.

La communication résume également les résultats de quelques expérimentations qui analysent l'interaction de l'instruction verbale et des techniques d'entraînement pour des opérateurs humains et des systèmes de simulation de l'apprentissage sur ordinateur.

1. Introduction

In the design of a system containing human operators there is no logical reason why the human components should not be treated in the same way as other sub-systems, but there are practical differences at present because the men in the system are not amenable to 'synthesis' by parameter adjustment as are conventional controllers. Selection procedures cannot pick out the astronaut with the skills necessary to fly a Gemini spacecraft, because men with the required specification do not exist. The requisite skills have to be acquired through learning, and the synthesis problem becomes one of making best use of the human adaptive capabilities through education and training.

The studies outlined in this paper have been concerned with the possibility of formalizing the synthesis procedures utilized in training and performing them automatically with an *automated feedback trainer*. The first section lays the foundations for the study by analysing adaptive behaviour and demonstrating that training is itself a control problem. The complexity and irreversibility of the human adaptive controller, however, makes a complete identification of its behavioural dynamics inherently impossible. Equally the structure of the human controller is not only unknown but unknowable, and recourse must be made to other information in the design of an optimum training system. This information comes from an analysis of the basic epistemological problem of attempting to control an environment whilst at the same time learning about it in order to improve the control policy.

These considerations suggest a suitable feedback training system, which is 'adaptive' in the sense that it varies the difficulty of the task to be learnt according to the operator's performance on it. Partial knowledge of the human control structure enables the stability of this training system to be analysed, and experimental studies of its behaviour with both humans and automatic controllers as operators confirm this analysis.

A suitable environment for the investigation of training is described which involves the acquisition of a novel tracking skill. Formal experiments have been performed on the learning of both human operators and intelligent artifacts in this environment, and a comparison is made between different training procedures, including feedback training. The importance of verbal communication with the human operator is demonstrated by the gross interaction between the effects of instructions and training procedures.

2. Training as a Control Problem

2.1. ADAPTIVE BEHAVIOUR

The expected behaviour of an adaptive controller when coupled to an environment is that, if its control policy is not satisfactory for that environment, then it will eventually become so. Thus it must be possible to segment the interaction of the controller with its environment into at least two phases, in the first of which it is not satisfactory and in the last of which it has become so.

This segmentation, which is inherent in the basic concept of adaption, may be extended to form a description of the full range of possible adaptive behaviour. A *task* is defined to be a segment of the interaction between controller and environment for which it is possible to say whether or not the controller has performed satisfactorily. At the beginning of a task the controller will be in some *state* which causes it to implement a particular control policy. If the controller is deterministic and the task is reasonably defined, then the final state of the controller will be uniquely determined by its initial state and the task given to it.

Hence the adaptive behaviour of a controller may be ascribed to an automaton whose inputs are tasks, whose states are controller states, and whose outputs are on one hand control policies, and on the other the satisfactoriness of the control policy for a given task. This is the *adaption-automaton* of the controller relative to the set of tasks, and training may be shown to be a problem of controlling this automaton.

2.2. MODES OF ADAPTION

The fundamental situation with which an adaptive controller is expected to cope, is to be coupled to a fixed environment and learn to control it satis-

factorily. This is equivalent to the *adaption-automaton* being given a sequence of inputs consisting of the same task repeated indefinitely. An interaction between controller and environment consisting of the repetition of a single task is defined to be *acceptable* if it is eventually always satisfactory. Thus, in an acceptable interaction, the initial performance of the controller does not matter, and for a number of repetitions of the task it may be satisfactory, unsatisfactory, or waver between the two. Eventually, however, it must become satisfactory and remain so; an acceptable interaction is one which reaches a *stable* condition of satisfactoriness. In this stable condition the controller is said to be *adapted* to the task.

These definitions, based on the behaviour of the adaption-automaton, may be extended to account for all the tremendous variety of adaptive phenomena. Only one mode of adaption, *potientiel adaption*, will be considered here. A controller in such a state that it will have an acceptable interaction with any one of a set of tasks is defined to be *potentially adaptive* to that set of tasks. Thus a controller which is potentially adaptive to a set of tasks is able to learn to perform any one of them satisfactorily. This definition is best illustrated by a diagram showing the state-trajectories of the adaption-automaton as the controller adapts.




Fig. 1.

The state-space of the automaton is shown as a rectangular region in Figure 1, and within it are delimited those states for which the controller is satisfactory given the task, t . The states for which the controller is adapted to t form a sub-set of these, since a trajectory starting in the adapted region must always remain satisfactory. The states for which the controller is potentially adaptive to t form another region enclosing the adapted one.

A trajectory through the state-space, generated by giving the controller the task, t , many times in succession, t^n , will show the following behaviour :

- started outside the potentially adaptive region, at Y, it may enter the region of satisfactory interaction but must eventually leave it;
- started within the potentially adaptive region, at X, it will remain within that region, eventually entering the adapted region and never leaving it.

2.3. MODES OF TRAINING

When the controller's initial state is outside the region of potential adaption to a task, learning will not take place if it is given that task alone. Given some other sequence of tasks, however, the controller will adapt to them and may, in so doing, become potentially adaptive to the original task—the sequence of tasks has *trained* it for the original task.

In Figure 1, for example, the point A is outside the region of potential adaption to the task, t , and repetition of this task will not lead to stable satisfactory performance. The sequence of tasks, u , however, gives rise to a trajectory which terminates within the region when the controller is potentially adaptive to t . If the training sequence, u , consisted of another task, t' , repeated, then we would say that there had been transfer of learning from the task t' to the task t .

The training sequence, u , will not necessarily be suitable for all initial conditions of the controller; e.g. the trajectory induced by u from the point B in Figure 1 terminates at B₁ which is outside the region of potential adaption. These are the regions in which the training sequence u gives complete transfer (trajectories enter adapted region) and partial transfer (trajectories enter potentially adaptive region) to the task, t . Outside these regions, although u is not a suitable training sequence, it may be possible to find an alternative sequence, v , which again causes the controller to become potentially adaptive or adapted. To choose between u and v , however, it is necessary to have some information about the initial condition of the controller.

These considerations lead to the definition of three types of training :

(i) Fixed training—to 'train' the controller for the task, t , it is given a sequence consisting of this task repeated a number of times. This involves no training effort and relied on the controller being potentially adaptive to the task, t .

(ii) Open-loop training—the trainer is given some training sequence, u , before proceeding to learn z . The training sequence is not adjusted for differences in controllers, but will obviously be chosen to maximize the region in state-space from which it gives partial or complete transfer.

(iii) Feedback training—the controller is given a sequence of tasks selected according to information about its state. Training has now become a problem of controlling the adaptation-automaton by varying its input according to some output received from it.

3. Problems of Learning

3.1. THE SUB-ENVIRONMENT PHENOMENON

Although theory of adaptive behaviour enables the rationale behind various forms of training to be defined and analysed, it does not itself indicate the techniques for setting up an open-loop training sequence or realizing a feedback trainer. Indeed it may appear that these are highly dependent on the structure of the specific controller to be trained, and cannot be considered in general. For example, the information flow from most natural environments is highly redundant, and different animals use different senses to identify, for instance, the route through a forest. Factors which increased the difficulty for an animal relying on its sense of smell would probably be irrelevant to one relying on its sense of sight.

These specific properties of the learning system and its environment are not the fundamental causes of difficulty in learning, however, and their importance lies in their influence on more basic and profound problems, inherent in the act of learning itself, which may be analysed in general. The study of adaptive artifacts, in particular, has led to an understanding of the basic epistemological problems involved in attempting to control an environment whilst at the same time learning about it in order to improve the control strategy (the so-called *dual control problem*).

The fundamental structure of all forms of adaptive controller is a two-level hierarchy in which the lower level implements one of the class of possible control policies, whilst the upper level selects the policy to be implemented. Difficulties in learning, or the purposeful variation of the policy selected as a function of experience, arise because any given control policy will generate some sub environment. That is, it will restrict the states and state transitions of the environment to some sub-set of the total possible behaviour. The sub-environment generated by an initial control policy may be entirely different from that generated by an optimum or satisfactory control policy, and learning in the initial sub-environment may then be irrelevant or deleterious to performance in the desired sub-environment. For example, if the environment is a high-order continuous control system to be regulated at zero output,

then the initial control policy will probably cause the closed-loop to be unstable, and not only will the output position and velocity never attain zero together, but also the system will limit and show non-linear behaviour.

3.2. TRAINING STRATEGIES

The obvious training strategy to alleviate the difficulties caused by the sub-environment phenomenon is to force the initial sub-environment of the controller to be the desired sub-environment. This may be represented, Figure 2, as the addition of a *training controller* to the environment, which, in the case of the regulation of high order dynamics, might apply auxiliary feedback to maintain the overall control loop marginally stable.

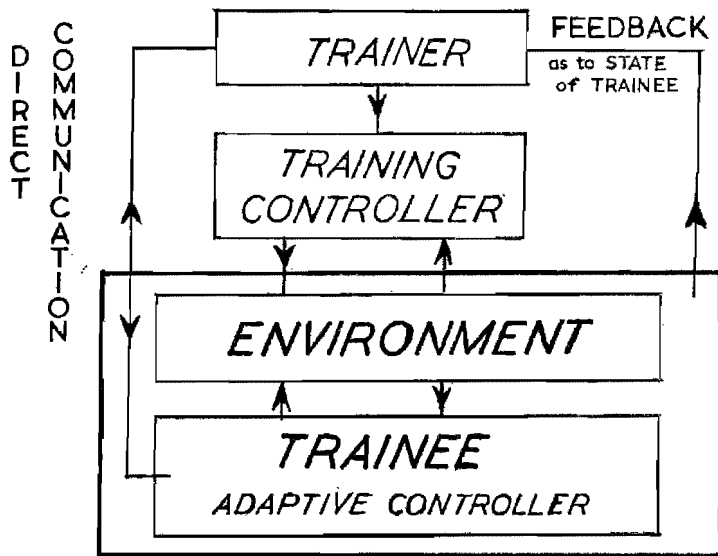


Fig. 2. — *Feedback training system*

Since the trainee in Figure 2 is assumed to be adaptive, the training controller need exert less and less control to maintain the desired sub-environment as time passes. Hence there may be a *trainer* which selects a suitable training controller either as a function of time (open-loop training), or according to information about the state of the trainee (feedback training). The structure of the training system is, in fact, an exact image of the structure of the trainee, regarded as an adaptive controller. With the human controller, and with recent learning machines, direct verbal communication is possible, and the trainer may have access to this channel for *priming* the trainee with control policies or adaptive strategies.

4. An Automated Feedback Trainer for a Tracking Skill

4.1. CHOICE OF CONTROL TASK

In choosing a control situation in which to investigate the learning of a perceptual-motor skill many factors were taken into account. It was required that the task be related to practical situations in which training was already employed, and the regulation of high order dynamics, such as those of the longitudinal motion of an aircraft (BLAKELOCK 1965) or submarine, was selected as being both realistic and of fundamental interest in manual and automatic control.

Preliminary experiments and comparison with aircraft dynamics indicated that a second-order, stable transfer-function, with undamped natural frequency in the range between 0 and .8 Hz, and a damping ratio in the range between 0 and 1, was most suitable. However, the human operator is capable of compensating such a system fairly easily, and the dynamics were increased to third-order by addition of a rate control. The overall transfer-function was thus of the form :

$$1/s(a^2s^2 + 2kabs + b^2)$$

which, by variation of k and a , may be swept from virtually first-order to pure third-order in a variety of trajectories through the natural-frequency/damping ratio plane. Variation of k and a thus constitutes a means of changing the degree of compensation required, and hence the difficulty of the task for the human operator.

The operator was provided with an input to the above transferfunction by means of a manual control, and a second disturbing input was provided within the system. The error in maintaining the output of the transfer-function at zero was shown to the operator on a cathode-ray tube display.

This control problem is similar to that used in some previous studies of human perceptual-motor skills, but, as commonly used, it suffers from two major defects which make it difficult to obtain meaningful results in experiments on learning. The first problem is operator fatigue which may affect tracking performance after intervals of as little as 90 seconds. This is a minor nuisance in studies of control policies since short tracking runs have to be used, but in the study of learning it causes artifacts and difficulties in experimental control which are virtually insuperable (artifacts, that is, when we are not trying to investigate the phenomenon of fatigue and its effects on learning). Fortunately, preliminary experiments had shown that the use of discrete push-button controls, rather than a conventional joy-stick, greatly reduced operator fatigue so that runs of 15 to 20 minutes became acceptable.

The second criticism of the basic control system, even with discrete controls, is that there is an unrealistic emphasis on the particular task of compensatory tracking. In most real-life situations where manual tracking is required,

the difficulties which the operator must learn to overcome are not based on requirements for a very high level of skill in a single task, but rather for competent performance of each of a number of interacting tasks. In terms of the discussion of section 3.1, the effect of interactions between tasks is that the sub-environment provided for the learning of one task by poor performance on another may be very different from the desired sub-environment when that task is being performed correctly. For a skilled operator the apparent interaction between tasks may be very slight—in terms of multivariable control theory he has de-coupled the control loops—and in the measurement of final control policies the interaction terms may usually be neglected. During learning, however, the interaction may be the predominant variable affecting adaptation, and should not be omitted from laboratory simulators for experiments on training.

Task interaction was introduced into the control system already described by the simple device of incorporating memory in the push-button controls. The operator had two buttons built into the arms of his chair, one for each thumb, and at any instant pressing one of the buttons would give a positive impulse at the input to the transfer-function, or pressing the other would give a negative impulse. At each push, however, the sense of the push-buttons reversed, so that the one which had previously been positive was next negative and vice versa. The chief problem introduced by this reversal was that an operator, on pushing a button and finding that it increased the error, had an innate tendency to then push the other button—increasing the error still further.

Finally, the disturbance at the input to the transfer function was chosen to be a square wave of 20 second period, so that its characteristics could also be learned. The overall control system is obviously very different from that required in work on operator dynamics, where sources of non-stationarity are to be avoided. The experimental paradigm is, however, one in which the phenomena of learning have every opportunity to appear.

4.2. FEEDBACK TRAINING STRATEGY

The third-order transfer-function described above is that of a linear system with three state-variables, the position, velocity and acceleration of a spot on the CRT display. The desired sub-environment of a regulatory controller is a region about zero in the state-space. Provided this region does not impinge on the boundaries of the state-space (the position, velocity, and acceleration, are each bounded in magnitude in any physical realization of the transfer-function), the system will behave within it in a linear manner. The desired sub-environment will be of finite size because of the disturbance, which even if it is predicted, cannot be cancelled instantaneously. The maximum value of the disturbance was chosen so that a skilled operator using the push-button controls could maintain the system in its linear region.

The control policy of a naive operator attempting to control a third-order system gives rise to an unstable loop, and the state-trajectory of the system tends to follow the boundaries of the state-space. Thus the initial sub-environment may lie entirely outside the desired sub-environment and will correspond to a non-linear, rather than linear, system. A suitable training controller will then be one which attempts to maintain the desired sub-environment by cancelling the disturbance and stabilizing the control-loop.

Figure 3 illustrates a training controller for a pure third-order system which has two feedback paths to stabilize the control loop and a third to vary the magnitude of the disturbance. The particular feedback paths chosen are, in fact, those which reduce the system to a pure integrator following a stable second-order transfer-function of variable natural-frequency and damping ratio—that is, the transfer-function described in section 4.1.

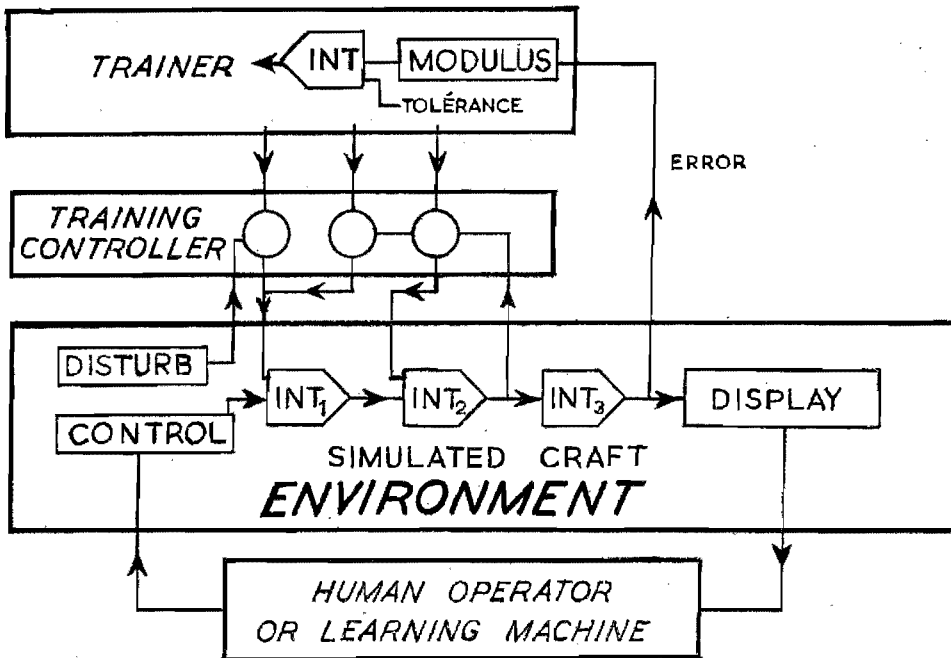


Fig. 3. — Feedback trainer for a control skill

There are thus three parameters of the training controller which affect the difficulty of regulating the environment. Within the range of values used one may say that the difficulty increases as :

the disturbance is increased from zero to its maximum value;

the undamped natural frequency is decreased from its maximum value to zero;

the damping ratio is decreased from its maximum value to zero.

In the experiments this three-dimensional space was reduced to a single dimension by fixing the values of one or more of the parameters and locking the others to a single continuum of difficulty. In fixed training and open-loop training the difficulty could be set at one or more levels during the training period, whereas in feedback training it had to be made to co-vary with the state of the operator.

Since the desired sub-environment is a region about zero in the state-space of the environment, it is possible for the trainer to detect by direct measurement whether or not this is being maintained. Under the experimental conditions, the bounds on the position of the spot were far more stringent than those on its velocity or acceleration, and hence the positional error was a sufficient indication of the effective sub-environment. A tolerated magnitude of error was fixed to define the boundary of the desired sub-environment, and the strategy of the trainer was such as to increase the difficulty of the task when the error was within tolerance, and to decrease it otherwise.

This feedback training strategy was realized in practice by taking modulus of the error, subtracting the tolerance from it, and feeding the result to an integrator. The output of the integrator drove a servo multiplier whose potentiometers set the magnitudes of disturbances and feedback around the integrators in the environment. Thus, when the mean error modulus was above tolerance, the output of the integrator tended to rise and decrease the difficulty of control, whilst when it was below tolerance the output of the integrator would fall and increase the difficulty of control.

4.3. STABILITY OF AUTOMATED FEEDBACK TRAINER

With a non-adaptive controller, the only stable value of difficulty will obviously be uniquely determined by the ability of the controller to regulate the control system. It is not obvious, however, that the feedback training loop is stable, and, indeed, it may be shown that with certain forms of controller instability may occur. The overall system is complex, since many feedback loops are operative and a major part is nonlinear, but a simple analysis may be based on linearization of the outer, parameter-adjusting loop. This demonstrates that the system is stable for an operator whose control strategy gives rise to a limit-cycle of monotonically increasing amplitude with difficulty. Because of the high-order of the control system, and the nature of the manual controls, this condition was satisfied in the experiments reported here.

The linearized characteristics of the outer loop depend on the form of control policy used by the operator—in particular, whether it induces a limit-cycle and at what rate this cycle is established. The experimental conditions

were such that both the human operators and automatic controllers induced limit-cycles whose change in amplitude, with changing environment, was rapid compared with the change of parameters in the training controller. The outer loop then behaves as a first-order system in which the difficulty approaches its asymptotic, stable value as a decaying exponential, whose time constant decreases with the gain of the integrator in the outer loop, and decreases with the rate of change of the amplitude of limit-cycle relative to the difficulty.

4.4. BEHAVIOUR OF FEEDBACK TRAINER

The obvious means of verifying the approximate stability analysis is by experimental study, and four examples of the variation of difficulty with time, using the feedback trainer of Figure 3 with non-adaptive controllers, are given in Figure 4. A and B were generated by human operators (skilled pilots using a joystick control), whilst C and D were generated by simple relay controllers. The parameter of difficulty, in this case, is the undamped natural frequency, with the damping-ratio fixed at unity and the disturbance at a low level. The asymptotic value of natural-frequency is a measure of the skill of the pilots, or the goodness of the automatic controllers. It may also be thought of as defining the stability boundary of the pilots and controllers with respect to the natural frequency of the controlled element. For frequencies much below the asymptotic value the loop is unstable, whilst above it the loop is stable.




Fig. 4. — Response of a feedback trainer

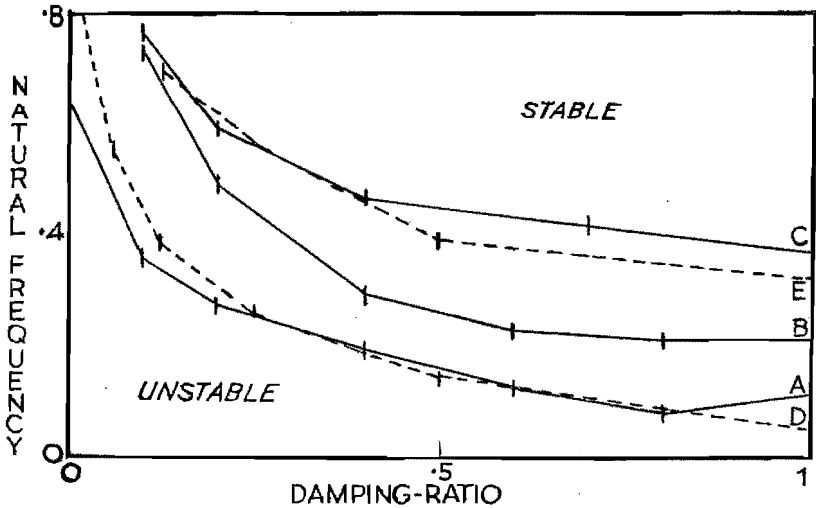


Fig. 5. — Stability boundaries for human and automatic control

By measuring the asymptotic value of natural frequency for different values of damping ratio, the stability boundary for a controller in the natural frequency/damping ratio may be measured. Figure 5 shows such boundaries for three human operators (A, B, C), and two relay controllers (D, E). Used in this way the feedback trainer may be regarded as a device for *testing* the skill of the human operator along two dimensions.

When the operator attached to the feedback trainer is adaptive and improves his control policy with experience, then the asymptotic value of difficulty will not be reached as rapidly, if it exists, as those of Figure 4. Instead, the difficulty will initially rise rapidly until the mean error modulus is at the prescribed level, and will thereafter vary to follow any changes in the operator's ability. Most importantly, it will maintain the desired sub-environment, whilst all the time minimizing the influence of the training controller and hence maximizing the extent to which the operator is performing the required task. Figure 6 shows the variation of difficulty with time for human operators learning the control task with push-buttons. Within 10 minutes B has adapted to a criterion of satisfactoriness corresponding to the $D = .5$ level of difficulty, whereas A takes over 20 minutes to become adapted. C never adapts to this criterion, although some learning can clearly be seen.

Figure 7 shows equivalent variations of difficulty with time for learning machines undergoing feedback training. A and B show learning to a high level and to a medium level, similar to that of human operators B and C in Figure 6. Machine C, however, shows (broken line) a new variety of learning behaviour, in that its control policy rapidly becomes satisfactory for a high level of difficulty, but then declines slowly in its effectiveness.




Fig. 6. — Feedback training of human operators




Fig. 7. — Feedback training of learning machines

Instability of adaption was also shown by the human operators, especially in the early stages of learning. As radical an example as that of Figure 7, however, was obtained only once, during some preliminary experiments. It was ascribed at the time to fatigue, boredom, or some such other convenient psychological variable. In retrospect, because it is not so easy to dismiss a machine's behaviour in this way, such negative learning, or mal-adaption, appears of great importance. The learning machine did not suffer from muscular fatigue, neither did it become bored or lose concentration.

One may only suppose that the changes in the sub-environment brought about by adaption of the control policy were such as to induce mal-adaption. In the human operator this phenomenon may be accompanied by complaints as to boredom or fatigue, but these do not *explain* the mal-adaption.

These experiments demonstrate that the automated feedback trainer behaves as expected, but do not show any necessary advantage over simpler training procedures. Formal experiments have been carried out to evaluate the relative merits of different training techniques, using the task and trainer described above. The results are summarized in the following section.

5. Summary of Experiments to Evaluate Feedback Training

Experiments have been carried out with a homogenous population comprising 72 R.A.F. pilots at an advanced stage of selection and training, in order to evaluate the relative merits of open-loop and feedback training. The effects of various forms of instruction and verbalization on the learning were also investigated.

Operators were trained under one of three conditions : high difficulty; low difficulty; or feedback, in which the level of difficulty was adjusted automatically to maintain their mean error constant. In each group, half the operators were given strong, or informative, instructions which explained to them the nature of the controls, whilst the others were given weak, or non-informative, instructions which told them only what they were required to achieve.

All the operators were tested (mean error measured) after training, first at the high level of difficulty and then at the low level. The test at the high level was in two stages, one incorporated at the end of the final training session without informing the operator, the other directly afterwards when the operator was warned of the test; this enabled the effect of instruction-induced stress to be measured. The operators also filled in questionnaires which enabled their stated interest and evaluation of the training situation to be measured, and which were open-ended as far as comments were concerned so that an estimate of the operator's verbalization could be obtained.

The main results of the experiments are as follows :

(a) The operators trained at a high level of difficulty show little or no learning and perform badly on all tests. The strong instructions have a significant effect in improving learning, but do not overcome the operator's basic difficulties. The high level of difficulty is not in itself unattainable, however, since over 65 % of the feedback grouped reached it, or a much higher level, during their training, and hence could perform well at the high level.

(b) The operators trained at the low level of difficulty split clearly according to the instructions given—those with weak instructions do not show appreciably better performance than those trained at a high level, whilst

those with strong instructions show a spread in performance from very good to very poor throughout the tests.

(c) The operators trained under feedback conditions all learn to a high standard. Those with weak instructions do not differ significantly in their performance from the group who trained at a low level with strong instructions. The feedback group with strong instructions are significantly better than all other groups.

(d) The overall effect of informative versus non-informative instructions is that strong instructions give significantly improved performance in all groups.

(e) The effect of instruction-induced stress is that operators trained at a high level of difficulty get worse, operators trained at a low level do not vary appreciably, whereas operators trained under feedback conditions get significantly better; there is no interaction with weak and strong instructions. This is the only difference in performance which differentiates the group trained at a low level with strong instructions from those trained under feedback conditions with weak instructions.

(f) The questionnaires show no appreciable difference in the interest expressed by the various groups. There is a marked difference in the variance of the estimates of the difficulty of the control task. The group trained at a high level with strong instructions show significantly more verbalization than the other groups.

(g) In terms of 'transfer of training' the results demonstrate clearly that the question as to whether easy-difficult or difficult-easy transfer is best is not meaningful. They indicate, however, that transfer from a difficult level to an easy one is better than learning solely on the easy one, provided that the difficult level is within the operator's ability to perform reasonably well.

(h) Since the operator's skill varies with learning, this also indicates that the optimum level of difficulty must be selected according to the operator's ability. The automatic feedback training loop used in these experiments is shown to be effective by the results, since, under the same conditions of instruction, the feedback group performed significantly better than the others under all test conditions.

(i) Informative instructions are shown by the results to be capable of compensating for poor training conditions, provided that these are not too poor. However, even though the feedback group with weak instructions could not be distinguished from the low level group with strong instructions on performance alone, they showed significantly less deterioration of performance under instruction-induced stress.

(j) The results with computer-simulated learning machines parallel those with human operators completely. This demonstrates not only the utility of these devices as dummy operators but also the generality of the learning and training phenomena investigated.

(k) The theoretical studies indicate that the effectiveness of feedback training in these experiments may be attributed to the interaction between learning to use the controls and learning to control the system dynamics. Training any task that involves learning at least two skills, of which one must be performed satisfactorily to provide a normal environment for the other, and vice versa, is liable to benefit from the use of an automated feedback trainer.

6. Conclusions and Acknowledgements

The studies and experiments outlined here obviously form only a preliminary approach to the development of *useful* feedback trainers which may be incorporated in practical training problems. They indicate, however, that the problems of controlling human learning are amenable to theoretical analysis within the bounds of our present knowledge, and that feedback training systems perform favourably compared with open-loop techniques in at least one training situation. The theory indicates that this advantage should increase when the feedback training technique is applied to the multidimensional, interacting sub-skills that must be acquired in learning to perform real control tasks, such as flying high-performance aircraft. The next stage of development is to investigate the application of feedback training techniques in the real world of flight simulators, rather than the academic world of tracking tasks.

This work was supported by funds from the Ministry of Defence, and I am grateful to Dr J. C. PENTON for his support and encouragement. I have omitted detailed references in this very much curtailed outline of the work, but would like to express my gratitude to Dr G. PASK and Dr E. M. HUDSON for access to their own work on adaptive training.

DISCUSSION GENERALE DE LA SESSION IV

Dr PASK : "The point was raised by Mr. DUKE of the practical application of assisted instructional systems and presumably of all complex adaptive training systems—the question being why don't we use the computer to control the total training strategy.

For example a circular system can be devised for specialised training, where students after following various fixed programmes and laboratory work are tested. The results are fed into a computer which according to several criterion selects the optimum strategy of instruction for the individual. A computer could in this way organize an institution or a school, students being tested periodically, daily, hourly even every minute if this was appropriate. This is a articulated approach—the computer organizing many different systems.