

by

Bernard Widrow

1933-61

Abstract

Adaptive control systems are capable of giving near optimum performance in the face of changing input command and noise characteristics, changing dynamics of the controlled processes, and changing goal or mission requirements. The principles of feedback control are used by the adaptation mechanism to control the structure of systems. Performance feedback makes automatic self-optimization possible.

A quantitative dimensionless measure of the quality of adaptation is given by the system "misadjustment", defined as the ratio of the mean increase in mean square error (resulting from making adjustments based on finite amounts of statistical data) divided by the minimum mean square error. Optimal use of performance measurements is made by "data-repeating". A feedback control system employing this scheme can adapt to a major change in input process statistics and/or a major change in the dynamic characteristics of the controlled member in about ten times the duration of the impulse response of the feedback system itself, and will exhibit a misadjustment of only 20 percent.

Faster adaptation with less misadjustment is possible with pattern-recognizing adaptive control systems. Pattern-recognizing systems use longer term experiences by associating current control problems with similar ones that have been seen previously and for which control strategies have already been worked out. In the case of vehicle control systems, these longer term experiences which are needed for the "training" of the pattern-recognizing filters could be obtained during a single flight, during many flights, or from telemetered data of previous flights of other similar vehicles.

Pattern-recognizing filters composed of adaptive ADALINE "neurons" have been devised for the classification of binary geometric patterns. The same principles are applicable to the design of filters that can be trained for classifying sequential analog signals.

A new electric circuit element called the "memistor" (a resistor with memory) has been devised to facilitate the realization of the ADALINE neuron and its adaptation procedure. It is a compact rugged electrochemical element whose resistance can be reversibly controlled by electroplating. The experiences of the neuron are stored in resistance values, in a simple and directly usable form.

Many research problems remain in the theory of pattern-recognizing adaptive control systems. Such systems promise to make possible certain unusual and very useful performance capabilities.

*Dept. of Electrical Engineering, Stanford University,
Stanford, California

A. Introduction

Adaptive control systems are capable of giving near optimum performance in the face of changing input command and noise characteristics (non-stationary input statistics), changing dynamics of the controlled process, and changing goal or mission requirements. These systems essentially consist of adjustable controllers, adaptation mechanisms for setting their adjustments, and the controlled processes themselves. Adaptive or self-optimizing systems automatically modify their own structures in order to achieve and maintain optimal performance. An adaptive system that continually searches for an optimum structure within an allowed class of possibilities by an orderly trial-and-error process would give performance superior to that of an optimum fixed system in many control situations.

Several ways of classifying adaptation schemes have been proposed in the literature. This author finds it convenient to think merely in terms of open-loop and closed-loop adaptation. The open-loop adaptation process involves making measurements of input or environmental characteristics, applying this information to a formula or a computational algorithm, and using the results to set the adjustments of the adaptive system. Closed-loop adaptation, on the other hand, involves automatic experimentation with these adjustments to optimize a measured system performance. An example of statistical prediction well illustrates the differences between these methods.

A Wiener(1) predictor is a linear filter designed to predict its input signal with minimum mean square error. The form of

this filter depends only upon the autocorrelation function of this input signal. An adaptive Wiener filter could be arranged as shown in Fig. 1(a). The correlator measures the input autocorrelation function. Continually updated measurements are applied to a computation process (which could be either digital or analog) which implements the requirements of the Wiener-Hopf equation and produces the form of the "best" filter. This structural information is then applied to an adjustable filter (which also could be analog or digital). The arrangement of Fig. 1(a) offers many of the advantages of adaptation. The basic system could be designed without the designer knowing the exact nature of the input statistics, and in addition, the system will adapt to changing input statistics and could thus provide highly effective operation in the non-stationary case.

A closed-loop adaptation process is shown in Fig. 1(b). The output signal of the adjustable predictor is delayed by an amount of time equal to the desired prediction time, and then compared with the system input signal. Perfect agreement implies perfect prediction, and indeed, the difference is a prediction error. The mean square of this error can automatically be minimized by an adaptation procedure, by a trial-and-error process. This form of adaptation utilizes "performance feedback" to achieve direct, automatic system synthesis, i.e., the selection of an "optimum" system from a predetermined class of possibilities.

The principles of closed-loop adaptation are further illustrated in Fig. 2. The input signal is applied to the "worker", an adjustable system that deals with real signal inputs to produce the system output signal. The human operator

adjusts the knobs to optimize the reading of the performance meter. This is done even though the knobs have interacting effects upon performance, even when the operator (or "boss") has no knowledge of what is inside the worker or what functions the knobs serve. The boss performs a purely mechanical service which could be automated. The combination of worker and boss are represented in the block diagram of Fig. 2(b). Feedback control is utilized here to determine and control the structure of systems. Closed-loop adaptation has the advantage of being usable where no analytic synthesis procedure is known or exists; for example, where error criteria other than mean-square are used and where systems are quasistatically nonlinear. In the event of a partial system failure, an adaptation system that continually monitors performance will optimize this performance by adjusting the intact parts. System reliability could be greatly enhanced by adaptation. (2)

B. The Performance Surface

It is the purpose of this section and the several succeeding sections to describe in greater detail how the worker-boss closed-loop adaptation process could be implemented, how the performance of systems embodying such principles is evaluated and predicted, and how performance is affected by the speed of adaptation.

According to the scheme of Fig. 2, the worker is an adjustable system which, in principle, could vary in complexity from a simple open-loop filter to an entire closed-loop feedback system. In the worker itself, the adjustments might directly control loop gain, individual time constants of transient components, pole-zero positions in the S-plane, or other characteristics that govern performance. No theory exists yet that

enables a designer in a general way to "optimize" his choice of the adjustment variables.

The types of workers that have been examined by this author are sampled-data filters, formed of tapped delay lines. Adjustment is made directly to their impulse responses by making the tap gains or weights be the adjustment variables. This type of filter can be readily implemented digitally.

Return now to the problem of making an adaptive predictor. The linear sampled-data system shown in Fig. 3 is intended to predict the next sample of the input sequence of samples. The present output sample $g(n)$ is a linear combination of present and past input samples. The constants in this combination are h_0, h_1, h_2, \dots , the predictor impulse-response samples, or the gains associated with the delay-line taps. Their choice constitutes the adjustable part of the predictor design. They may be adjusted in the following manner. Apply a mean square reading meter to $\epsilon(n)$, the difference between the present input and the delayed prediction. This meter will measure the mean square error in prediction. Adjust h_0, h_1, h_2, \dots , until the meter reading is minimized.

Suppose that the predictor has only two impulses in its impulse response, h_0 and h_1 . The mean square error for any setting of h_0 and h_1 can be readily derived.

$$\begin{aligned} \epsilon(n) &= f(n) - h_0 f(n-1) - h_1 f(n-2) \\ \overline{\epsilon^2} &= \phi_{ff}(0) h_0^2 + \phi_{ff}(0) h_1^2 - 2 \phi_{ff}(1) h_0 h_1 - 2 \phi_{ff}(1) h_1 \\ &\quad + 2 \phi_{ff}(1) h_0 h_1 + \phi_{ff}(0) \end{aligned} \quad (1)$$

The discrete autocorrelation function of the input is $\phi_{ff}(k)$. The mean square error is a parabolic function of the adjustments.

The optimum m-impulse predictor can be derived analytically by setting the partial derivatives of $\overline{\epsilon^2}$ with respect to the adjustment variables to zero. What results is the discrete analog of Wiener's optimization of continuous filters. Finding the optimum system experimentally (as in Fig. 2) is the same as finding the minimum of a paraboloid in m dimensions. Performance feedback is seen to be equivalent to the trial-and-error searching of a stochastic performance surface (which in a wide variety of cases is parabolic) for a minimum. In the case of a non-stationary statistical input, this surface is constantly changing in shape, orientation, and position. Performance feedback is seen to be a multidimensional "bottom-tracking" servo.

C. Surface-Searching; Performance Feedback

Iterative or trial-and-error surface-searching processes are integral parts of closed-loop adaptation systems. It is often convenient to represent such processes as feedback systems; the error of trial and error is analogous to the "error" of feedback control. Many of the relaxation and iterative methods commonly employed by numerical analysts appear to be linear feedback systems when represented in this manner.

Stationary points, or maxima and minima, are characterized by zero partial derivatives with respect to the independent variables. These partial derivatives generally increase with distance from the stationary point and, moreover, increase linearly for a quadratic surface. In using many of the gradient methods, (2) the surface is explored by making changes in the

Stuckow

independent variables (starting with an initial guess) in proportion to measured partial derivatives to obtain the next guess, and so forth. These methods give rise to geometric (exponential) decays in the independent variables as they approach a stationary point for second-degree or quadratic surfaces. This is illustrated by the one-dimensional Model of Fig. 4.

The "surface" being explored in Fig. 4 is given by Eq. 2. The first and second derivatives are given by Eqs. 3 and 4.

$$y = a(x-b)^2 + c \quad (2)$$

$$\frac{dy}{dx} = 2a(x-b) \quad (3)$$

$$\frac{d^2y}{dx^2} = 2a \quad (4)$$

Let the proportionality constant between change in guess and derivative be $-k$. This constant could be so chosen that the error in x decreases by one half with each iteration cycle, as illustrated in Fig. 4a. A sampled-data feedback model of the iterative process is shown in Fig. 4b. (4,5) The initial numerical guess is injected once at the beginning of the process, whereas the numerical reference or stationary value b is injected synchronously during each cycle. The numerical sequence at the point $x(n)$ begins with the initial guess and proceeds as a sampled transient that relaxes geometrically toward the stationary point, exactly like the sequence of guesses in the surface exploration. For the present, disregard the source of "derivative measurement noise."

The following is an explanation of the feedback model. The

next guess is equal to the present guess (this accounts for the unity feedback branch) plus a constant $(-k)$ times the derivative. From Eq. 3, the derivative equals $2a$ times $x(n)$ minus the constant $2ab$. Since the next guess will become the present guess for the next iteration cycle, it is stored by the unit delay (the feedforward branch of transfer function z) to supply the signal at node $x(n)$ at the proper time. It is clear from the flow-graph model that, if the iterative process is stable, equilibrium will be reached when $x(n)$ reaches the value b .

The flow-graph can be reduced, and the transfer function from any point to any other point can thus be found. The resulting characteristic equation is

$$(2ak - 1)z + 1 = 0 \quad (5)$$

The iterative process is stable when $0 < k < \frac{1}{2a}$. In order to choose the "loop gain" k to get a specific transient decay rate, one would have to measure the second derivative (2a) at some point on the curve.

Each time a guess in x is to be made, the derivative is physically measured (fig. 4a), whereas in the model (Fig. 4b) it is obtained as a quantity proportional to x . If the surface were of higher degree than second, the derivative would not be simply proportional to x but would be some polynomial in x . The model could still be made, but it would not be of a linear system, and transients would not be geometric. Nevertheless, the iteration process will locate the stationary point. In its vicinity, transients will be geometric because the second and lower degree terms of the Taylor expansion of any continuous surface become the dominating ones. For this and other reasons, exploration

AW/Draw

of the parabolic surface is given special attention.

The derivatives of a parabola and the partial derivatives of a parabolic surface could be measured in the manner illustrated in Fig. 5. The dimensionless ratio of γ to B is defined as the perturbation P of the measurement.

The first and second derivatives are given by Eqs. 6 and 7. These relations are precise for parabolas, and are approximate for higher degree curves.

$$\frac{dy}{dx} \Big|_{x_B} = \frac{1}{2\delta} (C-A) \quad (6)$$

$$\frac{d^2y}{dx^2} \Big|_{x_B} = \frac{1}{\delta^2} (C-2B+A) \quad (7)$$

Figure 6 shows a two-dimensional paraboloid and a plan view of a possible sequence of vector changes in the independent variables x_1 and x_2 while a minimum is being sought. Each component of a vector change is a linear combination of the local partial derivatives. The resulting transients are multidimensional geometric progressions.

The surface being searched is given by Eq. 8, the partial derivatives by Eqs. 9, and the second partial derivatives by Eqs. 10.

$$y = ax_1^2 + bx_2^2 + cx_1 + dx_2 + ex_1x_2 + f \quad (8)$$

$$\frac{\partial y}{\partial x_1} = 2ax_1 + c + ex_2 \quad \frac{\partial y}{\partial x_2} = 2bx_2 + d + ex_1 \quad (9)$$

$$\frac{\partial^2 y}{\partial x_1^2} = 2a \quad \frac{\partial^2 y}{\partial x_1 \partial x_2} = e \quad \frac{\partial^2 y}{\partial x_2^2} = 2b \quad (10)$$

A vector flow-graph model of the iterative process is given in Fig. 7(a). The branches in this graph are capable of carrying two-dimensional samples, indicated by column matrices, and the matrix gains of the branches signify that outputs equal premultiplied by gains. The two-dimensional flow graph is completely analogous to the one-dimensional graph. The feed-forward branch is merely a delay with no crosscoupling of the coordinates, and the unit feedback branch is simply a unit instantaneous transmission with no crosscoupling. The first partial derivatives are formed, as indicated by Eqs. 9, from the linear combination of the constants c and d and of the x 's premultiplied by the matrix of second partials.

This flow graph can be reduced straightforwardly by making use of the rules of matrix algebra. There are as many natural frequencies (decay rates) as there are independent coordinates. The multidimensional loop gain in this case is determined by choice of the matrix of k 's.

Among the more useful surface-searching methods are the method of steepest descent, Newton's method, and the Southwell Relaxation method. These methods can be represented by feedback models, like those of Fig. 7. They differ mainly in the choice of the k 's in the feedback matrix.

The method of steepest descent requires that vector changes in adjustment be made in the directions of the successive local gradient vectors, and that the magnitudes of these changes be proportional to gradient magnitudes. Here, the matrix of k 's is a diagonal one with equal elements on the main diagonal. Adjustment transients crosscouple from one coordinate to the other when this

Sweden

method is practiced.

Newton's method is a one-step process, wherein the matrix of k 's is the inverse of the matrix of second partials. Multidimensional transients die out completely in one step. A modified Newton's method has the same matrix of k 's, only scaled by a factor less than unity. Transients die out geometrically, not in one step, and are of a single time constant. Successive adjustments proceed along a straight line in multidimensional space from the initial guess to the stationary point. Cross-coupling of transients among the coordinates is not present with this method. Newton's method is mainly of interest for analytical purposes. Physical implementation in complicated situations has the disadvantage of requiring machinery for matrix inversion.

The Southwell method requires changes made to minimize y with each successive adjustment. In the two-dimensional feedback model, $k_{11} = \frac{1}{2a}$ and $k_{22} = \frac{1}{2b}$. Switches that close alternately must be included in the model (see Fig. 7(b)) to represent the one-component-of-change-at-a-time characteristic of this method. Although transients are of a single time constant, their effects crosscouple among the adjustment coordinates.

A basic limitation on the speed of adaptation of adaptive systems arises from statistical sample-size requirements in the measurement of environmental characteristics. Measurements of the mean square error surface are in general noisy, and this noise propagates by way of the iterative surface-searching process into the adjustment variables and causes loss in system performance. Study of the performance feedback processes leads

directly to a relation between system performance and its rate of adaptation.

D. Performance Measure - The Misadjustment

Noise enters the adaptation feedback system of Fig. 4 because the input process cannot be continued indefinitely for each measurement of mean square error (A,B,C, of Figure 5, needed for gradient measurement). This noise has the following effect upon adaptation: The slower the adaptation, the more precise it is; the faster the adaptation, the more noisy (and poor) are the adjustments.

Consider that the adaptive model has only a single adjustment. A plot of mean square error versus h_0 for this simplest system would be a parabola, analogous to the parabola of Figure 1. During each cycle of adjustment, the derivative of $y = \epsilon^2$ with respect to $x = h_0$ would have to be measured according to scheme of Figure 5.

Noise in the system adjustment causes loss in steady-state performance. It is useful to define a dimensionless parameter M , the 'misadjustment', as the ratio of the mean increase in mean square error to the minimum mean square error. It is a measure of how the system performs, on the average, after adapting transients have died out, compared with the fixed optimum system. With regard to the curve of Figure 4,

$$M = \frac{\bar{y} - c}{c} \quad (11)$$

From consideration of equation 2, it can be seen that $(\bar{y} - c)$, the average increase in y , is equal to the variance in x multiplied by a . The variance is due to derivative measurement noise which propagates by way of the iterative surface-searching process.

The noise propagation path is shown in the flow graph of Figure 4(b). Assuming that derivative measurement noises are statistically independent from one iteration cycle to the next, the variance in x equals the variance in derivative noise multiplied by $(1/8a^2\tau)$, a conservative approximation to the sum of squares of the impulses of the impulse response from the noise injection point to the adjustment x . The time constant τ is defined so that, if $\tau = 1$, adaptation transients decay by a factor $(1/\epsilon)$ with each iteration cycle.

A detailed derivation of the variance in derivative measurements is given in references 6 and 7. The result is that

$$(\text{variance in derivative measurement}) = \frac{\sigma_c}{N\tau} \quad (12)$$

The number of forward (See Fig. 5) or backward measurements per cycle is N . The perturbation is P . (See Fig. 5). Relation 12 is based upon several assumptions: that the adjustment x is in the vicinity of the minimum, that the prediction error signal is gaussian distributed (relation 12 is quite insensitive to the shape of this distribution density, however), and that the error samples are uncorrelated. The misadjustment can be deduced as $M = \frac{1}{8N\tau P}$.

If the nature of the physical process permits 'data repeating', i.e. if it is possible to apply the same strip of input data to the system for both forward and backward measurements (making sure that initial conditions are the same), the variance of the derivative measurement noise turns out not to depend upon the amplitude of the perturbation. Making the same assumptions as were made previously, the

expression for the variance with data repeating can be shown to be

$$(\text{variance in derivative measurement}) = \frac{4\sigma_c}{N} \quad (14)$$

In this case, N is the total number of error samples per cycle. Accordingly, the misadjustment is

$$M = \frac{1}{2(N\tau)} \quad (15)$$

The $(N\tau)$ product is related to the total number of samples 'seen' by the system in adapting to a step transient in input process statistics. Notice that a given effect could be achieved by using many samples per cycle (large N) and few cycles with large steps to adapt (small τ), or by using few samples per cycle (small N), and proceeding toward the optimum with small steps (large τ).

Let the number of samples that elapse in one time constant of adaptation be called the 'adaptation time constant' \bar{T} . Where data repeating is not practised, $\bar{T} = 2N\tau$. Where data is repeated, $\bar{T} = N\tau$. Expressions 13 and 15 become 16 and 17 respectively.

$$M = \frac{1}{4\bar{T}P} \quad (16)$$

$$M = \frac{1}{2\bar{T}} \quad (17)$$

Multidimensional adaptation processes can be analyzed by generalizing these methods. It can be shown that the misadjustment increases with m^2 when Newton's method is used, and that the misadjustment increases with m when data is repeated for any single time-constant method.

In the cases of one-step adaptation ($\tau=0$), or adaptation

Twidale

to-completion on a fixed body of N repeated input process samples, it can be shown that

$$M = \frac{m}{N} \quad (18)$$

E. Systems Applications

These principles may be applied in a variety of situations, two of which are illustrated in Figure 8. Performance feedback is used in the system of Figure 8(a) to achieve imitation of an unknown complex system. The adaptive system 'learns' of the characteristics of the unknown system by imitating its behavior as best it can. The mean square error is a parabolic function of the adjustments, if the input is stationary and the output of the unknown dynamic system is stationary. A combination of imitation and prediction enables an adaptive system to predict the output of an unknown dynamic system by making use of both its input and output signals. A conventional predictor would use only the output signal. In Figure 8(b), a scheme is shown which combines a low-noise, low-capacity link (for performance feedback) with a high-capacity, noisy communication link. An adaptive filter is used to separate noise and signal. The mean square error is again a parabolic function of the adjustments, and the rate of adaptation is limited by the capacity of the low-capacity link.

A fundamental form of adaptive feedback control system is shown in Fig. 9(a). Adaptation for the minimization of the mean square of the control-system error signal allows self-optimization in the event of change in input process statistics and in the event of unexpected change in "plant" characteristics. The mean-square error surface would not be

precisely parabolic in this case, and there is no guarantee that the surface would have only a single minimum. In the vicinity of a minimum, the surface can be represented by first and second degree terms in its multidimensional Taylor expansion and, as such, the misadjustment formulas which were derived for parabolic surfaces should at last describe "lock-on" behavior. In Fig. 9(b), a different error criterion is used. The feedback system is adapted to imitate a model. Such systems are called "model-reference" adaptive systems. (8)

The misadjustment gives a measure of the effectiveness of adaptation. It neither gives information on the magnitude of the minimum mean square error, nor does it give information on the effectiveness of the choice of adjustment variables (are there enough of these and are they the best to use). Simulation experiments have shown that the measured misadjustments rarely differ from their predicted values by as much as 20 or 30 per cent.

The misadjustment formulas are quite accurate when applied to the situations for which they have been derived. These formulas serve as 'rules of thumb' when other performance criteria then minimization of mean square error are used, and when the worker is non-linear. (9,10)

F. Speed Vs. Quality of Adaptation

Let the predictor of Fig. 3 have a total of five variable taps on its tapped delay line. Suppose that data-repeating is possible, and that a misadjustment of 10% is acceptable. The question is, how rapidly could this system adapt? The adaptation time constant can be derived from formula (17), generalized for m dimensions.

(19)

$$M = \frac{m}{2\sqrt{P}} = \frac{5}{2\sqrt{0.1}} = 0.1$$

$$\therefore \sqrt{P} = 25 \text{ samples.}$$

This system could adapt quite precisely to a major change in input process statistics after "seeing" several time constants worth of data (75 to 100 samples). In this case, the time constant of adaptation is five times the response time of the system itself.

If the predictor were operating "on line" in a real time process, data-repeating would not be possible. In addition, a cost resulting from perturbing the system adjustment when measuring gradient components would accrue to the misadjustment exactly equal to the perturbation. Making use of equation 16, the misadjustment is therefore

$$M = \frac{m^2}{4\sqrt{P}} + P \quad (20)$$

For any given M, the optimum choice of P (that minimizes \sqrt{P}) requires that

$$\frac{m^2}{4\sqrt{P}} = P$$

$$M = \frac{m}{\sqrt{P}} \quad (21)$$

With five adjustments and an allowable misadjustment of 10%,

$$.1 = \frac{5}{\sqrt{P}}$$

$$\therefore \sqrt{P} = 2500 \text{ samples} \quad (22)$$

The data repeating scheme is always the more efficient. It is 100 times as efficient for this example.

The benefits of data repeating can often be attained in on-line systems by using auxiliary off-line systems in conjunction with them. Gradient measurements are made by data repeating in the auxiliary systems; thus "dithering" of the adjustments of the on-line system is not required. An example of the application of these principles to the feedback control system of Fig. 9(a) is illustrated in Fig. 10. Increased speed of adaptation is possible with the system of Fig. 10, at the expense of more equipment. One problem is to give the auxiliary plant the flexibility to match the actual plant. Knowledge of the form of the actual plant can be put to good use here.

The system of Fig. 10 operates in the following fashion. The weights h_1, h_2, \dots in the auxiliary system controller are set equal to the like-numbered weights in the actual controller. The weights k_1, k_2, \dots in the plant model in the auxiliary system are then adjusted to minimize the mean square error in the difference between the actual system output and that of the auxiliary system. The weights k_1, k_2, \dots are then fixed for some time. The weights h_1, h_2, \dots in the auxiliary controller are varied to measure the performance surface gradient. The decision is then made by adaptation process No. 1 to change the adjustments in the actual controller according to the gradient measurements in the auxiliary system. Since both plant characteristics and input environmental characteristics would be continually varying in many control situations, both adaptation processes would be continually operating, first one, then the other, using and reusing the same data.

G. Pattern-Recognizing Control Systems

In the adaptive systems described above whose adjustable

Quidraw

portions are completely formed on the basis of recent past history, optimal use of performance measurements was seen to be made with data-repeating. The formulas developed above indicate that a feedback control system employing data-repeating could adapt to a major change in input process statistics and/or a major change in the dynamic characteristics of the controlled member in about ten times the duration of the impulse response of the system itself, and would exhibit a misadjustment of only 20%. To make possible even faster adaptation with less misadjustment, the use of pattern-recognizing adaptive control systems is hereby proposed. The schemes proposed are highly speculative. They have not yet been checked, either theoretically or experimentally.

Pattern-recognizing adaptive control systems would make use of other measurements for adaptation in addition to the direct performance measurements, such as environmental, input, noise, and where available at various points in the system, signal measurements. All of these measurements would be fed to pattern-recognizing filters (see Fig. 11). On-line system adjustments would be set by combinations of the outputs of the conventional adaptation systems and the outputs of the pattern-recognizing filters. The pattern-recognizing adaptive systems would have the ability to use longer-term experiences by associating current control problems with similar ones that have been seen previously and for which control strategies have already been worked out and learned.

Complete reliance upon the outputs of the pattern-recognizing

filter in the system of Fig. 11 for the setting of the knobs of the on-line controller would produce open-loop adaptation. Instead of this, it is proposed that adjustment control be divided between the dictates of the pattern-recognizing filter and the dictates of a slower conventional closed-loop adaptation process. The ultimate objective of adaptation in the example of Fig. 11 is the minimization of the control-system mean square error.

For pattern-recognizing adaptive systems to be feasible, it is essential that the pattern-recognizing filters themselves be adaptive. To design such filters manually would be quite difficult. Self-design would alleviate this problem. Also, it would permit the filter to be tailor-made to the individual control system, and to remain as such in spite of possible spontaneous internal system changes. In vehicle control systems, longer term experiences needed for the "training" of the pattern-recognizing filters could be obtained during a single flight, during many flights, or from telemetered data of previous flights of other similar vehicles.

Considerable progress has been made recently in the development of adaptive pattern-recognizing machines. One such machine which uses an artificial "neuron" called ADALINE (for adaptive linear) ⁽¹¹⁾ has been developed by this author and M.E. Hoff and is described in the next section. Although this machine (shown in Fig. 12) has been designed to classify binary geometric patterns, where input signals are quantized to be either +1 or -1, a similar machine has been used to classify ternary signals (+1, 0, -1), and in principle, could be used on true analog signals as well. ^{LOWDOWN}

H. The ADALINE Pattern-Classifying Machine

In Fig. 13, a block schematic of the ADALINE neuron is shown. This is actually a combinatorial logical circuit, and is a typical element in the adaptive pattern classifying circuits to be considered. This element bears some resemblance to a biological neuron, whence the name.

The binary input signals on the individual input lines have values of +1 or -1, rather than the usual binary values of 1 or 0. Within the neuron, a linear combination of the input signals is formed. The weights are the gains a_1, a_2, \dots , which could have both positive and negative values. The output signal is +1 if this weighted sum is greater than a certain threshold, and -1 otherwise. The threshold level is determined by the setting of a_0 , whose input is permanently connected to a +1 source. Varying a_0 varies a constant added to the linear combination of input signals.

For fixed gain settings, each of 2^5 possible input combinations would cause either a +1 or -1 output. Thus, all possible inputs are classified into two categories. The input-output relationship is determined by choice of the gains a_0, \dots, a_5 . In the adaptive neuron, these gains are set during an iterative learning process.

The adaptive pattern classification machine shown in Fig. 12 has been constructed for the purpose of illustrating the principles of adaptive behavior and artificial learning. During a training phase, crude geometric patterns are fed to the machine by setting the toggle switches in the 4x4 input switch array. Setting another toggle switch (the reference switch) tells the machine whether the desired output for the

particular input pattern is +1 or -1. The system learns something from each pattern and accordingly experiences a design change. The total experience gained in the learning process is stored in the values of the weights a_0, \dots, a_{16} . The machine can be trained on undistorted noise-free patterns by repeating them over and over until the iterative search process converges, or it can be trained on a sequence of noisy patterns on a one-pass basis such that the iterative process converges statistically. Combinations of these methods can be accommodated simultaneously. After training, the machine can be used to classify the original patterns and noisy or distorted versions of these patterns. ADALINE can be used to classify patterns into several categories by using a multi-level output quantizer, and by following exactly the same adaptive procedure.

The following is a description of a simple iterative searching routine. Many variations on this scheme are workable, and some give greater ultimate flexibility. A pattern is fed to the machine, and the reference switch is set to correspond to the desired output. The error ϵ (see Fig. 13) is then read (by switching the reference switch; the error voltage appears on the meter, rather than the neuron output voltage). All gains including the level are to be changed by the same absolute magnitude, such that the error is brought to zero. This is accomplished by changing each gain (which could be positive or negative) in the direction which will diminish the error magnitude by $1/17$. The 17 gains may be changed in any sequence, and after all changes are made, the error for the present input pattern is zero. Switching the reference back, the meter reads exactly the desired output. The next pattern, and its desired

output, is presented and the error is read. The same adjustment routine is followed and the error is brought to zero. If the first pattern were reapplied at this point, the error would be small but not necessarily zero. More patterns are inserted in like manner. Convergence is indicated by small errors (before adaption), with small fluctuations about a stable root mean square error value. This adaptation procedure may be readily modified to get slower (and smoother) adaptation by correcting only a fraction of the error with the insertion of each pattern.

The error signal measured and used in adaption of the neuron is the difference between the desired output and the sum before quantization. This error is indicated by ϵ in Fig. 13. The actual neuron error, indicated by ϵ_n in Fig. 13, is the difference between the neuron output and the desired output.

The objective of adaptation is the following. Given a collection of input patterns and the associated desired outputs, find the best set of weights a_0, a_1, \dots, a_m to minimize the mean square of the neuron error ϵ_n^2 . Individual neuron errors could only have the values of $+2, 0$, and -2 with a two-level quantizer. Minimization of ϵ_n^2 is therefore equivalent to minimizing the average number of neuron errors.

The simple adaption procedure described above minimizes ϵ_n^2 rather than ϵ_n . The measured error ϵ will be assumed to be gaussian-distributed with zero mean. Using certain geometric arguments, it can be shown that, under a wide variety of conditions, ϵ_n^2 is approximately half of ϵ^2 is approximately half of ϵ^2 , and that minimization of ϵ_n^2 is equivalent to minimization of ϵ^2 and therefore to minimization of the

probability of neuron error. The ratio of these mean squares has been calculated and is plotted in Fig. 13 as a function of the neuron error probability. This plot is a good approximation even when the error probability density differs considerably from the gaussian.

For any collection of input patterns and the associated desired outputs, the measured mean square error ϵ^2 can be shown to be a precisely parabolic function of the gain settings, a_0, \dots, a_n . Therefore, adjusting the a 's to minimize ϵ^2 is equivalent to searching a parabolic stochastic surface (having as many dimensions as there are a 's) for a minimum. How well this surface can be searched will be limited by sample size, i.e., by the number of patterns seen in the searching process.

A statistical analysis of the surface searching process, an analysis almost identical to that for the adaptive sampled-data system, has been presented in reference 11. This theory shows that the adaptation procedure implements the method of steepest descent. An extremely small sample size per iteration cycle is taken, namely one pattern. One-pattern-at-a-time adaptation has the advantages that derivatives are measured easily (with data-repeating accuracy), and that no storage is required within the adaptive machinery except for the weight values (which contain the past experiences of the neuron in a compact and directly usable form). The close association between mean square error and error probability allows misadjustment to be interpreted as a percentage of extra error probability resulting from training on a limited sample size or a limited number of patterns. Using formula (18), the theoretical misadjustment for a single neuron is (the

12 WIDROW

addition of 1 in the numerator accounts for the variable threshold level),

$$M = \frac{m+1}{N} \quad (23)$$

The number of input lines to the neuron is m and the number of patterns used in training is N .

Formula (23) has been verified experimentally by statistical generalization tests, similar to tests that psychologists use to measure animal learning. Neurons were trained with small numbers of patterns selected at random from a large collection. The percentage errors in classifying the large collection as a function of the number of training patterns used was able to be predicted. To get a reasonable misadjustment, say 20 or 30 percent, formula (23) leads to a simple "rule of thumb": the number of patterns required to train an adaptive classifier is equal to several times the number of input lines.

In the pattern-recognizing adaptive control system, patterns would be spacial and temporal, involving all the variables monitored by the recognition filter over significant episodes in time. More study will be required to determine how rapidly such filters could adapt.

I. Realization of Adaptive Circuits with Chemical "Memistors"

The structure of the Adaline neuron and its adaptation procedure is sufficiently simple, that an electronic fully-automatic neuron is being developed. To have such an adaptive neuron, it is necessary to be able to store the gain values, analog quantities which could be positive or negative, in such a manner that these values could be changed electronically.

A new circuit element called the memistor (a resistor with memory)⁽¹²⁾ has been devised by this author and M.E. Hoff for the realization of automatically-adapted Adaline neurons. A memistor provides a single variable gain factor. Each neuron therefore employs a number of memistors equal to the number of variable weights.

The memistor consists of a conductive substrate with insulated connection leads, and a metallic anode, all in an electrolytic plating bath. The conductance of the element is reversibly controlled by electroplating. Like the transistor, the memistor is a 3-terminal element. The conductance between two of the terminals is controlled by the time integral of the current in the third, rather than by its instantaneous value as in the transistor. Reproducible elements have been made which are continuously variable, which typically vary in resistance from 100Ω to 1Ω , and do this in about 10 seconds with several milliamperes of plating current. Adaptation is accomplished by dc current, while sensing the neuron logical structure is accomplished nondestructively by passing ac currents through the array of memistor cells.

None of the element values or memistor characteristics is critical, because performance feedback in the adaptation process automatically finds the best weights in any event. These neurons have been built and have adapted even with some defective memistor elements.

The first working memistors were made of ordinary pencil leads immersed in test tubes containing copper sulphate-sulphuric acid plating baths. Present elements are made by grinding down

Burrow

small $\frac{1}{10}$ watt carbon resistors so that a flat graphite surface is obtained with the resistor connections exposed. Light coats of rhodium provide smooth substrates for plating and protect the copper lead connections. These connections are insulated, and the substrates are sealed with their individual copper plating baths in polystyrene cells. These elements are small and rugged, and non-critical in manufacture. Improvements are being sought (by using different baths and different plating metals, different geometries and different substrate materials) in lifetime, and in electrical characteristics such as stability, relaxation, smoothness, and speed of adaptation.

It is expected that memistors and other components that will appear in the future will have a substantial effect in making possible cheap, simple, and reliable systems, both control and logical types.


References

1. N. Wiener, "Extrapolation, Interpolation, and Smoothing of Stationary Time Series with Engineering Applications," John Wiley and Sons, Inc., New York; 1949.
2. B. Widrow, W.H. Pierce, J.B. Angell, "Birth, Life and Death in Microelectronic Systems," IRE Transactions on Military Electronics, July 1961.
3. R.R. Brown, "Gradient methods for the computer solution of system optimization problems," WADC Technical Note 57-159, Servomechanisms Research Laboratory, M.I.T.; September, 1957.
4. J.R. Ragazzini and G.F. Franklin, "Sampled-Data Control System," McGraw-Hill Publishing Company, New York; 1958.


56

5. J.T. Tou, "Digital and Sampled-Data Control Systems," McGraw-Hill Publishing Company, New York, 1959.
6. B. Widrow, "Adaptive sampled-data systems - A statistical theory of adaptation, Wescon Conv. Rec., Inst. Radio Engrs., N.Y. Pt 4 (1959)
7. B. Widrow, "Adaptive Sampled-Data Systems", Proceedings of the First International Congress of the IFAC, Moscow, 1960.
8. H.P. Whitaker, J. Yamron, A. Kezer, "Design of Model-Reference Adaptive Control Systems for Aircraft," Report R-164, Instrumentation Laboratory, M.I.T., September, 1958.
9. D.F. DeLong, "Analysis of an adaptive sampled-data system", S.M. Thesis, Department of Electrical Engineering, M.I.T., January 1959.
10. D. Gabor, W.P.L. Wilby, and R. Woodcock, "A Universal Non-Linear Filter, Predictor and Simulator Which Optimizes Itself by a Learning Process Proceeding of the Institution of Electrical Engineers, Vol. 108, Part B., 1961.
11. B. Widrow and M.E. Hoff, "Adaptive Switching Circuits", IRE WESCON Convention Record, August 1960, part 4.
12. B. Widrow, "An Adaptive Adaline Neuron Using Chemical Memistors," 1960 Stanford University Technical Report No. 1553-2.

~~14 withdraw~~



(a) Open-Loop Adaptation



(b) Closed-Loop Adaptation

Fig. 1 Closed-Loop and Open-Loop Adaptive Predictors




Fig. 4 One-Dimensional Surface-Searching




Fig. 2 Closed-Loop Adaptation by "Performance Feedback"




Fig. 5 Derivative Measurement by Perturbation




Fig. 3 An Adjustable Sampled-Data Predictor


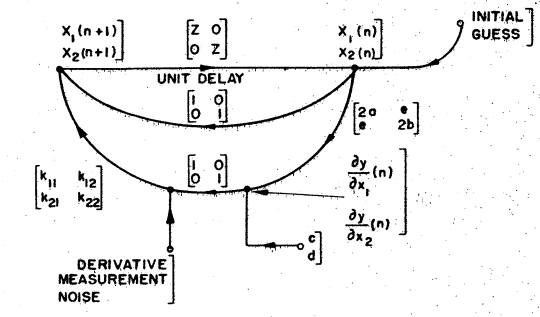
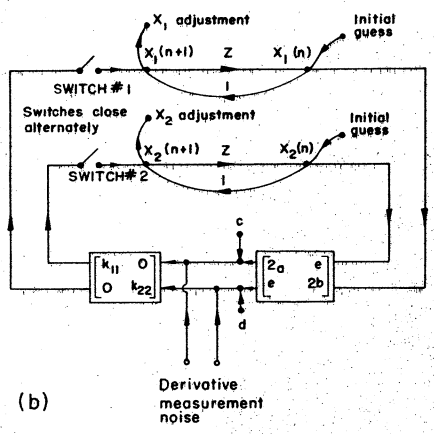


Fig. 6 Searching a Two-Dimensional Surface

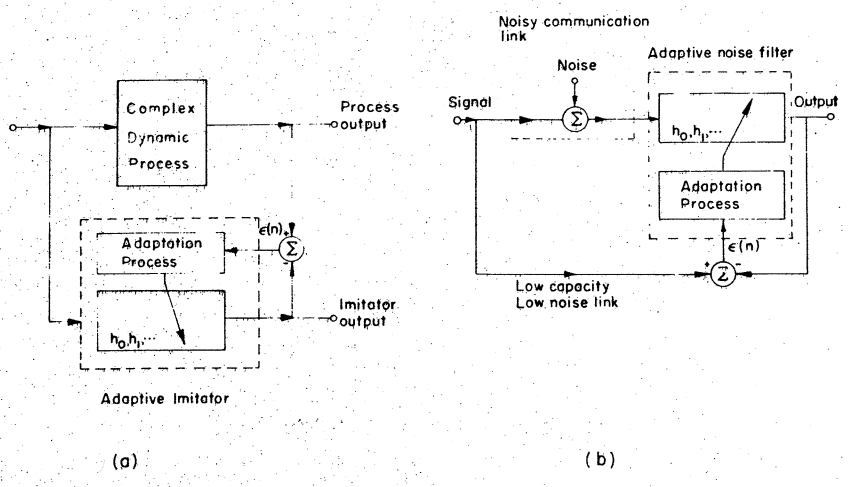


(a)



(b)

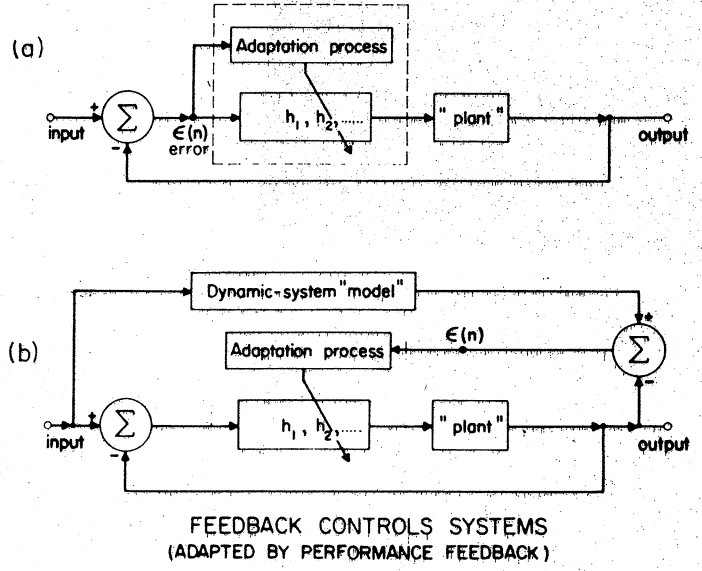
Fig. 7 Feedback Models of Two-Dimensional Surface-Searching



(a)

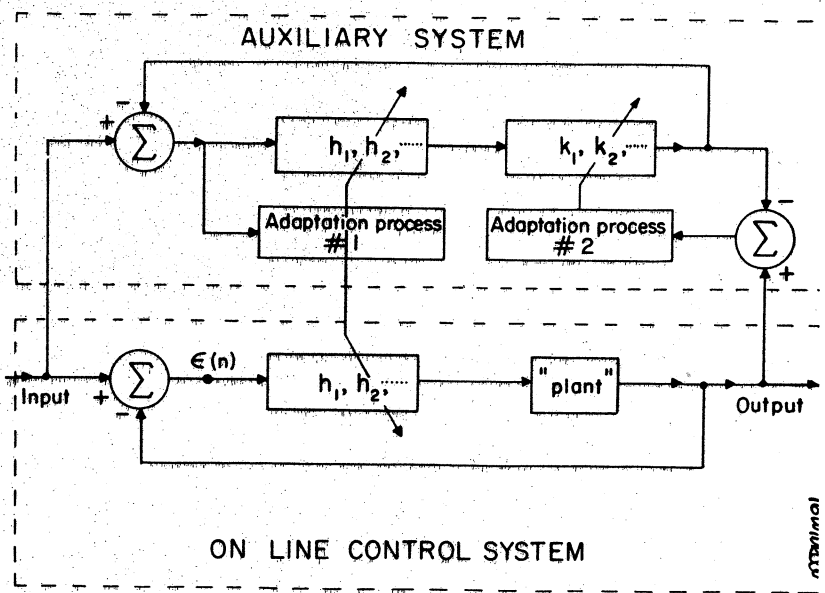
(b)

Fig. 8 Adaptive Noise Filter and Adaptive Learning Process



FEEDBACK CONTROL SYSTEMS
(ADAPTED BY PERFORMANCE FEEDBACK)

Fig. 9 Adaptive Feedback Systems; Model-Reference Adaptation



ON LINE CONTROL SYSTEM

Fig. 10 Data-Repeating With Auxiliary System

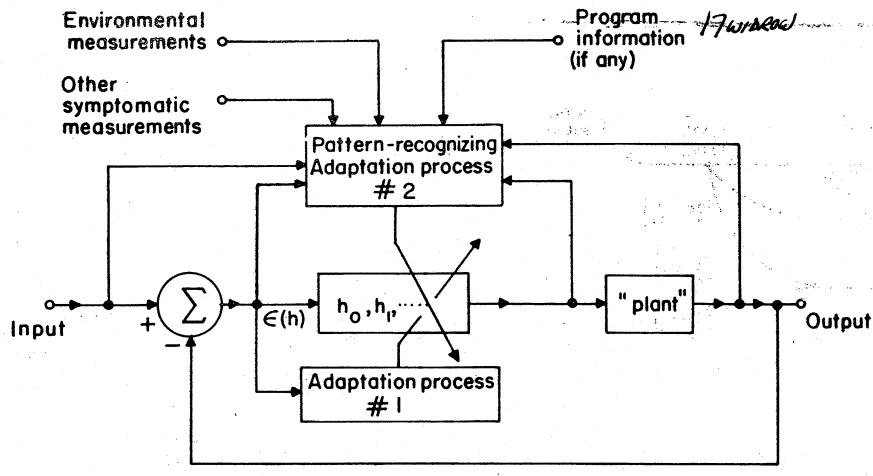


Fig. 11 Pattern-Recognizing Adaptive Feedback System




Fig. 12 ADALINE





Fig. 13 Block-Schematic of ADALINE; Error Relations

American Rocket Society

**GUIDANCE, CONTROL AND
NAVIGATION CONFERENCE**

**AUGUST 7-9, 1961
STANFORD UNIVERSITY
STANFORD, CALIF.**

THE SPEED OF ADAPTATION IN ADAPTIVE CONTROL SYSTEMS

by
Bernard Widrow
Department of Electrical Engineering
Stanford University
Stanford, California

1933-61

First publication rights reserved by
AMERICAN ROCKET SOCIETY, 500
Fifth Avenue, New York 36, N. Y. Ab-
stracts may be published without
permission if credit is given to the
author and to ARS.