

# Online learning from finite training sets and robustness to input bias

**Peter Sollich**<sup>1</sup>

Department of Physics, University of Edinburgh, Edinburgh EH9 3JZ, U.K.

**David Barber**<sup>2</sup>

Department of Medical Biophysics, University of Nijmegen,  
6525 EZ Nijmegen, The Netherlands.

**We analyse online gradient descent learning from finite training sets at non-infinitesimal learning rates  $\eta$ . Exact results are obtained for the time-dependent generalization error of a simple model system: a linear network with a large number of weights  $N$ , trained on  $p = \alpha N$  examples. This allows us to study in detail the effects of finite training set size  $\alpha$  on, for example, the optimal choice of learning rate  $\eta$ . We also compare online and offline learning, for respective optimal settings of  $\eta$  at given final learning time. Online learning turns out to be much more robust to input bias and actually outperforms offline learning when such bias is present; for unbiased inputs, online and offline learning perform almost equally well.**

## 1 Introduction

The analysis of online (gradient descent) learning, which is one of the most common approaches to supervised learning found in the neural networks community, has recently been the focus of much attention. The characteristic feature of online learning is that the weights of a network ('student') are updated each time a new training example is presented, such that the error on this example is reduced. In offline learning, on the other hand, the total error on all examples in the training set is accumulated before a gradient

---

<sup>1</sup>Royal Society Dorothy Hodgkin Research Fellow. Email P.Sollich@ed.ac.uk

<sup>2</sup>Email davidb@mbfys.kun.nl

descent weight update is made. For a given training set and starting weights, offline learning is entirely deterministic. Online learning, on the other hand, is a stochastic process due to the random choice of training example (from the given training set) for each update; in fact, it can essentially be viewed as a ‘noisy’ version of offline learning. The two are equivalent only in the limit where the learning rate  $\eta \rightarrow 0$  (see, *e.g.*, Heskes and Kappen, 1991). For both online and offline learning, the main quantity of interest is normally the evolution of the generalization error: After a given number of weight updates, how well does the student approximate the input-output mapping (‘teacher’) underlying the training examples?

Most analytical treatments of online learning assume either that the size of the training set is infinite, or that the learning rate  $\eta$  is vanishingly small. Both of these restrictions are undesirable: In practice, most training sets are finite<sup>3</sup>, and non-infinitesimal values of  $\eta$  are needed to ensure that the learning process converges after a reasonable number of updates. General results have been derived for the difference between online and offline learning to first order in  $\eta$ , which apply to training sets of any size (see, *e.g.*, Heskes and Kappen, 1991). However, these do not directly address the question of generalization performance. The most explicit analysis of the time evolution of the generalization error for finite training sets was provided by Krogh and Hertz (1992) for a scenario very similar to the one we consider below. Their  $\eta \rightarrow 0$  offline calculation will serve as a baseline for our work. For finite  $\eta$ , progress has been made in particular for so-called soft committee machine network architectures (see, *e.g.*, Saad and Solla, 1995; Biehl and Schwarze, 1995), but only for the case of infinite training sets.

In this paper, we give an exact analysis of online learning in a simple model system. Our aim is twofold: (1) to assess how the combination of non-infinitesimal learning rates  $\eta$  and finite training sets (containing  $\alpha$  examples per weight) affects online learning, and (2) to compare the generalization performance of online and offline learning. A priori, one

---

<sup>3</sup>Online learning can also be used to learn teacher rules that vary in time. The assumption of an infinite set (or ‘stream’) of training examples is then much more plausible, and in fact necessary for continued adaptation of the student. We do not consider this case in the following.

may expect online learning to perform worse due to its inherent randomness. We show that this disadvantage is actually negligible when online and offline learning are compared on an equal footing, *i.e.*, for their respective optimal learning rates. More importantly, we will see that online learning is much more robust to *input bias* than offline learning and actually performs *better* than the offline version in the case of biased inputs.

## 2 Model definition

We consider training of a linear student network with input-output relation

$$y = \frac{1}{\sqrt{N}} \mathbf{w}^T \mathbf{x}$$

Here  $\mathbf{x}$  is an  $N$ -dimensional vector of real-valued inputs,  $y$  the single real output and  $\mathbf{w}$  the weight vector of the network. ‘ $T$ ’ denotes the transpose of a vector and the factor  $1/\sqrt{N}$  is introduced for convenience. In online learning, whenever a training example  $(\mathbf{x}, y)$  is presented to the network, its weight vector is updated along the gradient of the squared error<sup>4</sup> on this example, *i.e.*,

$$\Delta \mathbf{w} = -\eta \nabla_{\mathbf{w}} \frac{1}{2} \left( y - \frac{1}{\sqrt{N}} \mathbf{w}^T \mathbf{x} \right)^2 = \eta \left( \frac{1}{\sqrt{N}} y \mathbf{x} - \frac{1}{N} \mathbf{x} \mathbf{x}^T \mathbf{w} \right)$$

where  $\eta$  is the learning rate. We are primarily interested in the case of online learning from finite training sets, where for each update an example is randomly chosen from a given set  $\{(\mathbf{x}^\mu, y^\mu), \mu = 1 \dots p\}$  of  $p$  training examples. If example  $\mu$  is chosen for update  $n$ , the weight vector is changed to

$$\mathbf{w}_{n+1} = \left\{ 1 - \frac{\eta}{N} [\mathbf{x}^\mu (\mathbf{x}^\mu)^T + \gamma] \right\} \mathbf{w}_n + \eta \frac{1}{\sqrt{N}} y^\mu \mathbf{x}^\mu \quad (\text{online}) \quad (1)$$

---

<sup>4</sup>We consider only squared error here, which is probably the most commonly used error measure. We also restrict our analysis to ‘vanilla’ gradient descent learning, excluding more sophisticated learning algorithms.

Here we have also included a weight decay  $\gamma$ . The update rule for *offline* learning is similar, but here the gradients for all  $p$  different training examples are accumulated before a weight update is made:

$$\mathbf{w}_{(r+1)p} = [1 - \eta(\lambda + \mathbf{A})]\mathbf{w}_{rp} + \frac{\eta}{\sqrt{N}} \sum_{\mu} y^{\mu} \mathbf{x}^{\mu} \quad (\text{offline}) \quad (2)$$

Here  $r$  is the number of offline weight updates; in order to compare online and offline learning at equal computational cost, we index the weight vectors for both cases by the number of gradient calculations, which is  $n = rp$  in the offline case. The matrix

$$\mathbf{A} = \frac{1}{N} \sum_{\mu} \mathbf{x}^{\mu} (\mathbf{x}^{\mu})^{\text{T}}$$

is the correlation matrix of the training inputs, and  $\lambda = \gamma\alpha$  is the weight decay rescaled by the number of examples per weight,  $\alpha = p/N$ . We will generally use  $\lambda$  (rather than  $\gamma$ ) to characterize the strength of the weight decay, for both online and offline learning. For simplicity, all student weights are assumed to be initially zero, *i.e.*,  $\mathbf{w}_{n=0} = \mathbf{0}$ .

The main quantity of interest to us is the *generalization error* of the student and its evolution during learning. We assume that the training examples are generated by a linear ‘teacher’, *i.e.*,  $y^{\mu} = \mathbf{w}_*^{\text{T}} \mathbf{x}^{\mu} / \sqrt{N} + \xi^{\mu}$ , where  $\xi^{\mu}$  is zero mean additive noise of variance  $\sigma^2$ . The teacher weight vector is taken to be normalized to  $\mathbf{w}_*^2 = N$  for simplicity. We first investigate the case of unbiased inputs ( $\langle \mathbf{x} \rangle = \mathbf{0}$ ), assuming that input vectors are sampled randomly from an isotropic distribution over the hypersphere  $\mathbf{x}^2 = N$  (biased inputs will be considered in Section 4). The generalization error, defined as the average of the squared error between student and teacher outputs for random inputs, is then

$$\epsilon_g = \frac{1}{2N} (\mathbf{w}_n - \mathbf{w}_*)^2 = \frac{1}{2N} \mathbf{v}_n^2 \quad \text{where} \quad \mathbf{v}_n = \mathbf{w}_n - \mathbf{w}_*.$$

In order to make the scenario analytically tractable, we focus on the limit  $N \rightarrow \infty$  of a large number of input components and weights, taken at constant number of examples

per weight  $\alpha = p/N$  and updates per weight ('learning time')  $t = n/N$ . In this limit, the generalization error  $\epsilon_g(t)$  becomes self-averaging (see however Section 4) and can be calculated by averaging both over the random selection of examples from a given training set and over all training sets. Our results can be straightforwardly extended to the case of perceptron teachers with a nonlinear transfer function, as in (Sollich, 1995).

### 3 Unbiased inputs

#### 3.1 Outline of calculation

We begin by deriving from the online learning weight update (1) an update equation for the 'selection' average of the generalization error (*i.e.*, its average with respect to the random choice of training examples for each update, denoted generically by  $\langle \dots \rangle$ ). In fact, it will turn out to be useful to consider a slightly generalized version of the generalization error,  $\epsilon_n = \frac{1}{2N} \mathbf{v}_n^T \mathbf{M} \mathbf{v}_n$ , with  $\mathbf{M}$  an arbitrary  $N \times N$  matrix. To get the update equation for  $\langle \epsilon_n \rangle$ , we first rewrite (1) in terms of  $\mathbf{v}_n$ , the difference between student and teacher weight vectors:

$$\mathbf{v}_{n+1} = \left\{ 1 - \eta \left[ \frac{1}{N} \mathbf{x}^\mu (\mathbf{x}^\mu)^T + \frac{\lambda}{p} \right] \right\} \mathbf{v}_n + \eta \frac{1}{\sqrt{N}} \xi^\mu \mathbf{x}^\mu - \frac{\eta \lambda}{p} \mathbf{w}_* \quad (3)$$

This can now be multiplied by its transpose, with the matrix  $\mathbf{M}$  inserted, and the selection average for update  $n$  performed. Discarding terms which become negligible in the large  $N$  limit, one finds after a little algebra

$$\begin{aligned} N (\langle \epsilon_{n+1} \rangle - \langle \epsilon_n \rangle) &= \frac{\tilde{\eta}}{N} (\mathbf{b} - \lambda \mathbf{w}_*)^T \mathbf{M} \langle \mathbf{v}_n \rangle - \frac{\tilde{\eta}}{N} \left\langle \mathbf{v}_n^T \left[ \lambda \mathbf{M} + \frac{1}{2} (\mathbf{A} \mathbf{M} + \mathbf{M} \mathbf{A}) \right] \mathbf{v}_n \right\rangle \\ &+ \frac{\tilde{\eta}^2 \alpha}{N} \sum_{\mu} \frac{1}{N} (\mathbf{x}^\mu)^T \mathbf{M} \mathbf{x}^\mu \left\{ \frac{1}{2} (\xi^\mu)^2 - \xi^\mu \frac{1}{\sqrt{N}} (\mathbf{x}^\mu)^T \langle \mathbf{v}_n \rangle + \frac{1}{2N} \langle \mathbf{v}_n^T \mathbf{x}^\mu (\mathbf{x}^\mu)^T \mathbf{v}_n \rangle \right\} \end{aligned} \quad (4)$$

where  $\tilde{\eta} = \eta/\alpha$  is a rescaled learning rate, and  $\mathbf{b} = \sum_{\mu} \xi^\mu \mathbf{x}^\mu / \sqrt{N}$ . We now want to transform (4) into a closed dynamical equation for  $\langle \epsilon_n \rangle$ . This means that all selection

averages need to be either eliminated or reduced to averages of the same form as  $\langle \epsilon_n \rangle$ . For the two terms linear in  $\langle \mathbf{v}_n \rangle$ , this is straightforward: The selection average of (1) yields directly

$$N (\langle \mathbf{v}_{n+1} \rangle - \langle \mathbf{v}_n \rangle) = \tilde{\eta} [ -(\lambda + \mathbf{A}) \langle \mathbf{v}_n \rangle + \mathbf{b} - \lambda \mathbf{w}_* ].$$

Starting from  $\mathbf{v}_0 = -\mathbf{w}_*$ , this can easily be solved, with the result (for  $N \rightarrow \infty$ )

$$\langle \mathbf{v}_n \rangle = (\lambda + \mathbf{A})^{-1} \{ \mathbf{b} - \lambda \mathbf{w}_* - \exp[-\tilde{\eta} t (\lambda + \mathbf{A})] (\mathbf{b} + \mathbf{A} \mathbf{w}_*) \} \quad (5)$$

from which the selection average has now disappeared. Learning rate and learning time enter only through the combination  $\tau = \tilde{\eta} t$ ; this rescaled time will be useful later on. In (4), the remaining terms quadratic in  $\mathbf{v}_n$  now present the main problem. The second term on the r.h.s. shows that the evolution of  $\epsilon_g = \epsilon_n(\mathbf{M} = \mathbf{1})$  depends on  $\epsilon_n(\mathbf{M} = \mathbf{A})$  which in turn depends on  $\epsilon_n(\mathbf{M} = \mathbf{A}^2)$  and so on, yielding an infinite hierarchy of order parameters. This problem was solved in (Sollich and Barber, 1997a) by introducing an auxiliary parameter  $h$  through  $\mathbf{M} = \exp(h\mathbf{A})$ ; all order parameters  $\epsilon_n(\mathbf{M} = \mathbf{A}^m)$ ,  $m = 1, 2, \dots$ , can then be obtained by differentiating  $\epsilon_n(h) = \frac{1}{2N} \mathbf{v}_n^T \exp(h\mathbf{A}) \mathbf{v}_n$ .

Here we choose a different route, which is somewhat more transparent and also more easily adapted to the case of biased inputs to be considered later. The main idea is to decompose the evolution of  $\mathbf{v}_n$  into components defined by eigenvectors of the input correlation matrix  $\mathbf{A}$ . (This is equivalent to changing to a coordinate system in which  $\mathbf{A}$  is diagonal, and then considering the components of  $\mathbf{v}_n$  separately.) More precisely, let us order the  $N$  eigenvalues of  $\mathbf{A}$  in ascending order and split them into  $K$  equal blocks, labeled by  $\kappa = 1 \dots K$ , each containing  $N/K$  eigenvalues. Let  $\mathbf{P}^\kappa$  be the projector matrices onto the spaces spanned by the eigenvectors of each block. Then  $\mathbf{v}_n = \sum_\kappa \mathbf{P}^\kappa \mathbf{v}_n$ ; likewise, the generalization error is decomposed as

$$\epsilon_g = \frac{1}{K} \sum_\kappa \epsilon_n^\kappa, \quad \epsilon_n^\kappa = \frac{K}{2N} \mathbf{v}_n^T \mathbf{P}^\kappa \mathbf{v}_n$$

Each of the generalization error components  $\epsilon_n^\kappa$  obeys the update equation (4), with  $\mathbf{M} = K\mathbf{P}^\kappa$ . But these equations now become closed, because

$$\mathbf{A}\mathbf{P}^\kappa = \mathbf{P}^\kappa\mathbf{A} \approx a^\kappa\mathbf{P}^\kappa$$

where  $a^\kappa$  is an eigenvalue from the  $\kappa$ -th block (formally, this approximation becomes exact in the limit  $K \rightarrow \infty$ , where the spread of eigenvalues within each block tends to zero). This immediately reduces the second term on the right-hand side of (4) to  $-2\tilde{\eta}(\lambda + a^\kappa)\langle\epsilon_n^\kappa\rangle$ . Only the very last term of (4) now remains to be brought into a similar form. This is achieved by noting that the factors  $c_\kappa^\mu = (K/N)(\mathbf{x}^\mu)^\top\mathbf{P}^\kappa\mathbf{x}^\mu$  are ‘within-sample self-averaging’ (Sollich and Barber, 1997a): Up to fluctuations which vanish as  $O(N^{-1/2})$  for large  $N$ , all  $c^\mu$  are equal to each other and hence to the training set (‘sample’) average

$$\frac{1}{p}\sum_\mu c_\kappa^\mu = \frac{K}{\alpha N}\text{tr}\mathbf{A}\mathbf{P}^\kappa \approx \frac{a^\kappa}{\alpha}$$

The last approximation again becomes exact<sup>5</sup> for  $K \rightarrow \infty$ . The factors  $c_\kappa^\mu = a^\kappa/\alpha$  can therefore be taken out of the sum over  $\mu$  in (4), leaving the selection average

$$\sum_\mu \frac{1}{2N}\langle\mathbf{v}_n^\top\mathbf{x}^\mu(\mathbf{x}^\mu)^\top\mathbf{v}_n\rangle = \frac{1}{2}\langle\mathbf{v}_n^\top\mathbf{A}\mathbf{v}_n\rangle \approx \frac{1}{K}\sum_\kappa a^\kappa\langle\epsilon_n^\kappa\rangle$$

We now have all the ingredients to write (4) as a closed system of evolution equations for the  $\epsilon_n^\kappa$ . In the large  $N$  limit, the change  $N(\langle\epsilon_{n+1}^\kappa\rangle - \langle\epsilon_n^\kappa\rangle)$  due to an update becomes the time derivative  $\partial_t\epsilon^\kappa$ , and  $\langle\epsilon_n^\kappa\rangle \rightarrow \epsilon^\kappa(t)$ . Using the rescaled time  $\tau = \tilde{\eta}t$  introduced above, one then has

$$[\partial_\tau + 2(\lambda + a^\kappa)]\epsilon^\kappa(\tau) = V^\kappa(\tau) + \tilde{\eta}W^\kappa(\tau) + \tilde{\eta}a^\kappa\frac{1}{K}\sum_{\kappa'} a^{\kappa'}\epsilon^{\kappa'}(\tau) \quad (6)$$

---

<sup>5</sup>The large  $K$  limit needs to be taken *after* the limit  $N \rightarrow \infty$  for ‘within-sample self-averaging’ to hold; this is why one cannot take  $K = N$  from the outset.

Here the functions  $V^\kappa(\tau)$  and  $W^\kappa(\tau)$  are

$$\begin{aligned} V^\kappa &= \frac{K}{N} (\mathbf{b} - \lambda \mathbf{w}_*)^T \mathbf{P}^\kappa \langle \mathbf{v}_n \rangle \\ W^\kappa &= a^\kappa \left[ \frac{1}{2N} \sum_\mu (\xi^\mu)^2 - \frac{1}{N} \mathbf{b}^T \langle \mathbf{v}_n \rangle \right] \end{aligned}$$

with  $\langle \mathbf{v}_n \rangle$  given by (5). Having derived (6), the rest of the calculation is fairly straightforward. Eq. (6) is formally solved using Laplace transforms with respect to  $\tau$ , for example  $\hat{\epsilon}^\kappa(z) = \int_0^\infty d\tau \exp(-z\tau) \epsilon^\kappa(\tau)$ :

$$\hat{\epsilon}^\kappa(z) = \frac{1}{z + 2(\lambda + a^\kappa)} \left[ \epsilon^\kappa(0) + \hat{V}^\kappa(z) + \tilde{\eta} \hat{W}^\kappa(z) + \tilde{\eta} a^\kappa \frac{1}{K} \sum_{\kappa'} a^{\kappa'} \hat{\epsilon}^{\kappa'}(z) \right] \quad (7)$$

with the initial condition  $\epsilon^\kappa(0) = \frac{K}{2N} \mathbf{w}_*^T \mathbf{P}^\kappa \mathbf{w}_*$ . Multiplying by  $a^\kappa$  and summing over  $\kappa$  then gives a self-consistency equation for  $K^{-1} \sum_\kappa a^\kappa \hat{\epsilon}^\kappa(z)$  which is easily solved. Inserting the solution into (7) then gives an explicit expression for  $\hat{\epsilon}^\kappa(z)$  and hence for the Laplace transform of the generalization error,  $\hat{\epsilon}_g(z) = K^{-1} \sum_\kappa \hat{\epsilon}^\kappa(z)$ . As a final step, the average over all training sets (*i.e.*, training inputs  $\mathbf{x}^\mu$  and output noises  $\xi^\mu$ ) is then carried out. In the end, everything can be written in terms of averages over the known eigenvalue spectrum (Hertz et al., 1989; Sollich, 1994) of the input correlation matrix  $\mathbf{A}$ . The explicit form of the final result (Sollich and Barber, 1997a) is rather cumbersome; we omit it here and note only the relatively simple dependence on  $\eta$ :

$$\hat{\epsilon}_g(z) = \hat{\epsilon}_0(z) + \frac{\eta \hat{\epsilon}_1(z)}{1 - \eta \hat{\epsilon}_2(z)} \quad (8)$$

The functions  $\epsilon_i(z)$  ( $i = 0 \dots 2$ ) depend on  $\alpha$ ,  $\sigma^2$  and  $\lambda$  (and, of course,  $z$ ), but are independent of  $\eta$ . The teacher weights do not appear explicitly: because of the isotropy of the input distribution, only the length of the teacher weight vector matters once an average over training sets has been taken, and this has already been fixed to  $\mathbf{w}_*^2 = N$ .

The calculation of the generalization error for offline learning is much simpler than that for the online case due to the absence of the selection average. In fact, the offline






Figure 1: Asymptotic generalization error  $\epsilon_\infty$  vs  $\eta$  and  $\lambda$ .  $\alpha$  as shown,  $\sigma^2 = 0.1$ .

weight update (2) can be iterated directly to yield

$$\mathbf{v}_{rp} = (\lambda + \mathbf{A})^{-1} \{ \mathbf{b} - \lambda \mathbf{w}_* - [1 - \eta(\lambda + \mathbf{A})]^r (\mathbf{b} + \mathbf{A} \mathbf{w}_*) \} \quad (9)$$

Multiplying this by its transpose gives directly the generalization error, and the average over training sets can then be carried out in the usual fashion (see, *e.g.*, Hertz et al., 1989). As expected on general grounds, for  $\eta \rightarrow 0$  (and only then) one obtains the same result as for online learning, corresponding to the term  $\hat{\epsilon}_0(z)$  in (8).

### 3.2 Discussion

We now briefly highlight some features of our exact result (8) for the generalization error achieved by online learning; a somewhat more detailed exposition can be found in (Sollich and Barber, 1997b). We discuss the asymptotic generalization error  $\epsilon_\infty$ , the convergence speed for large learning times, and the behaviour at small  $t$ ; finally, we compare online and offline learning. For numerical evaluations, we generally take  $\sigma^2 = 0.1$ , corresponding to a sizable noise-to-signal ratio of  $\sqrt{0.1} \approx 0.32$ .

The asymptotic generalization error is found directly from (8) using  $\epsilon_\infty = \epsilon_g(t \rightarrow \infty) = \lim_{z \rightarrow 0} z \hat{\epsilon}_g(z)$ . As expected, it coincides with the offline result (which is *independent* of  $\eta$ ) *only* for  $\eta = 0$ ; as  $\eta$  increases from zero, it increases monotonically. Reassuringly, our calculation reproduces existing  $O(\eta)$  results for this increase (Heskes and Kappen, 1991). In fig. 1 we plot  $\epsilon_\infty$  as a function of  $\eta$  and  $\lambda$  for  $\alpha = 0.5, 1, 2$ . We

observe that it is minimal for  $\lambda = \sigma^2$  and  $\eta = 0$ , as expected from corresponding results for offline learning (Krogh and Hertz, 1992)<sup>6</sup>. We also read off that for fixed  $\lambda$ ,  $\epsilon_\infty$  is an increasing function of  $\eta$ : The larger  $\eta$ , the more the weight updates tend to overshoot the minimum of the (total, *i.e.*, offline) training error. This causes a diffusive motion of the weights around their average asymptotic values (Heskes and Kappen, 1991) which increases  $\epsilon_\infty$ . In the absence of weight decay ( $\lambda = 0$ ) and for  $\alpha < 1$ , however,  $\epsilon_\infty$  is independent of  $\eta$ . In this case the training data can be fitted perfectly; every term in the total sum-of-squares training error is then zero and online learning does not lead to weight diffusion because all individual updates vanish. In general, the relative increase  $\epsilon_\infty(\eta)/\epsilon_\infty(\eta = 0) - 1$  due to nonzero  $\eta$  depends significantly on  $\alpha$ . For  $\eta = 1$  and  $\alpha = 0.5$ , for example, this increase is smaller than 6% for all  $\lambda$  (at  $\sigma^2 = 0.1$ ), and for  $\alpha = 1$  it is at most 13%. This means that in cases where training data is limited ( $p \approx N$ ),  $\eta$  can be chosen fairly large in order to optimize learning speed, without seriously affecting the asymptotic generalization error. In the large  $\alpha$  limit, on the other hand, one finds  $\epsilon_\infty = (\sigma^2/2)[1/\alpha + \eta/(2 - \eta)]$ . The relative increase over the value at  $\eta = 0$  therefore grows linearly with  $\alpha$ ; already for  $\alpha = 2$ , increases of around 50% can occur for  $\eta = 1$ .

Fig. 1 also shows that  $\epsilon_\infty$  diverges as  $\eta$  approaches a critical learning rate  $\eta_c$ : As  $\eta \rightarrow \eta_c$ , the ‘overshoot’ of the weight update steps becomes so large that the weights eventually diverge. From the Laplace transform (8), one finds that  $\eta_c$  is determined by  $\eta_c \hat{\epsilon}_2(z = 0) = 1$ ; it is a function of  $\alpha$  and  $\lambda$  only. As shown in fig. 2b-d,  $\eta_c$  increases with  $\lambda$ . This is reasonable, as the weight decay reduces the length of the weight vector at each update, counteracting potential weight divergences. In the small and large  $\alpha$  limits one has  $\eta_c = 2(1 + \lambda)$  and  $\eta_c = 2(1 + \lambda/\alpha)$ , respectively. For constant  $\lambda$ ,  $\eta_c$  therefore decreases<sup>7</sup> with  $\alpha$  (fig. 2b-d).

We now turn to the large  $t$  behaviour of the generalization error  $\epsilon_g(t)$ . For small  $\eta$ ,

---

<sup>6</sup>The optimal value of the *unscaled* weight decay decreases with  $\alpha$  as  $\gamma = \sigma^2/\alpha$ , because for large training sets there is less need to counteract noise in the training data by using a large weight decay.

<sup>7</sup>Conversely, for constant  $\gamma$ ,  $\eta_c$  *increases* with  $\alpha$  from  $2(1 + \gamma\alpha)$  to  $2(1 + \gamma)$ : For large  $\alpha$ , the weight decay is applied more often between repeat presentations of a training example that would otherwise cause the weights to diverge.




Figure 2: Sketch of definitions of  $\eta_{\text{min}}$  (minimal learning rate for slow mode),  $\eta_x$  (crossover to slow mode dominated convergence) and  $\eta_c$  (maximal (‘critical’) learning rate at which convergence still occurs), and their dependence on  $\alpha$ .

the most slowly decaying contribution to  $\epsilon_g(t)$ —the slowest ‘mode’—varies as  $\exp(-ct)$ , its decay constant  $c = \eta[\lambda + (\sqrt{\alpha} - 1)^2]/\alpha$  scaling linearly with  $\eta$ , the size of the weight updates, as expected (fig. 2a). For larger  $\eta$ , the picture changes due to a new slow mode arising from the denominator of (8). Interestingly, this mode exists only for  $\eta$  above a finite threshold  $\eta_{\text{min}} = 2/(\alpha^{1/2} + \alpha^{-1/2} - 1)$ . For finite  $\alpha$ , it could therefore not have been predicted from a small  $\eta$  expansion of  $\epsilon_g(t)$ . Its decay constant  $c_{\text{slow}}$  decreases to zero as  $\eta \rightarrow \eta_c$ , and crosses that of the normal mode at  $\eta_x(\alpha, \lambda)$  (fig. 2a). For  $\eta > \eta_x$ , the slow mode therefore determines the convergence speed for large  $t$ , and fastest convergence is obtained for  $\eta = \eta_x$ . However, it may still be advantageous to use lower values of  $\eta$  in order to lower the asymptotic generalization error (see below); values of  $\eta > \eta_x$  would deteriorate both convergence speed and asymptotic performance. Fig. 2b-d shows the dependence of  $\eta_{\text{min}}$ ,  $\eta_x$  and  $\eta_c$  on  $\alpha$  and  $\lambda$ . For  $\lambda$  not too large,  $\eta_x$  has a maximum at  $\alpha \approx 1$  (where  $\eta_x \approx \eta_c$ ), while decaying to  $\eta_x \approx \frac{1}{2}\eta_c$  for larger  $\alpha$ . This can be explained in terms of the anisotropy of the total training error surface (Sollich and Barber, 1997a), which is strongest for  $\alpha = 1$  and  $\lambda \rightarrow 0$ .

Consider now the small  $t$  behaviour of  $\epsilon_g(t)$ . Fig. 3 illustrates the dependence of  $\epsilon_g(t)$  on  $\eta$ ; comparison with simulation results for  $N = 50$  clearly confirms our calculations and demonstrates that finite  $N$  effects are not significant even for such fairly small  $N$ . For  $\alpha = 0.7$  (fig. 3a), we see that nonzero  $\eta$  acts as effective update noise, eliminating the minimum in  $\epsilon_g(t)$  which corresponds to over-training (Krogh and Hertz, 1992).  $\epsilon_\infty$  is also seen to be essentially independent of  $\eta$  as predicted for the small value of  $\lambda = 10^{-4}$

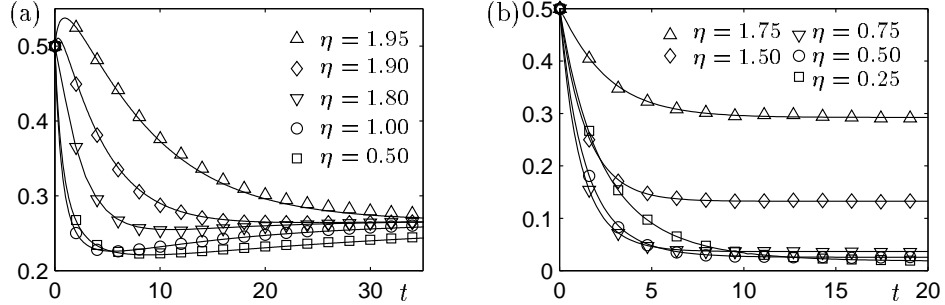


Figure 3:  $\epsilon_g$  vs  $t$  for different  $\eta$ . Simulations for  $N = 50$  are shown by symbols (standard errors less than symbol sizes).  $\lambda=10^{-4}$ ,  $\sigma^2=0.1$ . (a)  $\alpha=0.7$ , (b)  $\alpha=5$

chosen. For  $\alpha = 5$ , fig. 3b clearly shows the increase of  $\epsilon_\infty$  with  $\eta$ . It also illustrates how convergence first speeds up as  $\eta$  is increased from zero and then slows down again as  $\eta_c \approx 2$  is approached.

Above, we saw that the *asymptotic* generalization error  $\epsilon_\infty$  is minimal for  $\eta = 0$ . Fig. 4 shows what happens if we minimize  $\epsilon_g(t)$  instead for a given *final learning time*  $t$ , corresponding to a fixed amount of computational effort for training the network. As  $t$  increases, the optimal  $\eta$  decreases towards zero as required by the tradeoff between




Figure 4: (a) Optimal learning rate  $\eta$  vs. final learning time  $t$  for online (bold) and offline learning (thin lines), and (b) resulting generalization error  $\epsilon_g$ .  $\alpha = 1$ ,  $\sigma^2 = 0.1$ ,  $\lambda$  as shown. Note that although we plot offline results as continuous lines to avoid visual clutter, they are actually defined only at discrete values of the learning time,  $t = r\alpha$ , with  $r$  the number of offline weight updates.

asymptotic performance and convergence speed. For large  $t$ , the functional form of this decay is  $\eta_{\text{opt}} = (a + b \ln t)/t$  with  $t$ -independent coefficients  $a$  and  $b$  (Sollich and Barber, 1997a).

We now compare the performance of online learning to that of offline learning as calculated from (9). (The number of gradient calculations required for  $r$  offline weight updates is  $n = rp$ , corresponding to a learning time  $t = n/N = r\alpha$ ; the generalization error  $\epsilon_g(t)$  is therefore only defined for learning times  $t$  which are integer multiples of  $\alpha$ .) To compare online and offline learning on an equal footing, we again consider optimized values of  $\eta$  for given final learning time  $t$ . Fig. 4b shows that the performance loss from using online instead of offline learning is actually negligible. This may seem surprising given the stochasticity of weight updates in online learning, in particular for small  $t$ . However, fig. 4a shows that online learning can make up for this by allowing larger values of  $\eta$  to be used.

## 4 Biased inputs

### 4.1 Modifications to calculation

We now investigate how online and offline learning are affected by input bias  $\langle \mathbf{x} \rangle = \bar{\mathbf{x}} \neq \mathbf{0}$ . As a simple scenario of this kind, consider the case where the *deviations*  $\Delta \mathbf{x} = \mathbf{x} - \bar{\mathbf{x}}$  of the inputs from their average are still distributed isotropically over a hypersphere. We choose the radius  $R$  of this hypersphere such that the average value of  $\mathbf{x}^2$  is the same ( $N$ ) as for the unbiased case, *i.e.*,  $R^2 = N(1 - m^2)$  where  $m^2 = \bar{\mathbf{x}}^2/N$  measures the size of the bias. The generalization error (the squared deviation between student and teacher outputs averaged over all inputs) now has two components,

$$\epsilon_g = \frac{1}{2N} [(\bar{\mathbf{x}}^T \mathbf{v}_n)^2 + (1 - m^2) \mathbf{v}_n^2] \quad (10)$$

As before, we consider a teacher with weight vector of length  $\mathbf{w}_*^2 = N$ . In the presence of input bias, however, we also need to specify the average teacher output  $\bar{y} = \bar{\mathbf{x}}^T \mathbf{w}_*/\sqrt{N}$ .

This parameter is not constrained by our other assumptions; however, to limit the number of free parameters in the model, we choose it to have its typical root-mean-squared value when the directions of  $\mathbf{w}_*$  and  $\bar{\mathbf{x}}$  are uncorrelated:  $\bar{y}^2 = m^2$ .

As for the case of unbiased inputs, the evolution of the generalization error is largely determined by the eigenvalue spectrum of the input correlation matrix  $\mathbf{A}$ . This has been determined by a number of authors (LeCun et al., 1991; Wendemuth et al., 1993; Halkjær and Winther, 1997) and shows the following features: There is a ‘normal’ part of the spectrum, with eigenvalues which tend to finite values as  $N \rightarrow \infty$ ; the eigenvalues in this part of the spectrum are identical to those for the unbiased input case, except for a rescaling by the factor  $(1 - m^2)$ . Additionally, however, there is one isolated eigenvalue  $a_N = N\alpha m^2$  which is proportional to  $N$  and exists *only* in the presence of input bias. Intuitively, this corresponds to the fact that the component of the student weights along the direction of  $\bar{\mathbf{x}}$  is much more strongly determined by the training data because all input vectors have a component along  $\bar{\mathbf{x}}$ . Not surprisingly, therefore, the eigenvector corresponding to  $a_N$  is along the direction<sup>8</sup> of  $\bar{\mathbf{x}}$ .

We can see immediately that input bias has a drastic effect on offline learning by considering eq. (9): For the offline learning process to converge, the product of  $\eta$  and the largest eigenvalue of  $\lambda + \mathbf{A}$  must be less than two. In the presence of input bias, this gives the condition  $\eta < 2/(N\alpha m^2)$  (neglecting  $\lambda$ , which gives a negligible correction for  $N \rightarrow \infty$ ). The maximal learning rate is therefore *drastically reduced* from order unity to  $O(N^{-1})$ . A little reflection then shows that only the first contribution of the generalization error (10) decays for finite learning times; carrying out the average over training sets, one finds

$$\epsilon_g(t = r\alpha) = \frac{1}{2}m^2(1 - N\eta\alpha m^2)^{2r} + \frac{1}{2}(1 - m^2) \quad (11)$$

---

<sup>8</sup>In fact there is a small angle between this eigenvector and  $\bar{\mathbf{x}}$ , which however decreases as  $O((\alpha N)^{-1/2})$  as  $N$  grows large. LeCun et al. (1991) claimed that this angle is exactly zero; however, their argument cannot be quite correct as it would also entail that  $\mathbf{A}$  has only two different eigenvalues (whereas in reality it has a continuous spread of eigenvalues for any finite  $\alpha$ ).

The second contribution would only decay for learning times of  $O(N)$ , which are inaccessibly long in the limit  $N \rightarrow \infty$  that we consider.

Online learning, on the other hand, is not plagued by the same problem, as we now show. Consider the first contribution to the generalization error, which we write as  $\epsilon_{g,1} = \frac{1}{2}\delta_n^2$  with

$$\delta_n = \frac{1}{\sqrt{N}} \bar{\mathbf{x}}^T \mathbf{v}_n$$

From the update equation (3) one derives that

$$\delta_{n+1} = (1 - \eta m^2) \delta_n + \eta \xi^\mu m^2 \quad (12)$$

up to correction terms which vanish for  $N \rightarrow \infty$ . Starting from the initial value  $\delta_0 = -\bar{y}$ , this can easily be iterated and the selection average carried out to give

$$\langle \delta_n^2 \rangle = \bar{y}^2 (1 - \eta m^2)^{2n} + \eta^2 m^4 \frac{1 - (1 - \eta m^2)^{2n}}{1 - (1 - \eta m^2)^2} \frac{1}{p} \sum_{\mu} (\xi^\mu)^2$$

up to  $O(N^{-1})$  corrections; an average over training sets then gives  $p^{-1} \sum_{\mu} (\xi^\mu)^2 \rightarrow \sigma^2$ . For  $n = t = 0$ , only the first term is nonzero. On the other hand, for nonzero learning time  $t$  (and values of the learning rate such that convergence occurs, *i.e.*,  $0 < \eta < 2/m^2$ ) only the second term survives because  $n = Nt \rightarrow \infty$  for  $N \rightarrow \infty$ . We therefore have for the average value of the first contribution to the generalization error:

$$\langle \epsilon_{g,1}(t=0) \rangle = \bar{y}^2 = m^2, \quad \langle \epsilon_{g,1}(t > 0) \rangle = \frac{1}{2} \sigma^2 \frac{\eta m^2}{2 - \eta m^2} \quad (13)$$

The discontinuous change at  $t = 0$  reflects the fact that  $\langle \delta_n^2 \rangle$  changes from its initial to its asymptotic value after a number of updates  $n$  which does not increase with system size  $N$ .<sup>9</sup>

We still have to calculate the evolution of the second component  $\epsilon_{g,2} = (1 - m^2) \mathbf{v}_n^2 / (2N)$

---

<sup>9</sup>Note also that we have written the selection average in (13) explicitly because  $\epsilon_{g,1}$  is no longer self-averaging: Each weight update (12) causes a change in  $\delta_n$  and  $\epsilon_{g,1}$  of order unity, and hence the fluctuations of  $\epsilon_{g,1}$  remain nonzero even for  $N \rightarrow \infty$ .




Figure 5: Critical learning rate  $\eta_c$  for online learning vs input bias  $m^2$ , for weight decay  $\lambda$  as shown and training set size  $\alpha = 0, 1, \dots, 5$  (bottom to top). Compare fig. 2 for the case of unbiased inputs.

of the generalization error (10) for the case of online learning. At first sight, the  $O(N)$  eigenvalue of  $\mathbf{A}$  appears to complicate this task. However, the component of  $\mathbf{v}_n$  along  $\bar{\mathbf{x}}$ , the corresponding eigenvector, contributes only negligibly to  $\epsilon_{g,2}$ :

$$\frac{1}{2N} \left( \frac{1}{|\bar{\mathbf{x}}|} \bar{\mathbf{x}}^T \mathbf{v}_n \right)^2 = \frac{1}{Nm^2} \epsilon_{g,1} = O(N^{-1})$$

Thus only components of  $\mathbf{v}_n$  along directions corresponding to the  $O(1)$  eigenvalues of  $\mathbf{A}$  need to be considered; their evolution can be calculated exactly as in Section 3. The only change is the rescaled eigenvalue spectrum of  $\mathbf{A}$ ; in fact, one finds that  $\epsilon_{g,2}/(1-m^2)$  is exactly the same as  $\epsilon_g = \mathbf{v}_n^2/2N$  for *unbiased* inputs of length  $\mathbf{x}^2 = N(1-m^2)$ . It is easily checked that this change of effective input vector length can be effected by replacing  $\lambda$ ,  $\sigma^2$  and  $\eta$  in the expressions for  $\epsilon_g$  by the rescaled values  $\lambda' = \lambda/(1-m^2)$ ,  $(\sigma')^2 = \sigma^2/(1-m^2)$  and  $\eta' = \eta(1-m^2)$ , and so no new calculations need to be carried out.

## 4.2 Discussion

We have already mentioned that the critical learning rate for offline learning is drastically reduced to  $\eta_c = 2/N\alpha m^2$  by the presence of input bias. For online learning,  $\eta_c$  is affected in two ways: first through the ‘rescaling’ of  $\eta$  and  $\lambda$  explained above for the calculation of  $\epsilon_{g,2}$ , and secondly through the presence of the term  $\epsilon_{g,1}$ ; eq. (13) shows






Figure 6: (a) Optimal learning rate  $\eta$  vs. final learning time  $t$  for online learning in the presence of input bias  $m^2$  (values as shown;  $\alpha = 1$ ,  $\sigma^2 = 0.1$ ,  $\lambda = 0.001$ ). (b) Resulting generalization error  $\epsilon_g$ , with results for offline learning shown for comparison (thin lines). Note that while offline learning performs (marginally) better than online learning for unbiased inputs ( $m^2 = 0$ ), it is far worse as soon as the input bias is nonzero.

that for the latter to remain finite one requires  $\eta_c < 2/m^2$ . Fig. 5 illustrates the resulting variation of  $\eta_c$  with  $m^2$  for several values of  $\alpha$  and  $\lambda$ : As the bias increases from 0, the critical learning rate first increases until it reaches the value  $2/m^2$ ; from that point onwards, it follows the curve  $\eta_c = 2/m^2$  (independently of  $\alpha$  and  $\lambda$ ) until it reaches  $\eta_c = 2$  at<sup>10</sup>  $m^2 = 1$ . In marked contrast to the case of offline learning, the critical learning rate  $\eta_c$  for online learning therefore never decreases below values of order unity, and can actually be increased by the presence of input bias.

The different effects of input bias on the critical learning rates of online and offline learning are also reflected in the generalization performance for optimal values of  $\eta$  at given final learning time. For offline learning, eq. (11) shows that the optimal  $\eta = 1/(N\alpha m^2)$ , whatever the (integer) value of  $r = t/\alpha$ . This reduces the first contribution to the offline generalization error to zero for any  $r \geq 1$ , but still leaves a nonzero term  $\epsilon_g = (1 - m^2)/2$  (which as explained above would start to decay only for extremely long learning times  $t = O(N)$ ).

For online learning, on the other hand, the optimal learning rate remains of order one even in the presence of input bias. This was to be expected from the analogous results

<sup>10</sup>This is the maximal bias in our scenario since  $\langle \mathbf{x}^2 \rangle = N > \bar{\mathbf{x}}^2 = Nm^2$ .

for the critical learning rate, and can be seen explicitly in fig. 6(a). Fig. 6(b) shows the resulting generalization error, which is seen to *decrease* as the input bias increases. Online learning therefore successfully exploits the presence of the input bias to achieve better generalization performance<sup>11</sup>. This contrasts markedly with the case of offline learning, where generalization performance (at finite learning times  $t$ ) deteriorates as soon as an input bias is present<sup>12</sup>.

## 5 Conclusions

In this paper, we have obtained exact results for the generalization error achieved by online learning from finite training sets at non-infinitesimal learning rates. These apply directly only to the simple linear model that we have considered, but also exhibit generic features which we expect to be of general relevance. For example, the calculated dependence on  $\eta$  of the asymptotic generalization error  $\epsilon_\infty$  and the convergence speed shows that, in general, sizable values of  $\eta$  can be used for training sets of limited size ( $\alpha \approx 1$ ), while for larger  $\alpha$  it is important to keep learning rates small. More important from a practical point of view is probably the explicit comparison between online and offline learning that our results allow us to make. To make this comparison fair, we considered the generalization performance of both algorithms for the respective optimal values of the learning rate at a given final learning time  $t$ . For unbiased inputs, we found in this way that online learning performs only marginally worse than offline learning, whereas it is in fact vastly superior as soon as there is any kind of input bias. This suggests strongly that online learning should generally be preferred over offline learning in problems where

---

<sup>11</sup>Wendemuth et al. (1993) view the input bias as ‘additional information’ which leads to improved generalization. In our case, the same conclusion can be arrived at by considering the extreme limit of maximal bias,  $m^2 = 1$ : In this case, the distribution of input vectors collapses to the point  $\mathbf{x} = \bar{\mathbf{x}}$ , and so perfect generalization is obtained after only one training example has been presented. (For noisy training outputs, more examples would be needed; the generalization error then decays roughly as  $\epsilon_g \sim \sigma^2/n$ , which however still gives perfect generalization  $\epsilon_g = 0$  for any finite learning time  $t$ .)

<sup>12</sup>For biased inputs, we found an offline generalization error of  $\epsilon_g = (1 - m^2)/2$  for optimally chosen  $\eta$ , which is arbitrarily close to  $\frac{1}{2}$  for  $m^2$  sufficiently small. For unbiased inputs, on the other hand,  $\epsilon_g$  for optimal  $\eta$  is generally significantly smaller than a half, as illustrated by fig. 4, for example—it can never be greater than  $\frac{1}{2}$  since otherwise  $\eta = 0$  would give a lower  $\epsilon_g$ .

biased inputs cannot be a priori excluded.

In the future, we hope to extend our analysis to dynamic ( $t$ -dependent) optimization of  $\eta$ ; based on the results of (Luo, 1991; Heskes and Wiegerinck, 1996), however, one may suspect that performance improvements over optimal fixed  $\eta$  may be small. More importantly, more complicated network architectures need to be studied in which the crucial question of local minima can be addressed. We speculate that the superiority of online learning may be even more pronounced here, due to its stochastic weight updates which should facilitate the escape from local minima. We have done some exploratory work along those lines for soft-committee machine architectures, using a fairly simple-minded approximation scheme (Sollich and Barber, 1998). Considerable challenges remain, however, and there is much scope for future work in this direction.

## References

- Biehl, M. and Schwarze, H. (1995). Learning by online gradient descent. *Journal of Physics A*, 28:643–656.
- Halkjær, S. and Winther, O. (1997). The effect of correlated input data on the dynamics of learning. In Mozer, M. C., Jordan, M. I., and Petsche, T., editors, *Advances in Neural Information Processing Systems 9*, pages 169–175, Cambridge, MA. MIT Press.
- Hertz, J. A., Krogh, A., and Thorbergsson, G. I. (1989). Phase transitions in simple learning. *Journal of Physics A*, 22:2133–2150.
- Heskes, T. and Kappen, B. (1991). Learning processes in neural networks. *Physical Review A*, 44:2718–2762.
- Heskes, T. and Wiegerinck, W. (1996). A theoretical comparison of batch-mode, on-line, cyclic, and almost-cyclic learning. *IEEE Transactions on Neural Networks*, 7(4):919–925.
- Krogh, A. and Hertz, J. A. (1992). Generalization in a linear perceptron in the presence

- of noise. *Journal of Physics A*, 25:1135–1147.
- LeCun, Y., Kanter, I., and Solla, S. A. (1991). Eigenvalues of covariance matrices - application to neural- network learning. *Physical Review Letters*, 66(18):2396–2399.
- Luo, Z. (1991). On the Convergence of the LMS Algorithm with Adaptive Learning Rate for Linear Feedforward Networks. *Neural Computation*, 3(2):226–245.
- Saad, D. and Solla, S. A. (1995). Online learning in soft committee machines. *Physical Review E*, 52:4225.
- Sollich, P. (1994). Finite size effects in learning and generalization in linear perceptrons. *Journal of Physics A*, 27:7771–7784.
- Sollich, P. (1995). Learning unrealizable tasks from minimum entropy queries. *Journal of Physics A*, 28:6125–6142.
- Sollich, P. and Barber, D. (1997a). On-line learning from finite training sets. *Europhysics Letters*, 38:477–482.
- Sollich, P. and Barber, D. (1997b). Online learning from finite training sets: An analytical case study. In Mozer, M. C., Jordan, M. I., and Petsche, T., editors, *Advances in Neural Information Processing Systems 9*, pages 274–280, Cambridge, MA. MIT Press.
- Sollich, P. and Barber, D. (1998). Online learning from finite training sets in non-linear networks. In *Advances in Neural Information Processing Systems 10*, Cambridge, MA. MIT Press. To be published.
- Wendemuth, A., Opper, M., and Kinzel, W. (1993). The effect of correlations in neural networks. *Journal of Physics A*, 26(13):3165–3185.