



A Learning Framework for Neural Networks Using Constrained Optimization Methods

STAVROS J. PERANTONIS*, NIKOLAOS AMPAZIS and VASSILIS VIRVILIS sper@iit.demokritos.gr
*Institute of Informatics and Telecommunications, National Center for Scientific Research “Demokritos”,
GR-153 10 Agia Paraskevi, Greece*

Abstract. Conventional supervised learning in neural networks is carried out by performing unconstrained minimization of a suitably defined cost function. This approach has certain drawbacks, which can be overcome by incorporating additional knowledge in the training formalism. In this paper, two types of such additional knowledge are examined: Network specific knowledge (associated with the neural network irrespectively of the problem whose solution is sought) or problem specific knowledge (which helps to solve a specific learning task). A constrained optimization framework is introduced for incorporating these types of knowledge into the learning formalism. We present three examples of improvement in the learning behaviour of neural networks using additional knowledge in the context of our constrained optimization framework. The two network specific examples are designed to improve convergence and learning speed in the broad class of feedforward networks, while the third problem specific example is related to the efficient factorization of 2-D polynomials using suitably constructed sigma-pi networks.

Keywords: neural networks, supervised learning, constrained optimization

1. Introduction

Methods from the field of optimization have played an important role in developing training algorithms for connectionist systems. Indeed, the realization that simple gradient descent (back propagation – BP) can be applied with success to the training of multilayered feedforward networks (MFNs) [19] was responsible to a great extent for the resurgence of interest in this kind of networks during the mid 1980s. Most of the methods used for supervised learning originate from unconstrained optimization techniques. Obviously, this is related to the “black box” nature of connectionist systems: apart from the minimization of a cost function, no other information or knowledge is usually taken into account.

Nevertheless, recent research has shown that it is often beneficial to incorporate additional knowledge in the neural network architecture or learning rule [2,5,7,14,20]. Often, the additional knowledge can be encoded in the form of mathematical relations that have to be satisfied in addition to the demand for minimization of the cost function. A number of learning algorithms have been proposed that incorporate “target conditions”, i.e., additional conditions to be satisfied upon termination of learning. Notable

* Corresponding author.

examples include matching of hidden unit outputs with prespecified targets [8,18], satisfaction of box constraints for the synaptic weights that prevent hidden node saturation [21] or conditions that enable weight decay for the improvement of generalization properties [12,22]. Naturally, methods from the field of constrained optimization are essential for solving these modified neural network learning tasks. In the literature, the most common method for handling extra constraints is the penalty function formulation [22]. Recently, the application of interior point methods has also been investigated [21].

Apart from conditions that must be satisfied as best as possible upon termination of learning (target conditions), experience in training of feedforward networks shows that it is also useful to incorporate additional “incremental conditions”, i.e., conditions involving quantities that must be optimized incrementally at each epoch of the learning process. In this paper, we introduce a general constrained optimization framework for incorporating additional knowledge in the neural network learning rule. We thus formulate a general problem, whereby minimization of the cost function representing the distance of the networks outputs from preset target values is sought, subject to other relations that represent the additional knowledge. We show how to formulate the general problem of incorporating additional knowledge as a multiobjective optimization task, whose solution leads to a generic learning algorithm accounting for both target and incremental conditions.

To elucidate the usefulness of the approach, we present three examples of learning algorithms that incorporate additional knowledge about learning and show that they can be derived using the general framework introduced earlier.

In the first two examples, additional information about the specific type of the neural network, the nature and characteristics of its cost function landscape is used to facilitate learning in broad classes of problems. In the first example, a successful algorithm originally proposed in [14] is derived using our framework for constrained learning. Here the additional information is incorporated in one extra condition that seeks to facilitate navigation in long narrow valleys of the cost function landscape, thus accelerating learning. In the second example, the objective is to accelerate learning and improve convergence through successful negotiation of the cost function landscape in the vicinity of temporary minima. In previous related work this problem was studied for networks of two hidden nodes using a dynamical systems approach [1]. In this paper, the dynamical systems approach is extended to networks with a hidden layer of an arbitrary number of nodes. This extension provides us with information about the nature of temporary minima, whose successful negotiation can be achieved using the general constrained learning framework with multiple additional objectives.

In the final example, the objective is to solve a specific problem (polynomial factorization) using an appropriate neural network formulation. In previous work, we have used neural networks to solve this problem [13] taking into account conditions among the polynomial coefficients. In this paper we show that within the general constrained learning framework studied in this paper we can also derive a learning algorithm for solving the polynomial factorization problem effectively.

For all examples, we present simulations to compare the performance of the proposed constrained optimization learning method with other state of the art algorithms for multilayered feedforward networks.

2. A generic constrained learning algorithm

Conventional unconstrained supervised learning in neural networks involves minimization, with respect to the synaptic weights and biases, of a cost function of the form

$$E[d_i^{(p)} - y_i^{(p)}(\mathbf{w})]. \quad (1)$$

Here \mathbf{w} is a column vector containing all synaptic weights and biases of the network, p is an index running over the \mathcal{P} patterns of the training set, $y_i^{(p)}$ is the network output corresponding to output node i and $d_i^{(p)}$ is the corresponding target.

In this paper we suppose that there are additional relations to be satisfied, that represent the extra knowledge and involve the network's synaptic weights. Before introducing the form of the extra relations, we note that in order to minimize the cost function of equation (1) we adopt an epoch-by-epoch optimization framework with the following objectives:

- At each epoch of the learning process, the vector \mathbf{w} is to be incremented by $d\mathbf{w}$, so that the search for an optimum new point in the space of \mathbf{w} is restricted to a hypersphere of known radius δP centered at the point defined by the current \mathbf{w} :

$$d\mathbf{w}^T d\mathbf{w} = (\delta P)^2. \quad (2)$$

- At each epoch, the cost function must be decremented by a positive quantity δQ , so that at the end of learning E is rendered as small as possible. To first order, we can substitute the change in E by its first differential and demand that:

$$dE = \delta Q. \quad (3)$$

We now wish to introduce additional objectives in order to incorporate the extra knowledge into the learning formalism. We consider the following two cases of importance that involve additional mathematical relations representing knowledge about learning in neural networks:

Case 1. There are additional constraints $\Phi = \mathbf{0}$ which must be satisfied as best as possible upon termination of the learning process. Here Φ is a column vector whose components are known functions of the synaptic weights. We address this problem by introducing a function κ and demanding at each epoch of the learning process the maximization of κ subject to the condition that $d\Phi = -\kappa\Phi$. In this way, we ensure that Φ tends to $\mathbf{0}$ at a temporarily exponential rate.

Based on these considerations, the following optimization problem can be formulated for each epoch of the algorithm, whose solution will determine the adaptation rule for \mathbf{w} :

$$\text{Maximize}_{\mathbf{dw} \in R} \kappa \quad (4)$$

$$\text{s.t. } \mathbf{dw}^T \mathbf{dw} = (\delta P)^2, \quad (5)$$

$$dE = \delta Q, \quad (6)$$

$$d\Phi_m / \Phi_m = -\kappa, \quad m = 1, \dots, S, \quad (7)$$

where S is the number of components of Φ .

Case 2. This case involves additional conditions whereby there is no specific final target for the vector Φ , but rather it is desired that all components of Φ are rendered as large as possible at each individual epoch of the learning process. This is a multiobjective maximization problem, which we address by defining $\kappa = d\Phi_m$ and demanding that κ assume the maximum possible value at each epoch. Thus the constrained optimization problem is as before with equation (7) substituted by

$$\kappa = d\Phi_m, \quad m = 1, \dots, S. \quad (8)$$

The solution to the above constrained optimization problems can be obtained by a method similar to the constrained gradient ascent technique introduced by Bryson and Denham [4] and leads to a generic update rule for \mathbf{w} . We thus introduce suitable Lagrange multipliers L_1 and L_2 to take account of equations (6) and (5), respectively, and a vector of multipliers Λ to take account of equation (7) or equation (8). The Lagrangian thus reads

$$\kappa' = \kappa + L_1(\mathbf{G}^T \mathbf{dw} - \delta Q) + L_2[\mathbf{dw}^T \mathbf{dw} - (\delta P)^2] + \Lambda^T(\mathbf{F} \mathbf{dw} + \kappa \mathbf{1}), \quad (9)$$

where the following quantities have been introduced:

- A vector \mathbf{G} with elements $G_i = \partial E / \partial w_i$;
- A matrix \mathbf{F} whose elements are defined by $F_{mi} = (1/\Phi_m)(\partial \Phi_m / \partial w_i)$ (for case 1) or $F_{mi} = -\partial \Phi_m / \partial w_i$ (for case 2).

To maximize κ under the required constraints, we demand that

$$d\kappa' = d\kappa(1 + \Lambda^T \mathbf{1}) + (L_1 \mathbf{G}^T + \Lambda^T \mathbf{F} + 2L_2 \mathbf{dw}^T) d^2 \mathbf{w} = 0, \quad (10)$$

$$d^2 \kappa' = 2L_2 d^2 \mathbf{w}^T d^2 \mathbf{w} < 0. \quad (11)$$

Hence, the factors multiplying $d^2 \mathbf{w}$ and $d\kappa$ in equation (9) should vanish, and therefore we obtain

$$\mathbf{dw} = -\frac{L_1}{2L_2} \mathbf{G} - \frac{1}{2L_2} \mathbf{F}^T \Lambda, \quad (12)$$

$$\Lambda^T \mathbf{1} = -1. \quad (13)$$

Equation (12) constitutes the weight update rule for the neural network, provided that the Lagrange multipliers appearing in it have been evaluated in terms of known quantities. The result of this evaluation is summarized forthwith, whereas the full evaluation is carried out in the appendix. To complete the evaluation, it is necessary to introduce the following quantities:

$$\mathbf{I}_{GG} = \mathbf{G}^T \mathbf{G}, \quad (14)$$

$$\mathbf{I}_{GF} = \mathbf{F} \mathbf{G}, \quad (15)$$

$$\mathbf{I}_{FF} = \mathbf{F} \mathbf{F}^T, \quad (16)$$

$$\mathbf{R} = \frac{1}{I_{GG}} (\mathbf{I}_{GG} \mathbf{I}_{FF} - \mathbf{I}_{GF} \times \mathbf{I}_{GF}), \quad (17)$$

$$Z = 1 + \frac{(\mathbf{R}^{-1})_a (\mathbf{I}_{GF}^T \mathbf{R}^{-1} \mathbf{I}_{GF}) - (\mathbf{1}^T \mathbf{R}^{-1} \mathbf{I}_{GF})^2}{(\mathbf{R}^{-1})_a I_{GG}}, \quad (18)$$

where $(\cdot)_a$ denotes the sum of all elements of a matrix and $\mathbf{1}$ is a column vector whose elements are all equal to 1. In terms of these known quantities, the Lagrange multipliers are evaluated using the relations

$$L_2 = -\frac{1}{2} \left[\frac{I_{GG}}{(\mathbf{R}^{-1})_a [I_{GG} (\delta P)^2 - Z (\delta Q)^2]} \right]^{1/2}, \quad (19)$$

$$\mathbf{\Lambda} = -\frac{\mathbf{R}^{-1} \mathbf{1}}{(\mathbf{R}^{-1})_a} - \frac{2L_2 \delta Q}{I_{GG}} \left[\frac{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{I}_{GF}}{(\mathbf{R}^{-1})_a} \mathbf{R}^{-1} \mathbf{1} - \mathbf{R}^{-1} \mathbf{I}_{GF} \right], \quad (20)$$

$$L_1 = -\frac{2L_2 \delta Q + \mathbf{\Lambda}^T \mathbf{I}_{GF}}{I_{GG}}. \quad (21)$$

In the appendix, it is shown that δQ must be chosen adaptively according to the relation $\delta Q = -\xi \delta P \sqrt{I_{GG}}$. Here ξ is a real parameter with $0 < \xi < 1$. Consequently, the proposed generic weight update algorithm has two free parameters, namely δP and ξ .

3. Learning in MFNs: Negotiating long valleys

Consider a feedforward network with two layers of weights which has N external input signals with the addition of a bias input. The bias signal is identical for all neurons in the network. The hidden layer consists of M neurons and the output layer contains K neurons with sigmoid activation functions $f(s) = 1/(1 + \exp(-s))$. For a set of \mathcal{P} training patterns indexed by p , the off-line BP algorithm is obtained by performing gradient descent with respect to the mean square error (MSE) cost function

$$E = \frac{1}{2} \sum_{p=1}^{\mathcal{P}} \sum_{i=1}^K (d_i^{(p)} - y_i^{(p)})^2, \quad (22)$$

where $y_i^{(p)}$ denote the output activations and $d_i^{(p)}$ are the desired responses.

Learning in feedforward networks is usually hindered by specific characteristics of the landscape defined by the MSE cost function. The two most common problems arise because

- of the occurrence of long, deep valleys or troughs that force gradient descent to follow zig-zag paths;
- of the possible existence of temporary minima in the cost function landscape.

In cost function landscapes with long deep valleys, back propagation (gradient descent) is highly inefficient because it settles into zig-zag paths and is hopelessly slow [16]. Supplementing gradient descent with a momentum term proportional to the weight update vector of the previous epoch represents a compromise between the need to decrease the cost function and the need to proceed along relatively smooth paths in the weight space. By introducing the momentum term which is proportional to the previous epoch weight update, paths whereby current and previous weight update vectors are partially aligned are favored, smoother trajectories are followed and learning is accelerated. An illustration of this improved behaviour is given in [9, p. 123]. However, the selection of appropriate values for the learning rate and the momentum term coefficient is difficult and a method of adaptively determining these coefficients based on landscape characteristics is highly desirable. To this end, an iterative algorithm was introduced in [14] whose purpose is to maximize, at each epoch, the alignment between the current and previous weight update vectors without compromising the need for decrease of the cost function. This helps achieve the maximum possible alignment of successive weight update vectors, thus further suppressing zig-zagging and accelerating learning. To this end, we require satisfaction of an additional condition, amounting to maximization of the quantity $\Phi = (\mathbf{w} - \mathbf{w}_t)^T(\mathbf{w}_t - \mathbf{w}_{t-1})$ with respect to the synaptic weight vector \mathbf{w} at each epoch of the algorithm. Here \mathbf{w}_t and \mathbf{w}_{t-1} are the values of the weight vectors at the present and immediately preceding epoch respectively and are treated as known constant vectors. Since within our constrained learning framework $\mathbf{w} - \mathbf{w}_t$ and $\mathbf{w}_t - \mathbf{w}_{t-1}$ have constant moduli equal to δP (by equation (5)), maximization of Φ amounts to minimization of the angle between successive weight update vectors. Once the quantity to be maximized at each epoch has been specified, the learning rule can be derived readily from the generic constrained learning algorithm presented in the previous section. We have only one additional condition to satisfy (maximization of $d\Phi$ at each epoch), so that case 2 of the previous section is applicable. It is readily seen that in this case Λ has only one component equal to -1 (by equation (13)) and the weight update rule is quite simple:

$$d\mathbf{w} = -\frac{L_1}{2L_2}\mathbf{G} - \frac{1}{2L_2}\mathbf{u}, \quad (23)$$

where

$$\mathbf{u} = \mathbf{w}_t - \mathbf{w}_{t-1}, \quad L_1 = (I_{uG} - 2L_2\delta Q)/I_{GG}, \quad L_2 = \frac{1}{2} \left[\frac{I_{GG}I_{uu} - I_{uG}^2}{I_{GG}(\delta P)^2 - (\delta Q)^2} \right]^{1/2} \quad (24)$$

with

$$I_{uG} = \mathbf{u}^T \mathbf{G}, \quad I_{uu} = \mathbf{u}^T \mathbf{u}. \quad (25)$$

This is identical to the learning rule derived in [14]. Hence, weight updates are formed as linear combinations of the cost function derivatives \mathbf{G} with respect to the weights and of the weight updates \mathbf{u} at the immediately preceding epoch. This is similar to back propagation with a momentum term, with the essential difference that the coefficients of \mathbf{G} and \mathbf{u} are suitably adapted at each epoch of the learning process.

The resulting learning algorithm is herewith compared with BP and a host of other well known algorithms that have been used by other authors to train feedforward networks. Simulations quoted here and in the next section were conducted using Billnet, a locally developed neural network simulator (available at <http://www.iit.demokritos.gr/~vasvir+>). We report performance results for two well known benchmarks. In particular, results for the 11-11-1 multiplexer problem with 2048 patterns [10] are shown in table 1, whereas results for the 64-8-64 encoder problem with 64 patterns [6] are given in table 2. We assess the performance of a much wider range of algorithms than those tested in [14], namely the on-line and off-line versions of BP with momentum (BPM), resilient propagation (RPROP) [17], conjugate gradient methods of Fletcher-Reeves (CG/FR) and Polak-Ribière (CG/PR) with restarts [11], the quickprop algorithm of Fahlman (QP) [6] and the Delta Bar Delta algorithm of Jacobs (DBD) [10]. For each benchmark problem we performed 50 learning trials starting from different randomly chosen weights in the range -0.5 to 0.5 . The maximum number of epochs per trial was set to 1000 and learning was considered successful when Fahlman's "40-20-40" criterion was met [6]. Learning parameters were chosen to ensure the best possible performance for each algorithm. Symbols used for the parameters in tables 1 and 2 are the same used by the corresponding authors. The average and standard deviation of the number of epochs in the 50 trials is quoted, along with the success rate, i.e., the percentage of trials for which learning was successful according to Fahlman's criterion. Note that the

Table 1

Results in terms of number of epochs (mean value and standard deviation) and success rates for the 11-11-1 multiplexer task using the constrained learning algorithm of section 3 and other well known algorithms. F and NC denote failure of convergence in all trials.

	Proposed (sect. 3)	BPM off	BPM on	RPROP	CG (PR)	CG (FR)	QP	DBD
	$\delta P : 0.8$ $\xi : 0.5$	NC	$\eta : 0.6$ $\alpha : 0.7$	$\eta^+ : 1.2$ $\eta^- : 0.5$ $\Delta_M : 1.0$ $\Delta_m : 10^{-6}$ $\Delta_0 : 0.1$			NC	NC
Epochs (Mean)	145	F	287	284	148	F	F	F
Epochs (Std)	38	F	95	79	35	F	F	F
Successes (%)	100	0	94	90	70	0	0	0

Table 2

Results in terms of number of epochs (mean value and standard deviation) and success rates for the 64-6-64 encoder benchmark learning task using the constrained learning algorithm of section 3 and other well known algorithms.

	Proposed (sect. 3)	BPM off	BPM on	RPROP	CG (PR)	CG (FR)	QP	DBD
	$\delta P : 1.5$ $\xi : 0.5$	NC	$\eta : 0.5$ $\alpha : 0.7$	NC			$\varepsilon : 4.0$ $\mu : 1.75$ $\omega : -10^{-4}$ $\alpha : 0.0$	NC
Epochs (Mean)	161	F	285	F	170	F	410	F
Epochs (StD)	15	F	40	F	30	F	150	F
Successes (%)	100	0	100	0	100	0	98	0

proposed algorithm is the only method fully successful in all trials for both benchmarks. Moreover, it exhibits the lowest average number of epochs in the successful trials, comparable only to the average number of epochs achieved by CG/PR. In all methods, most computational burden lies in the forward pass (for all patterns) and in the backward pass through the MFN (both of $O(\mathcal{P}N_w)$, where N_w is the total number of weights and biases). For the multiplexer and encoder benchmarks considered here, CG utilized an average of 35 forward passes through the network per epoch during the line minimization phase. This is to be added to the cost of one forward and one backward pass needed in the gradient evaluation phase. On the other hand, all other algorithms examined here, including the proposed algorithm, just involve one forward and one backward pass per epoch.

4. Dynamics of learning near temporary minima

As mentioned in the previous section, the problem of slow learning in feedforward networks is also associated with the problem of temporary minima in the cost function landscape. In recent work [1], we have studied the problem of temporary minima in feedforward networks with just two hidden nodes using a method that originates from the theory of dynamical systems. One of the major results obtained are the analytical predictions for the characteristic dynamical transitions from flat plateaus (or temporary minima) of finite error to the desired levels of lower or even zero error. Our results have taken the form of closed dynamical laws for a finite set of characteristic observables describing the dynamical process of learning and the various transitions involved for MFNs. In this section we provide a sketchy derivation of the generalization of the work in [1] to networks with an arbitrary number of nodes in the hidden layer. A detailed derivation is out of the scope of this work and will be provided in a forthcoming paper. Our objective here is to show how the results of this analysis can be incorporated into the general constrained learning framework of section 2 in the form of multiple incremental constraints so as to improve the learning behaviour of feedforward networks.

Consider an MFN with one layer of hidden nodes. Given a small learning rate, the difference weight update equations of the off-line BP algorithm can be approximated by differential equations in time. The weight update rule for the hidden-to-output connections gives

$$\dot{w}_{ij} = \sum_p (d_i^{(p)} - y_i^{(p)}) y_i^{(p)} (1 - y_i^{(p)}) y_j^{(p)}. \quad (26)$$

For each hidden node j , the corresponding rule for the input-to-hidden connections reads as follows:

$$\dot{\mathbf{w}}_j = \sum_p \sum_i (d_i^{(p)} - y_i^{(p)}) y_i^{(p)} (1 - y_i^{(p)}) w_{ij} y_j^{(p)} (1 - y_j^{(p)}) \mathbf{x}^{(p)}. \quad (27)$$

In the above equations w_{ij} represents the weight connections between each hidden node j and output node i . We have also used the compact notation

$$\mathbf{w}_j = (w_{j0} \dots w_{jN})^T \quad (28)$$

with w_{jk} representing the weight connection between hidden node j and input node k . Finally, the input pattern vector corresponding to pattern p is denoted by $\mathbf{x}^{(p)}$ and the output of a hidden node j corresponding to a pattern p is denoted by $y_j^{(p)}$ ($j = 0$ for the bias signal).

It is well known that temporary minima result from the development of internal symmetries and from the subsequent building of redundancy in the hidden layer. In this case, one or more of the hidden nodes perform approximately the same function and therefore form clusters of redundant nodes. Due to the formation of these clusters, the network is trapped in a temporary minimum and it usually takes a very long time before the redundancy is broken and it finds its way down the cost function landscape.

Based on this observation, we can introduce suitable state variables formed by appropriate linear combinations of the synaptic weights, and derive a dynamical system model which describes the dynamics of the feedforward network in the vicinity of these temporary minima. We expect that appropriate state variables that map the temporary minimum to the origin of the phase plane are formed using the weights connected to hidden nodes j belonging to a particular cluster C . In particular, we consider the following state variables:

- the differences of each weight vector \mathbf{w}_j from the average of all \mathbf{w}_j ($j \in C$);
- the differences of each weight w_{ij} from the average of all weights w_{ij} ($j \in C$).

Hence we define

$$\boldsymbol{\omega}^c = \frac{\sum_{j \in C} \mathbf{w}_j}{M_c}, \quad (29)$$

$$\boldsymbol{\varepsilon}_j = \mathbf{w}_j - \boldsymbol{\omega}^c, \quad j \in C. \quad (30)$$

Similarly for each output node i , we define

$$v_i^c = \frac{\sum_{j \in C} w_{ij}}{M_c}, \quad (31)$$

$$\mu_{ij} = w_{ij} - v_i^c, \quad j \in C. \quad (32)$$

Using the BP weight update equations (26) and (27) and keeping only first order terms in the state variables μ_{ij} and ε_j , we can obtain from equations (32) and (30) the following linear differential equations that describe the dynamics of the system:

$$\dot{\mu}_{ij} = \sum_p (d_i^{(p)} - y_i^{0(p)}) y_i^{0(p)} (1 - y_i^{0(p)}) f'(\boldsymbol{\omega}^c \cdot \mathbf{x}^{(p)}) (\boldsymbol{\varepsilon}_j \cdot \mathbf{x}^{(p)}), \quad (33)$$

$$\begin{aligned} \dot{\boldsymbol{\varepsilon}}_j = & \sum_{pi} (d_i^{(p)} - y_i^{0(p)}) y_i^{0(p)} (1 - y_i^{0(p)}) f'(\boldsymbol{\omega}^c \cdot \mathbf{x}^{(p)}) \\ & \times \mathbf{x}^{(p)} \{ \mu_{ij} + v_i [1 - 2f(\boldsymbol{\omega}^c \cdot \mathbf{x}^{(p)})] (\boldsymbol{\varepsilon}_j \cdot \mathbf{x}^{(p)}) \}, \end{aligned} \quad (34)$$

where $f' = f(1 - f)$ and

$$y_i^{0(p)} = f \left[M_c v_i^c f(\boldsymbol{\omega}^c \cdot \mathbf{x}^{(p)}) + \sum_{J \notin C} w_{iJ} y_J^{(p)} \right]. \quad (35)$$

From equations (34) and (33) we observe that time derivatives of state variables corresponding to a certain hidden node j depend only on state variables corresponding to the same hidden node. Therefore, introducing the vector $\mathbf{u}_j = (\boldsymbol{\varepsilon}_j^T, \mu_{ij})^T$, $i = 1, \dots, K$ we obtain M_c equations of the form

$$\dot{\mathbf{u}}_j = \mathbf{J}_c \mathbf{u}_j. \quad (36)$$

The Jacobian matrix \mathbf{J}_c is given by

$$\mathbf{J}_c = \sum_p f'(\boldsymbol{\omega}^c \cdot \mathbf{x}^{(p)}) \begin{pmatrix} \sum_i A_i^{(p)} [1 - 2f(\boldsymbol{\omega}^c \cdot \mathbf{x}^{(p)})] \mathbf{x}^{(p)} \mathbf{x}^{(p)T} v^c & \mathbf{A}^{(p)} \mathbf{x}^{(p)} \\ (\mathbf{A}^{(p)} \mathbf{x}^{(p)})^T & \mathbf{0} \end{pmatrix}, \quad (37)$$

where

$$\mathbf{A}^{(p)} = ((d_i^{(p)} - y_i^{0(p)}) y_i^{0(p)} (1 - y_i^{0(p)})), \quad i = 1, \dots, K). \quad (38)$$

Due to the exponential nature of the solutions of equation (36), the followed trajectory obviously depends on the magnitude of the largest eigenvalue for each Jacobian matrix. As time passes, the magnitude of the largest eigenvalues grows, so that eventually a bifurcation of the eigenvalues occurs and the system follows a trajectory which allows it to move far away from the critical point. Asymptotically, the trajectories followed by the dynamical systems described by equation (36) are then parallel to the eigenvectors corresponding to the maximum Jacobian eigenvalues. Thus, after the bifurcation of the eigenvalues has occurred, weight updates approximately follow the rule

$$d\mathbf{u}_j \approx \text{sign}(\mathbf{u}_j^T \boldsymbol{\xi}^c) \boldsymbol{\xi}^c, \quad (39)$$

where $\boldsymbol{\xi}^c$ is the eigenvector of \mathbf{J}_c corresponding to the maximum eigenvalue λ_c .

Following this analysis, it is evident that if the maximum eigenvalues $\lambda_c, c = 1, \dots, S$ of the Jacobian matrices \mathbf{J}_c corresponding to each of the S clusters of hidden nodes are relatively large, then the network is able to escape from the temporary minimum. Hence, instead of waiting for the growth of the eigenvalues, the objective of our approach is to raise these eigenvalues more rapidly in order to facilitate learning. It is therefore beneficial to incorporate in the learning rule additional knowledge related to the desire for rapid growth of these eigenvalues. Since it is difficult in the general case to express the maximum eigenvalues in closed form in terms of the weights, we choose to raise the values of appropriate lower bounds $\Phi_c \leq \lambda_c$ for these eigenvalues, which are obtained as follows: it is well known from linear algebra that since \mathbf{J}_c is a real and symmetric matrix, then

$$\mathbf{z}^T \mathbf{J}_c \mathbf{z} \leq \lambda_c \mathbf{z}^T \mathbf{z} \quad \forall \mathbf{z} \in R^Q. \quad (40)$$

We have used the simplest choice for \mathbf{z} namely $\mathbf{z} = \mathbf{1} = (1 \ 1 \dots 1)^T$ which means that the product in the left hand side of equation (40) is simply the sum of the elements of the matrix.

Therefore we can apply the generic weight update rule (case 2 of section 2) using

$$\Phi_c = \mathbf{1}^T \mathbf{J}_c \mathbf{1}. \quad (41)$$

During learning, different clusters of hidden nodes may be formed and the number and structure of the Φ_c must be changed accordingly. Therefore, for the implementation of the algorithm, it is essential to have a method of detecting cluster formation during learning. This detection is performed periodically (at regular intervals of typically 50 epochs) using a standard subtractive clustering algorithm [23] which identifies the number of formed clusters and, therefore, the number of Φ_c . Hidden nodes are then assigned to a cluster using the fuzzy C -means algorithm [3], whereupon the analytical expressions for all Jacobian matrices and the corresponding Φ_c are constructed.

The proposed method is herewith compared with BP and a host of other well known algorithms that have been used by other authors to train feedforward networks. We report performance results for two benchmarks. The first benchmark is the well known 4 bit parity problem, which is characterized by the presence of temporary minima in its cost function. A network with a 4-4-1 architecture is employed to solve the problem. Fifty learning trials were performed starting from different randomly chosen weights in the range -0.5 to 0.5 and the maximum allowed number of epochs per trial was set to 1000. As before, learning was considered successful when Fahlman's "40-20-40" criterion was met. In table 3, the proposed method is compared with a host of other well known training algorithms. For the 4 bit parity problem, note that the successful negotiation of temporary minima by the proposed algorithm leads to a much higher success rate than all other methods. Moreover, learning is achieved in a relatively low average number of epochs (in this respect the proposed method is surpassed only by CG/PR, which, however, exhibits a very low success rate equal to just 30%).

Table 3

Results in terms of number of epochs (mean value and standard deviation) and success rates for the parity-4 benchmark using the constrained learning algorithm of section 4 and other well known algorithms.

	Proposed (sect. 4)	BPM OFF	BPM ON	RPROP	CG (PR)	CG (FR)	QP	DBD
	$\delta P : 0.2$ $\xi : 0.5$	$\eta : 0.7$ $\alpha : 0.5$	$\eta : 0.7$ $\alpha : 0.9$	$\eta^+ : 1.2$ $\eta^- : 0.5$ $\Delta_M : 1.0$ $\Delta_m : 10^{-6}$ $\Delta_0 : 0.1$			$\varepsilon : 1.0$ $\mu : 1.75$ $\omega : -10^{-4}$ $\alpha : 0.0$	$\varepsilon : 5.0$ $\kappa : 0.25$ $\phi : 0.12$ $\theta : 0.7$ $\alpha : 0.8$
Epochs (Mean)	236	660	826	462	150	327	489	532
Epochs (StD)	60	125	212	103	57	98	181	93
Successes (%)	90	7	30	22	30	30	44	55

Table 4

Results in terms of number of epochs (mean value and standard deviation) and success rates for the cancer benchmark using the constrained learning algorithm of section 4 and other well known algorithms.

	Proposed (sect. 4)	BPM OFF	BPM ON	RPROP	CG (PR)	CG (FR)	QP	DBD
	$\delta P : 0.2$ $\xi : 0.5$	NC	NC	$\eta^+ : 1.2$ $\eta^- : 0.5$ $\Delta_M : 1.0$ $\Delta_m : 10^{-6}$ $\Delta_0 : 0.1$			NC	NC
Epochs (Mean)	3272	F	F	346	F	F	F	F
Epochs (StD)	605	F	F	157	F	F	F	F
Successes (%)	100	0	0	20	0	0	0	0

The second benchmark is the “cancer3” classification problem of the PROBEN1 set [15]. It concerns the diagnosis of breast cancer, the task being to classify a tumor as benign or malignant based on cell descriptions gathered by a microscope. The problem has 9 real valued inputs, 2 binary outputs and consists of 699 examples, of which 350 are included in the training set. A network with 9 inputs, 5 hidden units and 2 outputs is employed here to solve this problem. For this benchmark, 10 learning trials were performed starting from different randomly chosen weights in the range -0.5 to 0.5 . The maximum allowed number of epochs per trial was set to 5000. Comparative results of different algorithms concerning number of epochs and success rates for the “cancer3” benchmark are shown in table 4. Note that the proposed algorithm is the only method fully successful in all trials with respect to Fahlman’s criterion. Of all the other algorithms examined, only RPROP was able to solve the problem in just 20% of the trials performed.

5. Problem specific example: Polynomial factorization

Polynomial factorization is an important problem with applications in various areas of mathematics, mathematical physics and signal processing. It is a difficult problem for polynomials of more than one variable, where the fundamental theorem of algebra is not applicable.

Consider, for example, a polynomial of two variables z_1 and z_2 :

$$A(z_1, z_2) = \sum_{i=0}^{N_A} \sum_{j=0}^{N_A} a_{ij} z_1^i z_2^j \tag{42}$$

with N_A even, and $a_{00} = 1$. For the above polynomial, we seek to achieve an exact or approximate factorization of the form

$$A(z_1, z_2) \approx \prod_{i=1,2} A^{(i)}(z_1, z_2), \tag{43}$$

where

$$A^{(i)}(z_1, z_2) = \sum_{j=0}^{M_A} \sum_{k=0}^{M_A} v_{jk}^{(i)} z_1^j z_2^k \tag{44}$$

with $M_A = N_A/2$. We can try to find the coefficients $v_{jk}^{(i)}$ by considering P training patterns selected from the region $|z_1| < 1, |z_2| < 1$. The primary purpose of the learning rule is thus to minimize with respect to the $v_{jk}^{(i)}$ a cost function of the form

$$E = \sum_p \left(\prod_{i=1,2} A^{(i)}(z_{1p}, z_{2p}) - A(z_{1p}, z_{2p}) \right)^2. \tag{45}$$

Note that this cost function corresponds to a sigma-pi neural network with the elements of $\mathbf{v}^{(1)}$ and $\mathbf{v}^{(2)}$ as its synaptic weights. Unconstrained minimization of the cost function has been tried, but often leads to unsatisfactory results, because it can be easily trapped in flat minima. However, there is extra knowledge available for this problem, in the form of constraints between the coefficients of the desired factor polynomials and the coefficients of the original polynomial. More explicitly, if we assume that $A(z_1, z_2)$ is factorable, then these constraints can be expressed as follows:

$$\Phi_{j+(N_A+1)i}^v = a_{ij} - \sum_{l=1}^i \sum_{m=1}^j v_{lm}^{(1)} v_{i-l, j-m}^{(2)} = 0 \tag{46}$$

with $0 \leq i \leq N_A, 0 \leq j \leq N_A$. Thus, the objective of the adaptation process is to reach a minimum of the cost function of equation (45) with respect to the variables $v_{jk}^{(i)}$, which satisfies as best as possible the constraints $\Phi^v = \mathbf{0}$, where $\Phi^v = (\Phi_{j+(N_A+1)i}^v, 0 \leq i \leq N_A, 0 \leq j \leq N_A)$.

Table 5

Factorization results for a 2-D polynomial. The coefficient matrix \mathbf{a} for the product polynomial, as well as the coefficient matrices $\mathbf{v}^{(1)}$ and $\mathbf{v}^{(2)}$ for the two factor polynomials are shown for the exact factorization, for the result obtained by the constrained optimization method described in section 5 and for the result obtained by the conventional method of BP.

	Exact	Constrained learning	Gradient descent
Product	$\begin{bmatrix} 1 & 1.5 & 0.5 \\ 4 & 5 & 2.25 \\ 3 & 6.5 & 1 \end{bmatrix}$	$\begin{bmatrix} 1.0000 & 1.5000 & 0.4997 \\ 4.0000 & 5.0038 & 2.2528 \\ 3.0020 & 6.5022 & 1.0092 \end{bmatrix}$	$\begin{bmatrix} 1.0000 & 1.2007 & 0.3604 \\ 0.8566 & 2.2905 & 1.0663 \\ 0.1834 & 0.7607 & 0.7888 \end{bmatrix}$
1st Factor	$\begin{bmatrix} 1 & 0.5 \\ 1 & 2 \end{bmatrix}$	$\begin{bmatrix} 1.0000 & 0.4995 \\ 1.0010 & 1.9997 \end{bmatrix}$	$\begin{bmatrix} 1.0000 & 0.5996 \\ 0.4282 & 0.8874 \end{bmatrix}$
2nd Factor	$\begin{bmatrix} 1 & 1 \\ 3 & 0.5 \end{bmatrix}$	$\begin{bmatrix} 1.0000 & 1.0005 \\ 2.9990 & 0.5047 \end{bmatrix}$	$\begin{bmatrix} 1.0000 & 0.6010 \\ 0.4284 & 0.8889 \end{bmatrix}$

Here we incorporate the extra relations using our constrained optimization framework. Since the constraints have to be satisfied as best as possible upon termination of the training process, it is appropriate to utilize case 1 of section 2. It turns out that the ensuing constrained learning algorithm can determine the factor polynomials in factorable cases and gives good approximate solutions in cases where the original polynomial is non-factorable.

Table 5 shows the coefficients of a factorable polynomial and the corresponding exact factor polynomials, as well as the solutions obtained by gradient descent with momentum and the constrained learning algorithm. Obviously, only the constrained learning algorithm results compare favorably with the exact result.

6. Conclusion

There are many types of *a priori* knowledge that can be incorporated into neural networks learning, in the form of additional relations that must be satisfied by the learning rule. These constraints are usually pointed out either by the selection of the network or learning rule itself, or from the specific problem at hand which the neural network tries to learn. In this paper, we have derived a generic learning rule in which many types of additional knowledge, codified as mathematical relations satisfied by the synaptic weights, can be incorporated. We have also given specific examples of its application to neural network learning. The major benefit of this learning approach is to help relax the “black box” nature of artificial neural networks and combine the merits of both connectionist and knowledge based approaches for designing and implementing efficient information systems.

Appendix: Derivation of constrained learning algorithm

In this appendix, evaluation of the Lagrange multipliers L_1 , L_2 and Λ involved in the general constrained learning framework of section 2 is carried out.

By multiplying both sides of equation (12) by \mathbf{G}^T and by taking into account equation (6) we obtain

$$\delta Q = -\frac{L_1}{2L_2} I_{GG} - \frac{1}{2L_2} \Lambda^T I_{GF}. \quad (47)$$

Solving for L_1 readily yields equation (21), which evaluates L_1 in terms of L_2 and Λ .

By left multiplication of both sides of equation (12) by \mathbf{F} and taking into account equations (7) and (21), we obtain

$$\kappa \mathbf{1} + \frac{\delta Q I_{GF}}{I_{GG}} = \frac{\mathbf{R} \Lambda}{2L_2}, \quad (48)$$

where the matrix \mathbf{R} is defined by equation (17). Solving equation (48) for Λ yields

$$\Lambda = 2L_2 \kappa \mathbf{R}^{-1} \mathbf{1} + \frac{2L_2 \delta Q}{I_{GG}} \mathbf{R}^{-1} I_{GF}. \quad (49)$$

By substituting this equation into equation (13) we arrive at

$$\kappa = -\frac{1 + \frac{2L_2 \delta Q}{I_{GG}} \mathbf{1}^T \mathbf{R}^{-1} I_{GF}}{2L_2 (\mathbf{R}^{-1})_a}. \quad (50)$$

We can now substitute this equation into equation (49) to obtain equation (20) which evaluates Λ in terms of L_2 .

To evaluate L_2 , we must substitute our expression for $d\mathbf{w}$ into equation (5). To make the algebra easier, we note that on account of equation (21), equation (12) can be written as

$$d\mathbf{w} = \frac{\delta Q}{I_{GG}} \mathbf{G} + \frac{1}{2L_2} \mathbf{A}, \quad (51)$$

where

$$\mathbf{A} = \frac{\Lambda^T I_{GF}}{I_{GG}} \mathbf{G} - \mathbf{F}^T \Lambda. \quad (52)$$

From the definition of \mathbf{A} we can readily derive the following properties:

$$\|\mathbf{A}\|^2 = \Lambda^T \mathbf{R} \Lambda, \quad \mathbf{A}^T \mathbf{G} = 0. \quad (53)$$

Substituting equation (51) into equation (5) and taking into account equation (53), we can obtain a relation involving only L_2 and Λ :

$$L_2 = -\frac{1}{2} \left[\frac{I_{GG} (\Lambda^T \mathbf{R} \Lambda)}{I_{GG} (\delta P)^2 - (\delta Q)^2} \right]^{1/2}, \quad (54)$$

where the negative square root sign has been selected on account of inequality (11).

By substituting equation (20) into equation (54) and solving for L_2 , equation (19) is obtained, with Z given by equation (18). Evaluation of all Lagrange multipliers in terms of known quantities is now complete.

As a final note, let us discuss our choice for δQ . This choice is dictated by the demand that the quantity under the square root in equation (54) be positive. It can readily be seen by the first of equation (53) that $\mathbf{A}^T \mathbf{R} \mathbf{A} \geq 0$. Since $I_{GG} = \mathbf{G}^T \mathbf{G} \geq 0$, it follows from equation (54) that care must be taken to ensure that $I_{GG}(\delta P)^2 > (\delta Q)^2$. The simplest way to achieve this is to set $\delta Q = -\xi \delta P \sqrt{I_{GG}}$ with $0 < \xi < 1$.

References

- [1] N. Ampazis, S.J. Perantonis and J. Taylor, Dynamics of multilayer networks in the vicinity of temporary minima, *Neural Networks* 12 (1999) 43–58.
- [2] D. Barber and D. Saad, Does extra knowledge necessarily improve generalization?, *Neural Computation* 8 (1996) 202–214.
- [3] J.C. Bezdek, *Pattern Recognition with Fuzzy Objective Function Algorithms* (Plenum, New York, 1981).
- [4] A.E. Bryson and W.F. Denham, A steepest ascent method for solving optimum programming problems, *Journal App. Mech.* 29 (1962) 247–257.
- [5] Y. le Cun, L.D. Jackel, B.E. Boser, J.S. Denker, H.-P. Graf, I. Guyon, D. Henderson, R.E. Howard and W. Hubbard, Handwritten digit recognition: Applications of neural network chips and automatic learning, *IEEE Communications Magazine* (1989) pp. 41–46.
- [6] S.E. Fahlman, Faster learning variations on back-propagation: An empirical study, in: *Proceedings of the Connectionist Models Summer School*, eds. D. Touretzky et al. (Morgan Kaufmann, 1988) pp. 29–37.
- [7] S. Gold, A. Rangarajan and E. Mjolsness, Learning with preknowledge: clustering with point and graph matching distance, *Neural Computation* 8 (1996) 787–804.
- [8] T. Grossman, The CHIR algorithm for feed forward networks with binary weights, The ‘moving targets’ training algorithm, in: *Advances in Neural Information Processing Systems*, ed. D.S. Touretzky (Morgan Kaufmann, 1990) pp. 516–523.
- [9] J. Hertz, A. Krogh and R.G. Palmer, *Introduction to the Theory of Neural Computation* (Addison-Wesley, 1991).
- [10] R.A. Jacobs, Increased rates of convergence through learning rate adaptation, *Neural Networks* 1 (1988) 295–307.
- [11] E.M. Johansson, F.U. Dowla and D.M. Goodman, Backpropagation learning for multilayer feedforward networks using the conjugate gradient method, *International Journal of Neural Systems* 2 (1992) 291–301.
- [12] A. Krogh, G.I. Thorbergsson and J.A. Hertz, A cost function for internal representations, in: *Advances in Neural Information Processing Systems*, ed. D.S. Touretzky (Morgan Kaufmann, 1990) pp. 733–740.
- [13] S.J. Perantonis, N. Ampazis, S. Varoufakis and G. Antoniou, Constrained learning in neural networks: Application to stable factorization of 2-D polynomials, *Neural Proc. Lett.* 7 (1998) 5–14.
- [14] S.J. Perantonis and D.A. Karras, An efficient learning algorithm with momentum acceleration, *Neural Networks* 8 (1995) 237–249.
- [15] L. Prechelt, PROBEN1-A set of neural network benchmark problems and benchmarking rules, Technical Report 21/94, Universität Karlsruhe, Germany (1994).
- [16] S.S. Rao, *Optimization Theory and Applications* (Wiley Eastern, New Delhi, 1984).

- [17] M. Riedmiller and H. Braun, A direct adaptive method for faster backpropagation learning: The RPROP algorithm, in: *Proceedings of the International Conference on Neural Networks*, San Francisco 1 (1993) pp. 586–591.
- [18] R. Rohwer, The ‘moving targets’ training algorithm, in: *Advances in Neural Information Processing Systems*, ed. D.S. Touretzky (Morgan Kaufmann, 1990) pp. 558–565.
- [19] D.E. Rumelhart, J.E. Hinton and R.J. Williams, Learning internal representations by error propagation, in: *Parallel Distributed Processing: Explorations in the Microstructures of Cognition*, Vol. 1, *Foundations*, eds. D.E. Rumelhart and J.L. McClelland (MIT Press, 1986) pp. 318–362.
- [20] P. Simard, Y. le Cun and J. Denker, Efficient pattern recognition using a new transformation distance, in: *Advances in Neural Processing Systems*, eds. S.J. Hanson et al. (Morgan Kaufmann, 1993) pp. 50–58.
- [21] T.B. Trafalis and N.P. Couellan, Neural network training via an affine scaling quadratic optimization algorithm, *Neural Networks* 9 (1996) 475–481.
- [22] A.S. Weigend, D.E. Rumelhart and B.A. Huberman, Generalization by weight elimination with application to forecasting, in: *Advances in Neural Information Processing Systems*, ed. D.S. Touretzky (Morgan Kaufmann, 1991) pp. 875–882.
- [23] R. Yager and D. Filev, Generation of fuzzy rules by mountain clustering, *Journal of Intelligent and Fuzzy Systems* 2 (1994) 209–219.