

NEURAL NETWORK TRAINING VIA QUADRATIC OPTIMIZATION

Michael A. Sartori¹ and Panos J. Antsaklis²

¹Carderock Division - Code 1941
Naval Surface Warfare Center
Bethesda, Maryland 20084

²Department of Electrical Engineering
University of Notre Dame
Notre Dame, Indiana 46556

Abstract - A new technique using quadratic optimization is proposed to find the weights of a single neuron, or a single-layer neural network, and extended to the multi-layer neural network. It is proposed here to find the weights for a neuron by minimizing a cost function that is quadratic with respect to the neuron's weights and to use these weights as an answer for minimizing a cost function that is quadratic with respect to the neuron's outputs. Previous methods, such as the least mean squares algorithm which is a gradient descent method and a precursor of the back-propagation algorithm, iteratively find weights for the neuron which minimize the cost function directly involving the nonlinearity of the neuron. By back-propagating the output error through the neural network's layers, the proposed method is extended to the multi-layer neural network. The described Quadratic Optimization Algorithm for the multi-layer neural network tends to work best for classification problems and tends to achieve successful results in a single iteration.

I INTRODUCTION

A new training method based on quadratic optimization is presented in this paper to find the weights of a single neuron, or a single-layer neural network, and is extended to a multi-layer neural network. Instead of minimizing a cost function that directly involves the nonlinearity of the neuron, a function which is quadratic with respect to the neuron's weights is minimized. The solution from this minimization problem is used as a solution for the original problem. In [1,2], the relationship between the two problems is established through a careful error analysis and an examination of the relationship between the minima. Due to the class of nonlinear functions often chosen for the neuron (e.g., the hyperbolic tangent function or the signum function), the error for using the solution from the quadratic minimization as a solution for the original problem is small, and even zero if the error from solving the quadratic minimization is zero, which is a case studied here. Furthermore, with the quadratic optimization procedure used to find the weights of the single neuron, it always converges and is faster than using a gradient descent algorithm on the original problem. By back-propagating the output error through a multi-layer neural network's hidden layers, the proposed method is extended to the multi-layer case. The so-called Quadratic Optimization Algorithm for the multi-layer neural network tends to work best for classification problems and tends to achieve successful results in a single iteration.

In Section II, the neuron considered in this paper and the problem of finding its weights, termed here the Neuron Training Problem (N), are defined. In Section III, the Neuron Quadratic Optimization Problem (NQ) is defined, and its solution is used as one for Problem (N). In Section IV, the single-layer neural network and the Single-Layer Neural Network Training Problem (L) are defined. The quadratic optimization procedure is then described for the single-layer neural network in terms of the Problem (L) and termed the Single-Layer Neural Network Quadratic Optimization Problem (LQ). In Section V, the multi-layer neural network and the Multi-Layer Neural Network Training Problem (M) are defined, and the back-propagation algorithm, one of the most common methods used to train the multi-layer neural network, is described. In Section VI, it is proposed to solve the Problem (LQ) for each layer of the multi-layer neural network by back-propagating the output layer's error to each hidden layer. The resulting procedure is termed the Quadratic Optimization Algorithm. Finally, in Section VII, two examples are given that illustrate the training procedure of this paper.

II THE NEURON

The neuron considered here is described by

$$y = f\left(\sum_{i=1}^m u_i w_i\right) = f(u'w), \quad (1)$$

where $f: \mathbb{R} \rightarrow \mathbb{R}$ is the nonlinearity of the neuron, $u = [u_1, \dots, u_m]^T \in \mathbb{R}^{m \times 1}$ is the input vector, $w = [w_1, \dots, w_m]^T \in \mathbb{R}^{m \times 1}$ is the weight vector, and $u_m = 1$ is the bias input for the neuron. The type of nonlinearity of the neuron is restricted to those functions commonly used in neuron models (e.g., the hyperbolic tangent function or the signum function).

Assume that a training set $\{u(j), d(j)\}$ for $1 \leq j \leq p$ consists of p pairs of input vectors and desired output scalars, where $u(j) \in \mathbb{R}^{m \times 1}$, $u_m(j) = 1$, and $d(j) \in \mathbb{R}$ for $1 \leq j \leq p$. The Neuron Training Problem (N) is defined as follows:

$$\min_w \hat{F}(w) \quad (N)$$

$$\hat{F}(w) = (d - \phi(U'w))(d - \phi(U'w))'$$

and where $d = [d(1), \dots, d(p)]^T \in \mathbb{R}^{p \times 1}$ is the desired output vector, $U = [u(1), \dots, u(p)] \in \mathbb{R}^{m \times p}$ is the matrix of input vectors, and $\phi(z) = [f(z_1), \dots, f(z_p)]^T \in \mathbb{R}^{p \times 1}$ with $z = [z_1, \dots, z_p]^T \in \mathbb{R}^{p \times 1}$. The notation $\phi(z)$ represents a map which takes a p -dimensional vector z and returns another p -dimensional vector with element $f(z_i)$, where f is the neuron's nonlinearity. In equation (N), $\hat{F}(w)$ is actually the sum of the squares of the error between the desired scalars and the output of the neuron:

$$\hat{F}(w) = \sum_{j=1}^p (d(j) - f(u(j)'w))^2. \quad (2)$$

If the popular gradient descent algorithm is used to solve (N), an iterative update equation is applied. In general, the gradient descent algorithm does not guarantee convergence to a global minimum due to the potential local minima entrapment. In addition, at regions of very low gradient, a gradient descent algorithm takes small orthogonal steps which result in a "zigzagging" effect of the updates and slow convergence [3,4].

III QUADRATIC PROBLEM FORMULATION

Since the Neuron Training Problem (N) is actually an unconstrained minimization problem, a variety of optimization techniques exist which may be employed to solve it. Unfortunately, due to the type of nonlinearities which are usually chosen for this problem (for example, the hyperbolic tangent function), the surface of the function $\hat{F}(w)$ is, in general, very complicated, and finding a w which minimizes the surface may be a very difficult task. It is proposed here that instead of finding a w which minimizes $\hat{F}(w)$, solve the following Problem (NQ) and use the solution of (NQ) as an answer for (N). The Neuron Quadratic Optimization Problem (NQ) is defined as follows:

$$\min_w F(w) \quad (NQ)$$

$$F(w) = (v - U'w)'(v - U'w)$$

and where v is such that $\phi(v) = d$. The function $F(w)$ can be re-written as:

$$F(w) = w'Aw - h'w + c, \quad (3)$$

where $A = UU' \in \mathbb{R}^{m \times m}$, $h = 2v'U' \in \mathbb{R}^{1 \times m}$, and $c = v'v \in \mathbb{R}^{1 \times 1}$. So, finding a w that minimizes $F(w)$ in (NQ) is equivalent to finding a w that minimizes (3).

Solving (NQ) is, in general, easier than solving (N). It is proposed here to solve (NQ) instead of (N) and to use the solution from (NQ) as an answer for (N). Intuitively speaking, if the w^* found from solving (NQ) also solves (N) with a small error, then this validates minimizing $F(w)$ instead of minimizing $\hat{F}(w)$.

In [1] and summarized in [2], both a careful analysis comparing the error for solving Problem (NQ) with the error for solving Problem (N) and the relationship between the solutions of Problems (NQ) and (N) are provided.

IV THE SINGLE-LAYER NEURAL NETWORK

The single-layer neural network considered here is comprised of n parallel neurons. With the appropriate extensions to the above, the Single Layer Neural Network Training Problem (L) is defined as follows:

$$\min_W \hat{F}(W) \quad (L)$$

$$\hat{F}(W) = \text{tr}((D - \Phi(U'W))(D - \Phi(U'W)))'$$

With $n = 1$, (L) reduces to (N). In equation (L), $\hat{F}(W)$ is actually a sum of the squares of the error between the individual desired output elements and the outputs of the neurons:

$$\hat{F}(W) = \sum_{k=1}^n \sum_{j=1}^p (d_k(j) - f(u(j)w_k))^2 \quad (4)$$

It is proposed here that instead of finding a W which minimizes $\hat{F}(W)$, solve the following Problem (LQ) and use the solution of (LQ) as an answer for (L). The **Single-Layer Neural Network Quadratic Optimization Problem (LQ)** is defined as follows:

$$\left. \begin{array}{l} \min_W F(W) \\ \text{where} \\ F(W) = \text{tr}((V - U'W)(V - U'W)) \end{array} \right\} \quad (LQ)$$

With $n = 1$, (LQ) reduces to (NQ). The function $F(W)$ can be re-written as:

$$F(W) = \sum_{i=1}^n w_i' A w_i - h_i' w_i + c_i \quad (5)$$

where $A = UU' \in \mathbb{R}^{m \times m}$, $h_i' = 2v_i'U' \in \mathbb{R}^{1 \times m}$, and $c_i = v_i'v_i \in \mathbb{R}^{1 \times 1}$ for $1 \leq i \leq n$. Thus, with A symmetric and positive definite, $F(W)$ is the sum of quadratics. The solving of (LQ) can be accomplished in many ways including either minimizing the n quadratics of (5) or finding a W that solves $UU'W = UV$. (6)

V THE MULTI-LAYER NEURAL NETWORK

The **multi-layer neural network** considered here consists of many layers of parallel neurons connected in a feedforward manner. Defining the symbol $\#k$ as the number of neurons in the k^{th} layer, the output of the k^{th} layer is described by

$$Y^k = \Phi(U^k W^k) \quad (7)$$

where $Y^k = [y_1^k, \dots, y_{\#k}^k] \in \mathbb{R}^{\#k}$ is the matrix of outputs, $y_i^k = [y_i^k(1), \dots, y_i^k(p)]' \in \mathbb{R}^{p \times 1}$ is the vector of outputs for the i^{th} neuron, $U^k = [u^k(1), \dots, u^k(p)] \in \mathbb{R}^{\#(k-1) \times p}$ is the matrix of input vectors, $u^k(j) = [y_1^{k-1}(j), \dots, y_{\#(k-1)}^{k-1}(j)]' \in \mathbb{R}^{\#(k-1) \times 1}$ is the vector of inputs for the j^{th} input pattern and is equal to the outputs from the previous layer plus the bias of one for the last term, $W^k = [w_{1,1}^k, \dots, w_{\#k, \#k}^k] \in \mathbb{R}^{\#(k-1) \times \#k}$ is the matrix of weight vectors, $w_{i,1}^k = [w_{i,1}^k, \dots, w_{i, \#(k-1)}^k]' \in \mathbb{R}^{\#(k-1) \times 1}$ is the vector of weights, and $\Phi(Z) \in \mathbb{R}^{\#k}$ is the same as defined previously for the single-layer neural network. Using $U^1 = U$, the output of the first hidden layer is described by

$$Y^1 = \Phi(U^1 W^1) \quad (8)$$

With $U^2 = [Y^1 \ 1] \in \mathbb{R}^{\#(k-1) \times 1}$ where $1 \in \mathbb{R}^{p \times 1}$, the output of the second hidden layer is described by

$$Y^2 = \Phi(U^2 W^2) \quad (9)$$

Continuing this inductive process, each successive layer is defined appropriately until the desired number of layers is reached. The last layer is called the **output layer** and is described by

$$Y^o = \Phi(U^o W^o) \quad (10)$$

where the superscript "o" denotes "output".

The **Multi-Layer Neural Network Training Problem (M)** is defined as follows:

$$\left. \begin{array}{l} \min_{W^1, \dots, W^o} \hat{F}(W^1, \dots, W^o) \\ \text{where} \\ \hat{F}(W^1, \dots, W^o) = \text{tr}((D - Y^o)(D - Y^o)) \end{array} \right\} \quad (M)$$

and where "tr" is the trace of a square matrix, (W^1, \dots, W^o) are the weight matrices of all the layers of the multi-layer neural network, $D = [d_1, \dots, d_n]' \in \mathbb{R}^{p \times n}$ is the desired output matrix, and Y^o is the output of the output layer of the multi-layer neural network. In relation to the previous training problem (L), the input matrix U is not directly in (M) since the input is "buried" beneath the hidden layers. If there are no hidden layers, then (M) reduces to (L). In equation (M), $\hat{F}(W^1, \dots, W^o)$ is actually the sum of the squares of the error between the individual desired output elements and the outputs of the neurons in the output layer:

$$\hat{F}(W^1, \dots, W^o) = \sum_{k=1}^n \sum_{j=1}^p (d_k(j) - y_k^o(j))^2 \quad (11)$$

One method to solve (M) is the back-propagation algorithm [5], which is a constant-step-size, gradient-descent algorithm that minimizes the least-squares cost function $\hat{F}(W^1, \dots, W^o)$. The sigmoid function is often considered to be the nonlinear function of the neurons (i.e., $f(z) = 1/(1 + e^{-z})$) and has the property that $f'(z) = f(z)(1 - f(z))$. Here, the hyperbolic tangent function, $f(z) = \tanh(z)$, is used as the neuron's nonlinearity, and $f'(z) = 1 - f(z)^2$. The weights of the neural network are adjusted after every **epoch** (i.e., one pass of the training set) by the constant-step-size gradient-descent rule:

$$w_{h_1}^k(t+1) = w_{h_1}^k(t) - \alpha \frac{\delta F}{\delta w_{h_1}^k} \quad (12)$$

After some manipulation of the partial derivative term, the back-propagation algorithm's rule for changing the weights of the multi-layer neural network is given by

$$w_{h_1}^k(t+1) = w_{h_1}^k(t) + \alpha \sum_{j=1}^p \delta_j^k(j) y_{h_1}^{k-1}(j) \quad (13)$$

where $\delta_j^k(j)$ is known as the **delta term**. If the k^{th} layer is the output layer, then

$$\delta_j^k(j) = [d_j(j) - y_j^o(j)] [1 - y_j^o(j)^2] \quad (14)$$

For the hidden layers,

$$\delta_j^k(j) = \sum_{r=1}^{\#(k+1)} \delta_r^{k+1}(j) w_{rj}^{k+1} [1 - y_j^k(j)^2] \quad (15)$$

VI THE QUADRATIC OPTIMIZATION ALGORITHM

In this section, the Multi-Layer Neural Network Training Problem (M) is solved by applying the solution of Problem (LQ) to each layer. The technique proposed here is based on the back-propagation algorithm, in which the propagation of the output error of the multi-layer neural network is used to form a desired output for each layer and hence to form a quadratic weight cost function at each layer. The solution of Problem (LQ) for each layer is then used as a solution for Problem (M). Several implementation considerations are discussed as well as the advantages and disadvantages of using this approach.

Instead of solving the Problem (M) using the back-propagation algorithm, it is proposed here to solve the Problem (LQ) for each layer of the multi-layer neural network and use this solution as one for (M). In solving (LQ) for the single-layer neural network, the desired output is known, and thus a matrix V can be found such that $\Phi(V) = D$. In solving (LQ) for the k^{th} layer of the multi-layer neural network, the matrix V^k needs to be found such that

$$\Phi(V^k) = D^k \quad (16)$$

where $V^k = [v_1^k, \dots, v_{\#k}^k] \in \mathbb{R}^{\#k}$, $v_i^k = [v_i^k(1), \dots, v_i^k(p)]' \in \mathbb{R}^{p \times 1}$, and $f(v_i^k(j)) = d_i^k(j)$ for $1 \leq i \leq \#k$, $1 \leq j \leq p$.

Assuming that all weights in the hidden layers have initial values, there is no problem in directly applying the methodology described for the single-layer neural network to the output layer. A matrix V^o can be chosen such that $\Phi(V^o) = D$, and Problem (LQ) can be applied to find the weights of the output layer. Unfortunately, there do not exist desired outputs D^k for the hidden layers, but by using the back-propagation of the output error, an approximation of these values can be obtained; the algorithm proposed here back-propagates the error between the desired output and the actual output of the neural network to all of the hidden layers to form an **approximated desired output** for each layer.

First, the errors at the output of each layer are defined. The error between the desired output and actual output for the i^{th} neuron of the k^{th} layer is given by

$$\hat{\epsilon}_i^k(j) = d_i^k(j) - y_i^k(j) \quad (17)$$

where $1 \leq i \leq \#k$, $1 \leq j \leq p$, and the error for the i^{th} neuron of the output layer is the quantity

$$\hat{\epsilon}_i^o(j) = d_i(j) - y_i^o(j) \quad (18)$$

where $1 \leq i \leq n$, $1 \leq j \leq p$. Next, in comparing the delta terms of (14) and (15) of the back-propagation algorithm, the error for the i^{th} neuron of the k^{th} layer (not equal to the output layer) can be viewed as

$$\hat{\epsilon}_i^k(j) = \sum_{h=1}^{\#(k+1)} \delta_h^{k+1}(j) w_{ih}^{k+1} \quad (19)$$

Combining (18) and (19) the desired output for the i^{th} neuron of the k^{th} layer can be viewed as

$$d_i^k(j) = y_i^k(j) + \sum_{h=1}^{\#(k+1)} \delta_h^{k+1}(j) w_{ih}^{k+1} \quad (20)$$

Using (20), the matrix V^k can be chosen such that

$$f(v_i^k(j)) = y_i^k(j) + \sum_{h=1}^{\#(k+1)} \delta_h^{k+1}(j) w_{ih}^{k+1} \quad (21)$$

where $1 \leq i \leq \#k$, $1 \leq j \leq p$. With V^k , Problem (LQ) can be solved to find the weights of the k^{th} layer. Since the hyperbolic tangent function is one-to-one and is assumed to be the nonlinear function of each neuron, $v_i^k(j)$ can be formed by applying the inverse of the function to both sides of (21). Thus, by back-propagating the error through the multi-layer neural network, a quadratic problem is formulated for each layer, and the results for the Problem (LQ) in relation to the Problem (L) are applicable here for each layer. This method does not guarantee convergence, but does tend to give good results with a fast computation time.

In implementing this quadratic optimization procedure for a multi-layer neural network, several observations are useful. First, in practice, limiting the neural network to two layers provides adequate results. Second, since a quadratic function is minimized for each layer, the hidden layer should be adjusted first,

and then the output layer can be updated using the newly found values for the hidden layer's weights. Third, when a one-to-one function is the hidden layer's nonlinearity and when the values $v_i^1(j)$ are found by inverting the one-to-one function, care must be taken to insure that $d_i^1(j)$ lies in the range of the function. For instance, if the hyperbolic tangent function is the nonlinearity for the hidden layer, its range is $(-1, 1)$ and hence $d_i^1(j) \in (-1, 1)$. To insure this, the output layer's weights need to be first initialized to small values around zero, for instance $w_{ih}^2 \in [-\frac{1}{\#1}, \frac{1}{\#1}]$. Next, when computing the desired output for the hidden layer, if the right-hand side of (21) is not in the range of the hidden layer's nonlinearity, the weights of the output layer can be scaled to insure this: if $\bar{w}_{ih}^2 = \max\{w_{ih}^2 \text{ for } 1 \leq i \leq \#1 + 1 \text{ and } 1 \leq h \leq \#2\} > 1/\#1$, then (21) is modified to

$$f(v_i^1(j)) = y_i^1(j) + \sum_{h=1}^{\#2} \delta_h^2(j) \frac{w_{ih}^2}{\#1 w^2} \quad (22)$$

for $1 \leq i \leq \#1 + 1$ and $1 \leq j \leq p$. If the right-hand side of (22) is still not in the range of the hidden layer's nonlinearity, the output layer's weights can continue to be scaled by $1/\#1$ until this occurs. Fourth, in practice, the weights for the output layer tend to be large in magnitude, which is attributed to the fitting of the mapping between U^2 and V^2 with the linear equation

$$U^2 U^2 W^2 = U^2 V^2. \quad (23)$$

To aid in avoiding the computational inaccuracies which may occur due to the large magnitudes of W^2 , it is suggested to choose $v_i^2(j) < 0.5$ for $1 \leq i \leq \#2$ and $1 \leq j \leq p$. This can be accomplished by scaling the desired outputs appropriately. This also aids in insuring that $d_i^1(j)$ is properly valued. With these observations, the Quadratic Optimization Algorithm used in practice to train a multi-layer neural network is as follows:

- 1) Given D , find V^2 such that $v_i^2(j) < 0.5$ for $1 \leq i \leq \#2$ and $1 \leq j \leq p$.
- 2) Initialize W^1 and W^2 with magnitudes less than $1/\#1$.
- 3) Test $\hat{F}(W^1, W^2)$. If small enough, then stop.
- 4) Find V^1 using either (21) or (22).
- 5) Find W^1 by solving $U^1 U^1 W^1 = U^1 V^1$.
- 6) Find U^2 for the new weights W^1 .
- 7) Find W^2 by solving $U^2 U^2 W^2 = U^2 V^2$.
- 8) *Go to 3).*

In applying the Quadratic Optimization Algorithm to various problems, several advantages and disadvantages are evident. First, finding the weights of a multi-layer neural network via the quadratic optimization approach described here tends to work well for classification problems in that the desired outputs are achieved and the generalization behavior of the neural network is accurate. The algorithm also tends to converge to a solution in a single step achieving a small value for $\hat{F}(W^1, W^2)$ and then to slowly vary around this value with more iterations. Thus, it is recommended to use the Quadratic Optimization Algorithm for a single iteration on classification problems.

Both of these properties are attributed to the finding of the weights via steps 5) and 7) of the algorithm. To achieve this type of performance, the choice of the number of hidden layer neurons is important. Since the overall mapping between the input patterns and the desired output patterns is accomplished via the solving of the linear system of equations in step 7), the number of hidden layer neurons needs to be large enough such that this linear approximation in the output layer succeeds. Clearly, the choice of $\#1$ is problem dependent. Thus, the choice of the number of hidden layer neurons is a design consideration and is dependent on the particular desired mapping of the training set. Furthermore, the initial values for the weights of the neural network are more important as the number of hidden layer neurons is reduced towards the level where the Quadratic Optimization Algorithm is unable to achieve the desired training set mapping. These properties of the algorithm are illustrated in Example 1.

The disadvantages of using the Quadratic Optimization Algorithm to train a multi-layer neural network are outlined next. First, the algorithm may not work well for non-classification problems in that the desired outputs may be approximately achieved but the generalization behavior of the neural network may be inaccurate. This behavior is attributed to the finding of the output layer's weights by the solving of the linear system of equations in step 7). Thus, it is suggested to restrict the use of the Quadratic Optimization Algorithm to classification problems. Secondly, in applying the algorithm, all of the training patterns need to be known and no values for the weights from previous iterations are saved. Thus, this quadratic optimization training procedure may not work for on-line learning. Thirdly, the algorithm also requires the solving of two linear systems of equations when there are two layers of weights: one in step 5) with m equations and m unknowns, and the other in step 7) with $\#1 + 1$ equations and $\#1 + 1$ unknowns. If these numbers are large, the solving of the linear systems may become burdensome, although there do exist many ways for solving such systems. Finally, the values for the output layer weights may be large, which is a potential disadvantage if implementation of the neural network is desired. This behavior is also attributed to the final calculation step of the algorithm, which attempts to form the desired mapping with the solving of a

linear system of equations to find the weights of the neural network's output layer.

The two examples in the following section illustrate these observations and some of the advantages and disadvantages of using the Quadratic Optimization Algorithm to train a multi-layer neural network.

VII EXAMPLES

Example 1

In this example, a comparison of the Quadratic Optimization Algorithm and the back-propagation algorithm for the training of a multi-layer neural network is presented for an extended XOR training set, where

$$U = \begin{bmatrix} -1 & -1 & 1 & 1 & -2 & -2 & 2 & 2 \\ -1 & 1 & -1 & 1 & -2 & 2 & -2 & 2 \end{bmatrix} \in \mathbb{R}^{3 \times 8}$$

and

$$d = [0.1 \ -0.1 \ -0.1 \ 0.1 \ 0.1 \ -0.1 \ -0.1 \ 0.1] \in \mathbb{R}^{8 \times 1}$$

Both the Quadratic Optimization Algorithm and the back-propagation algorithm are implemented in MATLAB on a Macintosh SE using programs that are not optimized. Steps 5) and 7) of the algorithm are solved using the pseudo-inverse function call for MATLAB. Using a two-layer neural network with the hyperbolic tangent function as the nonlinearity for both the hidden layer neurons and the output layer neuron, the number of hidden layer neurons is changed.

The training results for the two algorithms are compared in terms of the success of training the neural network for the desired classification, the values for the cost function \hat{F} , and the number of floating point operations, which is a function call in MATLAB. The multi-layer neural network is initialized to weights in the interval $[-1/\#1, 1/\#1]$ and first trained using the Quadratic Optimization Algorithm. Next, the same initial neural network is trained using the back-propagation algorithm until the same value for \hat{F} is achieved. (Note that at each iteration step of the back-propagation algorithm the gradient of (13) is not computed and is instead approximated using a single training pattern.) These results are shown in Table 1. The value $\hat{F}(t)$ denotes the value of $\hat{F}(W^1, W^2)$ after t iterations of the training procedure. The value "flops" indicates the number of floating point operations as counted by MATLAB for the t iterations. For the cases of 4 and 5 hidden layer neurons, the back-propagation trained neural network did not classify the input set correctly, and the training was continued until a lower \hat{F} was achieved. The resulting neural networks classified the input patterns correctly, and the results for the extended training using the back-propagation algorithm are shown in Table 2. The values for t and flops are for the total training time. As was described in the previous section, the multi-layer neural network trained with the Quadratic Optimization Algorithm requires enough hidden layer neurons such that the solution to the linear system of equations in step 7) is able to correctly approximate the desired mapping of the training set. For $\#1 = 2$ and $\#1 = 3$, the Quadratic Optimization Algorithm was unable to find values for the weights such that the desired mapping was achieved, while the back-propagation algorithm was able to successfully find such weights.

Table 1 Comparing the Quadratic Optimization Algorithm and the back-propagation algorithm.

#1	$\hat{F}(0)$	Quadratic Optimization		Back-Propagation			
		t	$\hat{F}(t)$	flops	t	$\hat{F}(t)$	flops
6	0.1063	1	0.0212	132024	1244	0.0208	1442790
5	0.1189	1	0.0244	10256	1169	0.0240	1158168
4	0.1330	1	0.0246	820	427	0.0240	351273

Table 2 Continuing training with the back-propagation algorithm.

#1	t	$\hat{F}(t)$	flops
4	1006	0.0096	476024

To illustrate some of the observations that are made at the end of Section VI, the weights and the outputs of a neural network found via the Quadratic Optimization Algorithm are presented. For the neural network with 5 hidden layer neurons, the output layer weights found using the Quadratic Optimization Algorithm for 1 iteration and those found using the back-propagation algorithm for 1729 iterations are

$$w_{QOA}^2 = [10.7469 \ 53.1767 \ 49.9710 \ 16.5357 \ 21.8128 \ 0.5822]$$

and

$$w_{BPA}^2 = [-0.3481 \ 0.3337 \ 0.0400 \ 0.0223 \ 0.0016 \ 0.1508]$$

As described previously, the weights for the output layer may be large, and they are for this case. However, in the simulations using the training sets of the other examples, these values may be several orders of magnitude larger than the ones for this extended XOR example.

As another comparison for the neural network with 5 hidden layer neurons, the outputs of the neural network found using the Quadratic Optimization

Algorithm for 1 iteration and those found using the back-propagation algorithm for 1729 iterations are

$$y_{QOA} = [0.0079 \quad -0.0256 \quad -0.0554 \quad 0.0170 \quad 0.1238 \quad -0.1056 \quad -0.0906 \quad 0.1233]$$

and

$$y_{BPA} = [0.0531 \quad -0.0802 \quad -0.0789 \quad 0.0514 \quad 0.1526 \quad -0.0809 \quad -0.0783 \quad 0.1165]$$

where the magnitude of the desired output is 0.1. Clearly, the neural network trained with the Quadratic Optimization Algorithm did not achieve the desired neural network outputs but did achieve the desired mapping for the classification training set. For this reason, it is recommended that the Quadratic Optimization Algorithm be used for classification problems and not for general function approximation problems. However, even though the exact desired outputs may not be achieved using the Quadratic Optimization Algorithm, the resulting neural network does have desirable generalization properties for classification training sets as illustrated in the following example.

Example 2

In a square of size $[0, 8] \times [0, 8]$, consider a circle of radius 2 centered at (4, 4). Let the input patterns be points inside the square. If the input pattern lies inside the circle, the corresponding desired output is 1, and if the input pattern lies outside the circle, the corresponding desired output is -1. 289 training patterns are chosen at random with a uniform distribution over the $[0, 8] \times [0, 8]$ square region and are depicted in Figure 1. Thus, $U \in \mathbb{R}^{3 \times 289}$ and $d \in \mathbb{R}^{289 \times 1}$. A two-layer neural network is provided with $\#1 = 30$. The hyperbolic tangent function is used as the nonlinearity for the hidden layer's neurons, and the signum function is used as the nonlinearity for the output layer's neurons. The neural network is trained with the Quadratic Optimization Algorithm, which is implemented in MATLAB on a Sun Sparc Station using a program that is not optimized. Steps 5) and 7) of the algorithm are solved using MATLAB's pseudo-inverse function. Since the signum function is the output layer's nonlinearity, the vector v^2 is chosen such that $v^2(j) = 0.1(d(j))$ for $1 \leq j \leq 289$. The weights for the neural network are chosen at random in the interval $[-1/30, 1/30]$ such that $\hat{F}(0) = 3.7221$. After applying the Quadratic Optimization Algorithm for one iteration requiring 14725850 floating point operations, $\hat{F}(1) = 0.3472$ and the training set is almost arbitrarily correctly classified. To test the generalization ability of the result, the neural network is probed with 1089 randomly chosen patterns. The resulting output is displayed in Figure 2. The trained neural network clearly generalizes well over the input space.

VIII CONCLUDING REMARKS

A new method based on the minimization of a quadratic function is presented for the training of a single neuron, a single-layer neural network, and a multi-layer neural network. The training procedure for the single neuron can be immediately applied to the single-layer neural network case but not to the multi-layer neural network case since there do not exist known desired outputs for the neural network's hidden layers. By using the concept of back-propagating the output error to the hidden layers, desired outputs are approximated for the hidden layers, and the results for the single-layer neural network are applied to each of the hidden layers. The training of a multi-layer neural network via the described Quadratic Optimization Algorithm tends to work best for classification problems and tends to achieve good results in a single iteration.

The results reported in this paper also appear in [1] and [2], which is an expanded version of this paper.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the partial support of the Jet Propulsion Laboratory (Contract Number 957856).

REFERENCES

- [1] Sartori M.A., *Feedforward Neural Networks and Their Application in the Higher Level Control of Systems*, Ph.D. Dissertation, Department of Electrical Engineering, University of Notre Dame, April 1991.
- [2] Sartori M.A., *Neural Network Training Via Quadratic Optimization*, Technical Report #90-05-01, Department of Electrical Engineering, University of Notre Dame, May 1990, Revised April 1991.
- [3] Gill P.E., Murray W., Wright M.H., *Practical Optimization*, Academic Press, New York, 1981.
- [4] Bazaraa M.S., Shetty C.M., *Nonlinear Programming: Theory and Applications*, Wiley, New York, 1979.
- [5] Rumelhart D.E., Hinton G.E., Williams R.J., "Learning Internal Representations by Error Propagation," in Rumelhart D.E., McClelland J.L., eds., *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, vol. 1: Foundation*, pp. 318-362, MIT Press, 1986.




Figure 1 Training set.




Figure 2 Testing the trained neural network.