1968 IEEE International Convention Digest
New York, March 18th - 21st

THE CONTROL OF HUMAN LEARNING

B.R.Gaines
Department of
Electrical Engineering Science,
University of Essex,
Colchester, Essex, U.K.

The designer of a system containing human controllers has a synthesis problem which, although similar in some respects to the classical problems of automatic controller synthesis, differs from them in that the human controller already exists as a complete object and is highly adaptive. Fabrication of an arbitrary controller is thus not possible, and even parameter-setting in any conventional sense is difficult (through the use of language). Techniques for the 'synthesis' of the human part of the control system have therefore developed in different directions to those of classical synthesis, and are based largely on the adaptive capabilities of the human operator.

This paper introduces the control problem involved in the regulation of human learning through training, suggests a suitable controller, or automated feedback trainer, for solving this problem automatically, and summarizes the results of some experiments with such trainers in which the interaction of verbal instruction and training techniques was investigated for both human operators and computer-simulated learning machines.

Training as a Control Problem The adaptive behaviour of a learning system or adaptive controller can be classified formally by supposing that its interaction with its environment is divided into tasks. Where a task might be interaction with a given environment for a prescribed length of time, and the purposefulness of learning is taken into account by postulating that the performance of a task is either satisfactory or unsatisfactory. Since the controller is adaptive the effect of performing a task will be to change its state. If the controller becomes adapted to the task then its state will eventually lie in some set for which its performance is always satisfactory for that task.

It may happen that given a task repeatedly the controller never becomes adapted to it, but that there is some sequence of alternative tasks which brings it to a state where it is adapted, or can adapt, to the required task. These tasks form a training sequence and the selection of appropriate tasks is obviously a control problem in the space of states of the adaptive controller.

Modes of Training According to the manner of selection of the training sequence various modes of training may be distinguished:-

(i) Fixed training - in which the trainee is immediately presented with the final task to be learnt. This relies entirely on the adaptive capability of the trainee and no control is exerted.

(ii) Open-loop training - in which some preliminary sequence of tasks is given to the trainee. Initial training on an 'easier' task is an example of this category - no account is taken of individual ability or state of learning.

(iii) Feedback training - in which the sequence of tasks given to the trainee is varied according to information about his state of adaption.

Structure of a Feedback Trainer The theory of adaptive behaviour does not, in itself, indicate a suitable structure for feedback training. This may be derived from consideration 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, the dual control problem. A given control policy restricts the states and state-transitions of the environment to some sub-set of the total possible behaviour, a sub-environment. The desired sub-environment, generated by the optimal policy of a satisfactory control policy, may be entirely different from that generated by the initial policy of a naive controller, and adaption which takes place in it may be irrelevant or even deleterious. The aim of the training system thus should be to maintain the desired sub-environment regardless of the control policy of the trainee.

The additional control loop around the environment, necessary to maintain the desired sub-environment, may be said to be closed by a training controller, the selection of which may be ascribed to a trainer having access to information about the state of the trainee. From an alternative viewpoint, the trainer, by changing the training controller, may be said to vary the difficulty of the environment according to the state of the trainee.

Experimental Feedback Trainer These considerations have been applied to the design of a feedback trainer for human operators and adaptive controllers learning a novel tracking skill. The environment chosen is a pure third-order transfer function which, when coupled to the training controller, becomes a single integration following a stable second-order term of variable natural-frequency and damping-ratio. The inputs to this are a disturbance of variable amplitude and impulses from two push-buttons acting as manual controls. The output is displayed on a CRT and the controller's task is to regulate the system so that its output is zero. The push-buttons incorporate memory so that the trainee has not only to learn to control a high-order system but also to use the manual inputs correctly.

The desired sub-environment is a region about zero in the state-space of this system, and within it the system behaves linearly. A control policy which makes the system unstable generates a sub-environment around the boundaries of the state-space, where the system is nonlinear. The trainer varies the training controller (thereby changing the damping-ratio and disturbance of the system) according to the modulus of the error. When the system is outside the desired sub-environment the error magnitude is large and the trainer decreases the difficulty of the environment - otherwise it increases it. This is a stable strategy for adjusting the training controller, provided certain reasonable conditions are met, and the difficulty rapidly levels out at a value where the desired sub-environment is just, but only just, maintained. As the trainee learns, the trainer reduces the effect of the training controller until satisfactory performance is taking place in the required environment.

Summary of Experimental Results The feedback trainer described above has been evaluated using 72 RAF pilots and a variety of learning machines as trainees. Operators were trained under one of three conditions: high difficulty, low difficulty, or feedback. In each group, half the operators were given strong, or informative instructions, which explained to them the nature of the controls, whilst the other half were given weak, or non-informative instructions. All the operators were tested both at the high and low levels of difficulty, and administered questionnaires to test their attitudes and verbalization.

(I) Operators trained at a high level of difficulty show virtually no learning; those at a low level show wide variation correlated with the strength of instruction; those under feedback training learn to a uniformly high level.
(II) Instruction-induced stress caused the performance of operators trained at a high level of difficulty to deteriorate; did not affect that of operators trained at a low level; and improved that of operators trained under feedback.
(III) There was no overall difference in the results with human and adaptive controllers.

References Gaines, B.R., Teaching Machines for Perceptual-Motor Skills, in Aspects of Educational Technology, Methuen 1966;
Adaptive Control Theory, in Encyclopaedia of Information, Linguistics and Control, Pergamon Press 1967;
Automated Feedback Trainers for Perceptual-Motor Skills, MoD Final Report, 1967.