Int. J. Man–Machine Studies (1972) **4**, 169–199

Axioms for Adaptive Behaviour

B. R. GAINES

Department of Electrical Engineering Science, University of Essex, U.K.

(*Received 28 March 1972*)

The study of adaptive behaviour, both natural and artificial, has developed without any generally accepted definitions of the terms " adaptive " or " adapted ", and it has been argued that a formal definition of these terms is not possible. However, if only for purposes of meaningful communication, some form of accepted definition would be useful, and it is possible that an investigation of the " impossibility " of definition might throw some light on fundamental problems concerning adaption and learning processes.

This paper commences with a critical discussion of previous attempts to define terms concerned with adaption and learning. It is concluded that these terms cannot meaningfully be defined absolutely, and do not have a single connotation, but that rigorous definitions of a variety of phenomena with connotations of adaption may be given relative to (arbitrarily) defined rules for describing behaviour. An extensive set of definitions of behavioural phenomena in these terms is given. Finally a machine, the adaption automaton, is defined whose behaviour is equivalent to that of the adaptive system and the various definitions are re-phrased in terms of the structure of the adaption automaton.

1. Introduction

The terms " adaptive ", " adapted ", and " learning ", are ones which originally arose in the biological sciences to denote the plasticity of behaviour shown by an organism in its struggle to survive in a novel or changeable environment. In recent years the same terms have come to be applied in the engineering sciences to systems designed to optimize their performance through interaction with their environment, and since these, or similar terms form an essential part of the statement of objectives of major areas of research into adaptive control and artificial intelligence, it is reasonable to expect the terms to be capable of fairly exact definition. This is not so, however, and the terms are used very loosely with tacit switches of connotation, particularly between " structural " and " behavioural " aspects of adaption. In this first section previous attempts at definition are critically analysed to demonstrate the variety of connotations of the term " adaptive ".

Stanier, at a symposium on " Adaptation in Micro-organisms " (1953), defined the term " adaptive " as " the totality of the various processes of change which confer on an organism fitness to its environment ". This definition contains all the elements essential to a taxonomy of adaptive behaviour. There is firstly the organism, which in the present discussion will be called a *controller* since both natural and artificial systems are being considered. There is secondly an *environment*, which term has entered the vocabularies of both psychology and engineering. There is thirdly the evaluative concept of " fitness " to the environment, which reflects a description of the organism's behaviour in terms of goal-seeking, and entails a *performance measure* of the interaction between controller and environment. Finally, there is the concept of " change " in that an organism is not initially fitted or its environment but becomes so through a *dynamic* process of adaption.

Adaption, viewed in this way, may be treated as a purely behavioural concept, since there is no necessity to introduce notions of how the organism adapts to its environment, or to argue *a priori* what factors will cause adaption to take place. Stanier's definition is informal, and non-operational in that it does not contain a decision procedure to determine when, or what form of, adaption takes place. Most attempts to treat learning and adaption more formally have gone beyond the observed behaviour and introduced structural or epistemological considerations, generally because their aims have been to model or predict behaviour. It is possible, however, to formalize the purely behavioural concept of adaption, and the following sections review the few attempts to establish operational definitions of adaptive, or goal-seeking, behaviour.

1.1. THE " ANALYTICAL BIOLOGY " OF SOMMERHOF

An early attempt to give a rigorous behavioural definition of the concepts of purpose and goal-seeking in biology and psychology was made by Sommerhof in his book *Analytical Biology* (1950). Sommerhof considers an environment which has a number of states, a regulator (controller) which also has a number of states, and a set of outcomes which are determined by a combination of the environment's state and the controller's state. Some of these outcomes are satisfactory and satisfy a " focal condition ", whilst others are not. Sommerhof proposes that goal-seeking may be said to occur when there is " directed correlation " between the states of the controller and those of the environment such that the outcome is satisfactory. He further proposes a " degree of goal-directedness " in terms of the set of states of the environment for which the outcome is satisfactory.

Sommerhof qualifies his definition with some constraints designed to

ensure non-triviality of the goal-directedness. The state of controller must not, in itself, ensure the satisfactoriness of the interaction—that is, the controller must in some sense, " take note " of the environment. The state of the controller and that of the environment must be " epistemically independent " —that is, there must not be a natural physical connection between them such that the interaction is bound to be satisfactory. Finally, single occurrences of satisfactory interactions do not show goal-directed behaviour, but the directed correlation must exist between a number of environment and controller states.

Although Sommerhof introduces the term " state " of the controller, this is irrelevant to his criterion for goal-directedness, which depends only upon the satisfactoriness of the interaction between controller and environment, and hence is purely behavioural. In control-theoretical terms, what he proposes is a " sensitivity analysis " (Radanovic, 1966) of the interaction between controller and environment, with goal-directedness being evinced by insensitivity of the satisfactoriness to disturbances of the environment.

Sommerhof's definition clearly applies to a simple servo-mechanism, and he gives a gun-aiming servo as an example of a goal-directed system. The constraints which he places upon the nature of the behaviour in order to distinguish between the " goal-directedness " of the simple servo, and the " non-goal-directedness " of, for example, a pendulum, do not appear to be essential, and detract from the main argument. The " physical law " relating position and acceleration of a pendulum is due to its constrained motion in a gravitational field, and is thus a function of the structure of the pendulum and its relationship to its environment. Similarly, the corresponding law between the position and acceleration of the armature of a servo motor is a consequence of the structure of the servo system and its connections to a load. Any behaviour of any system, no matter how complex, is physically determined. Before the advent of Newtonian dynamics, the behaviour of pendulums was equally as mysterious as that of living creatures. It is true that lack of understanding of the mechanism of a phenomenon may lead us to give it undue weight and importance, but this cannot form the basis of a logical distinction.

The main criticism of Sommerhof's approach is that it does not take into account the dynamics of adaptive behaviour. A controller does not generally change its state instantaneously according to the state of the environment, and the matter in which it changes state is a function not only of the immediate environment but also of its previous states. It is this sequential dependence, or memory, inherent in the behaviour of most adaptive systems, which gives rise to transfer effects in training, and, indeed, to most of the complexity of adaptive behaviour.

1.2. ASHBY'S FORMULATION OF " DIRECTED CORRELATION "

Ashby (1962) has given a set-theoretic formulation of Sommerhof's definition of goal-seeking behaviour, which is itself an advance in the theory of adaptive behaviour. Ashby considers a set of disturbances, D, which cause some changes in the environment. He defines these changes by a mapping, φ, from D into the set of possible values of the parameters of the environment, E. A disturbance, $d \varepsilon D$, at time t_0, produces an effect in the environment, $e \varepsilon E$, at time t_1, such that

$$e=\varphi(d). \tag{1}$$

The parameters of the goal-seeking system are similarly specified by a set, F, and its behaviour is a response to d, which may be represented by a mapping, μ, from D into F, such that f, the system response, satisfies—

$$f=\mu(d). \tag{2}$$

When the disturbance has evoked responses in the environment and goal-seeking system, $\varphi(d)$ and $\mu(d)$ respectively, then these two values interact to give some final outcome at time t_2. This corresponds to a mapping, η, from the product set, $E \times F$, into Z, where Z is the set of all possible outcomes when E and F range uncorrelatedly over all their values. Within Z is the sub-set, θ, of outcomes that satisfy the focal condition and are satisfactory.

Sommerhof's " directed correlation " is now defined as being shown by μ, in respect of D, φ, η and θ, if and only if

$$\forall d \varepsilon D, \eta(\varphi(d), \mu(d)) \varepsilon \theta \tag{3}$$

—that is, for any disturbance the outcome is satisfactory.

Ashby manipulates this result into a neater, and more intuitively satisfying, form, using his development of the set-theoretic terminologies of Bourbaki and Riguet (Ashby, 1962). In this notation a function-name, such as φ, standing on its own denotes the sub-set of the product set between its domain and range such that the functional relationship is satisfied. The inverse, φ^{-1}, denotes the corresponding sub-set of the product set between its range and domain. The composition of two such sets, $A \subset I \times J$ and $B \subset J \times K$, is denoted by $A.B$, and is a sub-set of the product set $I \times K$, such that

$$\forall (i,k) \varepsilon A.B, \exists j:(i,j) \varepsilon A, (j,k) \varepsilon B \tag{4}$$

—the set J is eliminated by this composition.

Using this notation, Ashby remarks that the expression for directed correlation can be simplified algebraically by noting that the set specified by

$$[\forall d \varepsilon D, (\varphi(d), \mu(d))] \tag{5}$$

—is identical with the set $\mu.\varphi^{-1}$, with D eliminated by composition. The criterion for directed correlation may be written

$$\forall d \varepsilon D, (\varphi(d), \mu(d)) \subset \eta^{-1}(\theta) \tag{6}$$

—which is thus equivalent to

$$\mu.\varphi^{-1} \subset \eta^{-1}(\theta) \tag{7}$$

—which in its turn may be re-written

$$\mu \subset \eta^{-1}(\theta).\varphi. \tag{8}$$

This is an elegant formulation of Sommerhof's definition of goal-directedness, in which the essential features stand out clearly and unnecessary use of the symbolism of differential calculus and continuous variables is avoided. In particular, although Ashby introduces times, t_0, t_1 and t_2, copied from Sommerhof's original argument, they are completely irrelevant and do not appear in the final result. The "directed correlation" defining goal-seeking is presented in a completely abstract form as an inclusion relationship between sets forming the domains and ranges of certain mappings.

Stripped of the time-dependencies, the definition appears so simple as to be almost trivial. Indeed, if one labels the parameters of the environment by the disturbances that cause them (or consider the parameters of the controller to reflect those of the environment rather than those of the disturbance) then φ becomes an identity mapping, and the criterion is merely that

$$\forall d \varepsilon D, \eta(\mu(d)) \subset \theta \tag{9}$$

—which states that the evaluation of the controller's policy is uniformly satisfactory. However, the simple and clear notation is in itself an invitation to introduce a state variable into the controller's parameter mapping. For example

$$f_{n+1} = \mu(f_n, d_n) \tag{10}$$

—which extends equation (2) to take account of an important aspect of the behaviour of the controller in time, previously obscured by the relevant introduction of t_0 etc., that an adaptive controller *becomes* satisfactory through its experience of the environment, and may possibly become unsatisfactory again—adaption is a *dynamic* process.

In summary, Sommerhof's analysis is a considerable step towards a rigorous definition of goal-directed behaviour, but does not take into account the dynamics of adaption. Ashby's set-theoretic formulation of this analysis is a further advance in clarity of exposition, and enables the defects, and possible extensions, of the original definition to be clearly seen.

1.3. BEHAVIOURAL DEFINITIONS OF ADAPTION IN CONTROL ENGINEERING

In the 1930's, about 10 years after Watson laid the foundations of behaviourism in psychology, engineers such as Bode and Nyquist were establishing techniques for the analysis of the structure and behaviour of automatic control systems. Requirements for such systems to operate with plant whose characteristics were not known in advance, e.g. drives with varying frictional loads, led engineers in the 1950's to develop self-optimizing controllers which changed their parameters to maximize their performance. These controllers were termed " adaptive ", since many of the concepts in both standard control systems and self-optimizing controllers were derived by analogy with biological systems.

It is interesting to note that control engineers reached similar disagreements over the definition and application of the term " adaptive " as have behavioural psychologists. In a panel discussion among many of the eminent pioneers of adaptive control, the position was summarized by the statement (Freeman, 1963) that " the best way we can define an adaptive system today is to ask 500 people and we'll get 500 different answers and perhaps if we ask the same people the next day we'll get 500 additional answers".

It is clear on reviewing the controversy that, although engineers were trying to achieve adaptive *behaviour*, in the sense of a changing control policy dependent upon the environment, the criticisms as to whether a controller was, or was not, adaptive were made largely on structural grounds. Thus, Raible (1963) claims that the Moe & Murphy (1962) " self-adaptive " controller is not adaptive because it only identifies plant parameters on-line, whilst " search and hill-climbing " are necessary for true adaption. Since control engineers are primarily involved in the design and fabrication of control systems, it was natural for the engineer to term a controller " adaptive " if its structure had been designed to be so, even if its actual behaviour was extremely mal-adaptive.

Occasionally a distinction between behavioural and structural definitions is noted, but not carried to its logical conclusion. For example, Clark (1963) remarks that most attempts to define adaptive systems " seem to be of an anatomical nature . . . in terms of the physical features of the control system itself . . . I feel that the definition might better be based on a functional definition ". He goes on to suggest—" where we have a plant whose dynamic characteristics vary in time . . . if we can design a controller that will solve our problem to the satisfaction of all who are concerned with the problem, might it not be a functional definition of an adaptive control system ".

However, although such functional, or behavioural, approaches were

proposed and distinguished from structural connotations of the term " adaptive ", the vagueness in the general use of the term continued and came to be accepted. Truxal (1963), in reviewing the field of adaptive control for a major Congress, re-iterated that an adaptive control system was one " designed from an adaptive viewpoint ", and stated that " While the literature of control theory is replete with arguments re the definition, progress in adaptive control theory has not been impeded by the failure of purists to reach universal agreement on an appropriate definition." The attitude engendered by failure to distinguish between structural and behavioural connotations is summed up by Florentin's (1962) wry remark at the end of a paper on an adaptive control system that, viewed in a certain way, his system appears " to be ordinary nonlinear feedback. It seems that any systematic formulation of the adaptive control problem leads to a meta-problem which is not adaptive."

1.4. ZADEH'S DEFINITION OF " ADAPTIVITY "

The distinction between structural and behavioural connotations of the term " adaptive " was first made, in a control engineering context, by Zadeh (1963) in a short paper entitled " On the definition of adaptivity ", which proposed a formal and completely behavioural intension of the term. He notes that " Most of the vagueness surrounding the notion of adaptivity is attributable to the lack of clear differentiation between the external manifestations of adaptive behaviour on the one hand and the internal mechanisms by which it is achieved on the other ". He goes on to propose a characterization of adaptive behaviour, stating that " our premise is that all systems are adaptive, and that the real question is what they are adaptive to, and to what extent ".

Zadeh frames his definition, as does Sommerhof, in terms of continuous functions and functionals. He considers a system subject to one of a set of specified input time functions, $u(t)$, defined on the semi-open interval, $t \geqslant 0$. A family of such inputs, which may have a probability measure defined over it, he defines as a *source*, $S\gamma$. For each source, he assumes that there is a performance criterion such that the behaviour of the system when connected to the source (that is, when receiving an input belonging to the family of inputs which defines the source) is, or is not, *acceptable*. Zadeh then defines " adaptive " in these terms—" A system is adaptive with respect to a family of sources, $(S\gamma)$, and a criterion of acceptability, if it performs acceptably well with every source in the family $(S\gamma)$."

This definition is very close indeed to that of Sommerhof, and lack of reference to the latter illustrates the neglect of his work whose publication

preceded that of Zadeh by some 13 years. Zadeh does not attempt to give the term " adaptive " an absolute meaning and exclude trivial cases, but states quite explicitly " Under this definition every system is adaptive with respect to some set of sources and performance criterion. Thus what matters is not whether the system is adaptive, or not but what are the sources and performance criteria to which it is adaptive."

The same criticism may be levelled at Zadeh's definition as at Sommerhof's —that he does not take account of the dynamics of adaptive behaviour. Indeed, this possibility is excluded by defining inputs on the interval, $t \geqslant 0$, rather than some finite interval. Zadeh takes into account the time variation of performance with learning only through the suggestion that " it is appropriate to use a performance function which assigns low weight to the performance in the initial stages of the learning process ".

It is interesting to note that Donaldson & Kishi (1965) in a textbook on automatic control, give a version of Zadeh's definition of " acceptability " in which they modify the input sources to be defined over a finite time interval. The reason for this is not stated but is inherent in their discussion, for they require an adaptive controller to modify its control action in an attempt to become an acceptably performing system. The acceptability of performance, in their structural definition, becomes something monitored by the controller, and hence regularly evaluated. This change of meaning, from an overall evaluation of the controller's ability to cope with its environment, to a local evaluation of whether it is, or is not, yet coping, is very important in the context of adaptive dynamics. Whilst their modification was necessary to the formulation of a structural definition, it could equally form the basis of a behavioural definition in which performance is monitored not by the controller, but rather by some outside observer.

Both Zadeh's and Donaldson & Kishi's definitions of " acceptability " are important and necessary to a theory of adaption, but a distinction must be made between them. In this paper, the term " acceptable " is applied to the interaction between controller and environment in a global sense, similar to that of Zadeh. Whilst the term " satisfactory " is used in a local sense equivalent to Donalson & Kishi's " acceptability ".

1.5. THE DYNAMICS OF ADAPTIVE BEHAVIOUR

The adaptive control systems studied in engineering have not been very complex, and, typically, they have had the property that either they are able to adjust their parameters to a level suitable to their environment and maintain them there, or they are not; and the ones which are not able to do so tend not to be reported. Any further complexity of behaviour, such as sometimes

adapting successfully and at other times not, or becoming satisfactory for a while and then becoming unsatisfactory again, would be regarded as a defect and removed if possible. Equally, complex behaviour based on interactions between environments, such as success in adapting to one environment only after having adapted to another, would not be noticed since the parameters of the controller would usually be reset upon changing the environment. Thus Zadeh's definition in terms of total success or total failure of adaption, for all time, is comprehensible in the light of engineering experience with adaptive systems.

In psychology, on the other hand, experience has been of systems whose level of complexity engineers are still many years from synthesizing. The human operator most certainly does not have a dichotomous capability in most tasks, and his performance of one is greatly affected by his experience of others. All animals show a similar complexity of behaviour, and it is interesting to compare some definitions of " learning " by behavioural psychologists with those more formal definitions discussed so far.

Guthrie (1952) gives a very simple definition—" These changes in behaviour which follow behaviour we call learning." Hunter (1934) attempts to exclude some behaviour, which is due to changes in peripheral structures, from the definition—" We may say that learning is taking place whenever behaviour shows a progressive change or trend with a repetition of the same stimulating situation and when the change cannot be accounted for on the basis of fatigue of receptor and effector changes." McGeoch & Irion (1952) require a change in the evaluation of the behaviour, rather than merely a change in the behaviour itself—" Learning, as we measure it, is a change in performance which occurs under conditions of practice." Thorpe (1965) requires these changes to be adaptive—" We can define learning as that process which manifests itself by adaptive changes in individual behaviour as a result of experience." Whilst Bush & Mosteller (1955) deliberately exclude this requirement—" We consider any systematic change in behaviour to be learning whether or not the change is adaptive, desirable for certain purposes, or in accordance with any such criteria."

There is one important aspect of adaptive behaviour, common to these diverse definitions, which does not appear in Sommerhof's or Zadeh's formulations, and that is the nature of learning as a *change* in behaviour with practice or experience. There is one engineering paper where the essential nature of the dynamics of adaption is clearly stated. Martens (1959) has proposed a definition of " machine learning " which is intended to provide a test to determine whether learning really does occur. He requires that a machine be able to adapt to two mutually incompatible criteria. Thus, in

terms of Zadeh's definition, the controller would have to have an acceptable performance with each of two source/performance-criterion pairs, such that it is physically impossible for a system with a fixed control strategy to have an acceptable performance with both.

This distinction is not as fundamental as it might appear, since it is clear that the machine must have some means of distinguishing between the two source/performance-criterion pairs, and if this information is considered as part of the source then the adaption is once again trivial on Marten's criterion. More important, however, is the rationale behind the situation which Martens requires to be established as a criterion of real learning. This is such that the controller cannot be adapted to both objectives at the same time, but must *become* adapted to one after being adapted to the other, and vice versa. This connotation of adaptive dynamics is clear enough in structural definitions of adaptive controllers, and in the definitions of learning by behavioural psychologists, but it seems to have been omitted in more formal analyses of adaptive behaviour.

The objective of the theoretical studies described in the remainder of this paper has been to extend the purely behavioural approach to the definition of adaption proposed by Sommerhof and Zadeh to a taxonomy of adaptive behaviour which takes full account of the dynamics of adaption.

2. The Behavioural Intension of Adaption

Given a situation which is to be regarded as that of a controller interacting with an environment, there are a number of decisions to be taken before a discussion of the adaptivity of the controller can begin. Within the theory of adaption these decisions have to satisfy certain logical constraints, but they are otherwise arbitrary. Within the context of human or animal behaviour, and its function, or of adaptive control practice, these decisions will certainly not be arbitrary, for the utility of adaptive concepts in specific situation will depend on them. Some of the decisions will be obvious and unmentioned, and others will be made explicitly. Much of the early controversy over conflicting usage of the term " adaptive " in engineering arose because the " obvious ", tacit decisions of one engineer were not those of another, or because disagreement over specific decisions was wrongly ascribed to the definition of adaption itself.

In attempting to quantify and formalize the application of the term " adaptive ", it is reasonable to expect that confusion is likely to arise over similar issues, and particular emphasis is placed in the following discussion on a clear exposition of those aspects of adaption which are arbitrary and depend

on agreed definitions. In the development proposed here, the arbitrary decisions are localized at the point where the interaction between controller and environment is divided into " tasks ", but in practice such extreme localization is unlikely to be apparent and the arbitrariness is spread more thinly.

Although the terminology used, of controller and environment, is particularly appropriate to control-like problems the theory applies, and is intended to be applied, to all aspects of behaviour involving purposeful interaction with some other system; the illustrative examples are chosen to emphasize this. It also applies to the analysis of any form of adaptive behaviour by animal or machine, and represents a unified approach to the study of adaption and learning.

2.1. THE ADAPTIVE CONTROL SITUATION

The three elements of a control situation are a *controller* interacting with an *environment* for a *purpose*. In biological systems none of these may be well-defined, and the reasonable assignment of roles and purposes is the subject of experimental study and theoretical analysis (Gaines, 1966, p. 339). Even man-made systems do not necessarily split into controller and environment in an obvious way, especially when the whole system is synthetic and contains local regulatory control loops. In the case of the human operator, a similar consideration applies, on the motor side, to such systems as the spinal reflexes controlling muscle contraction, and, on the perceptual side, to the dark/light adaption of the eye. In considering learning within the central nervous system, these peripheral systems are often better thought of as part of the " environment ".

However, the separation between controller and environment can generally be agreed upon, and they are, for the purposes of the theory, defined as black boxes with input and output such that the inputs of one are the outputs of the other. The following examples of environments for human controllers illustrate this concept together with the possible associated problems.

(i) The environment is a vehicle on a road. Its inputs and outputs clearly depend on whether it is being driven or being serviced. For a driver, the inputs to the environment will be through the steering wheel, clutch, etc., and the outputs from it will be visual images of the road ahead, speedometer, and so on, together with acceleration forces, etc.

(ii) The environment is a line figure in two dimensions subject to the transformations of Euclidian geometry. Its inputs are usually regarded as mathemetical operators transforming the figure, and its output as a

mathematical description of the transformed figure. Consideration of input/output at this level of abstraction is an example of neglecting peripheral dynamics, and can create problems. Although the mode of presentation is mathematically irrelevant, a figure presented pictorially may have an entirely different psychological effect from one presented as a set of incidence relationships. Whether the hand drawing a transformed figure is regarded as part of the environment of the human controller, or part of the controller itself, is an arbitrary decision, but it will have some effect on the analysis of the learning process in geometry.

(iii) The environment is another human, subject to natural language communication. Its inputs and outputs may be thought of as linguistic utterances, subject to the same considerations as in (ii).

In the context of a unified approach to natural and artificial systems, it is interesting to note that machines are also available for interaction with environments similar to these three. The first is the type of situation in which one reasonably expects to find some form of servo-mechanism with similar characteristics to the human operator. The second is a typical environment for Newell, Shaw & Simon's (1959) General Problem Solver, and the third is within the scope of Weizenbaum's (1966) ELIZA.

The purpose of each of the controllers in the environments above may be defined intensively, as motion from one location to another, or a proof that two angles are equal, or persuasion to perform some act. However, from the point of view of system behaviour, it is not necessary to define the purpose but only to give some prescription for saying when it has been achieved—that is, some performance measure is required. Sommerhof and Zadeh both define performance measures in numerical terms, and often they will be in that form, but, in fact, the only feature they use of the measures is a dichotic one—is the controller's performance up to the required standard or not? Thus the minimum requirement for the definition of a controller's purpose is a decision procedure which determines, at least occasionally, whether some segment of the interaction between controller and environment has been *satisfactory* or not.

2.2. SEGMENTATION OF THE INTERACTION INTO "TASKS"

For purposes of the behavioural definition of adaption, we are concerned with the manner in which the evaluation of a controller's interaction with its environment changes as a function of that interaction. The expected behaviour is that, if its control policy is not satisfactory for the environment, then it will eventually become so. Thus it must be possible to segment the

interaction between controller and environment into at least two phases, in the first of which it is not satisfactory, and in the second of which it has become so.

This segmentation of the interaction between controller and environment is inherent in the concept that an adaptive controller " becomes satisfactory ", rather than just " is satisfactory ", and is fundamental to the analysis of adaptive dynamics. A further dynamic feature of the behaviour is that the controller should remain satisfactory once it has become so; that is, it should reach a stable condition of satisfactoriness.

The segmentation is clearly present in Sommerhof's concept of a " focal condition " which must be attained, but he is only concerned with whether it is attained and not with what happens before or afterwards. This omission probably occurs because his theory develops from a consideration of such problems as firing a gun at a target, where the line of sight has to satisfy a terminal, or " focal ", condition only at the moment of firing. Zadeh, on the other hand although his definition strongly resembles that of Sommerhof, defines satisfactoriness over a semi-infinite interval rather than as a terminal condition, and makes some concession to the dynamics of adaption by suggesting that the performance criterion should attach less weight to the early stages of learning. These two, taken together, imply that the controller is expected to become adapted, and is, not " acceptable " unless it remains so.

In order to consider controllers which become satisfactory and then relapse, or to consider the effect of learning in one environment on later learning in another, it is necessary to extend the basic segmentation of the interaction which is inherent in the concept of becoming adapted, and analyse the changing evaluation of the controller's performance in greater detail. Donalson & Kishi in their unmarked variation of Zadeh's semi-infinite interval to a finite interval, do just this, primarily because they are considering the structure of a controller which is monitoring its own behaviour. The introduction of a defined time interval as a criterion for segmenting the interaction is unnecessarily restrictive, however, and the following definition of a " task " is a weak extension of the segmentation which is sufficient to form the basis of a taxonomy of adaptive behaviour.

2.3. DEFINITION OF A " TASK "

A *task* is a segment of the interaction between controller and environment for which it is possible to say whether or not the controller has performed satisfactorily. Equivalence relations between tasks (so that it is possible to say, for instance, that an interaction consists of the same task repeated several times) are arbitrary, but will generally follow the natural relationships

between different types of controlled system. A task will typically consist of some specification of plant parameters, initial conditions and period of interaction, together with a tolerable performance level above which a control policy is considered satisfactory.

Within the theory of adaption a " task " is restricted only by the necessity for some procedure to determine whether a particular interaction between controller and environment which constitutes a task is, or is not, satisfactory. In practice, it is convenient to choose the set of tasks in such a way that the segmentation of the interaction into tasks is unique. Any interaction may then be regarded as the performance by the controller of a set of well-defined tasks. Since the theory of adaption deals only with these tasks and the satisfactoriness of the controller in performing them, it is obviously of practical importance that they should be chosen to give adequate information about those aspects of the adaptive behaviour which are of interest. These are metatheoretic considerations, however, and the abstract concept of a " task " forms an adequate basis for an axiomatic theory of adaptive behaviour. The utility of the theory in any problems of practical interest is an empirical finding, and hence the informal justification given for the steps taken so far and the examples of the following sections.

In practice, the segmentation of an interaction into tasks may be performed in many ways—the time of interaction between controller and environment may be fixed—a criterion for the termination in terms of the behaviour itself may be given—the interaction may be terminated as soon as a decision can be made about the satisfactoriness of the controller. The " termination " itself may be purely conceptual, a convenient division of a continuous sequence of behaviour into separate sub-sequences, or it may have a physical reality in that the plant is modified at the termination of an interaction.

The following sections contain two examples of the segmentation of an interaction into tasks, the first from control engineering and the second from simple adaptive pattern recognition.

2.4. A SET OF TASKS FOR A SINGLE-INPUT, SINGLE-OUTPUT SYSTEM

Consider the stable, noiseless, second-order plant shown in Fig. 1, consisting of two integrators in cascade with feedback from the output. Its parameters are the gain, undamped natural frequency and damping ratio. Let a task be defined by ascribing values to these three parameters and to the initial values of the integrators, together with a time-varying demand signal, $f(t)$ for $0 \leqslant t \leqslant T$, and a decision procedure such that an interaction is satisfactory if, and only if,

$$\frac{1}{T}\int_{t_0}^{t_0+T}[f(t-t_0)-x(t)]^2\,dt<E^2 \tag{11}$$

where t_0 is the time at which the interaction starts, $x(t)$ is the output of the plant, and $E>0$ is some tolerance on the r.m.s. error.

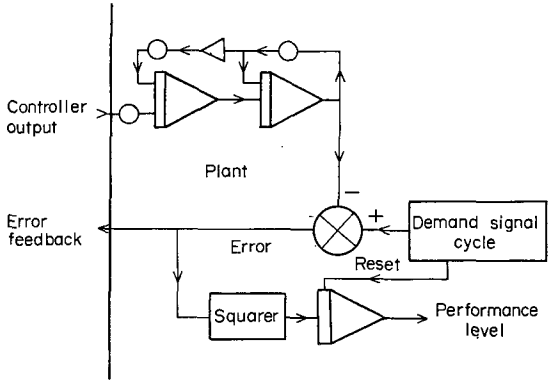

FIG. 1. A "task" generator for an adaptive controller.

To test the adaptivity of a controller to the plant and demand conditions specified by such a task, it is connected to the plant and the demand signal cycled with period T. After every cycle the task is complete, and the r.m.s. error during that cycle determines whether or not the controller has performed satisfactorily. If the controller is adaptive the r.m.s. error in each cycle might be expected to decrease, and hence (by suitable choice of E) the controller will be unsatisfactory initially, but after a number of repetitions of the task it will become satisfactory and remain so.

Many other forms of adaptive behaviour might arise, however—the controller could be always satisfactory or always unsatisfactory—it could start by being satisfactory and become unsatisfactory, never settling at one or the other. If other tasks with different values of the plant parameters or demand signals were interpolated, then the range of possible behaviour would become far greater still. It is the description of this variety of possible behaviours which concerns the behavioural theory of adaption.

2.5. EXAMPLE OF A SET OF TASKS FOR AN ADAPTIVE PATTERN-CLASSIFIER

The two integrator environment described in Section 2.4 is a typical continuous control system for both the human operator and automatic control systems. The analysis of adaption through segmentation of the interaction between controller and environment into " tasks ", however, applies equally

well to discrete, problem-solving environments, such as those involved in pattern-recognition. The following example of a simple, perceptron-like pattern-classifier, learning to dichotomize patterns represented by binary vectors, not only exemplifies an alternative form of task, but also demonstrates sufficient variety and complexity of adaptive behaviour to illustrate later definitions of modes of adaption.

The problem of the pattern-classifier is to assign each of the patterns in an input stream to one of two categories. During the learning phase, it is supplied with information (reward/punishment) as to whether each of its assignments is correct, and uses this to change its categorization policy. Thus the environment of the machine, as shown in Fig. 2, consists of a generator of patterns at its input, an acceptor of assignments at its output and a source of performance feedback at its reward/punishment input. The performance measure for a period of interaction will normally be based on the correctness of its decisions, e.g. the proportion correct.

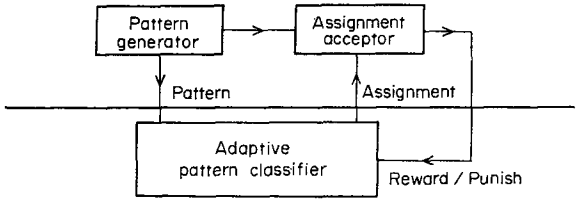

FIG. 2. A "task" generator for a pattern classifier.

A typical " task " for the pattern-classifier may be defined by a set of input patterns, such as pattern$_A$ followed by pattern$_B$, which may be written—$T_1=(B,A)$, together with a performance criterion. There are four possible, non-trivial, performance criteria, that the performance for T_1 is satisfactory if —the category to which A is assigned is correct—that to which B is assigned is correct—either is correct—or both are correct. Given a sufficient set of pattern sequences defining tasks, any stream of input patterns may be split up into segments corresponding to these tasks, and hence the interaction between the pattern-classifier and its environment may be segmented as previously described.

2.6. PROBABILITY AND INDETERMINACY IN THE DEFINITION OF A TASK

It has been tacitly assumed in the Section 2.4 that the initial conditions of the plant may be dropped from the definition of a task, even though this makes the satisfactoriness of the interaction indeterminate. It is often the case that certain potential parameters of a task may be regarded as irrelevant to its

description, because they do not appreciably affect the outcome of the decision procedure for evaluating the performance of the controller in the task. In the last example, if T is long compared with the period of the plant, then the initial conditions will have very little effect on the performance integral of inequality (11). In any practical application there will always be an effect of experimental error on the decision procedure, and indeterminacies having effects comparable with this may be neglected.

It may, however, be desirable to regard as irrelevant parameters which *do* affect the outcome of the decision procedure—for example, a task might be defined, in a similar way to Zadeh's " source ", as a set of fixed plant parameters together with a *class* of possible demand signals, perhaps with a probability measure. The decision as to a controller's satisfactoriness then becomes indeterminate or may be based on an ensemble measure of its performance—for example, the maximum or expected r.m.s. error. Indeterminacy may be accepted as a third form of evaluation, or may be removed by deciding that any unsatisfactory evaluation in an ensemble gives rise to an overall unsatisfactory evaluation. Any ensemble evaluation, although acceptable in theory, leads to acute experimental problems in the evaluation of adaptive behaviour, however, because it is impossible, in general, to reset the state of an adaptive system and repeat the same experiment a number of times. Attempts are made to overcome these problems in practice by the use of populations consisting of different adaptive systems assumed to be exemplars of a single system. This has obvious dangers in itself, but is again a metatheoretic procedure.

Thus, the concept of a " task " encompasses specific descriptions of the possible range of controlled systems, based on distributions over plant parameters, demand signals, disturbances and so on, and, through the segmentation of interactions, it enables the behaviour of a controller coupled to an environment to be regarded as the performance of a sequence of tasks, for each of which it is, or is not, satisfactory. A calculus of adaptive behaviour based on these concepts may be used to evaluate the relative adaptivity of various controllers only in terms of their ultimate satisfactoriness for given tasks, and the path-length (number of tasks) before this is achieved. More complex cost-functions for trajectories of adaptive behaviour could readily be defined, but at our present state of knowledge they are not justified.

3. The Definition of Modes of Adaption

The basic concept of a " task " introduced in the previous sections may be used as the foundation for definitions of different modes of adaptive behaviour

and relationships between them. 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 satisfactorily. The interaction is equivalent to the controller performing a sequence of tasks consisting of the same task repeated indefinitely. Such repetitive task sequences play a special role in the theory of adaptive behaviour because they correspond to the characterization of an adaptive system by its changing response to the same situation. The repetitive task sequence takes the place of the static " source ", or environment, in Zadeh's definition of adaptivity, but is exactly equivalent to it with the addition of segmentation. The importance of this segmentation in the study of adaptive dynamics begins to become apparent in the following version of Zadeh's definition of " acceptability ".

3.1. DEFINITION OF AN ACCEPTABLE INTERACTION

An interaction between controller and environment consisting of the repetition of a single task is *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. However, it must eventually become satisfactory and remain so—an acceptable interaction is one which reaches a *stable* condition of satisfactoriness. In this condition we may say that the controller has become " adapted " to the task.

3.2. DEFINITION OF A CONTROLLER ADAPTED TO A TASK

An interaction between controller and environment consisting of the repetition of a single task is *immediately acceptable* if it is always satisfactory —an immediately acceptable interaction is obviously acceptable. A controller in such a condition that it would have an immediately acceptable interaction with a task is *adapted* to that task.

3.3. DYNAMICS OF ADAPTION

The concepts of acceptability and adaptedness concern the performance by the controller of a single task, but the adaptivity of the controller is generally advantageous because the particular task it must perform is incompletely specified or may change. The possible varieties of adaptive behaviour in a control situation involving performance of any one of a number of tasks are many, but there are three modes of adaption of particular interest which are defined in the following sections.

Very often an adaptive controller will be required to perform a single task which will not change, but whose characteristics cannot be specified in

advance. It must be capable of having an acceptable interaction with any of a range of tasks, but need not necessarily be capable of adapting to a sequence of different tasks. This mode of behaviour is characterized in the following definition.

3.4. DEFINITION OF A POTENTIALLY ADAPTIVE CONTROLLER

A controller in such a condition that it will have an acceptable interaction with any one of a set of tasks is *potentially adaptive* to that set of tasks.

A potentially adaptive controller fulfils one function of an adaptive system in compensating for ignorance about the nature of its environment. It will not necessarily fulfil another major function by performing satisfactorily in a changing situation, since there is no implication that, having adapted to one task, it remains potentially adaptive to others. Potential adaption is implied in statements like, " a shoe adapts itself to the shape of a foot ", and is the weakest form of goal-attainment to merit the designation " adaptive ".

A controller which has to perform satisfactorily in a changing situation must not only adapt to its immediate task, but must also remain potentially adaptive to the other tasks it may meet. This mode of behaviour is characterized in the following definitions of compatible adaptivity.

3.5. DEFINITION OF A COMPATIBLY ADAPTED CONTROLLER

A controller is *compatibly adapted* to one task with respect to set of tasks if, in an interaction consisting of the repetition of that task, it is not always satisfactory but also potentially adaptive to the set of tasks.

A controller which is compatibly adapted to one task with respect to a set of tasks is clearly " adapted " to the first task, and, whilst performing it, never loses its capability for adapting to any of the set of tasks. If the controller is to function in a changing environment, it must always, when adapting to one task, become compatibly adapted with respect to the remaining tasks it may meet—that is, it must be potentially adaptive to all its possible tasks no matter which of them it has previously performed.

3.6. DEFINITION OF A COMPATIBLY ADAPTIVE CONTROLLER

A controller is *compatibly adaptive* to a set of tasks if, given any sequence of tasks from that set, it remains potentially adaptive to the set of tasks.

Thus, a controller which is compatibly adaptive to a set of tasks will have an acceptable interaction with any one of them, and, no matter how it becomes adapted to one of them, it will be compatibly adapted with respect to the remainder.

The phenomenon of potential, but not compatible, adaptivity is very interesting and quite common in biological adaption, both phylogenic and ontogenic, and in animal, even human behaviour. A microbe may adapt to a new culture to such an extent that it becomes dependent on it and dies when returned to its former environment. A species may evolve under the influence of a climatic change but come to an evolutionary dead-end such that it cannot reverse the process when the climate returns to its original form. A human subject may be able to learn to solve either of two problem types equally well, but, after having learnt one and performed it for some time, he acquires a " set " towards the processes involved in its solution and, when given the other, cannot discover how to solve it.

A controller which is compatibly adaptive to a set of tasks is not necessarily able to become adapted simultaneously to all of them. It is, however, quite possible for two tasks to be so similar that a controller which is adapted to one may also be adapted to the other. This mode of behaviour is characterized in the following definitions of joint adaptivity.

3.7. DEFINITION OF A JOINTLY ADAPTED CONTROLLER

A controller is *jointly adapted* to a set of tasks if, given any sequence of tasks from that set, it remains adapted to every member of the set.

Thus, a controller which is jointly adapted to a set of tasks will be always satisfactory given any sequence of these tasks. This is a very strong condition, and an even stronger one is that when a controller adapts to any one of a set of tasks it should eventually become jointly adapted to all of them.

3.8. DEFINITION OF A JOINTLY ADAPTIVE CONTROLLER

A controller is jointly adaptive to a set of tasks if it is both compatibly adaptive to the set and, during an acceptable interaction with any task in the set, it eventually becomes jointly adapted to the whole set.

Joint adaptivity would obviously be expected to be far rarer than compatible adaptivity, since it requires either that the same control policy be appropriate to two different environments, or that a very rapid change of policy be made when the environment changes. The former phenomenon arises, fairly trivially, when the environments are themselves very similar—for example, riding a scooter and then a motor-bike—and the latter is of greater interest. The occurrence of joint adaption when it is logically impossible for the same control policy to be satisfactory in both environments—for example, when the joy-stick in a tracking task is reversed in sense—implies that the learning of one skill does not greatly interfere with the skill already acquired.

3.9. INTERRELATIONSHIPS BETWEEN DIFFERENT MODES OF ADAPTION

The preceding definitions of different modes of adaptive behaviour have been given in order of increasing strength, for if a controller is jointly adaptive to a set of tasks it is also compatibly adaptive to them, and if a controller is compatibly adaptive to a set of tasks it is also potentially adaptive to them

$$\text{Jointly Adaptive} \rightarrow \text{Compatibly Adaptive} \rightarrow \text{Potentially Adaptive.}$$

These three modes of adaption are by no means exhaustive, and many variations are possible, defining other modes of adaptive behaviour. However, many of these would be regarded as pathological, serving no useful function and having no correspondence to a well-defined type of adaptive behaviour. Other, more interesting varieties of adaptive behaviour may be described in terms of the modes already defined. The questions which are usually of interest are whether a controller is potentially adaptive to a set of tasks, and whether it has any states in which it is compatibly or jointly adaptive to them. The modes of adaption which have been defined form explicata for all the common stereotypes of adaptive behaviour.

3.10. INTERRELATIONSHIPS BETWEEN TASKS

The definitions of *adapted*, *compatibly adapted*, *jointly adapted*, and *potentially*, *compatibly* and *jointly adaptive*, may be used to define binary relations on the set of tasks relative to a given controller—for example, task_1 is related to task_2 if, and only if, the controller is compatibly adapted to task_1 with respect to task_2. All six relations are reflexive, and only that induced by " compatibly adapted " is not symmetric. However, only " adapted " and " potentially adaptive " induce relations which are also transitive (and hence are equivalence relations). For instance, a controller may be jointly adapted to task_1 and task_2, and also jointly adapted to task_2 and task_3, but given a sequence containing both task_1 and task_3 there is no reason why even its potential adaptivity to both tasks should not disappear. It is this lack of equivalence relations which gives adaptive behaviour its extraordinary richness. A controller which showed no " pathological " behaviour would be very rare, although the more drastic forms would not be expected to occur. For example, the relation induced by " compatibly adaptive " ought to be one of equivalence, because no sequence of normal tasks should be able to destroy a controller's ability to adapt to one of them.

The binary relations over tasks, generated by consideration of a controller's adaptive behaviour, may be used as the basis for a taxonomy of environments according to the problems involved in adapting to them. For example, if a family of controllers, in adapting to one task, became jointly adapted to

another, then it would be reasonable to suppose that the two tasks were very similar in the control strategy required. If it were found that this was not so, but that potential adaption to one task implied potential adaption to the other, then it would be reasonable to suppose that the two tasks were related in the adaptive capabilities which they required—this is the basis of some " intelligence tests ".

3.11. ARBITRARINESS AND TRIVIALITY IN THE DEFINITIONS OF " ADAPTIVE "

The behavioural definitions of " adapted " and " adaptive " contain an arbitrary element because the classification of " tasks " is left undefined, and the segmentation of the interaction between controller and environment is at will within the (very weak) constraints of Section 2.3. This arbitrariness need not cause difficulty in the analysis of adaptive behaviour, provided it is accepted that at some stage in the discussion of an adaptive controller and its behaviour this classification must be agreed. Much of the early controversy over the application of the term " adaptive " arose because the " obvious ", tacit classification of one psychologist, or engineer, was not that of another, or because disagreement over such a classification was wrongly ascribed to the definition of adaption itself.

Even when the arbitrariness in the definitions is accepted, there remains the possibility that some types of adaptive behaviour may be " trivial ". For example, " jointly adaptive " is an apparently very strong condition which may be quite trivial in reality—for example, the tasks to which a controller becomes jointly adapted may be completely equivalent and need not be distinguished. " Potentially adaptive " is a very weak condition which may often be regarded as trivial, because it is shown by systems undergoing an irreversible descent to equilibrium. " Compatibly adaptive " adds the requirement of reversibility, and is closest to what is commonly regarded as being " really " adaptive. However, although a compatibly adaptive controller shows all the behaviour which one would expect of a " really " adapted to all of its tasks all the time, and hence show no adaptive dynamics—it is just a very good, but static, controller!

In testing the behaviour of an animal, or automatic controller, for " adaption " or " learning " it may be desirable to eliminate this " trivial " adaption. It is meaningful, for example, to ask whether an animal performs well in two different situations because it has a policy suited to both, or because it changes its policy according to the situation. To force a controller to show adaptive dynamics, one might say that it is " really adaptive " to a task if it was an acceptable, but not immediately acceptable, interaction consisting of the

repetition of that task. The controller cannot then be a static system which happens to have a suitable control policy.

Sommerhof (Section 1.1) formulates this non-triviality requirement by requiring that the state of the controller must not, in itself, ensure the satisfactoriness of the interaction. A particularly neat, and operational, formulation of the same requirement is given by Martens (1959) in his definition of " machine learning ". He requires that the controller have acceptable interactions judged by two incompatible performance criteria. It cannot then have a policy which satisfies both criteria at once, and hence it is not immediately acceptable for one task when it is adapted to the other. This may be rephrased that the controller is " really adaptive " if it is compatibly adaptive to a pair of tasks such that it is impossible for any controller to be jointly adapted to them both. However, it will be noted that all these specific formulations are encompassed in the present one.

4. An Automata-Theoretic Formulation of Adaptive Behaviour

In Sections 2 and 3 an axiomatic formulation of the nature of adaptive behaviour has been given in terms of tasks and the satisfactoriness of their performance. This formulation has been deliberately maintained in non-mathematical form, partly for purposes of presentation, but also because premature mathematization obscures the basic logical nature of the definitions which are independent of the notation used in their expression. However, it is useful in further extensions and applications of the theory to have some particular model in mind of the system which is giving rise to the adaptive behaviour. The definitions of modes of " adaption " may then be phrased in terms of the structure of this model.

One convenient model of the adaption of a system is that of a state-determined (or state-restricted) machine whose inputs are " tasks " and whose outputs are the " satisfactoriness " of the behaviour given the current state of the machine and the current task. The use of such a model need not necessarily detract from the essentially behavioural nature of the present study since it may be chosen to be *cybernetically equivalent* to the adaptive system defined by its behaviour (in the sense that the model emits all the observed behaviours and no others). However, the derivation of the model is of no concern in the present context, and its structure will be taken as given for the purpose of re-phrasing and illustrating the definitions of modes of adaption.

4.1. ADAPTION AUTOMATA

From Section 3, the adaptive behaviour of a system is completely described, in terms of this analysis, by a sequence of descriptions of the task given and

the satisfactoriness of the controller in performing each task. Consider an automaton defined by the quintuple

$$A=(T, P, S, \sigma, \pi) \tag{12}$$

where T is a set of *tasks* forming inputs to the automaton; $P=(p^+, p^-)$ is a set of two possible outputs from the automaton, p^+ indicating satisfactory performance and p^- unsatisfactory performance; S is a set of states for the automaton, and σ and π are the next state and output functions, respectively. Such a machine will be called an *adaption-automaton* representing the behaviour of some learning system.

In general the adaption automaton need not be determinate and both the next state and output functions may be assumed to be set-valued. Hence, if the automaton is in state, $s \varepsilon S$, is given the task, $t \varepsilon T$, then its next state, $s' \varepsilon S$, and its output, $p \varepsilon P$, belong to the transition, and output, sets, respectively:

$$s' \varepsilon \sigma(s,t) \tag{13}$$

$$p \varepsilon \pi(s,t). \tag{14}$$

Since it is the effect of sequences of tasks, especially those generated by the repetition of a single task, which are of interest, it is important to have a clear notation distinguishing between tasks, sets of tasks, sequences of tasks, and sets of sequences of tasks. A sequence consisting of the task t_1 followed by the task t_2 will be written t_1t_2 (with the obvious extension to longer sequences), a sequence consisting of the task, t, repeated n times will be written t^n. A typical set of tasks will be represented by the letter, $T(T_1, T_2$, etc.); a typical sequence of tasks will be represented by the letter, $u(u_1, u_2$, etc.). The set of sequences generated by the set of tasks, T, as free generators will be written $U(T)$—that is, the set of all possible sequences of tasks which may be formed using members of T. The function, σ, has an obvious extension from tasks to task sequences—if $u=tu'$, then:

$$\sigma(s,u) \equiv \{\sigma(s',u') : s' \varepsilon \sigma(s,t)\}. \tag{15}$$

Having established this structure, it is possible to give set-theoretic definitions of all the various modes of adaption defined in Section 3. These may now be characterized by the sets of states of the automaton in which the behaviour can arise which satisfies the constraints of the previous definitions.

4.2. ADAPTION SETS

Let $$W(T)=\{s : \forall t \varepsilon T, \pi(s,t)=p^+\} \tag{16}$$

that is, $W(T)$ is the set of states in which the controller will have a *satisfactory interaction* given any task from the set, T.

Let $$A(t)=\{s : \forall n, \sigma(s,t^n) \subset W(t)\} \tag{17}$$

that is, $A(T)$ is the set of states in which the controller is *adapted* to the task t, because, given a sequence of task consisting of t repeated, the interaction is always satisfactory.

Let $$P(T)=\{s: \forall t \varepsilon T, \exists N: \sigma(s,t^N) \subset A(t)\} \tag{18}$$

that is, $P(T)$ is the set of states in which the controller is *potentially adaptive* to the set of tasks, T, because, given any sequence of tasks consisting of a member of T repeated, the interaction is eventually always satisfactory.

Let $$C_A(t,T)=\{s: \forall t \varepsilon T, n, \sigma(s,t^n) \subset W(t) \cap P(T)\} \tag{19}$$

that is, $C_A(t,T)$ is the set of states in which the controller is *compatibly adapted* to the task, t, with respect to the set of tasks, T, because, given a sequence of tasks consisting of t repeated, the interaction is always satisfactory and the state remains with the sub-set which is potentially adaptive to the set of tasks, T.

Let $$C(T)=\{s: \forall u \varepsilon U(T), \sigma(s,u) \subset P(T)\} \tag{20}$$

that is, $C(T)$ is the set of states in which the controller is *compatibly adaptive* to the set of tasks, T, because, given any sequence of tasks consisting of members of T, its state remains in the sub-set which is potentially adaptive to the set of tasks, T.

Let $$J_A(T)=\{s: \forall t \varepsilon T, u \varepsilon U(T), \sigma(s,u) \subset A(t)\} \tag{21}$$

that is, $J_A(T)$ is the set of states in which the controller is *jointly adapted* to the set of tasks, T, because, given any sequence of tasks consisting of members of T, its state remains in the sub-set which is adapted to each member of T.

Let $$J(T)=\{s: \exists N, \forall t \varepsilon T, u \varepsilon U(T), \sigma(s,ut^N) \subset J_A(T)\} \tag{22}$$

that is, $J(T)$ is the set of states in which the controller is *jointly adaptive* to the set of tasks, T, because, given any sequence of tasks consisting of members of T, in further adapting to any one of these tasks it becomes jointly adapted to them all.

The definitions of this section include all the modes of adaption previously described in Section 3. The interrelationships between the various modes, briefly discussed in Section 3.9, now appear as inclusion relations between the adaption sets. The most important of these are:

$$W(T_1 \cup T_2)=W(T_1) \cap W(T_2) \tag{23}$$

$$P(T_1 \cup T_2)=P(T_1) \cap P(T_2) \tag{24}$$

$$C(T_1 \cup T_2) \subset C(T_1) \cap C(T_2) \tag{25}$$

$$J(T_1 \cup T_2) \subset J(T_1) \cap J(T_2) \tag{26}$$

$$A(t) \subset W(t) \cap P(t) \tag{27}$$

$$C_A(t,T) = A(t) \cap P(T) \tag{28}$$

$$P(t_1 \cup t_2) \supset A(t_1) \cap A(t_2) \subset J_A(t_1 \cup t_2) \tag{29}$$

$$P(T) \supset C(T) \supset J(T). \tag{30}$$

4.3. TRAJECTORIES IN THE STATE-SPACE OF AN ADAPTION-AUTOMATON

These relations, and the process of adaption itself, appear most clearly in diagrams of the adaption sets as sub-sets in the state-space of the adaption-automaton. In the same diagram, the dynamic behaviour of the adaptive system may be shown as the *trajectory* of states through which it passes in adapting to a sequence of tasks. The set of all states of adaption-automaton is shown as a rectangular region in Fig. 3. From an initial state, such as Y,

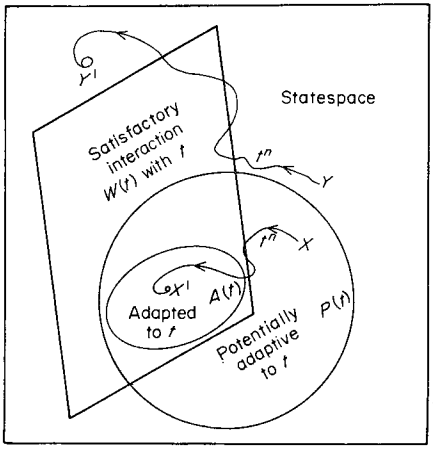

FIG. 3. Adaption to a single task.

repetition of a task, t, generates a trajectory of states within this region—the states are assumed to be assigned to points in the region such that, without loss of generality for a single task, the trajectory appears as a connected path. Within the region are delimited those states, $W(T)$, for which the controller is satisfactory when given the task, t. The states, $A(t)$, for which the controller is adapted to the task, t, from a sub-set of these, since a trajectory starting in the adapted region must always remain satisfactory. The states, $P(t)$, for which the controller is potentially adaptive to the task, t, form a third region,

enclosing the one in which it is adapted to t. Fig. 3 includes all the adaption sets relevant to a controller adapting to a single task.

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 always leave it; started within the potentially adaptive region, at X, it will remain that region, perhaps oscillating between being satisfactory, and being unsatisfactory, for a while, but eventually entering the adapted region, where it is always satisfactory, and never leaving it.

Some of the relations between potential, compatible, and joint adaption for two tasks, t_1 and t_2, are illustrated in Fig. 4. The space of all states has been split into the regions, $A(t_1), A(t_2), P(t_1 \cup t_2), C(t_1 \cup t_2), J(t_1 \cup t_2)$ and $J_A(t_1 \cup t_2)$ and inclusion relations, (29) and (30) are clearly illustrated. Since two tasks are involved, trajectories in the state-space may show a more complex variety of behaviour than for a single task. For example, since the state, X_0, is within the potentially adaptive region for both tasks, $P(t_1 \cup t_2)$, but outside the compatibly adaptive region, $C(t_1 \cup t_2)$, trajectories generated by the task sequences t_1^n or t_2^n, will eventually enter the adapted regions, $A(t_1)$ or $A(t_2)$, respectively, but, in so doing, they must also eventually leave $P(t_1 \cup t_2)$. This loss of potential adaptivity to one task may not take place immediately on first adapting to the other (as in the trajectory from X_0 to X_1), but may be dependent on adaption taking place to both tasks. For example, the trajectory from X_0 to X_2 under t_2^n enters $A(t_2)$ within $P(t_1 \cup t_2)$, and hence is within $C_A(t_1 \cup t_2)$. On taking adavantage of the residual adaptivity to t_1, however, the trajectory proceeds to X_3 which is not within $P(t_1 \cup t_2)$, so that the controller cannot again become adapted to t_2.

A trajectory starting from Y_1 within the compatibly adaptive region, $C(t_1 \cup t_2)$ cannot leave the region of potential adaption to both tasks, $P(t_1 \cup t_2)$, no matter what sequence of tasks of the form, $(t_1 + t_2)^*$, are given. Trajectories originating from within $C(t_1 \cup t_2)$ by t_1^n or t_2^n will drive the state of the adaption-automaton back and forth between the regions, $A(t_1) \cup A(t_2)$, respectively. If such trajectories should remain within the region $A(t_1) \cap A(t_2)$, so that the controller is always adapted to both tasks, then they are within the region of joint adaption, $J_A(t_1 \cup t_2)$, and *may* have originated in the region of joint adaptivity, $J(t_1 \cup t_2)$. They need not necessarily have done so, since, whilst all trajectories generated by sequences of the form $(t_1 + t_2)^*$ from within $J(t_1 \cup t_2)$ must eventually enter $J_A(t_1 \cup t_2)$, there are states outside $J(t_1 \cup t_2)$ from which *some* trajectories of this type will enter $J_A(t_1 \cup t_2)$.

For a given adaption-automaton, certain of the adaption sub-sets will generally be empty, and the definitions do not imply that the corresponding states exist, but merely classify them should they be present. Figures 3 and 4, are drawn as if the state-space had a convenient topology under the action of t_1 or t_2, such that the transitions occur to neighbouring points (in the Euclidean topology of the plane). A plot of trajectories in an arbitrary state-space will not show this convenient form in general, and the individual adaption sets may be partitioned into disjoint sub-sets, subject to the various inclusion relations. Neither do these figures illustrate the possible indeterminacy of the adaption-automaton—trajectories may fork in practice.

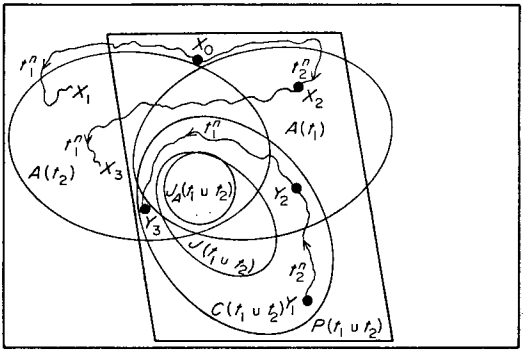

FIG. 4. Adaption to a pair of tasks.

4.4. TRAINING AS A CONTROL PROBLEM

The introduction of the concept of an adaption automaton, and consideration of adaptive behaviour as trajectories in the state space of the automaton, clearly do not in themselves advance the theory of adaptive behaviour. However, they provide an important conceptual foundation for applications of the theory. For example, if the initial state of the adaption automaton lies outside the region of potential adaption for the required task ($P(t)$ in Fig. 3) then no amount of learning experience with that task alone will give rise to ultimate satisfactory performance. However, there may be some sequence of tasks which will give rise to a trajectory from the initial state outside $P(t)$ to one within it. Such a sequence of tasks may be said to *train* the learning system for the performance of the task, t.

The selection of appropriate training sequences dependent upon information acquired about the structure and state of the adaption automaton of the trainee is clearly a *control problem* (Pask, 1965; Gaines, 1969). Indeed the previous definitions of adaption sets could be re-phrased in terms of stability

theory (Bhatia & Szego, 1967), with $A(t)$ as a weak attractor and $P(t)$ as its domain of attraction. The theory of adaptive systems then becomes one of the stability of a certain class of non-linear systems, and the theory of training becomes one of the control of the same class of systems. The theory of general systems is not sufficiently advanced for these considerations to have any immediate effect on studies of learning or training, but the integration of the concepts of adaption and stability, training and control, unifies major areas of current research in artificial intelligence and general systems theory.

5. Conclusions

The objective of this paper has been to provide an operational and purely behavioural definition of terms such as " adapted " and " adaptive ". The initial section was of a philosophical nature, analysing generally accepted concepts of learning and adaption, and previous definitions of these terms, in order to develop explicata for the various aspects of adaptive behaviour. The primitive construct derived was that of a " task ", being a segment of the interaction between controller and environment for which it is possible to say whether or not the controller has performed satisfactorily.

It was noted that terms such as " adaptive " are not absolute and have an inherent arbitrariness, so that a system termed " adaptive " by one observer may not be so evaluated by another. In the present analysis this arbitrariness is concentrated in the segmentation of the interaction between controller and environment into " tasks ". Once this segmentation is agreed, the different modes of adaption may be rigorously and unambiguously defined.

This approach differs from that taken in prevous attempts to define the term " adaptive " behaviourally, in that, not only is a performance criterion used as an explicatum of adaption giving rise to " good " or goal-seeking behaviour, but also the interaction between controller and environment is segmented in order to take an account of adaptive dynamics and to provide an explicatum of adapting being a *process* of becoming satisfactory.

This segmentation has the further advantage, very important in applications such as the study of training techniques, that it enables the effects of learning one task on the capability to acquire others to be taken into account. A variety of possible modes of adaptive behaviour relative to a set of tasks has been defined, each of which forms an explicatum of some common stereotype of adaptive behaviour, and the differentiation between which provides some gross account of the controller's overall adaptive capabilities.

If one assumes that the adaptive behaviour may be ascribed to an *adaption automaton*, whose inputs are tasks and whose outputs are the satisfactoriness

of performance of the tasks, then the definitions of modes of adaption may be used to classify the states of the automaton, and the adaptive behaviour may be viewed as a trajectory in the state-space of the automaton. This leads to an interesting unification of the fields of control theory and learning systems theory, in that adaption can be related to the *stability* of trajectories in the state-space of the adaption automaton and training becomes a *control* problem in the same state-space.

References

ASHBY, W. R. (1967). The set theory of mechanism and homestasis. In *Automaton Theory and Learning Systems*. London, New York: Academic Press.

BHATIA, N. P. & SZEGO, G. P. (1967). *Dynamical Systems: Stability Theory and Applications*. Berlin, New York: Springer-Verlag.

BUSH, R. R. & MOSTELLER, F. (1955). *Stochastic Models for Learning*. New York: Wiley.

CLARK, R. M. (1963). Panel discussion. In *Adaptive Control Systems*. Oxford: Pergamon Press.

DONALSON, D. D. & KISHI, F. H. (1965). Review of adaptive control system theories and techniques. In *Modern Control Systems Theory*. New York: McGraw Hill.

FLORENTIN, J. J. (1962). Optimal, probing, adaptive control of a simple Bayesian system. *J. Electron. Control*, **13**, 165.

FREEMAN, H. (1963). Panel discussion. In *Adaptive Control Systems*. Oxford: Pergamon Press.

GAINES, B. R. (1967). Teaching machines for perceptual-motor skills. In *Aspects of Educational Technology*. London: Methuen.

GAINES, B. R. (1969). Adaptively controlled instruction for a tracking skill. In *Programmed Learning Research*. Paris: Dunod.

GUTHRIE, E. R. (1952). *The Psychology of Learning*. New York: Harper.

HUNTER, W. S. (1934). Learning—IV. Experimental studies and learning. In *A Handbook of Experimental Psychology*. Worcester, Mass.: Clark University Press.

MCGEOCH, J. A. & IRION, A. L. (1952). *The Psychology of Human Learning*. McKay.

MARTENS, H. H. (1959). Two notes on machine " learning ". *Inf. Contr.*, **2**, 364–379

MOE, M. L. & MURPHY, G. J. (1962). An approach to self-adaptive control based. on the use of time moments and a model reference. *IRE Trans.*, AC-7, 82.

NEWELL, A., SHAW, J. C. & SIMON, H. A. (1959). Report on the general problem solving problem. *Proc. Int. Conf. Inf. Process*, Paris.

PASK, G. (1965). Teaching as a control process. *Control.*

RADANOVIC, L. (Ed.) (1966). *Sensitivity Methods in Control Theory*. Oxford: Pergamon Press.

RAIBLE, R. H. (1963). Comments on " An approach to self-adaptive control based on the use of time moments and a model reference ", *IEEE Trans.* AC-**8**, 270.

SOMMERHOF, G. (1950). *Analytical Biology*. London, Oxford: Oxford University Press.

STANIER, R. Y. (1953). Adaption, evolutionary and physiological. In *Adaption in Micro-Organisms*, Third Symp. Soc. Gen. Microbiol. London: Cambridge University Press.

THORPE, W. H. (1956). *Learning and Instinct in Animals*. London: Methuen.

TRUXAL, J. G. (1963). Adaptive control. *Survey Paper 2nd Int. Congr. IFAC*. Basle.

WEIZENBAUM, J. (1966). ELIZA—A computer program for the study of natural language communication between man and machine. *Comm. ACM*, **9**, 36.

ZADEH, L. A. (1963). On the definition of adaptivity. *Proc. IEEE*, 469.