

# On the Statistical Efficiency of the LMS Family of Adaptive Algorithms

Bernard Widrow and Max Kamenetsky

ISL, Department of Electrical Engineering, Stanford University, Stanford CA, USA  
e-mail: {widrow, maxk}@stanford.edu

**Abstract**—Two gradient descent adaptive algorithms are compared, the LMS algorithm and the LMS/Newton algorithm. LMS is simple and practical, and is used in many applications worldwide. LMS/Newton is based on Newton's method and the LMS algorithm. LMS/Newton is optimal in the least squares sense. It maximizes the quality of its adaptive solution while minimizing the use of training data. No other linear least squares algorithm can give better performance.

LMS is easily implemented, but LMS/Newton, although of great mathematical interest, cannot be implemented in most practical applications. Because of its optimality, LMS/Newton serves as a benchmark for all least squares adaptive algorithms. The performances of LMS and LMS/Newton are compared, and it is found that under many circumstances, both algorithms provide equal performance. For example, when both algorithms are tested with statistically nonstationary input signals, their average performances are equal. When adapting with stationary input signals and with random initial conditions, their respective learning times are on average equal. However, under worst-case initial conditions, the learning time of LMS can be much greater than that of LMS/Newton, and this is the principal disadvantage of the LMS algorithm. But the strong points of LMS are ease of implementation and optimal performance under important practical conditions. For these reasons, the LMS algorithm has enjoyed very widespread application.

## I. INTRODUCTION

LEARNING systems take on many forms. Of special interest here are the adaptive linear combiner of Fig. 1 and the adaptive transversal filter of Fig. 2. The linear combiner is the basic building block of almost all neural networks and adaptive filters. The linear combiner and the adaptive filter have found very wide application in practice [1]–[6]. For example, the present-day Internet would not exist without adaptive filters in every modem.

A learning system is generally characterized as an operator on signals, images, sounds, etc. that has adjustable parameters and that has a mechanism, an adaptive algorithm, for automatically adjusting the parameters in order to optimize the operator's performance. The adjustable parameters in Figs. 1 and 2 are adjustable weights, indicated by circles with arrows through them. The input signals are stochastic, and information obtained from the inputs is used by the adaptive algorithm to adjust the weights. The algorithm is thus a consumer of data. An efficient algorithm minimizes the usage of data while maximizing the quality of the solution, i.e. achieving parameter adjustments close to optimum. Minimizing data usage and maximizing solution quality are antagonistic. Minimizing data usage corresponds to fast adaptive convergence. But fast convergence could provide a poor quality of solution. This tradeoff

is present in all learning systems.

## II. THE LMS AND LMS/NEWTON ALGORITHMS

When referring to the linear combiner of Fig. 1, the set of weights is designated by the weight vector

$$\mathbf{w}_k = \begin{bmatrix} w_{1k} \\ w_{2k} \\ \vdots \\ w_{nk} \end{bmatrix}, \quad (1)$$

and the set of input signals to the weights is the input vector

$$\mathbf{x}_k = \begin{bmatrix} x_{1k} \\ x_{2k} \\ \vdots \\ x_{nk} \end{bmatrix}. \quad (2)$$

The output signal  $y_k$  is the inner product of  $\mathbf{x}_k$  and  $\mathbf{w}_k$ :

$$y_k = \mathbf{x}_k^T \mathbf{w}_k = \mathbf{w}_k^T \mathbf{x}_k. \quad (3)$$

The subscripts  $k$  are a time index.

The desired response  $d_k$  is an input training signal that is obtained in practice from the physical context of the application. Many examples of how desired responses may be obtained are shown in [1] for a variety of practical applications, such as prediction, noise cancelling, sensor array processing, and so forth. The error signal in Fig. 1 is  $\epsilon_k$ , the difference between the desired response and the actual output signal:

$$\epsilon_k = d_k - y_k. \quad (4)$$

Certain statistical properties of the inputs to the linear combiner are of importance. Assuming that the input and the desired response are stationary, the input autocorrelation matrix, designated by  $\mathbf{R}$ , is defined as

$$\mathbf{R} \triangleq E[\mathbf{x}_k \mathbf{x}_k^T], \quad (5)$$

and the crosscorrelation vector between the input  $\mathbf{x}_k$  and the desired response  $d_k$  is defined as

$$\mathbf{p} \triangleq E[d_k \mathbf{x}_k]. \quad (6)$$

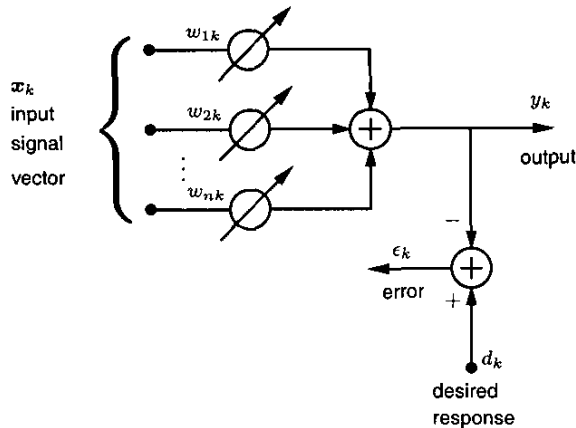


Fig. 1. Adaptive linear combiner.

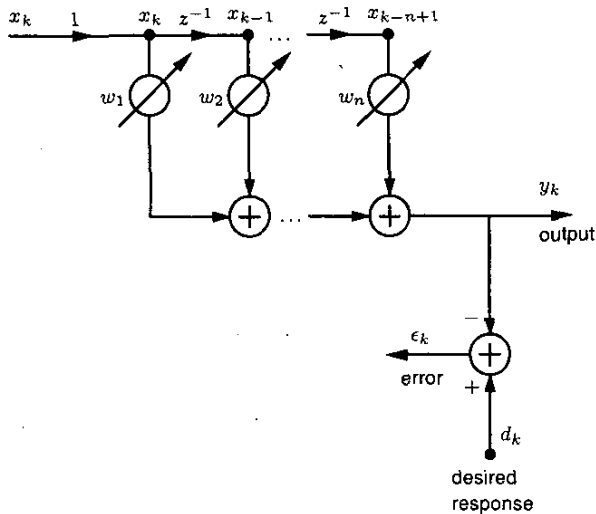


Fig. 2. Adaptive transversal digital filter.

For a given weight vector, with stationary input  $x_k$  and desired response  $d_k$  flowing into the linear combiner, the mean square error (MSE) is

$$\xi \triangleq E[\epsilon_k^2] = E[d_k^2] - 2p^T w + w^T R w. \quad (7)$$

The MSE is thus a quadratic function of the weights. If there were only two weights, the MSE could be plotted as in Fig. 3. With many weights, the surface, known as the performance surface, is a hyperparaboloid. The gradient of this surface, is a vector obtained by differentiation of Eq. (7):

$$\nabla \triangleq \begin{bmatrix} \frac{\partial \xi}{\partial w_1} \\ \vdots \\ \frac{\partial \xi}{\partial w_n} \end{bmatrix} = -2p + 2Rw. \quad (8)$$

Assuming  $R$  is not singular, the optimal solution  $w^*$  is found




Fig. 3. Sample MSE for a two-weight system.

by setting the gradient to zero:

$$w^* = R^{-1} p. \quad (9)$$

This weight vector is the best linear least squares solution and is commonly known as the Wiener solution.

The Wiener solution is of great theoretical interest, but it cannot be used directly in many practical circumstances because one would not know the statistics  $R$  and  $p$ . Input data samples are available, and they can be used on an iterative basis as they arrive for the adjustment or adaptation of the weights. The weight adjustment process that is the simplest and most widely used in the world today is the Widrow-Hoff LMS algorithm, derived in 1959 [7]. This algorithm is based on the method of steepest descent, using instantaneous gradients. It is interesting to note that the backpropagation algorithm of Werbos [8] also utilizes the method of steepest descent using instantaneous gradients in adapting the weights of a neural network. Backpropagation is a remarkable generalization of the LMS algorithm. It is the most widely used algorithm for adapting neural networks.

Steepest descent can be written as

$$w_{k+1} = w_k + \mu(-\nabla_k). \quad (10)$$

The next weight vector,  $w_{k+1}$ , equals the present weight vector,  $w_k$ , plus a change which is proportional to the negative gradient. The proportionality constant is  $\mu$ , and this is a design parameter that controls stability and rate of convergence. The LMS algorithm is obtained from Eqs. (10) and (4) as

$$\begin{cases} w_{k+1} = w_k + 2\mu\epsilon_k x_k, \\ \epsilon_k = d_k - x_k^T w_k. \end{cases} \quad (11)$$

The gradient for each iteration is instantaneous and is given by  $-2\epsilon_k x_k$ , as shown in [1].

When using steepest descent to find the minimum of a quadratic function of the weights, the weights progress geometrically toward the Wiener solution. In fact, each weight converges toward its Wiener value with a progression that is a sum of geometric progressions. Each of the geometric progressions has an individual "time constant" (the unit of time is the iteration cycle). It has been shown [1] that there are as many distinct time constants as there are distinct eigenvalues of the  $\mathbf{R}$  matrix. These time constants depend on the eigenvalues of  $\mathbf{R}$  and correspond to natural modes of the adaptive algorithm. The relative amplitudes of the modes are different from one weight to another and depend on the initial conditions of the weight vector, i.e. its initial setting.

Since one rarely has *a priori* knowledge of the orientation of the initial weight vector setting with respect to the eigenvectors of the  $\mathbf{R}$  matrix, it is difficult to predict these time constants and therefore difficult to predict the rate of convergence of the LMS algorithm. In spite of this drawback, LMS is very widely used.

A more predictable algorithm is Newton's method. Assuming that  $\mathbf{R}$  is not singular, it can be written as

$$\mathbf{w}_{k+1} = \mathbf{w}_k + \mu \lambda_{\text{ave}} \mathbf{R}^{-1} (\nabla_k). \quad (12)$$

Here, the negative gradient is premultiplied by  $\mu \lambda_{\text{ave}} \mathbf{R}^{-1}$ , where  $\lambda_{\text{ave}}$  is the average of the eigenvalues of  $\mathbf{R}$ . Based on Eq. (12) and using the instantaneous gradient  $-2\epsilon_k \mathbf{x}_k$ , the LMS/Newton algorithm can be written as

$$\begin{cases} \mathbf{w}_{k+1} = \mathbf{w}_k + 2\mu \lambda_{\text{ave}} \mathbf{R}^{-1} \epsilon_k \mathbf{x}_k, \\ \epsilon_k = d_k - \mathbf{x}_k^T \mathbf{w}_k. \end{cases} \quad (13)$$

With Newton's method, there is only one natural mode, corresponding to one time constant that is independent of the weight vector's initial conditions. It is interesting to note that Eq. (13) is identical to Eq. (11) when all of the eigenvalues of  $\mathbf{R}$  are equal.

So, the good news about LMS/Newton is that its rate of convergence is predictable and does not depend on initial conditions. The bad news is that one cannot implement this algorithm in practice because  $\mathbf{R}^{-1}$  is generally unknown. LMS, based on steepest descent, has disadvantages, but it is simple and easy to implement. Also, as will be shown, it performs equivalently to LMS/Newton under many important conditions.

### III. LEARNING WITH A FINITE SET OF DATA SAMPLES

The LMS/Newton algorithm is not only very predictable in its convergence behavior, but it is also highly efficient in its use of input data. It will be shown that in the least squares sense, no other algorithm can be more efficient in its data usage than LMS/Newton.

In order to study the question of efficiency for the LMS/Newton algorithm, it is useful to contemplate training the linear combiner with a finite number of data samples. One data sample consists of an  $\mathbf{x}$  vector and its associated desired response.

Assume that a set of  $N$  training input samples and associated desired responses is drawn from a given distribution. Define a matrix  $\mathbf{X}$  for these training samples as the set of  $N$   $\mathbf{x}$  vectors:

$$\mathbf{X} \triangleq [\mathbf{x}_1 \quad \mathbf{x}_2 \quad \cdots \quad \mathbf{x}_N]^T. \quad (14)$$

Define a desired response vector for these samples as

$$\mathbf{d} \triangleq [d_1 \quad d_2 \quad \cdots \quad d_N]^T. \quad (15)$$

Define an error vector for these samples as

$$\boldsymbol{\epsilon} \triangleq [\epsilon_1 \quad \epsilon_2 \quad \cdots \quad \epsilon_N]^T. \quad (16)$$

The objective is to find a set of weights that will minimize the sum of the squares of the errors for the training sample, i.e. minimize  $\boldsymbol{\epsilon}^T \boldsymbol{\epsilon}$ . Assuming  $\mathbf{X}$  is full rank, one obtains the optimal least squares solution as [1], [9]

$$\mathbf{w}_{\text{LS}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{d}. \quad (17)$$

One could directly calculate this solution from the data, or the same result could be obtained by training the weights with the LMS algorithm or any other least squares algorithm, recirculating the training data over and over until the weights converge and stabilize. It should be noted that if the number  $N$  were increased without bound, the finite least squares solution of Eq. (17) would converge in the mean-square sense to the true Wiener solution for the given distribution of training samples and desired responses.

In practice, one would have only a finite number of data samples available for training. The question is, given  $N$  training samples, how well does the optimal solution of Eq. (17) perform compared to the Wiener solution when operating on an infinite set of samples drawn from the same distribution? The true Wiener solution yields the minimum MSE. The least squares solution based on  $N$  training samples produces more MSE than the true Wiener solution. There is an excess MSE. If a different set of  $N$  data samples were randomly selected from the same distribution, a different optimal solution (17) would result, with a different excess MSE. Given an ensemble of randomly selected data sets, each with  $N$  samples of data, there would be an ensemble of optimal solutions (17) with an ensemble of excess MSEs. One could take an ensemble average of the excess MSEs and normalize with respect to the minimum MSE of the true Wiener solution. This dimensionless ratio is called the misadjustment due to training with  $N$  data samples:

$$\begin{aligned} M &\triangleq \frac{\text{(average excess MSE)}}{\text{(minimum MSE)}} \\ &= \frac{\text{(number of weights)}}{\text{(number of training samples)}} = \frac{n}{N}. \end{aligned} \quad (18)$$

A derivation of this result is given in [10] and [11]. It was first reported by Widrow and Hoff at a WESCON conference

in 1960 [7]. Years of experience with adaptive filters and neural networks lead one to accept a misadjustment of 10% as a reasonable design value. This gives a performance that is only 10% worse than that of the optimal Wiener solution. When training a linear combiner, a 10% misadjustment is obtained when the number of training samples is equal to 10 times the number of weights. When training a neural network having only one neuron, the number of training patterns would be equal to 10 times the number of weights. When training an adaptive filter with a steady flow of stationary input data, the parameter  $\mu$  should be chosen so that the training time or convergence time (several time constants of the learning curve) would be equal to 10 times the length of its impulse response in order to have a misadjustment of 10%. There is no similar law for multilayered neural networks trained with backpropagation, but one may speculate that for a network with many inputs and a single output, regardless of the number of neurons and layers, training with a number of patterns equal to 10 times the number of inputs should give good performance.

#### IV. LEARNING WITH LMS/NEWTON

Consider the linear combiner, with stationary input data flowing in real time. Starting from an initial (non-optimal) weight vector and adapting with the LMS/Newton algorithm, the weight vector will geometrically (exponentially) relax toward the Wiener solution, with noise superposed. The noise originates from adapting with instantaneous gradients. (Estimating the gradient with a single data sample produces a noisy gradient.) As the weights relax toward the Wiener solution, the MSE also relaxes exponentially toward the level of the minimum mean square error  $\xi_{\min}$ , with noise superposed. Because of gradient noise, the weights do not converge on the Wiener solution, but undergo Brownian motion about it. The MSE is almost always greater than  $\xi_{\min}$  and never goes below it, because of the noise in the weights while adapting. A plot of MSE versus number of iterations is shown in Fig. 4, where  $\xi_{\text{fin}}$  is the asymptotic MSE, and  $\xi_{\text{excess}}$  is the difference between it and the minimum MSE  $\xi_{\min}$ . This type of plot is called the "learning curve". Note that, in general,  $\xi_{\text{excess}} > 0$  because the misadjustment is not zero, and thus, the asymptotic MSE  $\xi_{\text{fin}}$  is greater than  $\xi_{\min}$ , which is reflected in Fig. 4.

After gross adaptive transients have died out (after about four exponential time constants), the weights are in a steady state Brownian motion and the MSE hovers randomly above  $\xi_{\min}$ , exhibiting "excess mean square error". The misadjustment for this situation is defined as the average of the excess mean square error divided by  $\xi_{\min}$ :

$$M \triangleq \frac{\text{(average excess MSE)}}{\text{(minimum MSE)}}. \quad (19)$$

The bigger one makes  $\mu$ , the faster the algorithm converges, and the greater is the misadjustment. It has been shown [12], [1], [11] that a good approximation for the misadjustment

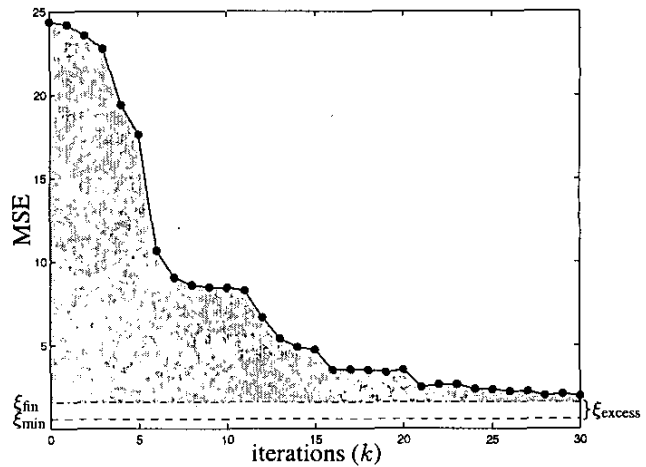


Fig. 4. Sample learning curve with gradient noise.

is

$$M = \mu \text{Tr}(\mathbf{R}) = \frac{n}{4\tau_{\text{MSE}}}, \quad (20)$$

where  $\text{Tr}(\mathbf{R})$  is the trace of  $\mathbf{R}$ , and  $\tau_{\text{MSE}}$  is the time constant of the mean square error learning curve. The unit of time is the iteration cycle. Time is marked by number of iteration cycles, and since a new input training sample is used with each iteration cycle, time can equivalently be marked by number of training samples used.

#### V. OPTIMALITY OF LMS/NEWTON IN A STATIONARY ENVIRONMENT

The LMS/Newton algorithm exponentially weights its input data over time as it establishes its weight values. The settling time of the adaptive process is of the order of four time constants of the MSE learning curve. At any moment, the weights are determined by adaptation that has taken place over essentially the last four time constants worth of data. Thus, in a steady flow situation, the training data "consumed" or "absorbed" at any time by the LMS/Newton algorithm is essentially comprised of the most recent  $4\tau_{\text{MSE}}$  samples. From Eq. (20), the misadjustment of the LMS/Newton algorithm can therefore be expressed as

$$M = \frac{n}{4\tau_{\text{MSE}}} = \frac{\text{(number of weights)}}{\text{(number of independent training samples)}}. \quad (21)$$

When learning with a finite set of data samples, the optimal weight vector is the best least squares solution for that set of samples, and it is often called the "exact least squares solution". This solution, given by Eq. (17), makes the best use of the finite number of data samples, in the least squares sense. All of the data are weighted equally in affecting the solution. This solution will of course vary from one set of samples to

another, and it has a misadjustment of

$$M = \frac{(\text{number of weights})}{(\text{number of independent training samples})} \quad (22)$$

For the same consumption of data, it is apparent that LMS/Newton and exact least squares yield the same misadjustment. Although we are comparing “apples with oranges” by comparing a steady flow algorithm with an algorithm that learns with a finite number of data samples, we nevertheless find that LMS/Newton is as efficient as exact least squares when we relate the quality of the weight-vector solution to the amount of data used in obtaining it. Since the exact least squares solution makes optimal use of the data, so does LMS/Newton.

### VI. TRANSIENT LEARNING WITH STATIONARY INPUT DATA

The LMS/Newton algorithm makes optimal use of its training data. Its learning curve is exponential with a single time constant. For practical purposes, its convergence time is of the order of four time constants, although in principle, convergence would take forever. Fast convergence is desirable because, during the initial learning transient, the MSE is excessive. The LMS algorithm has a learning curve which is a sum of exponentials. But, when all of the eigenvalues of the  $R$  matrix are equal, LMS has a single exponential learning curve that is identical to that of LMS/Newton; and LMS is therefore optimal. Generally, the eigenvalues are not equal, and LMS has a different kind of learning curve than LMS/Newton.

After learning transients die out, the steady state misadjustment for LMS/Newton is


$$M = \mu \text{Tr}(R), \quad (23)$$

and the steady state misadjustment for LMS is


$$M = \mu \text{Tr}(R). \quad (24)$$

The derivations of Eq. (23) and Eq. (24) are given in [1] and [11]. A fair and reasonable way to compare two algorithms is to adjust their rates of convergence so that they would have the same steady state misadjustment. One could always make an algorithm “look good” with fast adaptation, but this causes excessive misadjustment. To compare LMS with LMS/Newton, their  $\mu$ -parameters should be set to the same value.

Fig. 5 shows plan views of a quadratic MSE surface, Fig. 5(a) indicating adaptive steps for Newton’s method and Fig. 5(b) showing corresponding steps for the method of steepest descent with equivalent initial conditions. These steps correspond to three adaptive transient experiments, each starting from a different point on the same contour of constant MSE and operating with the same value of  $\mu$ . The steps using Newton’s method are always directed toward the bottom of the quadratic bowl, whereas those of steepest descent follow the local gradient, orthonormal to the contours of constant MSE.



(a)



(b)

Fig. 5. Illustration of Newton’s method versus steepest descent: (a) Newton’s method, (b) steepest descent. The Wiener solution is indicated by \*. The three initial conditions are indicated by O.

Fig. 6 shows learning curves corresponding to the adaptive steps illustrated in Fig. 5. All three learning curves derived from Newton’s method are identical, since the initial starting conditions are located on the same constant MSE contour, and all three time constants are the same. Fig. 6 shows all three learning curves as a single curve labeled “Newton’s method”. The three steepest descent curves are distinct, having individual time constants. The curves corresponding to initial conditions falling on an eigenvector (a principal axis of the elliptical contours) are pure exponentials, whereas the curve corresponding to the initial condition between the eigenvectors is a sum of two exponentials.

Which algorithm converges faster? For some initial condi-

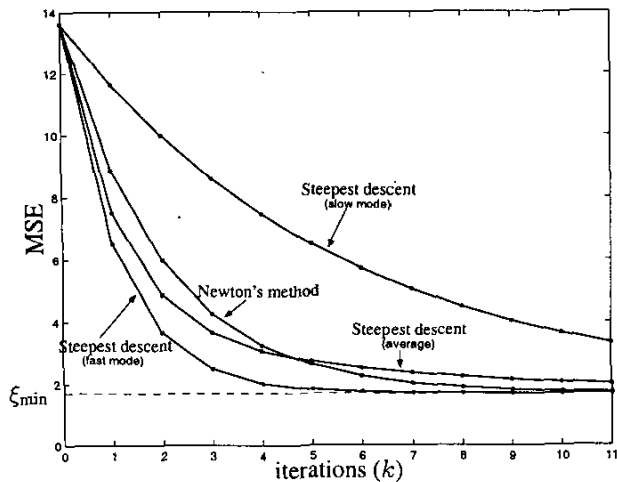


Fig. 6. Steepest descent and Newton's method learning curves.

tions, LMS converges faster than LMS/Newton. For other initial conditions, LMS is slower than LMS/Newton. Yet for other initial conditions, LMS has multiple modes, some faster than LMS/Newton and some slower than it. The question is: which is faster, LMS or LMS/Newton?

#### VII. EXCESS ERROR ENERGY

The question "which is faster" prompts a new look at the learning curve and at the issue of learning time. Regarding the learning curves of Fig. 7, it is clear that during the transient, the MSE is excessive. One strives to reduce this MSE to  $\xi_{\min}$  as fast as possible. The area under the curve and above  $\xi_{\min}$ , the shaded area, is defined as the excess error energy. This area is a sum of MSE over time. One would like this area to be as small as possible.

The learning curve shown in Fig. 7(a) is an exponential having a single time constant. The area under this curve, the excess error energy, is equal to the amplitude of the exponential (its initial excess MSE) multiplied by the time constant. The learning time for this curve may be considered to be four time constants. This is somewhat arbitrary, but reasonable. Accordingly, the learning time for the single exponential is defined as

$$\left( \text{learning time} \right) \triangleq 4 \times \frac{(\text{excess error energy})}{(\text{initial excess MSE})}. \quad (25)$$

The learning curve shown in Fig. 7(b) is a sum of two exponentials. Its excess error energy is the same as that of the single exponential learning curve of Fig. 7(a), and the initial excess MSE of both curves is the same. The learning times of both curves are therefore the same and are both given by Eq. (25). In general, the learning time of a learning curve which is a sum of any number of exponentials will be defined hereby as given by Eq. (25). Starting from the same initial MSE, the convergence times of two different learning curves are defined as being identical if their respective excess error energies are equal.





Fig. 7. Idealized learning curves (no gradient noise): (a) single-exponential, (b) two-exponential. The shaded area represents the excess error energy.

Excess error energy is proportional to learning time as it is defined here.

The question remains, which has the greater learning time, LMS or LMS/Newton? Referring to Fig. 6, one can see that LMS/Newton has a fixed learning time, but the learning time of LMS depends on initial conditions. Sometimes LMS is faster than LMS/Newton, sometimes slower. Under worst-case initial conditions, LMS can have a much greater learning time than LMS/Newton. However, it has been shown [11] that with random initial conditions, the average learning time with LMS is identical to the unique learning time of LMS/Newton. LMS/Newton is optimal and its performance is the benchmark. It is important to realize that the average performance of LMS is identical to that of LMS/Newton, when comparing both algorithms starting from the same randomly chosen initial con-

ditions. One must be careful to note that this does *not* mean that the learning curves of LMS and LMS/Newton are identical, even when averaged over initial conditions. On the contrary, it can be shown that the average initial convergence of LMS is faster than that of LMS/Newton, whereas the average final convergence of LMS is slower than that of LMS/Newton [11]. However, their average learning times and average excess error energies are the same.

### VIII. LMS AND LMS/NEWTON IN A NONSTATIONARY ENVIRONMENT

With statistically stationary inputs, the quadratic performance surface is fixed, and the Wiener solution is fixed. With nonstationary inputs, this surface changes randomly, and the Wiener solution is not fixed but is a randomly moving target. There is an analogy between, on the one hand, transient adaptation from random initial conditions toward a fixed Wiener target and, on the other hand, steady state adaptation toward a randomly moving Wiener target. In this section, a comparison is made of the performances of the LMS algorithm and the LMS/Newton algorithm when adapting with nonstationary inputs of a simple form.

Filtering nonstationary signals is a major area of application for adaptive systems. When the statistical character of an input signal changes gradually, randomly, and unpredictably, a filtering system that can automatically optimize its input-output response in accord with the requirements of the input signal could yield superior performance relative to that of a fixed, non-adaptive system. The performance of the conventional steepest descent LMS algorithm is compared here with LMS/Newton (which, as demonstrated above, possesses optimality qualities), when both algorithms are used to adapt transversal filters with nonstationary inputs. The nonstationary situations to be studied are highly simplified, but they retain the essence of the problem that is common to more complicated and realistic situations.

The example considered here involves modeling or identifying an unknown time-varying system by an adaptive LMS transversal filter of length  $n$ . The unknown system is assumed to be a transversal filter of the same length  $n$  whose weights (impulse response values) vary as independent stationary ergodic first-order Markov processes, as indicated in Fig. 8. The input signal  $x_k$  is assumed to be stationary and ergodic. Additive output noise, assumed to be stationary and ergodic, of mean zero and of variance  $\xi_{\min}$ , prevents a perfect match between the unknown system and the adaptive system. The minimum MSE is, therefore,  $\xi_{\min}$ , and it is achieved whenever the weights of the adaptive filter,  $w_k$ , match those of the unknown system. The latter are at every instant the optimal values for the corresponding weights of the adaptive filter and are designated as  $w_k^*$ , the time index indicating that the unknown "target" to be tracked is time-varying.

The components of  $w_k^*$  are generated by passing independent white noises of variance  $\sigma^2$  through identical one-pole low-pass filters. The components of  $w_k^*$  therefore vary as in-

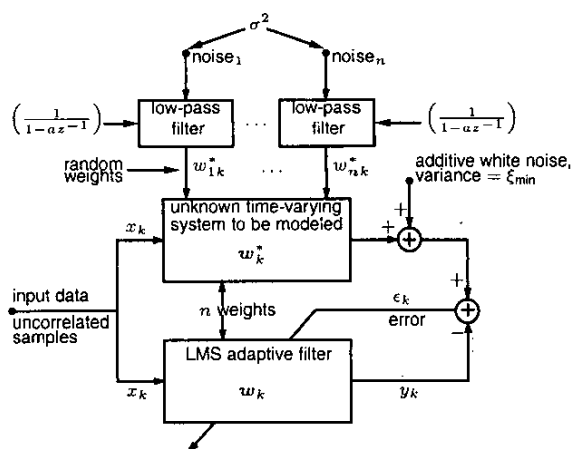


Fig. 8. Modeling an unknown time-varying system.

dependent first-order Markov processes. The formation of  $w_k^*$  is illustrated in Figs. 8 and 9.

According to the scheme of Fig. 8, minimizing the MSE causes the adaptive weight vector  $w_k$  to attempt to best match the unknown  $w_k^*$  on a continual basis. The  $R$  matrix, dependent only on the statistics of  $x_k$ , is constant even as  $w_k^*$  varies. The desired response of the adaptive filter,  $d_k$ , is nonstationary, being the output of a time-varying system. The minimum MSE,  $\xi_{\min}$ , is constant. Thus the MSE function, a quadratic bowl, varies in position while its eigenvalues, eigenvectors, and  $\xi_{\min}$  remain constant.

In order to study this form of nonstationary adaptation both analytically and by computer simulation, a model comprising an ensemble of nonstationary adaptive processes has been defined and constructed as illustrated in Fig. 9. Throughout the ensemble, the unknown filters to be modeled are all identical and have the same time-varying weight vector  $w_k^*$ . Each ensemble member has its own independent input signal going to both the unknown system and the corresponding adaptive system. The effect of output noise in the unknown systems is obtained by the addition of independent noises of variance  $\xi_{\min}$ . All of the adaptive filters are assumed to start with the same initial weight vector  $w_0$ ; each develops its own weight vector over time in attempting to pursue the moving Markovian target  $w_k^*$ .

For a given adaptive filter, the weight-vector tracking error at the  $k$ th instant is  $v_k \triangleq w_k - w_k^*$ . This error is due to both the effects of gradient noise and weight-vector lag and may be expressed as

$$v_k = w_k - w_k^* = \underbrace{(w_k - E[w_k])}_{\text{weight-vector noise}} + \underbrace{(E[w_k] - w_k^*)}_{\text{weight-vector lag}}. \quad (26)$$

The expectations are averages over the ensemble. Eq. (26) identifies the two components of the error. Any difference be-

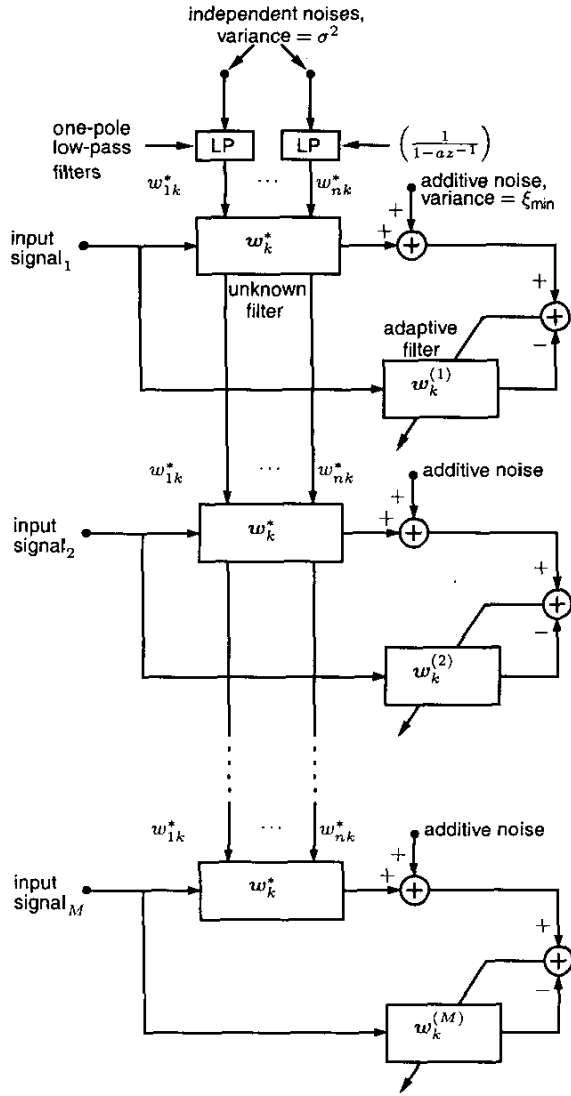


Fig. 9. An ensemble of nonstationary adaptive processes.

tween the ensemble mean of the adaptive weight vectors and the target value  $w_k^*$  is due to lag in the adaptive process, while the deviation of the individual adaptive weight vectors about the ensemble mean is due to gradient noise.

Weight-vector error causes an excess MSE. The ensemble average excess MSE at the  $k$ th instant is

$$\left( \begin{array}{c} \text{average} \\ \text{excess MSE} \end{array} \right)_k = E \left[ (w_k - w_k^*)^T R (w_k - w_k^*) \right]. \quad (27)$$

Using Eq. (26), this can be expanded as follows:

$$\begin{aligned} \left( \begin{array}{c} \text{average} \\ \text{excess MSE} \end{array} \right)_k &= E \left[ (w_k - E[w_k])^T R (w_k - E[w_k]) \right] \\ &\quad + E \left[ (E[w_k] - w_k^*)^T R (E[w_k] - w_k^*) \right] \\ &\quad + 2E \left[ (w_k - E[w_k])^T R (E[w_k] - w_k^*) \right]. \end{aligned} \quad (28)$$

Expanding the last term of Eq. (28) and simplifying since  $w_k^*$  is constant over the ensemble,

$$\begin{aligned} 2E \left[ w_k^T R E[w_k] - w_k^T R w_k^* - E[w_k]^T R E[w_k] \right] \\ + E \left[ w_k^T R w_k^* \right] = 2 \left( E[w_k]^T R E[w_k] - E[w_k]^T R E[w_k] \right) \\ - E[w_k]^T R w_k^* + E[w_k]^T R w_k^* = 0. \end{aligned} \quad (29)$$

Therefore, Eq. (28) becomes

$$\begin{aligned} \left( \begin{array}{c} \text{average} \\ \text{excess MSE} \end{array} \right)_k &= E \left[ (w_k^* - E[w_k])^T R (w_k - E[w_k]) \right] \\ &\quad + E \left[ (E[w_k] - w_k^*)^T R (E[w_k] - w_k^*) \right]. \end{aligned} \quad (30)$$

The average excess MSE is thus a sum of components due to both gradient noise and lag:

$$\begin{aligned} \left( \begin{array}{c} \text{average excess} \\ \text{MSE due to lag} \end{array} \right)_k &= E \left[ (E[w_k] - w_k^*)^T R (E[w_k] - w_k^*) \right] \\ &= E \left[ (E[w'_k] - w_k^*)^T \Lambda (E[w'_k] - w_k^*) \right] \end{aligned} \quad (31)$$

$$\begin{aligned} \left( \begin{array}{c} \text{average excess MSE} \\ \text{due to gradient noise} \end{array} \right)_k &= E \left[ (w_k - E[w_k])^T R (w_k - E[w_k]) \right] \\ &= E \left[ (w'_k - E[w'_k])^T \Lambda (w'_k - E[w'_k]) \right], \end{aligned} \quad (32)$$

where  $w'_k \triangleq Q^T w_k$ ,  $w_k^* \triangleq Q^T w_k^*$ ,  $\Lambda$  is a diagonal matrix of the eigenvalues of  $R$ , and  $Q$  is the modal matrix of  $R$ . The total misadjustment is therefore a sum of two components, that due to lag and that due to gradient noise. These components of misadjustment have been evaluated by Widrow *et al.* [10]. The total misadjustment for adaptation with the LMS algorithm is

$$\begin{aligned} M_{\text{sum}} &= \left( \begin{array}{c} \text{misadjustment} \\ \text{due to gradient noise} \end{array} \right) + \left( \begin{array}{c} \text{misadjustment} \\ \text{due to lag} \end{array} \right) \\ &= \mu \text{Tr}(R) + \frac{n\sigma^2}{4\mu\xi_{\min}}. \end{aligned} \quad (33)$$

Since  $M_{\text{sum}}$  is convex in  $\mu$ , an optimal choice of  $\mu$  that minimizes  $M_{\text{sum}}$  can be obtained by differentiating  $M_{\text{sum}}$  with respect to  $\mu$  and setting the derivative to zero. Optimization takes



place when the two terms of Eq. (33) are made equal. When this happens, the loss in performance from adapting too rapidly (due to gradient noise) is equal to the loss in performance from adapting too slowly (due to lag).

It is interesting to note that  $M_{\text{sum}}$  in Eq. (33) depends on the choice of the parameter  $\mu$  and on the statistical properties of the nonstationary environment but does not depend on the spread of the eigenvalues of the  $\mathbf{R}$  matrix. It is no surprise, therefore, that when the components of misadjustment are evaluated for the LMS/Newton algorithm operating in the very same environment, the expression for  $M_{\text{sum}}$  for LMS/Newton turns out to be

$$M_{\text{sum}} = \left( \begin{array}{c} \text{misadjustment} \\ \text{due to gradient noise} \end{array} \right) + \left( \begin{array}{c} \text{misadjustment} \\ \text{due to lag} \end{array} \right) \\ = \mu \text{Tr}(\mathbf{R}) + \frac{n\sigma^2}{4\mu\xi_{\min}}, \quad (34)$$

which is the same as Eq. (33). From this we may conclude that the performance of the LMS algorithm is equivalent to that of the LMS/Newton algorithm when both are operating with the same choice of  $\mu$  in the same nonstationary environment, wherein they are tracking a first-order Markov target. The optimum value of  $\mu$  is the same for LMS/Newton as for LMS. Since LMS/Newton is optimal, we may conclude that the LMS algorithm is also optimal when operating in a first-order Markov nonstationary environment. And it is likely optimal or close to it when operating in many other types of nonstationary environments, although this has not yet been proven.

## IX. CONCLUSION

An adaptive algorithm is like an engine whose fuel is input data. Two algorithms adapting the same number of weights and operating with the same misadjustment can be compared in terms of their consumption of data. The more efficient algorithm consumes less data, i.e. converges faster. On this basis, the LMS/Newton algorithm has the highest statistical efficiency that can be obtained. The LMS/Newton algorithm therefore can serve as a benchmark for statistical efficiency against which all other algorithms can be compared.

The role played by LMS/Newton in adaptive systems is analogous to that played by the Carnot engine in thermodynamics. They both do not exist physically. But their performances limit the performances of all practical systems, adaptive and thermodynamic, respectively.

The LMS/Newton algorithm uses learning data most efficiently. No other learning algorithm can be more efficient. The LMS algorithm performs identically to LMS/Newton when all eigenvalues of the input autocorrelation matrix (the  $\mathbf{R}$  matrix) are equal. In most practical cases, however, the eigenvalues are not equal. Regardless of eigenvalue spread, the LMS algorithm performs equivalently, on average, to LMS/Newton in nonstationary environments and under transient learning conditions with random initial conditions. However, under "worst case" initial conditions, LMS can converge much more slowly

than LMS/Newton. Under "best case" initial conditions, LMS converges much faster than LMS/Newton. On average, their convergence rates are equivalent in terms of their excess error energies. Along with its simplicity, ease of implementation, and robustness, the equivalent performance between LMS and LMS/Newton is one of the major reasons for the popularity of the LMS algorithm. Dr. Widrow has discussed this with Dr. Paul Werbos, and both conclude that the backpropagation algorithm is popular for similar reasons.

## REFERENCES

- [1] B. Widrow and S. D. Stearns, *Adaptive Signal Processing*, Prentice Hall, Upper Saddle River, NJ, 1985.
- [2] B. Widrow, P. Mantey, L. Griffiths, and B. Goode, "Adaptive antenna systems," *Proc. IEEE*, vol. 55, no. 12, pp. 2143–2159, Dec. 1967.
- [3] B. Widrow, J. R. Glover, Jr., J. M. McCool, J. Kaunitz, C. S. Williams, R. H. Hearn, J. R. Zeidler, E. Dong, Jr., and R. C. Goodlin, "Adaptive noise cancelling: Principles and applications," *Proc. IEEE*, vol. 63, no. 12, pp. 1692–1716, Dec. 1975.
- [4] J. G. Proakis, *Digital Communications*, chapter 11, McGraw-Hill, New York, 4th edition, 2001.
- [5] M. Ghogho, M. Ibnkahla, and N. J. Bershad, "Analytic behavior of the LMS adaptive line enhancer for sinusoids corrupted by multiplicative and additive noise," *IEEE Trans. Signal Process.*, vol. 46, no. 9, pp. 2386–2393, Sep. 1998.
- [6] B. Widrow and E. Walach, *Adaptive Inverse Control*, Prentice Hall, Upper Saddle River, NJ, 1996.
- [7] B. Widrow and M. E. Hoff, "Adaptive switching circuits," in *IRE WESCON Conv. Rec.*, 1960, vol. 4, pp. 96–104.
- [8] P. J. Werbos, *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*, Ph.D. thesis, Harvard University, Cambridge, MA, August 1974.
- [9] T. Kailath, A. H. Sayed, and B. Hassibi, *Linear Estimation*, Prentice Hall, Upper Saddle River, NJ, 2000.
- [10] B. Widrow, J. McCool, M. G. Larimore, and C. R. Johnson, Jr., "Stationary and nonstationary learning characteristics of the LMS adaptive filter," *Proc. IEEE*, vol. 64, no. 8, pp. 1151–1162, Aug. 1976.
- [11] B. Widrow and M. Kamenetsky, "On the efficiency of adaptive algorithms," in *Least-Mean-Square Adaptive Filters*, S. Haykin and B. Widrow, Eds. John Wiley & Sons, New York, 2003.
- [12] B. Widrow and E. Walach, "On the statistical efficiency of the LMS algorithm with nonstationary inputs," *IEEE Trans. Inf. Theory*, vol. IT-30, no. 2, pp. 211–221, Mar. 1984.