# Second Order Properties of Error Surfaces : Learning Time and Generalization 

Yann Le Cun<br>AT\&T Bell Laboratories<br>Crawfords Corner Rd.

Holmdel, NJ 07733, USA Ramat Gan, 52100 Israel

Ido Kanter<br>Department of Physics<br>Bar Ilan University

Sara A. Solla<br>AT\&T Bell Laboratories<br>Crawfords Corner Rd.<br>Holmdel, NJ 07733, USA


#### Abstract

The learning time of a simple neural network model is obtained through an analytic computation of the eigenvalue spectrum for the Hessian matrix, which describes the second order properties of the cost function in the space of coupling coefficients. The form of the eigenvalue distribution suggests new techniques for accelerating the learning process, and provides a theoretical justification for the choice of centered versus biased state variables.


## 1 INTRODUCTION

Consider the class of learning algorithms which explore a space $\{W\}$ of possible couplings looking for optimal values $W^{*}$ for which a cost function $E(W)$ is minimal. The dynamical properties of searches based on gradient descent are controlled by the second order properties of the $\mathrm{E}(\mathrm{W})$ surface. An analytic investigation of such properties provides a characterization of the time scales involved in the relaxation to the solution $W^{*}$.
The discussion focuses on layered networks with no feedback, a class of architectures remarkably successful at perceptual tasks such as speech and image recognition. We derive rigorous results for the learning time of a single linear unit, and discuss their generalization to multi-layer nonlinear networks. Causes for the slowest time constants are identified, and specific prescriptions to eliminate their effect result in practical methods to accelerate convergence.

## 2 LEARNING BY GRADIENT DESCENT

Multi-layer networks are composed of model neurons interconnected through a feedforward graph. The state $x_{i}$ of the $i$-th neuron is computed from the states $\left\{x_{j}\right\}$ of the set $S_{i}$ of neurons that feed into it through the total input (or induced local field) $a_{i}=\sum_{j \in S_{i}} w_{i j} x_{j}$. The coefficient $w_{i j}$ of the linear combination is the coupling from neuron $j$ to neuron $i$. The local field $a_{i}$ determines the state $x_{i}$ through a nonlinear differentiable function $f$ called the activation function : $x_{i}=f\left(a_{i}\right)$. The activation function is often chosen to be the hyperbolic tangent or a similar sigmoid function.
The connection graph of multi-layer networks has no feedback loops, and the stable state is computed by propagating state information from the input units (which receive no input from other units) to the output units (which propagate no information to other units). The initialization of the state of the input units through an input vector $X$ results in an output vector $\mathcal{O}$ describing the state of the output units. The network thus implements an input-output map, $\mathcal{O}=\mathcal{O}(X, W)$, which depends on the values assigned to the vector $W$ of synaptic couplings.

The learning process is formulated as a search in the space $W$, so as to find an optimal configuration $W^{*}$ which minimizes a function $E(W)$. Given a training set of $p$ input vectors $X^{\mu}$ and their desired outputs $D^{\mu}, 1 \leq \mu \leq p$, the cost function

$$
\begin{equation*}
E(W)=\frac{1}{2 p} \sum_{\mu=1}^{p}\left\|D^{\mu}-\mathcal{O}\left(X^{\mu}, W\right)\right\|^{2} \tag{2.1}
\end{equation*}
$$

measures the discrepancy between the actual behavior of the system and the desired behavior. The minimization of $E$ with respect to $W$ is usually performed through iterative updates using some form of gradient descent:

$$
\begin{equation*}
W(k+1)=W(k)-\eta \nabla E \tag{2.2}
\end{equation*}
$$

where $\eta$ is used to adjust the size of the updating step, and $\nabla E$ is an estimate of the gradient of $E$ with respect to $W$. The commonly used Back-Propagation algorithm popularized by (Rumelhart, Hinton, and Williams, 1986), provides an efficient way of estimating $\nabla E$ for multi-layer networks.
The dynamical behavior of learning algorithms based on the minimization of $E(W)$ through gradient descent is controlled by the properties of the $E(W)$ surface. The goal of this work is to gain better understanding of the structure of this surface through an investigation of its second derivatives, as contained in the Hessian matrix H.

## 3 SECOND ORDER PROPERTIES

We now consider a simple model which can be investigated analytically: an N dimensional input vector feeding onto a single output unit with a linear activation function $f(a)=a$. The output corresponding to input $X^{\mu}$ is given by

$$
\begin{equation*}
\mathcal{O}^{\mu}=\sum_{i=1}^{N} w_{i} x_{i}^{\mu}=W^{T} X^{\mu} \tag{3.1}
\end{equation*}
$$

where $x_{i}^{\mu}$ is the $i$-th component of the $\mu$-th input vector, and $w_{i}$ is the coupling from the $i$-th input unit to the output.

The rule for weight updates

$$
\begin{equation*}
W(k+1)=W(k)-\frac{\eta}{p} \sum_{\mu=1}^{p}\left(\mathcal{O}^{\mu}-d^{\mu}\right) X^{\mu} \tag{3.2}
\end{equation*}
$$

follows from the gradient of the cost function

$$
\begin{equation*}
E(W)=\frac{1}{2 p} \sum_{\mu=1}^{p}\left(d^{\mu}-\mathcal{O}^{\mu}\right)^{2}=\frac{1}{2 p} \sum_{\mu=1}^{p}\left(d^{\mu}-W^{T} X^{\mu}\right)^{2} \tag{3.3}
\end{equation*}
$$

Note that the cost function of Eq. (3.3) is quadratic in $W$, and can be rewritten as

$$
\begin{equation*}
E(W)=\frac{1}{2}\left(W^{T} \mathbf{R} W-2 Q^{T} W+c\right) \tag{3.4}
\end{equation*}
$$

where $\mathbf{R}$ is the covariance matrix of the input, $R_{i j}=1 / p \sum_{\mu=1}^{p} x_{i}^{\mu} x_{j}^{\mu}$, a symmetric and nonnegative $N \times N$ matrix; the $N$-dimensional vector $Q$ has components $q_{i}=$ $1 / p \sum_{\mu=1}^{p} d^{\mu} x_{i}^{\mu}$, and the constant $c=1 / p \sum_{\mu=1}^{p}\left(d^{\mu}\right)^{2}$. The gradient is given by $\nabla E=\mathbf{R} W-Q$, while the Hessian matrix of second derivatives is $\mathbf{H}=\mathbf{R}$.

The solution space of vectors $W^{*}$ which minimize $E(W)$ is the subspace of solutions of the linear equation $\mathbf{R} W=Q$, resulting from $\nabla E=0$. This subspace reduces to a point if $\mathbf{R}$ is full rank. The diagonalization of $\mathbf{R}$ provides a diagonal matrix $\boldsymbol{\Lambda}$ formed by its eigenvalues, and a matrix $\mathbf{U}$ formed by its eigenvectors. Since $\mathbf{R}$ is nonnegative, all eigenvalues satisfy $\lambda \geq 0$.

Consider now a two-step coordinate transformation: a translation $V^{\prime}=W-W^{*}$ provides new coordinates centered at the solution point; it is followed by a rotation $V=\mathbf{U} V^{\prime}=\mathbf{U}\left(W-W^{*}\right)$ onto the principal axes of the error surface. In the new coordinate system

$$
\begin{equation*}
E(V)=\frac{1}{2} V^{T} \boldsymbol{\Lambda} V+E_{0} \tag{3.5}
\end{equation*}
$$

with $\boldsymbol{\Lambda}=\mathbf{U}^{T} \mathbf{R U}$ and $E_{0}=E\left(W^{*}\right)$. Then $\partial E / \partial v_{j}=\lambda_{j} v_{j}$, and $\partial^{2} E / \partial v_{j} \partial v_{k}=$ $\lambda_{j} \delta_{j k}$. The eigenvalues of the input covariance matrix give the second derivatives of the error surface with respect to its principal axes.

In the new coordinate system the Hessian matrix is the diagonal matrix $\boldsymbol{\Lambda}$, and the rule for weight updates becomes a set of $N$ decoupled equations:

$$
\begin{equation*}
V(k+1)=V(k)-\eta \mathbf{\Lambda} V(k) \tag{3.6}
\end{equation*}
$$

The evolution of each component along a principal direction is given by

$$
\begin{equation*}
v_{j}(k)=\left(1-\eta \lambda_{j}\right)^{k} v_{j}(0) \tag{3.7}
\end{equation*}
$$

so that $v_{j}$ will converge to zero (and thus $w_{j}$ to the solution $w_{j}^{*}$ ) provided that $0<\eta<2 / \lambda_{j}$. In this regime $v_{j}$ decays to zero exponentially, with characteristic time $\tau_{j}=\left(\eta \lambda_{j}\right)^{-1}$. The range $1 / \lambda_{j}<\eta<2 / \lambda_{j}$ corresponds to underdamped dynamics: the step size is large and convergence to the solution occurs through
oscillatory behavior. The range $0<\eta<1 / \lambda_{j}$ corresponds to overdamped dynamics: the step size is small and convergence requires many iterations. Critical damping occurs for $\eta=1 / \lambda_{j}$; if such choice is possible, the solution is reached in one iteration (Newton's method).
If all eigenvalues are equal, $\lambda_{j}=\lambda$ for all $1 \leq j \leq N$, the Hessian matrix is diagonal: $\mathbf{H}=\boldsymbol{\Lambda}$. Convergence can be obtained in one iteration, with optimal step size $\eta=1 / \lambda$, and learning time $\tau=1$. This highly symmetric case occurs when cross-sections of $E(W)$ are hyperspheres in the $N$-dimensional space $\{W\}$. Such high degree of symmetry is rarely encountered: correlated inputs result in nondiagonal elements for $\mathbf{H}$, and the principal directions are rotated with respect to the original coordinates. The cross-sections of $E(W)$ are elliptical, with different eigenvalues along different principal directions. Convergence requires $0<\eta<2 / \lambda_{j}$ for all $1 \leq j \leq N$, thus $\eta$ must be chosen in the range $0<\eta<1 / \lambda_{\max }$, where $\lambda_{\max }$ is the largest eigenvalue. The slowest time constant in the system is $\tau_{\max }=\left(\eta \lambda_{\min }\right)^{-1}$, where $\lambda_{\min }$ is the lowest nonzero eigenvalue. The optimal step size $\eta=1 / \lambda_{\max }$ thus leads to $\tau_{\max }=\lambda_{\max } / \lambda_{\min }$ for the decay along the principal direction of smallest nonzero curvature. A distribution of eigenvalues in the range $\lambda_{\min } \leq \lambda \leq \lambda_{\max }$ results in a distribution of learning times, with average $\langle\tau\rangle=\lambda_{\max }\langle 1 / \lambda\rangle$.
This analysis demonstrates that learning dynamics in quadratic surfaces are fully controlled by the eigenvalue distribution of the Hessian matrix. It is thus of interest to investigate such eigenvalue distribution.

## 4 EIGENVALUE SPECTRUM

The simple linear unit of Eq. (3.1) leads to the error function (3.4), for which the Hessian is given by the covariance matrix

$$
\begin{equation*}
R_{i j}=\mathrm{i} / p \sum_{\mu=1}^{p} x_{i}^{\mu} x_{j}^{\mu} \tag{4.1}
\end{equation*}
$$

It is assumed that the input components $\left\{x_{i}^{\mu}\right\}$ are independent, and drawn from a distribution with mean $m$ and variance $v$. The size of the training set is quantified by the ratio $\alpha=p / N$ between the number of training examples and the dimensionality of the input vector . The eigenvalue spectrum has been computed (Le Cun, Kanter, and Solla, 1990), and it exhibits three dominant features:
(a) If $p<N$, the rank of the matrix $\mathbf{R}$ is $p$. The existence of $(N-p)$ zero eigenvalues out of $N$ results in a delta function contribution of weight $(1-\alpha)$ at $\lambda=0$ for $\alpha<1$.
(b) A continuous part of the spectrum,

$$
\begin{equation*}
\rho(\lambda)=\frac{\left[4 \alpha v^{2}-(\lambda \alpha-v(1+\alpha))^{2}\right]^{1 / 2}}{2 \pi \lambda v} \tag{4.2}
\end{equation*}
$$

within the bounded interval $\lambda_{-}<\lambda<\lambda_{+}$, with $\lambda_{ \pm}=(1 \pm \sqrt{\alpha})^{2} v / \alpha$ (Krogh and Hertz, 1991). Note that $\rho(\lambda)$ is controlled only by the variance $v$ of the distribution from which the inputs are drawn. The bounds $\lambda_{ \pm}$are well defined, and of order one. For all $\alpha<1, \lambda_{-}>0$, indicating a gap at the lower end of the spectrum.
(c) An isolated eigenvalue of order $N, \lambda_{N}$, present in the case of biased inputs ( $m \neq 0$ ).
True correlations between pairs $\left(x_{j}, x_{k}\right)$ of input components might lead to a quite different spectrum from the one described above.
The continuous part (4.2) of the eigenvalue spectrum has been computed in the $N \rightarrow \infty$ limit, while keeping $\alpha$ constant and finite. The magnitude of finite size effects has been investigated numerically for $N \leq 200$ and various values of $\alpha$. Results for $N=200$, shown in Fig. 1, indicate that finite size effects are negligible: the distribution $\rho(\lambda)$ is bounded within the interval $\left[\lambda_{-}, \lambda_{+}\right]$, in good agreement with the theoretical prediction, even for such small systems. The result (4.2) is thus applicable in the finite $p=\alpha N$ case, an important regime given the limited availability of training data in most learning problems.


Figure 1: Spectral density $\rho(\lambda)$ predicted by Eq. (4.2) for $m=0, v=1$, and $\alpha=0.6,1.2,4$, and 16. Experimental histograms for $\alpha=0.6$ (full squares) and $\alpha=4$ (open squares) are averages over 100 trials with $N=200$ and $x_{i}^{\mu}= \pm 1$ with probability $1 / 2$ each.

The existence of a large eigenvalue $\lambda_{N}$ is easily understood by considering the structure of the covariance matrix $\mathbf{R}$ in the $p \rightarrow \infty$ limit, a regime for which a detailed analysis is available in the adaptive filters literature (Widrow and Stearns, 1985). In this limit, all off-diagonal elements of $R$ are equal to $m^{2}$, and all diagonal elements are equal to $v+m^{2}$. The eigenvector $U_{N}=(1 \ldots 1)$ thus corresponds to the eigenvalue $\lambda_{N}=N m^{2}+v$. The remaining $(N-1)$ eigenvalues are all equal to $v$ (note that the continuous part of the spectrum collapses onto a delta function at $\lambda_{-}=\lambda_{+}=v$ as $\left.p \rightarrow \infty\right)$, thus satisfying $\operatorname{tr} \mathbf{R}=N\left(m^{2}+v\right)$. The large part of $\lambda_{N}$ is eliminated for centered distributions with $m=0$, such as $x_{i}^{\mu}= \pm 1$ with probability $1 / 2$, or $x_{i}^{\mu}=3,-1,-2$ with probability $1 / 3$. Note that although $m$ is
crucial in controlling the existence of an isolated eigenvalue of order $N$, it plays no role in the spectral density of Eq. (4.2).

## 5 LEARNING TIME

Consider the learning time $\tau=\alpha\left(\lambda_{\max } / \lambda_{\min }\right)$. The eigenvalue ratio ( $\lambda_{\max } / \lambda_{\min }$ ) measures the maximum number of iterations, and the factor of $\alpha$ accounts for the time needed for each presentation of the full training set.

For $m=0, \lambda_{\max }=\lambda_{+}$, and $\lambda_{\min }=\lambda_{-}$. The learning time $\tau=\alpha\left(\lambda_{+} / \lambda_{-}\right)$can be easily computed using Eq. (4.2): $\tau=\alpha(1+\sqrt{\alpha})^{2} /(1-\sqrt{\alpha})^{2}$. As a function of $\alpha, \tau$ diverges at $\alpha=1$, and, surprisingly, goes through a minimum at $\alpha=(1+\sqrt{2})^{2}=$ 5.83 before diverging linearly for $\alpha \rightarrow \infty$. Numerical simulations were performed to estimate $\tau$ by counting the number $T$ of presentations of training examples needed to reach an allowed error level $\tilde{E}$ through gradient descent. If the prescribed error $\tilde{E}$ is sufficiently close to the minimum error $E_{0}, T$ is controlled by the slowest mode, and it provides a good estimate for $\tau$. Numerical results for $T$ as a function of $\alpha$, shown in Fig. 2, were obtained by training a single linear neuron on randomly generated vectors. As predicted, the curve exhibits a clear maximum at $\alpha=1$, as well as a minimum between $\alpha=4$ and $\alpha=5$. The existence of such optimal training set size for fast learning is a surprising result.


Figure 2: Number of iterations $T$ (averaged over 20 trials) needed to train a linear neuron with $N=100$ inputs. The $x_{j}^{\mu}$ are uniformly distributed between -1 and +1 . Initial and target couplings $W$ are chosen randomly from a uniform distribution within the $[-1,+1]^{N}$ hypercube. Gradient descent is considered complete when the error reaches the prescribed value $\tilde{E}=0.001$ above the $E_{0}=0$ minimum value.

Biased inputs $m \neq 0$ produce a large eigenvalue $\lambda_{\max }=\lambda_{N}$, proportional to $N$ and responsible for slow convergence. A simple approach to reducing the learning time is to center each input variable $x_{j}$ by subtracting its mean. An obvious source of systematic bias $m$ is the use of activation functions which restrict the state variables to the $[0,1]$ interval. Symmetric activation functions such as the hyperbolic tangent are empirically known to yield faster convergence than their nonsymmetric counterparts such as the logistic function. Our results provide an explanation to this observation.

The extension of these results to multi-layer networks rests on the observation that each neuron $i$ receives state information $\left\{x_{j}\right\}$ from the $j \in S_{i}$ neurons that feed into it, and can be viewed as minimizing a local objective function $E_{i}$ whose Hessian involves the the covariance matrix of such inputs. If all input variables are uncorrelated and have zero mean, no large eigenvalues will appear. But states with $\overline{x_{j}}=m \neq 0$ produce eigenvalues proportional to the number of input neurons $N_{i}$ in the set $S_{i}$, resulting in slow convergence if the connectivity is large. An empirically known solution to this problem, justified by our theoretical analysis, is to use individual learning rates $\eta_{i}$ inversely proportional to the number of inputs $N_{i}$ to the $i$-th neuron. Yet another approach is to keep a running estimate of the average $\overline{x_{j}}$ and use centered state variables $\tilde{x_{j}}=x_{j}-\overline{x_{j}}$. Such algorithm results in considerable reductions in learning time.

## 6 CONCLUSIONS

Our results are based on a rigorous calculation of the eigenvalue spectrum for a symmetric matrix constructed from the outer product of random vectors. The spectral density provides a full description of the relaxation of a single adaptive linear unit, and yields a surprising result for the optimal size of the training set in batch learning. Various aspects of the dynamics of learning in multi-layer networks composed of nonlinear units are clarified: the theory justifies known empirical methods and suggests novel approaches to reduce learning times.

## References

A. Krogh and J. A. Hertz (1991), 'Dynamics of generalization in linear perceptrons', in Advances in Neural Information Processing Systems 3, ed. by D. S. Touretzky and R. Lippman, Morgan Kaufmann (California).
Y. Le Cun, I. Kanter, and S. A. Solla (1990), 'Eigenvalues of covariance matrices: application to neural-network learning', Phys. Rev., to be published.
D. E. Rumelhart, G. E. Hinton, and R. J. Williams (1986), 'Learning representations by back-propagating errors', Nature 323, 533-536.
B. Widrow and S. D. Stearns (1985), Adaptive Signal Processing, Prentice-Hall (New Jersey).

